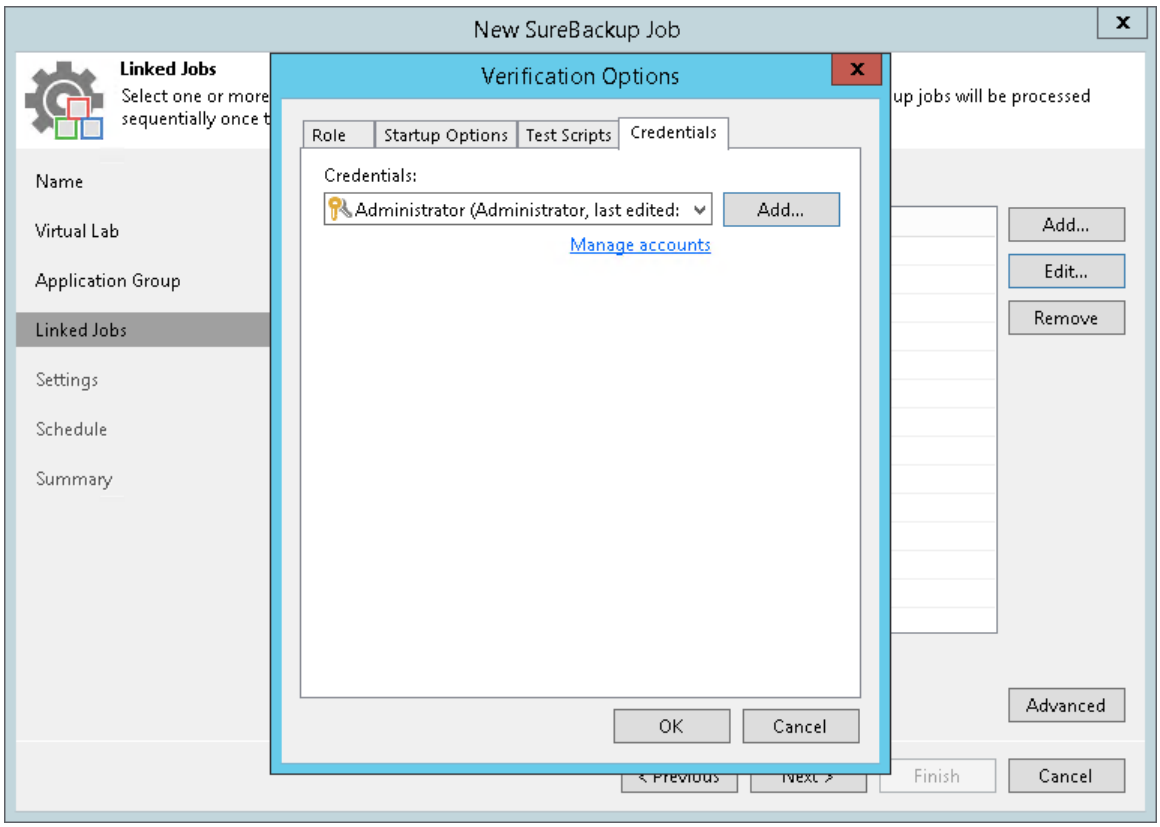
2. From the **Credentials** list, select credentials for the account under which you want to run the script.

   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 7. Specify Additional Job Settings

On the **Settings** step of the wizard, specify additional settings for the SureBackup job:

1. [For VM backups only] If you want to validate the backup file with a CRC check and make sure that the file is not corrupted, select the **Validate entire virtual disk contents** check box. You can optionally exclude VMs being a part of the application group from this test. To do this, select the **Skip validation for application group VMs** check box. For more information, see Recovery Verification Tests.
2. If you want Veeam Backup & Replication to scan VM data with antivirus software, select the **Scan the selected restore point for malware** check box. For more information, see Secure Restore.
   - if you want the antivirus software to continue scanning VM data after the first malware is found, select the **Scan the entire image** check box. For information on how to view results of the malware scan, see Viewing Recovery Verification Job Statistics.
   - If you do not want to scan VMs from the application group, select the **Skip application group machines from malware scan** check box. In this case, the antivirus will only scan VMs from linked jobs.

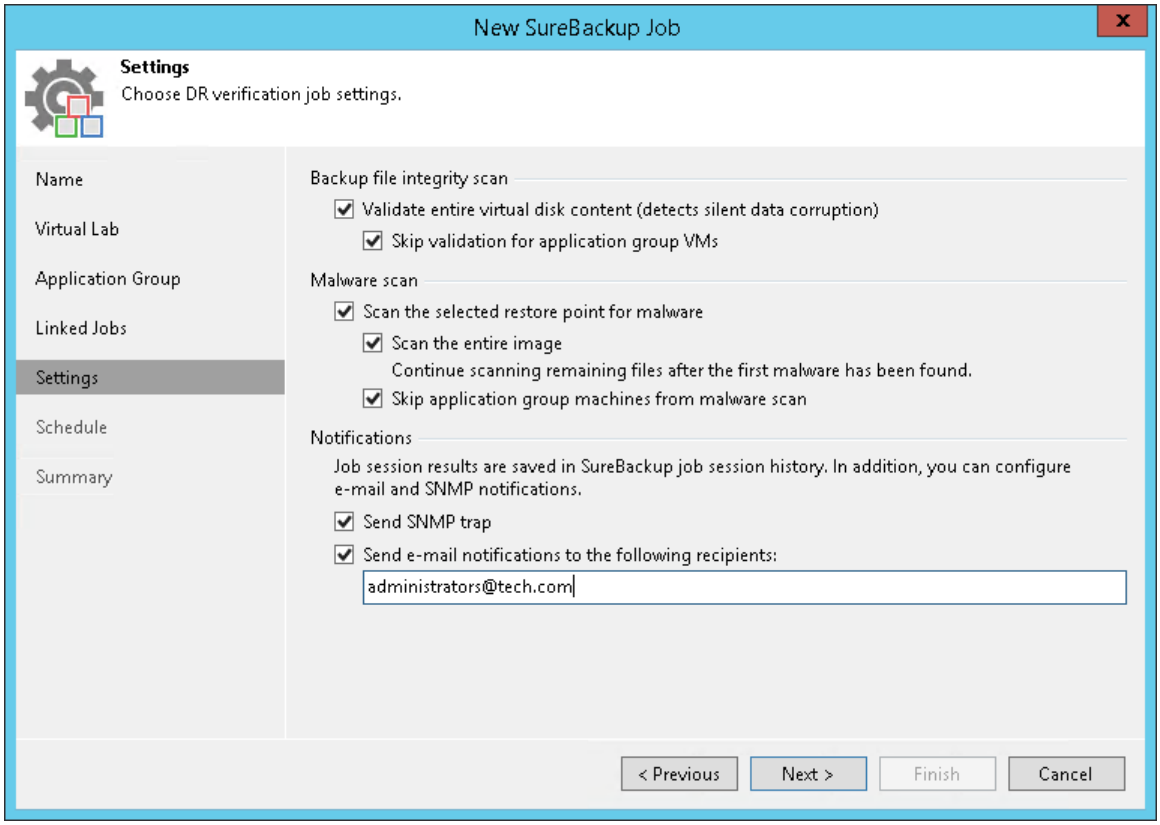
   Veeam Backup & Replication scans VM data with antivirus before running verification tests. Mind that the SureBackup job may take considerable time to complete if you are verifying backups of large sized VMs.
3. If you want to receive SNMP traps, select the **Send SNMP trap** check box.

   SNMP traps will be sent only if you configure global SNMP settings in Veeam Backup & Replication and on recipient's computer. For more information, see Specifying SNMP Settings.
4. If you want to receive notifications by email, select the **Send email notifications to the following recipients** check box. In the field below, specify recipient's email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent only if you configure global email notification settings in Veeam Backup & Replication. For more information, see Specifying Email Notification Settings.

NOTE

If you enable the **Keep the application group running after the job completes** option at the **Application Group** step of the wizard, the **Skip validation for application group VMs** option will be automatically enabled.



Step 8. Specify Job Schedule

At the **Schedule** step of the wizard, select to manually run the SureBackup job or schedule the job at specific time, for example, after the corresponding backup or replication job completes.

1. To define a job schedule, select the **Run the job automatically** check box. If this check box is not selected, you will have to manually start the job to perform recovery verification.
2. Choose the necessary schedule option for the job:
   - **Daily at** to start the job at specific time every day, on week days or on specific days.
   - **Monthly at** to start the job once a month on the specified day.

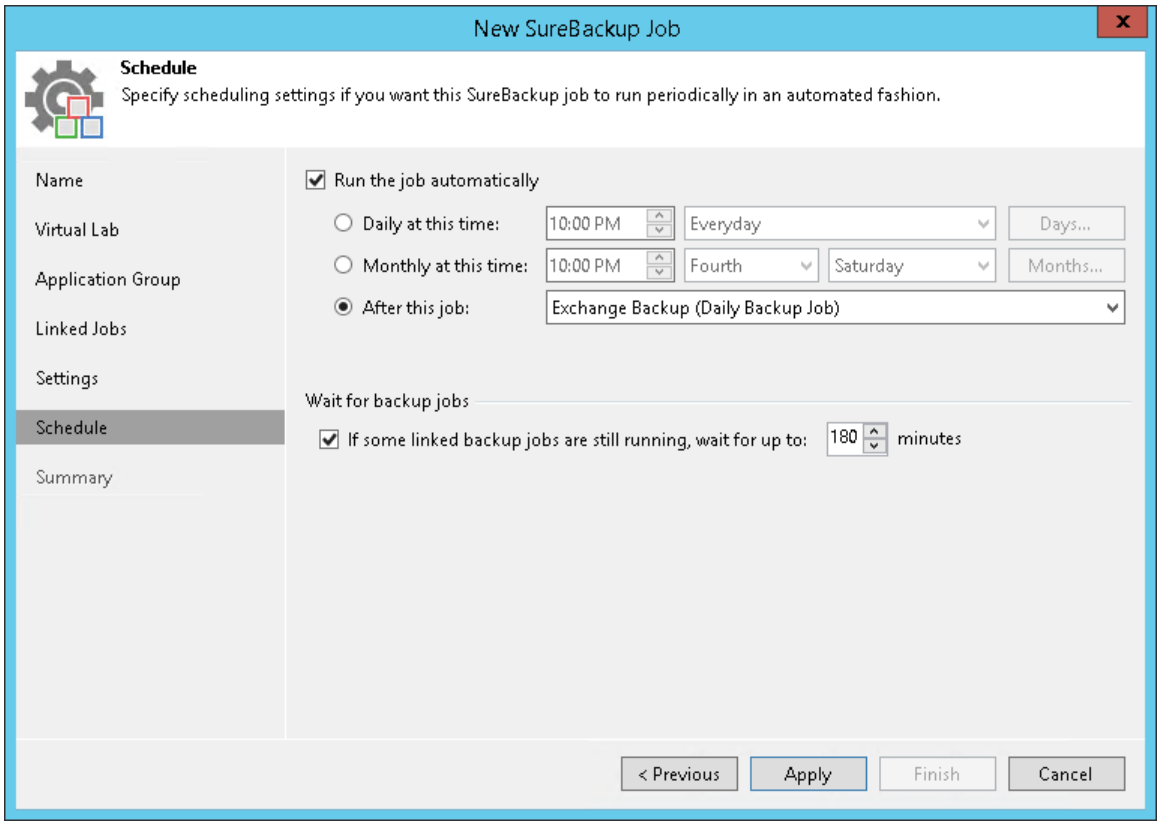
> **NOTE**
>
> When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - **After this job** to chain the job. Typically, a SureBackup job should run after the linked backup or replication job completes. In this case, the SureBackup job will verify the VM backup or VM replica created by the source backup or replication job.

     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, at the **Schedule** step of the wizard, select the **After this job** option and choose the preceding job from the list.
3. In some cases, the linked backup or replication job may not complete until the SureBackup job starts. If Veeam Backup & Replication finds out that the linked job is still running, the SureBackup job will fail to start. To overcome this situation, select the **If some linked backup jobs are still running, wait up to <N> minutes** check box and specify the necessary time period in the field on the right. If the linked job is still running, Veeam Backup & Replication will wait for the defined period of time and check the linked job after this period elapses.
   - If the linked job is finished within the specified period, the SureBackup job will start.
   - If the linked job is still running, the SureBackup job will not start.

NOTE

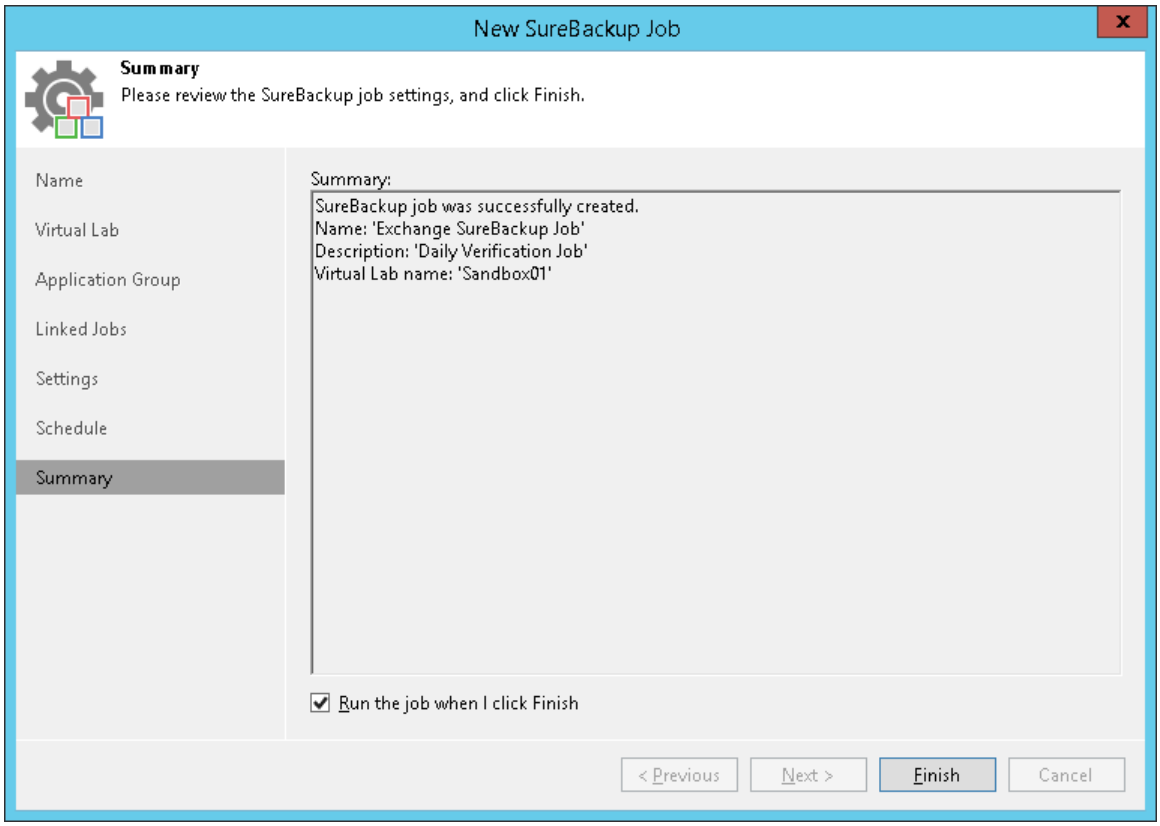
The **After this job** function will only start a job if the first job in the chain is started automatically by schedule. If the first job is started manually, jobs chained to it will not be started.



Step 9. Review Job Summary and Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of SureBackup job configuration.

1. Review details of the SureBackup job.
2. If you want to start the job right after you finish working with the wizard, select the **Run the job when I click Finish** check box.
3. Click **Finish** to save the job settings and close the wizard.



Starting and Stopping SureBackup Job

You can instruct the SureBackup job to verify the latest restore point of a VM backup or VM replica or select a specific restore point to which the VM from the backup or VM replica must be started.

To start a VM from the latest restore point:

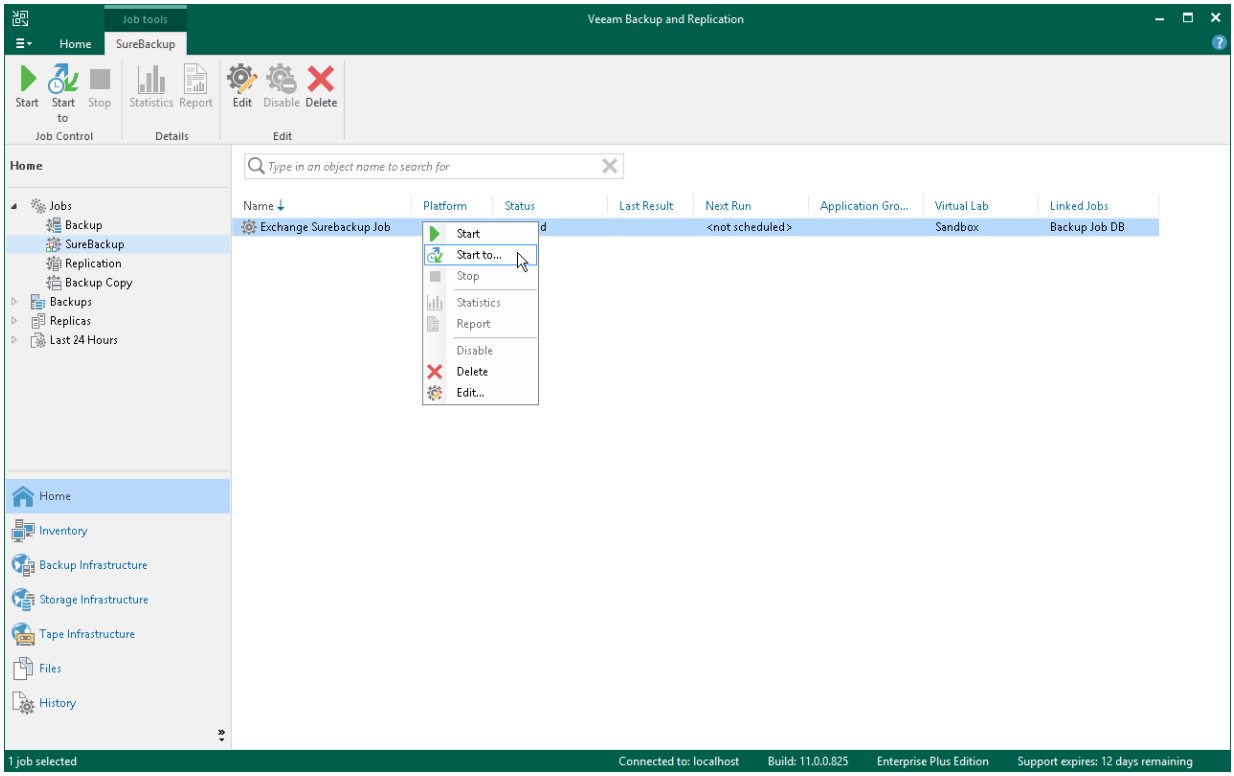
1. Open the **Home** view.
2. In the inventory pane, click **SureBackup** under **Jobs**.
3. In the working area, select the SureBackup job and click **Start** on the ribbon. You can also right-click the SureBackup job and select **Start**. Veeam Backup & Replication will start VMs in the application group and verified VMs from the latest restore point and perform necessary tests for them.

To start VMs from a specific point in time:

1. Open the **Home** view.
2. In the inventory pane, select **SureBackup** under **Jobs**.

3. In the working area, select the SureBackup job and click **Start to** on the ribbon. You can also right-click the SureBackup job and select **Start to**.

4. In the **Restore Point** window, select an approximate date of the restore point creation. Veeam Backup & Replication will pick the most recent restore point prior to the selected day and start VMs from the application group and verified VMs from this restore point.

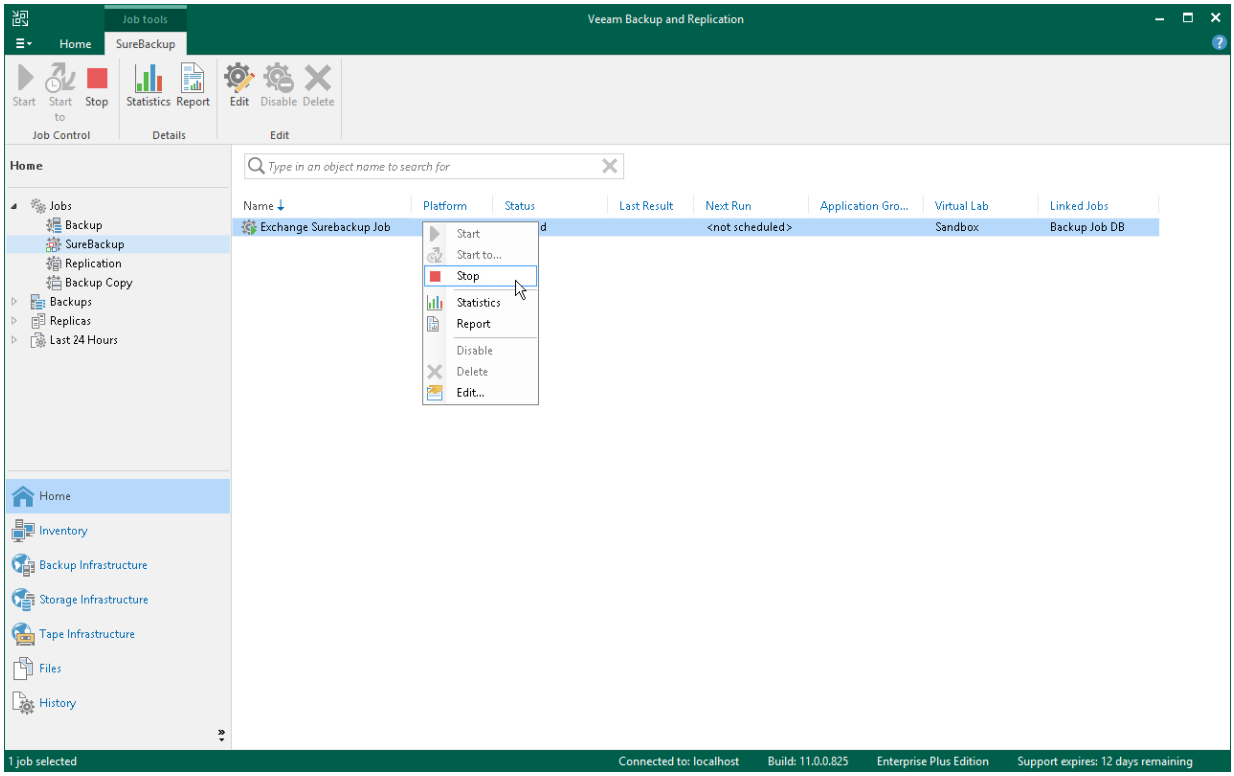


To stop a running SureBackup job session:

1. Open the **Home** view.

2. In the inventory pane, select **SureBackup** under **Jobs**.

3. In the working area, select the SureBackup job and click **Stop** on the ribbon. You can also right-click the SureBackup job and select **Stop**.



Viewing Recovery Verification Job Statistics

You can monitor how tests for verified VMs are performed while a recovery verification job is running.

To see the status of VM tests:

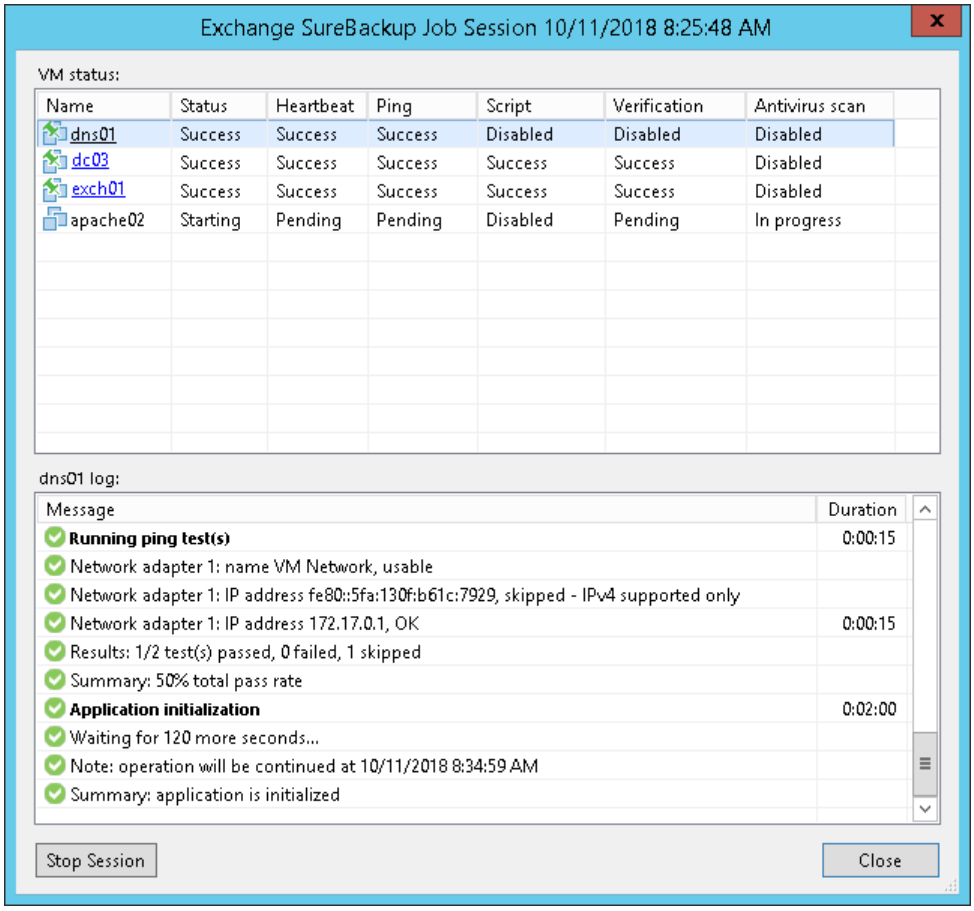
1. Open the **Home** view.
2. In the inventory pane, select **SureBackup** under **Jobs**.
3. In the working area, right-click a recovery verification job and select **Statistics**. You can also double-click the job in the list.

The job session window displays statistics for all VMs that are started during the SureBackup job: VMs from the application group in the specified order and VMs from linked jobs. For your convenience, these VMs are marked with different icons.

After the verified VM is powered on, its name is displayed as a hyperlink. You can click the link to open the VM console to see what is happening inside the VM or perform manual testing.

If some VM fails to be verified automatically, you can start it manually when this VM is powered off. To start a VM, right-click the VM in the list and select **Start**. If the application group has already been powered off by that time, it will be started again. After that, you can open the VM console and perform verification and testing manually.

If you enabled malware scan at the Settings step of the SureBackup job wizard, you can view the detailed logging of the scan process. To view logs, click the **Scan Log** button that will appear at the bottom of the job session window after the scan is complete.



Creating SureBackup Session Reports

You can generate HTML reports with statistics on the SureBackup job. A report contains detailed data on job sessions: job status, start and end time, details of the session performance, status of verified VMs and test results. You can generate a report for the whole SureBackup job or a specific job session/sessions.

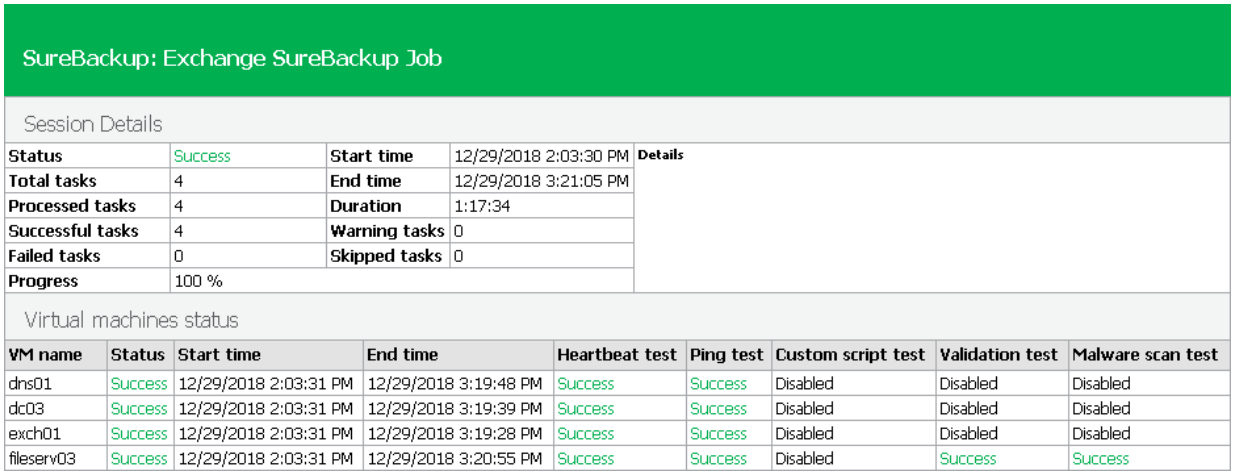
The SureBackup job report contains data on all sessions initiated for a specific job. To generate a SureBackup job report:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the SureBackup job and click **Report** on the ribbon. You can also right-click the SureBackup job and select **Report**.

The session report contains data on a single job session. To generate a session report:

1. Open the **History** view.
2. In the inventory pane, select **Jobs**.

3. In the working area, select the session and click **Report** on the ribbon. You can also right-click the session and select **Report**.



SureBackup: Exchange SureBackup Job

Session Details

Status	Success	**Start time**	12/29/2018 2:03:30 PM	**Details**
Total tasks	4	**End time**	12/29/2018 3:21:05 PM	
Processed tasks	4	**Duration**	1:17:34	
Successful tasks	4	**Warning tasks**	0	
Failed tasks	0	**Skipped tasks**	0	
Progress	100 %			

Virtual machines status

VM name	Status	Start time	End time	Heartbeat test	Ping test	Custom script test	Validation test	Malware scan test
dns01	Success	12/29/2018 2:03:31 PM	12/29/2018 3:19:48 PM	Success	Success	Disabled	Disabled	Disabled
dc03	Success	12/29/2018 2:03:31 PM	12/29/2018 3:19:39 PM	Success	Success	Disabled	Disabled	Disabled
exch01	Success	12/29/2018 2:03:31 PM	12/29/2018 3:19:28 PM	Success	Success	Disabled	Disabled	Disabled
fileserv03	Success	12/29/2018 2:03:31 PM	12/29/2018 3:20:55 PM	Success	Success	Disabled	Success	Success

XML Files with VM Roles Description

VM roles that you can assign to verified VMs and VMs from the application group are described in XML files. These XML files are stored in the `%ProgramFiles%\Veeam\Backup and Replication\Backup\SbRoles` folder on the backup server.

To add a new role, you must create a new XML file and save it to the SbRoles subfolder on the backup server. Do not save the XML file on the machine where the Veeam Backup & Replication console is installed — this will not affect the list of roles in Veeam Backup & Replication.

XML files describing VM roles have the following structure:

```
<SbRoleOptions>
 <Role>
   <SbRole>
     <Id>4CDC7CC4-A906-4de2-979B-E5F74C44832F</Id>
     <Name>Web Server</Name>
   </SbRole>
 </Role>
 <Options>
   <SbVerificationOptions>
     <ActualMemoryPercent>100</ActualMemoryPercent>
     <MaxBootTimeoutSec>300</MaxBootTimeoutSec>
     <AppInitDelaySec>120</AppInitDelaySec>
     <TestScripts>
       <TestScripts>
         <TestScript>
           <Name>Web Server</Name>
           <Type>Predefined</Type>
           <TestScriptFilePath>Veeam.Backup.ConnectionTester.exe</TestScriptFilePath>
           <Arguments>%vm ip% 80</Arguments>
         </TestScript>
       </TestScripts>
     </TestScripts>
     <HeartbeatEnabled>True</HeartbeatEnabled>
     <PingEnabled>True</PingEnabled>
   </SbVerificationOptions>
 </Options>
</SbRoleOptions>
```

The XML file with the role description contains the following tags and parameters:

Tag	Required/ Optional	Description
<SbRoleOptions>	Required	Encapsulates the VM role file.
<Role>	Required	Parent tag for a role assigned to a VM. *<SbRole>*, *<Id>* and *<Name>* are children of this tag.
<SbRole>	Required	Encapsulates basic information for a VM role: ID and name.
<Id>	Required	Unique identifier of a VM role.
<Name>	Required	Name of a VM role. The VM role name is displayed in the roles list on the **Role** tab.

Tag	Required/ Optional	Description
<Options>	Required	Parent tag for startup and test script options to be used for the defined role. *<SbVerificationOptions>*, *<ActualMemoryPercent>*, *<MaxBootTimeoutSec>*, *<AppInitDelaySec>*, *<TestScripts>*, *<Name>*, *<Type>*, *<TestScriptFilePath>*, *<Arguments>*, *<HeartbeatEnabled>*, *<PingEnabled>* are children of this tag.
<SbVerificationOptions>	Required	Encapsulates options data for a VM role.
<ActualMemoryPercent>	Optional	Percent of the original memory level that must be pre-allocated to a verified VM on the system boot.
<MaxBootTimeoutSec>	Optional	Maximum allowed time to boot a VM.
<AppInitDelaySec>	Optional	Duration of time for which Veeam Backup & Replication must wait after the VM is successfully booted in the virtual lab. After this time elapses, Veeam Backup & Replication will run test scripts. Time is specified in seconds.
<TestScripts>	Optional	Encapsulates test script data for a VM role.
<Name>	Optional	Name of a VM role. The VM role name is displayed on the **Test Scripts** tab.
<Type>	Optional	Type of the test script: *Predefined* or *Custom*.
<TestScriptFilePath>	Optional	Path to an executable file of the test script to be performed. The path can be absolute or relative.
<Arguments>	Optional	Arguments to be passed to the script. You can use the following variables: • *%vm_ip%* — IP address of a verified VM. or %vm_fqdn% — a fully qualified domain name of a verified VM. • *%log_path%* — path to a log file to which verification results are stored.
<HeartbeatEnabled>	Required	Must a heartbeat test be enabled for this VM role: *True* or *False*.

Tag	Required/ Optional	Description
<PingEnabled>	Required	Must a ping test be enabled for this VM role: *True* or *False*.

Manual Recovery Verification

Beside automatic recovery verification, you can perform manual verification of VM backups. Manual verification can be performed with all editions of Veeam Backup & Replication.

Boot Test

To perform a VM boot test, perform Instant Recovery for the verified VM. Power on the VM but do not connect the VM to the production network to avoid conflicts with the original VM.

Application Test

To perform an application test:

1. Create an isolated network.
2. Use the **Instant Recovery** wizard to restore the verified VM. At the **Ready to Apply** step of the wizard, clear the **Connect VM to network** check box.
3. When the VM is started, connect it to the isolated network.

The same procedure must be performed for all VMs that run applications on which the verified VM is dependent such as domain controller and DNS. All VMs must be connected to the same isolated network and started in the correct order: for example, DNS > domain controller > verified VM.

SureReplica

To guarantee recoverability of your data, Veeam Backup & Replication complements the SureBackup recovery verification technology with SureReplica.

SureReplica is in many respects similar to the SureBackup recovery verification. It lets you validate your disaster recovery environment without impacting the production infrastructure. You can automatically verify every created restore point of every VM replica and ensure that they are functioning as expected.

The SureReplica technology is not limited only to VM replica verification. Just like SureBackup, it provides the following capabilities:

- SureReplica: automated VM replica verification
- On-Demand Sandbox: an isolated environment for testing VM replicas, training and troubleshooting
- U-AIR: recovery of individual items from applications running on VM replicas

How SureReplica Works

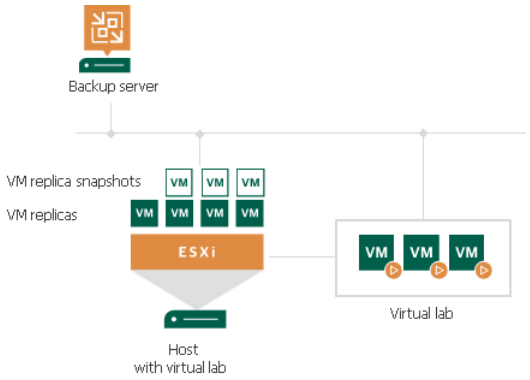
SureReplica is Veeam's technology that lets you test a VM replica for recoverability. To ensure that the VM replica is functioning properly, Veeam Backup & Replication performs its "live" verification.
Veeam Backup & Replication automatically boots the VM replica from the necessary restore point in the isolated environment, performs tests against the VM replica, powers it off and creates a report on the VM replica state.

The SureReplica technology does not require the vPower engine. A VM replica is essentially an exact copy of a VM with a set of restore points. The VM replica data is stored in the raw decompressed format native to VMware. Therefore, to start a VM replica in the virtual lab, you do not need to present its data via the vPower NFS datastore to the ESXi host. Veeam Backup & Replication re-configures the VM replica settings for recovery verification, connects the VM replica to the isolated virtual lab and powers it on.

As there is no need to publish the VM from the backup file, the SureReplica processing is typically faster than SureBackup. Subsequently, the U-AIR and On-Demand Sandbox operations are faster, too.

During VM replica verification, Veeam Backup & Replication performs the following actions:

1. Veeam Backup & Replication triggers a VMware snapshot for a VM replica. The snapshot protects the VM replica from changes while it is running. All changes made to the VM replica are written to the delta file.
2. Veeam Backup & Replication starts the VM replica in the virtual lab.
3. Veeam Backup & Replication performs tests against the verified VM replica.
4. When the verification process is over, Veeam Backup & Replication removes the delta file of the VM replica snapshot, powers off the VM replica and creates a report on its state. The report is sent to the backup administrator by email.



NOTE

Veeam Backup & Replication verifies only VM replicas in the *Ready* state. If a VM replica is in the *Failover* or *Failback* state, the verification process fails.
When Veeam Backup & Replication verifies the VM replica, it puts the VM replica to the *SureBackup* state. You cannot perform failback and failover operations for a VM replica in the *SureBackup* state until recovery verification or the U-AIR process is over and the VM replica returns to the *Ready* state.

To perform VM replica verification, you need to create the following objects:

1. Application group. During recovery verification, the VM replica is not started alone: it is started together with VMs on which the VM replica is dependent. Starting a VM replica in conjunction with other VMs enables full functionality of applications running inside the VM replica 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 VM replica and VMs from the application group are started and tested.
3. SureBackup job. The SureBackup job is a task for VM replica verification process. You can run the SureBackup job manually or schedule it to run automatically by schedule.

Replica Recovery Verification Tests

To verify a VM replica, Veeam Backup & Replication performs the same tests as for VM backup verification, except the backup validation test. You can run predefined tests or perform your own tests against VMs. The predefined tests include the following ones:

- Heartbeat test
- Ping test
- Application test

For more information, see Backup Recovery Verification Tests.

Application Group

You can add to the same application groups both VMs from backups and VMs from replicas. Keep in mind that all VMs from the application group must have at least one valid restore point created by the time the SureBackup job starts.

For more information, see Application Group.

Virtual Lab Configuration

Veeam Backup & Replication offers three types of the virtual lab configuration for VM replica verification:

- Basic single-host virtual lab
- Advanced single-host virtual lab
- Advanced multi-host virtual lab

Basic Single-Host Virtual Labs

The basic single-host virtual lab configuration can be used if your disaster recovery (DR) site is configured in the following way:

- All VM replicas that you want to verify are registered on the same ESXi host.
- All VM replicas that you want to verify are connected to the same network.

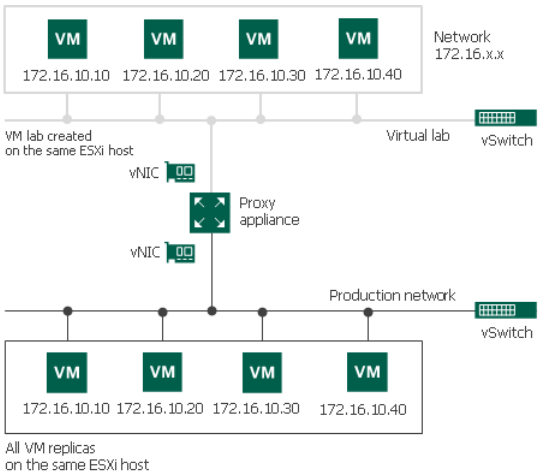
> **IMPORTANT**
>
> For this configuration type, the virtual lab must be created on the same ESXi host where VMs replicas are located. If you create the virtual lab on some other ESXi host, the SureBackup job will fail.

For the basic single-host virtual lab, Veeam Backup & Replication creates one virtual network that is mapped to the production network. Additionally, 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.

Veeam Backup & Replication automatically configures all settings for the basic single-host virtual lab. The proxy appliance is also created and configured automatically and placed to the virtual lab folder and resource pool on the ESXi host.



Advanced Single-Host Virtual Labs

The advanced single-host virtual lab configuration can be used if your virtual environment is configured in the following way:

- All VM replicas that you want to verify are located on the same ESXi host.
- VM replicas you want to verify are connected to different networks.

> **IMPORTANT**
>
> For this configuration type, the virtual lab must be created on the same ESXi host where VMs replicas are located. If you create the virtual lab on some other ESXi host, the SureBackup job will fail.

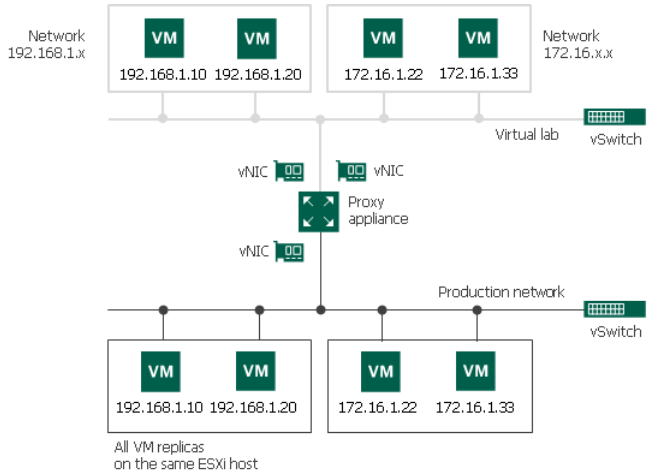
In the advanced single-host virtual lab, Veeam Backup & Replication creates several virtual networks. The number of virtual networks corresponds to the number of production networks to which verified VM replicas are connected. Networks in the virtual lab are mapped to corresponding production networks.

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.

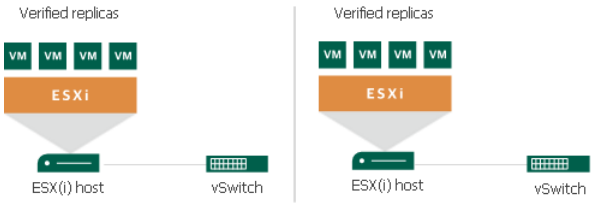


Limitations of Single-Host Virtual Labs

If VM replicas are located on different hosts, you cannot use the single-host virtual lab configuration (basic or advanced). A single-host virtual lab uses standard vSwitches that have specific configuration limitations.

When you create or edit a virtual lab, Veeam Backup & Replication creates a new port group for each isolated network in the virtual lab. All VMs in the isolated network are added to this port group. Such configuration helps differentiate the traffic passing through the standard vSwitch to the isolated network in the virtual lab.

However, the standard vSwitch has a restriction: it is "limited" to one ESXi host. A standard vSwitch is configured on a specific ESXi host. The configuration of the standard vSwitch, such as information about port groups, resides on the ESXi host where the vSwitch is configured. Other ESXi hosts in the virtual environment do not have access to this information.



For this reason, the single-host configuration can only be used if all VM replicas are registered on the same ESXi host. If attempt to verify VM replicas registered on different ESXi hosts in the single-host virtual lab, VMs from different port groups will not be able to "see" each other and communicate with each other.

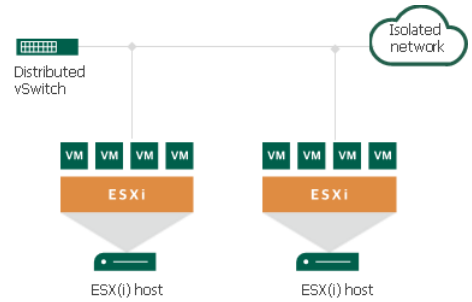
To overcome this limitation and verify VM replicas that are registered on different ESXi hosts, you can use advanced multi-host virtual labs.

Advanced Multi-Host Virtual Labs

The advanced multi-host virtual lab configuration can be used if your disaster recovery (DR) site is configured in the following way:

- VM replicas that you want to verify are located on different ESXi hosts.
- VM replicas that you want to verify are connected to one or several networks.

The advanced multi-host virtual lab leverages the VMware Distributed vSwitch (DVS) technology. For more information, see VMware Docs.



When you configure an advanced multi-host virtual lab, you must select an ESXi host on which the proxy server will be created and DVS on which Veeam Backup & Replication will create isolated networks.

Veeam Backup & Replication does not offer an option to automatically configure the DVS. The DVS must be preconfigured in your virtual environment.

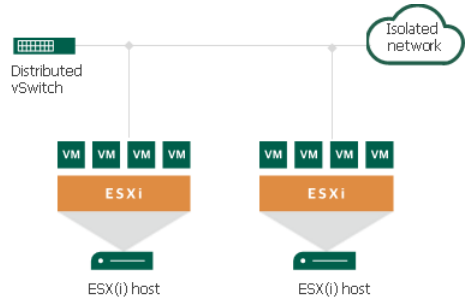
> **IMPORTANT**
>
> DVS is limited to one datacenter. For this reason, all verified VM replicas and VM replicas from the application group must belong to the same datacenter. If VM replicas belong to different datacenters, you will be able to start them in the virtual lab but Veeam Backup & Replication will not be able to automatically verify them.

Isolated Networks on DVS

For every isolated network in the virtual lab, Veeam Backup & Replication adds a new DVS port group to the DVS. The added port group is named after the isolated network.

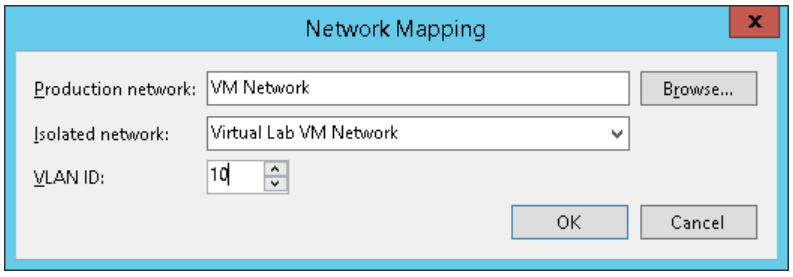
The port groups created on the DVS must be isolated from the production environment. To isolate port groups, you can use one of the following methods:

1. **Connect DVS uplinks to the isolated network**. You can link the DVS that you plan to use for recovery verification to an external isolated network using uplink adapters. This operation must be performed manually by the backup administrator.



2. **Use VLAN tagging**. This method can be used only if your router supports VLAN ID tagging. When you specify settings for isolated networks in the **New Virtual Lab** wizard, you can define different VLAN IDs for different isolated networks. Setting VLAN IDs restricts communication of VM replicas in the isolated network from the production environment.

> **IMPORTANT**
>
> If your network does not support VLAN ID tagging or the virtual lab is configured incorrectly, VM replicas will be started in the virtual lab but Veeam Backup & Replication will not be able to automatically verify them.



Port Groups and VLAN IDs

You need to be extremely careful when specifying port group and VLAN ID settings for the advanced multi-host virtual lab.

Port Groups in Advanced Multi-Host Virtual Labs

For the advanced multi-host virtual lab, Veeam Backup & Replication uses an existing DVS that was configured by the backup administrator beforehand. Veeam Backup & Replication creates a number of new port groups on the DVS, one per isolated network in the virtual lab.

When Veeam Backup & Replication creates a new port group, it performs a check for the DVS selected for the virtual lab:

- If a port group with the specified name already exists, Veeam Backup & Replication starts using it for the virtual lab. However, in this case, Veeam Backup & Replication will not be the owner of this port group.
- If a port group with the specified name does not exist, Veeam Backup & Replication creates it and becomes the owner of the created port group.

When a virtual lab is removed, Veeam Backup & Replication checks the ownership of the port group:

- If Veeam Backup & Replication is not the owner of the port group, the port group remains on the DVS. Veeam Backup & Replication stops using it.
- If Veeam Backup & Replication is the owner of the port group, Veeam Backup & Replication removes this port group from the DVS.

Several virtual labs can use the same port group. For this reason, you should be extremely careful when removing virtual labs. If Veeam Backup & Replication is the owner of the virtual lab and the port group is removed, other virtual labs using this port group may fail to start.

VLAN IDs in Advanced Multi-Host Virtual Labs

A DVS port group has VLAN ID settings. If you select an existing port group for the virtual lab, you must specify its VLAN ID in the virtual lab settings.

- If VLAN ID settings are specified correctly, Veeam Backup & Replication will be able to configure the virtual lab and verify VM replicas in it.
- If VLAN ID settings are specified not correctly, Veeam Backup & Replication will report an error informing that the selected port group exists but cannot be used due to incorrect VLAN ID settings.

SureBackup Job for VM Replicas

You can verify VM replicas with the SureBackup job. The SureBackup job aggregates all settings and policies of a recovery verification task, such as application group and virtual lab to be used, VM replicas that must be verified in the virtual lab and so on. The SureBackup job can be run manually or scheduled to be performed automatically.

When a SureBackup job runs, Veeam Backup & Replication first creates an environment for VM replica verification:

1. Veeam Backup & Replication starts the virtual lab.
2. In the virtual lab, it starts VMs from the application group in the required order. VMs from the application group remain running until the verified VM replicas are booted and tested. If Veeam Backup & Replication does not find a successful VM replica or backup for any of VMs from the application group, the SureBackup job will fail.

When the virtual lab is ready, Veeam Backup & Replication starts VM replicas from the necessary restore point, tests and, depending on the specified settings, verifies them one by one or creates several streams and tests VM replicas simultaneously. If Veeam Backup & Replication does not find a successful restore point for any of verified VM replicas, verification of this VM replica fails, but the job continues to run.

By default, you can start and test up to three VM replicas 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 may decrease.

When 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 replication job linked to it. The VM replica files may be locked by the replication job and the SureBackup job will be unable to verify such replica. In this situation, Veeam Backup & Replication will not start the SureBackup job until the replication job is over.

To overcome the situation of job overlapping, you may chain the replication and SureBackup jobs or define the timeout period for the SureBackup job. For more information, see Specifying Job Schedule.

> **NOTE**
>
> You can mix VM backups and replicas in the recovery verification job. For example, the application group may contain VMs that will be started from backup files and the job linked to the recovery verification job may be a replication job. Veeam Backup & Replication supports any type of a mixed scenario. Note that VMs that you verify with a SureBackup job must belong to the same platform — VMware or Hyper-V.

SureBackup Job for VM Replicas Processing

The recovery verification process for VM replicas includes the following steps:

1. **Getting virtual lab configuration**. Veeam Backup & Replication gets information about configuration of the virtual lab where verified VM replicas must be started.
2. **Starting virtual lab routing engine**. Veeam Backup & Replication starts a proxy appliance. The proxy appliance is used as a gateway that provides access to VM replicas the virtual lab.
3. **Publishing**. Veeam Backup & Replication triggers a protective VMware snapshot for the verified VM replica.

4. **Reconfiguring**. Veeam Backup & Replication updates configuration files of the VM replica to connect the VM replica to the isolated network in the virtual lab.
5. **Configuring DC**. If the VM replica has the Domain Controller or Global Catalog role, the VM replica is reconfigured.
6. **Powering on**. Veeam Backup & Replication powers on the VM replica in the isolated network.
7. **Heartbeat test**. Veeam Backup & Replication checks whether the green or yellow VMware Tools heartbeat signal is coming from the VM replica or not. If the VM replica has no VMware Tools, the test is not performed and a notification is written to the session details.
8. **Running ping tests**. Veeam Backup & Replication checks if the VM replica responds to the ping requests or not. If the VM replica 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.
9. **Application initialization**. Veeam Backup & Replication waits for applications installed in the VM replica, for example, Microsoft SQL Server, to start. The application initialization period is defined in the properties of the SureBackup job and by default is equal to 120 sec. Depending on the software installed in a VM, the application initialization process may require more time than specified in the SureBackup 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 error situation occurs, you need to increase the Application initialization timeout value and start the job once again.
10. **Running test scripts**. Veeam Backup & Replication runs scripts to test whether the application installed in the VM replica is working correctly or not. If the VM replica has no VMware Tools installed and/or there are no NICs and mapped networks for them, Veeam Backup & Replication skips tests that use variables %vm_ip% and %vm_fqdn%, as the IP address of the VM cannot be determined. Test results are written to the 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.
11. **Powering off**. After all tests have been performed, Veeam Backup & Replication powers off the verified VM replica.
12. **Unpublishing**. Veeam Backup & Replication deletes the protective VMware snapshot and rollbacks all changes made to the VM replica while it was running in the virtual lab.
13. **Stopping virtual lab engine**. Veeam Backup & Replication powers off the proxy appliance in the virtual lab.

Stabilization Algorithm

To perform tests for a VM replica without errors, Veeam Backup & Replication needs to know that the VM replica is ready for testing. To determine this, Veeam Backup & Replication waits for the VM replica to reach a "stabilization point": — the moment when the VM replica booted and reports it is ready for tests. After the stabilization point has been reached, Veeam Backup & Replication can start heartbeat tests, ping tests and test scripts against the VM replica.

Veeam Backup & Replication establishes the stabilization point with the help of VMware parameters that it gets from the VM replica. Depending on the VM replica configuration, it uses one of the three algorithms:

- **Stabilization by IP**. This algorithm is used if the VM replica 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 replica for mapped networks that is sent by VMware Tools running in the VM replica. The sent IP address must be valid and must not change for a specific period of time.
- **Stabilization by heartbeat**. This algorithm is used if the VM replica has VMware Tools installed but there are no NICs and mapped networks for them. In this case, Veeam Backup & Replication waits for the green or yellow heartbeat signal from VMware Tools installed inside the VM replica.

- **Stabilization by Maximum allowed boot time**. This algorithm is used if the VM replica has neither VMware Tools installed, nor NICs and mapped networks for them. In this case, Veeam Backup & Replication will wait for the time specified in the **Maximum allowed boot time** field, which is considered to be a stabilization period for the VM replica. Once this time interval is exceeded, Veeam Backup & Replication considers that the VM replica is successfully booted and is ready for testing.

Once the stabilization point has been established, Veeam Backup & Replication runs ping, heartbeat tests and test scripts against the verified VM replica.

The stabilization process cannot exceed the value specified in the **Maximum allowed boot time** field. If the stabilization point cannot be determined within the **Maximum allowed boot time**, the recovery verification process will be finished with the timeout error. For this reason, you should be careful when specifying this value. Typically, a VM replica started by a recovery verification job requires more time to boot that a VM started regularly. When such error situation occurs, you need to increase the **Maximum allowed boot time** value and start the job again.

On-Demand Sandbox

If you need to perform tests for production VMs, you can use an On-Demand Sandbox™. 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

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.

To create the On-Demand Sandbox, you must configure the following objects:

- Virtual lab in which VMs will be started. For more information, see Virtual Lab.
- Application group. The application group must include all VMs and/or VM replicas that you want to start in the On-Demand Sandbox. This can be one VM or a group of VMs that work together. For more information, see Application Group.
- SureBackup job. The virtual lab and application group must be linked to this job. For more information, see SureBackup Job.

On-Demand Sandbox for Storage Snapshots

In On-Demand Sandbox, you can start VMs from snapshots existing on the production storage array. You can use On-Demand Sandbox to test VMs, troubleshoot issues, perform training and so on.

Veeam Backup & Replication offers On-Demand Sandbox functionality for the following storage systems:

- Dell EMC VNX(e), Unity XT/Unity
- HPE 3PAR StoreServ and Primera, including secondary volumes — HPE 3PAR Peer Persistence
- HPE StoreVirtual/LeftHand/P4000 series and HPE StoreVirtual VSA (Virtual Storage Appliance)
- IBM Spectrum Virtualize, including secondary IBM volumes — IBM Spectrum Virtualize HyperSwap
- NetApp, including secondary arrays — NetApp SnapMirror and NetApp SnapVault
- HPE Nimble storage, including secondary arrays — Nimble Snapshot Replicated Copy
- Universal Storage API Integrated Systems

On-Demand Sandbox configuration where VMs from storage snapshots are started is similar to configuration of the regular On-Demand Sandbox. To start a VM from the storage snapshot in the isolated environment, you must configure the following objects:

- **Virtual lab**. The virtual lab must mirror the networking scheme of the production environment. You can configure a new virtual lab or use an existing virtual lab. Any type of the virtual lab configuration is supported: basic single-host, advanced single-host or advanced multi-host. For more information, see Virtual Lab.
- **Application group**. The application group must contain one or several VMs that you want to start in the On-Demand Sandbox. You can select VMs from volumes or LUNs on the storage system. During the SureBackup job, Veeam Backup & Replication will detect the latest snapshot for this volume or LUN and start the VM from this snapshot. For more information, see Application Group.
- **SureBackup job**. You must link the application group with VMs and virtual lab to the SureBackup job. For more information, see SureBackup Job.

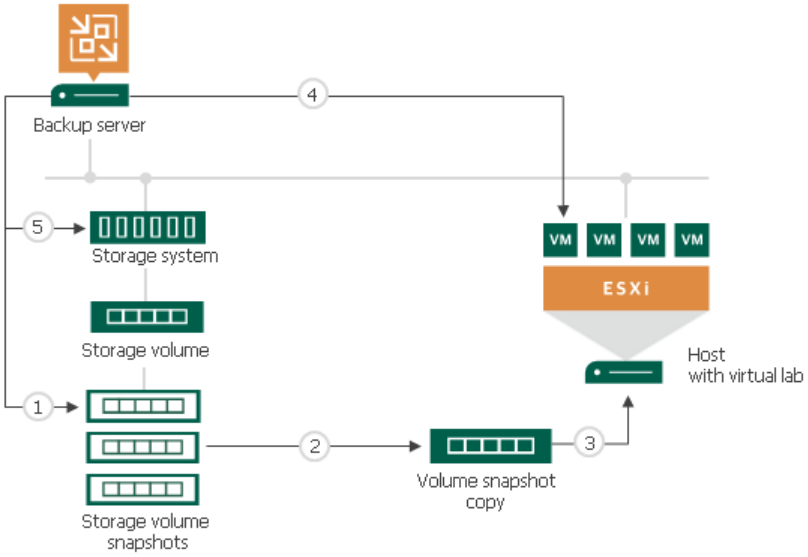
How On-Demand Sandbox for Storage Snapshots Works

To start a VM from a storage snapshot in On-Demand Sandbox, Veeam Backup & Replication needs to present this storage snapshot to an ESXi host as a datastore. To do this, Veeam Backup & Replication performs the following actions:

1. Veeam Backup & Replication detects the latest storage snapshot for the VM whose disks are located on the storage system.
2. Veeam Backup & Replication triggers the storage system to create a copy of the storage snapshot. The snapshot copy helps protect the storage snapshot from changes.

   To create a snapshot copy, Veeam Backup & Replication uses the same technology as for Data Recovery from Storage Snapshots. The technology choice depends on licenses installed on the storage system. For more information, see the Data Recovery from Storage Snapshots section in the Veeam Backup & Replication User Guide.
3. The snapshot copy is presented as a new datastore to the ESXi host on which the virtual lab is registered.

4. Veeam Backup & Replication performs regular operations required for On-Demand Sandbox: reconfigures the VMX file, starts the VM, performs necessary tests for it and so on.
5. After you finish working with VMs and power off On-Demand Sandbox, Veeam Backup & Replication performs cleanup operations: powers off the VM and the proxy appliance in the virtual lab, unmounts the datastore from the ESXi host and triggers the storage system to remove the snapshot copy.



Number of Mounted NFS Datastores

You can add to the application group several VMs that reside on different storage snapshots. In this case, Veeam Backup & Replication will trigger several snapshot copies (one per each storage snapshot) and present the equal number of datastores to the ESXi host.

The number of NFS datastores that can be mounted to the ESXi host is limited by VMware vSphere. If number of snapshot copies is great, Veeam Backup & Replication may fail to present all of them as datastores to the ESXi host. In this case, VMs in the application group will not be started and the SureBackup job will fail. For more information about limitations, see this VMware KB article.

To overcome this situation, Veeam Backup & Replication offers the mechanism of the snapshot copy re-mounting:

1. If Veeam Backup & Replication detects that there are not enough resources to mount a datastore, it displays a warning and offers you to free up resources on the ESXi host.
2. During the next 20 minutes, Veeam Backup & Replication attempts to mount the datastore with the time interval of 2 minutes.
3. If resources are freed and Veeam Backup & Replication manages to mount the datastore, VMs in the application group are started and the SureBackup job continues to run. If resources on the ESXi hosts are not freed within 20 minutes, the SureBackup job fails.

Limitations for On-Demand Sandbox for Storage Snapshots

Before you start using On-Demand Sandbox for storage snapshots, check limitations for Data Recovery from Storage Snapshots. For more information, see the General Limitations section in the Veeam Backup & Replication User Guide.

Mixed Scenarios

You can start VMs from different sources in the On-Demand Sandbox:

- VM backups
- VM replicas
- VMs from storage snapshots

For example, you can add VMs from backups and VMs from storage snapshots to the same application group and link a replication job to the SureBackup job.

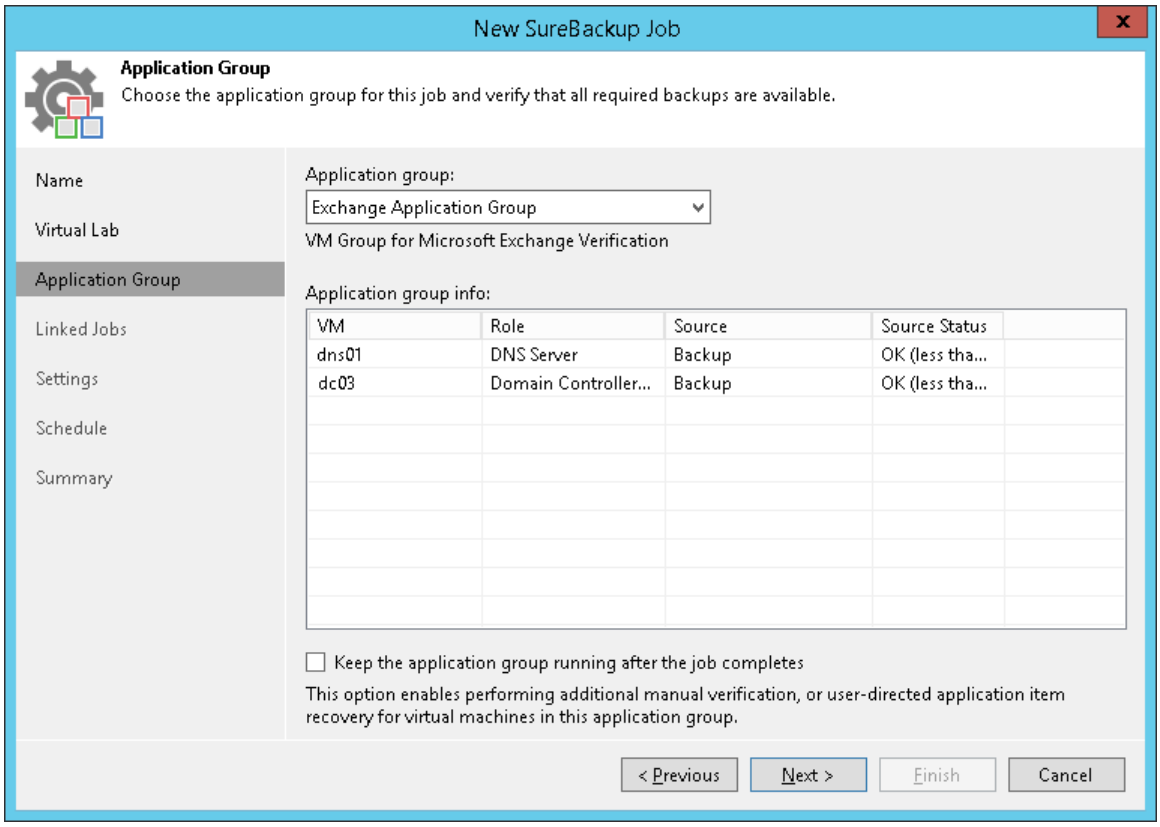
You cannot link jobs that trigger snapshots on storage arrays to the SureBackup job. This option is not supported.

Type of Job/Object	SureBackup	SureReplica	SureSnap
Application group	●	●	●
Linked job	●	●	○

Configuring On-Demand Sandbox

To configure the On-Demand Sandbox, perform the following steps:

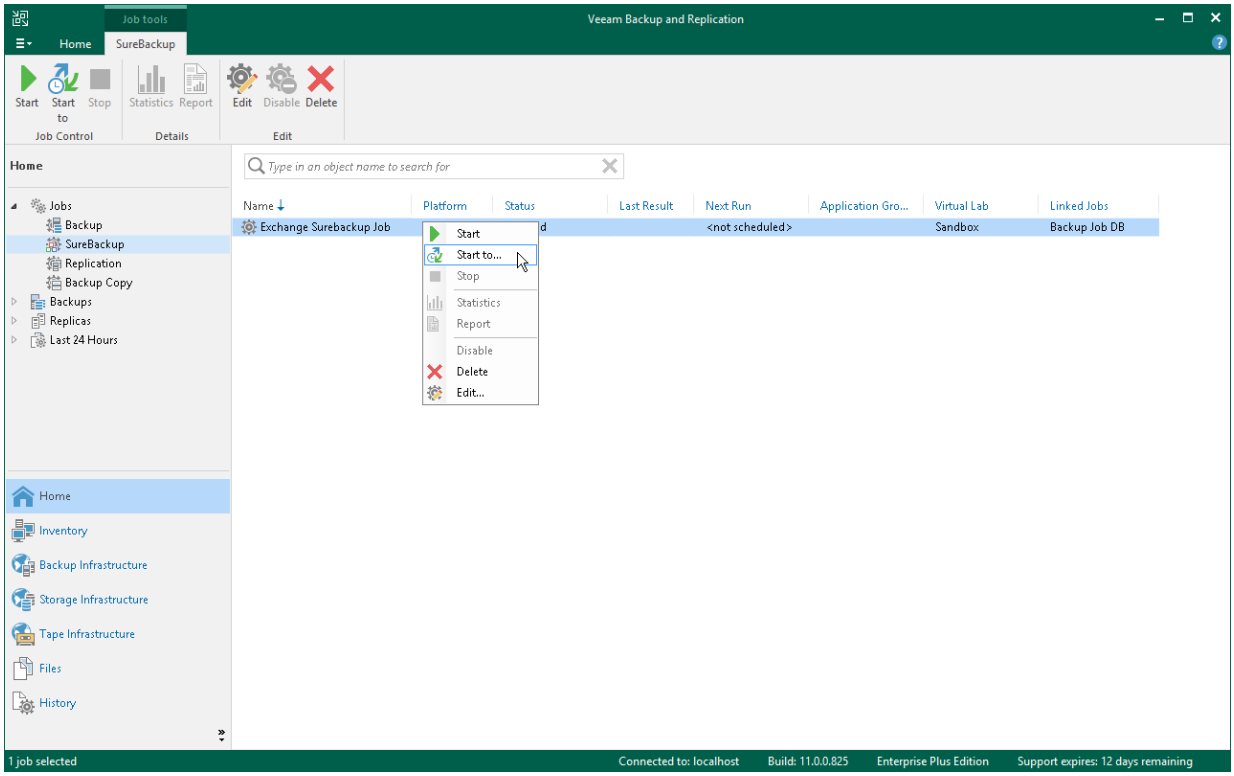
1. Configure a virtual lab in which you plan to start VMs. For more information, see Creating Virtual Lab.
2. Configure an application group. The application group must contain all VMs that you plan to start in the On-Demand Sandbox and all VMs on which these VMs are dependent. For more information, see Creating Application Group.
3. Configure a SureBackup job:
   a. Launch the **New SureBackup Job** wizard.
   b. At the **Virtual Lab** step of the wizard, select the configured virtual lab.
   c. At the **Application Group** step of the wizard, select the configured application group.
   d. Select the **Keep the application group running after the job completes** check box.
   e. Configure other job settings as required and save the job settings.



To start VMs in the On-Demand Sandbox, run the SureBackup job:

1. Open the **Home** view.
2. In the inventory pane, select **SureBackup**.
3. In the working area, right-click the configured SureBackup job and select **Start** or **Start to**.

Veeam Backup & Replication will start the virtual lab and power on VMs from the application group in the virtual lab. You will be able to connect to VMs and perform tests for them.



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.

If you want to restore data directly from storage snapshots, see the Data Recovery from Storage Snapshots section in the Integration with Storage Systems Guide.

> **NOTE**
>
> Veeam Backup & Replication provides backward compatibility: backups created with previous product versions can be restored with later product versions. However, backups created with later product versions cannot be restored with previous product versions.

VM Recovery

VM recovery includes the following methods:

- Instant Recovery to VMware vSphere — to instantly recover workloads (VMs, EC2 instances, physical servers and so on) directly from compressed and deduplicated backup files as VMware vSphere VMs. When you perform Instant Recovery, Veeam Backup & Replication mounts recovered VM images to a host directly from backups stored on backup repositories.

  Instant Recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production workloads. However, Instant Recovery provides for VMs "temporary spares" with limited I/O performance. To provide the recovered VMs full I/O performance, you must finalize Instant Recovery — migrate the recovered VMs to production environment. If you do not want to migrate the recovered VM, you can stop publishing it. This removes the recovered VM.

  Use Instant Recovery for tier 1 VMs with little tolerance for business interruption and downtime. Besides disaster recovery matters, Instant VM Recovery can also be used for testing purposes.

- Instant Recovery to Microsoft Hyper-V — to instantly recover workloads directly from compressed and deduplicated backup files as Microsoft Hyper-V VMs. In many respects, this method is similar to Instant Recovery to VMware vSphere.

- Entire VM restore — to recover entire VMs. When you recover VMs, you extract VM images from backups to the production storage. Entire VM restore takes more resources and time to complete than Instant Recovery but recovers VMs with full I/O performance. You also do not need to perform additional steps to finalize the recovery process.

  Use entire VM restore for VMs that require full I/O performance as soon as they are recovered and that tolerate some downtime.

- Staged restore — to run executable scripts for VMs before recovering them to the production environment. Staged restore is a part of entire VM restore.

  Use this option when you need to make sure that recovered VMs do not contain any personal or sensitive data.

- Restore to Microsoft Azure — to restore workloads of different data protection environments as Microsoft Azure VMs.

- Restore to Amazon EC2 — to restore workloads of different data protection environments as EC2 instances.

- Restore to Google Compute Engine (GCE) — to restore workloads of different data protection environments as Google VM instances.

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.

When you perform Instant Recovery, Veeam Backup & Replication mounts workload images to a host directly from backups stored on backup repositories. This means that Veeam Backup & Replication creates fully functioning "temporary spares" with limited I/O performance. To provide full I/O performance, you must migrate these "temporary spares" to the production site. For more information, see Migration of Recovered VMs to Production Site.

Besides disaster recovery matters, Instant Recovery can also be used for testing purposes. Instead of extracting workload images to production storage to perform regular disaster recovery (DR) testing, you can run a workload directly from a backup file, boot it and make sure the guest OS and applications are functioning properly. For more information, see Finalizing Instant Recovery to VMware vSphere.

Instant Recovery supports bulk processing so you can immediately recover multiple workloads at once. If you perform Instant Recovery for several workloads, Veeam Backup & Replication uses the resource scheduling mechanism to allocate and use optimal resources required for Instant Recovery. For details, see Resource Scheduling.

Supported Backup Types

You can recover workloads from the following types of backups:

- Backups of VMware vSphere virtual machines created by Veeam Backup & Replication
- Backups of vCloud Director virtual machines created by Veeam Backup & Replication
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud Platform
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of RHV VMs created by Veeam Backup for RHV

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 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.

To improve I/O performance for a recovered VM, you can redirect VM changes to a specific datastore that is closer to the host where the VM resides. In this case, Veeam Backup & Replication will trigger a snapshot and will put it to the *Veeam IR* directory on the selected datastore, together with metadata files holding changes to the VM image.

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. In this case, original VM data will be pulled from the NFS datastore to the production storage and consolidated with VM changes while the VM is still running. Storage vMotion, however, can only be used if you select to keep VM changes on the NFS datastore without redirecting them. Note that to use Storage vMotion, you need an appropriate VMware license.
- Use Quick Migration. In this case, Veeam Backup & Replication will perform a two-stage migration procedure — instead of pulling data from the vPower NFS datastore, it will recover the VM from the backup file on the production server, then move all changes and consolidate them with the VM data.

For more information on the relocation methods, see Quick Migration. For more information on how to launch the migration for workloads recovered with Instant Recovery, see Finalizing Instant Recovery to VMware vSphere.

Performing Instant Recovery to VMware vSphere

With Instant Recovery, you can recover different workloads from backups and register them as VMware vSphere VMs. For the list of backups that you can use for recovery, see Supported Backup Types.

To perform Instant Recovery to VMware vSphere, use the **Instant Recovery to VMware** wizard.

> **NOTE**
>
> If you want to recover workloads as Hyper-V VMs, see Performing Instant Recovery to Hyper-V.

Before You Begin

Before you perform Instant Recovery, consider the following:

- You can recover a workload from a backup that has at least one successfully created restore point.
- If you recover a workload to the production network, make sure that the original workload is powered off.
- Restore of CSV (Cluster Shared Volumes) is not supported. Cluster disks used as CSV are automatically excluded from restore.
- We strongly recommend to have `dracut` and `mkinitrd` installed on Linux workloads that will be restored. Otherwise, they may not boot after restore.
- If you want to scan recovered VM data for viruses, check the secure restore requirements and limitations.

- You must provide enough free disk space in vPower NFS datastore. The minimum amount of free space must equal the RAM capacity of the recovered VM plus 200MB. For example, if the recovered VM has 32 GB of virtual RAM, 32.2 GB of free space is required.

  By default, vPower NFS datastore is located in the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`. The vPower NFS datastore is not used when you select to redirect virtual disk updates to a VMware vSphere datastore when configuring the job.

- [For Veeam Quick Migration with Smart Switch] In addition to the disk space mentioned above, you need to provide more disk space in vPower NFS datastore. The minimum amount of free space must equal the RAM capacity of the recovered VM.

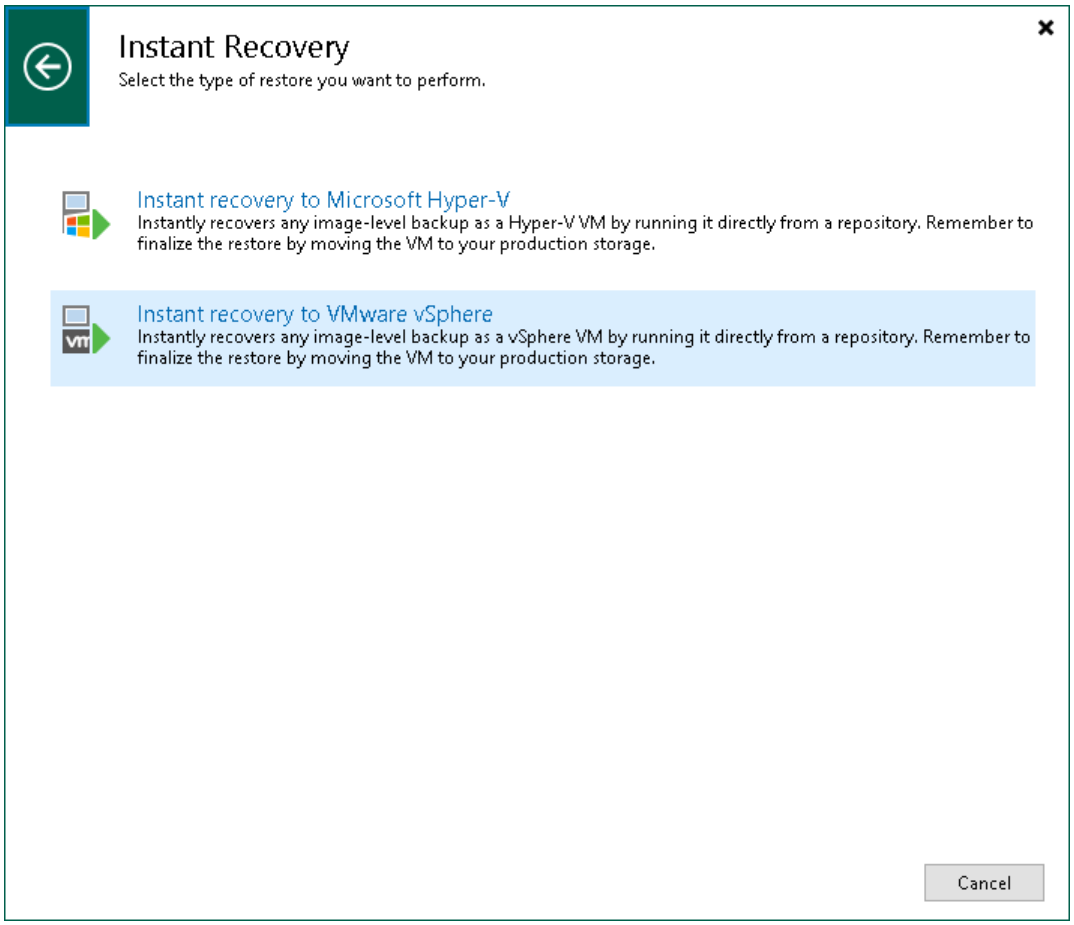
- [For Nutanix AHV VMs] Instantly recoverd VM will have default virtual hardware settings: 2 CPU cores, 4GB RAM and one network adapter. If you want to change the default settings, turn off the VM and set the required virtual resources. Note that you must not switch off the instant recovery session before turning off the VM.

Step 1. Launch Instant Recovery Wizard

To launch the **Instant Recovery to VMware** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select one of the following:
  - **VMware vSphere** > **Restore from backup** > **Entire VM restore** > **Instant Recovery** – if you want to recover VMware vSphere VMs from a VM backup created by Veeam Backup & Replication.
  - **VMware vCloud Director** > **Restore from backup** > **VM restore** > **Entire VM restore** > **Instant Recovery** > **Instant Recovery to VMware vSphere** – if you want to recover vCloud Director VMs from a VM backup created by Veeam Backup & Replication.
  - **Microsoft Hyper-V** > **Restore from backup** > **Entire VM restore** > **Instant Recovery to VMware vSphere** – if you want to recover Hyper-V VMs from a VM backup created by Veeam Backup & Replication.
  - **Agent** > **Entire machine restore** > **Instant Recovery to VMware vSphere** – if you want to recover physical machines from a backup created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.
  - **AWS EC2 backup** > **Entire machine restore** > **Instant Recovery to VMware vSphere** – if you want to recover EC2 instances from a backup created by Veeam Backup for AWS.
  - **Azure IaaS** > **Entire machine restore** > **Instant Recovery to VMware vSphere** – if you want to recover Azure VMs from a backup created by Veeam Backup for Microsoft Azure.
  - **GCE backup** > **Entire machine restore** > **Instant Recovery to VMware vSphere** – if you want to recover VM instances from a backup created by Veeam Backup for Google Cloud Platform.
  - **Nutanix backup** > **Entire machine restore** > **Instant Recovery to VMware vSphere** – if you want to recover VMs from backups created by Veeam Backup for Nutanix AHV.
  - **RHV Proxy** > **Entire machine restore** > **Instant Recovery to VMware vSphere** – if you want to recover VMs from backups created by Veeam Backup for RHV.

- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, select workloads that you want to recover and click **Instant Recovery** > **VMware vSphere** on the ribbon. Alternatively, you can right-click one of the selected workloads and select **Instant recovery** > **VMware vSphere**.



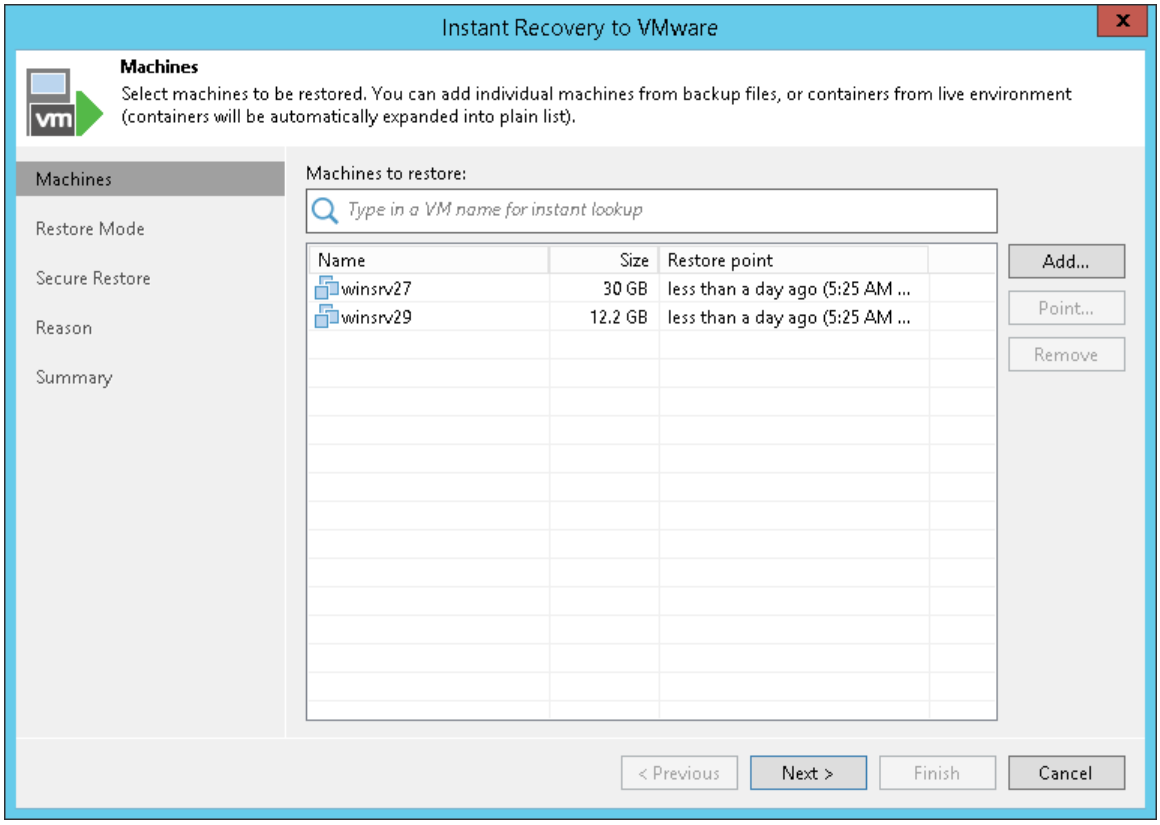
Step 2. Select Workloads

At the **Machines** step of the wizard, select one or multiple workloads that you want to recover:

1. Click **Add**.
2. In the **Backup Browser** window, do the following:
   a. [For VMware vSphere VMs and vCloud Director VMs] Select where to browse for VMs:
      - **From infrastructure** – browse the virtual environment and select VMs or VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) to recover. If you select a VM container, Veeam Backup & Replication will expand it to a plain VM list.

        When you add a VM to the list, Veeam Backup & Replication displays information about the most recent restore point in the **Restore point** column. If no restore point is available for the added VM, Veeam Backup & Replication displays a warning next to this VM.
      - **From backup** – browse existing backups and select virtual machines from backups.

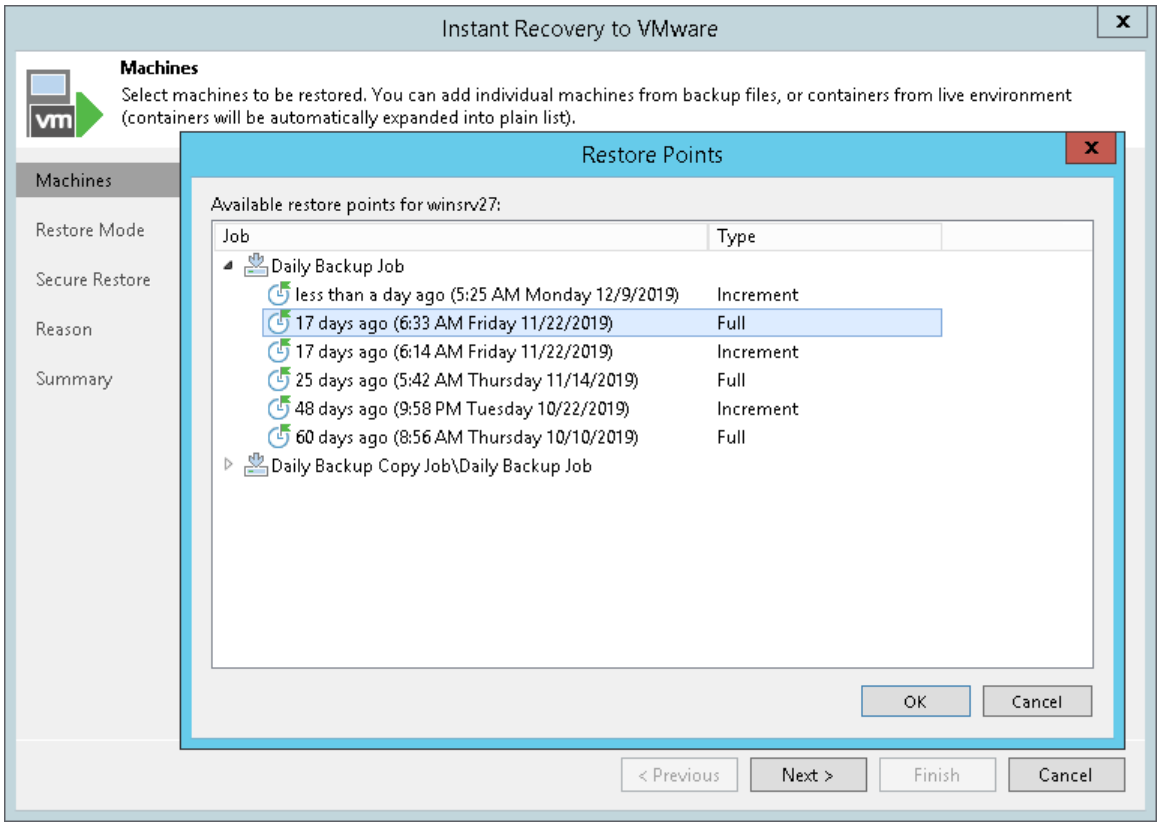
   b. [For other workloads] In the list of backup jobs, expand a job and select workloads.
   c. Click **Add**.



Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore point to recover workloads. You can recover a workload to an earlier state, if necessary:

1. In the **Machines to restore** list, select a workload.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point from which you want to recover the workload.



Step 4. Select Restore Mode

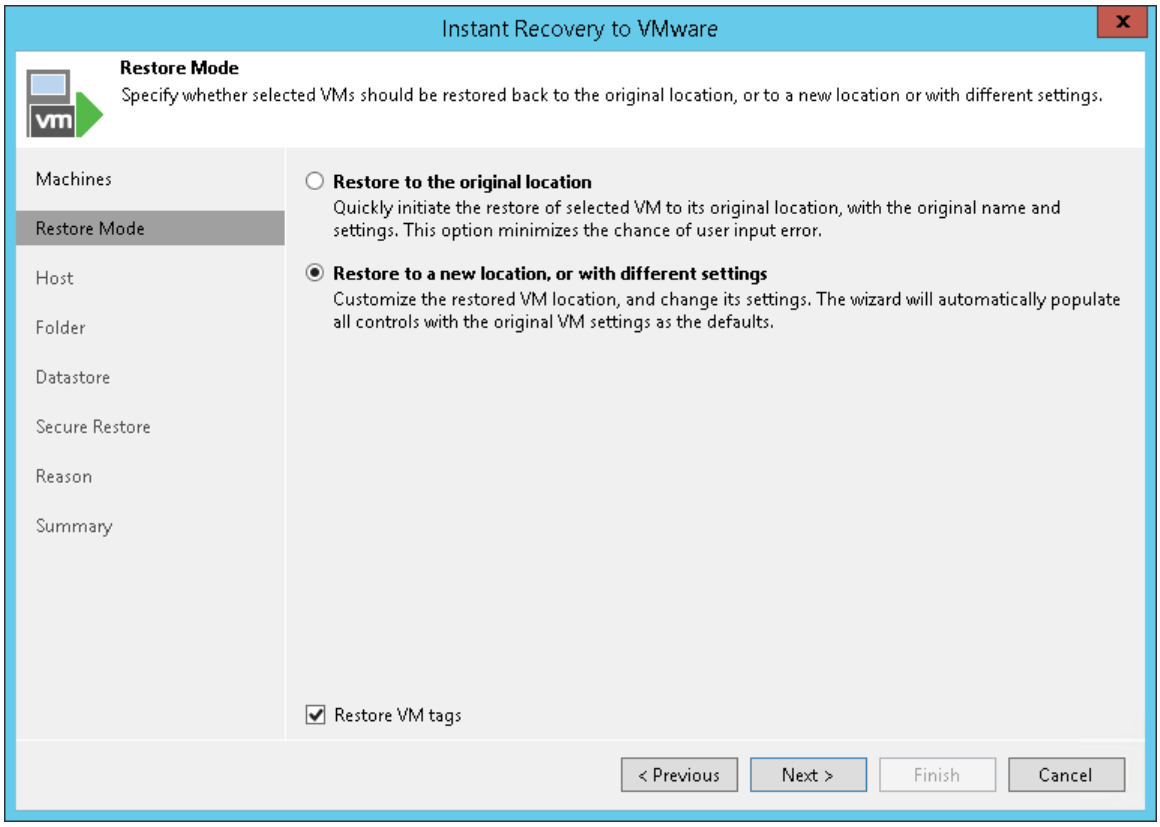
This step is available only if you recover VMware vSphere VMs or vCloud Director VMs.

At the **Restore Mode** step of the wizard, specify a destination for VM recovery and whether you want to recover VM tags:

1. Select a destination for recover:
   - **Restore to the original location** – select this option if you want to recover VMs with their initial settings and to their original location. If this option is selected, you will pass directly to the Reason step of the wizard.

     > **IMPORTANT**
     >
     > If you recover a VM with the original settings, and the original VM still exists in the virtual infrastructure, the original VM will be removed.

   - **Restore to a new location, or with different settings** – select this option if you want to recover VMs to a new location, or to any location but with different settings (such as VM location, network settings, format of recovered virtual disks and so on). If this option is selected, the **Instant Recovery to VMware** wizard will include additional steps for customizing VM settings.
2. If you want to recover tags that were assigned to the original VMs and assign them to the recovered VMs, select the **Restore VM tags** check box. Veeam Backup & Replication will recover the VMs with original tags if the following conditions are met:
   a. You recover VMs to their original location.
   b. The original VM tags are available on the source vCenter Server.



Step 5. Specify Destination for Restored VMs

Specify a destination where the recovered VMs will reside. The destination settings differ depending on the number of workloads that you recover:

- Specifying Destination for One VM
- Specifying Destination for Multiple VMs

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.

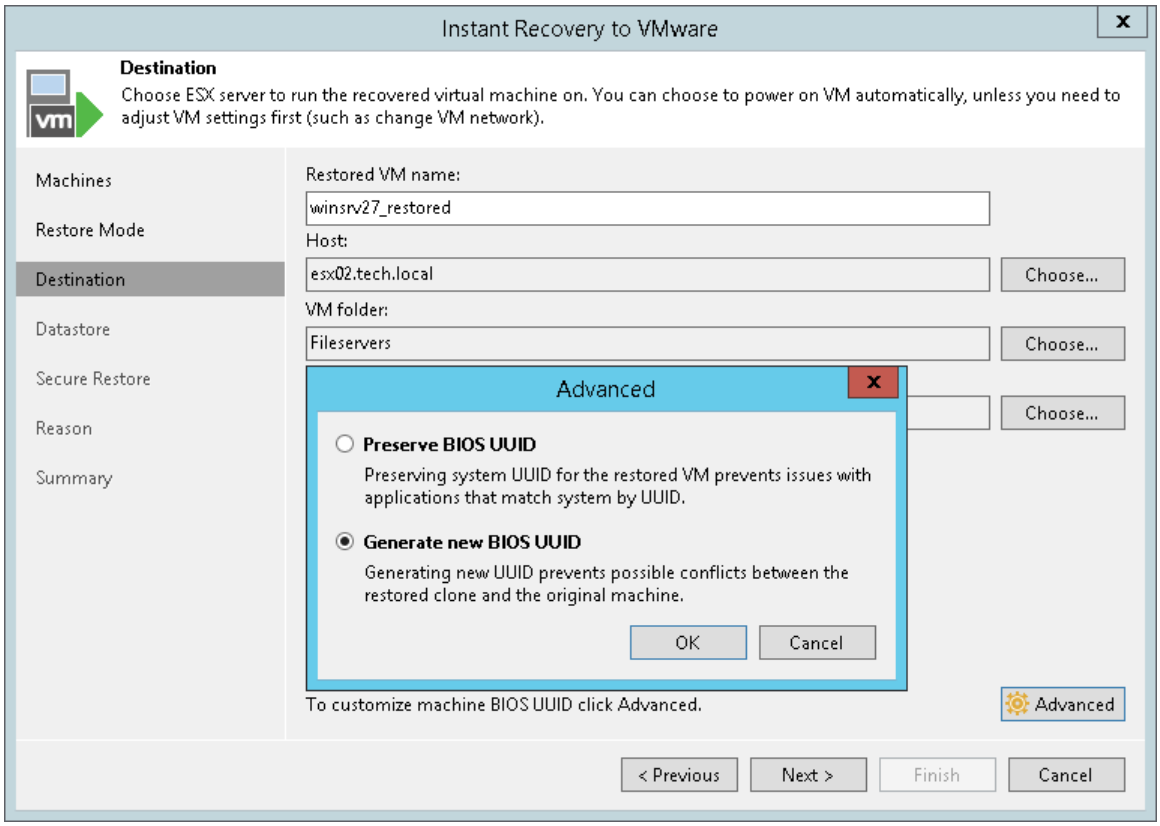
At this step of the wizard, you configure destination settings such as the recovered VM name, target host, VM folder and so on. For workloads other than VMware vSphere VMs and vCloud Director VMs, you must also configure a network mapping table. This table maps networks in the original site to networks in the target site. When the job starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update VM configuration files to replace the original networks with the specified networks in the target site. As a result, you will not have to re-configure network settings manually.

To specify the destination and network mapping settings for a recovered VM:

1. In the **Restored VM name** field, specify a name under which the workload will be recovered.
2. In the **Host** field, specify a host on which the VM will run.
3. In the **VM folder** field, specify a folder to which the recovered VM files will be placed.
4. In the **Resource pool** list, select a resource pool to which the VM will be placed.
5. [For workloads other than VMware vSphere VMs and vCloud Director VMs] In the **Networks** section, configure the network mapping table. To configure one row of the table, select a network in the list and click **Choose**. The **Select Network** window displays all networks to which the target host or cluster is connected. In the list of networks, select a network to which the recovered VM will be connected instead of the original network.

6. Click the **Advanced** button and choose whether to preserve the BIOS UUID or generate a new BIOS UUID.

   We recommend that you select to generate a new BIOS UUID for the recovered VM to prevent conflicts if the original workload still resides in the production environment. The BIOS UUID change is not required if the original VM no longer exists, for example, it was deleted.



Specifying Destination for Multiple VMs

The following steps are available if you recover multiple workloads and recover them to a new location or with different settings.

If you recover multiple workloads, specify the following settings for the destination:

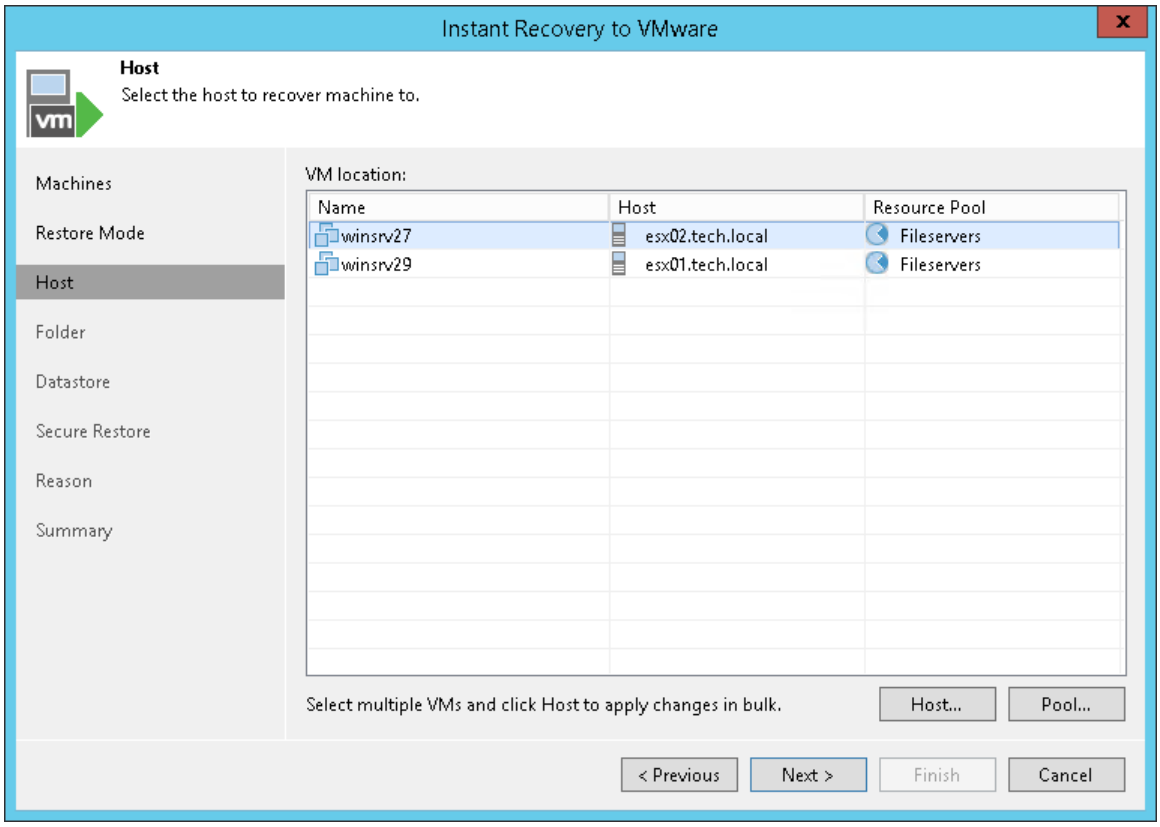
1. At the Host step of the wizard, select a target host.
2. At the Folder step of the wizard, specify VM settings.
3. [For workloads other than VMware vSphere VMs and vCloud Director VMs] At the Network step of the wizard, specify network settings.

Selecting Target Host

At the **Host** step of the wizard, specify a target host and resource pool for recovered VMs:

1. In the list, select the necessary VMs and click the **Host** button.
2. From the virtual environment, select a standalone or clustered host where the selected VMs will be registered.
3. Select one or multiple VMs and click the **Pool** button.

4. In the list, select a resource pool where the selected VMs will be stored.



Specifying VM Settings

For each recovered VM, you can change a VM name, BIOS UUID and folder where VM files must be stored. It is recommended that you specify a new name and generate a new BIOS UUID to prevent conflicts if the original workload still resides in the production environment. The name and BIOS UUID change is not required if the original workload no longer exists, for example, it was permanently deleted.

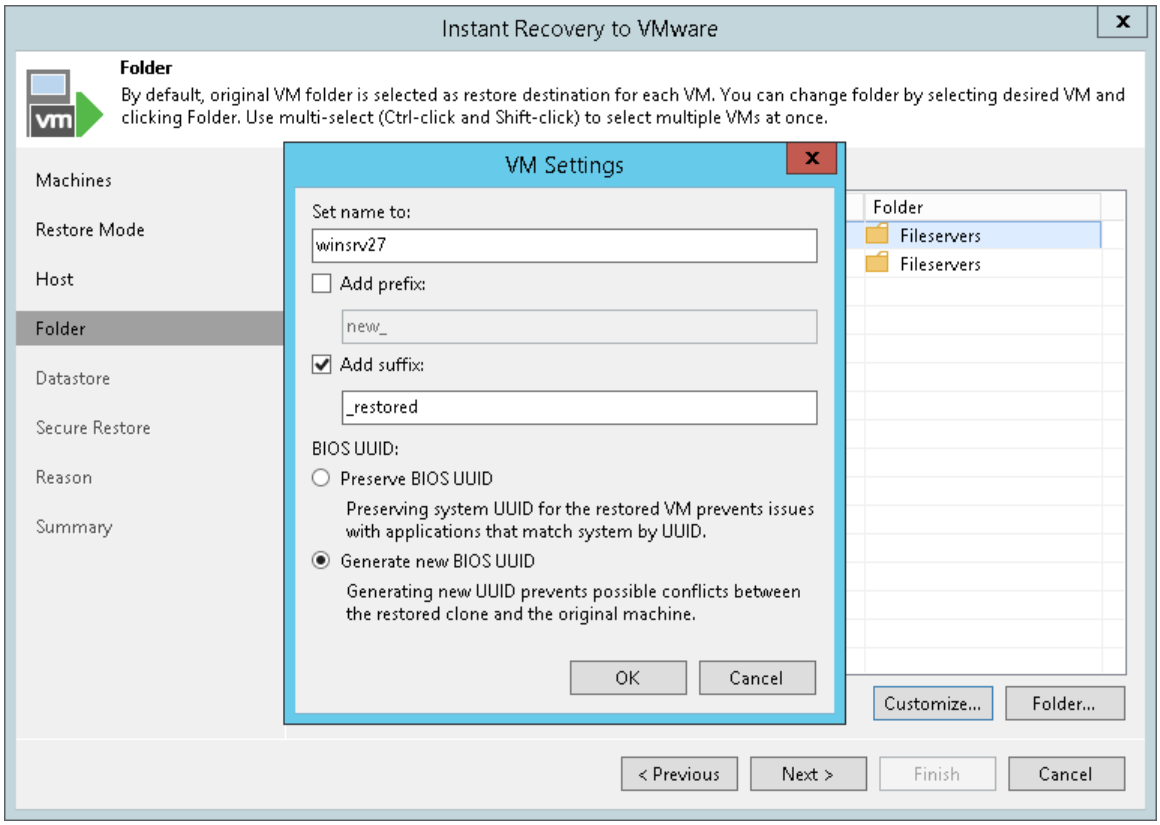
To specify a folder:

1. In the **VM settings** list, select the necessary VMs.
2. Click the **Folder** button.
3. Select a folder where VM files must be stored.

To change a recovered VM name and BIOS UUID:

1. In the **VM settings** list, select one VM.
2. Click the **Customize** button.
3. In the **VM Settings** window, do the following:
   - In the **Set name to** field, specify a new VM name.
   - To add a prefix and suffix to the name specified in the **Set name to** field, select **Add prefix** and **Add suffix** check boxes.

- In the **BIOS UUID** section, specify whether to preserve or generate a new BIOS UUID.



Specifying Network Settings

This step is available if you recover workloads other than VMware vSphere VMs and vCloud Director VMs.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the original site to networks in the target site (site where VMs will be recovered). When the job starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update VM configuration files to replace the original networks with the specified networks in the target site. As a result, you will not have to re-configure network settings manually.

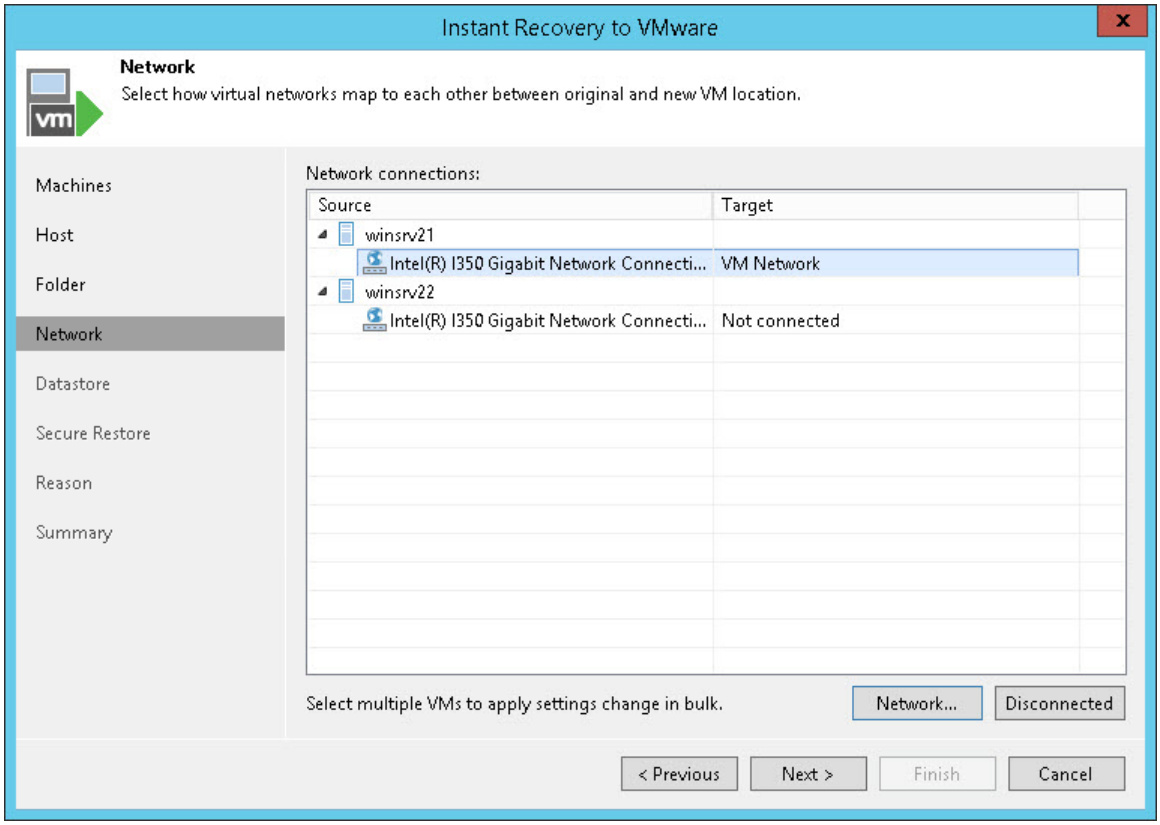
To change networks to which the recovered VMs will be connected:

1. In the list, select one or multiple workloads and click the **Network** button.

   If a workload is connected to multiple networks, you can select a network to map and click **Network**.

2. The **Select Network** window displays all networks to which the target host or cluster is connected. In the list, select a network to which the recovered VM will be connected after recover.

If you do not want to connect a recovered VM to any virtual network, select the original workload and click **Disconnected**.



Step 6. Select Destination for Virtual Disk Updates

This step is available if you recover workloads to a new location or with different settings.

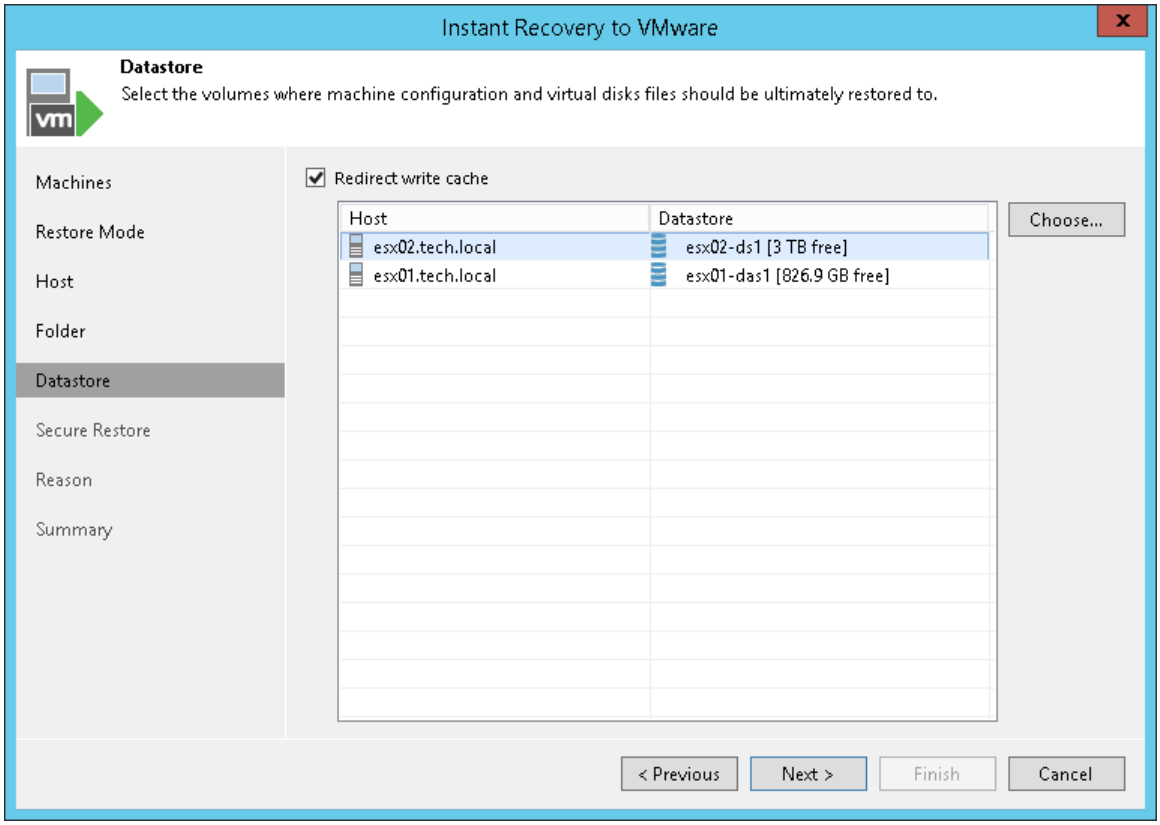
At the **Datastore** step of the wizard, you can select where to store redo logs when a VM is running from a backup. Redo logs are auxiliary files used to keep changes that take place while the recovered VM runs.

By default, redo logs are stored on the vPower NFS server. You can store redo logs on any datastore in the virtual environment if necessary. Redirecting redo logs improves recovery performance but makes Storage vMotion not possible for ESXi 5.5. 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.
2. Click **Choose** and select a datastore from the list.

> **IMPORTANT**
>
> If the size of recovered VM disks is greater than 2 TB, you must not place redo logs on a VSAN datastore. Otherwise, Veeam Backup & Replication will fail to create a snapshot for the recovered VMs. For more information, see VMware Docs.



Step 7. Configure Helper Appliance

This step is available if you recover workloads with Linux OS, recover them to a new location or with different settings, and VIX API is not available.

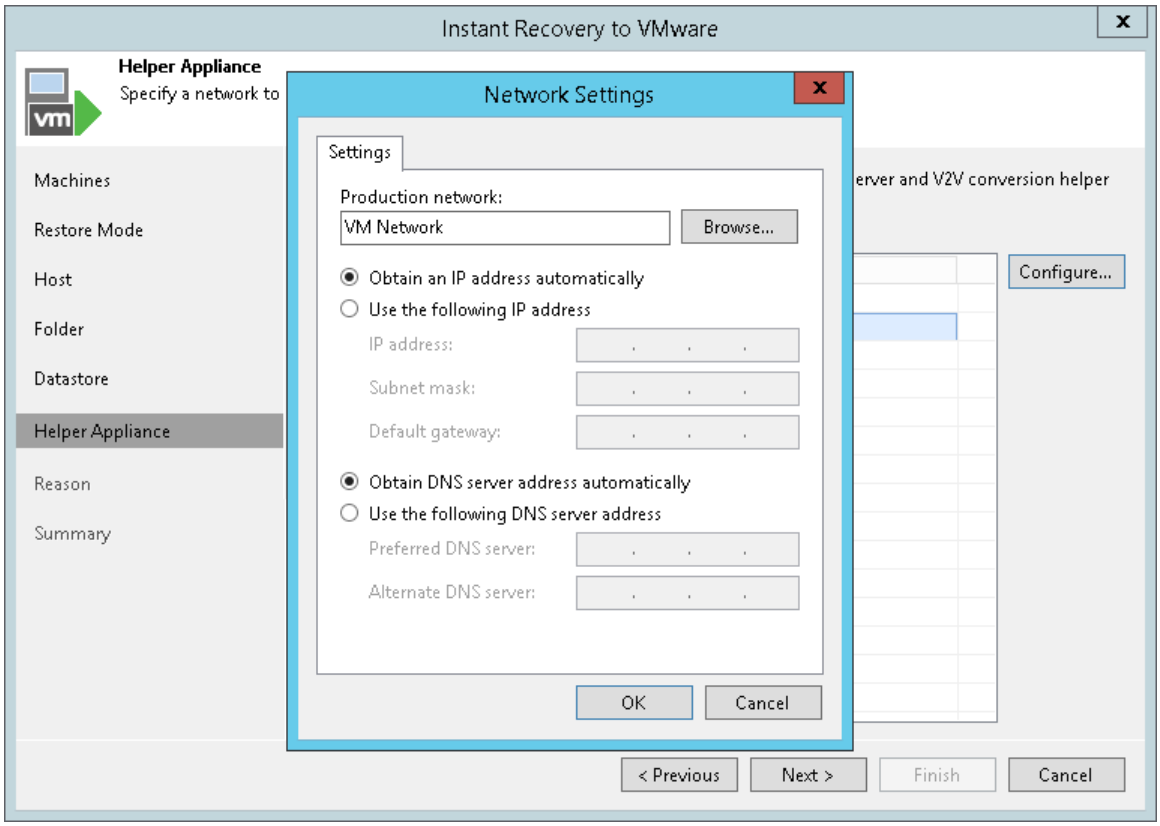
Veeam Backup & Replication recovers Linux machines to a standalone host using a helper appliance. The helper appliance is an auxiliary Linux-based VM registered by Veeam Backup & Replication. The appliance is quite small — around 50 MB. It requires 1024 MB RAM and takes around 10 seconds to boot.

To configure the helper appliance:

1. [For multiple machines] In the **Network** list, expand a host and select one machine.
2. Click the **Configure** button.
3. In the **Network Settings** window, select a network for the helper appliance.
   a. Click the **Browse** button near the **Production network** field.
   b. In the **Select Network** window, Veeam Backup & Replication shows a list of networks to which the target host is connected. In this list, select a network to which the helper appliance must be connected.

      Mind that the backup server and the mount server must have access to the helper appliance over the network.
4. Specify IP address settings for the helper appliance:
   - If you use a DHCP server in the network and want to obtain IP address for the helper appliance automatically, make sure that the **Obtain an IP address automatically** option is selected.
   - To manually specify the IP address of the helper appliance, select the **Use the following IP address** option and enter the IP address, subnet mask and default gateway address in the corresponding fields.
5. Specify DNS server settings:
   - If you use a DHCP server in the network and the IP address for the DNS server is obtained automatically, make sure that the **Obtain DNS server address automatically** option is selected.

- To manually specify the IP address of the DNS server, select the **Use the following DNS server address** option and enter the IP addresses of the preferred and alternate DNS servers in the corresponding fields.



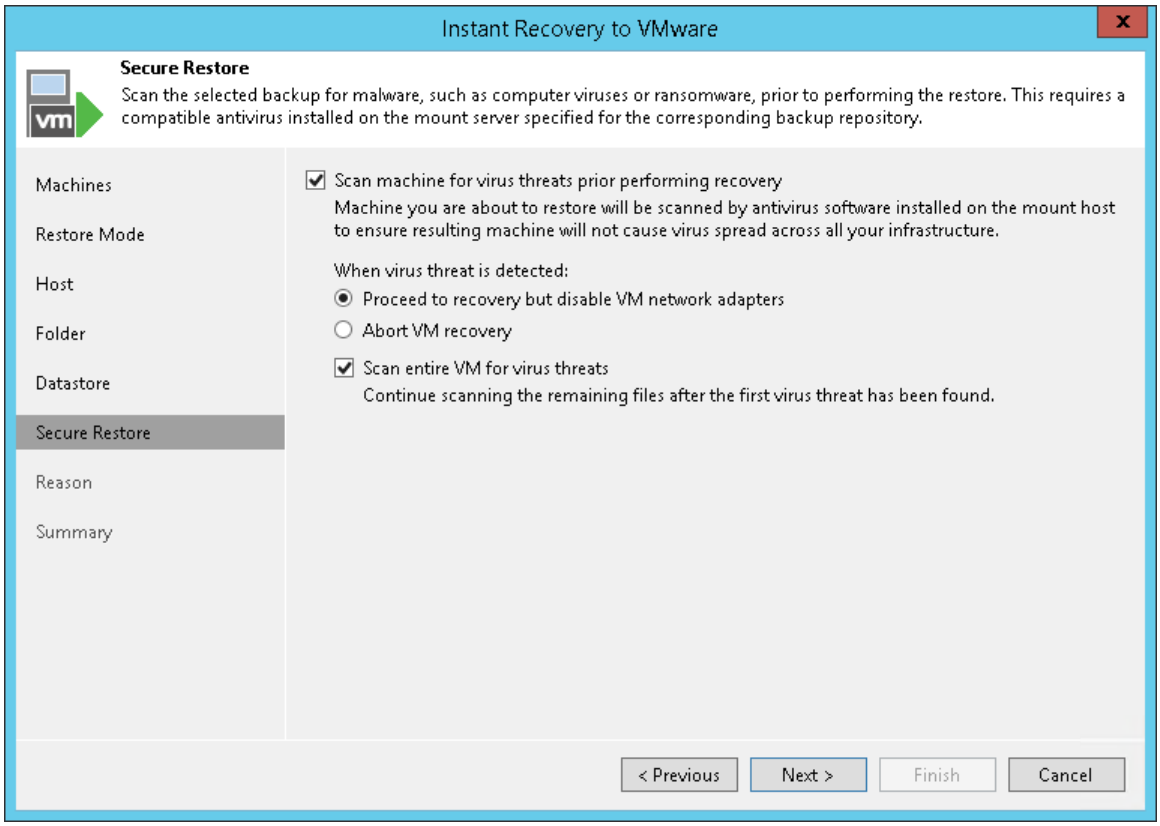
Step 8. Specify Secure Restore Settings

This step is available if you recover workloads with Microsoft Windows OS and recover them to a new location or with different settings.

You can instruct Veeam Backup & Replication to perform secure restore — scan restored VM data with antivirus software before restoring VMs to the production environment. For more information on secure restore, see Secure Restore.

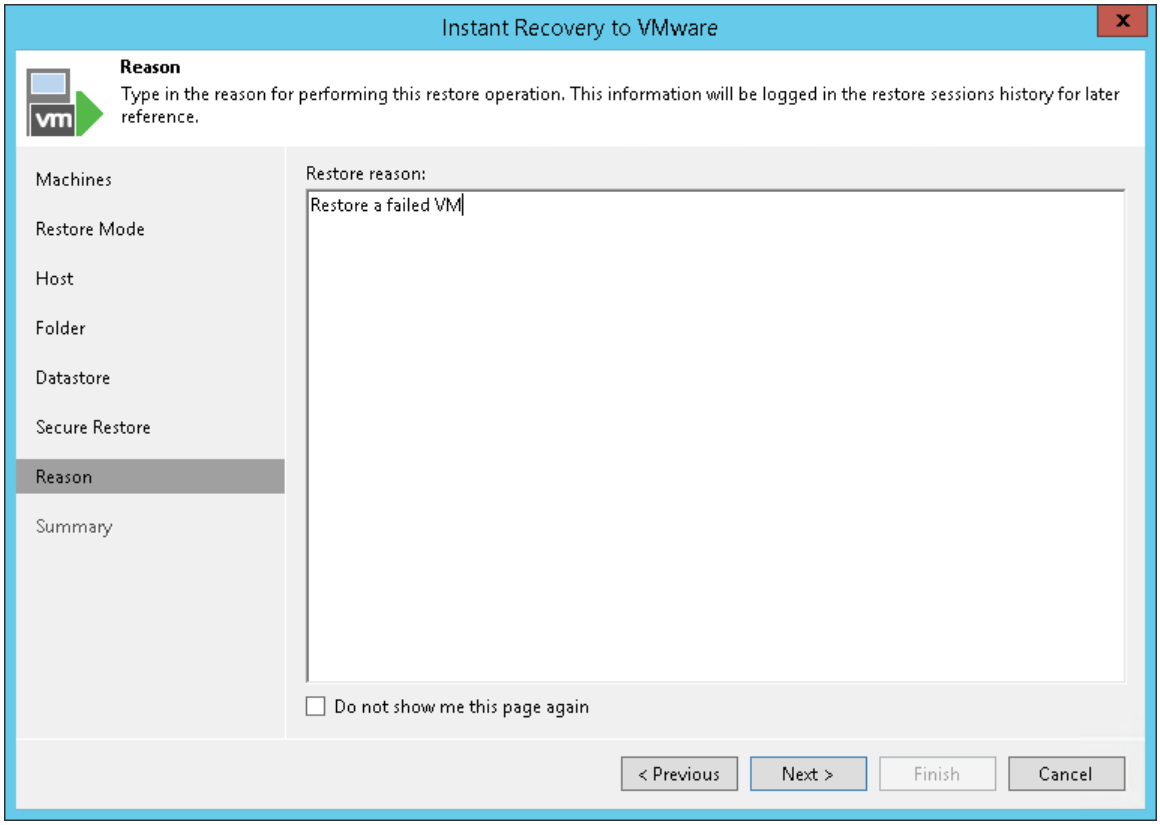
To specify secure restore settings:

1. At the **Secure Restore** step of the wizard, select the **Scan machine for virus threats prior performing recovery** check box.
2. Select which action Veeam Backup & Replication will take if an antivirus software finds a virus threat:
   - **Proceed to recovery but disable VM network adapters** — if you want to restore VMs with disabled network adapters (NICs).
   - **Abort VM recovery** — if you want to cancel the restore session.
3. Select the **Scan entire VM for virus threats** check box if you want the antivirus software to continue scanning VM data after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



Step 9. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for performing Instant Recovery. The information you provide will be saved in the session history and you can reference it later.



Step 10. Verify Instant Recovery Settings

At the **Summary** step of the wizard, specify additional settings for Instant Recovery:

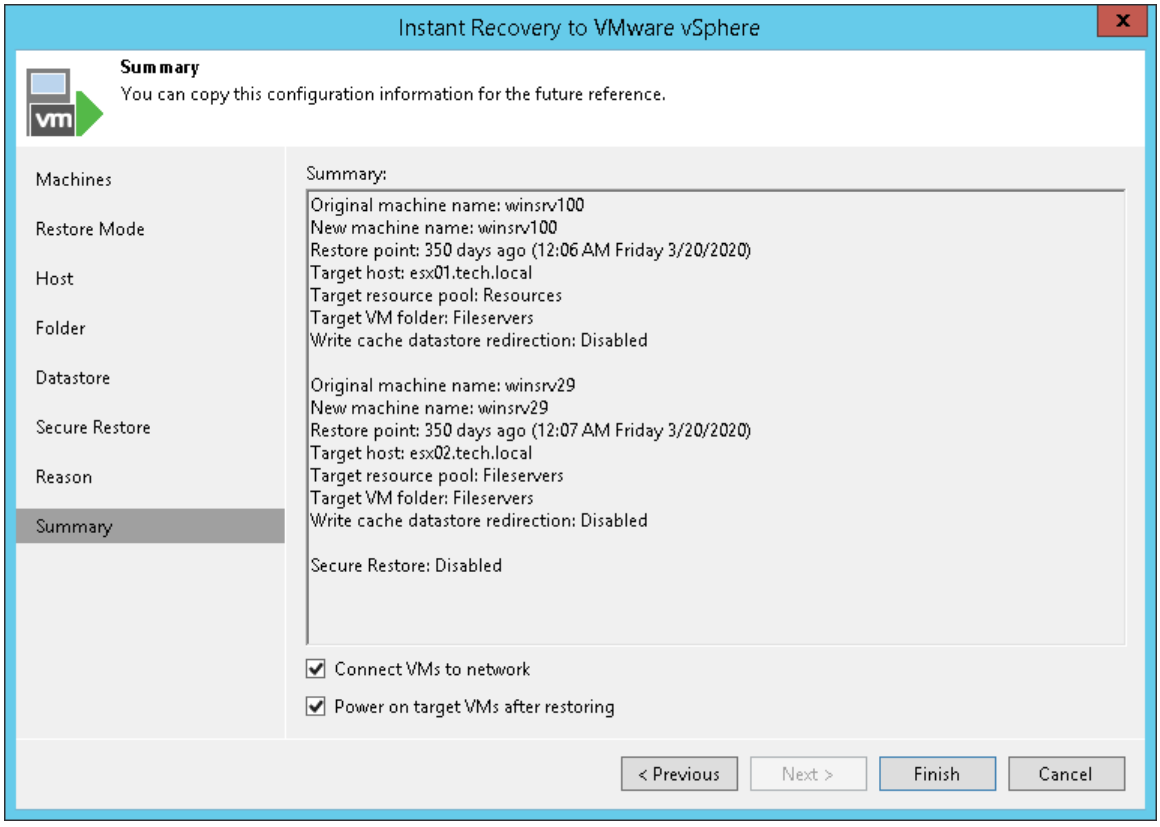
1. If you recover production workloads that have failed and want to recover them with initial network settings, select the **Connect VMs to network** check box.

   If you recover workloads for testing disaster recovery while the original workloads are still running, leave this check box unselected. Before you power on the recovered VMs, you must disconnect them from the production network and connect to a non-production network to avoid conflicts.

2. To start the VMs right after recovery, select the **Power on target VMs after restoring** check box. If you recover the workloads to the production network, make sure that the original workloads are powered off.

3. Check settings that you have specified for Instant Recovery and click **Finish**.

4. Check that the publishing process has started and click **Close**.



What You Do Next

Finalizing Instant Recovery to VMware vSphere

Finalizing Instant Recovery to VMware vSphere

After the VMs have been successfully recovered, you must finalize the process. For this, test the recovered VMs and decide whether to migrate them to production environment or stop publishing.

- Testing recovered VMs
- Migrating recovered VMs

- Stop publishing recovered VMs

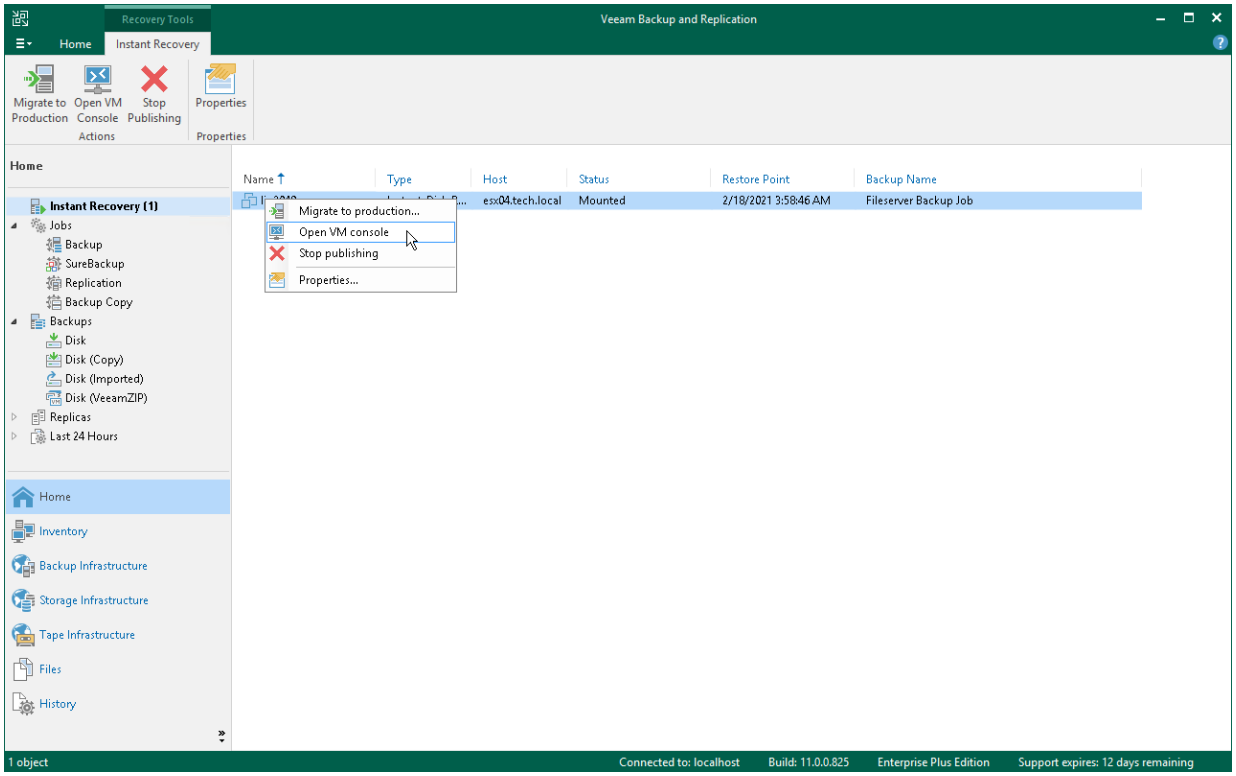
Testing Recovered VMs

To test the recovered VMs before you migrate them to production, you can launch VMware Remote Console software from the Veeam Backup & Replication console.

> **IMPORTANT**
>
> Before you launch VMware Remote Console, make sure that this software is installed on the machine where the Veeam Backup & Replication console runs.

To open a VM console in Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click a VM and select **Open VM console**.



Migrating Recovered VMs

If the VMs were recovered successfully, you can migrate them to the production environment.

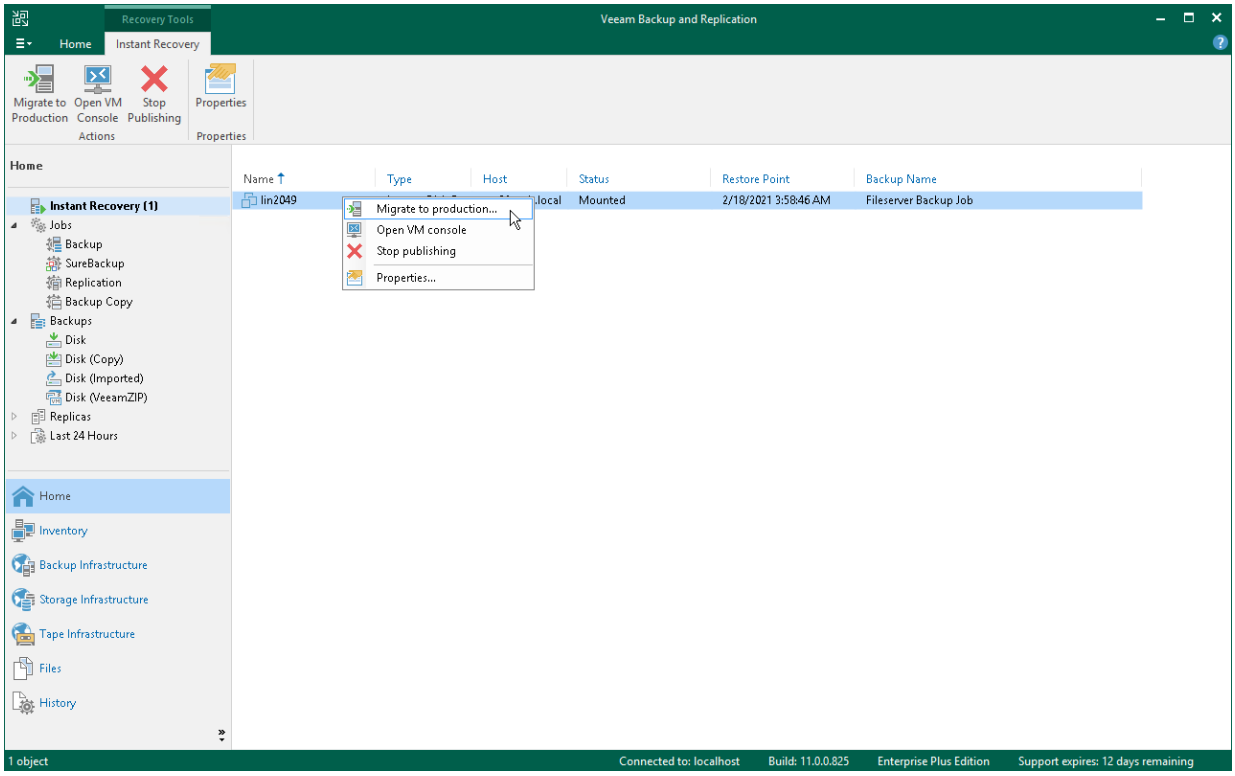
To migrate a recovered VM to the production environment:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Migrate to production**. Veeam Backup & Replication will launch the Quick Migration wizard.

   During migration, Veeam Backup & Replication will recover the VM from the backup file and additionally move all changes that were made while the VM was running from the backup in the Instant Recovery mode.

If you have launched Instant Recovery to a different location and you want to protect the recovered VM after migration finishes, you need to add the recovered VM to a backup job manually. If you have launched Instant Recovery to the original location, your actions depend on the method used for migration and whether the **Delete source VM files upon successful migration** check box is enabled in the migration wizard. For more information, see Finish Working with the Quick Migration Wizard.



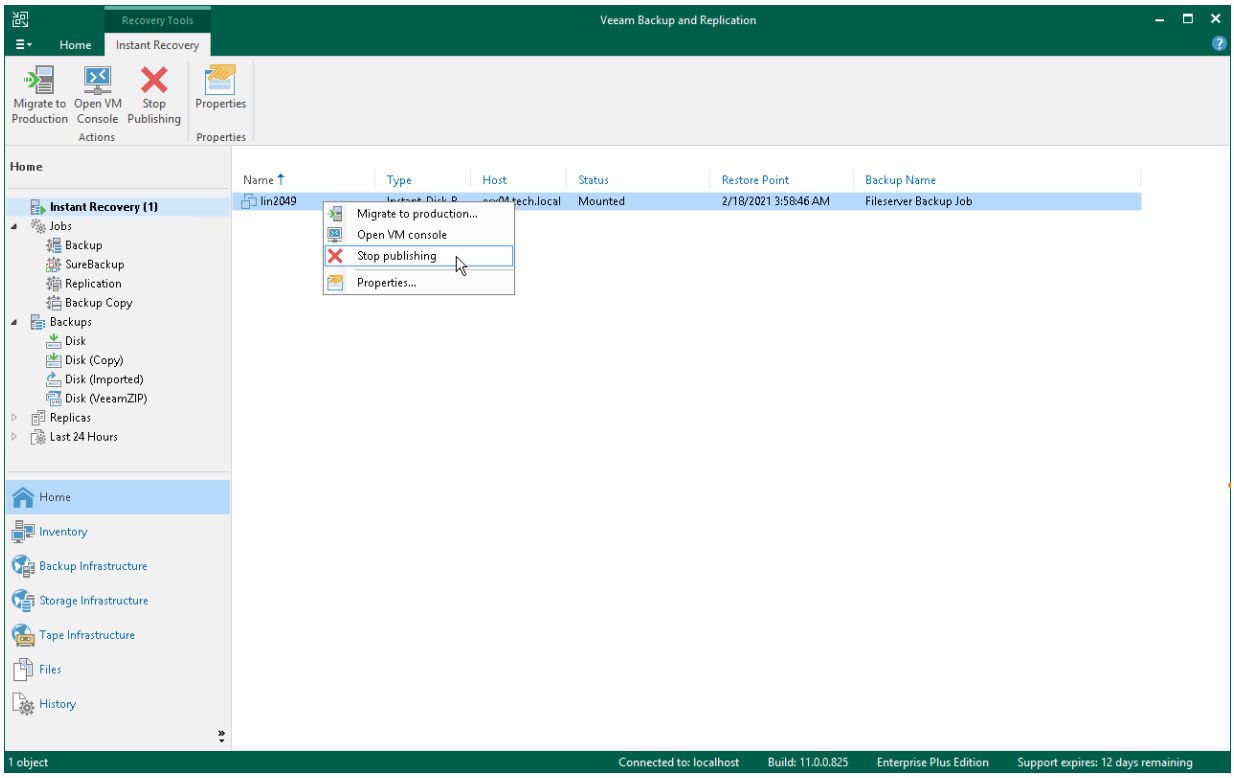
Stop Publishing Recovered VMs

If your tests have failed, you can stop publishing the recovered VMs. This will remove the recovered VMs from the host that you selected as the destination for recovery. Note that all changes made in the recovered VMs will be lost.

To stop publishing a recovered VM:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Stop publishing**.



Instant Recovery to Microsoft Hyper-V

With Instant Recovery to Microsoft Hyper-V, you can immediately recover different workloads (VMs, EC2 instances, physical servers and so on) as Microsoft Hyper-V 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.

When you perform Instant Recovery, Veeam Backup & Replication creates dummy VMs and mounts to VMs workload disks directly from backups stored on backup repositories. These dummy VMs have limited I/O performance. To provide full I/O performance, you must migrate the VMs to the production site. For more information, see Migration of Recovered VMs to Production Site.

Besides disaster recovery matters, Instant Recovery can also be used for testing purposes. Instead of extracting workloads to production storage to perform regular disaster recovery (DR) testing, you can run a workload directly from a backup file, boot it and make sure the guest OS and applications are functioning properly. For more information, see Finalizing Instant Recovery to Microsoft Hyper-V.

Instant Recovery supports bulk processing so you can immediately recover multiple workloads at once. If you perform Instant Recovery for several workloads, Veeam Backup & Replication uses the resource scheduling mechanism to allocate and use optimal resources required for Instant Recovery. For details, see Resource Scheduling.

Supported Backup Types

You can recover workloads from the following types of backups:

- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of VMware vSphere virtual machines created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud Platform
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of RHV VMs created by Veeam Backup for RHV

How Instant Recovery Works

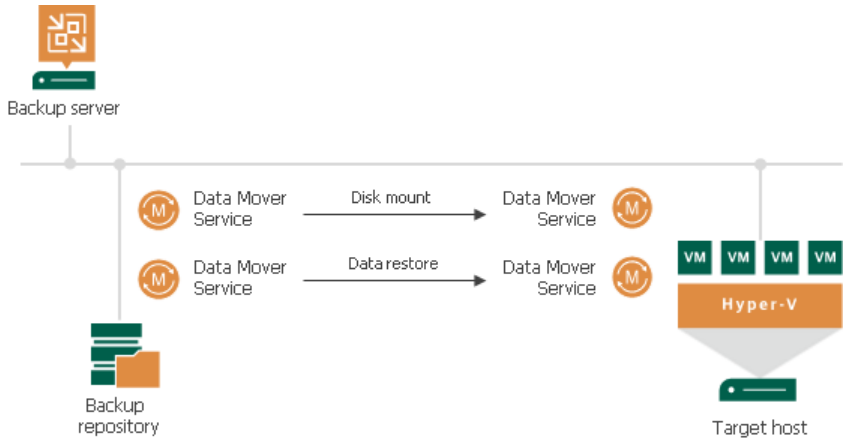
Instant Recovery is performed in the following way:

1. Veeam Backup & Replication reads the workload configuration from the backup file in the backup repository and creates a dummy VM with empty disks on the target host. The created VM has the same settings as the workload in the backup file. Note that Veeam Backup & Replication pre-allocates disk space required for the recovered VM at the beginning of the Instant Recovery process.
2. Veeam Backup & Replication initiates creation of a protective snapshot for the dummy VM and the VM is started. If the Instant Recovery process fails for some reason, the protective snapshot guarantees that no data is lost.

3. On the backup repository and on the target host, Veeam Backup & Replication starts a pair of Veeam Data Movers that are used to mount the VM disks from the backup file to the dummy VM.
4. On the target host, Veeam Backup & Replication starts a proprietary Veeam driver. The driver redirects requests to the file system of the recovered VM (for example, when a user accesses some application) and reads necessary data from the backup file in the backup repository using the pair of Veeam Data Movers that maintain the disk mount.

Migration of Recovered VMs to Production Site

When you begin the migration process, Veeam Backup & Replication starts one more pair of Veeam Data Movers — one Veeam Data Mover on the backup repository and one on the target host. This pair of Veeam Data Movers copies data of the recovered VM from the backup repository to the target host in the background, and populates disks of the VM started on the target host.



The driver on the target host knows which data has already been recovered permanently and does not redirect requests to such data, reading it directly from the disks of the recovered VM. Thus, performance of the instantly recovered VM will increase as more of the data is copied. When the VM is recovered completely, all Veeam Data Movers are stopped.

Performing Instant Recovery to Microsoft Hyper-V

With Instant Recovery, you can recover different workloads from backups and register them as Microsoft Hyper-V VMs. For the list of backups that you can use for recovery, see Supported Backup Types.

To perform Instant Recovery, use the **Instant Recovery** wizard.

> **NOTE**
>
> If you want to recover workloads as VMware vSphere VMs, see Performing Instant Recovery to VMware vSphere.

Before You Begin

Before you perform Instant Recovery to Microsoft Hyper-V, consider the following:

- You must add the Hyper-V target host to which you want to recover machines to your backup infrastructure.

- The following configuration is not supported:
  - The storage I/O balancing is enabled on a Hyper-V host. Starting from Microsoft Windows Server 2012, I/O balancing is enabled by default.
  - The backup repository role is assigned to a VM on this Hyper-V host. The backup from which you recover is stored on this repository.
  - You plan to recover a VM to this Hyper-V host or SureBackup is targeted to this host.

  For information on how to recover with such configuration, see this Veeam KB article.
- Make sure that the **Disable changed block tracking for this host** option is not selected for a host to which you plan to recover a workload. If this option is selected for the host, the driver required for work of Instant Recovery will be disabled. For more information, see the Configuring Connected Volumes section in the User Guide for Microsoft Hyper-V.
- You can recover a workload from a backup that has at least one successfully created restore point.
- If you recover a workload to the production network, make sure that the original workload is powered off to avoid conflicts.
- We strongly recommend to have `dracut` and `mkinitrd` installed on Linux workloads that will be restored. Otherwise, they may not boot after restore.
- [If you recover Hyper-V VMs] The version of the target host on which a VM is recovered must be the same or later than the version of the source host where the original VM was registered.

  For example, you can restore a VM from the host that runs *Microsoft Windows Server 2016* to the target host that runs Microsoft Windows Server 2016 (including version 1809), Microsoft Windows Server 2019 or later.

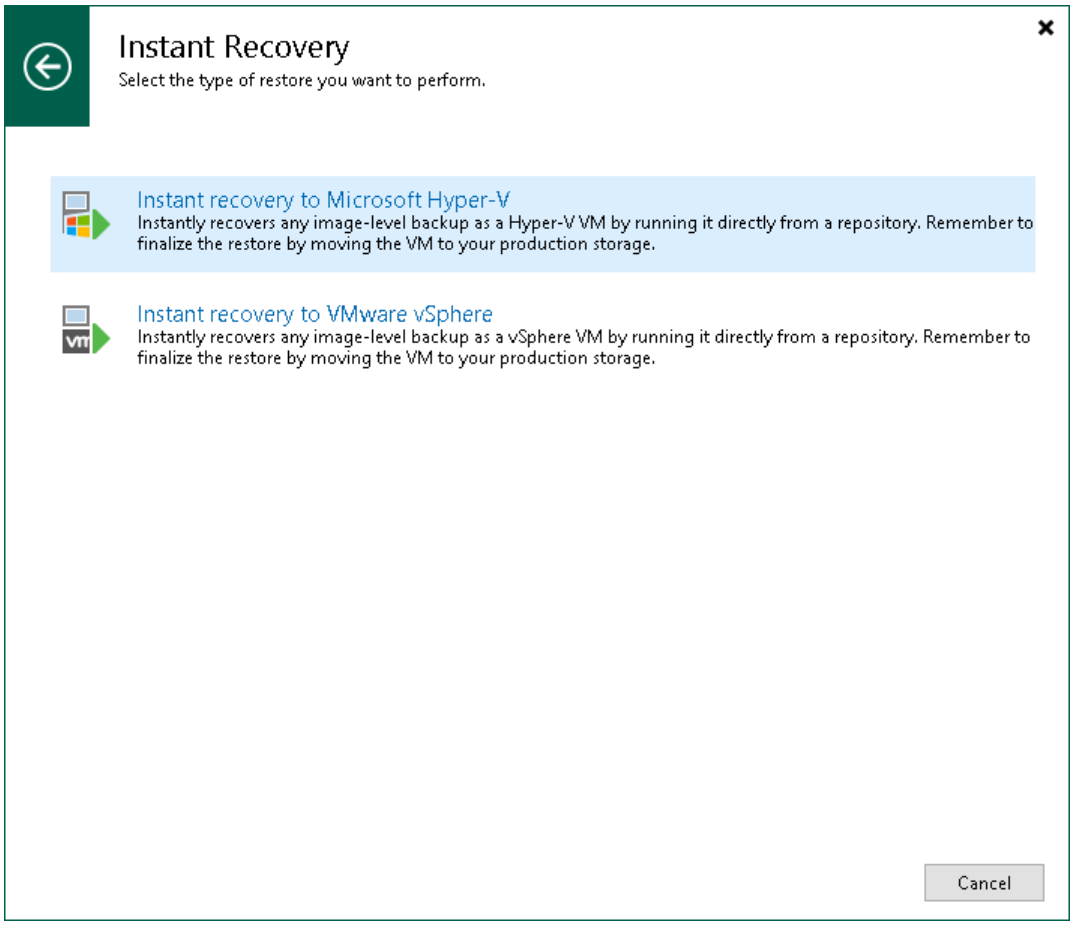
  The Hyper-V role must be enabled on both source and target hosts.
- If you want to scan recovered VM data for viruses, check the secure restore requirements and limitations.
- On non-Microsoft Windows SMB3 storage, for example, Tintri, Veeam Backup & Replication may display the "*Failed to disable integrity bit on disk N*" warning during the restore process. You can ignore this warning for non-Microsoft Windows SMB3 storage.
- The recovered VM will have the same MAC address as the original workload. Therefore, if you recover the workload to the same Hyper-V host where the original workload is running, a MAC address conflict may occur. To overcome this situation, power off the original workload before you start the recovery process.
- [For Nutanix AHV VMs] The recovered VM will not be connected to a network. You must connect to the network manually.
- [For Nutanix AHV VMs, Amazon EC2 instances and Microsoft Azure virtual machines] Instantly recovered VM will have default virtual hardware settings: 2 CPU cores, 4GB RAM and one network adapter. If you want to change the default settings, turn off the VM and set the required virtual resources. Note that you must not switch off the instant recovery session before turning off the VM.

Step 1. Launch Instant Recovery Wizard

To launch the **Instant Recovery to Hyper-V** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select one of the following:
  - **Microsoft Hyper-V** > **Restore from backup** > **Entire VM restore** > **Instant Recovery** — if you want to recover Hyper-V VMs from a VM backup created by Veeam Backup & Replication.
  - **VMware vSphere** > **Restore from backup** > **Entire VM restore** > **Instant Recovery to Microsoft Hyper-V** — if you want to recover VMware vSphere VMs from a VM backup created by Veeam Backup & Replication.
  - **Agent** > **Entire machine restore** > **Instant Recovery to Microsoft Hyper-V** — if you want to recover physical machines from a backup created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.
  - **AWS EC2 backup** > **Entire machine restore** > **Instant Recovery to Microsoft Hyper-V** — if you want to recover EC2 instances from a backup created by Veeam Backup for AWS.
  - **Azure IaaS** > **Entire machine restore** > **Instant Recovery to Microsoft Hyper-V** — if you want to recover Azure VMs from a backup created by Veeam Backup for Microsoft Azure.
  - **GCE backup** > **Entire machine restore** > **Instant Recovery to Microsoft Hyper-V** — if you want to recover VM instances from a backup created by Veeam Backup for Google Cloud Platform.
  - **Nutanix backup** > **Entire machine restore** > **Instant Recovery to Microsoft Hyper-V** — if you want to recover VMs from backups created by Veeam Backup for Nutanix AHV.
  - **RHV Proxy** > **Entire machine restore** > **Instant Recovery to Microsoft Hyper-V** — if you want to recover VMs from backups created by Veeam Backup for RHV.

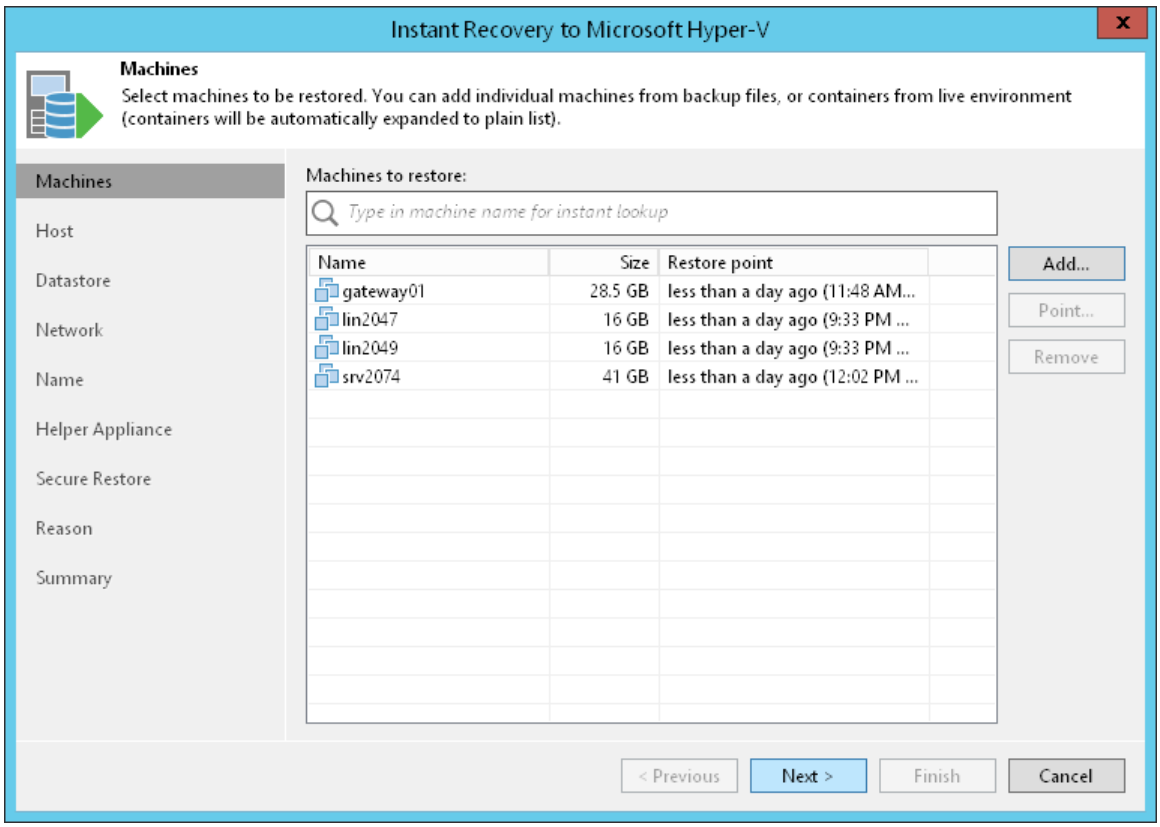
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, select workloads that you want to recover and click **Instant Recovery** > **Microsoft Hyper-V** on the ribbon. Alternatively, you can right-click one of the selected workloads and select **Instant recovery** > **Microsoft Hyper-V**.



Step 2. Select Workloads

At the **Machines** step of the wizard, select workloads that you want to recover:

1. Click **Add**.
2. In the **Backup Browser** window, do the following:
   a. [For Hyper-V, VMware vSphere VMs and vCloud Director VMs] You can browse for necessary VMs using one of the following way:
      - **From infrastructure** – use this option to browse the virtual environment and select VMs or VM containers to recover. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
      - **From backup** – browse existing backups and select VMs under backup jobs.
   b. [For other workloads] In the list of backup jobs, expand a job and select workloads.
   c. Click **Add**.

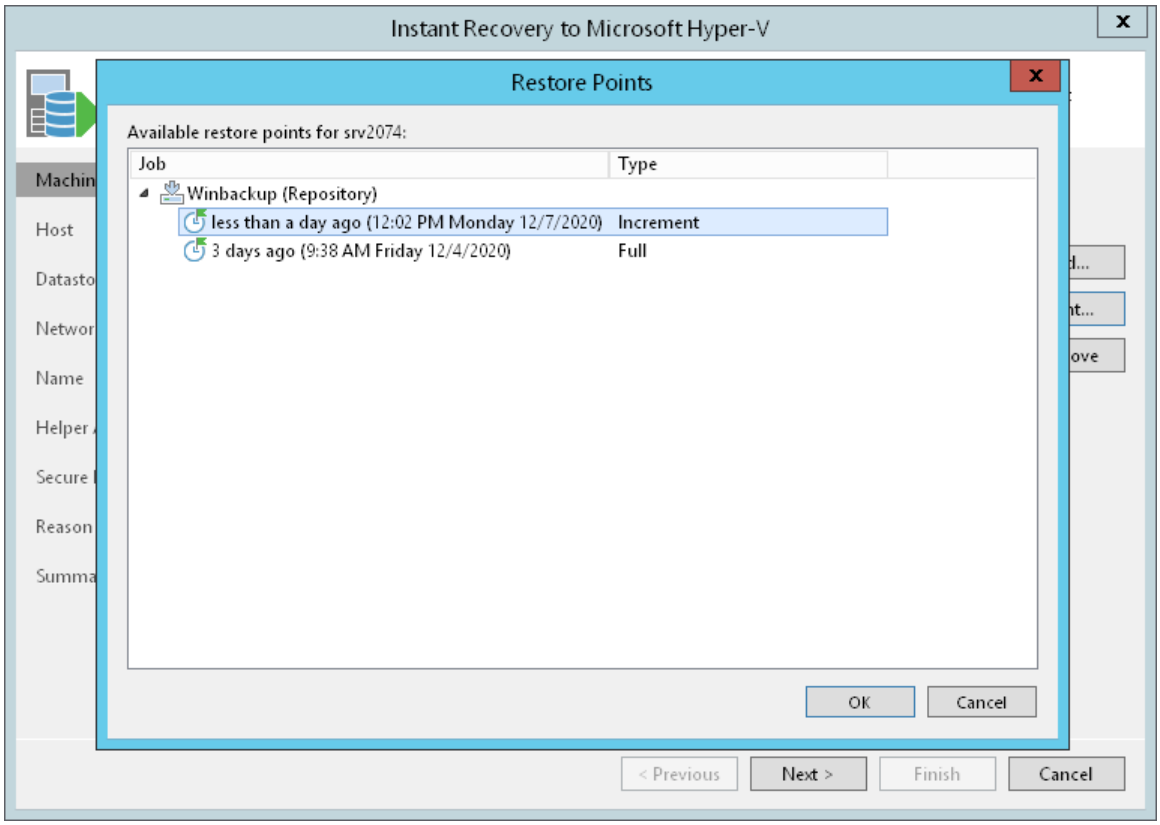


Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore points to recover the workloads. You can recover a workload to an earlier state, if necessary. If you have chosen to recover multiple workloads, you can select a restore point for each workload in the list.

To select a restore point:

1. In the **Machines to restore** list, select a workload.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point from which you want to recover the workload.



Step 4. Select Recovery Mode

This step is available only if you recover Hyper-V VMs.

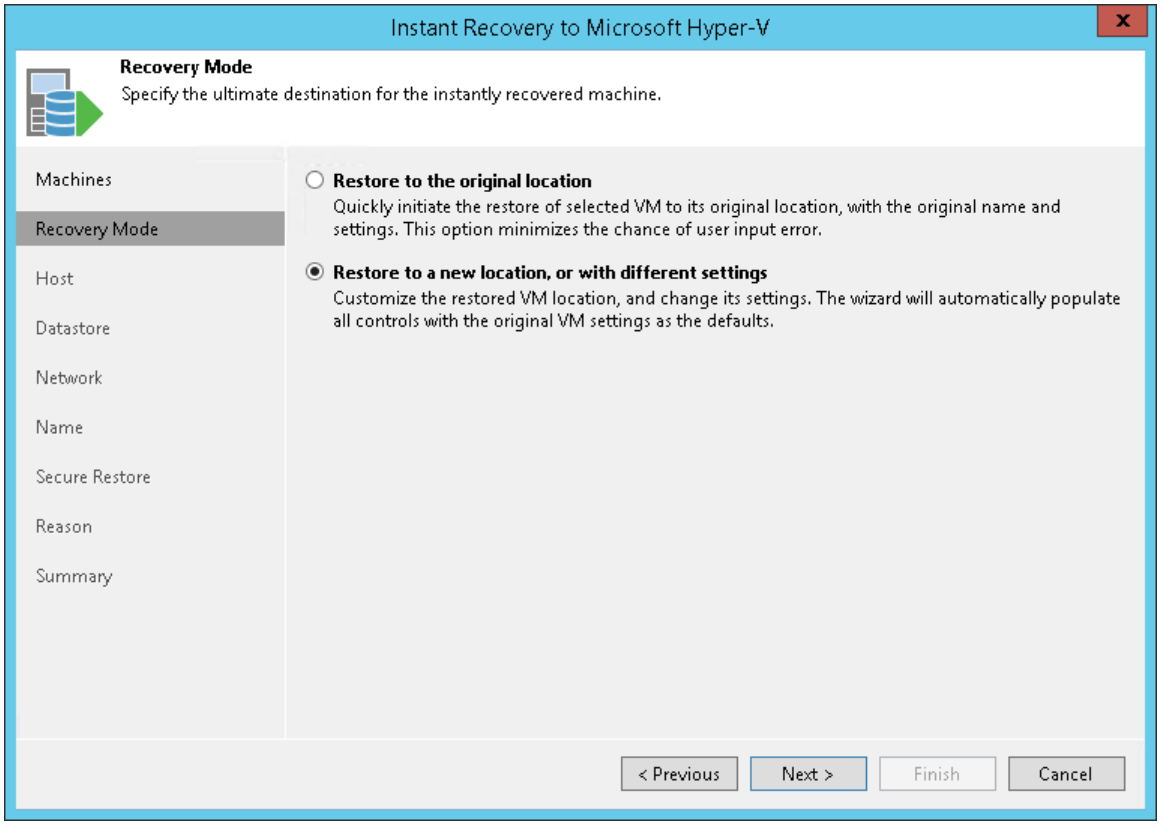
At the **Recovery Mode** step of the wizard, choose the necessary restore mode:

- Select **Restore to the original location** if you want to recover VMs with initial settings and to original location. If this option is selected, you will pass directly to the Reason step of the wizard.
- Select **Restore to a new location, or with different settings** if you want to recover VMs to a different location and/or with different settings (such as location, network settings and so on). If this option is selected, the **Instant Recovery to Hyper-V** wizard will include additional steps for customizing VM settings.

IMPORTANT

If you recover a VM to the original location, consider the following:

- If the original VM still exists in the virtual infrastructure, its disks will be removed. Make sure that other VMs in the virtual environment do not use these disks.
- The VM settings contain the ID of the VM group to which the machine belongs. To recover the VM to the original VM group, you must not delete the original VM group or change the hierarchy of its parent VM groups.



Step 5. Select Target Host

This step is available if you recover workloads to a new location or with different settings.

At the **Host** step of the wizard, specify a target Hyper-V host or a Hyper-V cluster where you want to locate the recovered VMs:

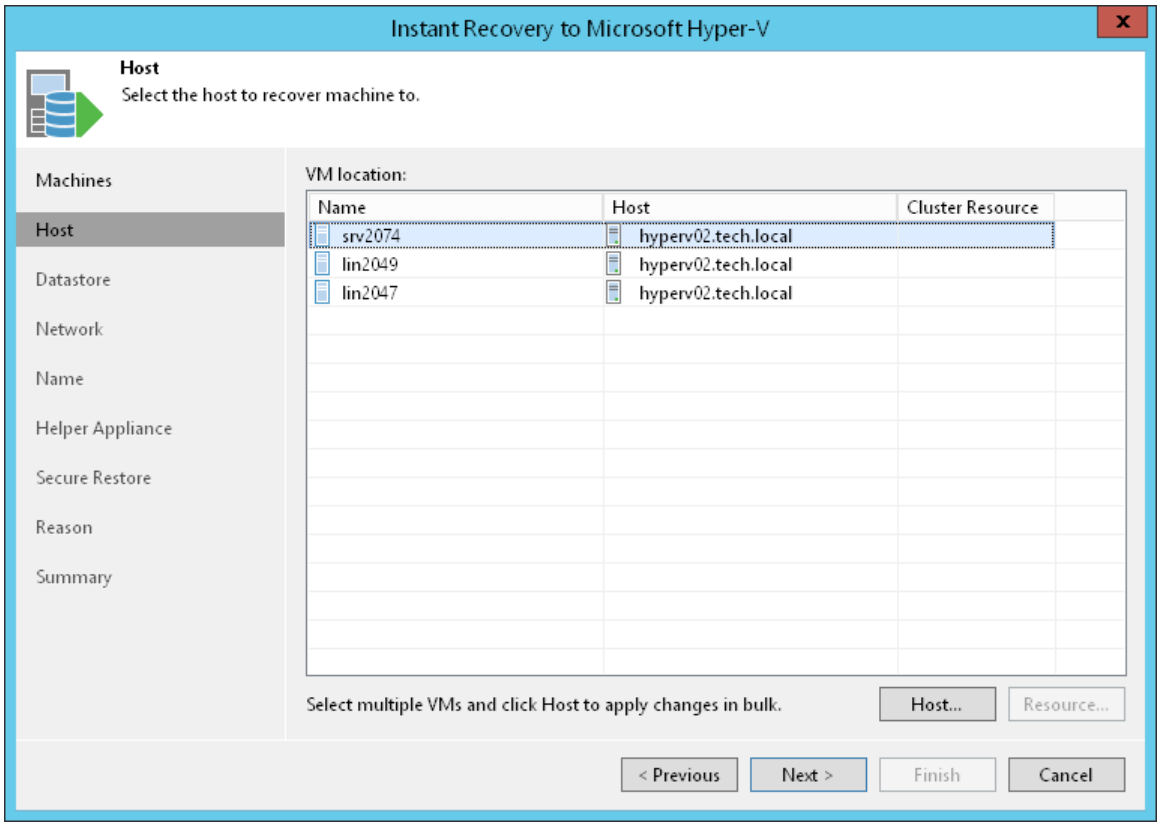
1. In the **VM location** list, select the necessary VMs and click **Host**.
2. In the **Select Server** window, select a standalone host or cluster where the selected VM will be registered.
3. If you have selected a Hyper-V cluster, you can specify the cluster resource settings. Click **Resource** and select one of the following options in the **Cluster Resource Settings** window:
   - **Register VM as a cluster resource** — if you want to assign a cluster role to the recovered VM.
   - **Do not register VM as a cluster resource** — if you do not want to assign a cluster role to the recovered VM.

IMPORTANT

You cannot recover workloads to a Microsoft Hyper-V 2008 R2 Cluster due to Microsoft limitations. You can recover workloads to Microsoft Hyper-V 2012 Failover Cluster or Microsoft Hyper-V 2012 R2 Failover Cluster.

If you need to recover a workload to a Microsoft Hyper-V 2008 R2 Cluster, you can use the following workaround:

1. Recover the workload to any node of the cluster.
2. Configure the recovered VM as highly available. For more information, see Microsoft Docs.



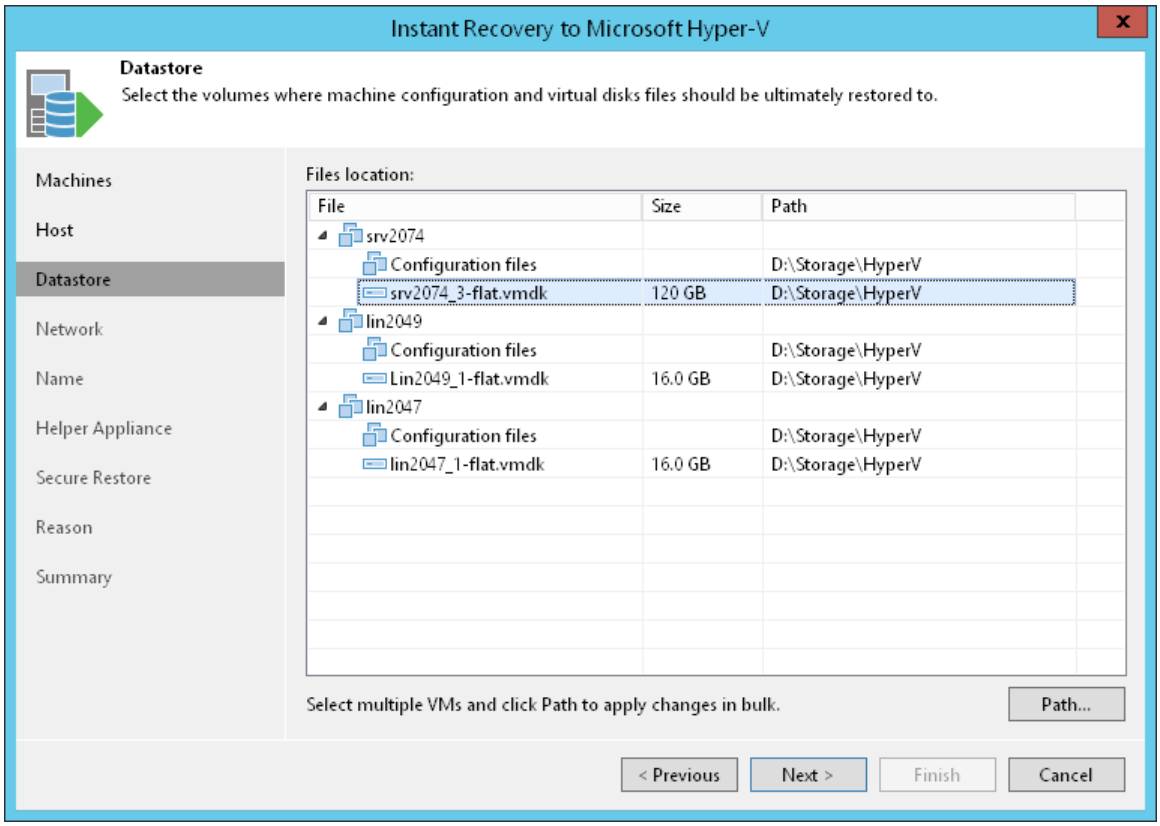
Step 6. Select Target Datastore

This step is available if you recover workloads to a new location or with different settings.

At the **Datastore** step of the wizard, specify a path to the folder where VM configuration files and disks will be stored:

1. Select the necessary workloads in the **Files location** list and click **Path**. Alternatively, you can expand a workload in the list and select individual files. Use this method if you want to place configuration and disk files to different locations.
2. In the **Select Folder** window, do one of the following:
   - Select an existing folder where VM files will be stored.
   - Create a new folder by clicking **New Folder** at the bottom of the window.
   - Type a path to an SMB3 shared folder in the search field at the bottom of the **Select Folder** window. The path must be specified in the UNC format, for example: *||172.16.11.38|Share01*.

> **IMPORTANT**
>
> The host or cluster on which you register VMs must have access to the specified SMB3 shared folder. If you are using SCVMM 2012 or later, the server hosting the Microsoft SMB3 shared folder must be registered in SCVMM as a storage device. For more information, see Microsoft Docs.



Step 7. Specify Network Mapping

This step is available if you recover workloads to a new location or with different settings.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the original site to networks in the target site (site where VMs will be recovered). When the job starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update VM configuration files to replace the original networks with the specified networks in the target site. As a result, you will not have to re-configure network settings manually.

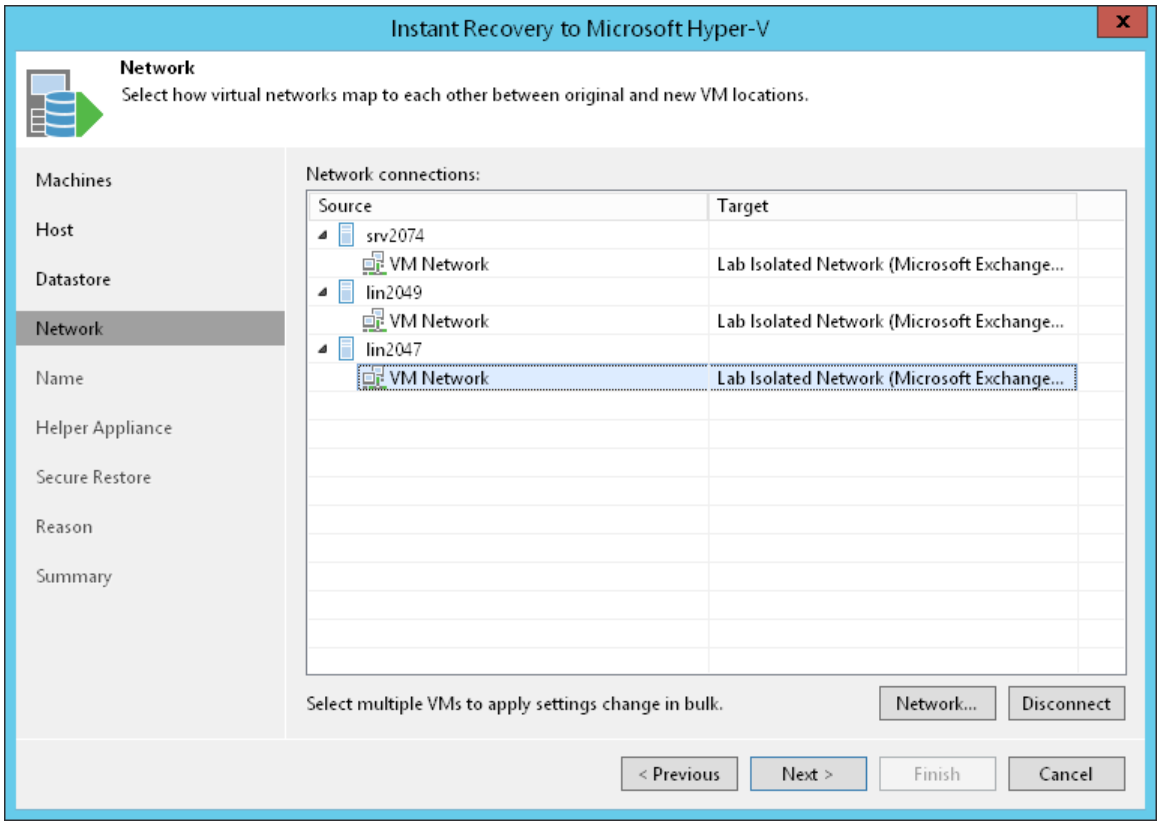
To specify network mapping settings:

1. In the **Network connections** list, select the necessary workloads and click **Network**.

   If a workload is connected to multiple networks, expand the workload, select a network to map and click **Network**.

2. In the **Select Network** window, select a network to which the selected workload must be connected after recovery.

If you do not want to connect the recovered VM to any virtual network, select the original workload in the list and click **Disconnect**.



Step 8. Change VM Name and UUID

This step is available if you recover workloads to a new location or with different settings.

At the **Name** step of the wizard, specify names under which VMs will be recovered and select whether you want to preserve VM UUIDs or change them. By default, Veeam Backup & Replication preserves the original names and UUIDs.

> **NOTE**
>
> We recommend that you specify a new name and generate a new UUID for a VM to prevent conflicts if the original workload still resides in the production environment. The name and UUID change is not required if the original workload no longer exists, for example, it was permanently deleted.

Changing Names

To change a VM name:

1. In the **Virtual machines** list, select the necessary workloads and click **Name**.
2. In the **Change Name** section, enter a new name explicitly or specify a change name rule by adding a prefix or suffix to the original workload name.

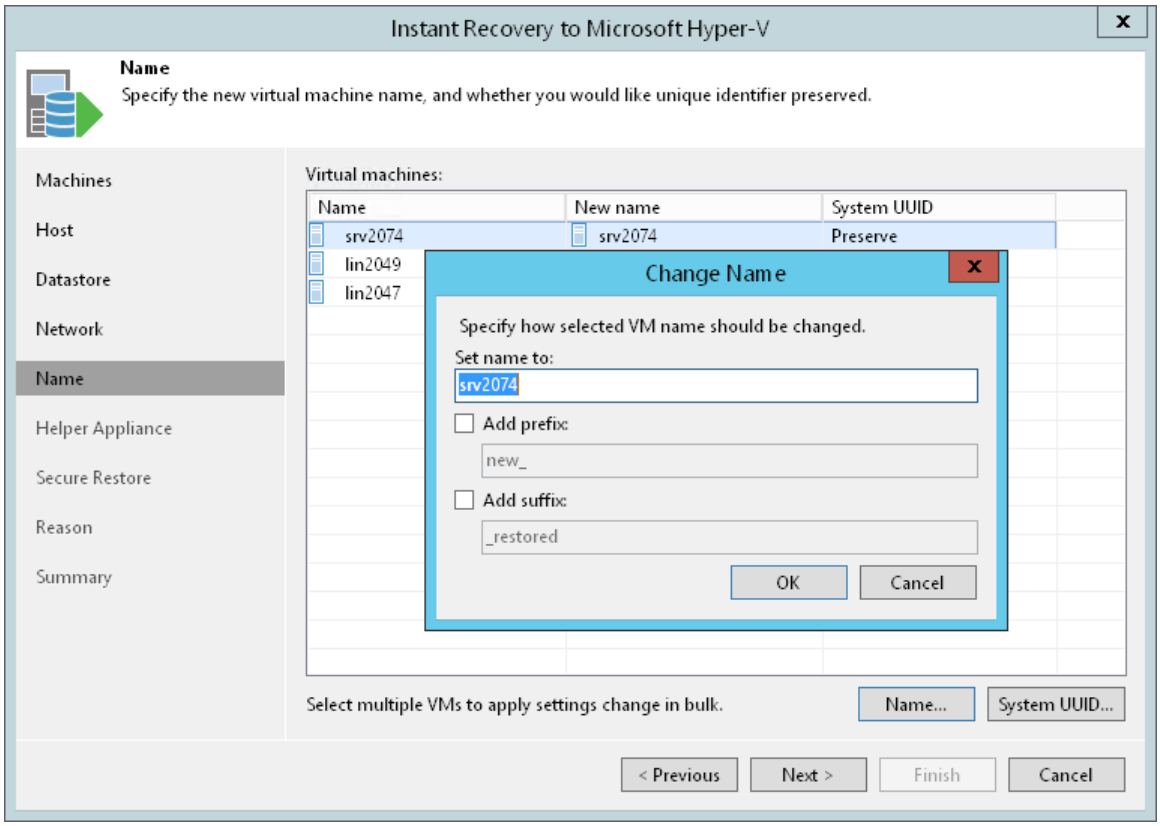
Alternatively, you can change a VM name directly in the **Virtual machines** list. To do this, click the **New Name** field and enter the name to be assigned to the recovered VM.

Changing UUIDs

To change VM identification settings:

1. Select the necessary workloads in the list and click **System UUID**.

2. In the **BIOS UUID Settings** window, select to generate a new UUID.



Step 9. Configure Helper Appliance

This step is available if you recover workloads with Linux OS and recover them to a new location or with different settings.

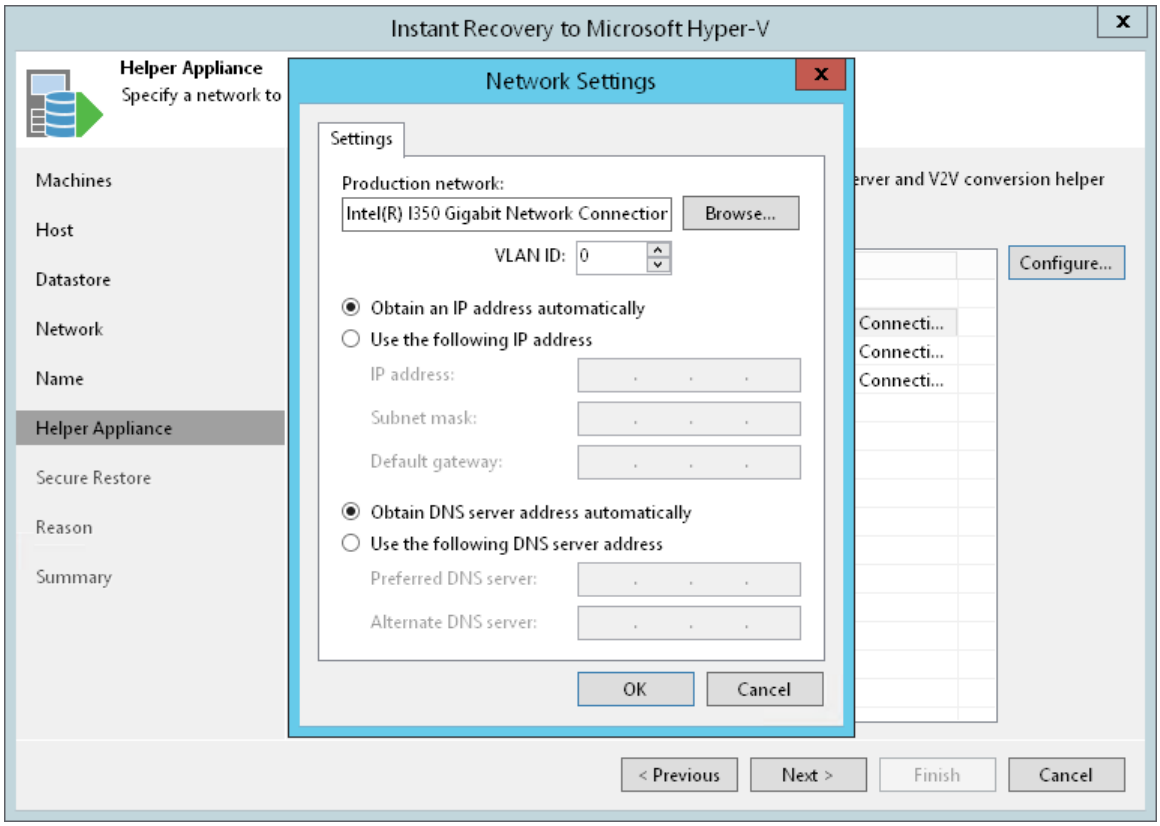
Veeam Backup & Replication recovers Linux machines using a helper appliance. The helper appliance is an auxiliary Linux-based VM registered by Veeam Backup & Replication. The appliance is quite small — around 50 MB. It requires 1024 MB RAM and takes around 10 seconds to boot.

At the **Helper Appliance** step of the wizard, configure the helper appliance network settings:

1. [For multiple machines] In the **Network** list, expand a host and select one machine for which you want to configure the helper appliance.
2. Click the **Configure** button.
3. In the **Network Settings** window, select a network for the helper appliance.
   a. Click the **Browse** button to the right of the **Production network** field.
   b. In the **Select Network** window, Veeam Backup & Replication shows a list of networks to which the target host is connected. In this list, select a network to which the helper appliance must be connected.

      Mind that the backup server and the mount server must have access to the helper appliance over the network.
4. Specify IP address settings for the helper appliance:
   - If you use a DHCP server in the network and want to obtain IP address for the helper appliance automatically, make sure that the **Obtain an IP address automatically** option is selected.
   - To manually specify the IP address of the helper appliance, select the **Use the following IP address** option and enter the IP address, subnet mask and default gateway address in the corresponding fields.
5. Specify DNS server settings:
   - If you use a DHCP server in the network and the IP address for the DNS server is obtained automatically, make sure that the **Obtain DNS server address automatically** option is selected.

- To manually specify the IP address of the DNS server, select the **Use the following DNS server address** option and enter the IP addresses of the preferred and alternate DNS servers in the corresponding fields.



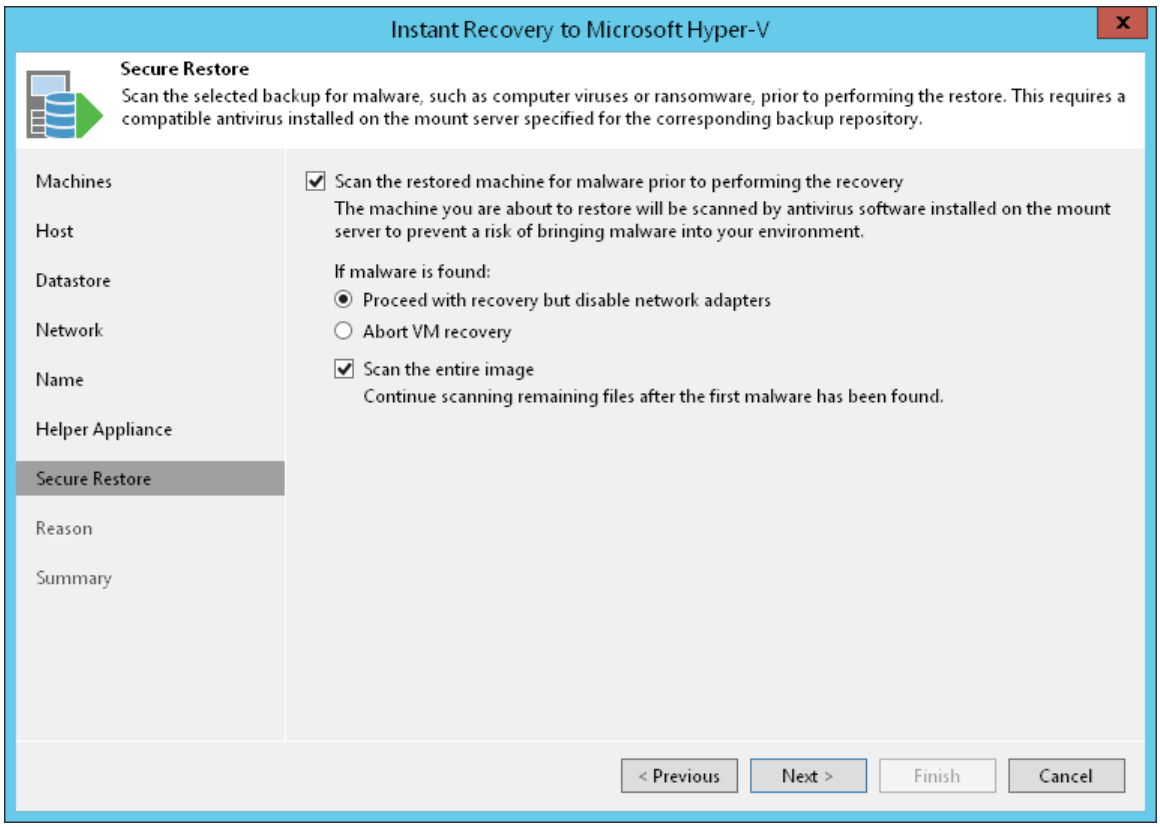
Step 10. Specify Secure Restore Settings

This step is available if you recover workloads with Microsoft Windows OS and recover them to a new location or with different settings.

You can instruct Veeam Backup & Replication to perform secure restore — scan VM data with antivirus software before restoring the VM to the production environment. For more information on secure restore, see Secure Restore.

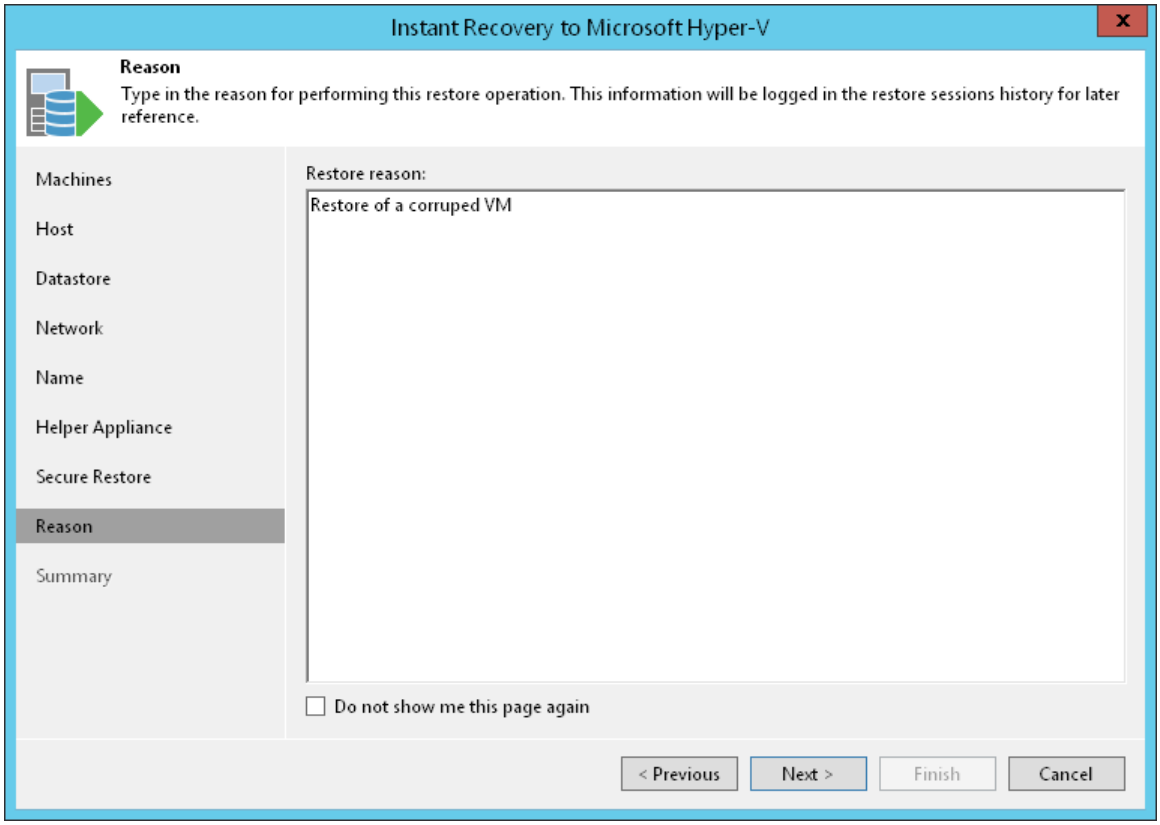
To specify secure restore settings:

1. Select the **Scan the restored machine for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if an antivirus software finds a virus threat:
   - **Proceed with recovery but disable network adapters** — if you want to restore the VM with disabled network adapters (NICs).
   - **Abort VM recovery** — if you want to cancel the restore session.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue the VM data scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



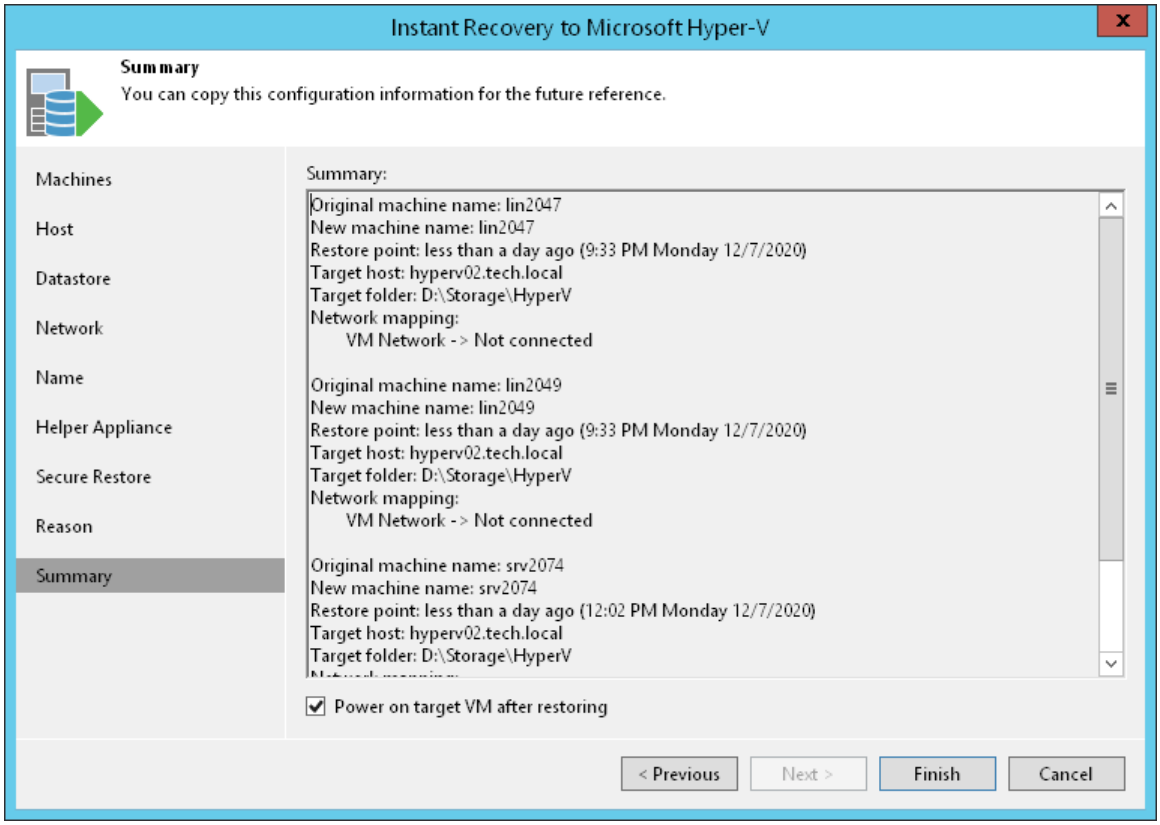
Step 11. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for performing Instant Recovery of the workloads. The information you provide will be saved in the session history and you can reference it later.



Step 12. Verify Instant VM Recovery Settings

At the **Summary** step of the wizard, check settings of Instant Recovery and click **Finish**. If you want to start the recovered VMs on the target host, select the **Power on target VM after restoring** check box.



What You Do Next

Finalizing Instant Recovery to Microsoft Hyper-V

Finalizing Instant Recovery to Microsoft Hyper-V

After the VMs have been successfully recovered, you must finalize the process. For this, test the recovered VMs and decide whether to migrate them to production environment or stop publishing.

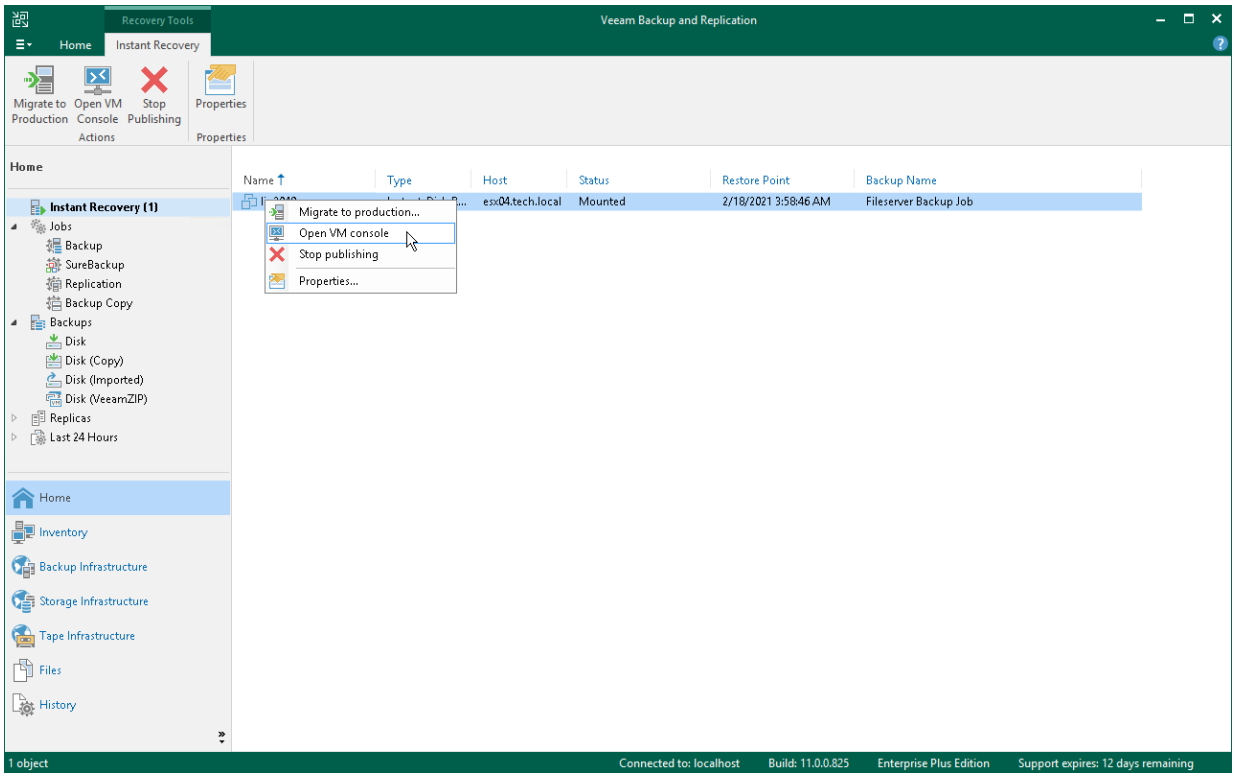
Testing Recovered VMs

To test the recovered VMs before you migrate them to production, you can launch VM consoles from Veeam Backup & Replication or open the consoles in the Hyper-V client.

To launch a VM console from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Open VM console**.



If the test fails, you can stop publishing the recovered VMs. For details, see Stop Publishing Recovered VMs.

Migrating Recovered VMs

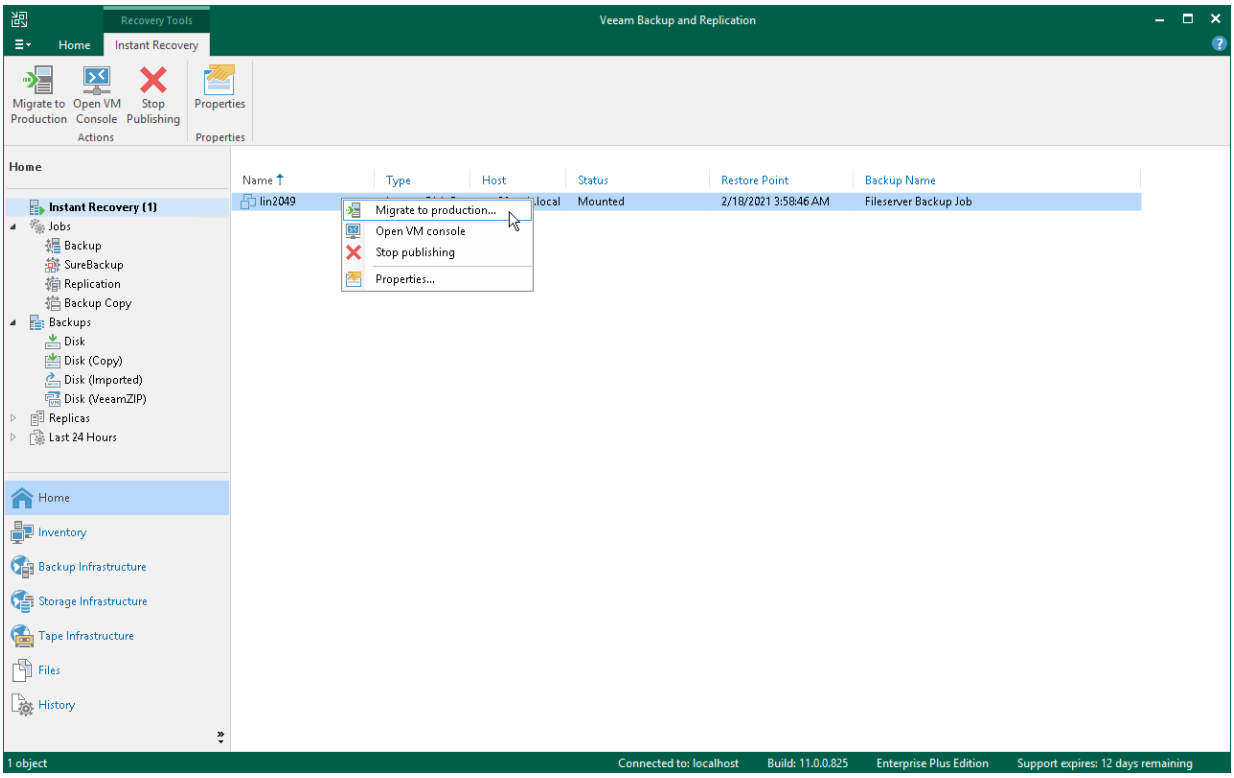
When Veeam Backup & Replication migrates VMs, it transfers VM disks data to the production storage that you have selected as a destination for the recovered VMs.

> **NOTE**
>
> [If you recover Hyper-V VMs] After the migration is finished, and if the selected destination differs from the original location, the original VMs still exist. If you do not need them, you must manually remove them in the Hyper-V client.

To migrate a recovered VM to production:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Migrate to production**.



Stop Publishing Recovered VMs

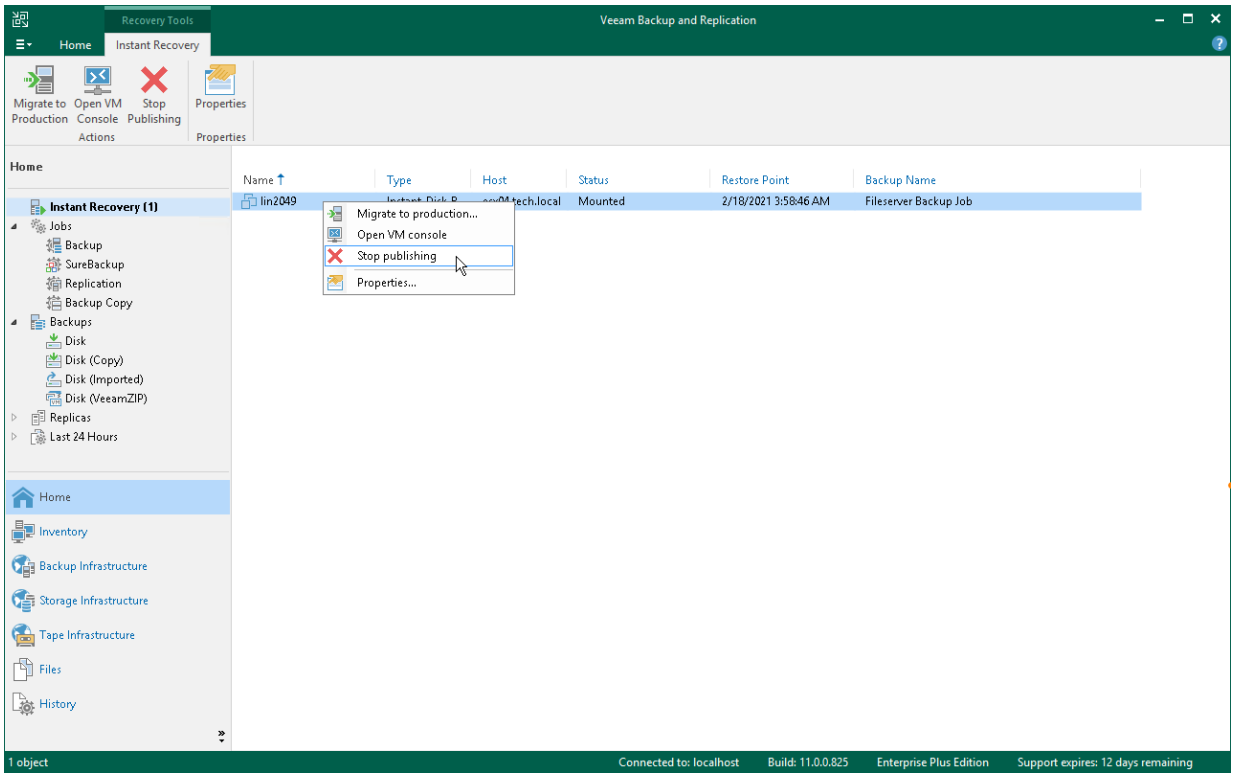
If you have ensured that the VM is working and you do not need it anymore, or your tests have failed, you can stop publishing the recovered VMs. This will remove the recovered VMs from the storage that you selected as the destination for recovery. Note that all changes made in the recovered VMs will be lost.

> **IMPORTANT**
>
> If the destination is the original location, both the original and recovered VMs are removed.

To remove a recovered VM:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Stop publishing**.



Entire VM Restore

With Veeam Backup & Replication, you can restore an entire VM from a backup file to the latest state or to a previous point in time if the original VM fails.

When you restore an entire VM, Veeam Backup & Replication extracts the VM image from a backup to the production storage.Then Veeam Backup & Replication pulls the VM data from the backup repository to the selected storage, registers the VM on the chosen ESXi host and, if necessary, powers it on.

Entire VM restore enables full disk I/O performance while Instant Recovery provides a "temporary spare" for a VM as the vPower NFS throughput is limited.

How Entire VM Restore Works

A VM can be restored to its original location or to a new location. When you restore a VM to its original location, Veeam Backup & Replication powers off the original VM and restores only those disks that are included in the backup file. All other disks remain unchanged.

When you restore a VM to a new location, you can specify new VM settings such as the new VM name, the host and datastore where the VM will reside, disk format (thin or thick provisioned) and network properties. Veeam Backup & Replication will change the VM configuration file and store the VM data to the location of your choice.

To perform entire VM restore, Veeam Backup & Replication uses one of the following transport modes:

- If the backup proxy is connected directly to the SAN fabric or has access to NFS datastores, Veeam Backup & Replication uses the Direct storage access transport mode. Veeam Data Movers, which are deployed on the backup repository and backup proxy, retrieve VM data from the backup file and put it directly to the necessary datastore.

  Veeam Backup & Replication can restore only thick VM disks using the Direct storage access transport mode. For thin VM disks restore, Veeam Backup & Replication will use the Virtual appliance or Network transport modes. Alternatively, you can instruct Veeam Backup & Replication to restore VM disks as thick.

- If the backup proxy is virtualized and resides on the ESXi host to which the VM must be restored, Veeam Backup & Replication uses the Virtual appliance transport mode. The Virtual appliance mode utilizes VMware ESXi capabilities of HotAdding disks to the VM and thus eliminates the need to transfer the backup data across the network. Veeam Data Movers deployed on the backup repository and backup proxy retrieve VM data from the backup file and put it directly to the necessary datastore via the ESXi I/O stack.

- If the Direct storage access and Virtual appliance transport modes cannot be used, Veeam Backup & Replication uses the Network transport mode.

> **NOTE**
>
> If a VM has several VM disks, Veeam Backup & Replication restores VM disks in parallel.

Quick Rollback

When you restore a full VM or VM hard disk to the original location, you can instruct Veeam Backup & Replication to perform quick rollback — incremental data restore. Instead of restoring an entire VM or VM disk from a backup file, Veeam Backup & Replication will recover only those data blocks that are necessary to revert the VM or VM disk to an earlier point in time. Quick rollback significantly reduces the recovery time and has little impact on the production environment.

For quick rollback, Veeam Backup & Replication uses the Changed Block Tracking technology. Veeam Backup & Replication gets information about the current VM state and compares it with the CBT information in the backup file. This way, Veeam Backup & Replication detects what data blocks must be transported back to the production datastore to rebuild the VM or VM disk to the necessary point in time.

It is recommended that you use quick rollback if you restore a VM or VM disk after a problem that has occurred at the level of the VM guest OS — for example, there has been an application error or a user has accidentally deleted a file on the VM 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:

- VM or VM disk must be restored to its original location.
- CBT must be enabled for the VM disk or all disks of a VM that you plan to restore.
- The backup file from which you plan to restore a VM or a VM disk must be created with the **Use changed block tracking data** option enabled.

Limitations for Quick Rollback

Quick rollback has the following limitations:

-
- 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.
- Use quick rollback and VM guest OS file exclusion wisely. If you exclude specific files and folders from the VM guest OS during backup and use quick rollback to restore the VM or VM disk from such backup, Veeam Backup & Replication will restore only the content of the backup file. The excluded data will not be restored. For example, if you exclude C:\Folder from the backup, data in this folder will not be backed up and will not be available in the resulting backup file. After some time, data in `C:\Folder` may change but the folder will still not be backed up (since the job excludes this folder). For this reason, when you perform quick rollback, Veeam Backup & Replication will restore all data that have changed except the excluded `C:\Folder`.

Restoring Entire VMs

To restore an entire VM, use the **Full VM Restore** wizard.

Before You Begin

Before you restore a VM from a backup, consider the following:

- You can restore a VM from a backup that has at least one successfully created restore point.
- If you back up a VM with vRDM disks, Veeam Backup & Replication converts the disks into VMDK files. Thus, when you restore a VM with a vRDM disk, Veeam Backup & Replication restores this disk as a VMDK file. If you want to preserve the vRDM format for restored disks, use Quick Rollback. For more information, see Quick Rollback.
- If you want to scan VM data for viruses, check the secure restore requirements and limitations.

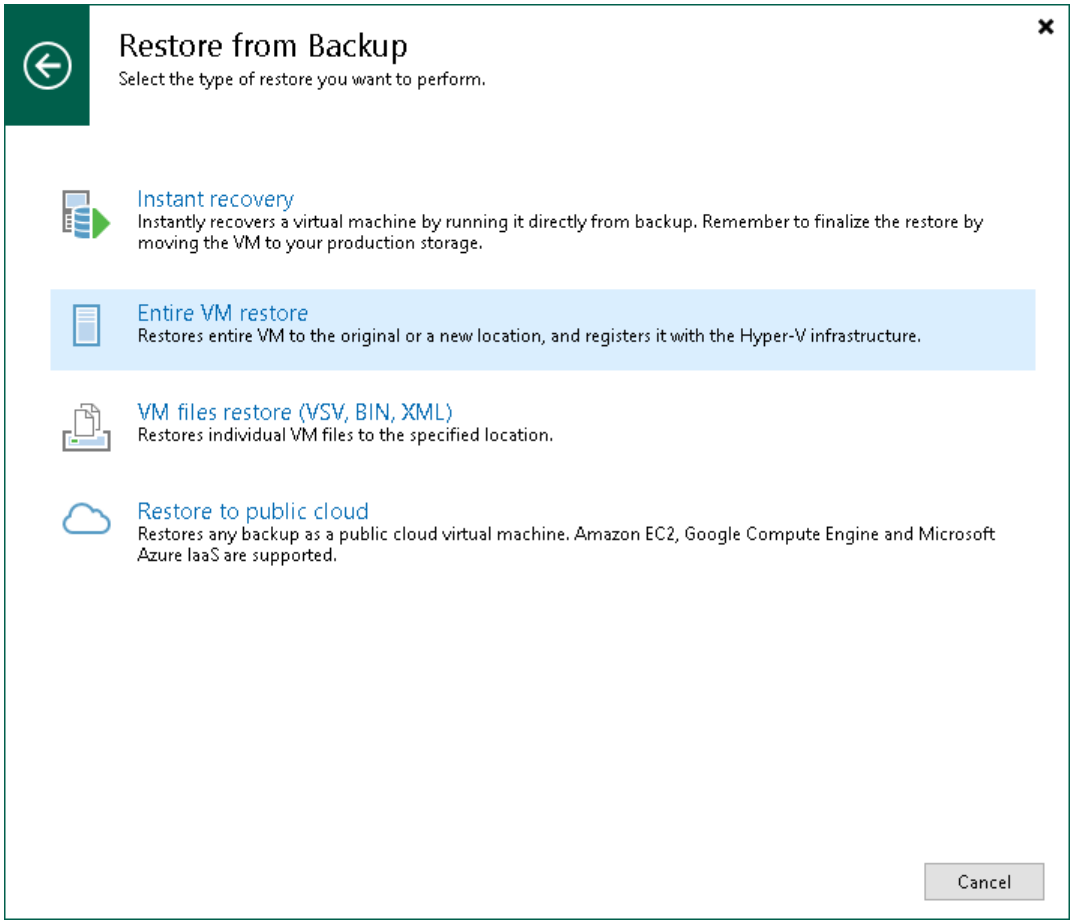
- If you want to run an executable script for a VM, check the staged restore requirements and limitations.
- When you restore a VM, mind the Virtual Hardware version compatibility. For more information, see this VMware KB article.

Step 1. Launch Full VM Restore Wizard

To launch the **Full VM Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Entire VM restore** > **Entire VM restore**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and select the machine that you want to restore and click **Entire VM** on the ribbon. Alternatively, right-click the machine that you want to restore and select **Restore entire VM**.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Entire VM**.

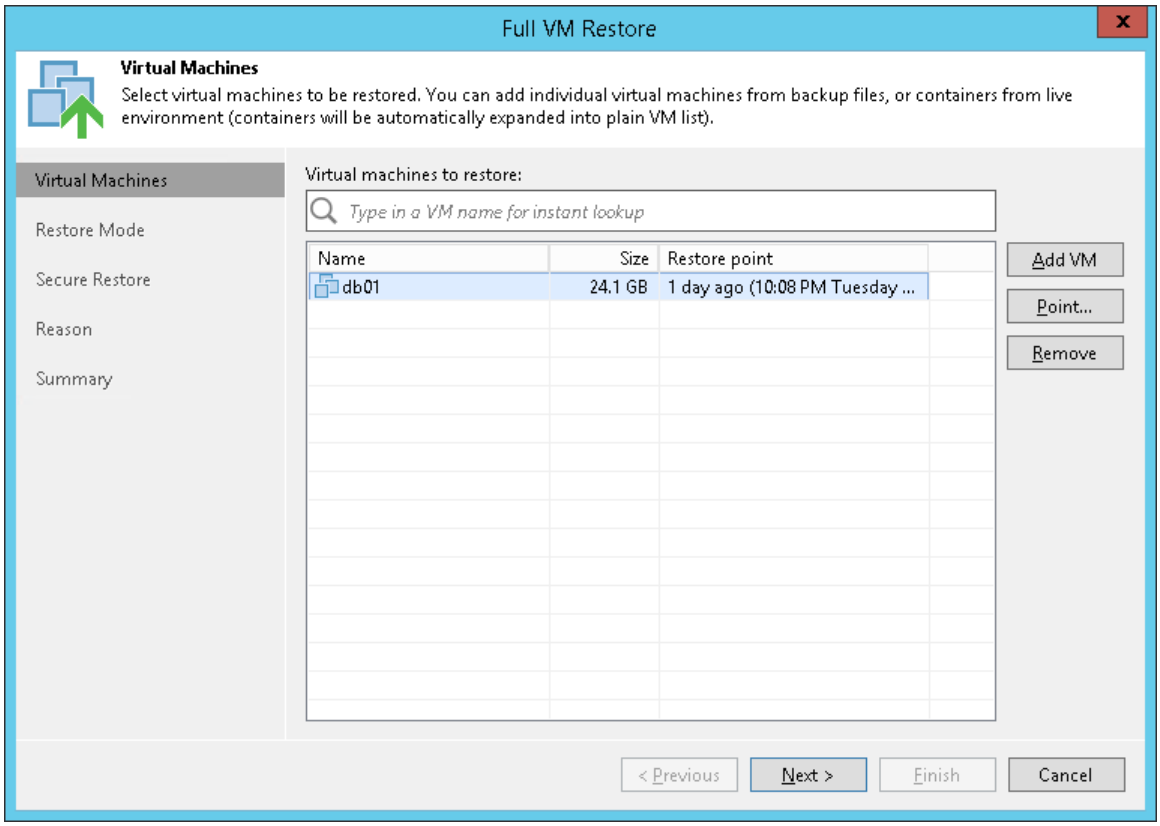
  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.



Step 2. Select VMs

At the **Virtual Machines** step of the wizard, select VMs that you want to restore:

1. Click **Add VM**.
2. Select where to browse for VMs:
   - **From infrastructure** – browse the virtual environment and select VMs or VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) to restore. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
   - **From backup** – browse existing backups and select VMs under backup jobs.

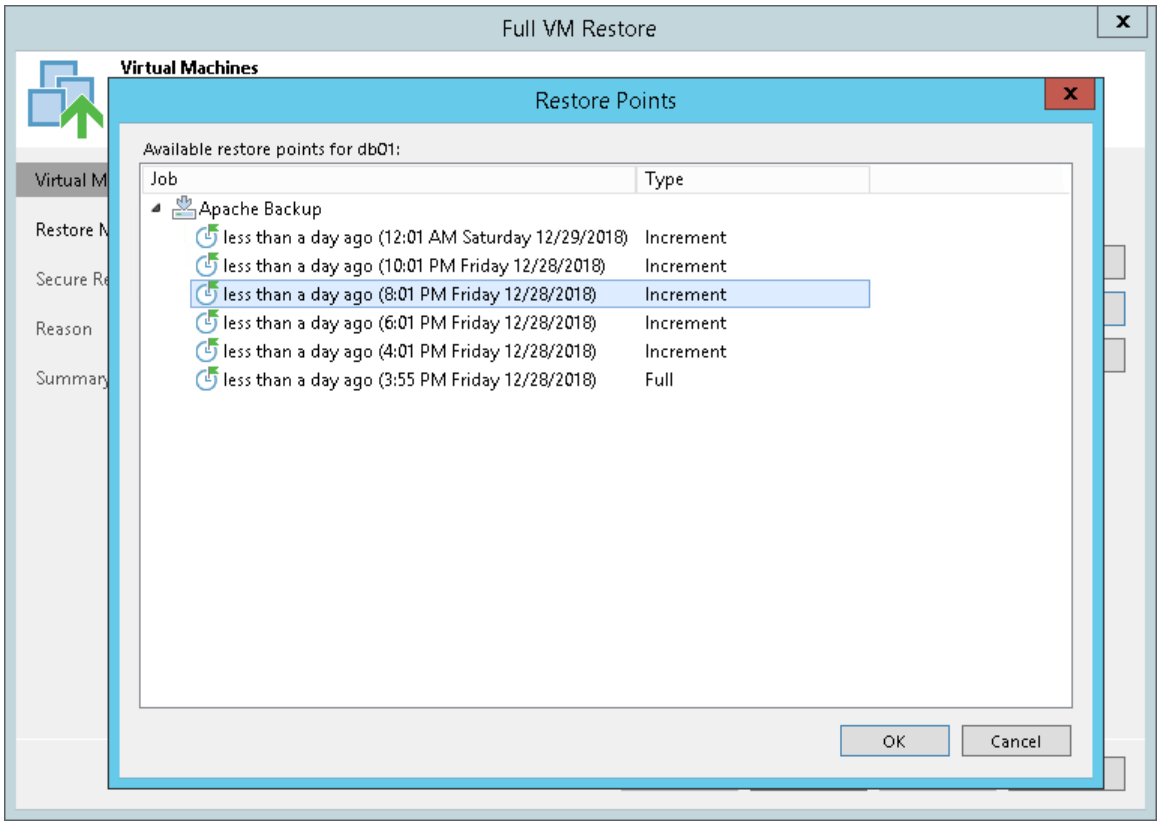


Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore point. However, you can restore the VM to an earlier state. If you have chosen to restore several VMs, you can select the necessary restore point for each VM in the list.

To select a restore point for a VM:

1. In the **Virtual machines to restore** list, select a VM.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point from which you want to restore the VM.



Step 4. Select Restore Mode

At the **Restore Mode** step of the wizard, choose the necessary restore mode and backup proxy to transfer VM data:

1. Choose a restore mode:
   - Select **Restore to original location** to restore VMs with their initial settings and to their original location. If this option is selected, you will immediately pass to the Reason step of the wizard.

     During restore to the original location, Veeam Backup & Replication restores only those disks that are included in the backup file. This means that after the restore finishes, you do not have to update existing jobs which process the original VMs.
   - Select **Restore to a new location, or with different settings** to restore VMs to a different location or with different settings (such as VM location, network settings, format of restored virtual disks and so on). If this option is selected, the **Full VM Restore** wizard will include additional steps for customizing VMs settings.

     During restore to a new location, Veeam Backup & Replication creates new VMs. If you want to process the restored VMs, you must edit existing jobs or create new jobs to process the restored VMs. If you restore VMs with the same name and to the same folder as the original VMs, Veeam Backup & Replication deletes the original VMs. In this case, you must edit existing jobs to exclude original VMs from them.
   - Select **Staged restore** to run an executable script for VMs before restoring them to the production environment. If this option is selected, the **Full VM Restore** wizard will include an additional step for customizing staged restore settings.

     During staged restore to the original location, that is, when you leave the original settings on the next steps of the wizard, Veeam Backup & Replication removes the original VMs. However, Veeam Backup & Replication automatically updates the existing jobs to process the restored VMs and to exclude the original VMs.

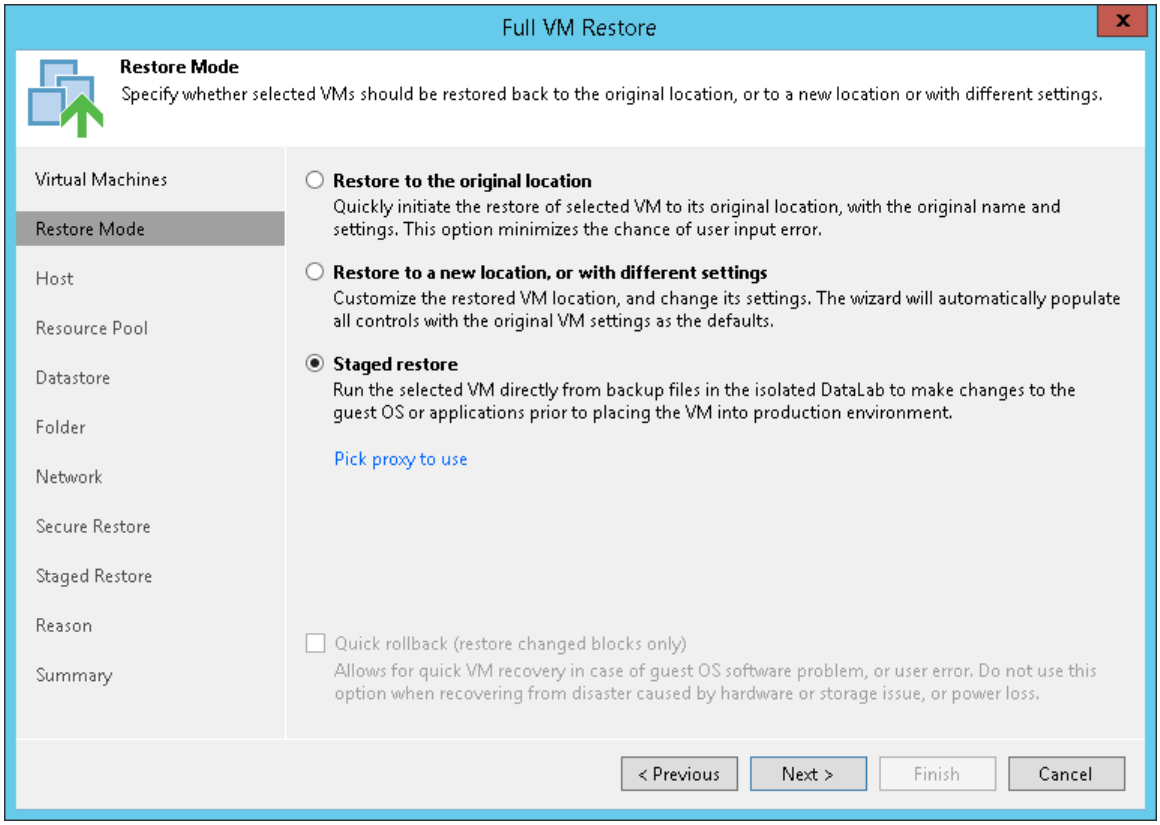
     During staged restore to a new location, Veeam Backup & Replication creates new VMs. If you want to process the restored VMs, you need to edit existing jobs or create new jobs.
2. [For VM restore to the original location] Select the **Quick rollback** check box to perform incremental restore for the VM. Veeam Backup & Replication will query CBT to get data blocks that are necessary to revert the VM to an earlier point in time, and will restore only these data blocks from the backup. Quick restore significantly reduces the restore time and has little impact on the production environment.

   It is recommended that you enable this option if you restore a VM after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.

   For more information on quick rollback, its requirements and limitations, see Quick Rollback.
3. Click the **Pick proxy to use** link to select backup proxies over which VM data must be transported to the source datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing VM data.

     During the restore process, VMs are processed simultaneously. 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 writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VMs processing.

- If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that will be used for restore. It is recommended that you select at least two proxies to ensure that VMs are recovered should one of backup proxies fail or lose its connectivity to the source datastore during restore.



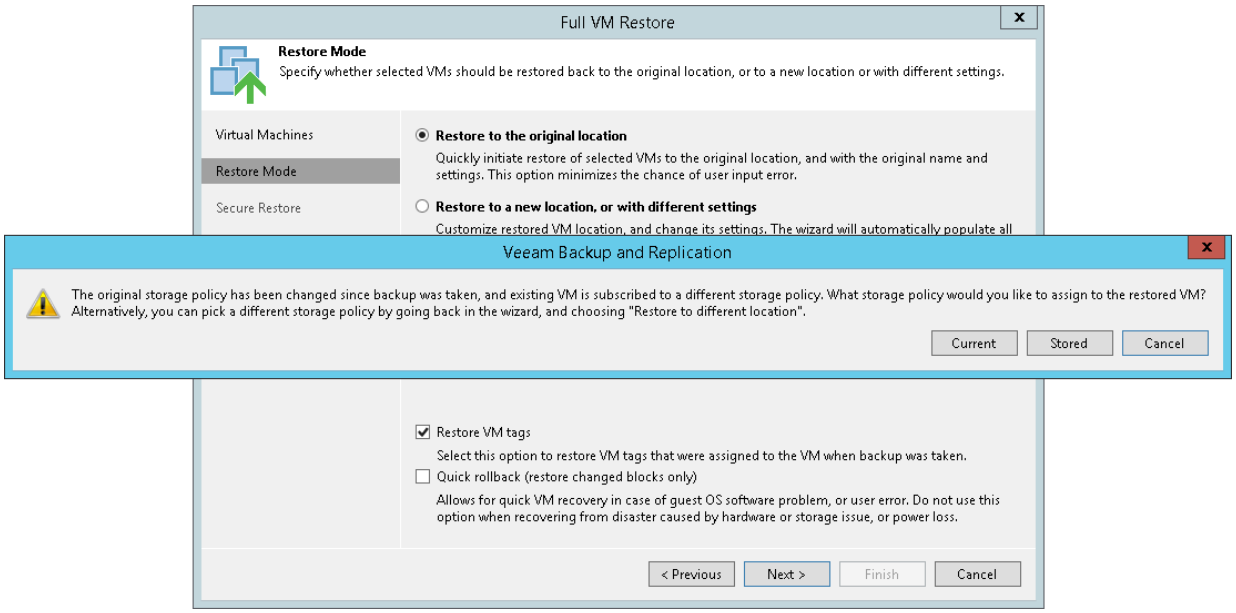
Restoring Storage Policies

If the backed-up VM was associated with the storage policy, in the restore 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 storage policy in the backup file. 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 backup, and whose information is stored in the backup file.

For more information, see Storage Policies.

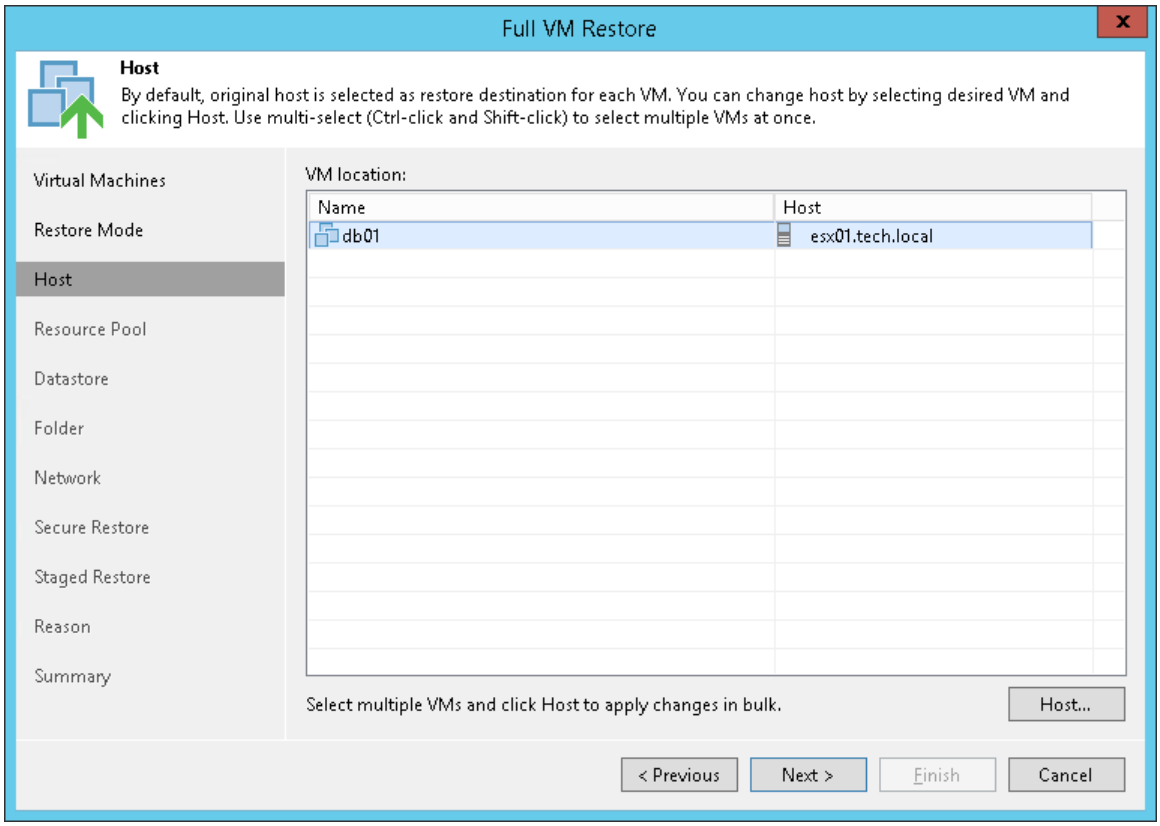


Step 5. Select Target Host

The **Host** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

To specify a target host:

1. Select the necessary VMs in the list and click **Host**.
2. Choose a host or cluster where the selected VMs must be registered.

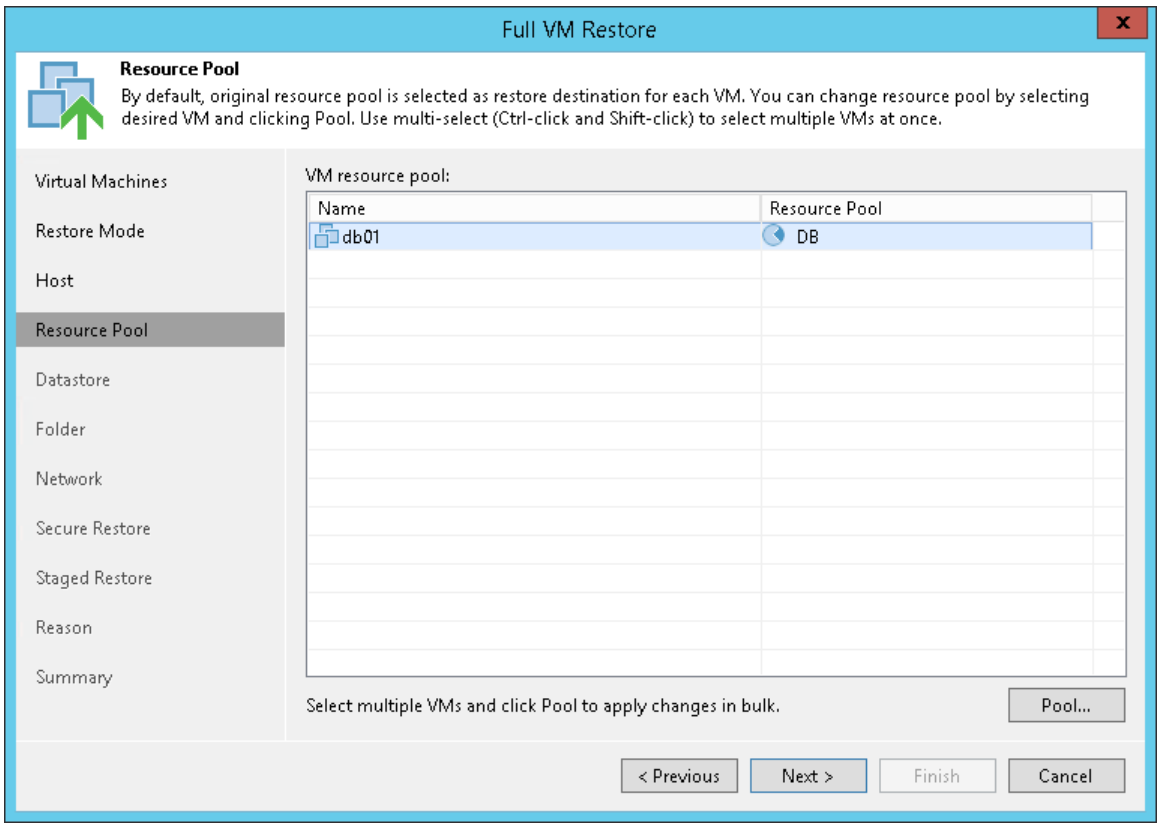


Step 6. Select Target Resource Pool

The **Resource Pool** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

To specify a destination resource pool:

1. Select the necessary VMs in the list and click **Pool**.
2. Select a resource pool to which the VMs must be placed.
3. If necessary, select a vApp in which the VMs must be included.



Step 7. Select Target Datastore and Disk Type

The **Datastore** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

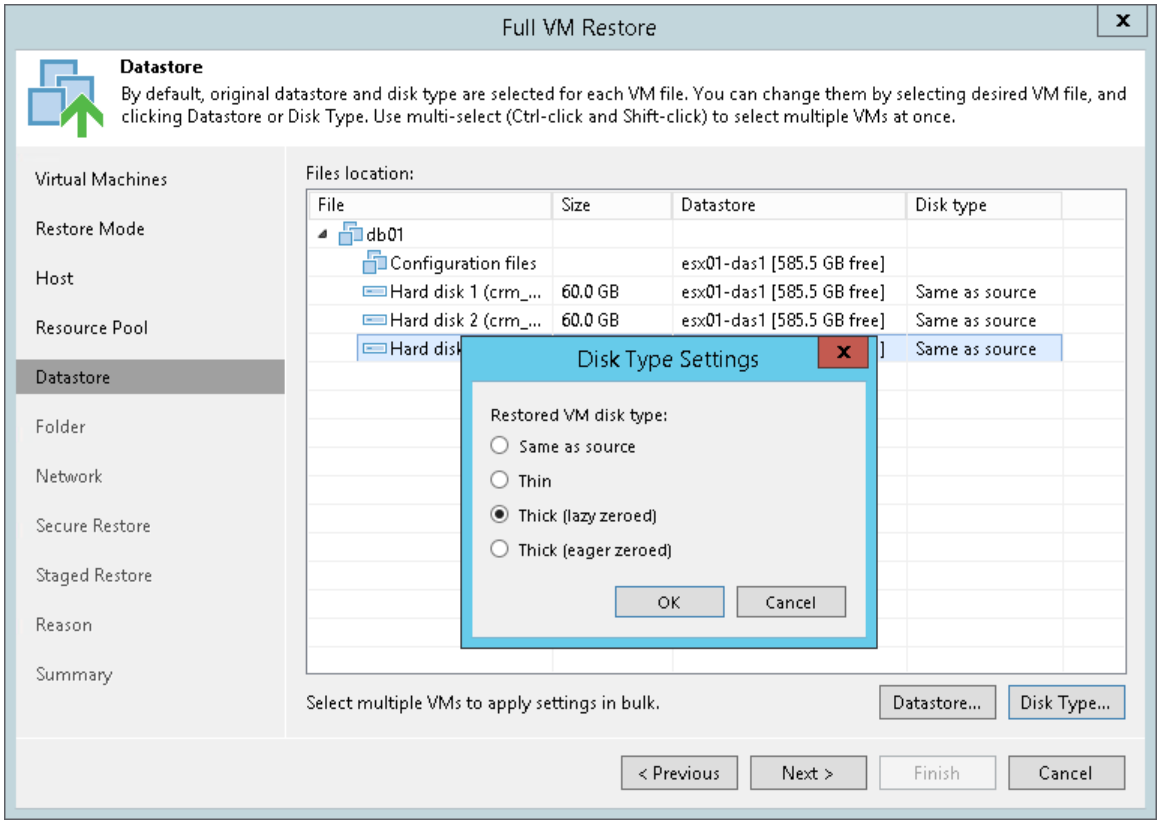
You can place an entire VM to a particular datastore or choose to store configuration files and disk files of the restored VM in different locations.

1. To place VMs to a datastore, select the necessary VMs in the **Files location** list and click **Datastore**. In the **Disk Type Settings** window, select a datastore and click **OK**.

   If you want to place configuration and disk files to different datastores, expand a VM in the **File location** list and select individual files. Then specify the necessary datastore.

2. By default, Veeam Backup & Replication preserves the format of restored VM disks. To change the disk format, expand a VM in the **Files location** list, select the necessary disks and click **Disk Type**. In the **Disk Type Settings** section, choose the format that will be used to restore virtual disks of the VM: same as source, thin, thick lazy zeroed or thick eager zeroed. For more information about disk types, see VMware Docs. Click **OK**.

> **NOTE**
>
> Disk format change is supported only for VMs with Virtual Hardware version 7 or later.



Step 8. Select Target Folder and Change VM Settings

The **Folder** step of the wizard is available if you have chosen to change the location and settings for restored VMs.

At the **Folder** step of the wizard, specify a destination VM folder, choose whether you want to restore VM tags and change VM names. By default, Veeam Backup & Replication preserves the original names.

Specifying Destination VM Folder

To specify a destination VM folder:

1. Select a VM in the list and click **Folder**.
2. Choose a folder to which the VM will be placed.

> **NOTE**
>
> Mind the following:
>
> - If you restore a VM to a standalone ESXi host that is not managed by the vCenter Server, you cannot select a destination folder: this option will be disabled.
> - During entire VM restore, Veeam Backup & Replication preserves the UUID of the original VM.

Changing Names

To change a VM name:

1. Select a VM in the **VM folder** list and click **Name**.
2. In the **Change Name** section, enter a new name explicitly or specify a change name rule by adding a prefix or suffix to the original VM name.

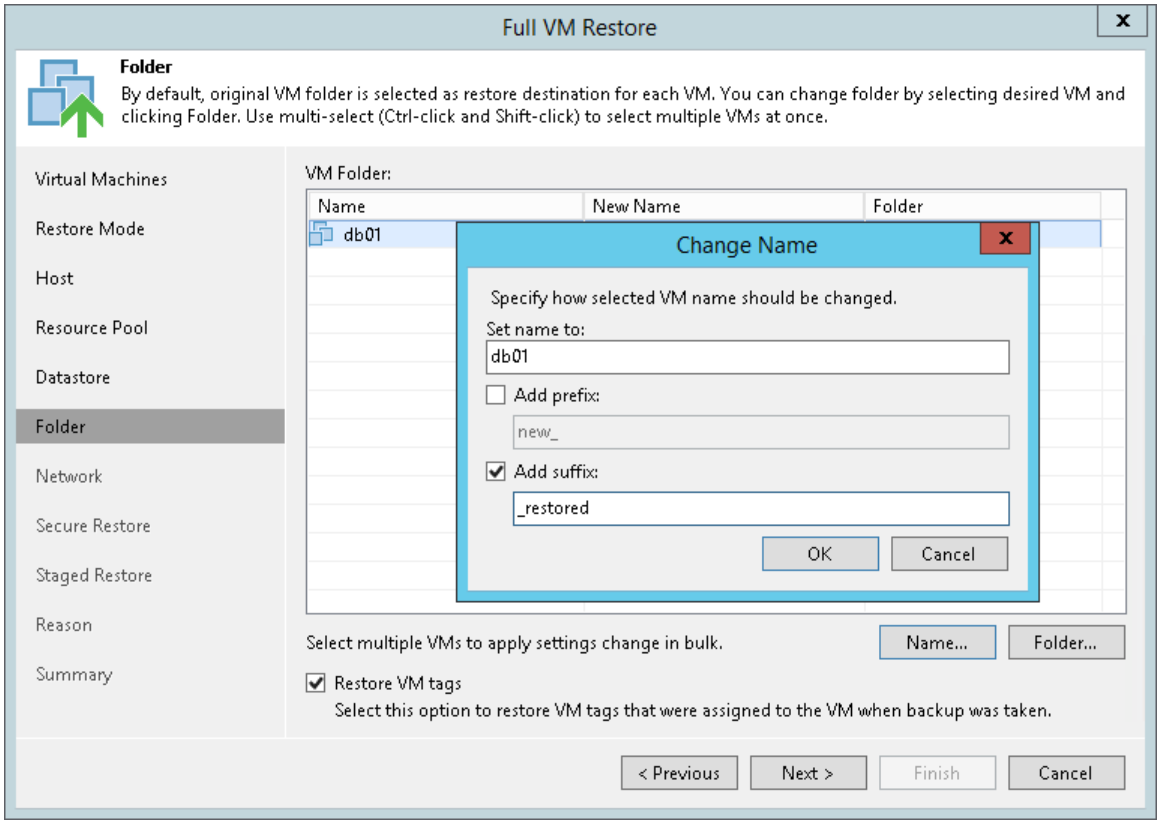
Alternatively, you can change a VM name directly in the **VM folder** list. To do this, click the **New Name** field and enter the name to be assigned to the recovered VM.

Restoring VM Tags

Select the **Restore VM tags** check box if you want to restore tags that were assigned to the original VM and to assign them to the restored VM. Veeam Backup & Replication will restore the VM with original tags if the following conditions are met:

- The VM is restored to its original location.

- The original VM tag is still available on the source vCenter Server.



Step 9. Specify Network Mapping

The **Network** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

If you plan to restore a VM to a new location, for example, another site with a different set of networks, you can map source site networks to target site networks. Veeam Backup & Replication will use the network mapping table to update configuration files of the VM on the fly, during the restore process.

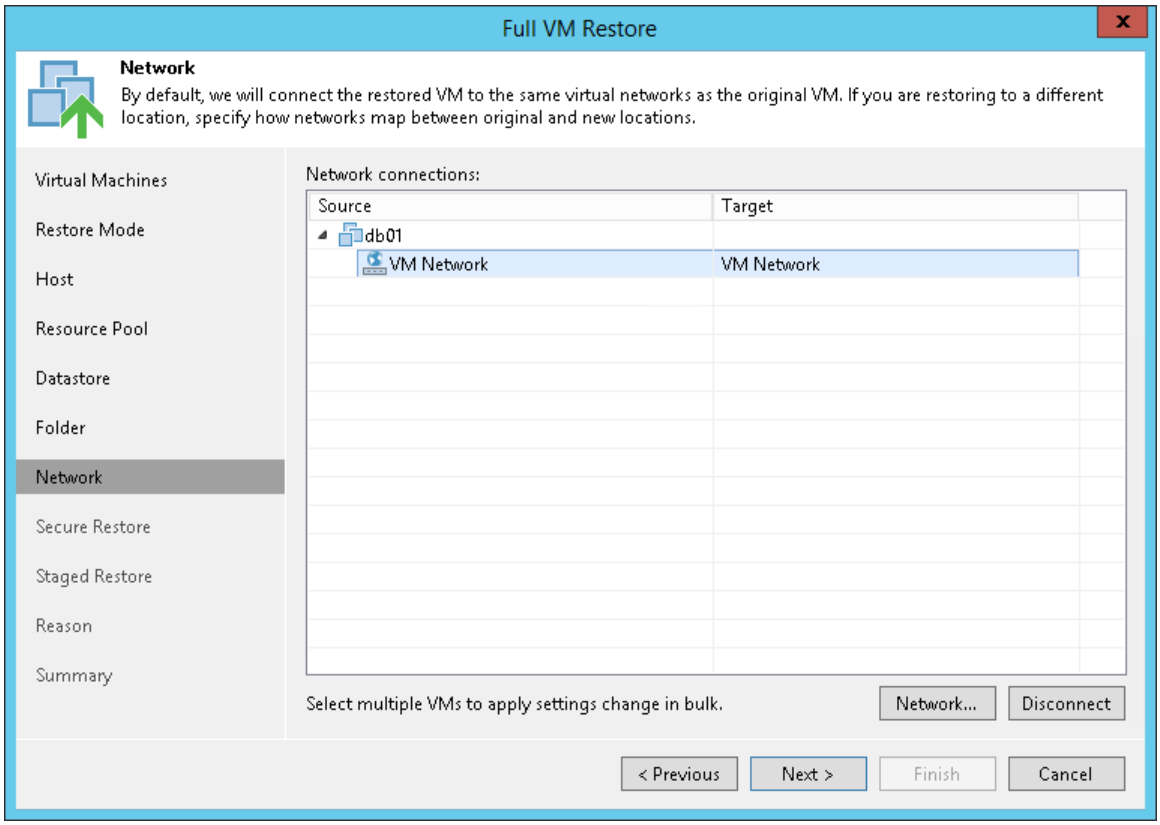
To change networks to which the restored VM will be connected:

1. Select a VM in the list and click **Network**. To apply changes in bulk, select several VMs in the list and click **Network**.

   If a VM is connected to multiple networks, expand the VM, select the network to map and click **Network**. The **Select Network** section displays all networks to which the target host or cluster is connected.

2. From the list of available networks, choose a network to which the VM must have access upon restore.

If you do not want to connect a restored VM to your virtual networks, select the VM in the list and click **Disconnected**.



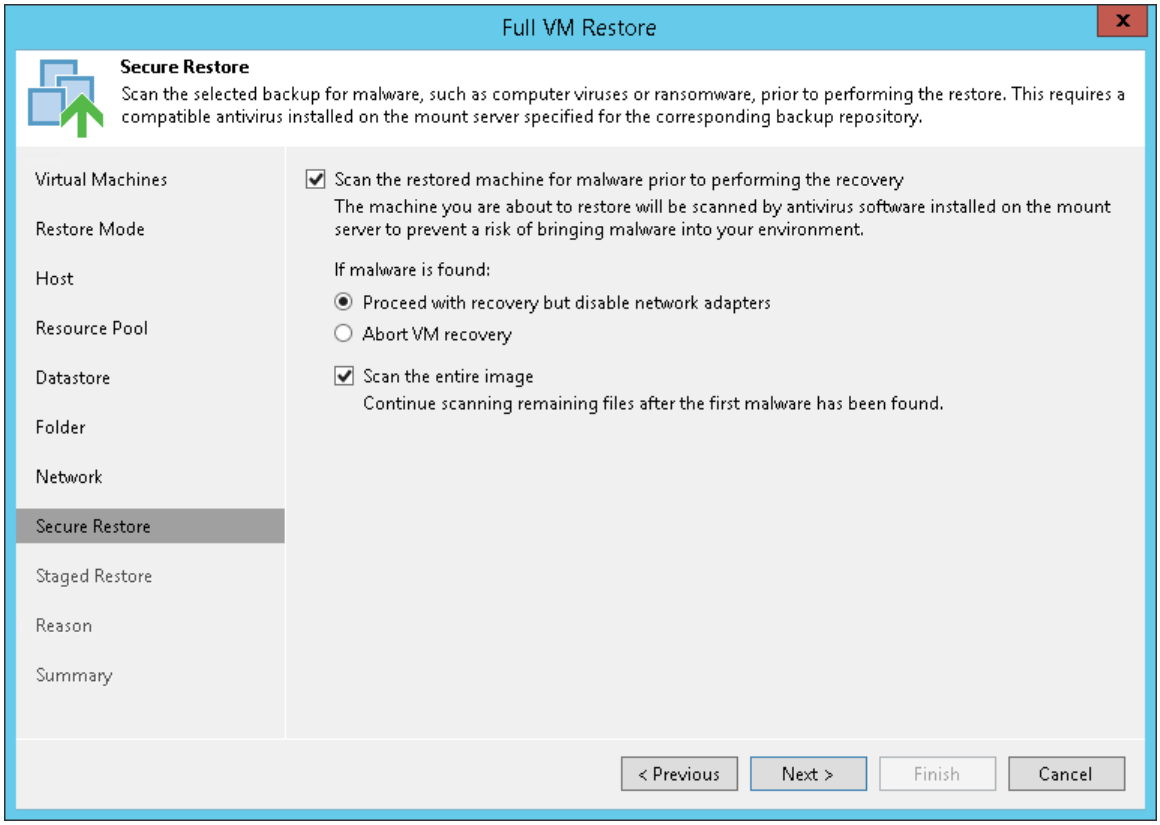
Step 10. Specify Secure Restore Settings

This step is available if you restore Microsoft Windows VMs.

You can instruct Veeam Backup & Replication to perform secure restore — scan machine data with antivirus software before restoring the machine to the production environment. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

1. At the **Secure Restore** step of the wizard, select the **Scan the restored machine for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software finds a virus threat:
   - **Proceed with recovery but disable network adapters**. Select this action if you want to restore the machine with disabled network adapters (NICs).
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
3. Select the **Scan the entire image** check box if you want the antivirus to continue machine scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



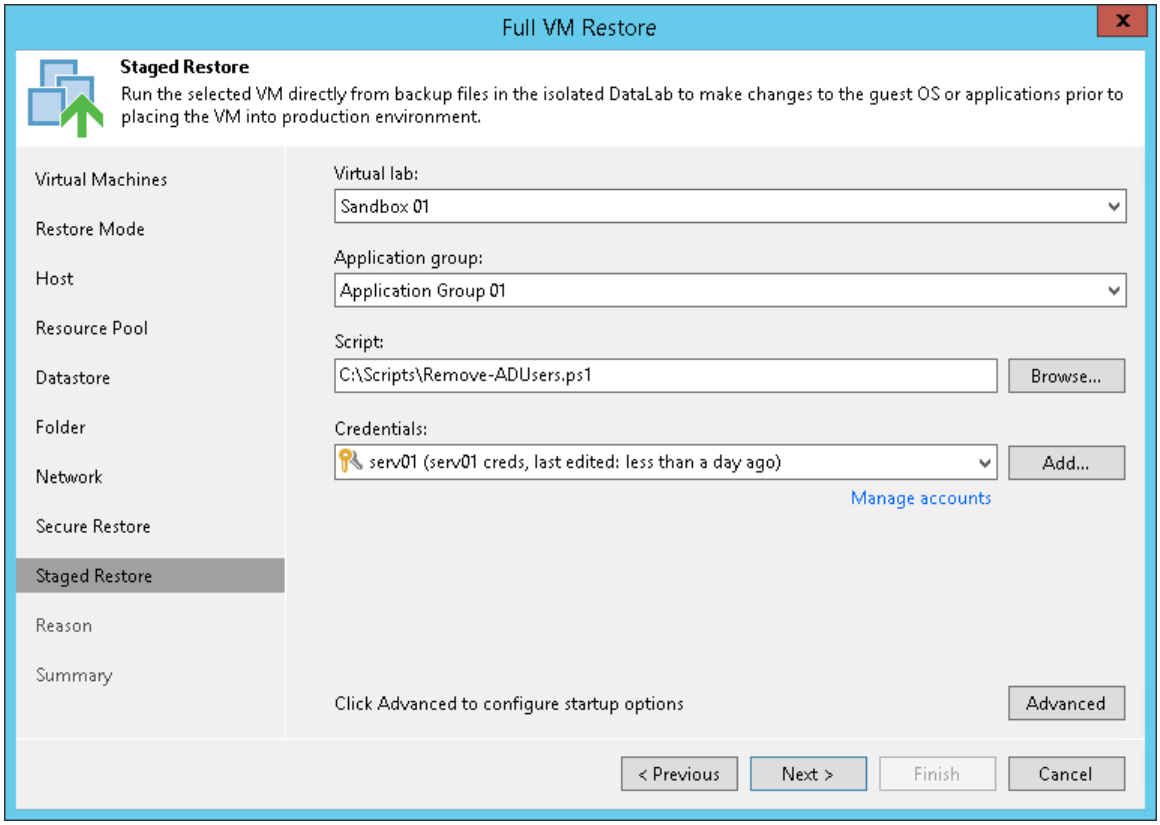
Step 11. Specify Staged Restore Settings

The **Staged Restore** step of the wizard is available if you have chosen the **Staged restore** option at the Restore Mode step of the wizard.

Staged restore to run an executable script for VMs before restoring them to the production environment. For more information, see Staged Restore.

To specify staged restore settings:

1. From the **Virtual lab** list, select a virtual lab that will be used to start VMs. The list contains all virtual labs that are created or connected to the backup server.
2. From the **Application group** list, select an application group if script execution requires other VMs to be powered on. In the virtual lab during staged restore, Veeam Backup & Replication will start VMs from the selected application group in the required order. The **Application group** list contains all application groups that are created on the backup server. For more information, see Application Group.
3. On the right of the **Script** field, click **Browse** to choose the script from a local folder on the backup server.
4. From the **Credentials** list, select credentials for the account that has administrator privileges on VMs for 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 of the **Credentials** field to add the credentials. For more information, see Managing Credentials.



VM Startup Settings

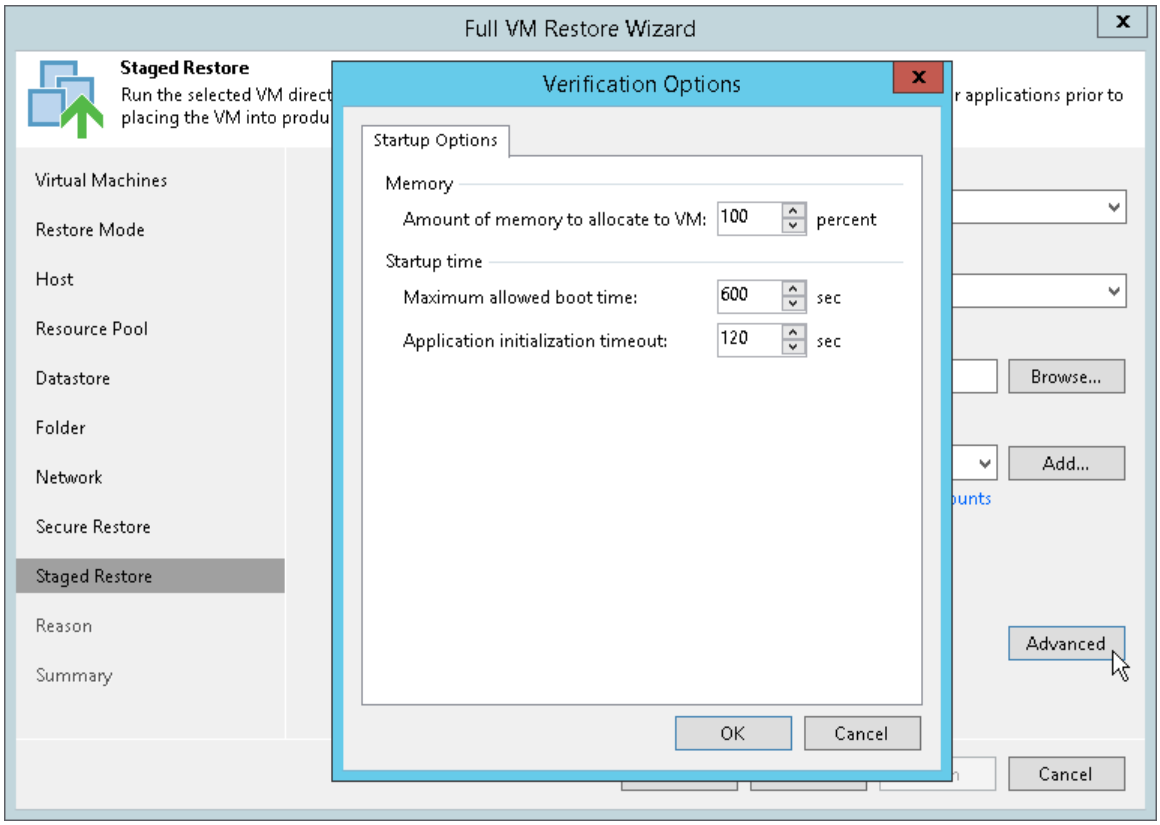
If you want to start VMs after recovery, perform the following steps:

1. Click **Advanced**.

2. In the **Memory** section, specify the amount of memory that you want to pre-allocate to a VM when it 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 4096 MB of RAM is allocated to the VM in the production environment and you specify 50% as a memory rate, 2048 MB of RAM will be allocated to the VM on startup.

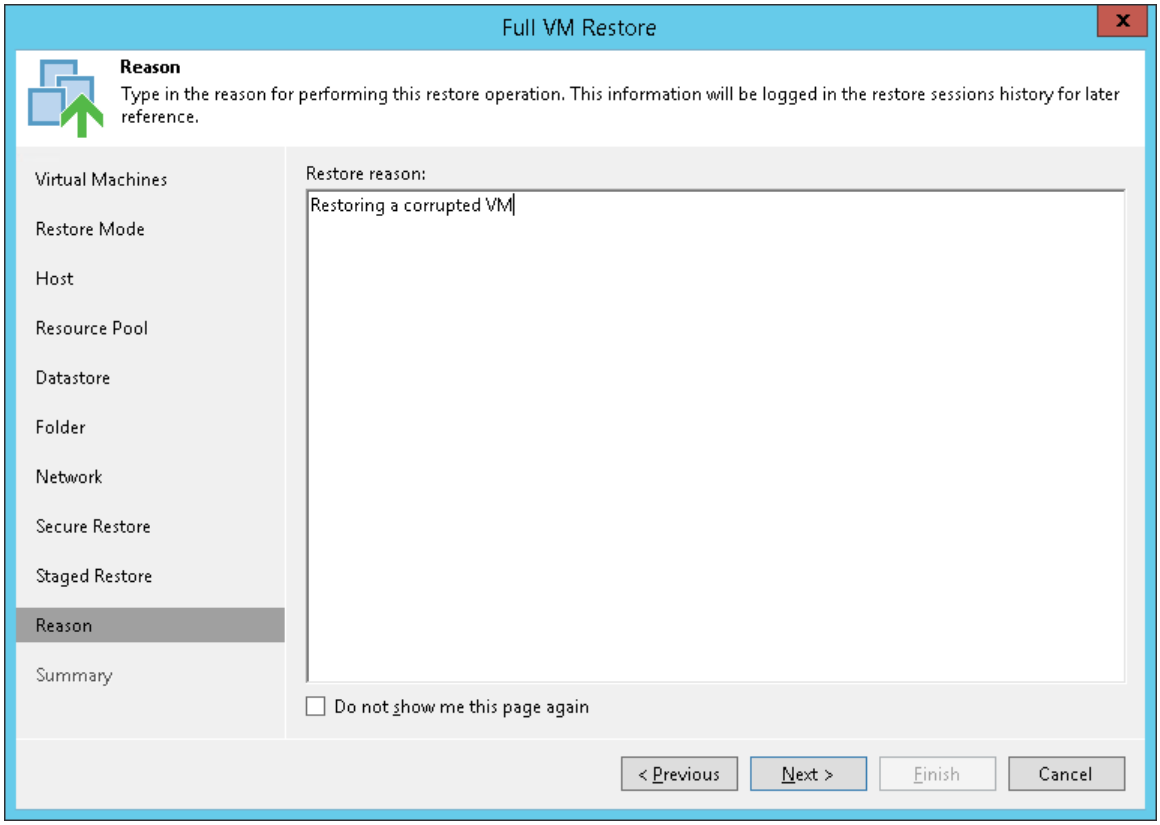
3. In the **Startup time** section, specify the allowed boot time for the VM and timeout to initialize applications on the VM.

   Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started in a virtual lab 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 process fails with the timeout error. If such error occurs, you need to increase the **Maximum allowed boot time** value and perform VM restore again.



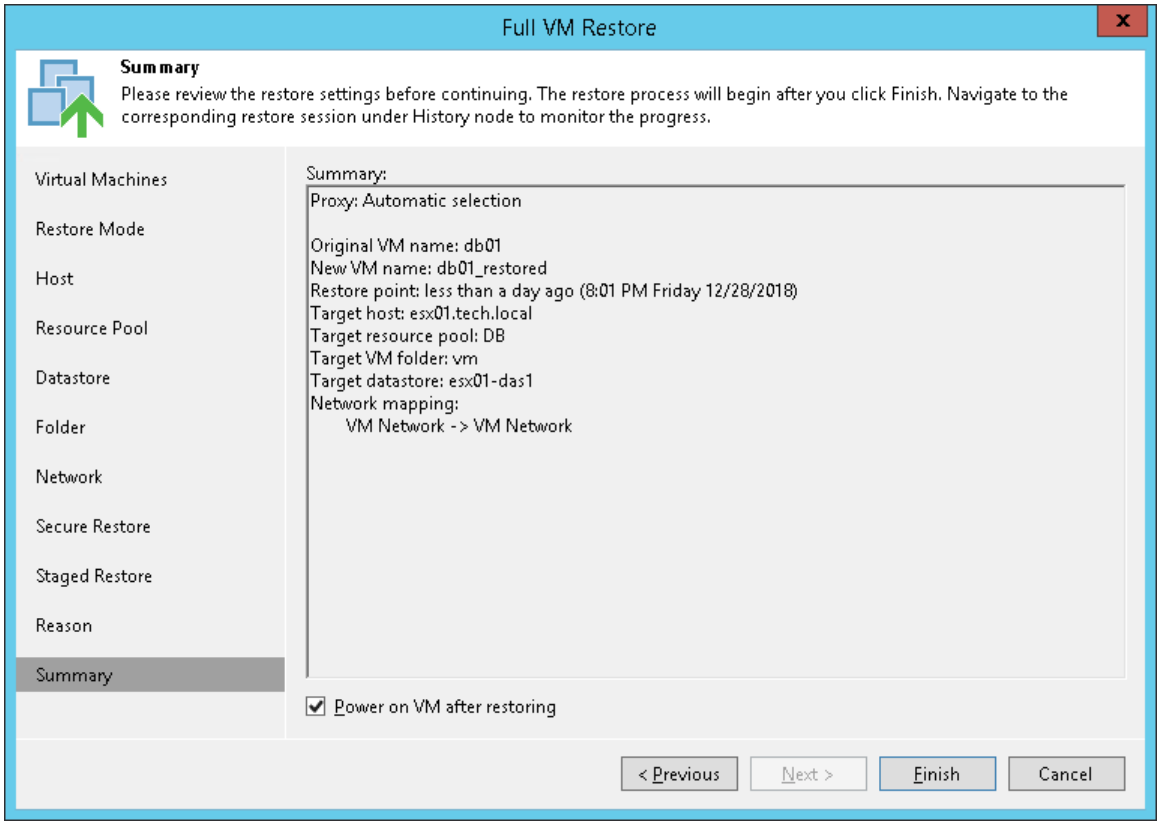
Step 12. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the selected VMs. The information you provide will be saved in the session history and you can reference it later.



Step 13. Verify Restore Settings

At the **Summary** step of the wizard, check restore settings and click **Finish**. If you want to start the recovered VMs on the target host, select the **Power on target VM after restoring** check box.



Staged Restore

Staged restore allows you to run an executable script for VMs before recovering them to the production environment. Staged restore is a part of the entire VM restore operations. To perform staged restore, you must select the **Staged Restore** mode in the Full VM Restore wizard and specify staged restore settings.

Staged restore can help you ensure that recovered VMs do not contain any personal or sensitive data. For example, you can instruct Veeam Backup & Replication to run a Windows PowerShell script that removes Active Directory users:

```
$UserName = "John.Smith"
$ADUser = Get-ADUser -Filter 'Name -like $UserName'
if (!$ADUser)
   {
       [Environment]::Exit(1)
   }
Remove-ADUser -Identity $UserName -Confirm:$false
```

NOTE

The staged restore functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

Requirements and Limitations for Staged Restore

Before you perform staged restore, check the following prerequisites:

- You must have a preconfigured virtual lab in your backup infrastructure. For more information, see Virtual Lab.
- Scripts that you plan to run must reside in a local folder on a backup server.
- If you plan to perform staged restore for several VMs within one restore session, make sure these VMs run OS of the same type: either Microsoft Windows or Linux. In the current version of Veeam Backup & Replication, you cannot specify credentials and scripts for each VM individually.
- When restoring VMs, Veeam Backup & Replication uses the Veeam Quick Migration method. vMotion and Storage vMotion methods cannot be used. For more information on the Quick Migration method, see Quick Migration.

How Staged Restore Works

For staged restore, Veeam Backup & Replication uses a preconfigured virtual lab, an executable script located on the backup server, and credentials to connect to VMs and run the script. Veeam Backup & Replication performs staged restore in the following way:

1. In the virtual lab, Veeam Backup & Replication starts VMs directly from compressed and deduplicated backup files that reside in the backup repository. To achieve this, Veeam Backup & Replication uses the Veeam vPower NFS Service.

   If you selected to use an application group to run a script, Veeam Backup & Replication first starts VMs from the application group in the required order.

2. Veeam Backup & Replication copies the script from the backup server to VMs that you plan to restore.

   To connect to VMs, Veeam Backup & Replication uses credentials specified in staged restore settings.

3. Veeam Backup & Replication runs the copied script on every VM.

   To run the script, Veeam Backup & Replication uses the same technology as for pre-freeze and post-thaw scripts. For more information, see Pre-Freeze and Post-Thaw Scripts.

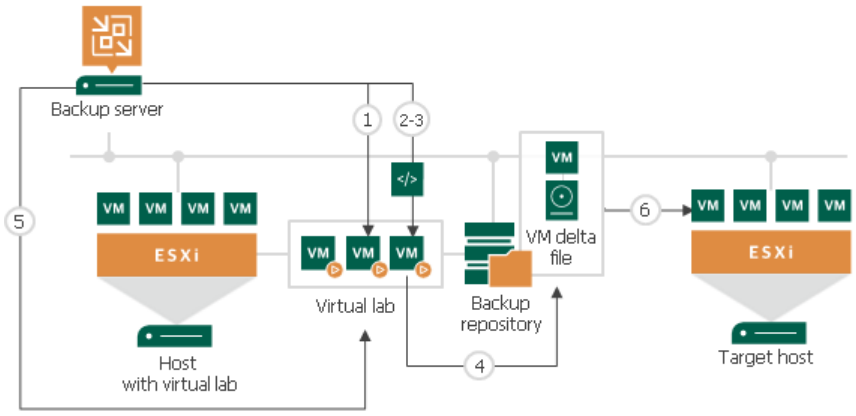
4. All VM changes that take place during script execution are written to VM delta files.

   By default, Veeam Backup & Replication stores delta files on the vPower NFS server. You can change the destination for VM delta files in virtual lab settings.

5. After the script execution is complete, Veeam Backup & Replication makes a safe shutdown of VMs in the virtual lab.

6. Veeam Backup & Replication restores VMs in a changed state to the production environment.

   To achieve that, Veeam Backup & Replication copies VM data from the backup repository and delta files to the target host using Veeam Quick Migration.



Restore to Microsoft Azure

Veeam Backup & Replication allows you to restore different workloads (VMs, Google VM instances, physical servers and so on) from backups to Microsoft Azure.

You can use Veeam Backup & Replication to complete the following tasks:

- Restore workloads from Veeam backups to Microsoft Azure.
- Migrate workloads from the on-premises infrastructure to the cloud.
- Create a test environment in the cloud for troubleshooting, testing patches and updates and so on.

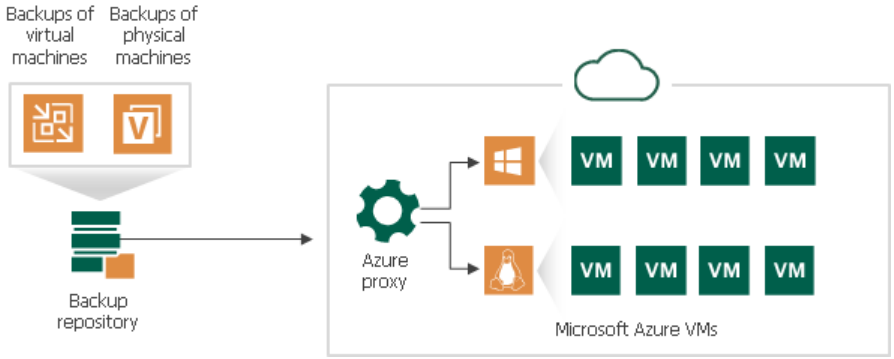
Supported Backup Types

You can restore workloads from the following types of backups:

- Backups of Microsoft Windows and Linux VMs created by Veeam Backup & Replication. You can use backups of VMware vSphere VMs or vCloud Director VMs.
- Backups of Microsoft Windows machines created by Veeam Agent for Microsoft Windows. Backups must be created at the entire machine level or volume level.
- Backups of Linux machines created by Veeam Agent for Linux. Backups must be created at the entire machine level or volume level.
- Backups of Amazon EC2 instances created by Veeam Backup for AWS.
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure.
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud Platform.
- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV.
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of RHV VMs created by Veeam Backup for RHV.

You can restore a workload to the latest restore point or any previous restore point in the backup chain.

Veeam Backup & Replication employs the Microsoft Azure Resource Manager deployment model.
Veeam Backup & Replication supports batch restore — you can launch the restore process for several VMs at a time.



> **IMPORTANT**
>
> Starting from Veeam Backup & Replication version 9.5 Update 4, the Classic deployment model is deprecated. Thus, you cannot add Classic Azure accounts. You can restore VMs in the Classic model only if you have added the Classic Azure account before upgrading to Veeam Backup & Replication 9.5 Update 4.

Generation 2 VM Support

By default, Veeam Backup & Replication restores workloads as Generation 1 VMs. Veeam Backup & Replication also offers experimental support for Generation 2 VMs within restore to Microsoft Azure feature.

Generation 2 VMs are virtual machines with advanced functionality. For more information about Generation 2 VMs, see Microsoft Docs. For more information about Microsoft Azure support for Generation 2 VMs, see Microsoft Docs.

To learn how to enable Generation 2 VM support, see this Veeam KB article. To learn Generation 2 VM support limitations, see Limitations for Restore to Microsoft Azure.

Limitations for Restore to Microsoft Azure

When planning to restore workloads to Microsoft Azure, mind the following general limitations:

- Veeam Backup & Replication supports restore to Microsoft Azure for the following workloads:
  - Microsoft Windows workloads that run Windows Server 2008/Windows Vista and later.
  - Linux workloads (see the Supported Distributions & Versions section in Microsoft Docs).

    > **IMPORTANT**
    >
    > We strongly recommend to have `dracut` and `mkinitrd` installed on Linux machines that will be restored to Azure. Otherwise, they may not boot after restore.

- When you select a storage account whose resources you want to use to store disks of the restored workload, mind the following:
  - Storage accounts with the zone-redundant storage (ZRS), geo-zone-redundant storage (GZRS) and geo-redundant storage (GRS) replication options are not supported. However, read-access geo-redundant storage (RA-GRS) option is supported. For details on replication options, see Microsoft Docs.
  - If you plan to use a premium storage account and want to store unmanaged disks there, the restore speed for such disks will be limited to 30 MB/s (approximately).
- [Azure VMs] Veeam Backup & Replication does not support restore of Azure Ultra Disks.
- [Azure VMs] Mind the following limitations for disk sizes:
  - [Unmanaged VM disks] Veeam Backup & Replication supports restoring of disks equal to or less than 4093 GB. This is due to the following reasons: VM disks can increase in size up to 2 GB because of conversion during the restore process; Azure supports disk up to 4095 GB. For more information on all disk sizes that Azure supports, see Microsoft Docs.

- [Managed VM disks] Veeam Backup & Replication supports restoring disks equal or less than 4093 GB for OS disks and equal to or less than 32765 GB for other disks. During the restore process, VM disks can increase in size up to 2 GB because of conversion. For more information on all managed disk sizes that Azure supports, see Microsoft Docs. For more information on OS disk size that Azure supports, see Microsoft Docs. Note that supported disk sizes for Azure and Veeam Backup & Replication differ.

> **IMPORTANT**
>
> The price of a restored VM disk can become higher because of the increase in disk size during the restore process. For more information on pricing, see Managed Disks pricing and Unmanaged Disk and Page Blob pricing.

- [Azure Stack VMs] Veeam Backup & Replication supports restoring of managed and unmanaged disks equal to or less than 1021 GB. This is due to the following reasons: VM disks can increase in size up to 2 GB because of conversion during the restore process; Azure Stack supports disk up to 1023 GB. For more information on all disk sizes that Azure supports, see Microsoft Docs.

  You can change the maximum supported size for unmanaged VM disks with a registry key. For more information, contact Veeam Customer Support.

  > **IMPORTANT**
  >
  > The price of a restored VM disk can become higher because of the increase in disk size during the restore process. For more information on pricing, see Azure Stack Hub Pricing.

- If the system disk of an initial workload uses the GPT partitioning scheme, the number of partitions on the disk cannot exceed 4. During restore such disk will be converted to a disk with the MBR partitioning scheme.
- The restore to Microsoft Azure functionality does not support the Azure Hybrid Use Benefit program.
- [For restore from backups created with Veeam Agent for Microsoft Windows] Workloads from a backup that contains a failover cluster are restored as separate VMs, not as a cluster. Shared cluster disks of these VMs are restored as regular disks.

Requirements and Limitations for Generation 2 VM Support

Before enabling Generation 2 VM support, mind the following requirements and limitations:

- When you select VM sizes at the Specify VM Size step of the **Restore to Azure** wizard, make sure that the selected size is compatible with Generation 2 VMs.
- Generation 2 VMs support only managed disks. Thus, you will need to select the *managed* storage type from the **Storage type** list at the Specify VM Size step of the **Restore to Azure** wizard.
- Make sure that OSes of backed-up workloads that you plan to restore have UEFI boot. Otherwise, the restored VMs may be unbootable.

How Restore to Microsoft Azure Works

The restore process differs for Microsoft Windows and Linux workloads.

NOTE

Mind the following:

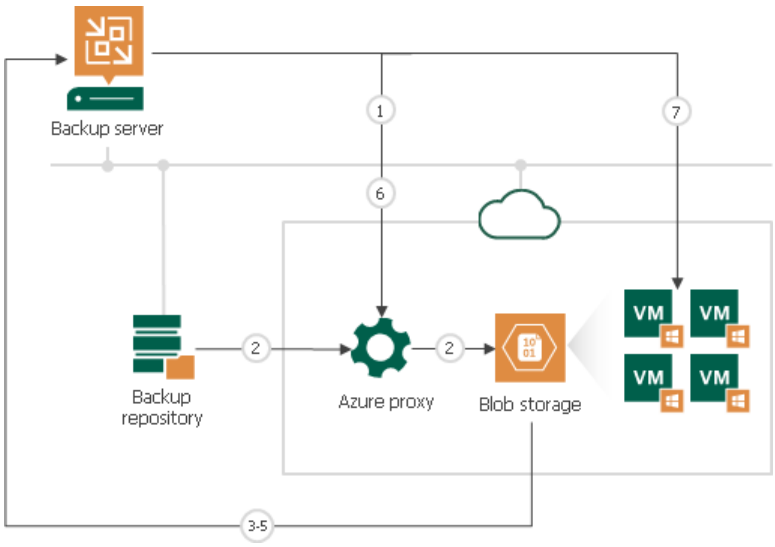
- If you added the Veeam Backup for Microsoft Azure appliance to the Veeam Backup & Replication infrastructure and you plan to restore Microsoft Azure VMs from restore points that were created using the appliance, the restore works as described in the Performing VM Restore section in Veeam Backup for Microsoft Azure User Guide.
- [If you restore from backups created by products other than Veeam Backup for Microsoft Azure] By default, during restore, Veeam Backup & Replication creates Generation 1 VMs. Such VMs support only Basic Input/Output System (BIOS) firmware interface. If you restore workloads with Unified Extensible Firmware Interface (UEFI), Veeam Backup & Replication converts UEFI into the BIOS firmware interface.

  For information on Generation 2 VM support, see Generation 2 VM Support.

Restore of Microsoft Windows Workloads

To restore a Microsoft Windows workload, Veeam Backup & Replication performs the following steps:

1. If you use an Azure proxy for restore, Veeam Backup & Replication powers on the Azure proxy. For more information about the Azure proxy, see Managing Azure Proxies.
2. Veeam Backup & Replication converts disks of a backed-up workload to the VHD format and uploads converted disks to Blob storage in Microsoft Azure.
3. Veeam Backup & Replication mounts uploaded disks to the backup server.
4. Veeam Backup & Replication prepares disks for restore. As part of this process, it enables Remote Desktop rules, configures firewall rules, prepares disks for Microsoft Azure agent installation and so on.
5. Veeam Backup & Replication unmounts prepared disks from the backup server.
6. If you use an Azure proxy for restore, Veeam Backup & Replication powers off the Azure proxy after a timeout.
7. Veeam Backup & Replication registers a Microsoft Azure VM with the prepared workload disks. After the registration process is complete, the Microsoft Azure VM is powered on immediately. Then the Microsoft Azure agent is installed on the VM.

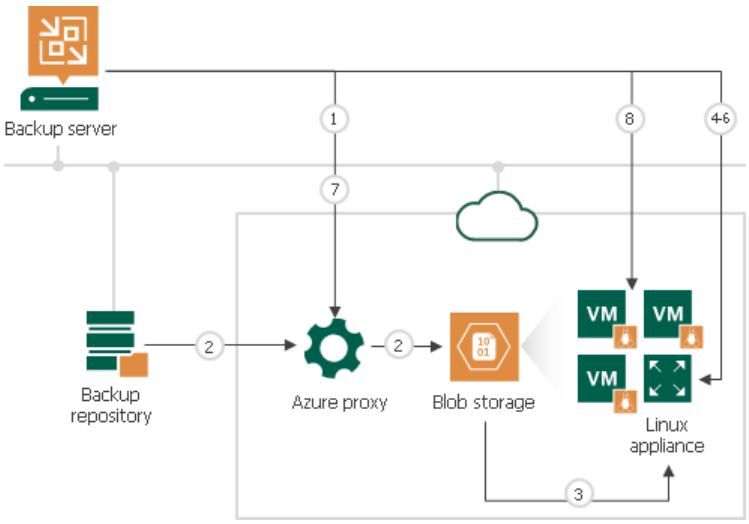


Restore of Linux Workloads

For restore of Linux workloads, Veeam Backup & Replication uses a helper appliance. The helper appliance is a small auxiliary Linux-based VM in Microsoft Azure registered by Veeam Backup & Replication. During the restore process, Veeam Backup & Replication mounts disks of a backed-up workload to the helper appliance to prepare disks for restore. For more information, see Managing Helper Appliances.

To restore a Linux workload, Veeam Backup & Replication performs the following steps:

1. If you use an Azure proxy for restore, Veeam Backup & Replication powers on the Azure proxy. For more information about the Azure proxy, see Managing Azure Proxies.
2. Veeam Backup & Replication converts disks of a backed-up workload to the VHD format and uploads converted disks to Blob storage in Microsoft Azure.
3. Veeam Backup & Replication mounts uploaded disks to the helper appliance that resides in the location to which you restore the Linux workload.
4. Veeam Backup & Replication starts the helper appliance with mounted disks.
5. Veeam Backup & Replication prepares disks for restore. As part of this process, it enables remote connection rules, configures firewall rules and so on.
6. Veeam Backup & Replication unmounts prepared disks from the helper appliance and powers off the helper appliance.
7. If you use an Azure proxy for restore, Veeam Backup & Replication powers off the Azure proxy after a timeout.
8. Veeam Backup & Replication registers a Microsoft Azure VM with the prepared workload disks. After the registration process is complete, the VM is powered on immediately.



Configuring Components and Accounts for Restore

Before you restore workloads, you must first add an account to be used for restore and then configure the required components:

1. Add a restore account:
   - Add a Microsoft Azure compute account if you want to restore workloads to Microsoft Azure.

- Add a Microsoft Azure Stack Hub compute account if you want to restore workloads to Azure Stack Hub.

These accounts must have specific built-in Azure roles (the roles are listed in the sections about adding the accounts). If you do not want to use built-in roles, you can create a custom role with granular permissions. For more information, see Creating Custom Role for Azure Account.

2. [For restore of Linux workloads] Configure helper appliances in Microsoft Azure.

> **NOTE**
>
> [If you have not changed default credentials for helper appliance] Before you configure the helper appliances, we recommend you to change the default credentials used during the helper appliance deployment. For more information on how to do this, see Changing Credentials for Helper Appliances.

3. [For restore process speed-up] Configure an Azure proxy.

> **NOTE**
>
> If you added the Veeam Backup for Microsoft Azure appliance to the Veeam Backup & Replication infrastructure and you plan to restore VMs from restore points that were created using the appliance, you do not need to configure the helper appliance and Azure proxy. Also, restore to Microsoft Azure works as described in the Performing VM Restore section in Veeam Backup for Microsoft Azure User Guide.

Creating Custom Role for Azure Account

If you do not want to use built-in Azure roles, you can create a custom role with granular permissions:

1. Run one of the following scripts in Azure PowerShell:

   Script for Az PowerShell

```
Import-Module Az.Resources
$role = [Microsoft.Azure.Commands.Resources.Models.Authorization.PSRoleDefiniti
on]::new()
$role.Name = 'Veeam Restore Operator'
$role.Description = 'Permissions for Veeam Restore to Microsoft Azure'
$role.IsCustom = $true
$permissions = @(
'Microsoft.Storage/storageAccounts/listkeys/action',
'Microsoft.Storage/storageAccounts/read',
'Microsoft.Network/locations/checkDnsNameAvailability/read',
'Microsoft.Network/virtualNetworks/read',
'Microsoft.Network/virtualNetworks/subnets/join/action',
'Microsoft.Network/publicIPAddresses/read',
'Microsoft.Network/publicIPAddresses/write',
'Microsoft.Network/publicIPAddresses/delete',
'Microsoft.Network/publicIPAddresses/join/action',
'Microsoft.Network/networkInterfaces/read',
'Microsoft.Network/networkInterfaces/write',
'Microsoft.Network/networkInterfaces/delete',
'Microsoft.Network/networkInterfaces/join/action',
'Microsoft.Network/networkSecurityGroups/read',
'Microsoft.Network/networkSecurityGroups/write',
'Microsoft.Network/networkSecurityGroups/delete',
'Microsoft.Network/networkSecurityGroups/join/action',
'Microsoft.Compute/locations/vmSizes/read',
'Microsoft.Compute/locations/usages/read',
'Microsoft.Compute/virtualMachines/read',
'Microsoft.Compute/virtualMachines/write',
'Microsoft.Compute/virtualMachines/delete',
'Microsoft.Compute/virtualMachines/start/action',
'Microsoft.Compute/virtualMachines/deallocate/action',
'Microsoft.Compute/virtualMachines/instanceView/read',
'Microsoft.Compute/virtualMachines/extensions/read',
'Microsoft.Compute/virtualMachines/extensions/write',
"Microsoft.Compute/virtualMachines/convertToManagedDisks/action",
'Microsoft.Compute/disks/read',
'Microsoft.Compute/disks/write',
"Microsoft.Compute/disks/beginGetAccess/action",
"Microsoft.Compute/disks/delete",
"Microsoft.Compute/disks/endGetAccess/action"
'Microsoft.Resources/checkResourceName/action',
'Microsoft.Resources/subscriptions/resourceGroups/read',
'Microsoft.Resources/subscriptions/resourceGroups/write',
'Microsoft.Resources/subscriptions/locations/read')
$role.Actions = $permissions
$role.NotActions = (Get-AzRoleDefinition -Name 'Virtual Machine Contributor').N
otActions
$subs = '/subscriptions/00000000-0000-0000-0000-000000000000'  #use your subscr
iption ID
$role.AssignableScopes = $subs
New-AzRoleDefinition -Role $role
```

NOTE

Mind the following:

- You must use `Connect-AzAccount` and `Get-AzSubscription` to input the subscription ID within the script.
- The script is provided for Az PowerShell module 6.0.0. The naming of commands may vary for other versions of Az PowerShell module.

Script for legacy AzureRM PowerShell

```
Import-Module AzureRm.Resources
$role = [Microsoft.Azure.Commands.Resources.Models.Authorization.PSRoleDefiniti
on]::new()
$role.Name = 'Veeam Restore Operator'
$role.Description = 'Permissions for Veeam Restore to Microsoft Azure'
$role.IsCustom = $true
$permissions = @(
'Microsoft.Storage/storageAccounts/listkeys/action',
'Microsoft.Storage/storageAccounts/read',
'Microsoft.Network/locations/checkDnsNameAvailability/read',
'Microsoft.Network/virtualNetworks/read',
'Microsoft.Network/virtualNetworks/subnets/join/action',
'Microsoft.Network/publicIPAddresses/read',
'Microsoft.Network/publicIPAddresses/write',
'Microsoft.Network/publicIPAddresses/delete',
'Microsoft.Network/publicIPAddresses/join/action',
'Microsoft.Network/networkInterfaces/read',
'Microsoft.Network/networkInterfaces/write',
'Microsoft.Network/networkInterfaces/delete',
'Microsoft.Network/networkInterfaces/join/action',
'Microsoft.Network/networkSecurityGroups/read',
'Microsoft.Network/networkSecurityGroups/write',
'Microsoft.Network/networkSecurityGroups/delete',
'Microsoft.Network/networkSecurityGroups/join/action',
'Microsoft.Compute/locations/vmSizes/read',
'Microsoft.Compute/locations/usages/read',
'Microsoft.Compute/virtualMachines/read',
'Microsoft.Compute/virtualMachines/write',
'Microsoft.Compute/virtualMachines/delete',
'Microsoft.Compute/virtualMachines/start/action',
'Microsoft.Compute/virtualMachines/deallocate/action',
'Microsoft.Compute/virtualMachines/instanceView/read',
'Microsoft.Compute/virtualMachines/extensions/read',
'Microsoft.Compute/virtualMachines/extensions/write',
"Microsoft.Compute/virtualMachines/convertToManagedDisks/action",
'Microsoft.Compute/disks/read',
'Microsoft.Compute/disks/write',
"Microsoft.Compute/disks/beginGetAccess/action",
"Microsoft.Compute/disks/delete",
"Microsoft.Compute/disks/endGetAccess/action"
'Microsoft.Resources/checkResourceName/action',
'Microsoft.Resources/subscriptions/resourceGroups/read',
'Microsoft.Resources/subscriptions/resourceGroups/write',
'Microsoft.Resources/subscriptions/locations/read')
$role.Actions = $permissions
$role.NotActions = (Get-AzureRmRoleDefinition -Name 'Virtual Machine Contributo
r').NotActions
$subs = '/subscriptions/00000000-0000-0000-0000-000000000000'  #use your subscr
iption ID
$role.AssignableScopes = $subs
New-AzureRmRoleDefinition -Role $role
```

NOTE

Mind the following:

- You must use `Connect-AzureRmAccount` and `Get-AzureRmSubscription` to input the subscription ID within the script.
- The script is provided for Microsoft Azure PowerShell version 5.1.1 bundled with Veeam Backup & Replication. The naming of commands may vary for other versions of Microsoft Azure PowerShell.

2. Assign the created role to the required Azure User. For details, see the Manage access to Azure resources using RBAC and the Azure portal section in the RBAC for Azure resources documentation.
3. In the **Subscription** step of the **Microsoft Azure Compute Account** wizard, select **Use existing account** and select the Azure user with the assigned role. For details, see Adding Microsoft Azure Compute Account.

Reference

Create Custom Roles Using Azure PowerShell

Managing Helper Appliances

To restore Linux workloads to Microsoft Azure, Veeam Backup & Replication requires helper appliances. A helper appliance is a small auxiliary virtual machine that Veeam Backup & Replication automatically deploys in Microsoft Azure. During the restore process, Veeam Backup & Replication mounts disks of the restored workload to the helper appliance to prepare these disks for restore.

Helper appliances are persistent. After the restore process finishes, helper appliances get powered off and remain in Microsoft Azure. The appliances remain in the powered off state until you start a new restore process. Note that Microsoft Azure will bill you for storing helper appliances disks in the storage account. To remove a helper appliance, follow the instruction provided in Removing Helper Appliances.

Configuring Helper Appliances

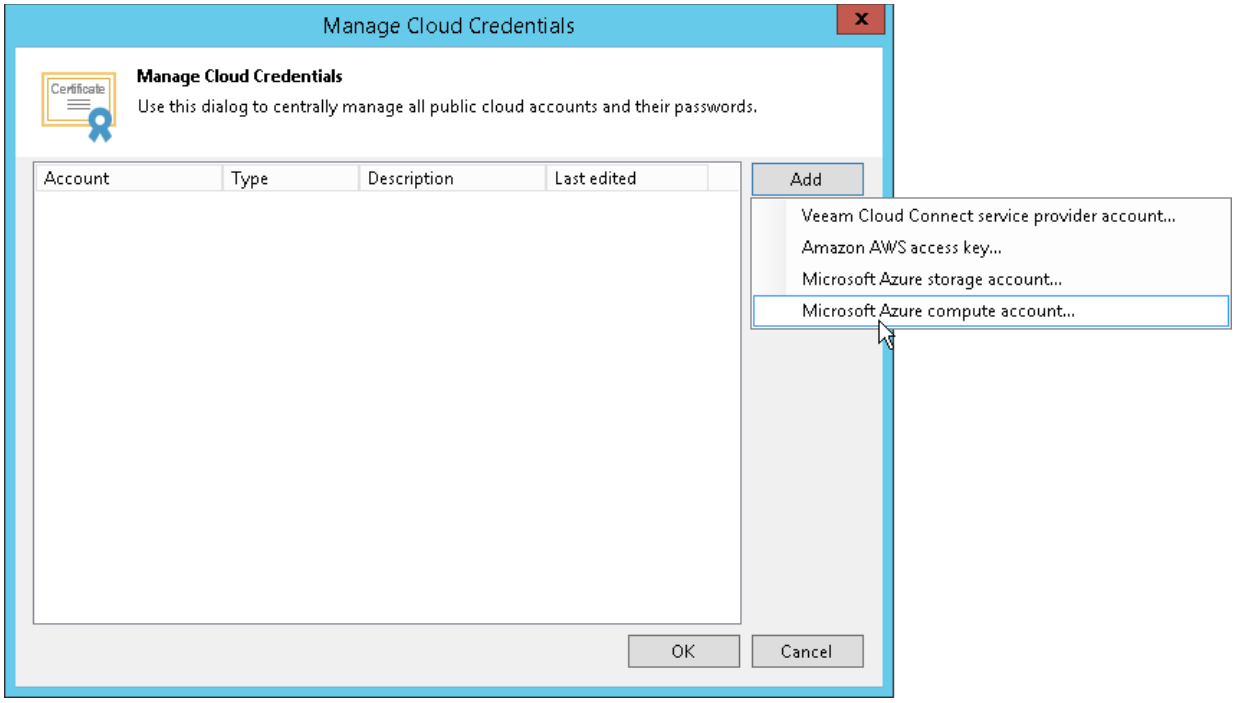
Before you configure helper appliances, mind the following:

- Veeam Backup & Replication uses its built-in credentials record to work with all helper appliances. For security reasons, we recommended that you change a password for this account before you set up the helper appliances. Changing credentials is required only once. For more information, see Changing Credentials for Helper Appliances.
- If you plan to restore Linux workloads to different locations, you must configure a helper appliance in each location to which workloads will be restored.

To configure a helper appliance:

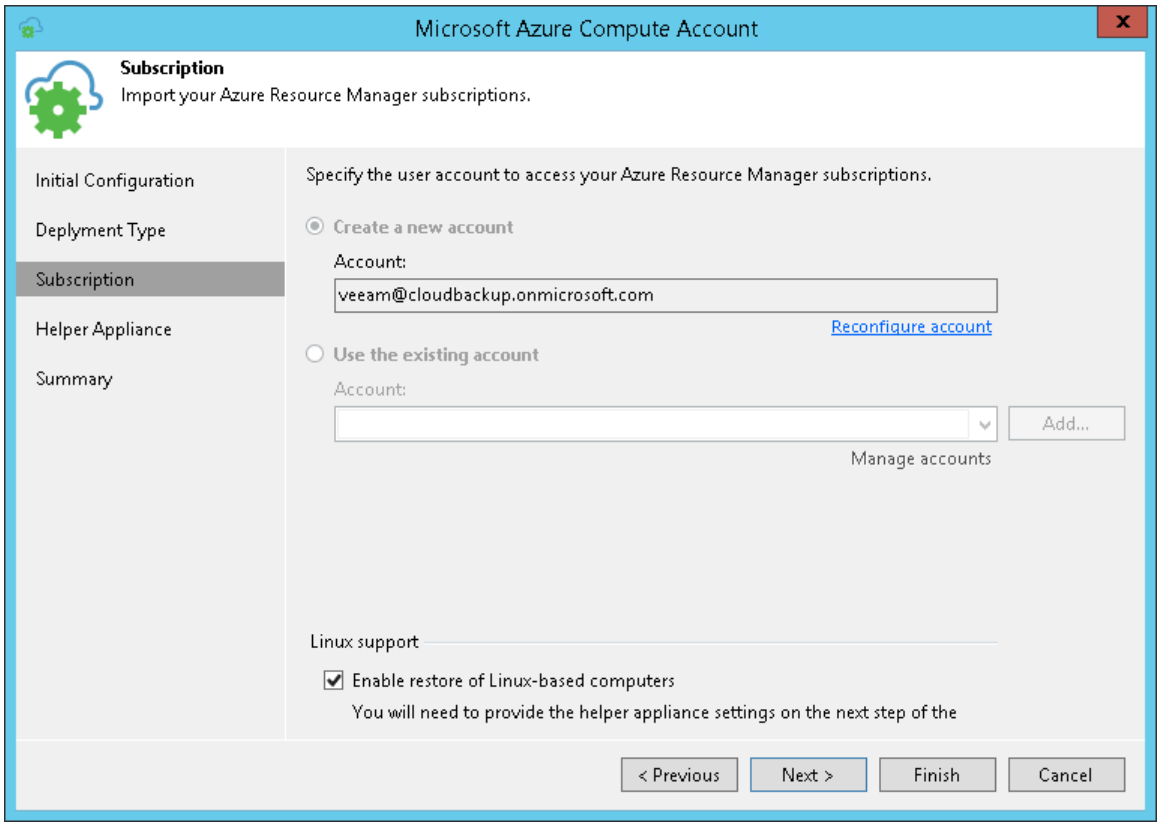
1. From the main menu, select **Manage Cloud Credentials**.

2. In the **Manage Cloud Credentials** window, click **Add** > **Microsoft Azure compute account**.



3. At the **Initial Configuration** step of the **Microsoft Azure Compute Account** wizard, click **Next**.
4. At the **Deployment Type** step of the wizard, select the necessary deployment type.

   For more information on when each deployment type is used, see Microsoft Azure Compute Accounts and Microsoft Azure Stack Hub Compute Accounts.
5. At the **Subscription** step of the wizard, select the **Enable restore of Linux-based computers** check box.



6. At the **Helper Appliance** step of the wizard, configure settings of a helper appliance. Note that if you plan to restore Linux workloads to multiple locations, you must configure a helper appliance in each location.

   To configure one helper appliance, do the following:

   a. On the right of the **Helper appliances** list, click **Add**.

   b. From the **Subscription** list, select a subscription whose resources you want to use to configure the helper appliance. The subscription list contains all subscriptions that are associated with the Azure compute account or Azure Stack Hub account.

   c. From the **Location** list, select a location where you want to configure a helper appliance. Make sure that you select a geographic region with which at least one storage account of the subscription is associated.

   d. From the **Storage account** list, select a storage account whose resources you want to use to store disks of the helper appliance.

   > **NOTE**
   >
   > You cannot use a storage account with the ZRS or GZRS replication option. For details, see Microsoft Docs.

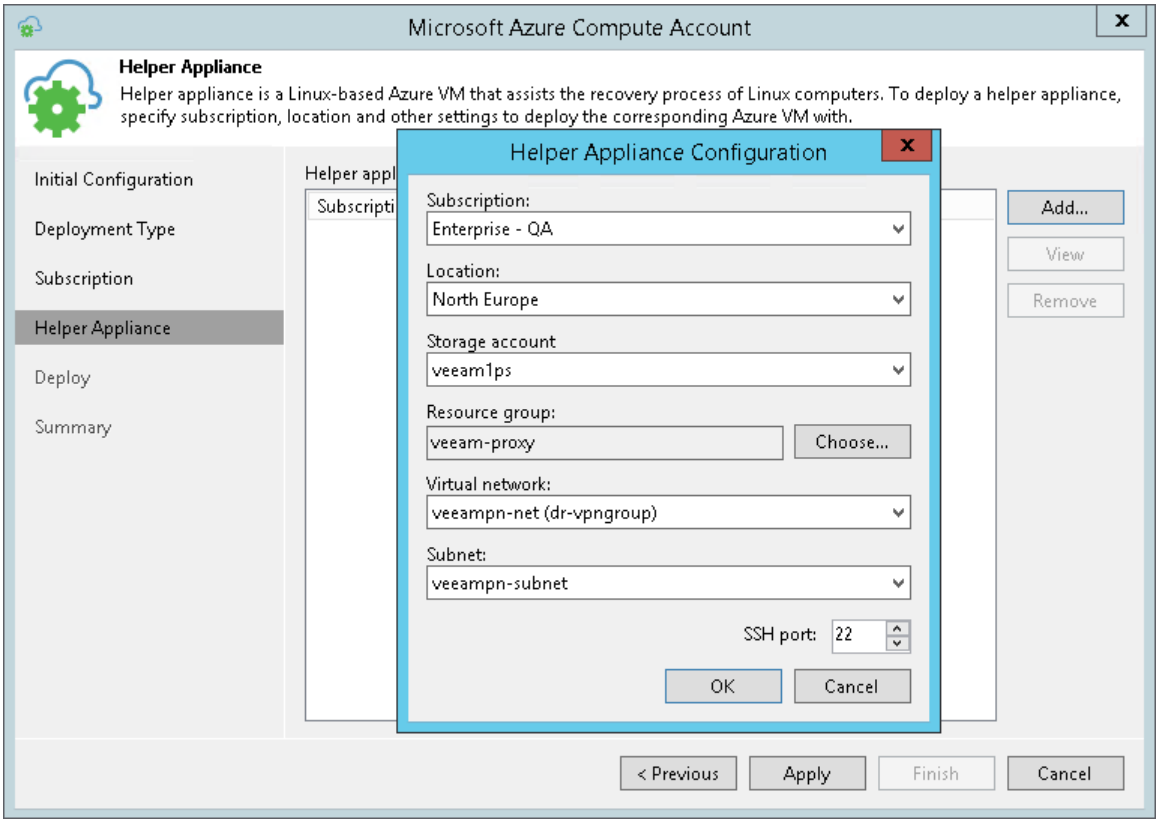
   e. Click **Choose** if you do not want Veeam Backup & Replication to create a new resource group.

   f. From the **Virtual network** list, select a network to which the helper appliance must be connected.

   g. From the **Subnet** list, select a subnet for the helper appliance.

   h. At the **SSH port** field, specify a port over which Veeam Backup & Replication will communicate with the helper appliance.

   i. Click **OK**.



7. At the **Deploy** step of the wizard, wait until Veeam Backup & Replication configures the helper appliance in Microsoft Azure. The appliance will remain powered off until you start restoring Linux workloads to the necessary location in Microsoft Azure.
8. At the **Summary** step of the wizard, review summary information and click **Finish** to close the wizard.

Changing Credentials for Helper Appliances

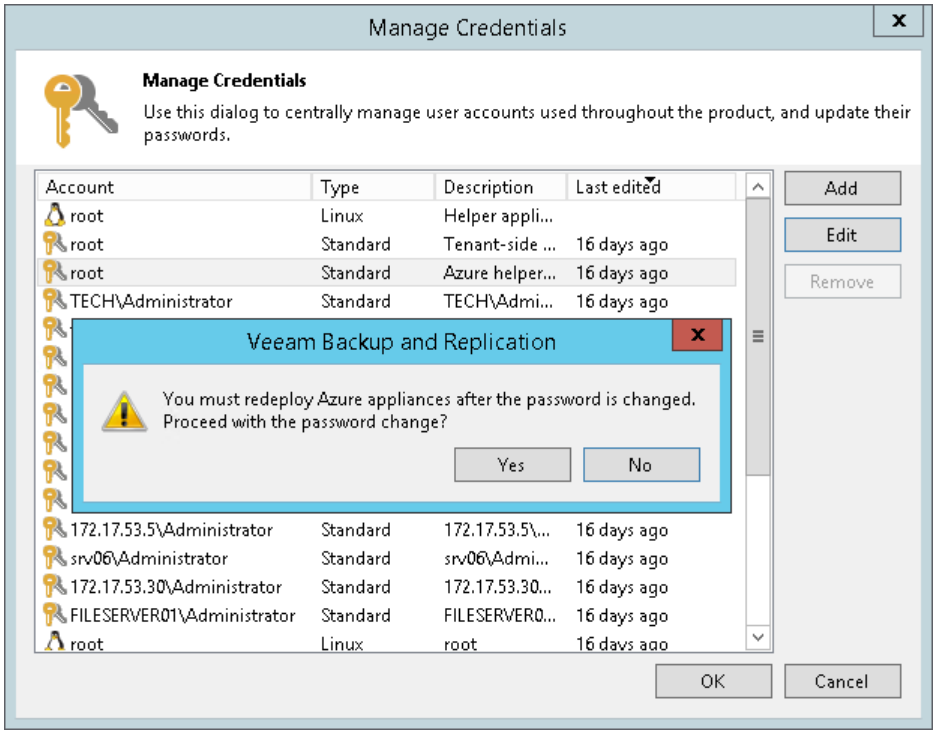
By default, Veeam Backup & Replication uses its built-in credentials record to work with all helper appliances in Microsoft Azure and Azure Stack Hub. You can find this credential record in the credentials manger in the Veeam Backup & Replication console: the *root* account with "Azure helper appliance credentials" in the **Description** column.

For security reasons, we recommend you to change the password for this credentials record before you set up helper appliances.

> **IMPORTANT**
>
> When you change the password in the built-in credentials record, you must re-deploy all existing helper appliances in Microsoft Azure and Azure Stack Hub. To redeploy appliances, you must remove all configured appliances and then configure them once again.

To change the password in the credentials record for the helper appliances:

1. From the main menu, select **Manage Credentials**.
2. In the **Manage Credentials** window, click the built-in credentials record for the Azure helper appliances.
3. Click **Edit**.
4. In the **Password** field, specify a new password.
5. Click **OK** to save changes.

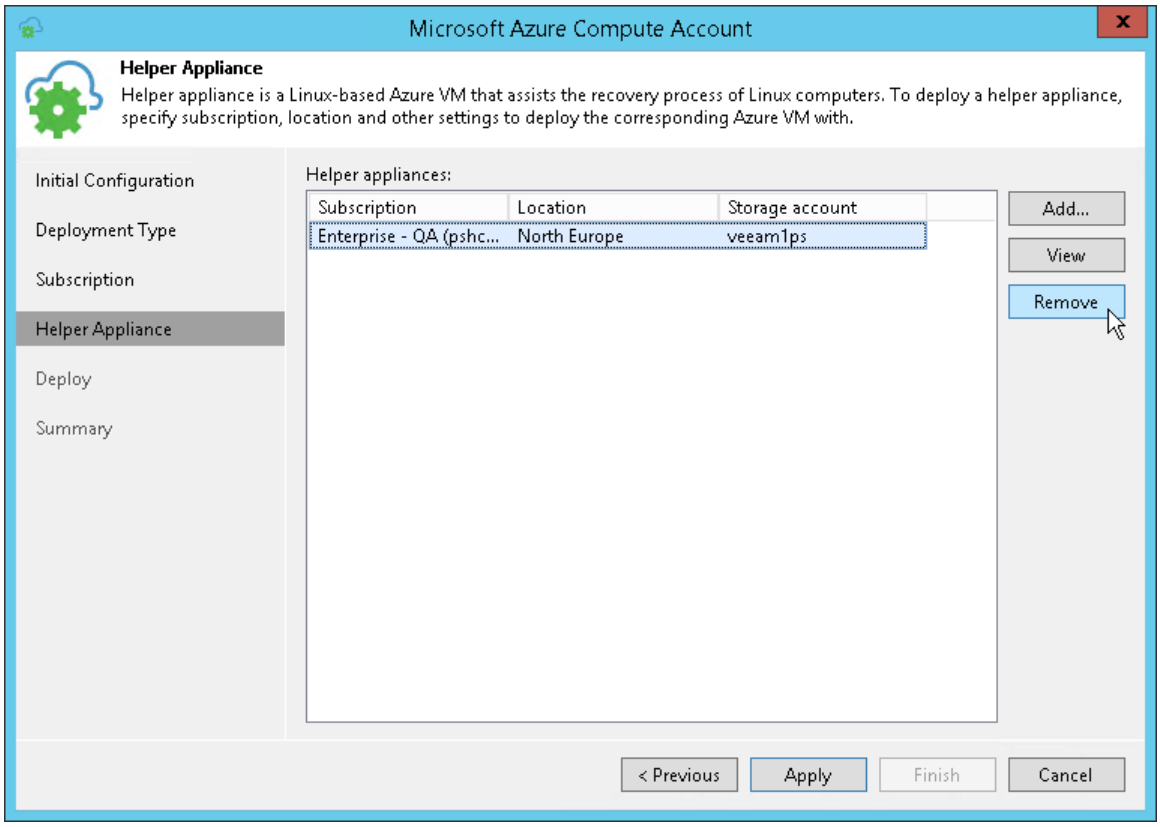


Removing Helper Appliances

If you no longer need to restore Linux workloads to Azure or Azure Stack Hub, you can remove helper appliances:

1. From the main menu, select **Manage Cloud Credentials**.
2. In the accounts list, select the Azure account and click **Edit**.
3. Pass to the **Helper Appliance** step of the **Microsoft Azure Compute Account** wizard.
4. In the **Helper appliances** list, select the helper appliance and click **Remove**.

> **IMPORTANT**
>
> Do not clear the **Enable restore of Linux-based computers** check box at the **Subscription** step of the wizard to remove helper appliances. In this case, the **Microsoft Azure Compute Account** wizard will not display the **Helper Appliance** step. Helper appliances themselves will remain in Microsoft Azure.



Managing Azure Proxies

Microsoft Azure proxies help speed up the restore process especially if you restore workloads to a distant location or the network connection is slow. The Azure proxy is a small auxiliary machine in Microsoft Azure over which Veeam Backup & Replication transports VM disk data to Blob storage. Veeam components installed on the Azure proxy compress and deduplicate disk data, which helps reduce network traffic.

Although Azure proxies are optional, we recommend you to configure them as Azure proxies do not require a lot of resources but can significantly improve restore performance. Configure an Azure proxy in a location to which you plan to restore workloads or close to this location. If you plan to restore workloads to different locations, configure at least one Azure proxy in each location.

The process of Azure proxy deployment takes some time. It is recommended that you configure the Azure proxy in advance, before you start the restore process. To configure an Azure proxy, use the **Azure Proxy** wizard as described in Configuring Azure Proxies. Veeam Backup & Replication will deploy a Microsoft Windows Server 2012 R2 machine in Microsoft Azure and assign the role of the Azure proxy to this machine. You can then instruct Veeam Backup & Replication to use the Azure proxy for restore tasks.

Configuring Azure Proxies

To configure an Azure proxy, use the **Azure Proxy** wizard.

Before You Begin

Before you configure an Azure proxy, check the following prerequisites:

- You must add information about your Microsoft Azure compute account or Microsoft Azure Stack Hub compute account to Veeam Backup & Replication. For more information, see Adding Microsoft Azure Compute Accounts or Adding Azure Stack Hub Compute Accounts.
- You must configure the following objects in Microsoft Azure beforehand:
  - Storage account whose resources you plan to use to store disks of the Azure proxy.
  - Networks to which you plan to connect the Azure proxy.

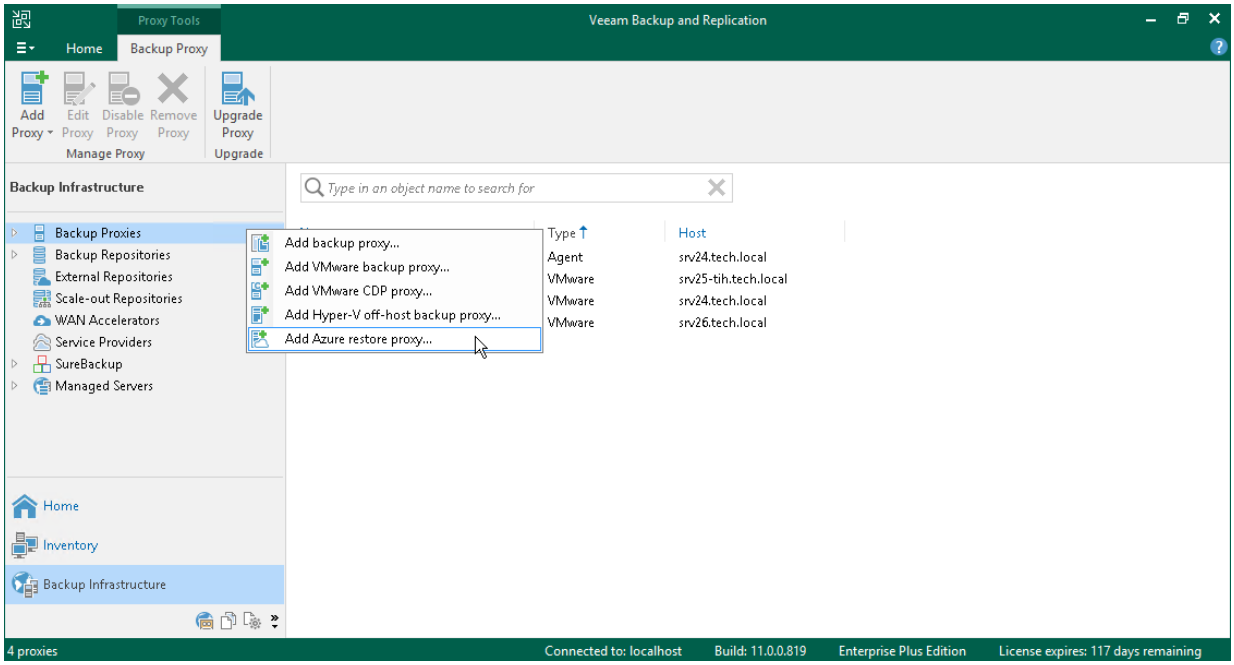
  For storage accounts and network configuration, you must use the same deployment model that you plan to use for Azure proxy creation.

> **IMPORTANT**
>
> When you deploy Azure proxy for Azure Stack Hub, make sure that Windows Server 2012 R2 is available in Azure marketplace.

Step 1. Launch Azure Proxy Wizard

To launch the **Azure Proxy** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane, select the **Backup Proxies node. On the ribbon,** click **Add Proxy** > **Azure restore proxy**.
- Open the **Backup Infrastructure** view. In the inventory pane, right-click the **Backup Proxies** node and select **Add Azure restore proxy**.

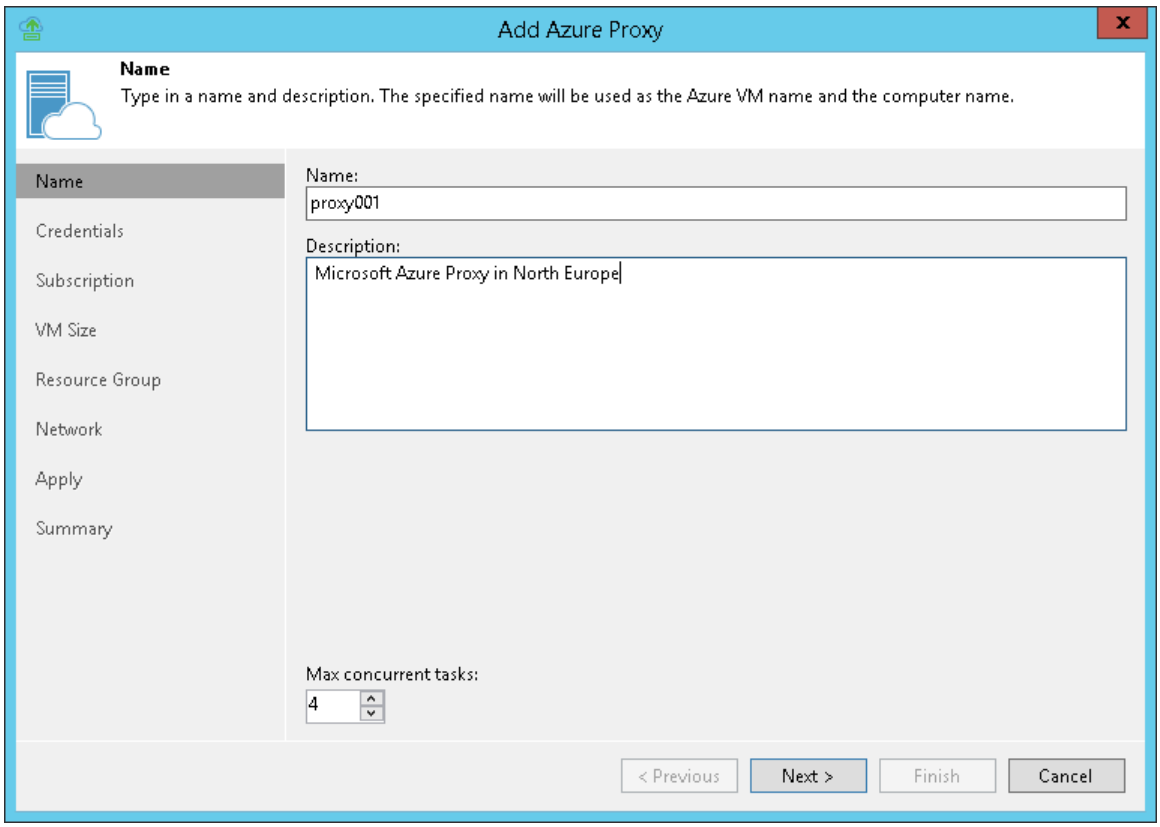


Step 2. Specify Azure Proxy Name

At the **Name** step of the wizard, specify a name and description for the Azure proxy:

1. In the **Name** field, specify a name for the Azure proxy. The name must meet the following requirements:
   - The name must not be longer than 15 characters.
   - The name must contain only alphanumeric characters and hyphens.
   - The name must start with a letter and end with a letter or number.
   - The name must not contain only numeric characters.
   - The name must not contain special characters: *`!@#$%^&*()+=_[]{}|;:.'",<>/?*.
2. In the **Description** field, provide a description for the Azure proxy.
3. At the **Max concurrent tasks** field, specify the number of tasks that the Azure proxy must handle in parallel. If the **Max concurrent tasks** value is exceeded, the Azure proxy will not start a new task until one of current tasks finishes.

   Veeam Backup & Replication creates one task per one workload disk. By default, Azure proxy handles 4 concurrent tasks.



Step 3. Specify Credentials and Transport Port

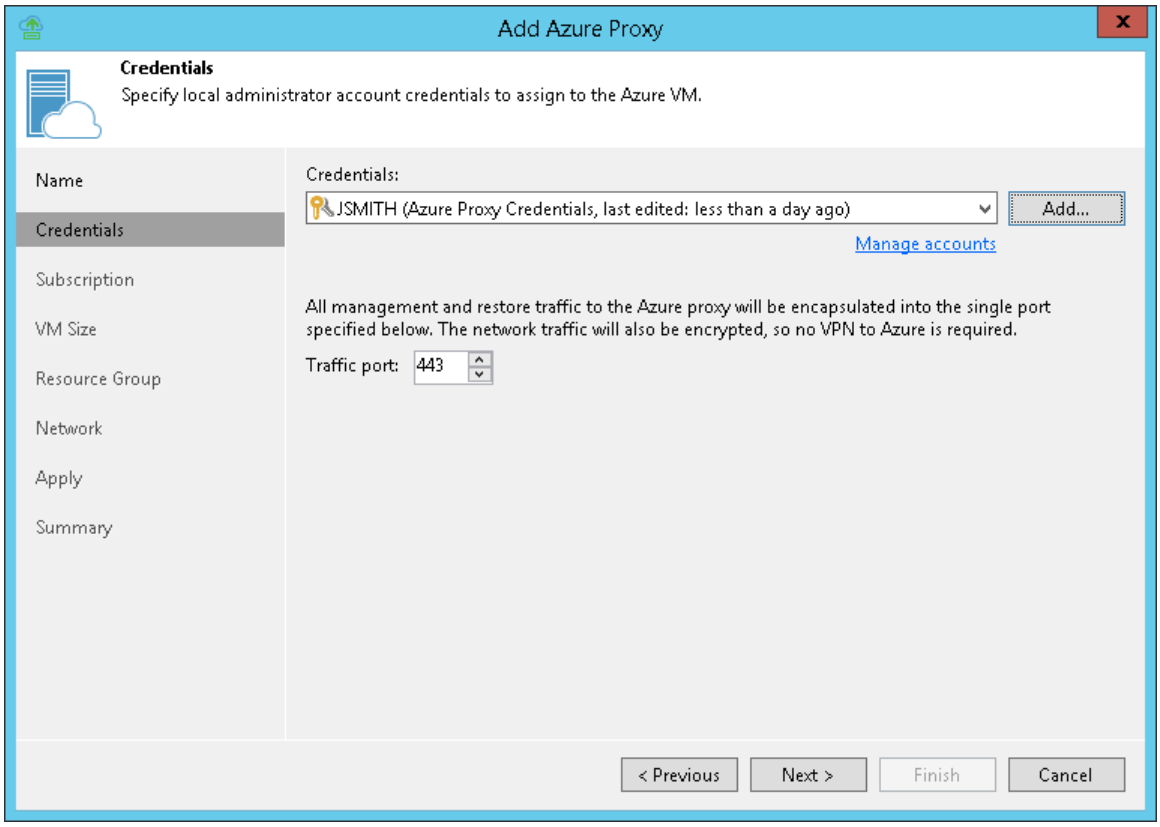
At the **Credentials** step of the wizard, specify credentials of the local administrator account on the Azure proxy and define the transport port:

1. From the **Credentials** list, select credentials of a user that will be assigned the Local Administrator permissions on the Azure proxy.

   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.

   > **IMPORTANT**
   >
   > Mind the following:
   >
   > - You cannot use reserved names such as 'administrator', 'admin', 'user', 'abc@123', 'P@$$w0rd' and so on as a user name and password of the local administrator account.
   > - You must specify the user name without a domain or Microsoft Azure machine name.
   > - The password must be at least 8 characters long, and must contain at least 1 uppercase character, 1 lowercase character, 1 numeric character and 1 special character.

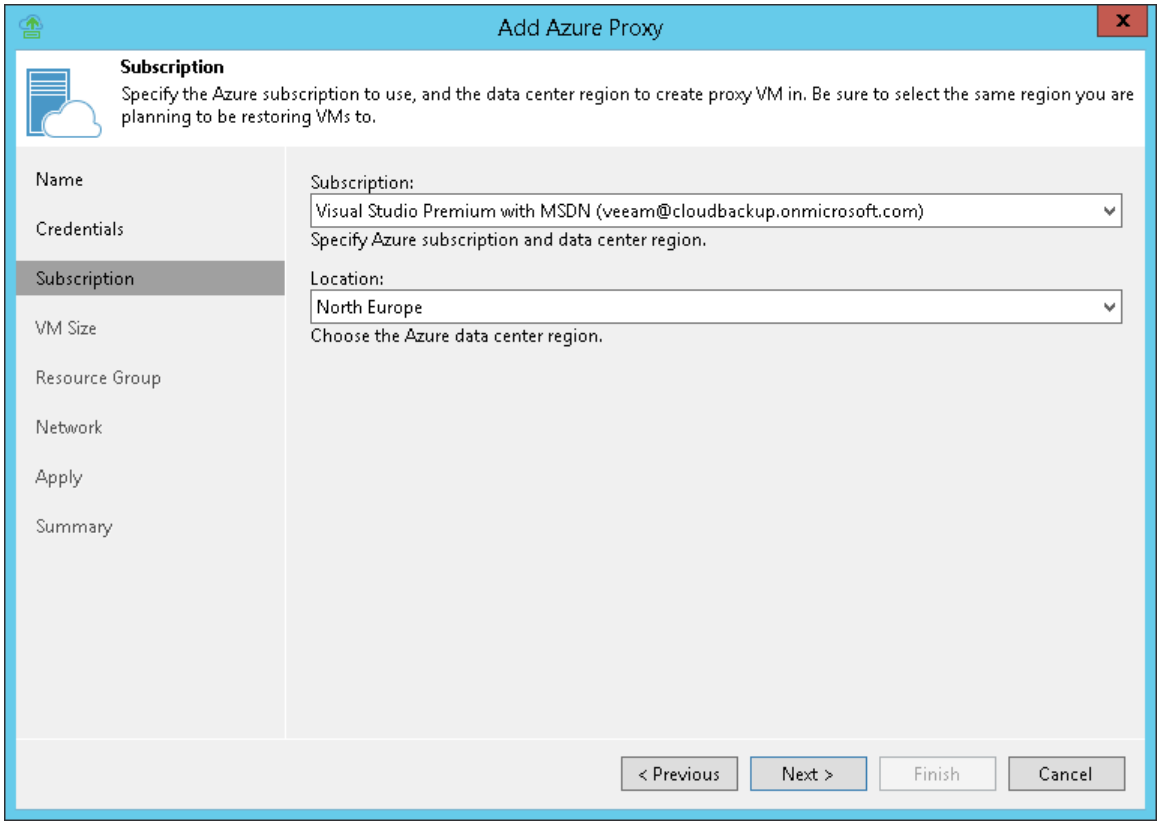
2. In the **Traffic port** field, specify a port over which Veeam Backup & Replication will control components installed on the Azure proxy and transport workload disks data to Blob storage. The port must be opened on the backup server and backup repository that stores backups.



Step 4. Select Subscription and Location

At the **Subscription** step of the wizard, select a subscription and location for the Azure proxy:

1. From the **Subscription** list, select a subscription whose resources you want to use to deploy the Azure proxy. The subscription list contains all subscriptions associated with Azure or Azure Stack Hub compute accounts that you added to Veeam Backup & Replication.
2. From the **Locations** list, select a geographic region to which you want to place the Azure proxy. Make sure that you select a geographic region with which at least one storage account of the subscription is associated.



Step 5. Select VM Size

At the **VM size** step of the wizard, you can select the size for the Azure proxy VM and specify which storage account you want to use to deploy the Azure proxy VM:

1. From the **Size** list, select the size for the Azure proxy.

   The default size is *Standard_F4s_v2*. If *Standard_F4s_v2* is not available, we recommend to select a similar compute optimized VM size of the previous generation — F-series. For example, *Standard_F4*, *Standard_F4s*. These sizes are sufficient to transport VM disks data to Blob storage. If necessary, you can select a greater size for the Azure proxy.

   If you select a premium storage account, make sure that the VM size is compatible with the selected account.

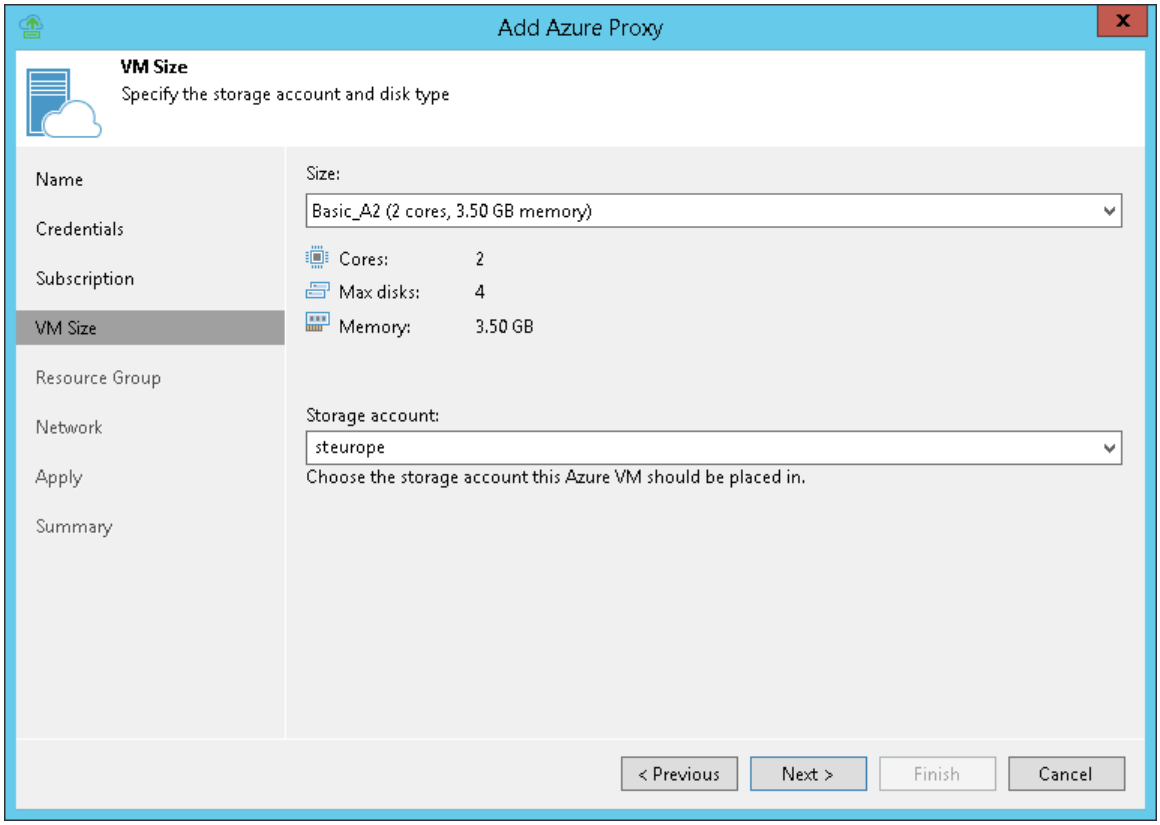
   > **NOTE**
   >
   > Azure proxy VMs created in Veeam Backup & Replication version prior 10a have smaller sizes — *Basic_A2*. We recommend you to change sizes of such proxies to the sizes listed above. This will enhance the performance of restore to Azure.
   >
   > You can change VM sizes in Microsoft Azure Portal or deploy new proxies with the required sizes in the Veeam Backup & Replication.

2. From the **Storage account** list, select a storage account whose resources you want to use to store disks of the Azure proxy. The storage account must be compatible with the VM size you select.

   The list of storage accounts will contain only general purpose storage accounts. Blob storage accounts will not be displayed in the list of subscriptions. For more information about account types, see Microsoft Docs.

   > **NOTE**
   >
   > You cannot use a storage account with the ZRS or GZRS replication option for the Azure proxy. For details, see Microsoft Docs.

   > **TIP**
   >
   > Microsoft Azure subscriptions have default limits on the number of CPU cores. Make sure that the VM size you select does not exceed limits of the subscription.


Add Azure Proxy
VM Size
Specify the storage account and disk type
Name
Credentials
Subscription
VM Size
Resource Group
Network
Apply
Summary
Size:
Basic_A2 (2 cores, 3.50 GB memory)
Cores: 2
Max disks: 4
Memory: 3.50 GB
Storage account:
steurope
Choose the storage account this Azure VM should be placed in.
< Previous
Next >
Finish
Cancel

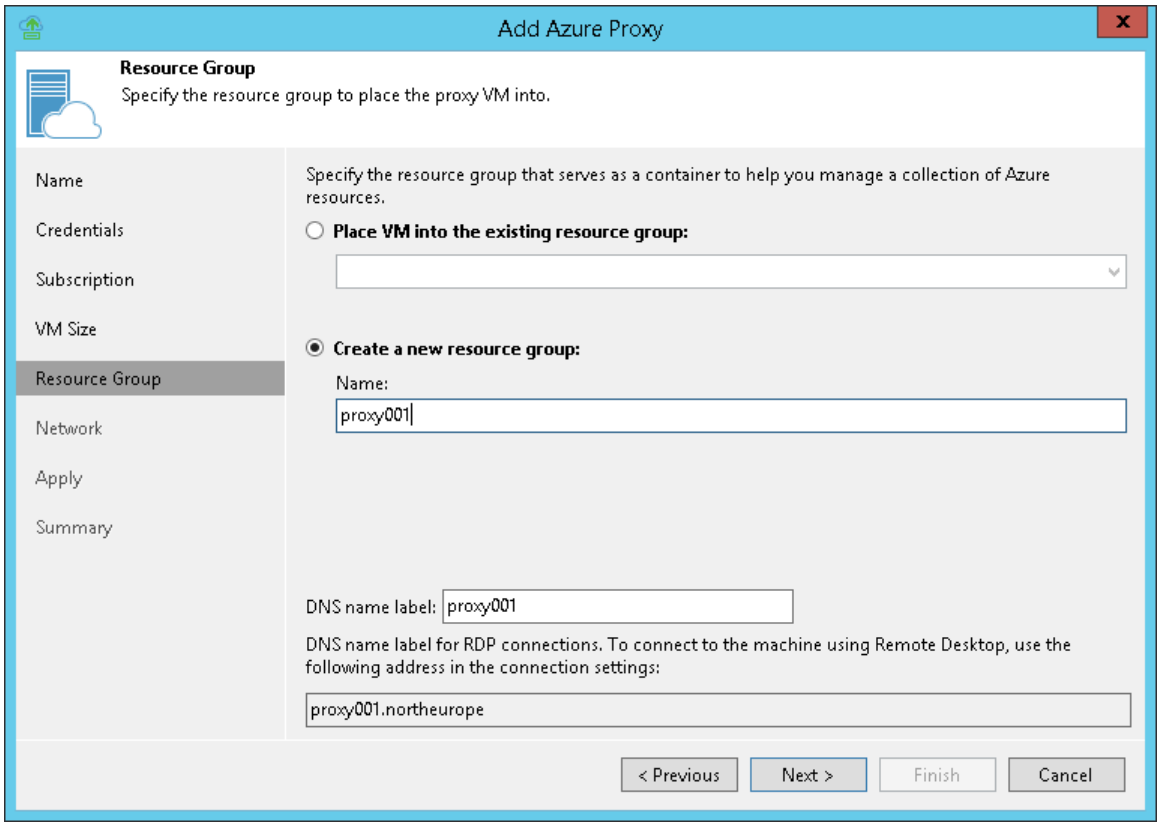
Step 6. Select Resource Group

At the **Resource Group** step of the wizard, you specify the resource group to which the Azure proxy must be placed and configure DNS name label:

1. You can place the Azure proxy to an existing or new resource group:
   - Select **Place VM into the existing resource group** to place the Azure proxy to an existing resource group. From the drop-down list, select the necessary resource group.
   - Select **Create a new resource group** to create a dedicated resource group for the Azure proxy. In the **Name** field, enter a name for the new resource group. The resource group name can be up to 64 characters long and can contain only alphanumeric, underscore and hyphen characters.

     For the new resource group, Veeam Backup & Replication automatically creates a network security group, dynamic public IP and network interface.
2. In the **DNS name label** field, enter a name of the dynamic public IP. The DNS name label can be up to 80 characters long, and can contain only alphanumeric, dash and underscore characters. For more information, see Microsoft Docs.

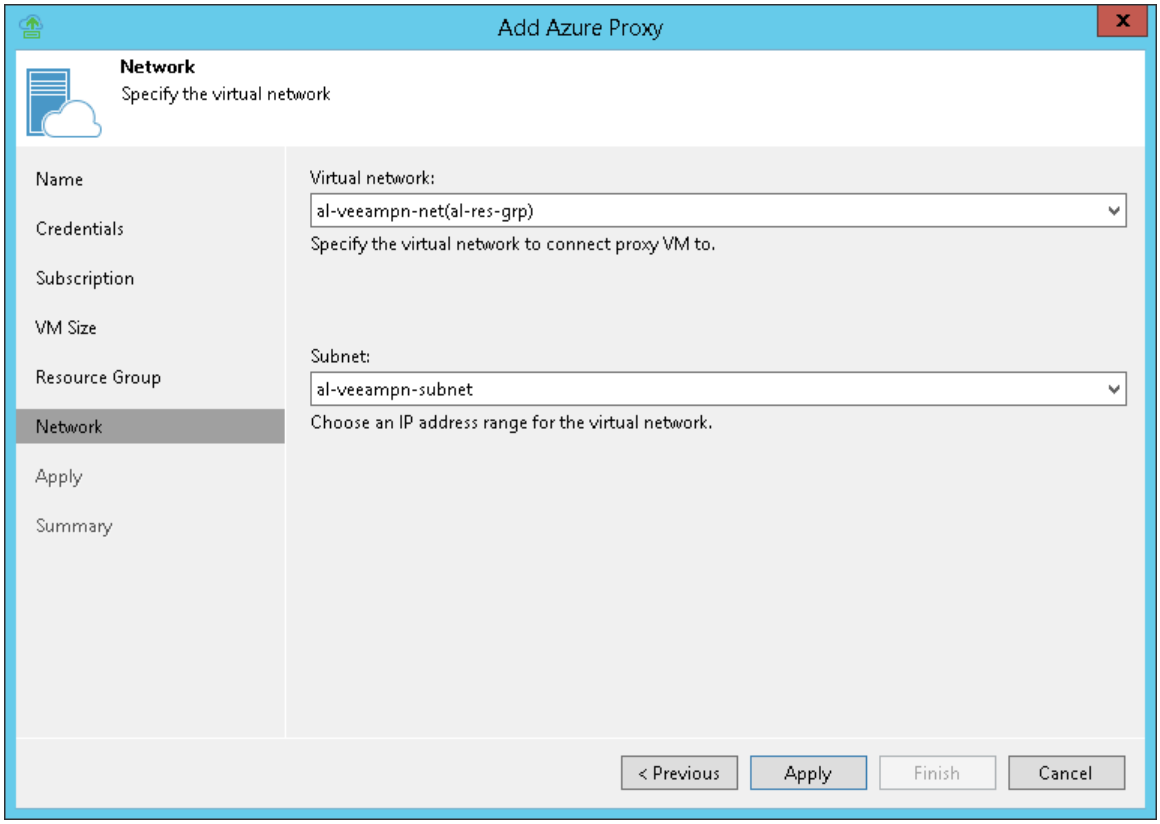
> **TIP**
>
> Microsoft Azure subscriptions have default limits on the number of resource groups. If you decide to create a new resource group, make sure that you do not exceed limits of the subscription.



Step 7. Select Virtual Network

At the **Network** step of the wizard, you select to which network and subnet the Azure proxy will be connected.

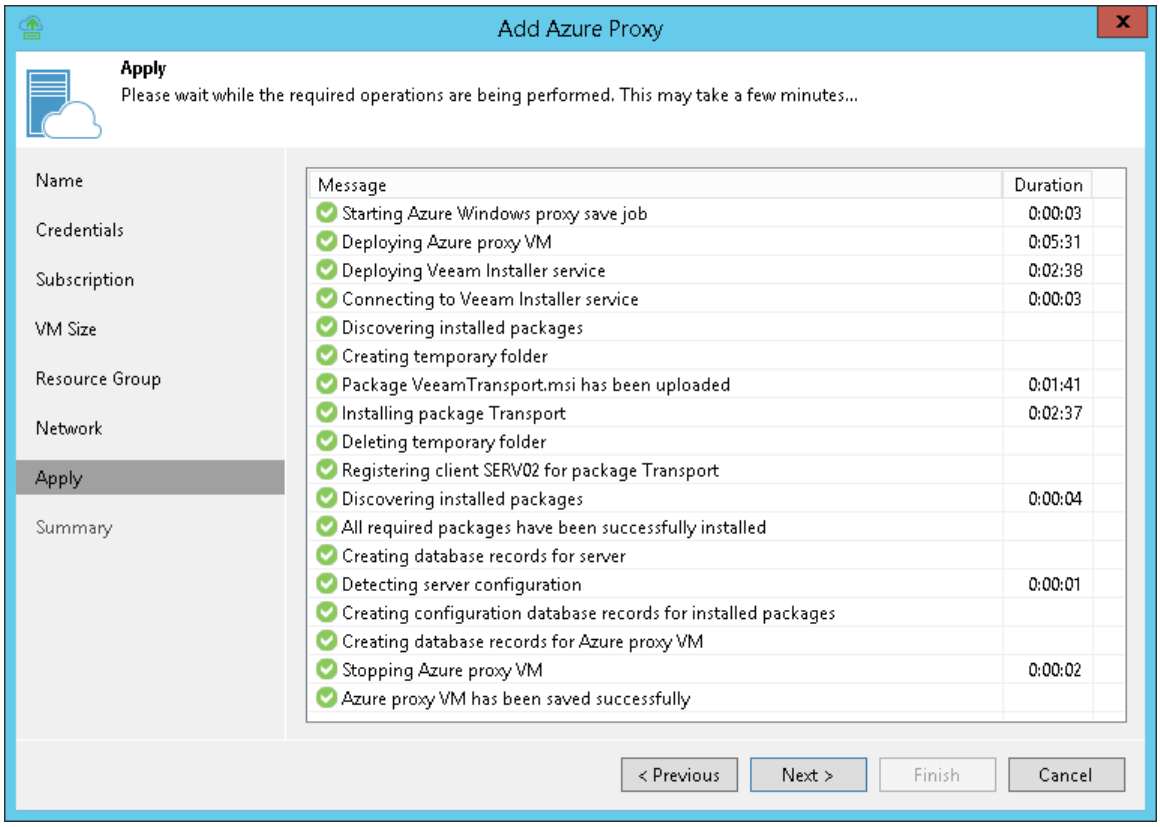
> **IMPORTANT**
>
> If you want to restore from backups in an on-premises object storage repository, the selected virtual network must have access to the source object storage repository. To provide access to object storage repositories, you can use VPN or Azure ExpressRoute.



Step 8. Start Azure Proxy Deployment

At the **Apply** step of the wizard, Veeam Backup & Replication deploys the Azure proxy with the specified settings. You can view the deployment progress in the real-time mode. When the configuration process is over, click **Next**. At the **Summary** step of the wizard, click **Finish** to close the wizard.

> **TIP**
>
> The Azure proxy deployment may take several minutes. You can close the **Azure Proxy** wizard and continue working with Veeam Backup & Replication while the proxy is being deployed. To view the deployment progress, open the **History** view, in the inventory pane select **System**, and double-click the task of the proxy deployment in the working area.



Removing Azure Proxies

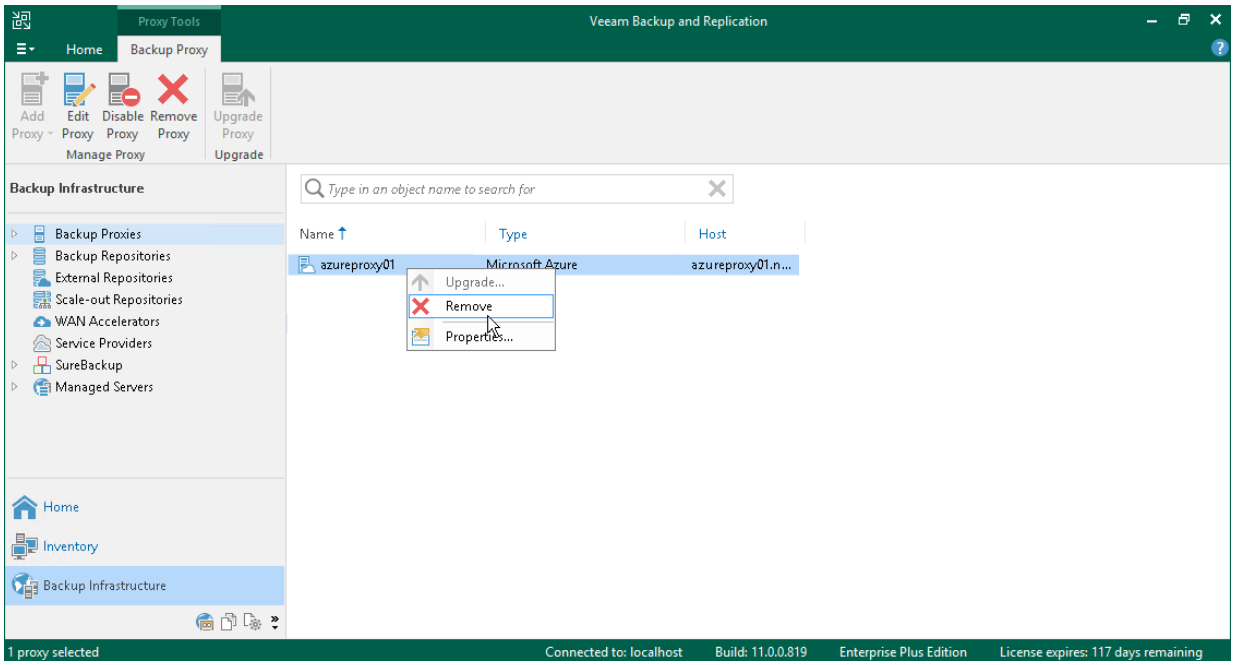
Veeam Backup & Replication does not provide a possibility to edit settings of deployed Azure proxies. If you want to change Azure proxy configuration, remove the Azure proxy and create a new proxy.

To remove an Azure proxy, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, right-click the Azure proxy and select **Remove**.

IMPORTANT

If you want to remove an Azure or an Azure Stack Hub compute account from Veeam Backup & Replication, you must remove all Azure proxies first.



Restoring to Microsoft Azure

Before you restore workloads to Microsoft Azure, you must configure an account to be used for restore and required components. For more information, see Configuring Components and Accounts for Restore. Then use the **Restore to Azure** wizard to restore the workloads.

NOTE

If you added the Veeam Backup for Microsoft Azure appliance to the Veeam Backup & Replication infrastructure and you plan to restore Microsoft Azure VMs from restore points that were created using the appliance, the steps of the restore wizard differ from the steps described in this guide. For more information, see the Restoring to Microsoft Azure section in the Integration with Veeam Backup for Microsoft Azure Guide.

Before You Begin

Before you begin restore to Microsoft Azure, check the following prerequisites:

- Check limitations listed in Limitations for Restore to Microsoft Azure.
- You must create a backup of the workload that you want to restore in Microsoft Azure. For the list of supported backups, see Restore to Microsoft Azure.
- A backup chain from which you plan to restore a workload must reside in a backup repository added to the backup infrastructure. You can also import a backup to the Veeam Backup & Replication console. For more information, see Importing Backups.

- You must configure the following objects in Microsoft Azure beforehand:
  - Storage account whose resources you plan to use to store disks of the restored workload.
  - Networks to which you plan to connect the restored workload.

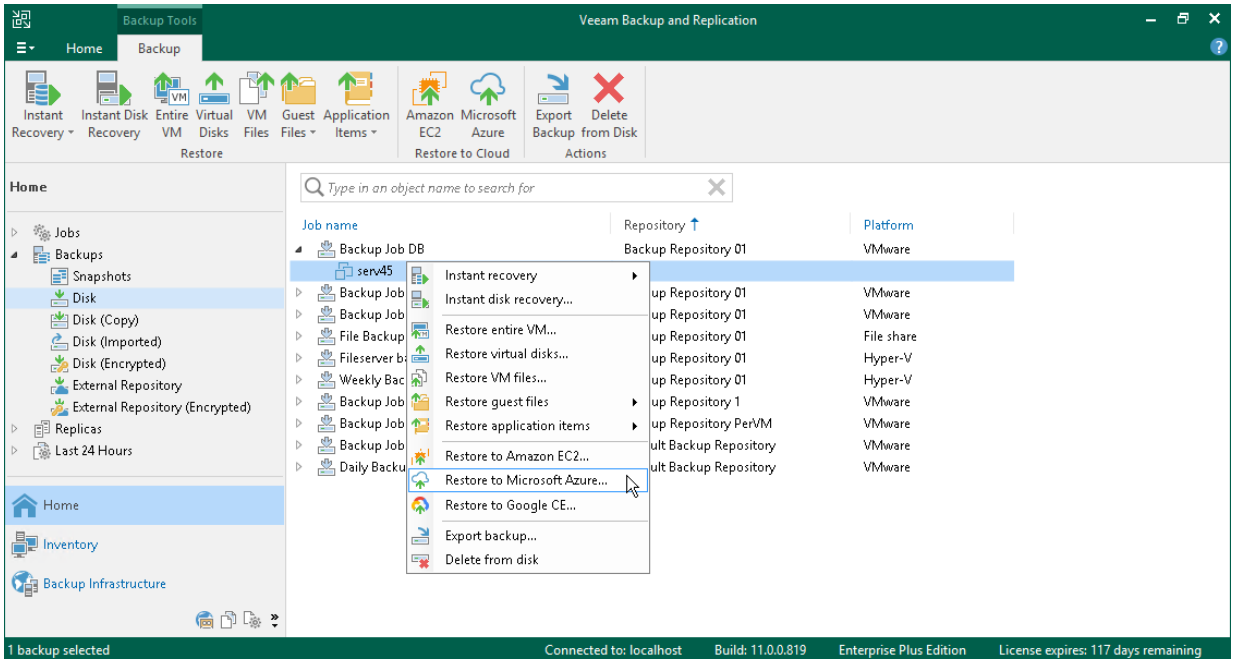
  For storage accounts and network configuration, you must use the same deployment model that you plan to use for workload restore.
- Make sure that you configured all the required components and accounts in Veeam Backup & Replication as described in Configuring Components and Accounts for Restore.
- [For speeding up restore from Capacity Tier] It is strongly recommended to use Azure proxy when you restore from backups residing on a Capacity Tier. For more information on Azure proxies, see Managing Microsoft Azure Proxies.
- You must set up correct time on the backup server. Otherwise you may not be able to add a Microsoft Azure compute account or Microsoft Azure Stack Hub compute account to Veeam Backup & Replication, or the restore process may fail.

Step 1. Launch Restore to Azure Wizard

To launch the **Restore to Azure** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select the type of backups from which you want to restore:
  - VMware vSphere or vCloud Director
  - Microsoft Hyper-V
  - Agent
  - AWS EC2 backup
  - GCE backup
  - Nutanix AHV
  - RHV Proxy

  In the displayed window, click **Entire VM restore** > **Restore to public cloud** > **Restore to Microsoft Azure**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select workloads that you want to restore and click **Restore to Microsoft Azure** on the ribbon. Alternatively, you can right-click one of the workloads that you want to restore and select **Restore to Microsoft Azure**.
- Double-click a full backup file (VBK) or backup metadata file (VBM) in a file browser. Veeam Backup & Replication will start its console. In the **Backup Properties** window, select the necessary workload and click **Restore** > **Restore to Microsoft Azure**.



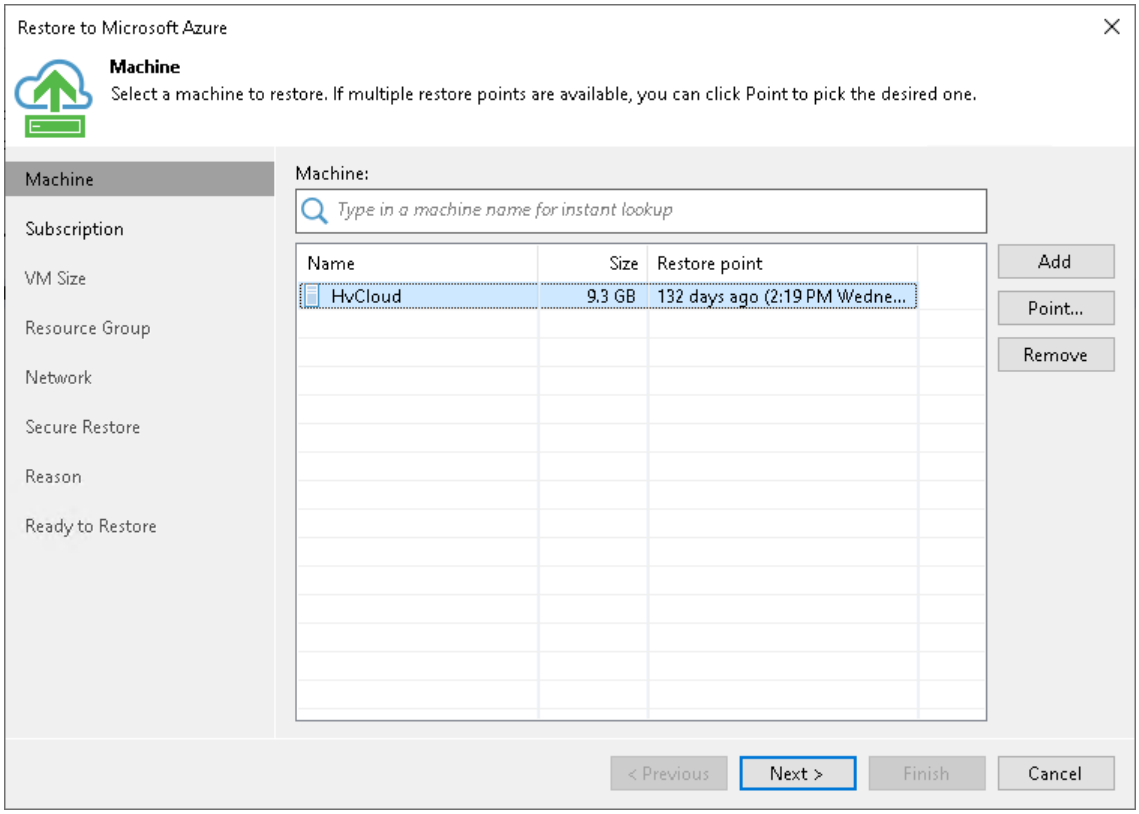
Step 2. Select Workloads and Restore Points

At the **Machine** step of the wizard, specify workloads that you want to restore and specify restore points to which you want to restore the workloads. By default, Veeam Backup & Replication restores workloads to the latest valid restore point in the backup chain.

Selecting Workloads

To select workloads to restore:

1. On the right of the **Machine** list, click **Add**.
2. In the **Backup Browser** window, expand the necessary backup, select workloads and click **Add**.

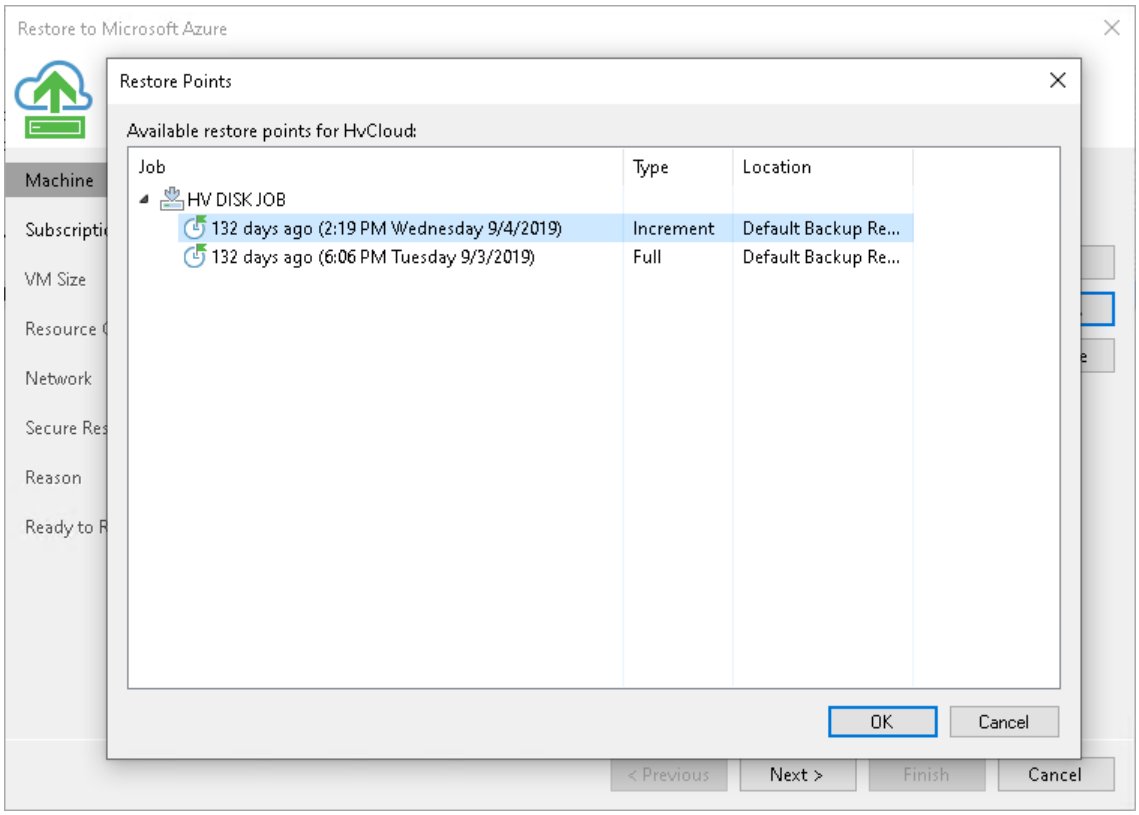


Selecting Restore Points

To select a restore point for a workload, do the following:

1. In the **Machine** list, select a workload.
2. Click **Point** on the right.

3. In the **Restore Points** window, select a restore point to which you want to restore the workload.



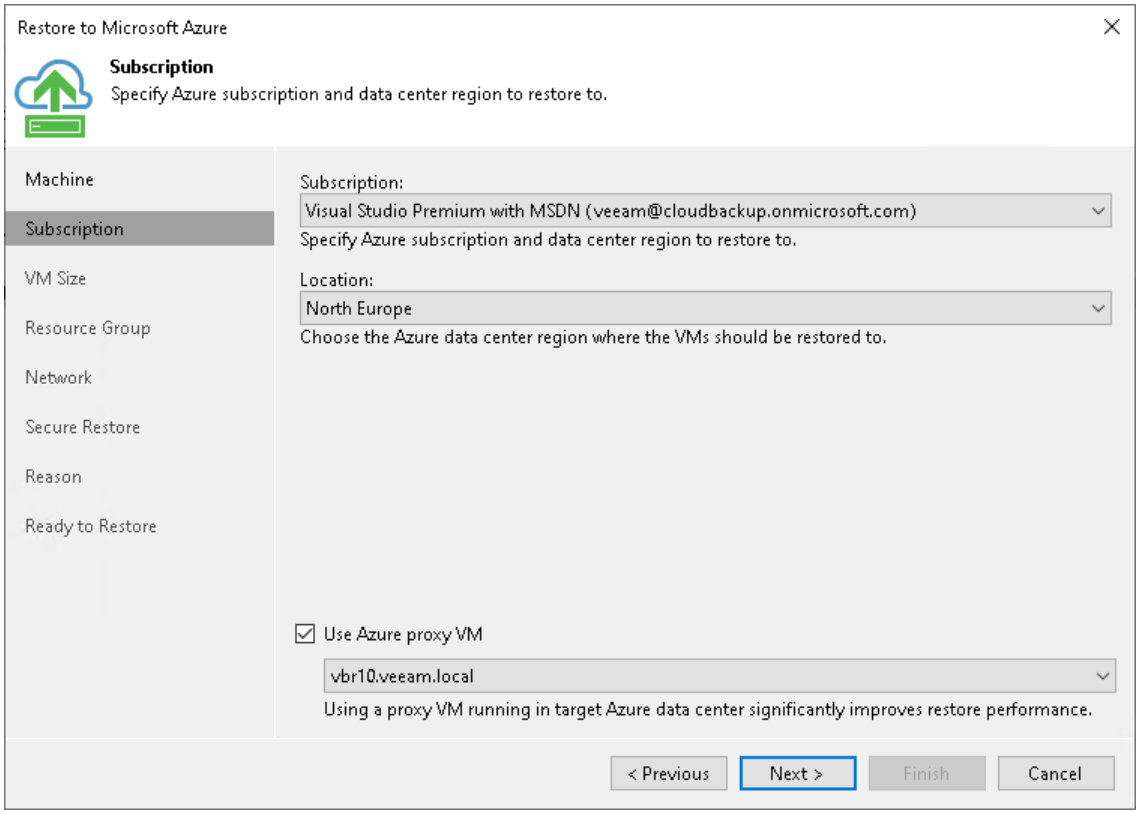
Step 3. Select Subscription and Location

At the **Subscription** step of the wizard, select a subscription, location for the restored workloads and define how workload data must be transported to Microsoft Azure or Azure Stack Hub:

1. From the **Subscription** list, select a subscription whose resources you want to use. The subscription list contains all subscriptions associated with the Azure compute or Azure Stack Hub compute accounts that you have added to Veeam Backup & Replication.
2. From the **Locations** list, select a geographic region to which you want to place the restored workloads. Make sure that you select a geographic region with which at least one storage account of the subscriptions is associated.
3. If you are restoring the workloads to a distant location and want to speed up the restore process, select the **Use Azure proxy VM** check box. From the **Proxy VM** list, select a Microsoft Azure proxy.

   It is recommended that you configure the Azure proxy in the same location to which you plan to restore the workload. For more information, see Managing Azure Proxies.

> **IMPORTANT**
>
> [For restore of Linux workloads] You must have a preconfigured helper appliance in the location to which you restore Linux workloads. If the appliance is not configured, Veeam Backup & Replication will display the **Microsoft Azure Compute Account** wizard so that you can configure the appliance in the selected location.



Step 4. Specify VM Size and Disks

At the **VM Size** step of the wizard, you can select VM sizes, storage accounts where to store disks of the restored workloads, select disks to restore and change their type. By default, Veeam Backup & Replication selects the smallest size that can support the number of disks for the restored workload and restores all workload disks.

Selecting VM Size and Storage Account

To select a size and storage account:

1. In the **Azure VM Configuration** list, select a workload and click **Edit**.
2. From the **Size** list, select a size for the restored workload.

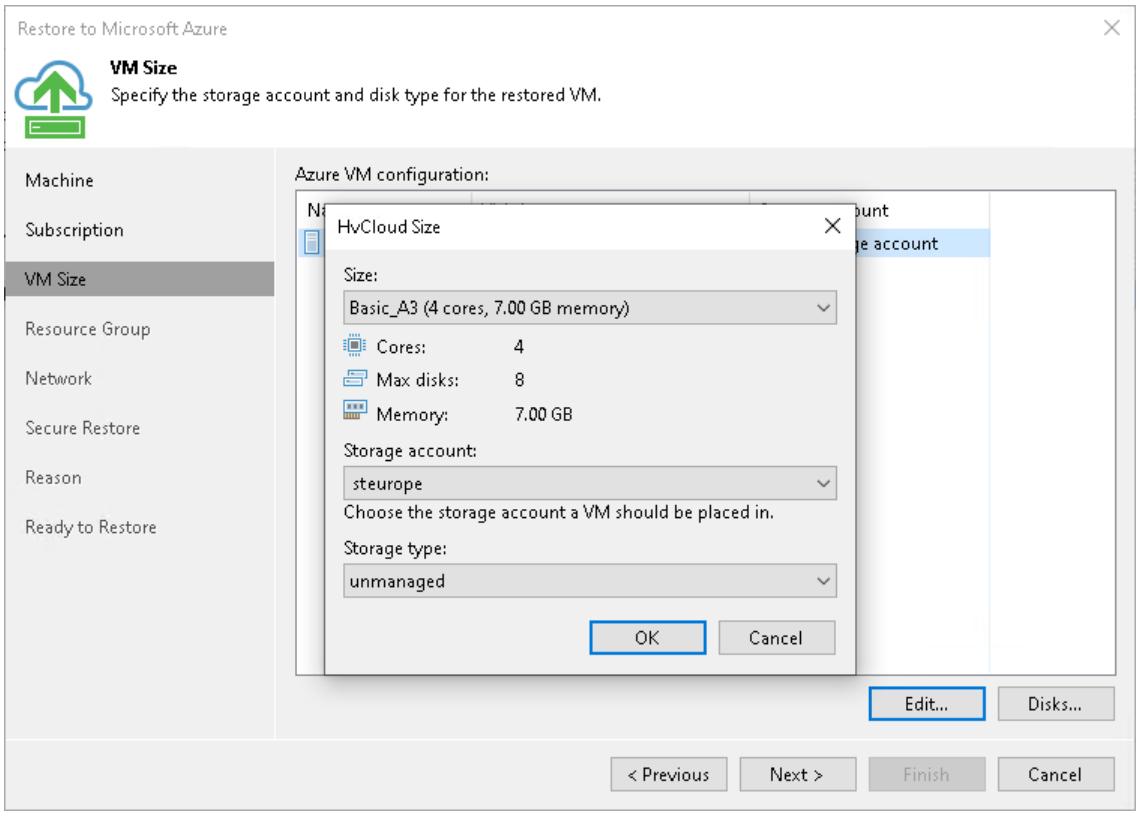
   Make sure that you select the right workload size that corresponds to the initial workload configuration. The size affects the number of CPU cores, memory and disk resources that will be allocated to the restored workload. For more information, see Microsoft Docs.
3. From the **Storage account** list, select a storage account whose resources you want to use to store disks of the restored workload. The storage account must be compatible with the workload size you select. Note the limitations for storage accounts in Limitations for Restore to Microsoft Azure.

   The list of storage accounts contains only general purpose storage accounts. Blob storage accounts are not be displayed in the list of subscriptions. For more information about account types, see Microsoft Docs.

   If you select a premium storage account, make sure that the restored workload size is compatible with the selected account.
4. From the **Storage type** list, select the *managed* or *unmanaged* disk type. Note the limitations for disk sizes in Limitations for Restore to Microsoft Azure.

NOTE

Microsoft Azure subscriptions have default limits on the number of CPU cores. Make sure that the restored workload size that you select does not exceed limits of the subscription.



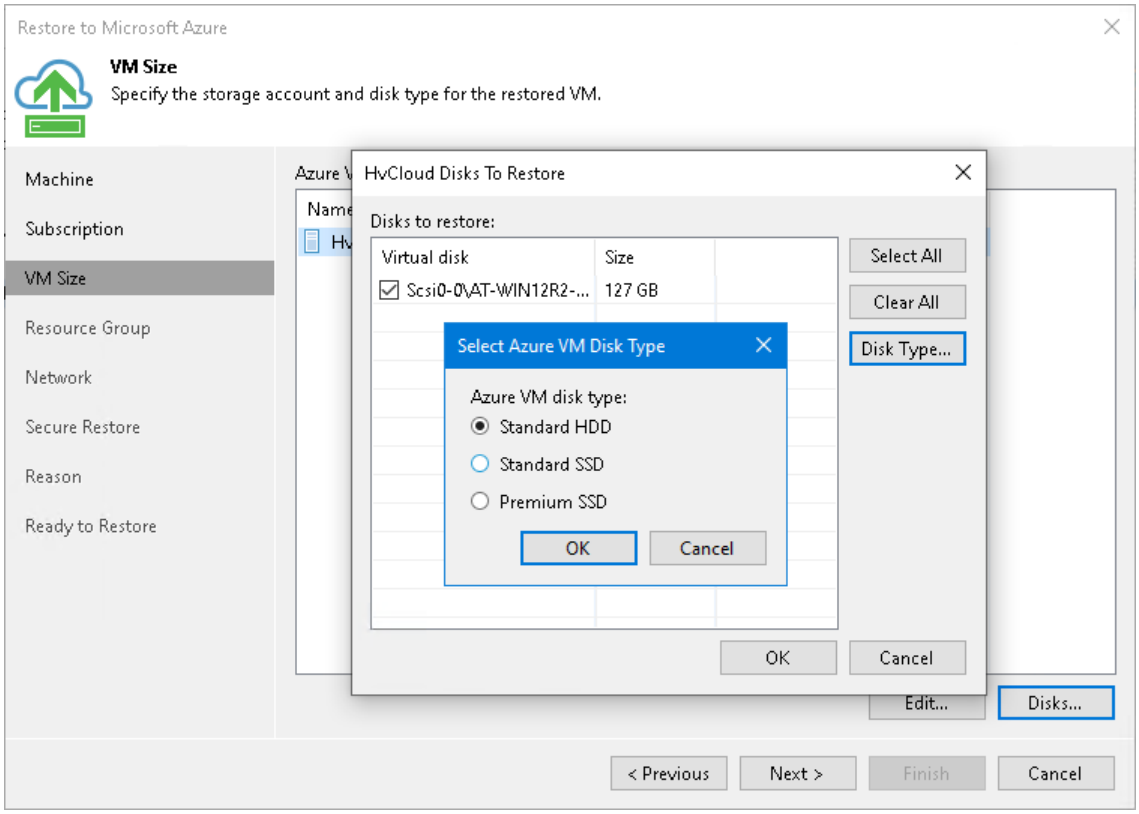
Selecting Disks to Restore and Changing Their Types

You can restore all disks or specific disks of a workload. You can also change disk types of the restored disks if you have selected the managed storage type.

To select specific disks and change their type, do the following:

1. In the **Azure VM Configuration** list, select a workload and click **Disks**.
2. In the **Disks to restore** window, check that check boxes next to disks that you want to restore are selected. Clear check boxes next to disks that you do not want to restore.
3. [For managed storage type] Select a disk and click **Disk Type**. In the **Select Azure VM Disk Type** window, select the necessary type. For more information on disk types, see Microsoft Docs.

> **IMPORTANT**
>
> The selected disk type must be compatible with the selected workload size.



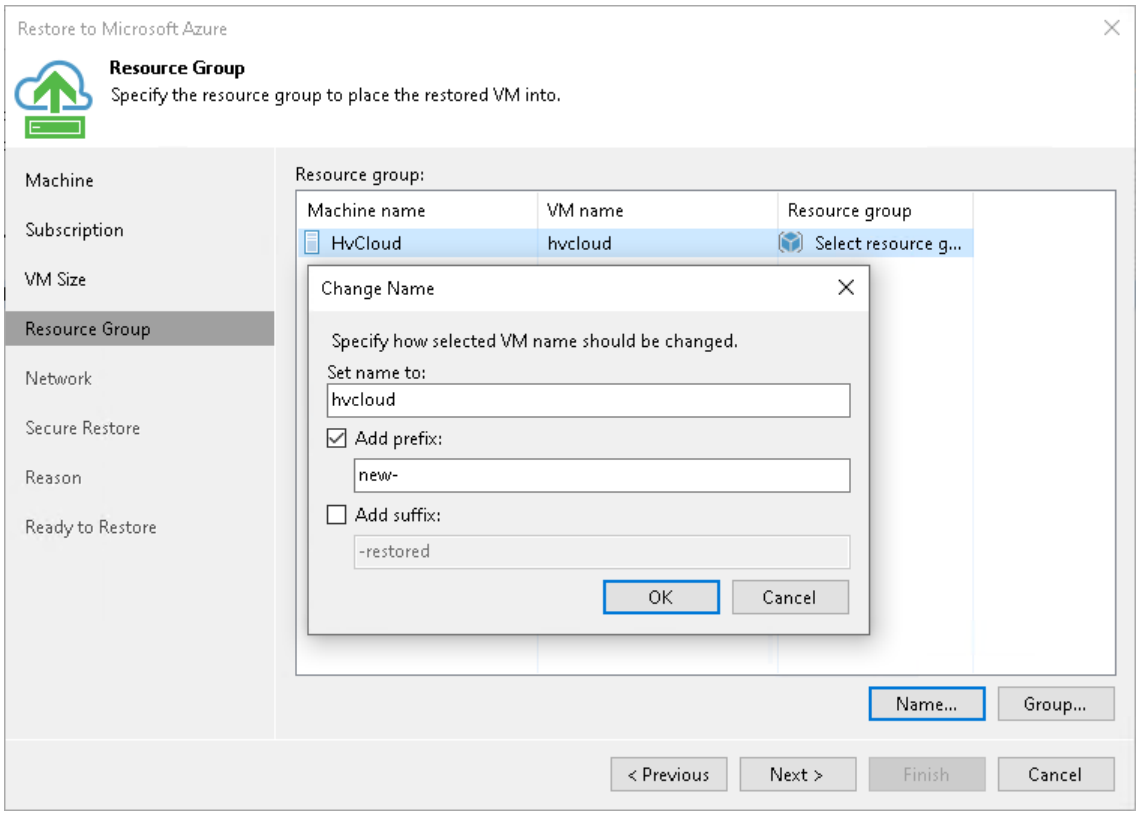
Step 5. Specify VM Name and Resource Group

At the **Resource Group** step of the wizard, you can specify new names for the restored workloads and select resource groups for them. By default, Veeam Backup & Replication restores workloads with their original names and creates a new resource group for the restored workloads and places them to it.

Specifying New Name

To define a new name for a workload:

1. In the **Resource group** list, select a workload and click **Name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule — add a prefix and/or suffix to the original workload name.



Selecting Resource Group

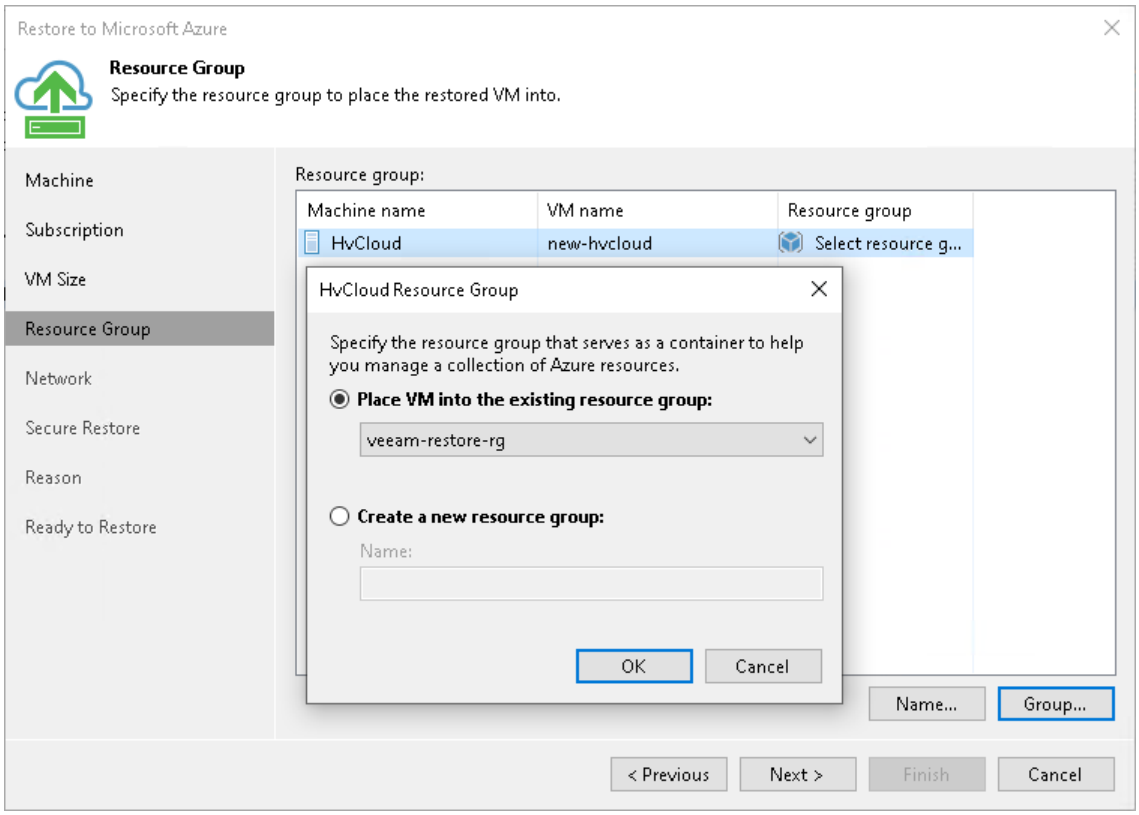
To change a resource group to which a workload will be restored:

1. In the **Resource group** list, select a workload and click **Group**.
2. In the **VM Resource Group** window, select the necessary option for the workload:
   - Select **Place VM into the existing resource group** if you want to place the workload to an existing resource group. Then from the drop-down list, select the necessary resource group.
   - Select **Create a new resource group** if you want to create a dedicated resource group for the restored workload. In the **Name** field, enter a name for the new resource group.

     In the new resource group, Veeam Backup & Replication automatically creates a network security group, a dynamic public IP and network interface.

NOTE

Microsoft Azure subscriptions have default limits on the number of resource groups. If you decide to create a new resource group, make sure that you do not exceed limits of the subscription.



Step 6. Configure Network Settings

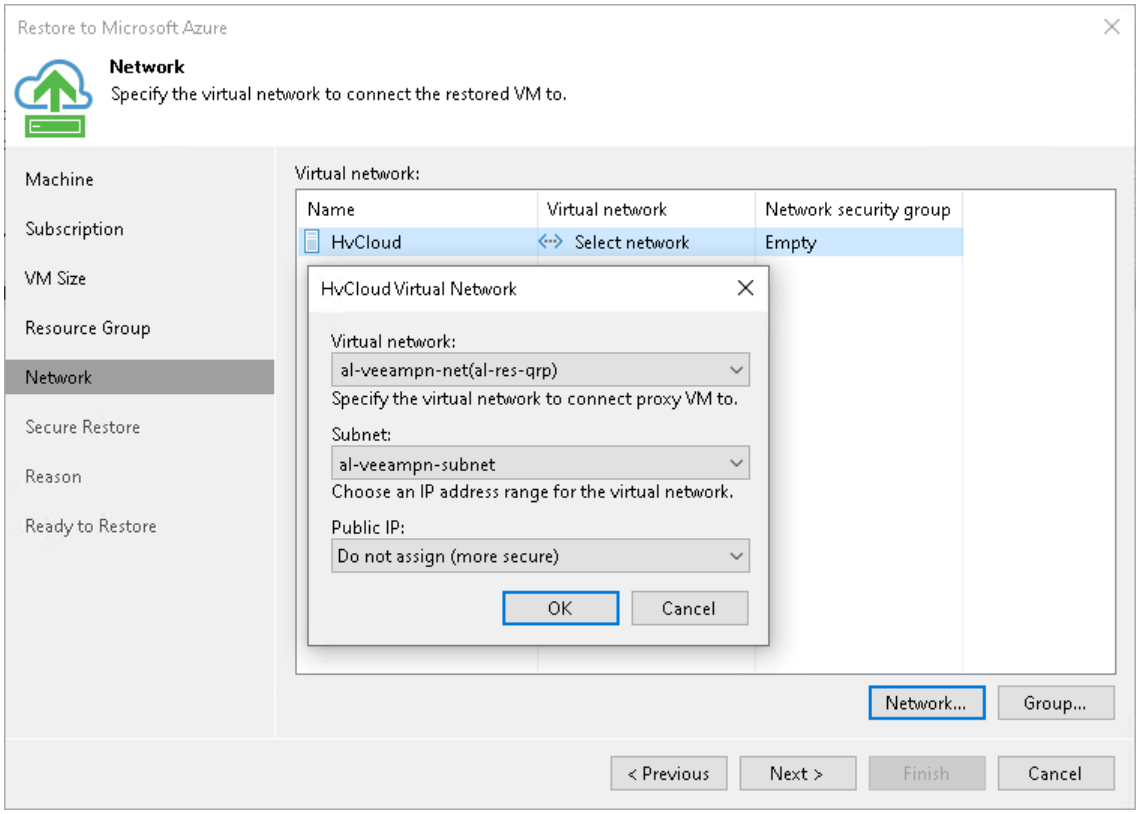
At the **Network** step of the wizard, select to which networks and subnets the restored workloads will be connected. By default, Veeam Backup & Replication creates a new security group for the restored workloads.

Configuring Network

To define network settings for a workload, do the following:

1. In the **Virtual network** list, select a workload and click **Network**.
2. From the **Virtual network** drop-down list, select a network to which the workload must be connected.
3. From the **Subnet** drop-down **list**, select a subnet for the workload.
4. In the **Public IP** field, specify whether to assign a public IP to the workload. You have two options:
   - **Assign (restored VM will be accessible from the Internet).** The public IP will be assigned to the restored workload. For security reasons, make sure traffic filtration rules are properly configured in the security group.
   - **Do not assign (more secure).** The public IP will not be assigned. You can assign a public IP later in the settings of the restored workload.

> **NOTE**
>
> Veeam Backup & Replication can connect a workload only to one virtual network. If necessary, you can manually configure additional network connections in Microsoft Azure after the workload is restored.

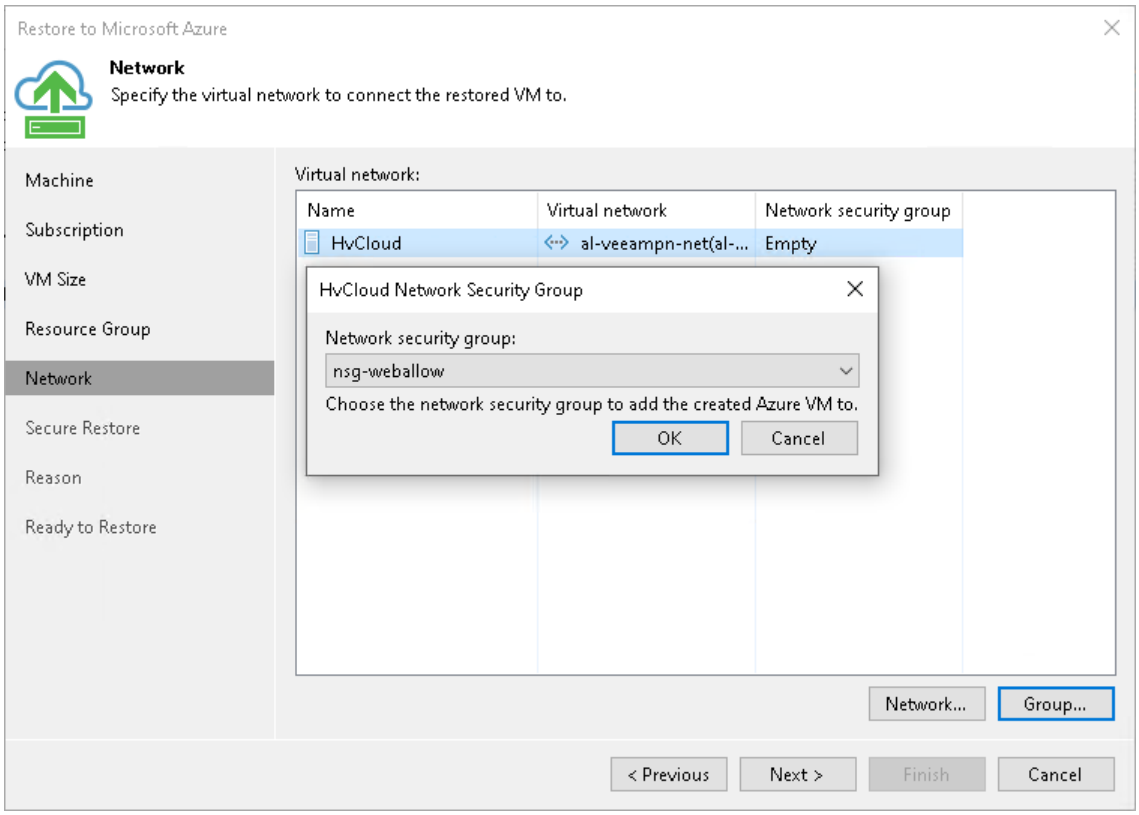


Configuring Security Group

To change a security group to which a workload will be restored:

1. From the **Virtual network** list, select a workload and click **Group**.
2. Select the network security group from the **Network security group** list.

   If you leave the *Empty* value, Veeam Backup & Replication will create a new network security group.



Step 7. Specify Secure Restore Settings

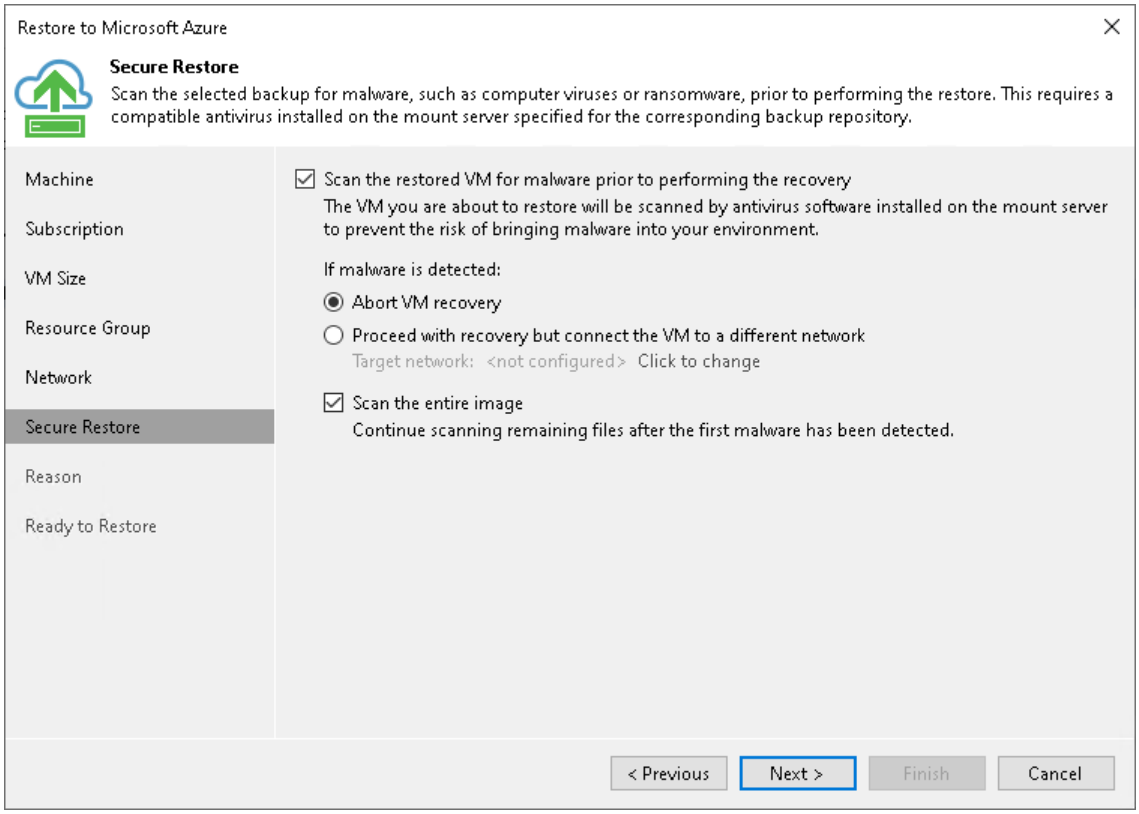
This step is available if you restore Microsoft Windows workloads.

You can instruct Veeam Backup & Replication to perform secure restore — scan workload data with antivirus software before restoring the workload to Microsoft Azure or Azure Stack Hub. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

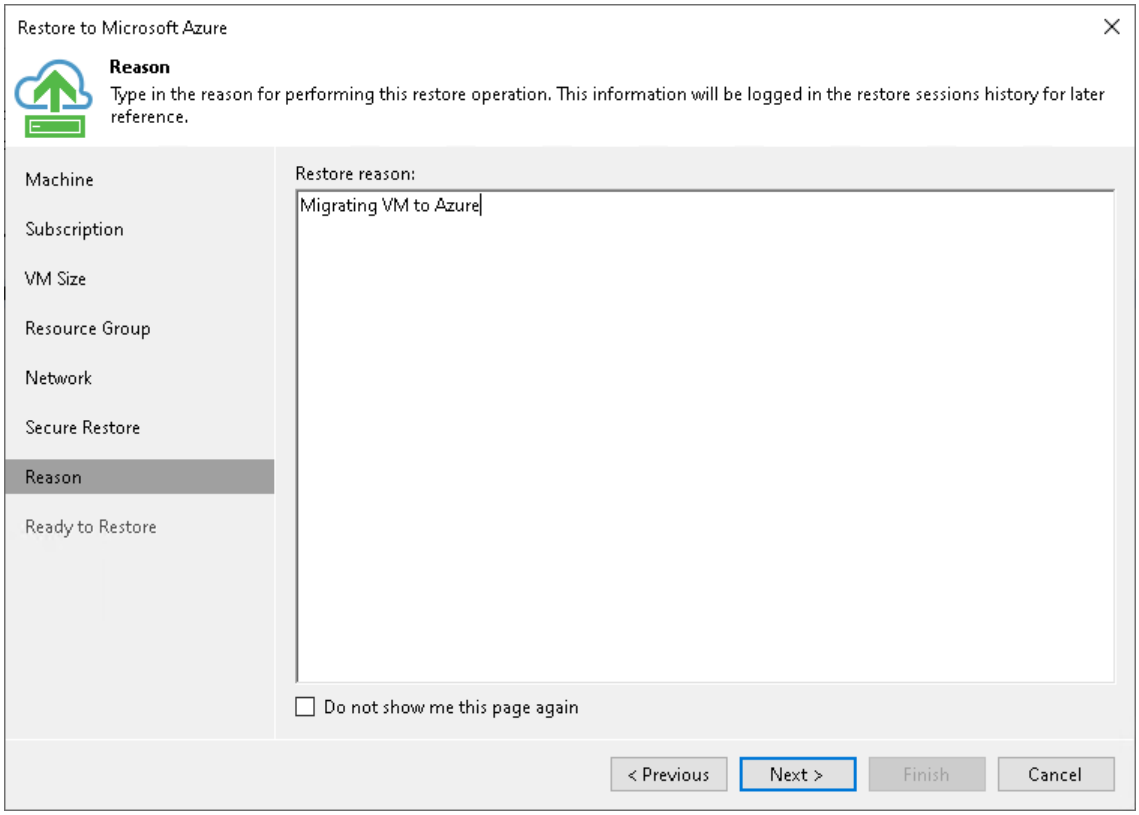
1. At the **Secure Restore** step of the wizard, select the **Scan VM for virus threats prior performing recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software finds a virus threat:
   - **Abort VM recovery.** Select this action if you want Veeam Backup & Replication to cancel the restore session.
   - **Proceed recovery but connect VM to a different network**. Select this action if you want to restore the workload to a different Microsoft Azure virtual network.

     Click the **Click to change** link to select the virtual network.
3. Select the **Scan entire VM for virus threats** check box if you want the antivirus software to continue workload scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



Step 8. Specify Restore Reason

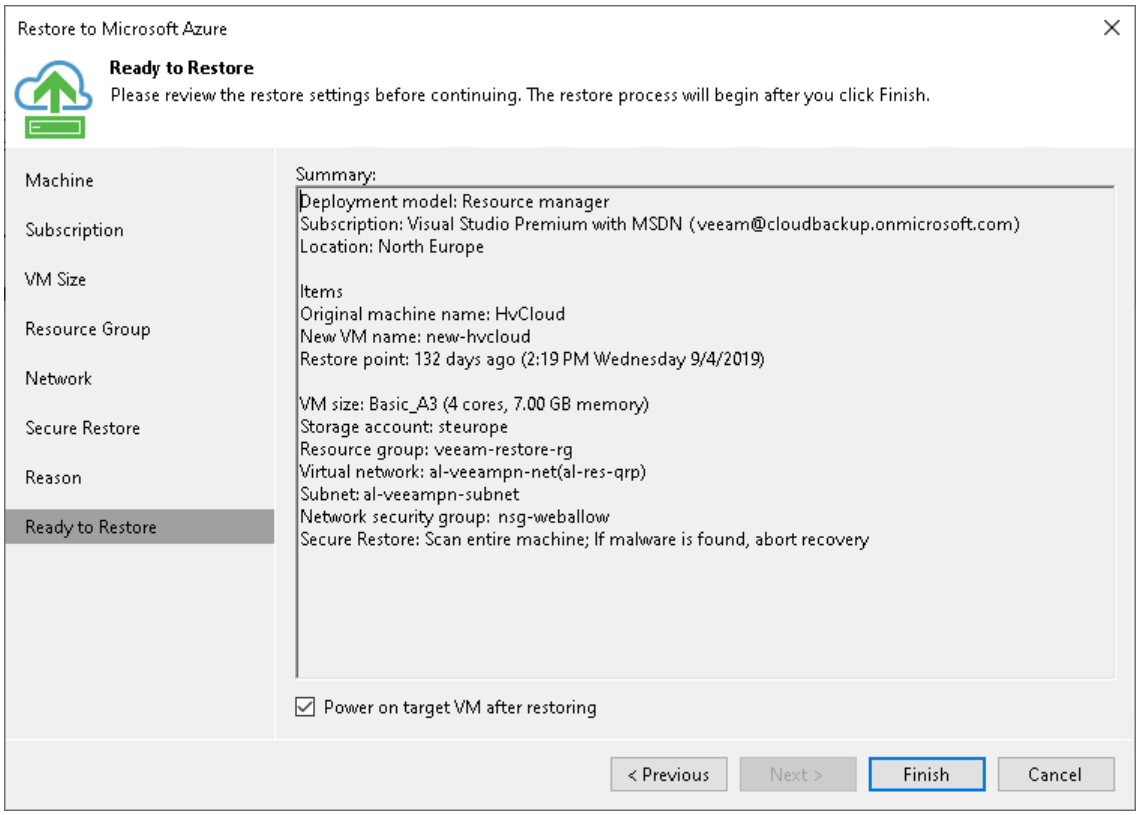
At the **Reason** step of the wizard, enter a reason for restoring the workload. The information you provide will be saved in the session history in Veeam Backup & Replication, and you can view it later.



Step 9. Verify Restore Settings

At the **Ready to Restore** step of the wizard, check the specified settings and click **Finish**. If you want to start the Azure VM right after restore, select the **Power on VM after restoring** check box.

You can trace the restore process in the **Restore Session** window. If you need to cancel the workload restore, click the **Cancel** restore task link.



Restore to Amazon EC2

Veeam Backup & Replication allows you to restore different workloads (VMs, Google VM instances, physical servers and so on) to Amazon Elastic Compute Cloud (Amazon EC2) as EC2 instances. An EC2 instance is a virtual machine in Amazon EC2 with a preconfigured combination of computing resources.

You can use Veeam Backup & Replication to perform the following operations:

- Restore workloads to Amazon EC2 from backups.
- Migrate workloads from the on-premises infrastructure to the cloud.
- Create a test environment in the cloud for troubleshooting, testing patches and updates, and so on.

Supported Backup Types

You can restore workloads from the following types of backups:

- Backups of VMware vSphere or vCloud Director virtual machines created by Veeam Backup & Replication.
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication.
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.

  Backups must be created at the entire machine level or volume level.
- Backups of Amazon EC2 instances created by Veeam Backup for AWS.
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure.
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud Platform.
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV.
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of RHV VMs created by Veeam Backup for RHV.

Proxy Appliances

Proxy appliance is an auxiliary Linux-based EC2 instance. It is used to upload backed-up data to Amazon EC2. Veeam Backup & Replication automatically deploys the proxy appliance in Amazon EC2 only for the duration of the restore process and removes it immediately after that.

Depending on the type of backups you are restoring from and their location, the proxy appliance may be required or optional. The proxy appliance is required when you restore from:

- Backups of EC2 instances that are stored in external repositories.
- Backups of virtual and physical machines that are stored in object storage repositories.

The proxy appliance is optional when you restore from backups of virtual and physical machines stored in backup repositories, or backups of EC2 instances copied to backup repositories with backup copy jobs. It is recommended, however, to use the proxy appliance in scenarios where it is optional, as the proxy appliance can significantly improve restore performance. You can specify the proxy appliance settings at the Proxy Appliance step of the **Restore to Amazon EC2** wizard.

Requirements for Proxy Appliance

When configuring a proxy appliance, mind the following:

- If you want to restore from backups in an on-premises object storage repository, the proxy appliance machine must have access to the source object storage repository. To provide access to object storage repositories, you can use VPN or AWS Direct Connect.
- To upload one machine disk to Amazon EC2, the proxy appliance requires 1 GB RAM. Make sure that the type of EC2 instance selected for the proxy appliance offers enough memory resources to upload all machine disks. Otherwise, the restore process may fail.
- A subnet and security group that you select for the proxy appliance must meet the following requirements:
  - Auto-assignment of public IPv4 addresses must be enabled in the subnet. For more information on how to enable this option, see the AWS Documentation.
  - The subnet route table must contain a default route to an active AWS internet gateway. For more information on internet gateways and how to create route tables, see the AWS Documentation.
  - The subnet must have no network access control lists (ACLs) or a network ACL that allows inbound and outbound traffic on the ports listed in the Used Ports section.
  - The security group must allow inbound and outbound traffic on the ports listed in the Used Ports section.

How Restore to Amazon EC2 Works

The workflow of the restore process depends on whether the proxy appliance is used or not. For more information on the proxy appliance, see Proxy Appliance.

> **NOTE**
>
> Mind the following:
>
> - If you added the Veeam Backup for AWS appliance to the Veeam Backup & Replication infrastructure and you plan to restore Amazon EC2 instances from restore points that were created using the appliance, you do not need to configure the proxy appliance. Also, restore to Amazon EC2 works as described in the Entire EC2 Instance Restore section in the Veeam Backup for AWS User Guide.
> - Amazon EC2 does not support GUID partition table (GPT) disks. That is why during restore, Veeam Backup & Replication converts GPT disks into Master Boot Record (MBR) disks.

Restoring to Amazon EC2 with Proxy Appliance

If a proxy appliance is required or you selected to use it for restore to Amazon EC2, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a proxy appliance in Amazon EC2.

   During the restore process, the proxy appliance communicates with backup infrastructure components over the SSH protocol and the network redirector that is deployed on the proxy appliance.

2. For every disk of a backed-up workload, Veeam Backup & Replication creates an empty EBS volume in Amazon EC2.

3. Veeam Backup & Replication hot-adds empty disks to the proxy appliance and restores backed-up data to the EBS volumes.
4. Veeam Backup & Replication creates a target instance in Amazon EC2.
5. Veeam Backup & Replication detaches the EBS volumes from the proxy appliance and attaches them to the target instance.
6. After the restore process is complete, Veeam Backup & Replication removes the proxy appliance from Amazon EC2.

Restoring to Amazon EC2 without Proxy Appliance

If you selected not to use a proxy appliance for restore to Amazon EC2, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication uploads disks of a backed-up workload to Amazon S3.

   In Amazon S3, the uploaded disks are stored to the temporary bucket in the RAW format.
2. Veeam Backup & Replication imports the backed-up data from the temporary bucket in Amazon S3 to EBS volumes in Amazon EC2.
3. Veeam Backup & Replication creates a target instance in Amazon EC2 and attaches the EBS volumes to it.
4. After the import process is complete, Veeam Backup & Replication removes the temporary bucket from Amazon S3.

AWS IAM User Permissions

To restore to Amazon EC2, it is recommended that the IAM user whose credentials you plan to use to connect to AWS has administrative permissions — access to all AWS actions and resources.

If you do not want to provide full access to AWS, you can grant to the IAM user a minimal set of permissions that will be sufficient for restore. To do that, create the following policy in the JSON format and attach it to the IAM user:

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Action": [
   "ec2:DescribeInstances",
   "ec2:RunInstances",
   "ec2:TerminateInstances",
   "ec2:StartInstances",
   "ec2:StopInstances",
   "ec2:ModifyInstanceAttribute",
   "ec2:DescribeImages",
   "ec2:ImportImage",
   "ec2:DeregisterImage",
   "ec2:DescribeVolumes",
   "ec2:CreateVolume",
   "ec2:ModifyVolume",
   "ec2:ImportVolume",
   "ec2:DeleteVolume",
   "ec2:AttachVolume",
   "ec2:DetachVolume",
   "ec2:CreateSnapshot",
```

```
      "ec2:DescribeSnapshots",
      "ec2:DeleteSnapshot",
      "ec2:DescribeSubnets",
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeSecurityGroups",
      "ec2:DescribeKeyPairs",
      "ec2:CreateKeyPair",
      "ec2:DeleteKeyPair",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeVpcs",
      "ec2:DescribeConversionTasks",
      "ec2:DescribeImportImageTasks",
      "ec2:DescribeVolumesModifications",
      "ec2:CancelImportTask",
      "ec2:CancelConversionTask",
      "ec2:CreateTags",
      "ec2:DescribeAccountAttributes",
      "ec2:DescribeDhcpOptions",
      "ec2:DescribeVpcAttribute",
      "iam:GetRole",
      "iam:CreateRole",
      "iam:PutRolePolicy",
      "iam:DeleteRolePolicy",
      "s3:CreateBucket",
      "s3:ListBucket",
      "s3:ListAllMyBuckets",
      "s3:DeleteBucket",
      "s3:PutObject",
      "s3:DeleteObject",
      "s3:GetBucketLocation",
      "s3:PutLifeCycleConfiguration",
      "s3:GetObject",
      "s3:RestoreObject",
      "s3:AbortMultiPartUpload",
      "s3:ListBucketMultiPartUploads",
      "s3:ListMultipartUploadParts"
    ],
    "Effect": "Allow",
    "Resource": "*"
   }]
  }
```

Alternatively, you can attach the created policy to the IAM group or role to which the IAM user is assigned.

For information on how to create and attach a policy to an IAM user, see the Creating IAM Policies and Adding and Removing IAM Identity Permissions sections in the AWS IAM User Guide.

Restoring to Amazon EC2

To restore workloads to Amazon EC2, use the **Restore to Amazon EC2** wizard.

> **NOTE**
>
> If you added the Veeam Backup for AWS appliance to the Veeam Backup & Replication infrastructure and you plan to restore Amazon EC2 instances from restore points that were created using the appliance, the steps of the restore wizard differ from the steps described in this guide. For more information, see the Restoring to Amazon section in the Integration with Veeam Backup for AWS Guide.

Before You Begin

Before you restore workloads to Amazon EC2, mind the following requirements and limitations:

- Check whether a proxy appliance must be configured for restore. For more information, see Proxy Appliances.
- The backup server and repositories with workload backup files must have access to the Internet.

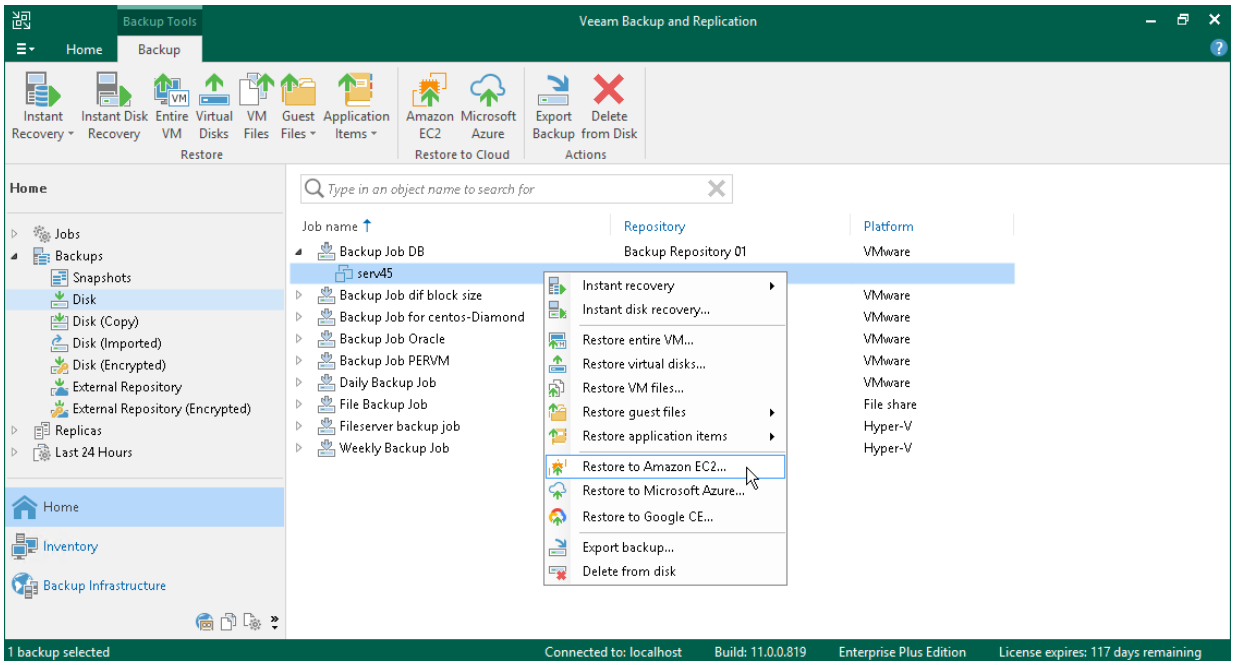
  If backup files are located on deduplicating storage appliances or shared folder repositories, the Internet connection is required for gateway servers that communicate with these repositories.
- You must have a backup of the workload that you plan to restore to Amazon EC2.
- Make sure that a user whose credentials you plan to use to connect to AWS has permissions to restore to Amazon EC2. For more information, see AWS IAM User Permissions.
- If you restore workloads other than EC2 instances, check the supported OS, EC2 instance and file system types in the AWS Documentation.
- If you plan to assign AWS tags to the restored EC2 instance, check limitations for tags in the AWS Documentation.

Step 1. Launch Restore to Amazon EC2 Wizard

To begin the restore process, do one of the following.

- On the **Home** tab, click **Restore** and select the type of backups from which you want to restore:
  - VMware vSphere
  - VMware vCloud Director
  - Microsoft Hyper-V
  - Agent
  - AWS EC2 backup
  - Microsoft Azure backup
  - GCE backup
  - Nutanix AHV
  - RHV Proxy

  In the displayed window, click **Entire VM restore** > **Restore to public cloud** > **Restore to Amazon EC2**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select workloads that you want to restore and click **Restore to Amazon EC2** on the ribbon. Alternatively, you can right-click one of the workload that you want to restore and select **Restore to Amazon EC2**.
- Double-click a full backup file (VBK) or backup metadata file (VBM) in a file browser. Veeam Backup & Replication will start its console. In the **Backup Properties** window, select the necessary workload and click **Restore** > **Restore to Amazon EC2**.



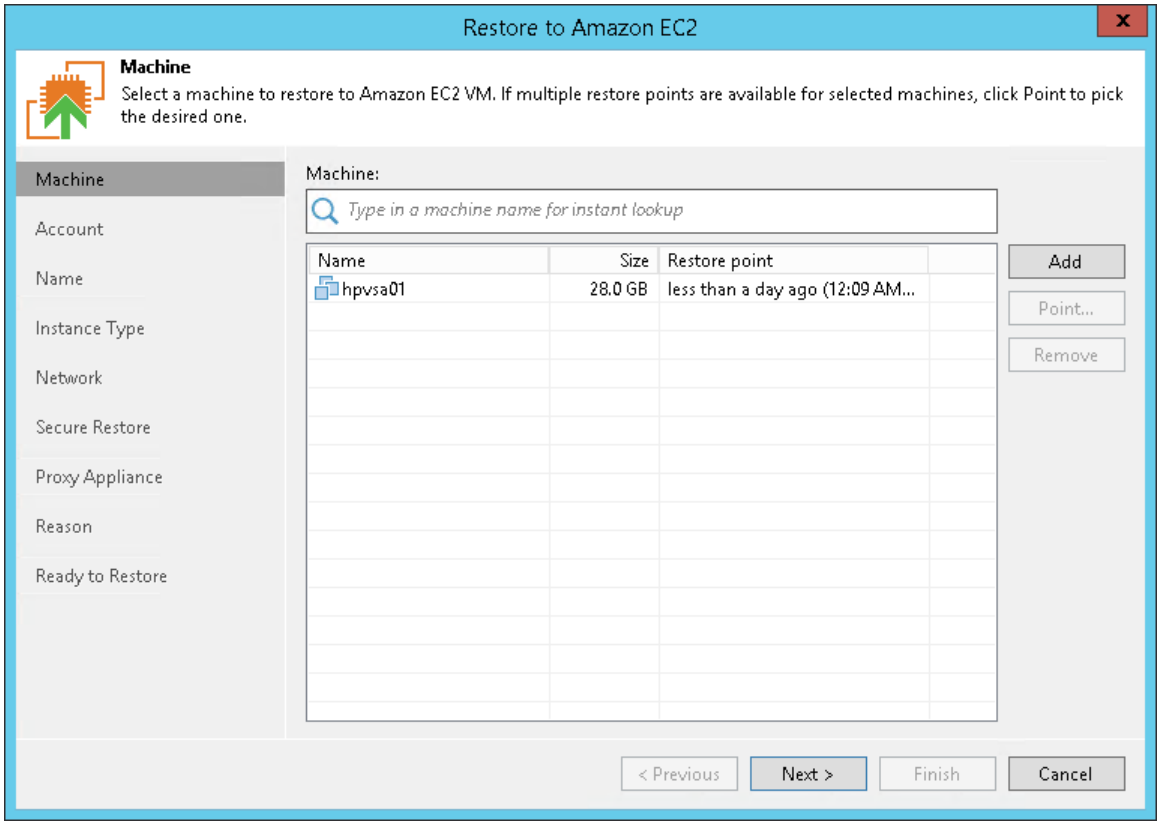
Step 2. Select Workloads and Restore Points

At the **Machine** step of the wizard, specify workloads that you plan to restore and restore points to which you want to restore the workloads. By default, Veeam Backup & Replication restores workloads to the latest valid restore point in the backup chain.

Selecting Workloads

To select workloads to restore:

1. On the right of the **Machine** list, click **Add**.
2. In the **Backup Browser** window, expand the necessary backup, select the workloads and click **Add**.

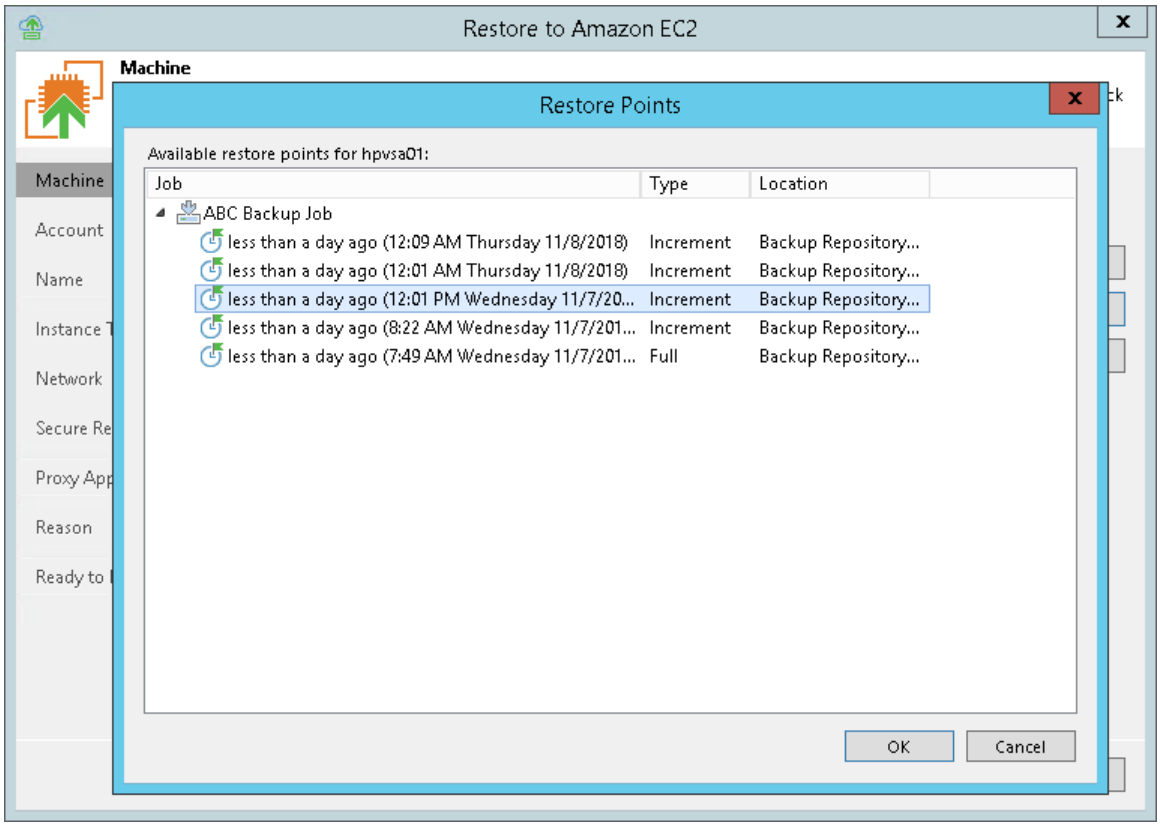


Selecting Restore Point

To select a restore point for a workload:

1. In the **Machine** list, select a workload.
2. Click **Point** on the right.

3. In the **Restore Points** window, select a restore point to which you want to restore the workload.



Step 3. Specify Credentials and Region Settings

At the **Account** step of the wizard, specify AWS user credentials and region:

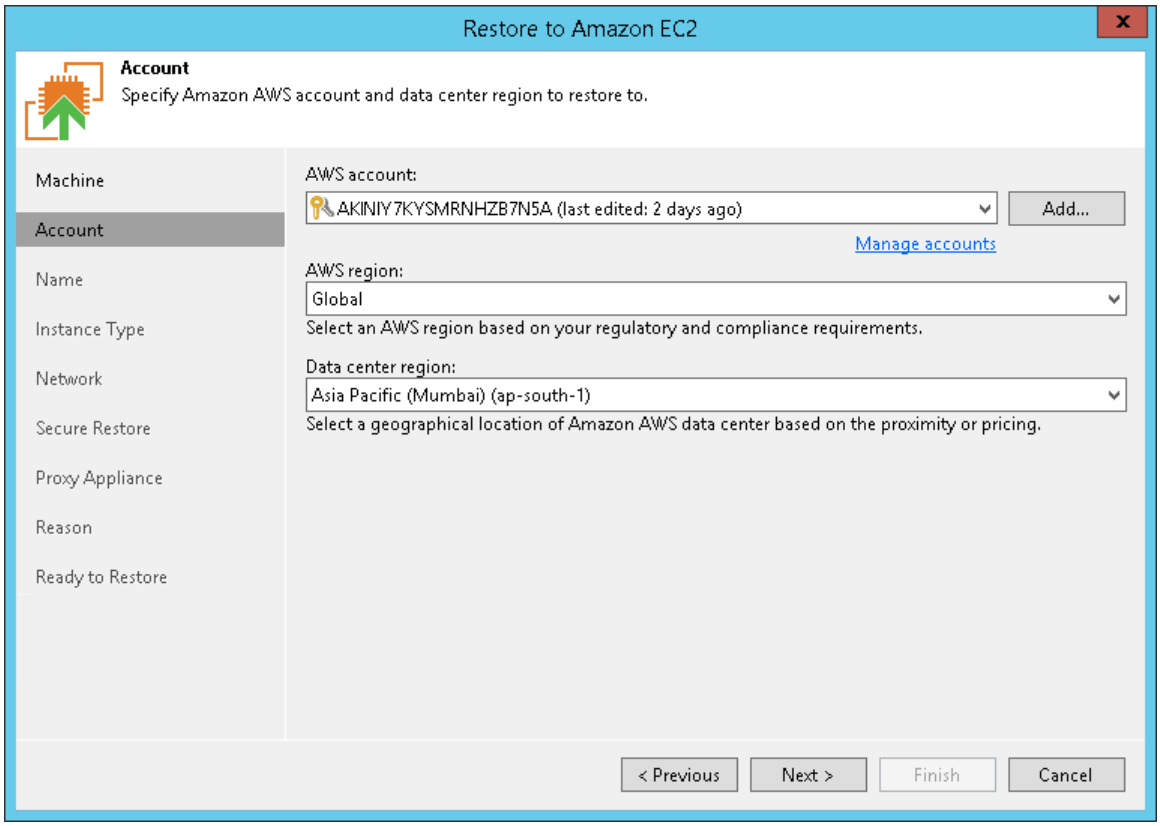
1. From the **AWS account** list, select credentials of a user account that will be used to connect to AWS. This user account must have permissions listed in AWS IAM User Permissions.

   When you add AWS user credentials, Veeam Backup & Replication imports information about resources associated with this user. During the restore process, Veeam Backup & Replication accesses these resources and uses them to create a target instance in Amazon EC2.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage accounts** link or click **Add** on the right to add the necessary credentials.

2. From the **AWS region** list, select the AWS region in which Veeam Backup & Replication will restore workloads as Amazon EC2 instances.

3. From the **Data center region** list, select the geographic region where Veeam Backup & Replication will restore the workloads.



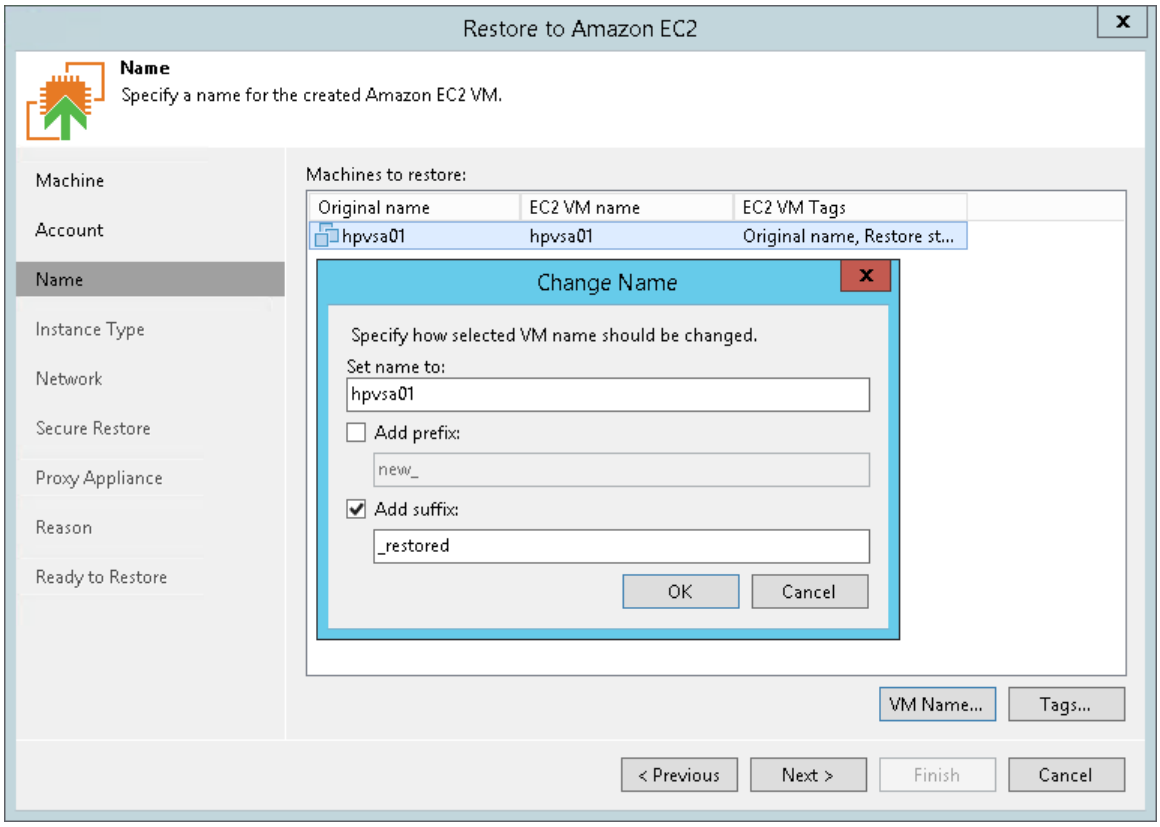
Step 4. Specify Name and Tags

At the **Name** step of the wizard, you can specify names and manage AWS tags for the restored workloads. By default, Veeam Backup & Replication uses the original workload names and adds the **Original name** and **Restore start time** tags.

Specifying New Name

To define a new name for a workload that will be restored:

1. In the **Machines to restore** list, select a workload and click **VM name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule — add a prefix and/or suffix to the original workload name.



Managing AWS Tags

You can use AWS tags to categorize instances in Amazon EC2. A tag is a label with metadata that includes two properties: a key and a value. For more information on AWS tags, see the AWS Documentation.

You can modify or delete these tags, or add new ones.

> **NOTE**
>
> If you restore a workload from backups of an Amazon EC2 instance, Veeam Backup & Replication displays tags that were assigned to this instance. You can modify or delete these tags as well.

Adding Tag

To add a new tag:

1. In the **Machines to restore** list, select a workload and click **Tags**.
2. In the **Tags** window, click **Add**.
3. In the **EC2 VM Tag** window, specify the **Key** and **Value** properties.

   Note that you cannot add the tag with the *Name* key. Veeam Backup & Replication uses the *Name* tag to set the name for the restored EC2 instance in Amazon EC2.

Modifying Tag

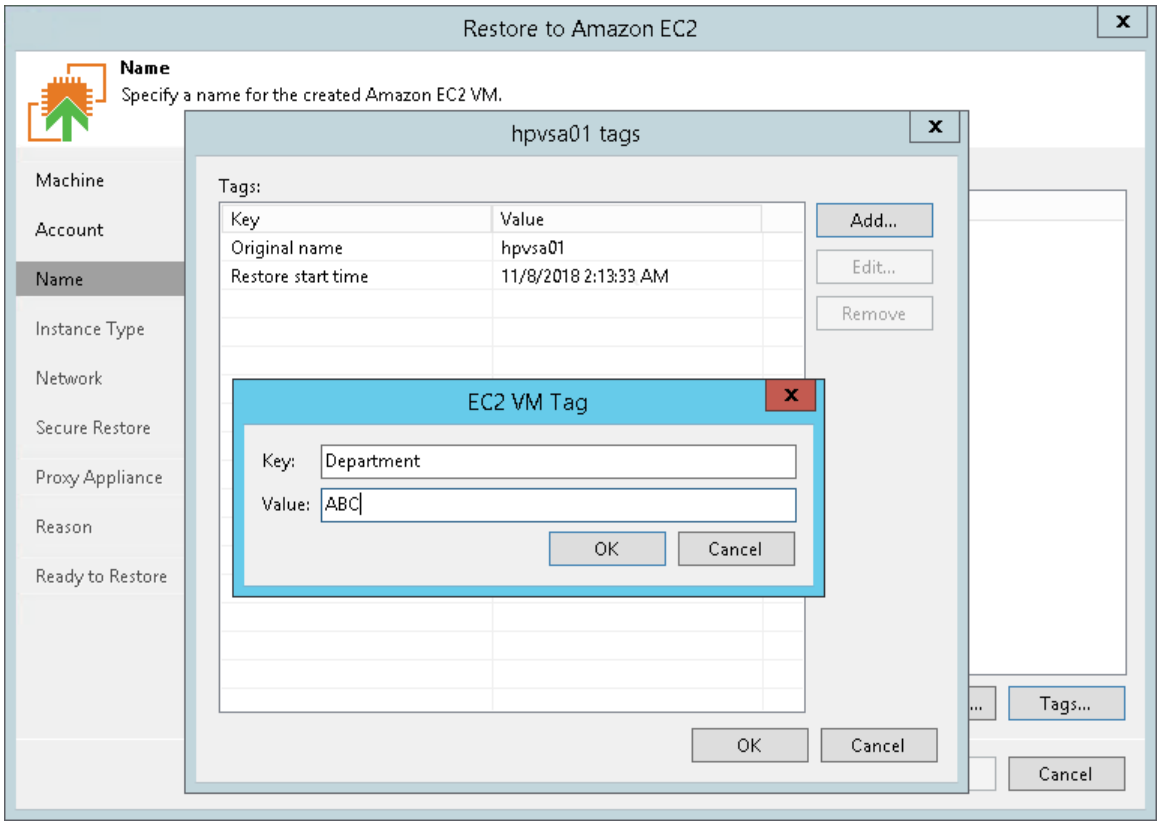
To modify a tag:

1. In the **Machines to restore** list, select a workload and click **Tags**.
2. In the **Tags** window, select the necessary tag and click **Edit**.
3. In the **EC2 VM Tag** window, edit the **Key** or **Value** properties.

Deleting Tag

To delete a tag:

1. In the **Machines to restore** list, select a workload and click **Tags**.
2. In the **Tags** window, select the necessary tag and click **Remove**.



Step 5. Specify Instance Type and Disks

At the **Instance Type** step of the wizard, you can configure the instance type for the restored workload, select which disks to restore and change their type. By default, Veeam Backup & Replication restores all disks as Amazon Elastic Block Store (EBS) volumes of the General Purpose SSD type. For information on types of EBS volumes, see the AWS Documentation.

Selecting Instance Type

You can select the amount of computing resources that AWS will provision for your restored workload — an Amazon EC2 instance type. Each instance type offers a unique combination of CPU, memory, storage, and networking resources.

To select an instance type for a workload that will be restored:

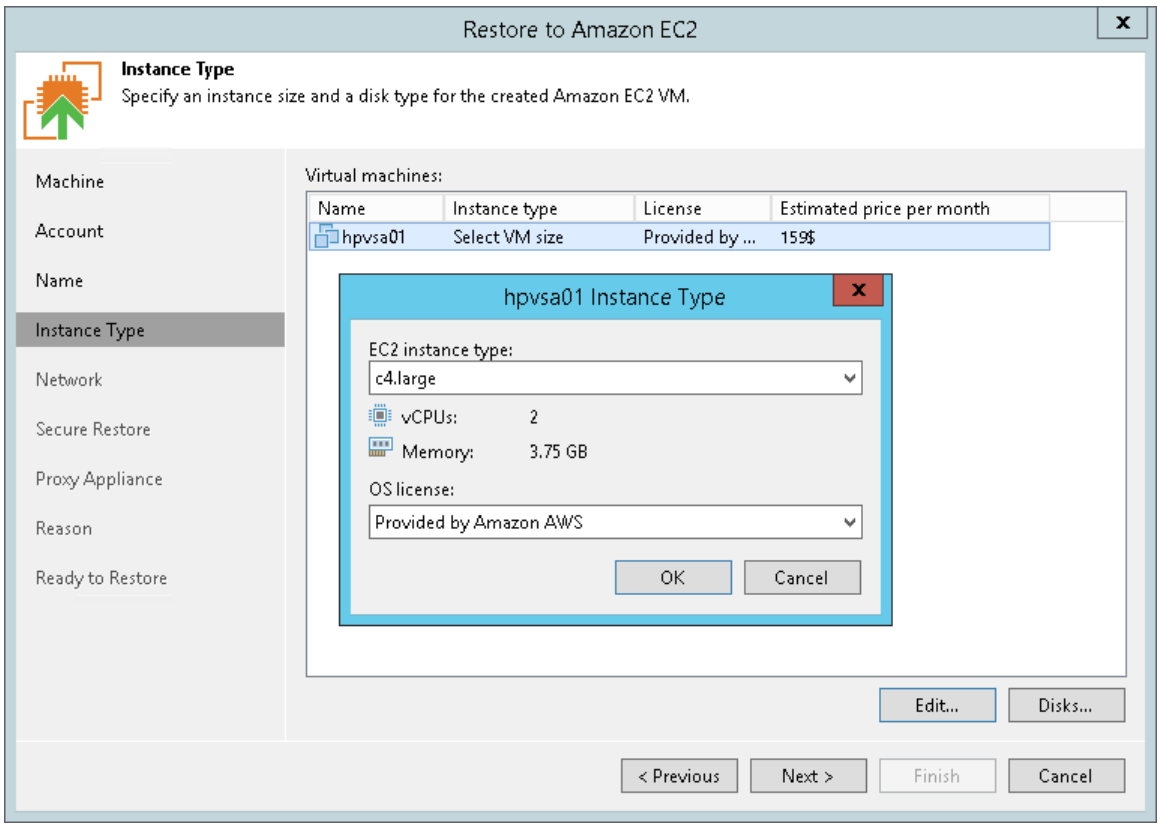
1. In the **Virtual machines** list, select a workload and click **Edit**.
2. From the **EC2 instance type** list, select the instance type for the restored workload.

   Make sure that you select the right instance type that corresponds to the initial workload configuration. For information on instance types, see Amazon EC2 Instance Types.

   If you restore an EC2 instance from backups created with Veeam Backup for AWS, Veeam Backup & Replication will identify the type of a backed-up instance and select it by default.
3. From the **OS license** list, select the license policy that AWS will apply for software on the restored workload:
   - [For Linux workloads] The Bring Your Own License (BYOL) policy is used.
   - [For Microsoft Windows workloads] Select one of the following license policies:
     - **Provided by Amazon AWS**. Select this option if you want to obtain licenses for Microsoft software from AWS.

- **Bring Your Own License (BYOL)**. Select this option if you want to use your existing licenses for Microsoft software.

For more information on Microsoft software licensing in AWS, see Microsoft Licensing on AWS.



Selecting Workload Disks to Restore and Changing Their Types

You can restore all disks or specific disks of a workload. You can also change disk types of the restored disks.

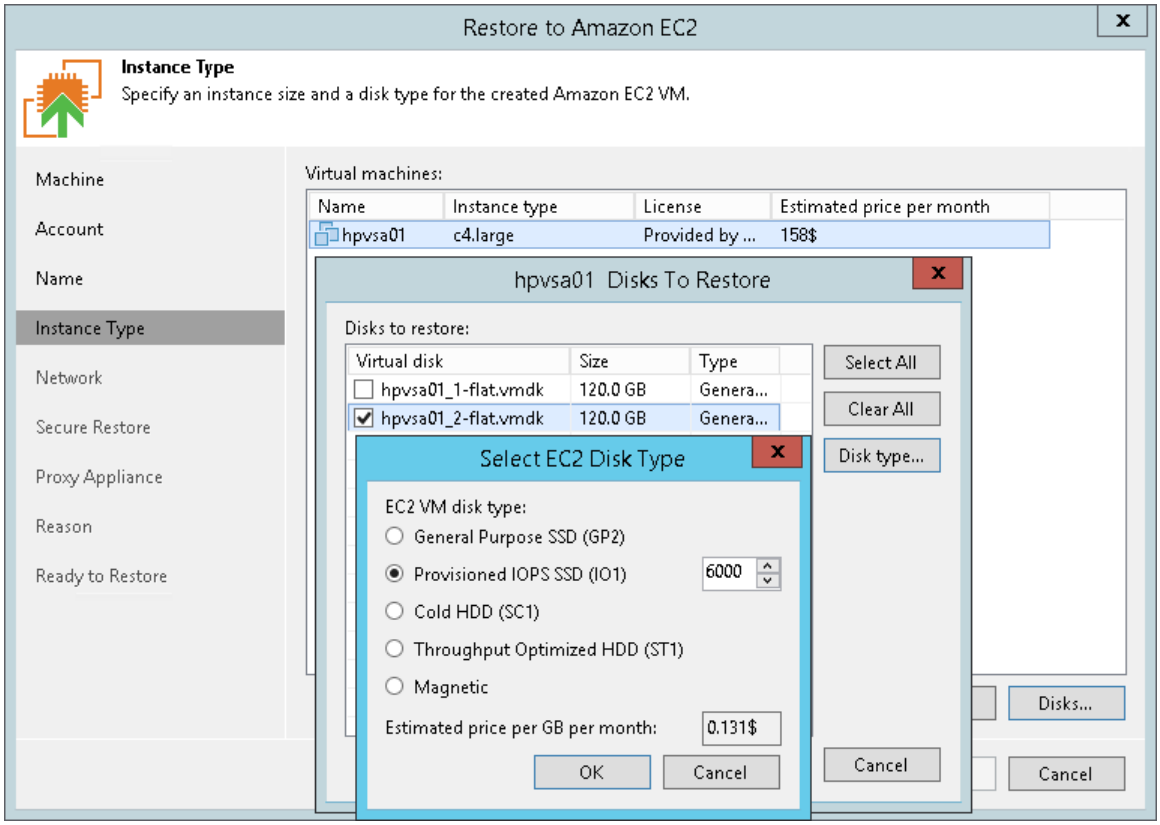
To select specific disks and change their type, do the following:

1. In the **Virtual machines** list, select a workload and click **Disks**.
2. In the **Disks To Restore** window, make sure that check boxes next to disks that you want to restore are selected. Clear check boxes next to disks that you do not want to restore.
3. Select a disk whose type you want to change and click **Disk type**.
4. In the **Select EC2 Disk Type** window, choose the disk type.

   If you selected the **Provisioned IOPS SSD (IO1)** type, you can also specify the maximum number of input/output operations per second (IOPS) for the volume. For more information on IOPS, see the AWS Documentation.

TIP

For your convenience, Veeam Backup & Replication uses the AWS Simple Monthly Calculator tools to estimate an approximate price per month for using a selected instance. The estimated price is calculated based on the instance type, license policy and disk configuration.

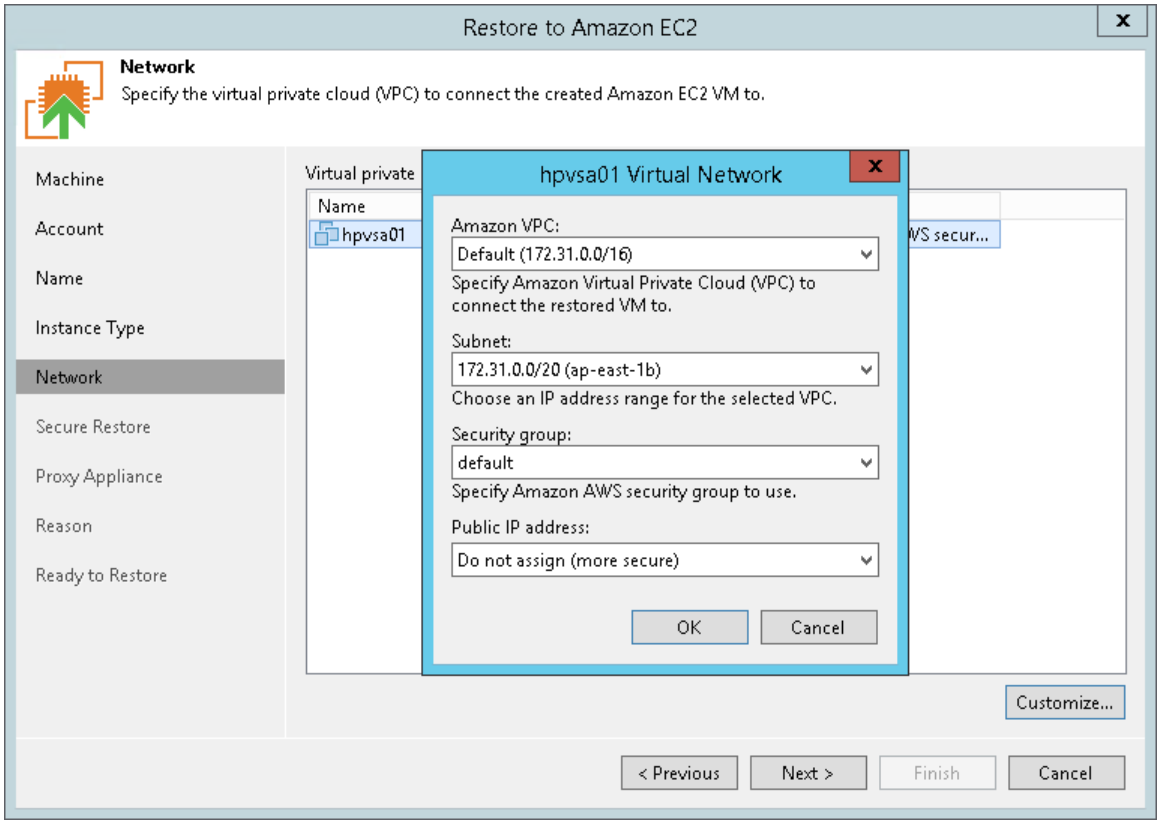


Step 6. Configure Network Settings

At the **Network** step of the wizard, you can select to which Amazon Virtual Private Cloud (Amazon VPC) the restored workload must be connected. You can also specify a subnet, and a security group — a virtual firewall for the restored EC2 instance. For more information on Amazon VPC, see the AWS Documentation.

To configure network settings for the restored workload, do the following:

1. From the **Amazon VPC** list, select the Amazon VPC where the restored workload will be launched.
2. From the **Subnet** list, select the subnet for the restored workload.
3. From the **Security group list**, select a security group that will be associated with your restored workload.
4. From the **Public IP address** list, select one of the following:
   - **Assign (restored VM will be accessible from the Internet)** — if you want to assign a public IP to the restored workload. For security reasons, make sure traffic filtration rules are properly configured in the security group.
   - **Do not assign (more secure)** — if you do not want to assign a public IP.



Step 7. Specify Secure Restore Settings

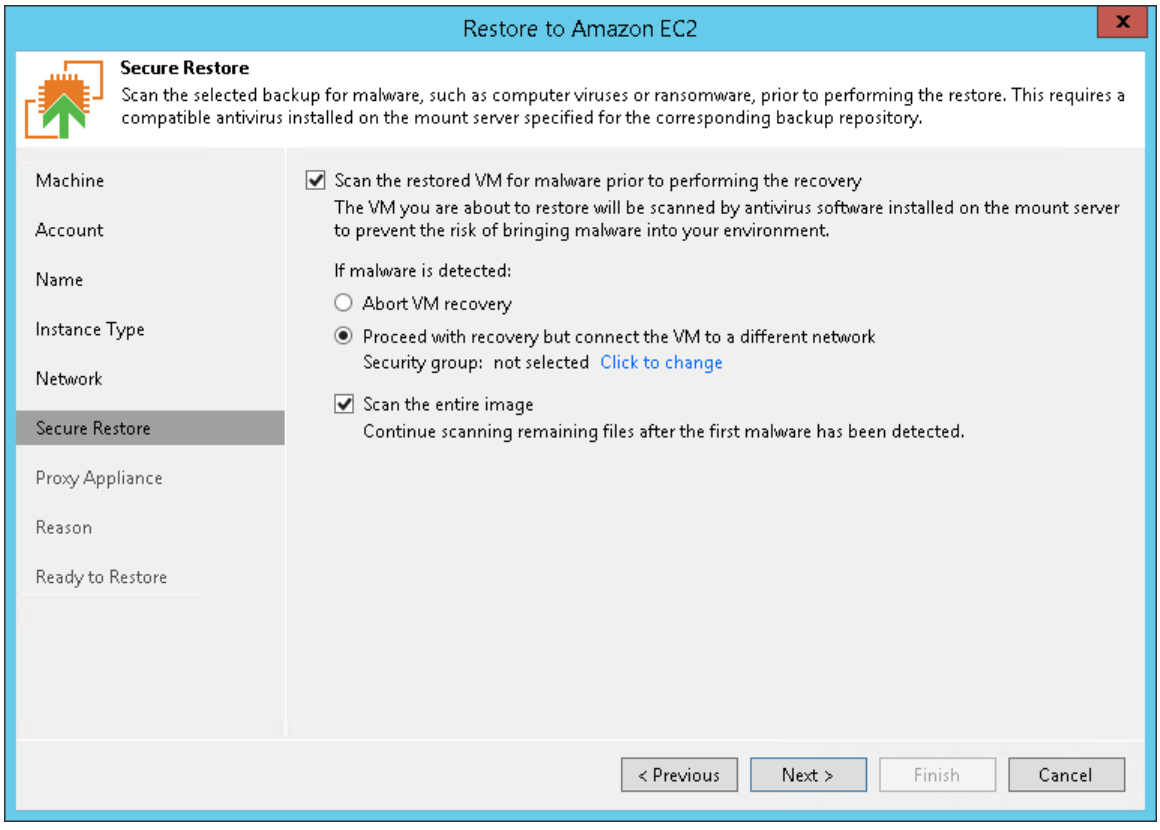
This step is available if you restore Microsoft Windows-based workloads.

You can instruct Veeam Backup & Replication to perform secure restore — scan workload data with antivirus software before restoring the workload to Amazon EC2. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

1. At the **Secure Restore** step of the wizard, select the **Scan the restored VM for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software finds a virus threat:
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
   - **Proceed with recovery but connect the VM to a different network**. Select this action if you want to restore the workload to a different AWS security group.

     Click the **Click to change** link to select the security group.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue workload scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

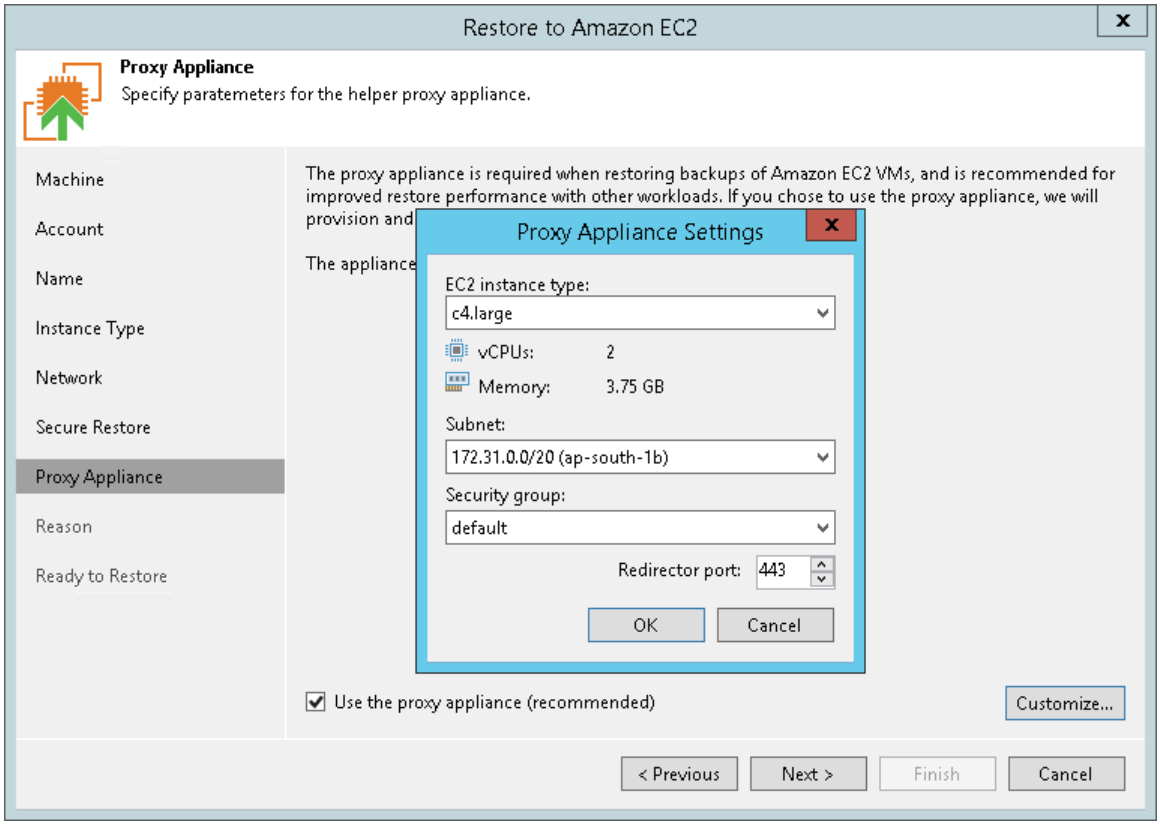


Step 8. Specify Proxy Appliance

At the **Proxy Appliance** step of the wizard, you can specify proxy appliance settings. A proxy appliance is an auxiliary Linux-based instance used to upload disks of a backed-up workload to Amazon EC2. For more information on the proxy appliance and requirements for it, see Proxy Appliance.

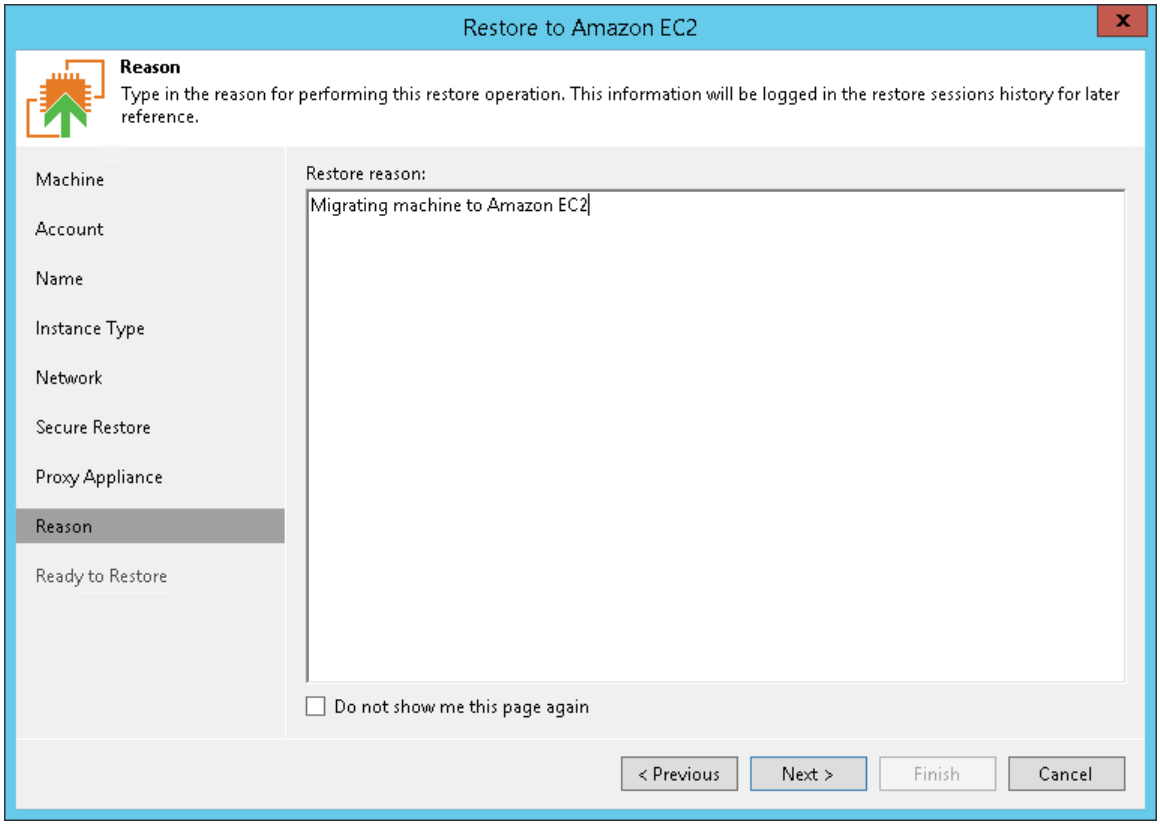
To specify proxy appliance settings, do the following:

1. Select the **Use the proxy appliance** check box.
2. Click **Customize**.
3. From the **EC2 instance type** list, select the instance type for the proxy appliance.
4. From the **Subnet** list, select the subnet for the proxy appliance.
5. From the **Security group list**, select a security group that will be associated with the proxy appliance.
6. 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.



Step 9. Specify Restore Reason

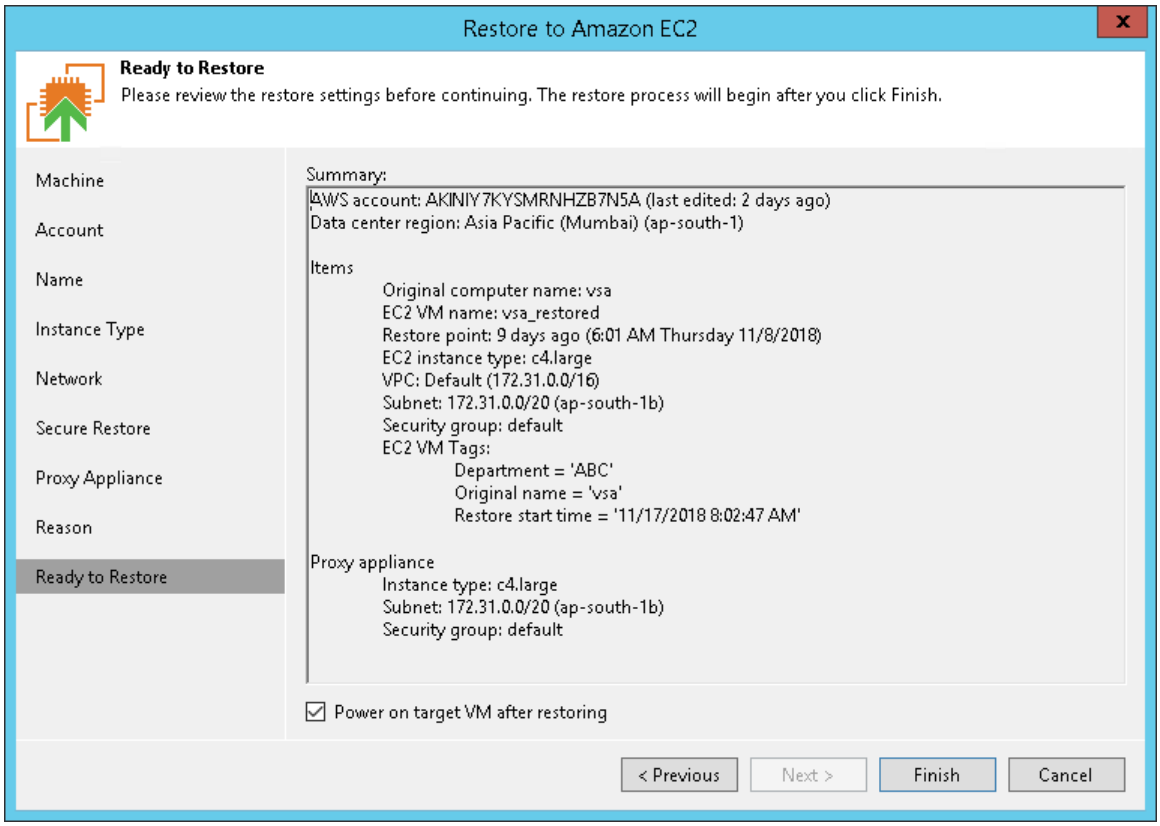
At the **Reason** step of the wizard, enter a reason for restoring the workload. The information you provide will be saved in the session history in Veeam Backup & Replication, and you can view it later.



Step 10. Verify Restore Settings

At the **Ready to Restore** step of the wizard, check the specified settings and click **Finish**. If you want to start the EC2 instance right after restore, select the **Power on VM after restoring** check box.

After the wizard closes, you can track the restore process in the **Restore Session** window. If you need to cancel the workload restore, click the **Cancel** restore task link.



Restore to Google Compute Engine

> **NOTE**
> This feature is available in Veeam Backup & Replication starting from version 11a (build 11.0.1.1261).

Veeam Backup & Replication allows you to restore different workloads (VMs, Google VM instances, physical servers and so on) to Google Compute Engine as VM instances. A VM instance is a virtual machine in Google Compute Engine with a preconfigured combination of computing resources.

You can use Veeam Backup & Replication to perform the following operations:

- Restore machines to Google Compute Engine from backups.
- Migrate machines from the on-premises infrastructure to the cloud.
- Create a test environment in the cloud for troubleshooting, testing patches and updates, and so on.

Supported Backup Types

You can restore machines from the backups of the following types:

- Backups of VMware vSphere or vCloud Director VMs created by Veeam Backup & Replication.
- Backups of Microsoft Hyper-V VMs created by Veeam Backup & Replication.
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.

  Backups must be created at the entire machine level or volume level.
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud Platform.
- Backups of Amazon EC2 VM instances created by Veeam Backup for AWS.
- Backups of Microsoft Azure VMs created by Veeam Backup for Microsoft Azure.
- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV.
- Backups of RHV VMs created by Veeam Backup for RHV.

Proxy Appliances

Proxy appliance is an auxiliary Linux-based VM instance. It is used to upload backed-up data to Google Compute Engine. Veeam Backup & Replication automatically deploys the proxy appliance in Google Compute Engine only for the duration of the restore process and removes it immediately after that.

Depending on the type of backups you are restoring from and their location, the proxy appliance may be required or optional. The proxy appliance is required when you restore from:

- Backups of Google Compute Engine VM instances that are stored in external repositories.
- Backups of virtual and physical machines that are stored in object storage repositories.

The proxy appliance is optional when you restore from backups of virtual and physical machines stored in backup repositories, or backups of Google Compute Engine virtual machines copied to backup repositories with backup copy jobs. It is recommended, however, to use the proxy appliance in scenarios where it is optional, as the proxy appliance may significantly improve restore performance. You can specify the proxy appliance settings at the Proxy Appliance step of the **Restore to Google Compute Engine** wizard.

Requirements for Proxy Appliance

When configuring a proxy appliance, mind the following:

- If you want to restore from backups in an on-premise object storage repository, the proxy appliance machine must have access to the source object storage repository. To provide access to object storage repositories, you can use VPN or Google Dedicated Interconnect. For more information, see the Google Cloud documentation.
- To upload one machine disk to Google Compute Engine, the proxy appliance requires 1 GB RAM. Make sure that the type of Google Compute Engine instance selected for the proxy appliance offers enough memory resources to upload all machine disks. Otherwise, the restore process may fail.
- The VPC route table must contain a route from the IP address of the Veeam Backup & Replication server to an active Google Cloud internet gateway. For more information on internet gateways and how to create route tables, see the Google Cloud documentation.

How Restore to Google Compute Engine Works

The workflow of the restore process depends on whether the proxy appliance is used or not. For more information on the proxy appliance, see Proxy Appliance.

> **NOTE**
>
> If you add the Google Compute Engine appliance to the Veeam Backup & Replication infrastructure and plan to restore Google Compute Engine virtual machines from restore points that were created using the appliance, you do not need to configure the proxy appliance. Also, restore to Google Compute Engine works as described in the Performing Instance Restore section in the Veeam Backup for Google Cloud Platform User Guide.

Restoring to Google Compute Engine with Proxy Appliance

If the proxy appliance is used for restore to Google Compute Engine, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a proxy appliance in Google Compute Engine.

   During the restore process, the proxy appliance communicates with backup infrastructure components over the SSH protocol and the network redirector that is deployed on the proxy appliance.
2. For every disk of a backed-up workload, Veeam Backup & Replication creates a disk in Google Compute Engine.
3. Veeam Backup & Replication hot-adds empty disks to the proxy appliance and restores backed-up data to the disks.
4. Veeam Backup & Replication creates a target instance in Google Compute Engine.
5. Veeam Backup & Replication detaches the disks from the proxy appliance and attaches them to the target instance.
6. After the restore process is complete, Veeam Backup & Replication removes the proxy appliance from Google Compute Engine.

Restoring to Google Compute Engine without Proxy Appliance

If the proxy appliance is not used for restore to Google Compute Engine, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication uploads disks of a backed-up workload to Google Cloud Storage bucket.

   In Google Cloud Storage bucket, the uploaded disks are stored to the temporary bucket in the RAW format.

2. Veeam Backup & Replication imports the backed-up data from the temporary bucket in Google Cloud Storage to disks in Google Compute Engine.

3. Veeam Backup & Replication creates a target instance in Google Compute Engine and attaches disks to the target instance.

4. After the import process is complete, Veeam Backup & Replication removes the temporary bucket from Google Cloud Storage.

Google Compute Engine IAM User Permissions

To enable restore of workloads to Google Compute Engine, do the following:

1. Grant the following roles to the IAM user whose credentials you plan to use to connect to Google Compute Engine:
   - Compute Admin role (roles/compute.admin)

     To avoid granting the Compute Admin role to the IAM user Compute Engine service account for security reasons, you can create a custom role with the following Compute Engine IAM permissions and grant it instead:

     ```
     compute.addresses.list
     compute.disks.create
     compute.disks.delete
     compute.disks.get
     compute.disks.use
     compute.disks.useReadOnly
     compute.firewalls.create
     compute.firewalls.delete
     compute.firewalls.list
     compute.globalOperations.get
     compute.images.create
     compute.images.delete
     compute.images.get
     compute.images.useReadOnly
     compute.instances.attachDisk
     compute.instances.create
     compute.instances.delete
     compute.instances.detachDisk
     compute.instances.get
     compute.instances.getGuestAttributes
     compute.instances.list
     compute.instances.setLabels
     compute.instances.setMetadata
     compute.instances.setTags
     compute.instances.stop
     compute.machineTypes.list
     compute.networks.get
     compute.networks.list
     compute.networks.updatePolicy
     compute.projects.get
     compute.regions.list
     compute.subnetworks.get
     compute.subnetworks.list
     compute.subnetworks.use
     compute.subnetworks.useExternalIp
     compute.zoneOperations.get
     compute.zones.get
     compute.zones.list
     ```

   - Cloud Build Editor role (roles/cloudbuild.builds.editor)
   - Project IAM Admin role (roles/resourcemanager.projectIamAdmin)
   - Storage Admin role (roles/storage.admin)
   - Storage HMAC Key Admin (roles/storage.hmacKeyAdmin)

- Viewer role (roles/viewer)

For more information, see the *Prerequisites for importing and exporting VM images* section in the Google Cloud documentation.

2. Make sure that the Cloud Build API is enabled. Then grant the following roles to the Cloud Build service account in Google Compute Engine:
   - Compute Admin role (roles/compute.admin)

     To avoid granting the Compute Admin role to the Cloud Build service account for security reasons, you can use the custom role that you created for the IAM user Compute Engine service account and grant it instead.
   - Service Account Token Creator role (roles/iam.serviceAccountTokenCreator)
   - Service Account User role (roles/iam.serviceAccountUser)
   - [Optional: to export or import images that use shared VPCs] Compute Network User role (roles/compute.networkUser)

   For more information, see the *Prerequisites for importing and exporting VM images* section in the Google Cloud documentation.

Restoring to Google Compute Engine

To restore workloads to Google Compute Engine, use the **Restore to Google Compute Engine** wizard.

> **NOTE**
>
> If you add the Google Compute Engine appliance to the Veeam Backup & Replication infrastructure and plan to restore Google Compute Engine virtual machines from restore points that were created using the appliance, you do not need to configure the proxy appliance. Also, restore to Google Compute Engine works as described in the Performing Instance Restore section in the Veeam Backup for Google Cloud Platform User Guide.

Before You Begin

Before you restore workloads to Google Compute Engine, mind the following requirements and limitations:

- Check whether a proxy appliance must be configured for restore. For more information, see Proxy Appliances.
- The backup server and repositories with workload backup files must have access to the Internet.

  If backup files are located on deduplicating storage appliances or shared folder repositories, the Internet connection is required for gateway servers that communicate with these repositories.
- You must have a backup of the workload that you plan to restore to Google Compute Engine.
- Make sure that the Cloud Build API is enabled. For more information on enabling the Cloud Build API and other requirements for importing virtual disks into Google Compute Engine, see the Google Cloud Documentation.
- Make sure the IAM service account that you plan to use to restore workloads to Google Compute Engine has permissions to restore to Google Compute Engine. For more information, see Google Compute Engine IAM User Permissions.

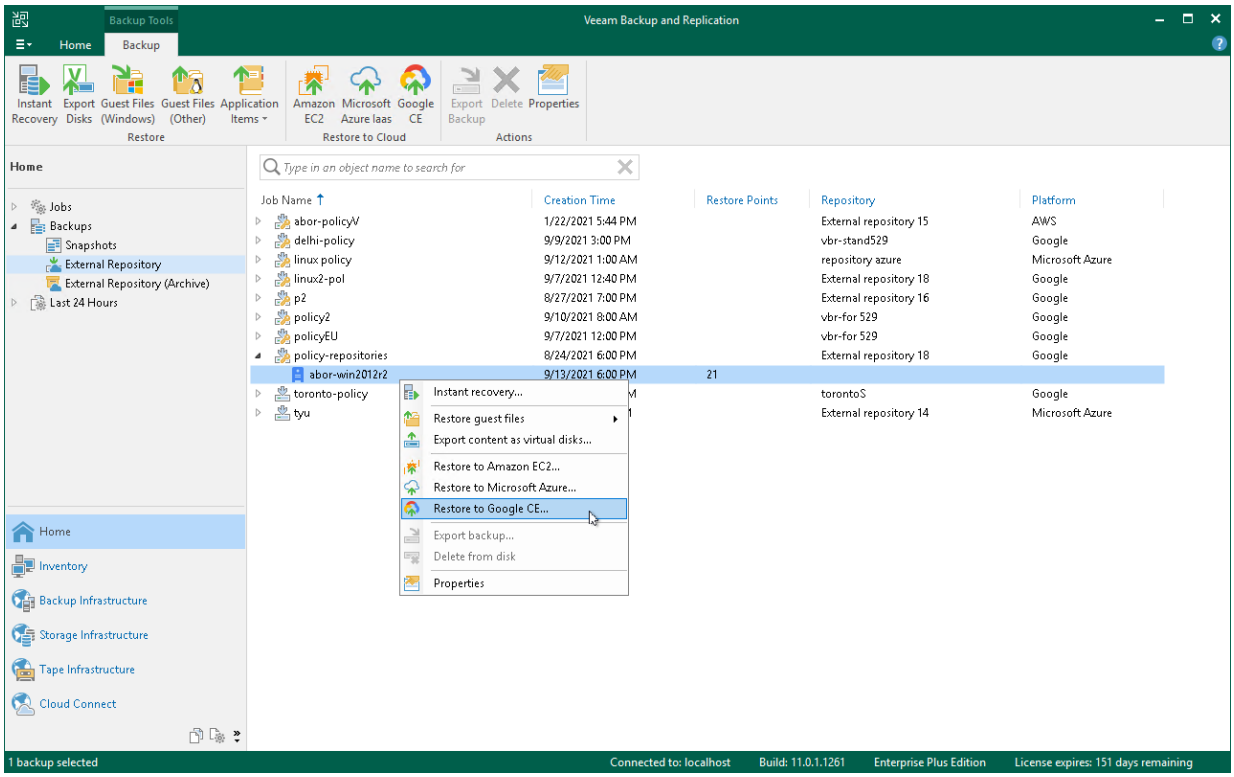
- If you restore workloads from backups of virtual and physical machines (non-Google Compute Engine virtual machines), check the supported operating systems and their differences from standard images in the Google Cloud documentation.
- If you plan to assign Google labels to the restored workload, check limitations for labels in the Google Cloud documentation.

Step 1. Launch Restore to Google Compute Engine Wizard

To begin the restore process, do one of the following.

- On the **Home** tab, click **Restore** and select the type of backups from which you want to restore:
  - VMware vSphere or vCloud Director
  - Microsoft Hyper-V
  - Agent
  - AWS EC2 backup
  - Microsoft Azure backup
  - GCE backup
  - Nutanix AHV
  - RHV Proxy

  In the displayed window, click **Entire VM restore** > **Restore to public cloud** > **Restore to Google Compute Engine**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select workloads that you want to restore and click **Restore to Google CE** on the ribbon. Alternatively, you can right-click one of the workloads that you want to restore and select **Restore to Google CE**.
- Double-click a full backup file (VBK) or backup metadata file (VBM) in a file browser. Veeam Backup & Replication will start its console. In the **Backup Properties** window, select the necessary workload and click **Restore** > **Restore to Google CE**.



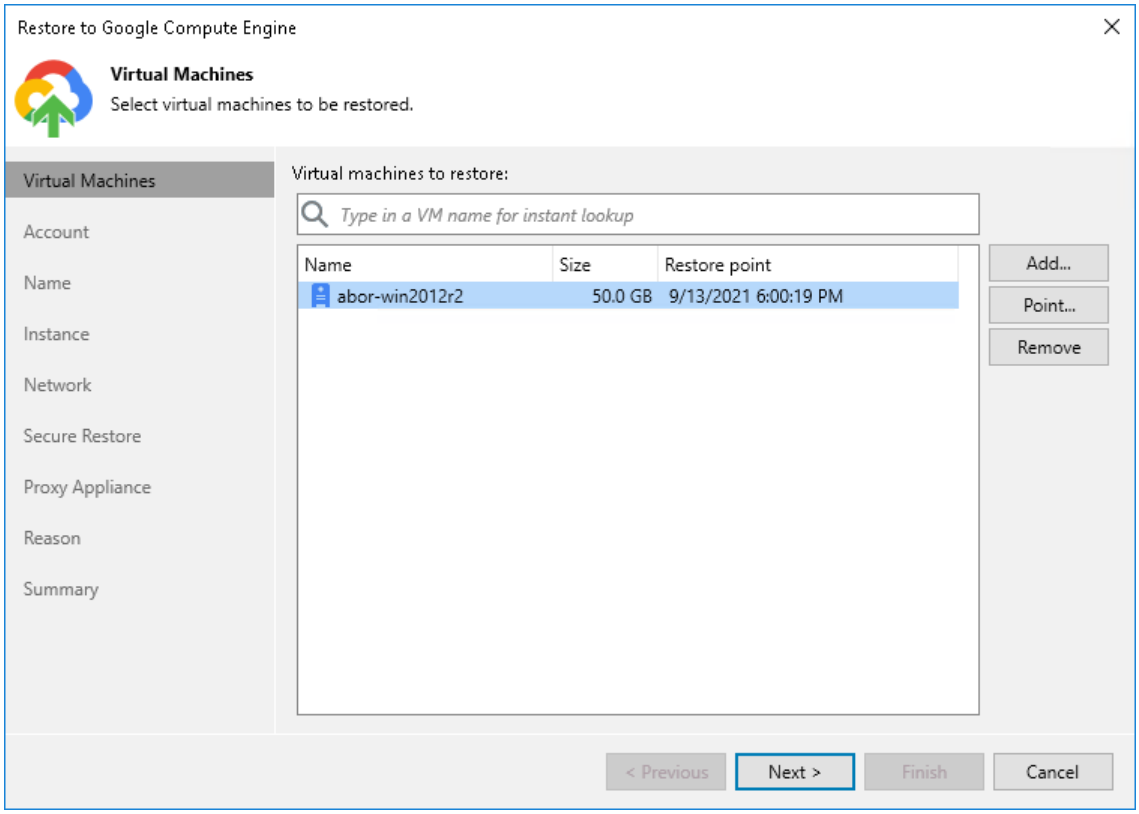
Step 2. Select Workloads and Restore Points

At the **Virtual Machines** step of the wizard, specify the workload that you plan to restore and specify a restore point to which you want to restore the workload. By default, Veeam Backup & Replication restores workloads to the latest valid restore point in the backup chain.

Selecting Workloads

To select workloads to restore:

1. On the right of the **Virtual machines to restore** list, click **Add**.
2. In the **Backup Browser** window, expand the necessary backup, select workloads and click **Add**.

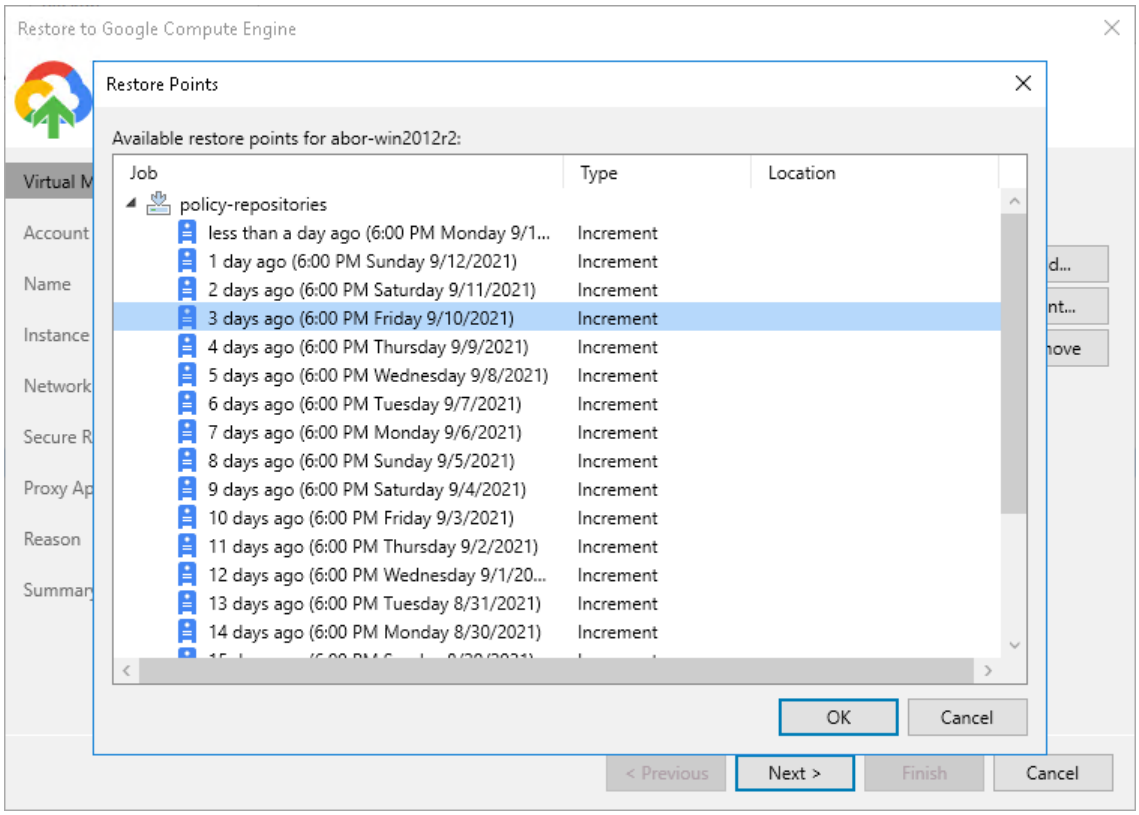


Selecting Restore Points

To select a restore point, for a workload do the following:

1. In the **Virtual machines to restore** list, select a workload.
2. Click **Point** on the right.

3. In the **Restore Points** window, select a restore point to which you want to restore the workload.



Step 3. Specify Credentials and Datacenter Settings

At the **Account** step of the wizard, specify a Google Cloud Platform service account, data center and availability zone to use for restore:

1. From the **Google Cloud Platform service account** list, select user credentials to connect to Google Compute Engine.

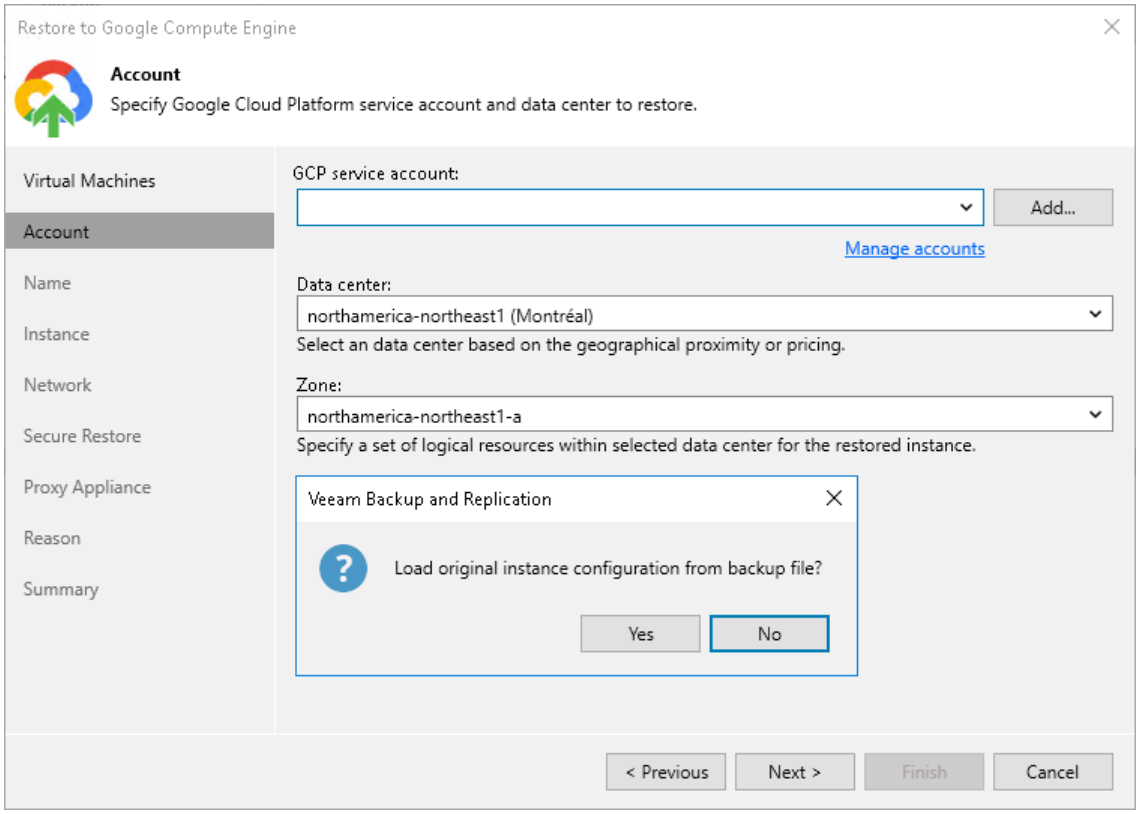
   When you add credentials of the Google Cloud Platform service account, Veeam Backup & Replication imports information about resources associated with this service account. During the restore process, Veeam Backup & Replication accesses these resources and uses them to create a target VM instance in Google Compute Engine.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage accounts** link or click **Add** on the right to add the necessary credentials, as described in Google Cloud Platform Service Accounts.

2. From the **Data center** list, select the Google Cloud datacenter where Veeam Backup & Replication will restore your workload as a VM instance.

3. From the **Zone** list, select the availability zone inside the Google Cloud datacenter where the restored workload will reside.

If you restore a Google Compute Engine VM instance from a backup created with Veeam Backup for Google Cloud Platform to the same Google Cloud region where the instance is placed, after you click **Next**, the wizard will offer you to use region settings associated with this instance.



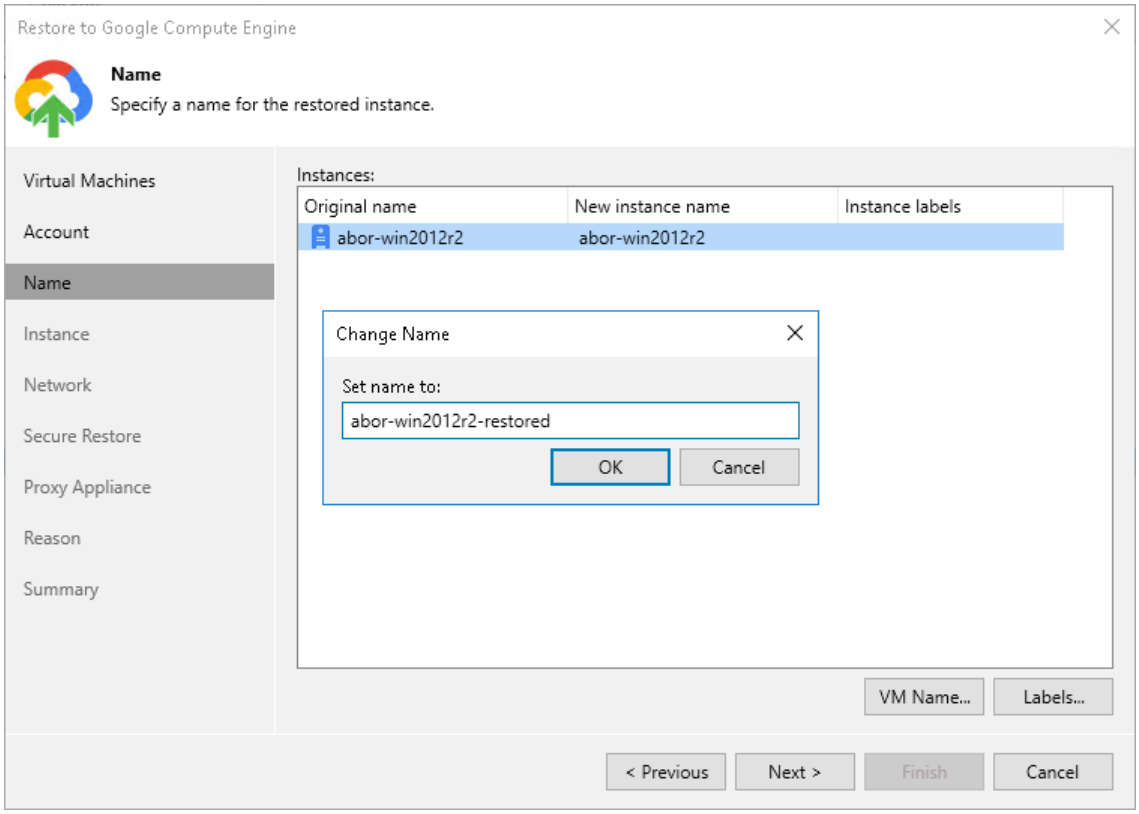
Step 4. Specify Name and Labels

At the **Name** step of the wizard, you can specify names and manage Google labels for the restored workloads. By default, Veeam Backup & Replication uses original workload names.

Specifying New Name

To define a new name for a restored workload:

1. In the **Instances** list, select a workload and click **VM Name**.
2. In the **Set name to** field of the **Change Name** window, enter a new name for the restored workload.



Managing Google Cloud Labels

You can use Google Cloud labels to categorize instances in Google Compute Engine. A label is a tag with metadata that includes two properties: a key and a value. For more information on Google Cloud labels their format and limitations, see the Google Cloud documentation.

Adding Label

To add a new label:

1. In the **Instances** list, select a workload and click **Labels**.
2. In the **Labels** window, click **Add**.
3. In the **GCE Instance Label** window, specify the **Key** and **Value** properties.

Modifying Label

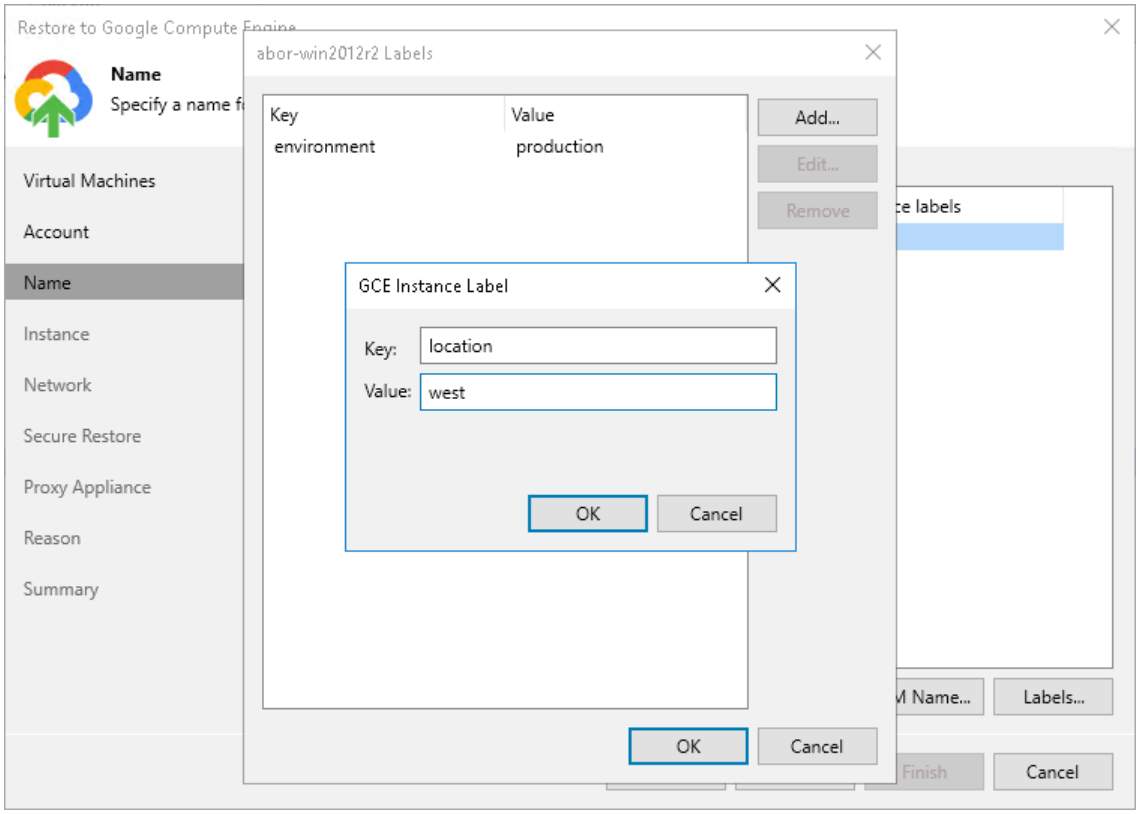
To modify a label:

1. In the **Instances** list, select a workload and click **Labels**.
2. In the **Labels** window, select the required label and click **Edit**.
3. In the **GCE Instance Label** window, edit the **Key** or **Value** properties.

Deleting Label

To delete a label:

1. In the **Instances** list, select a workload and click **Labels**.
2. In the **Labels** window, select the required label and click **Remove**.



Step 5. Specify Instance Type and Disks

At the **Instance Type** step of the wizard, select instance types and disk types for the restored workloads. By default, Veeam Backup & Replication restores all disks as Google Compute Engine (GCE) disks of the Balanced persistent disk type. For information on types of GCE disks, see the Google Cloud documentation.

Selecting Instance Type

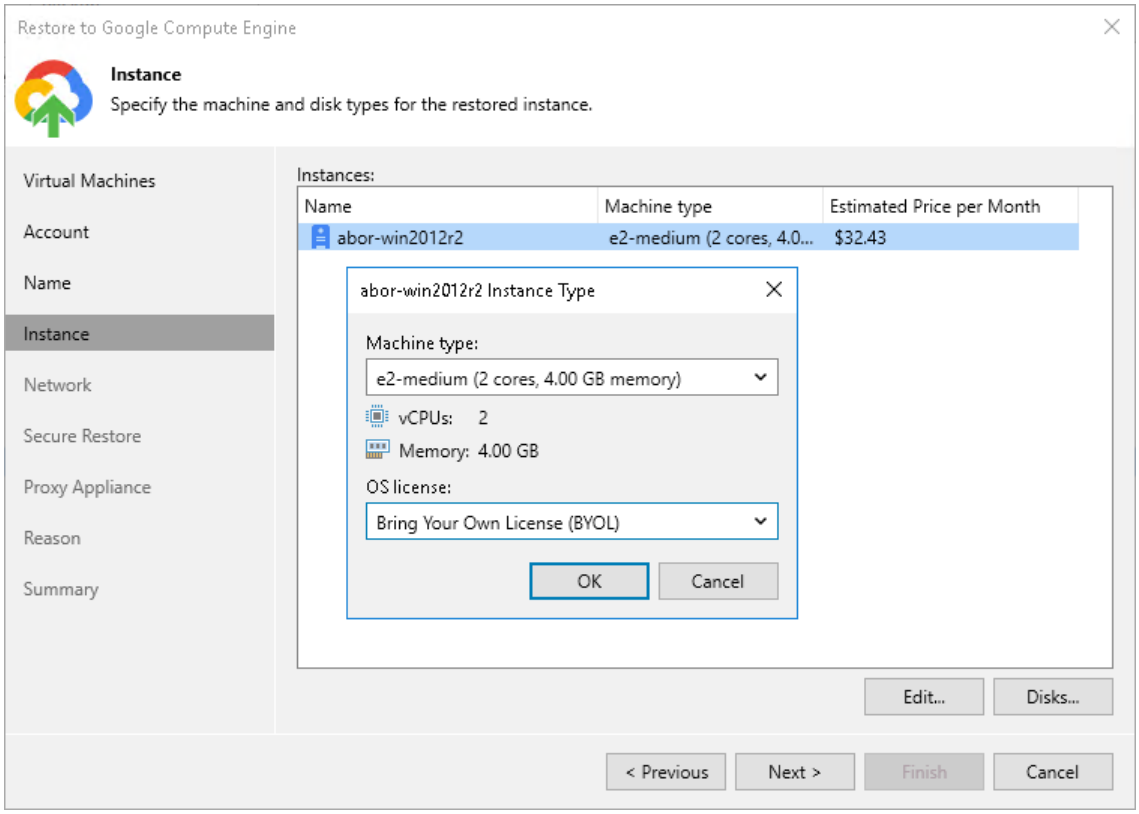
You can select the amount of computing resources that Google Compute Engine will provision for your restored workload — a Google Compute Engine instance type. Each instance type offers a unique combination of CPU and memory resources.

To select an instance type for a workload:

1. In the **Instances** list, select a workload and click **Edit**.
2. From the **Machine type** list, select the instance type for the restored workload.

   Make sure that you select the right instance type that corresponds to the initial workload configuration. For the information on instance types, see Google Cloud Documentation.

   Note that if you restore a Google Compute Engine instance from the backup created with Veeam Backup for Google Cloud Platform, Veeam Backup & Replication will identify the type of a backed-up instance and select it by default.
3. From the **OS license** list, select an option that will define what license Google Compute Engine will use for the OS on the restored workload:
   - **Provided by GCE** — the OS license will be provided by Google Compute Engine.
   - **Bring Your Own License (BYOL)** — the OS license will be restored from the backup. For more information, see Google Cloud documentation.

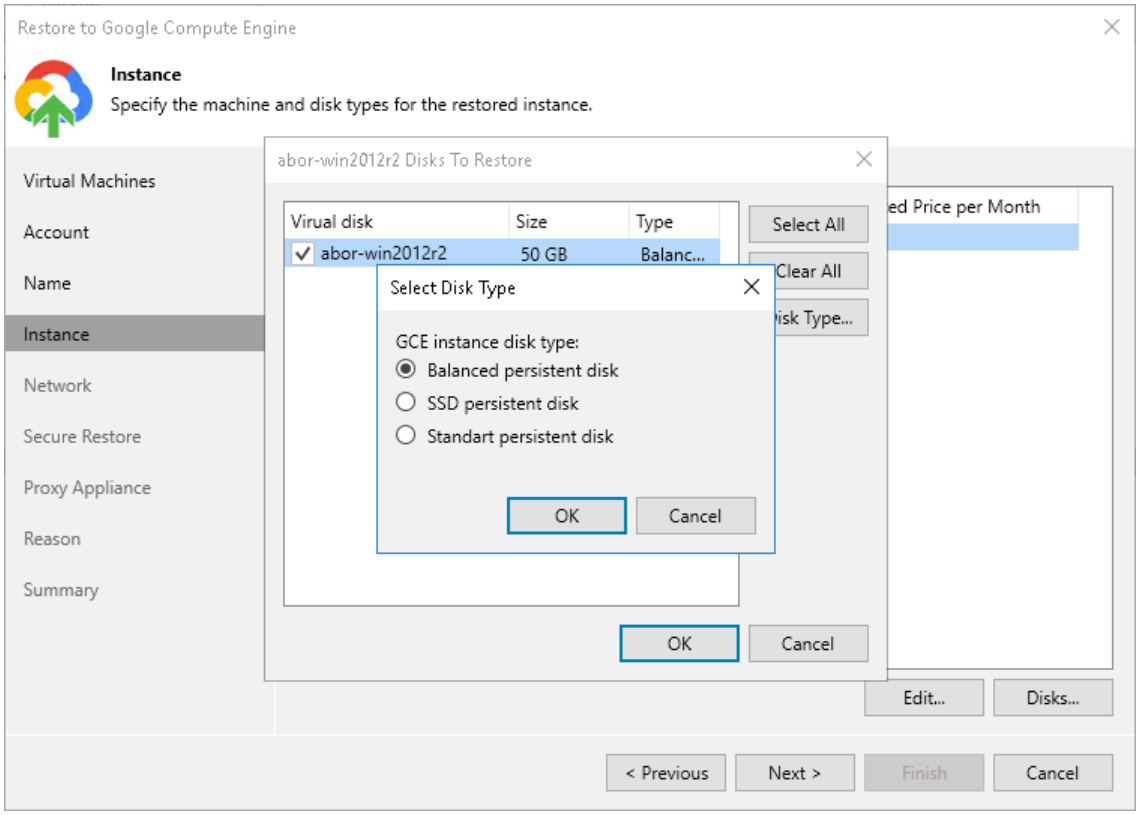


Selecting Disk Type

You can restore all disks or specific disks of a workload. You can also change disk types of the restored disks.

To select workload disks for restore:

1. In the **Virtual machines** list, select a workload and click **Disks**.
2. In the **Disks To Restore** window, make sure that check boxes next to disks that you want to restore are selected. Clear check boxes next to disks that you do not want to restore.
3. Select a disk whose type you want to change and click **Disk type**.
4. In the **Select EC2 Disk Type** window, choose the disk type.

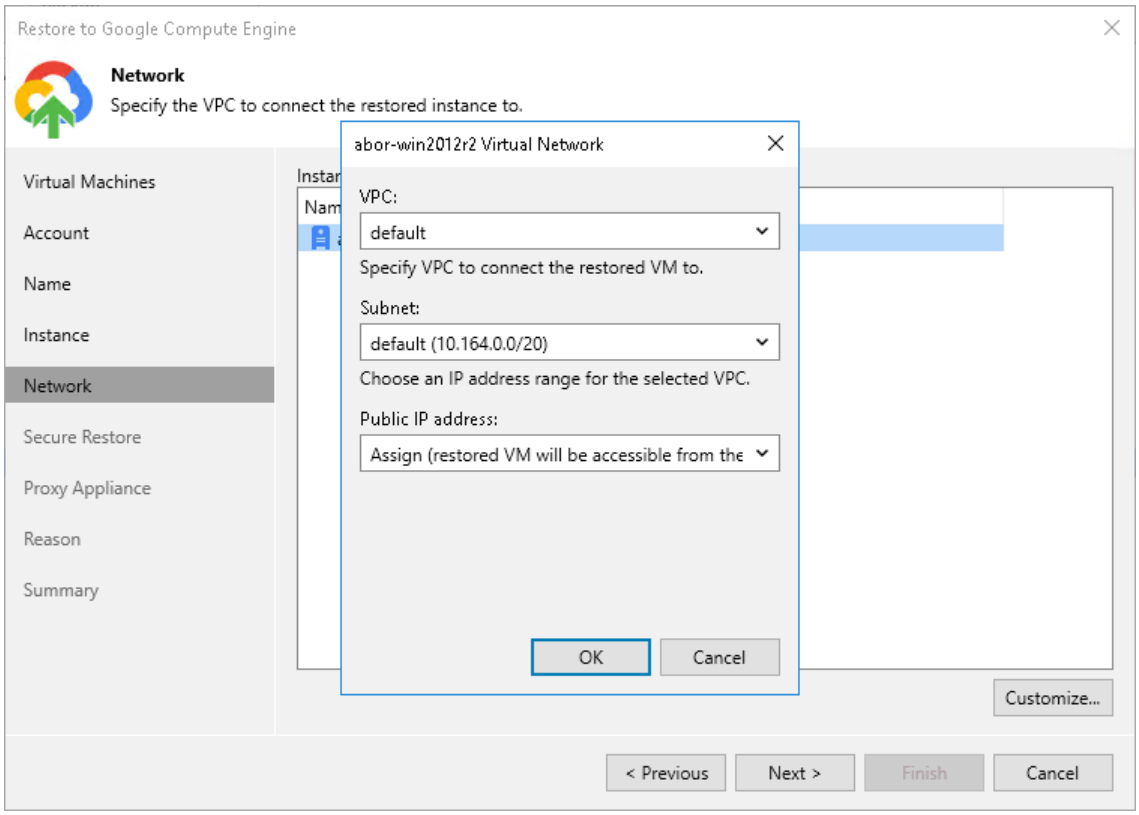


Step 6. Select Google VPC

At the **Network** step of the wizard, you can select to which Google Virtual Private Cloud (Google VPC) network the workload must be connected after restore. You can also specify a subnet. For the information on Google VPC, see the Google Cloud documentation.

To define network settings for the restored workload, do the following:

1. From the **VPC** list, select the VPC where the restored workload will be launched.
2. From the **Subnet** list, select an IP address range for the selected VPC.
3. From the **Public IP address** list, select one of the following:
   - **Assign (restored VM will be accessible from the Internet)** — if you want to assign a public IP to the restored workload. For security reasons, make sure firewall rules are properly configured in the target VPC.
   - **Do not assign (more secure)** — if you do not want to assign a public IP.



Step 7. Specify Secure Restore Settings

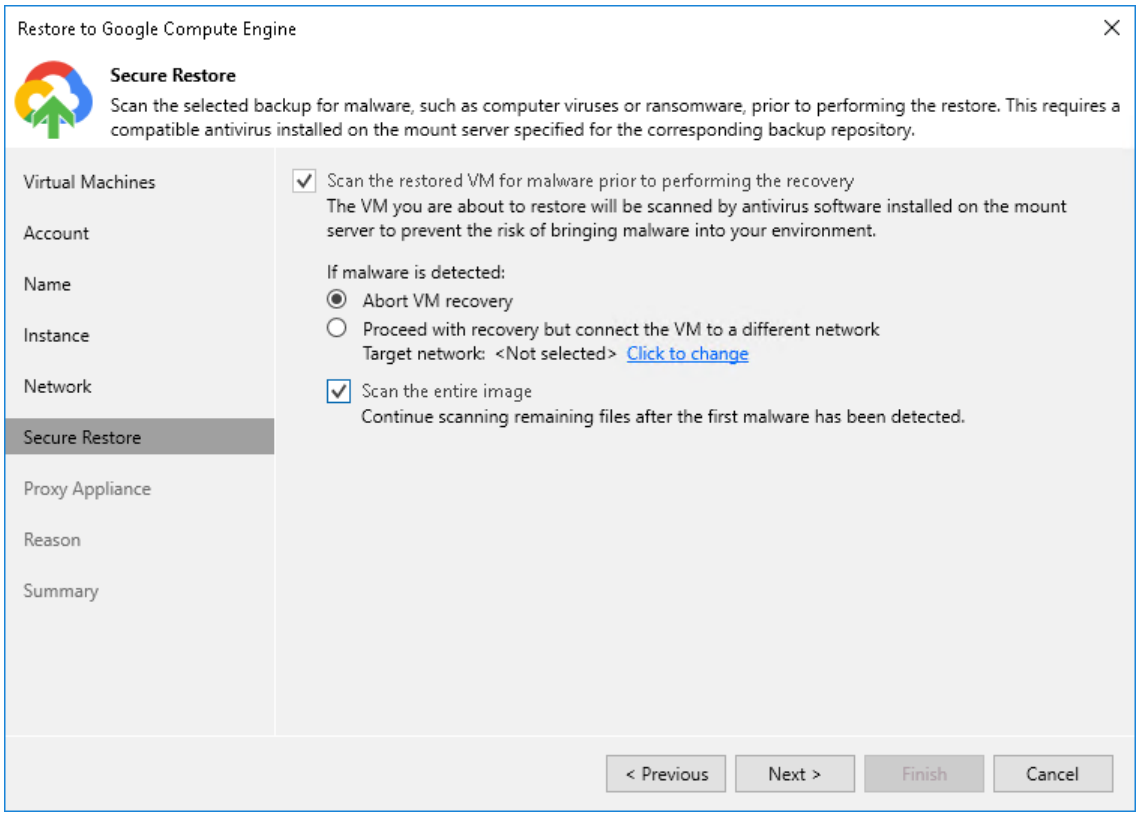
This step is available if you restore Microsoft Windows workloads.

You can instruct Veeam Backup & Replication to perform secure restore — scan workload data with antivirus software before restoring the workload to Google Compute Engine. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

1. At the **Secure Restore** step of the wizard, select the **Scan the restored VM for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software finds a virus threat:
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
   - **Proceed with recovery but connect the VM to a different network**. Select this action if you want to restore the workload to a different Google Cloud network.

     Click the **Click to change** link to specify the VPC and subnet for this network.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue workload scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

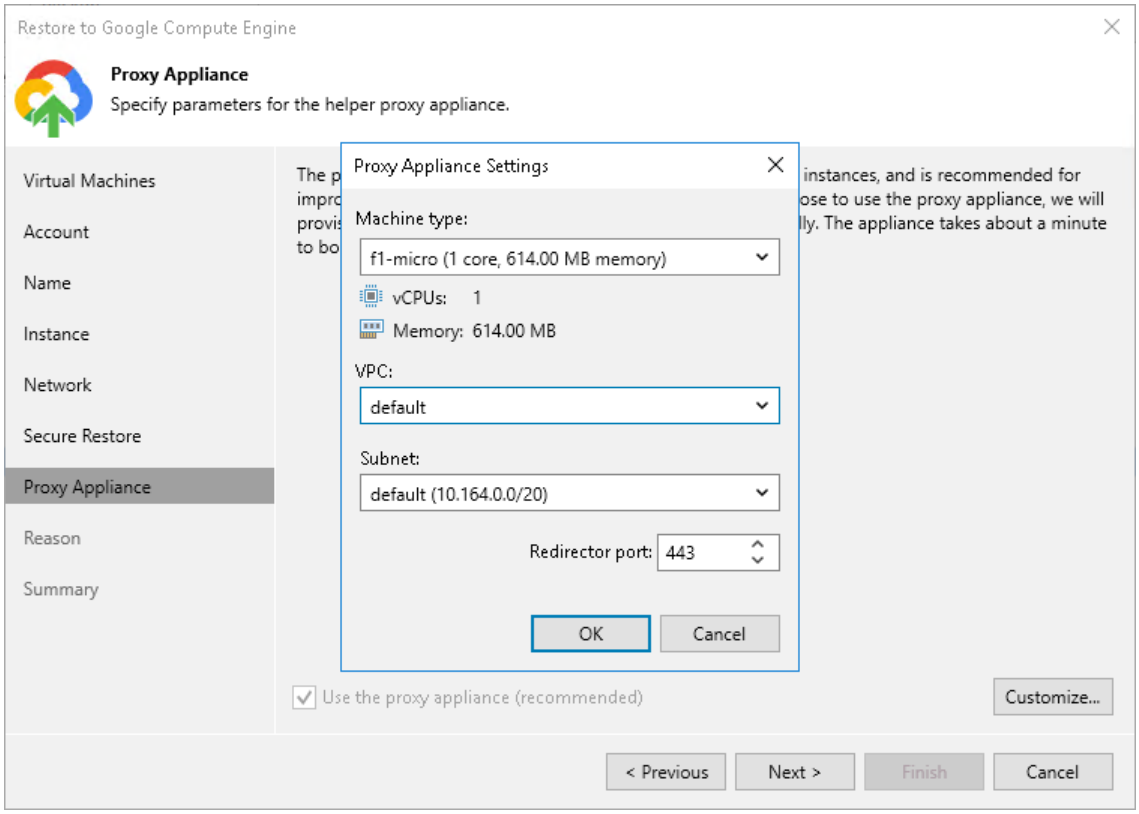


Step 8. Specify Proxy Appliance

At the **Proxy Appliance** step of the wizard, you can specify proxy appliance settings. A proxy appliance is an auxiliary Linux-based instance used to upload disks of a backed-up workload to Amazon EC2. For more information on the proxy appliance and requirements for it, see Proxy Appliances.

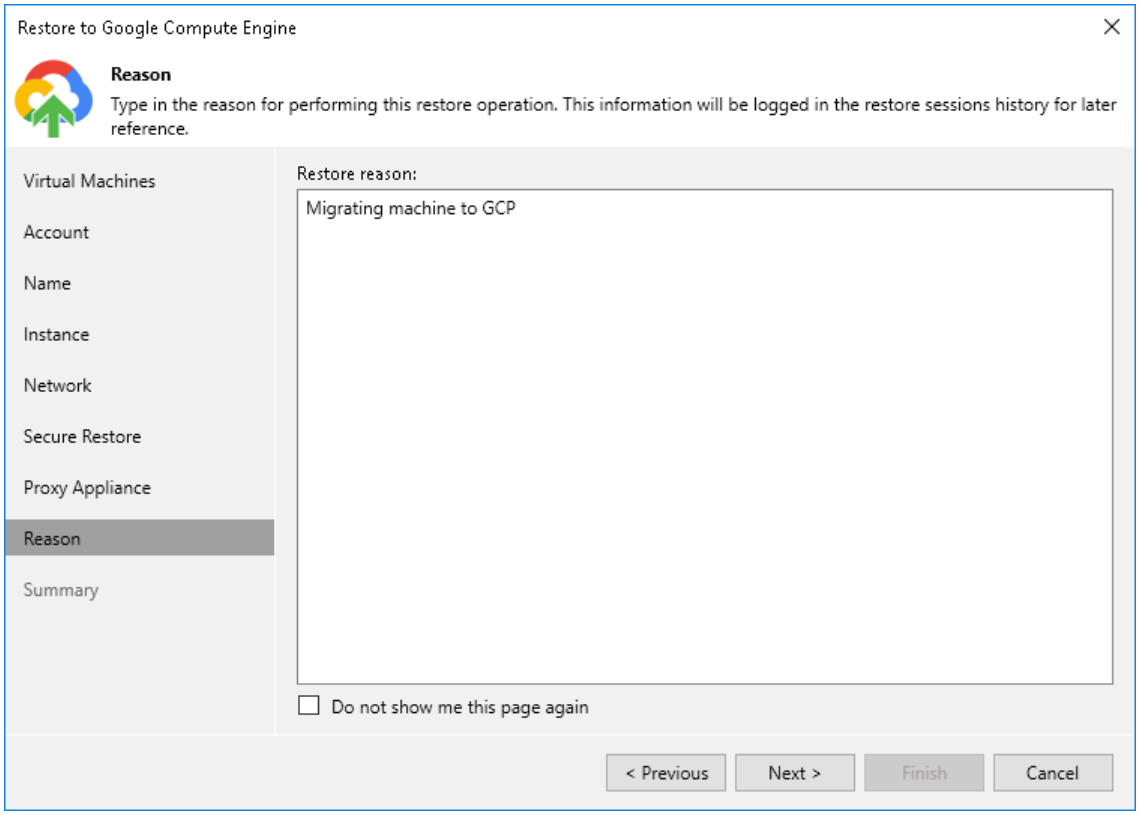
To specify proxy appliance settings, do the following:

1. Select the **Use the proxy appliance** check box.
2. Click **Customize**.
3. From the **Machine type** list, select the instance type for the proxy appliance.
4. From the **VPC** list, select the VPC network for the proxy appliance.
5. From the **Subnet** list, select the subnet for the proxy appliance.
6. 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.



Step 9. Specify Restore Reason

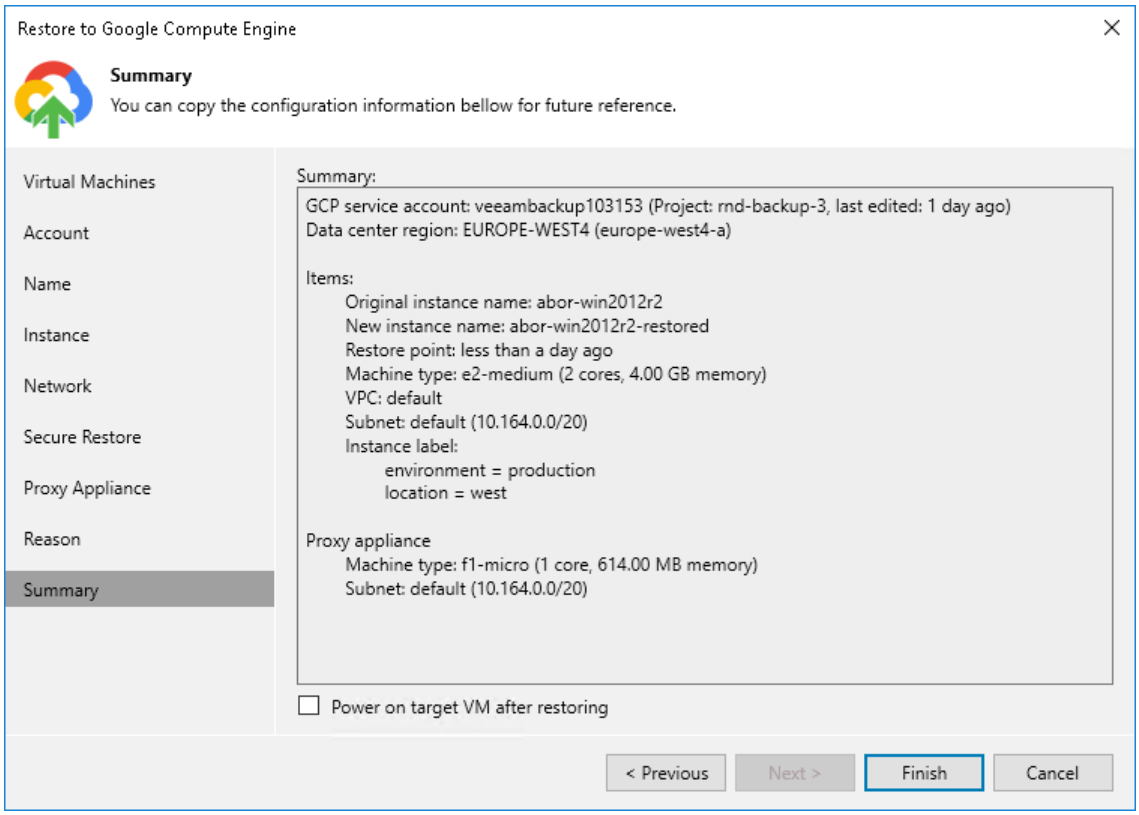
At the **Reason** step of the wizard, enter a reason for restoring the workload. The information you provide will be saved in the session history in Veeam Backup & Replication, and you can view it later.



Step 10. Verify Restore Settings

At the **Ready to Restore** step of the wizard, check the specified settings and click **Finish**. If you want to start the VM instance right after restore, select the **Power on VM after restoring** check box.

You can track the restore process in the **Restore Session** window. If you need to cancel the workload restore, click the **Cancel** restore task link.



Disk Recovery

Disk recovery includes the following methods:

- Instant Disk Recovery — to instantly recover VM disks from backup files. The disks are recovered in their initial format.

  Instant Disk Recovery helps improve recovery time objectives (RTO) but publishes disks with the reduced I/O performance, that is, provides "temporary spares". To provide the full I/O performance, you must finalize Instant Disk Recovery — migrate the recovered disks to the production environment. If you do not want to migrate the recovered disks, you can stop publishing them. This removes the recovered disks.

  Use Instant Disk Recovery if you want to keep the target VM (VM to which you want to attach recovered disks) turned on. If the VM can be turned off, you can use virtual disk restore.

- Instant First Class Disk Recovery — to instantly recover VM disks from backup files and register them as First Class Disks (FCDs). FCDs are a type of improved VMDK that allow you to perform different tasks: create snapshots, delete and restore data stored on VMDK without attaching these disks to a VM.

  This method is similar to Instant Disk Recovery except for the format in which disks are recovered.

  Use Instant FCD Recovery if you want to recover recovered disks as FCDs.

- Virtual disk restore — to restore VM disks. When you restore disks, you extract them from backups to the production storage. Virtual disk restore takes more resources to complete than Instant Disk Recovery but restores disks with full I/O performance. You also do not need to perform additional steps to finalize the restore process.

  Use virtual disk restore if the target VM can be turned off. If the VM must stay turned on, perform Instant Disk Recovery.

- Disk export — to convert disks of different workloads (EC2 instances, Microsoft Azure VMs and so on) in the VMDK, VHD or VHDX formats. Then you can export these disks to the backup infrastructure or place them on a datastore connected to an ESXi host (for VMDK disk format only).

  Use disk export to convert disks to the VMDK, VHD or VHDX formats.

Instant Disk Recovery

With Instant Disk Recovery, you can immediately restore VM disks from a backup file and publish them in the initial format. If you want to register disks as First Class Disks (FCD), see Instant First Class Disk (FCD) Recovery. For the list of all disk recovery methods and their brief descriptions, see Disk Recovery.

Use Instant Disk Recovery if you need to:

- Recover VM disks, not an entire VM. Otherwise, use Instant Recovery to VMware vSphere.
- Recover VM disks and keep the target VM (to which you want to attach recovered disks) turned on. If the VM can be turned off, use Virtual Disks Restore.

The Instant Disk Recovery is a process that must be finalized. First, test the recovered disks and then decide whether to migrate them to the production environment or stop publishing the disks. To migrate the disks, Veeam Backup & Replication uses the Quick Migration mechanism. For more information on how to finalize Instant Disk Recovery, see Finalizing Instant Disk Recovery.

How Instant Disk Recovery Works

Veeam Backup & Replication performs Instant Disk Recovery using vPower technology.
Veeam Backup & Replication performs the following steps to recover disks:

1. Checks whether the vPower NFS datastore is mounted to the ESXi host and contains VMDK files of the recovered disks.
2. [If you replace an existing disk] Unmounts and deletes the existing disk.
3. Connects the recovered disks to the target VM.
4. Initiates a creation of a protective snapshot for the target VM. If Instant Disk Recovery process fails, the protective snapshot guarantees no data loss.
5. Writes changes made to the recovered disks to redo logs on the vPower NFS server.

Performing Instant Disk Recovery

To recover disks using Instant Disk Recovery, use the **Instant Disk Recovery** wizard.

Before You Begin

Before you perform Instant Disk Recovery, consider the following:

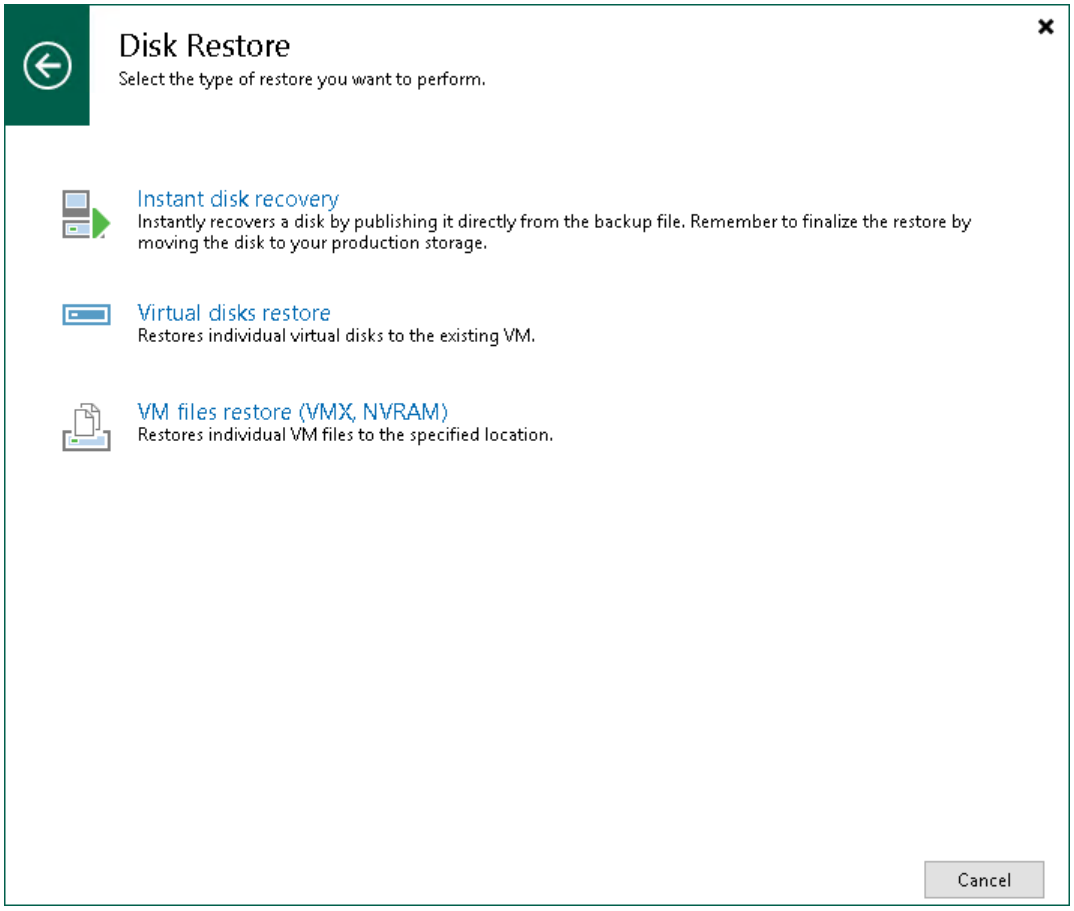
- You can restore a virtual disk from a backup that has at least one successfully created restore point.
- If you want to scan disk data for viruses, check the secure restore requirements and limitations.
- Instant Disk Recovery to vCloud Director is not supported.
- You must have at least 10 GB of free disk space on the datastore where write cache folder is located. This disk space is required to store virtual disk updates for the restored VM.

  By default, Veeam Backup & Replication writes virtual disk updates to the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`.

Step 1. Launch Disk Recovery Wizard

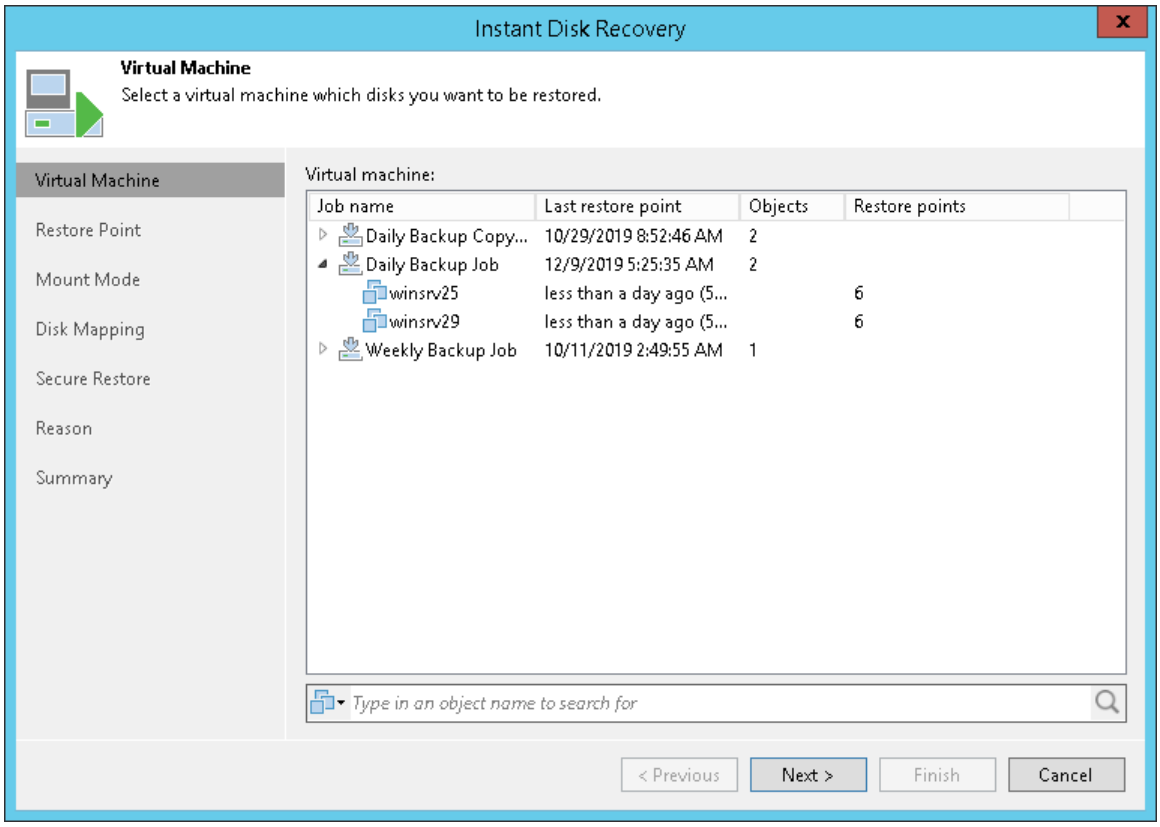
To launch the **Instant Disk Recovery** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **Instant disk recovery**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand a backup, select a VM whose disks you want to restore and click **Instant Disk Recovery** on the ribbon. Alternatively, you can right-click a VM whose disks you want to restore and select **Instant disk recovery**.



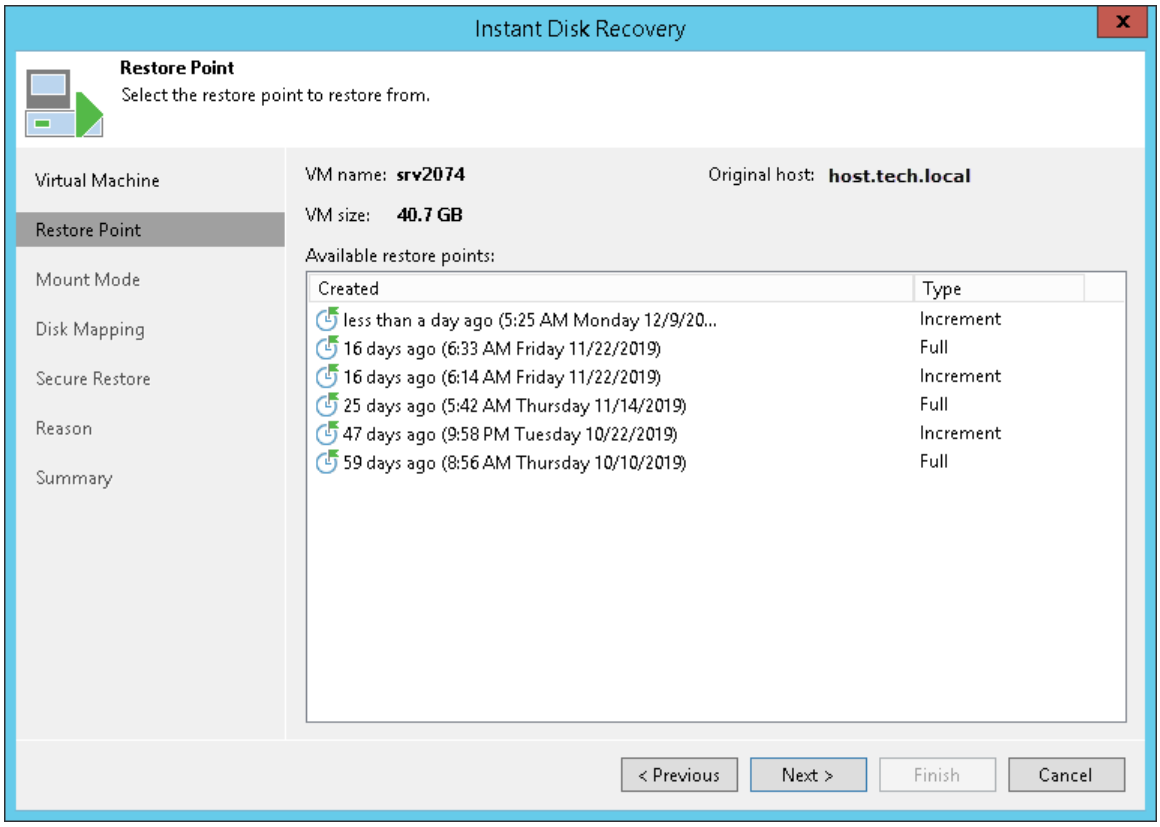
Step 2. Select Source VM

At the **Virtual Machine** step of the wizard, select a VM whose disks you want to recover:



Step 3. Select Restore Point

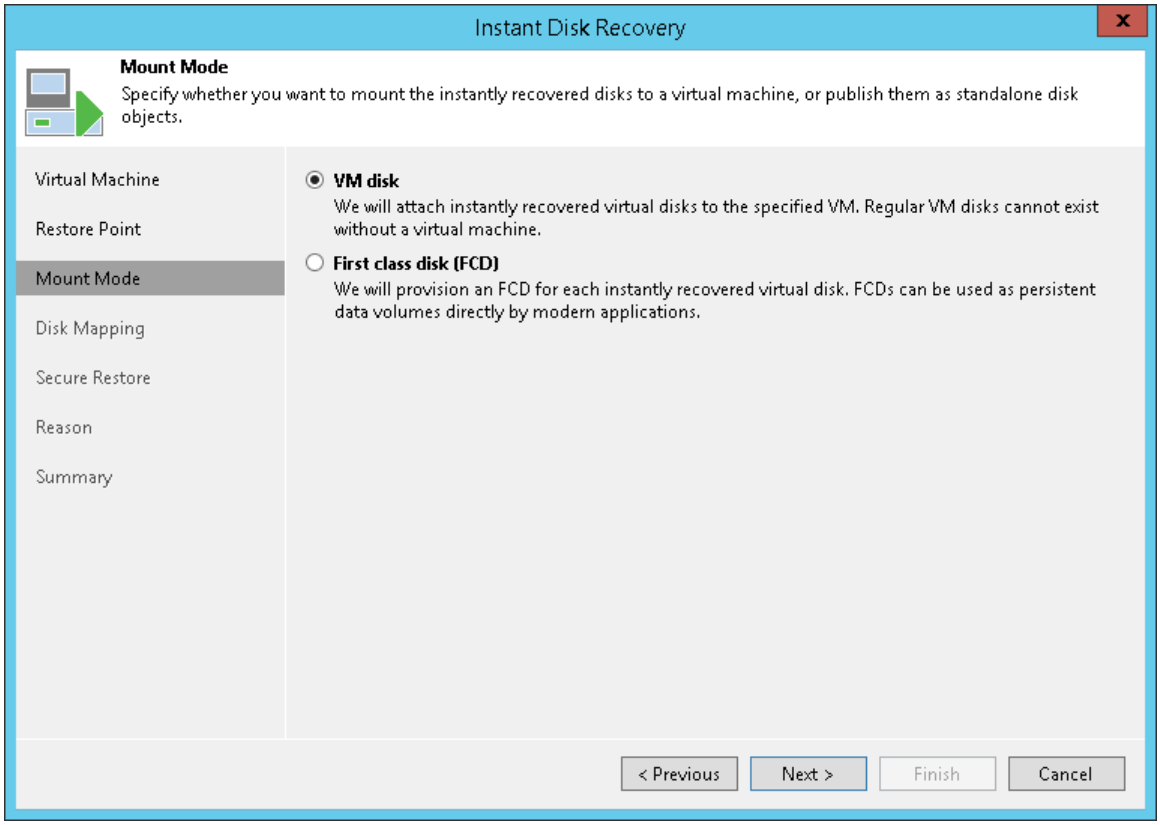
At the **Restore Point** step of the wizard, select a restore point from which you want to recover VM disks.



Step 4. Select Mount Mode

At the **Mount Mode** step of the wizard, select the **VM disk** option to register virtual disks on a VM added to an ESXi host.

If you want to register virtual disks on a cluster as FCDs, select the First class disk (FCD) option. In this case, steps of the wizard differ and Veeam Backup & Replication performs Instant FCD Recovery as described in Performing Instant First Class Disk (FCD) Recovery.



Step 5. Select Virtual Disks to Restore

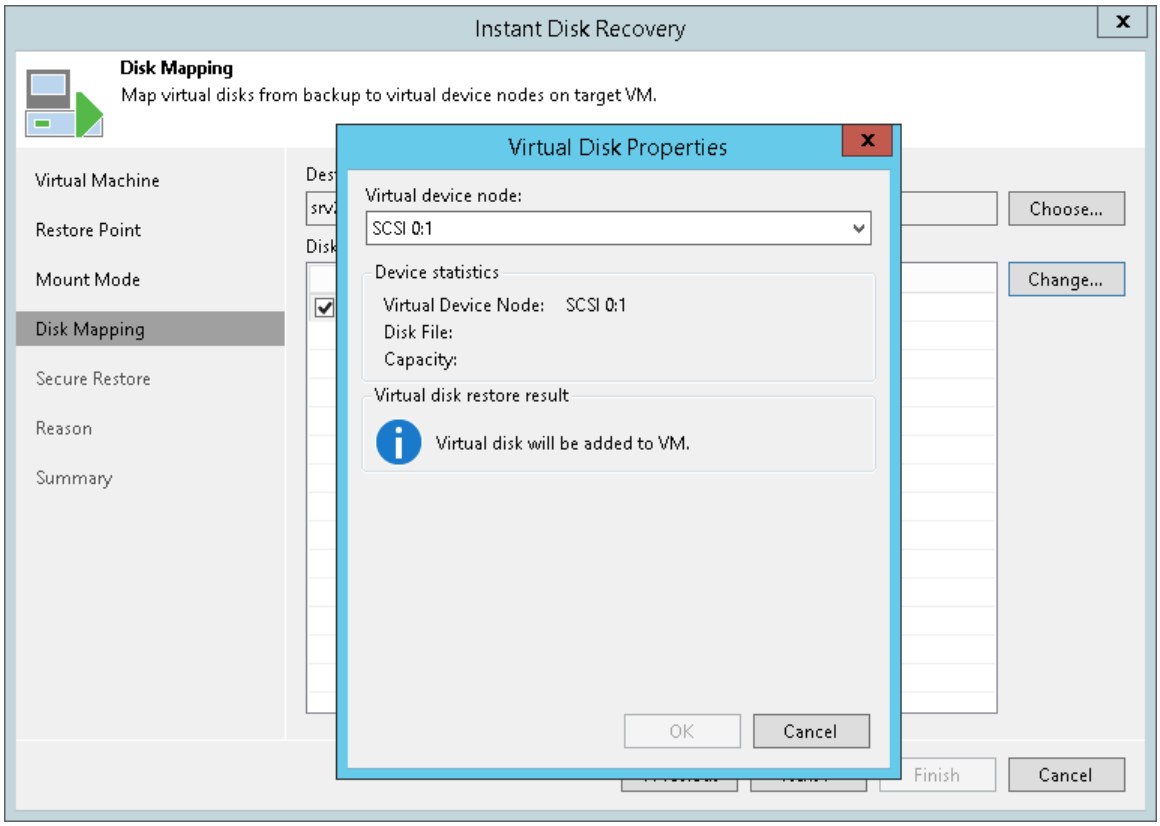
At the **Disk Mapping** step, select virtual disks that you want to restore and choose a VM to which the disks must be attached:

1. By default, Veeam Backup & Replication maps the recovered disks to the original VM. If the original VM was relocated or if you want to attach disks to another VM, select a target VM. For this, click **Choose** and select a VM from the virtual environment.

   > **IMPORTANT**
   >
   > Do not use the backup server or a VM where the vPower NFS Service is installed as a target VM.

2. In the **Disk mapping** section, select virtual disks that you want to restore.
3. To change a virtual device node where a virtual disk will be restored, select a disk in the **Disk mapping list** and click **Change**. In the **Virtual Disk Properties** window, select the node:
   - If you want to replace an existing virtual disk, select an occupied virtual node. The original disk will be deleted.
   - If you want to attach the restored disk to the VM as a new drive, select a node that is not occupied yet.

     > **NOTE**
     >
     > If you restore a virtual disk to an unoccupied node, Veeam Backup & Replication restores the disk in the offline state. To work with it, you need to bring the disk online.



Step 6. Specify Secure Restore Settings

This step is available if you restore disks of Microsoft Windows VMs.

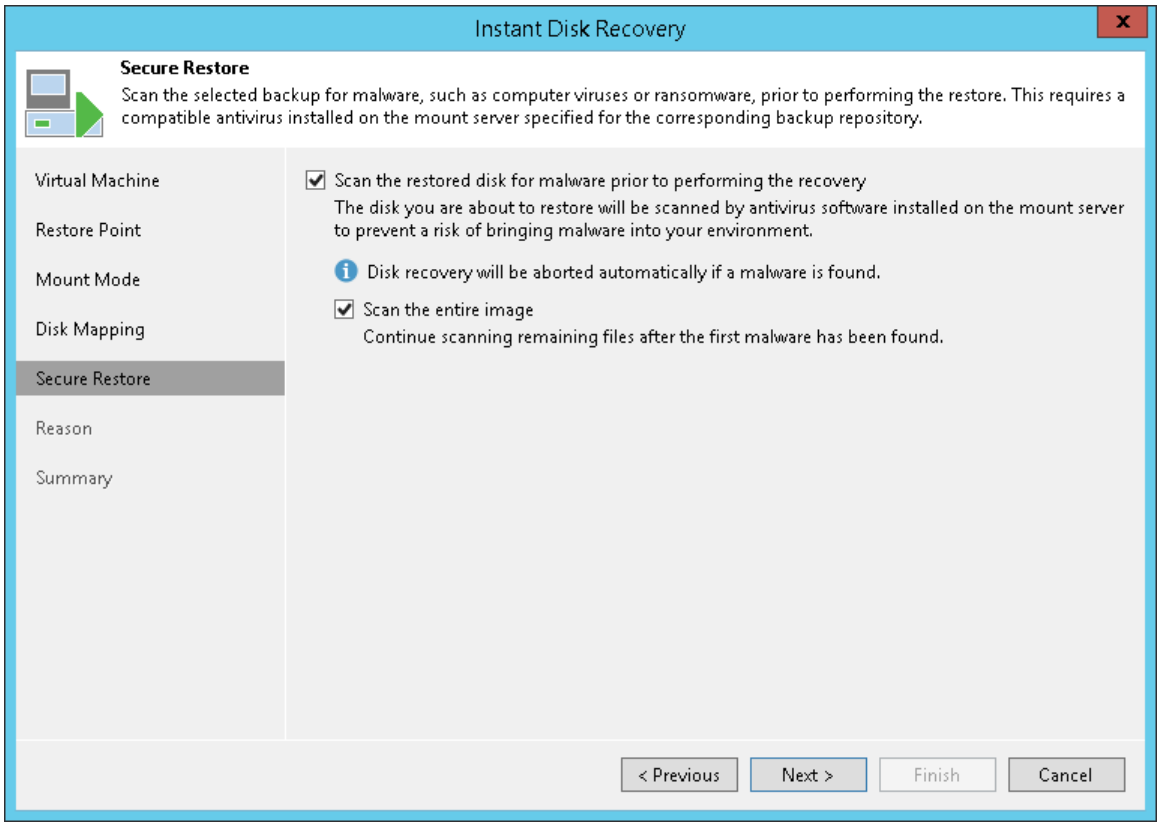
At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to perform secure restore — scan virtual disk data with antivirus software before restoring the disk. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

1. At the **Secure Restore** step of the wizard, select the **Scan the restored disk for malware prior to performing the recovery** check box.

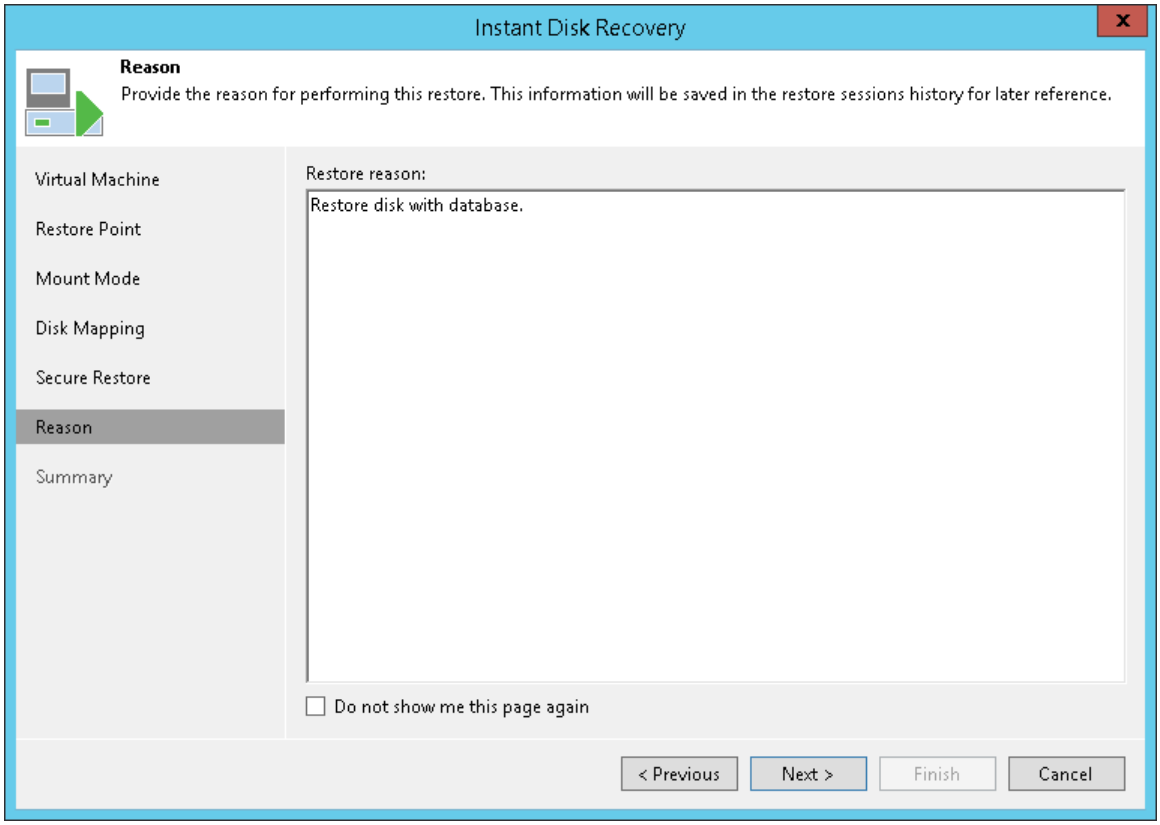
   If malware is found, Veeam Backup & Replication will cancel the disk restore process.

2. Select the **Scan the entire image** check box if you want the antivirus software to continue the virtual disk scan after a virus threat is detected. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



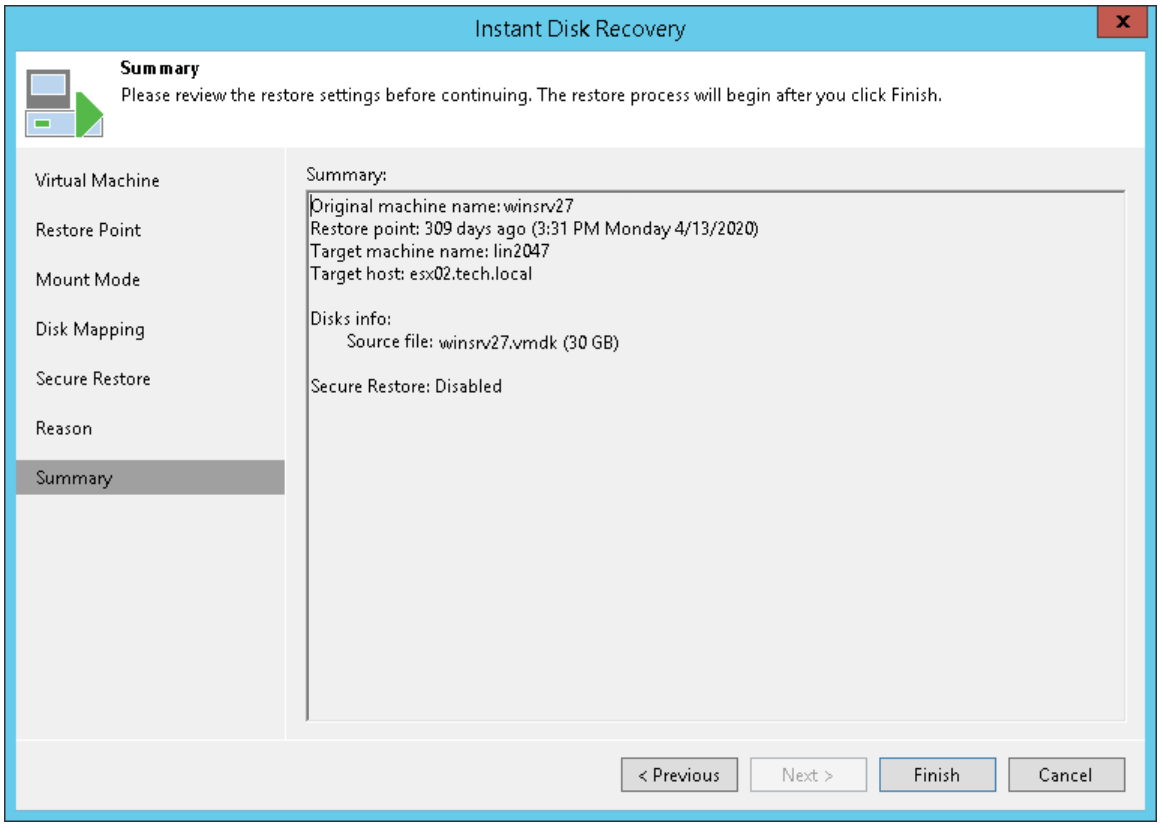
Step 7. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM disks. This information will be saved in the session history so that you can reference it later.



Step 8. Verify Recovery Settings

At the **Summary** step of the wizard, check settings of Instant Disk Recovery and click **Finish**.



Finalizing Instant Disk Recovery

After the disks have been successfully recovered, you must finalize the process. For this, test the recovered disks and decide whether to migrate them to production environment or stop publishing them.

Testing Recovered VM Disks

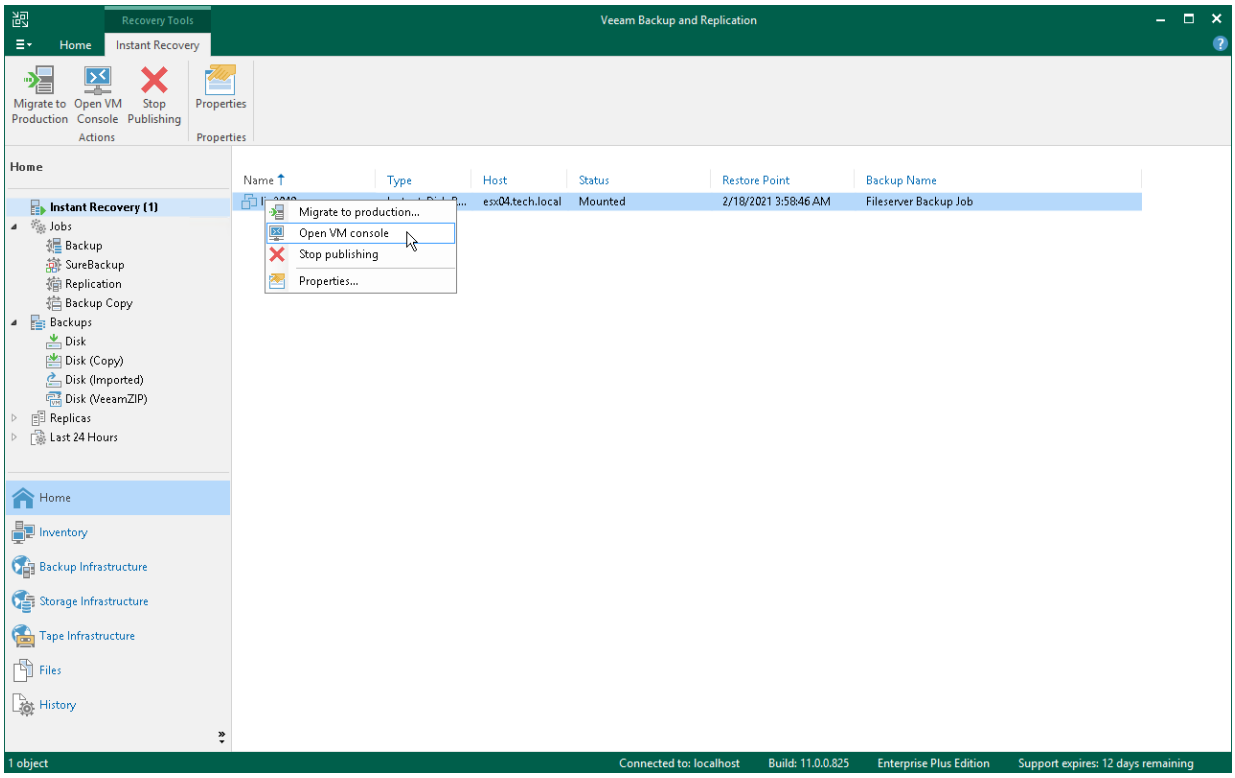
To test the recovered disks before you migrate them to production, you can launch from the Veeam Backup & Replication console the VMware Remote Console software of the VM to which the disks were recovered.

> **IMPORTANT**
>
> Before you launch VMware Remote Console, make sure that this software is installed on the machine where the Veeam Backup & Replication console runs.

To open a VM console in Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Open VM console**.



Migrating Recovered VM Disks

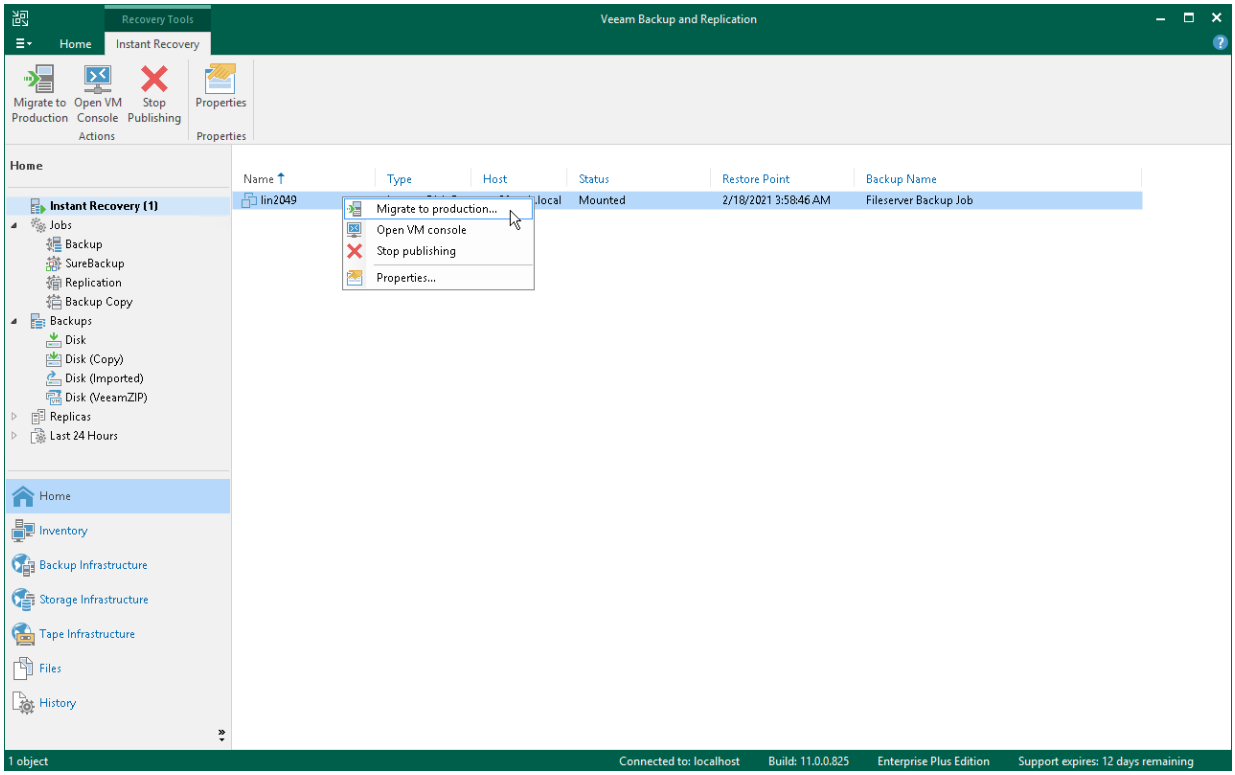
To migrate recovered disks to the production environment:

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 Quick Migration wizard.

   The wizard will be opened on the **Destination** step. At this step, you can change the datastore where the virtual disks will be placed. By default, Veeam Backup & Replication places the disks in the datastore where the VM configuration file is stored.

   Other fields (host, resource pool and VM folder) are populated with data of the VM where you restore the disks. You cannot change this data.

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.



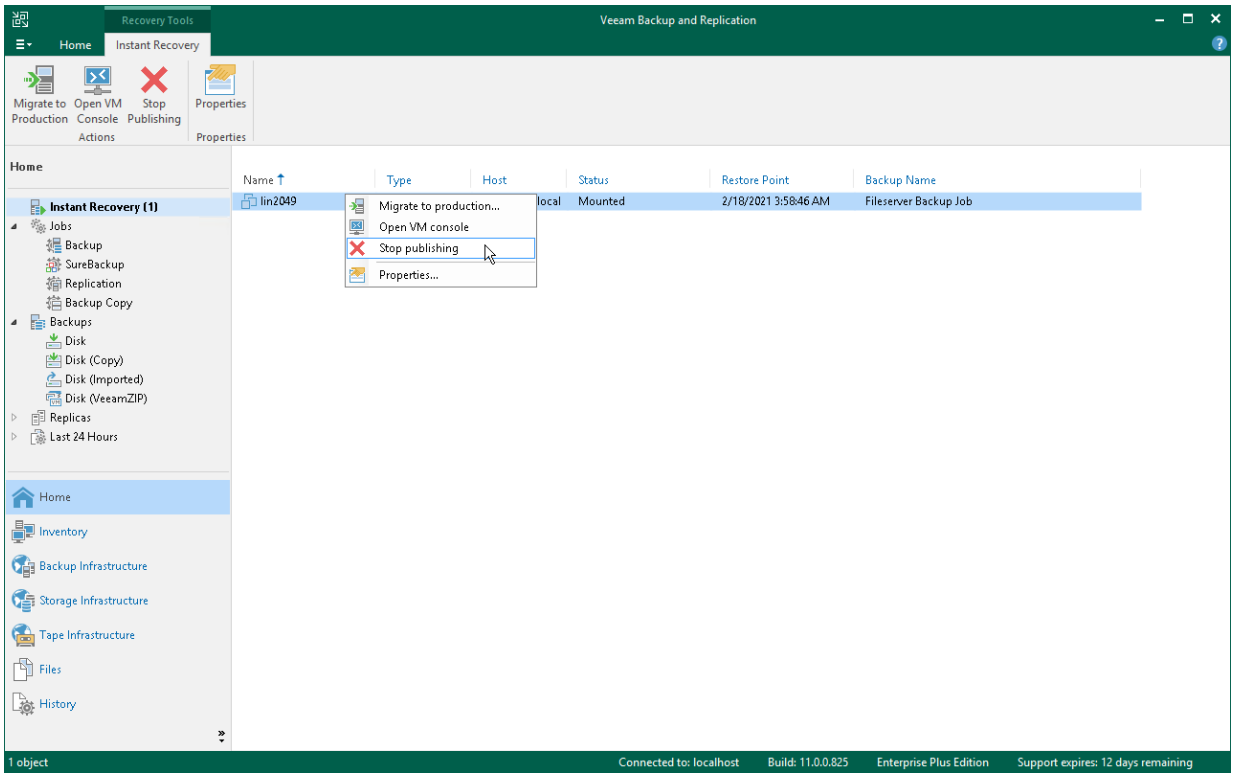
Stop Publishing Recovered VM Disks

If your tests have failed, you can stop publishing the recovered disks. This will remove the recovered disks from the VM that you selected as the destination for recovery. Note that all changes made to the recovered disks will be lost.

To remove the recovered disks:

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 the disks were recovered and select **Stop publishing**.



Instant First Class Disk (FCD) Recovery

With Instant First Class Disk (FCD) Recovery, you can immediately restore VM disks from a backup file and register them as FCDs. If you want to recover disks in their initial format, see Instant Disk Recovery. For the list of all disk recovery methods and their brief descriptions, see Disk Recovery.

The Instant FCD Recovery is a process that must be finalized. You must migrate FCDs to the production environment or stop publishing them. To migrate the disks, Veeam Backup & Replication uses the Quick Migration mechanism. For more information on how to finalize Instant FCD Recovery, see Finalizing Instant FCD Recovery.

How Instant FCD Recovery Works

Instant FCD Recovery workflow is divided into two parts: recovery of disks and publishing them as FCDs.

Recovering Disks

Veeam Backup & Replication performs the following steps to recover disks:

1. Initiates a recover session.
2. Connects to a target VMware cluster and mounts vPower NFS datastore to all ESXi hosts added to the target VMware cluster.
3. Mounts recovered disks and publishes them to the target vPower NFS datastore.

Publishing Recovered Disks as FCDs

Veeam Backup & Replication performs the following steps to publish the recovered disks as FCDs:

1. Registers the published disks as FCDs to the target vPower NFS datastore.
2. Initiates a creation of a protective snapshot for every FCD. If the Instant FCD Recovery process fails, the protective snapshot guarantees no data loss.

After the snapshots are created, Veeam Backup & Replication unmounts the recovered disks from the vPower NFS datastore.

Veeam Backup & Replication performs Instant FCD Recovery using vPower technology.

> **NOTE**
>
> Mind the following:
>
> - If you specify the custom datastore to keep the redo logs, information on the FCD snapshots will be stored in the instant recovery session logs.
> - Veeam Backup & Replication does not remove the mounted vPower NFS datastore from a VMware cluster.

Performing Instant FCD Recovery

To register VM disks as FDCs, use the **Instant Disk Recovery** wizard.

Before You Begin

Before you perform Instant FCD Recovery, consider the following:

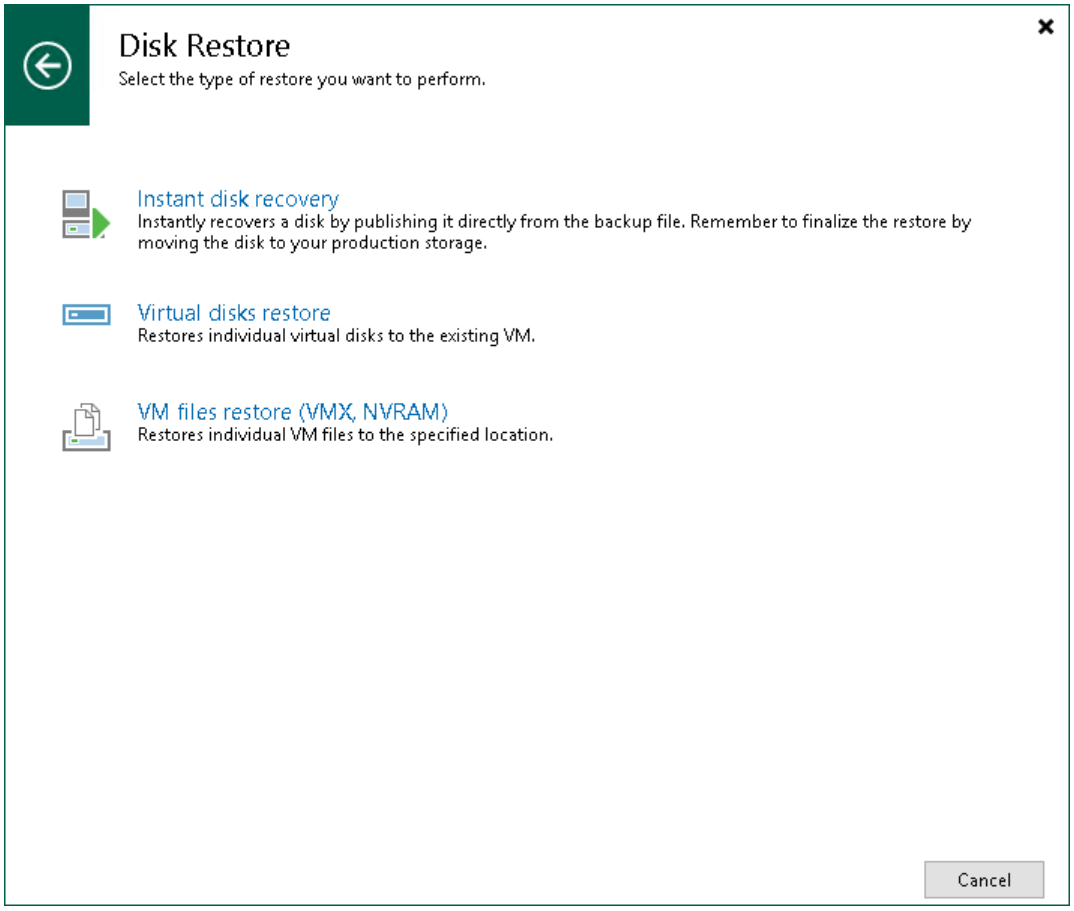
- The vCenter Server where the target cluster (cluster to which you plan to recover disks) is located must be version 6.7.U3 or higher.
- The target cluster must be configured beforehand and added to the backup infrastructure. The hosts located in this cluster must meet the following requirements:
  - ESXi hosts must be powered on.
  - Veeam Backup & Replication must be able to connect to ESXi hosts.
  - API version of ESXi hosts must be 6.7.U3 or higher.
  - ESXi must have the VMkernel network interface enabled.
- Veeam Backup & Replication does not support restore to a single ESXi host.
- You can restore a virtual disk from a backup that has at least one successfully created restore point.
- Instant FCD Recovery to vCloud Director is not supported.
- If you attach a FCD to a VM after restore and attach the FCD to an unoccupied node, Veeam Backup & Replication restores the disk in the offline state. To work with it, you need to bring the disk online.
- The datastore which you specify to store redirect logs must meet the following requirements:
  - The datastore must be added to the vCenter infrastructure and must be available for Veeam Backup & Replication to set up a connection to it.
  - The datastore must be located in the same vCenter, where a cluster to which Veeam Backup & Replication will register FCDs, is located.
  - The datastore must be mounted to all ESXi hosts of a cluster to which Veeam Backup & Replication will register FCDs.
  - You must have at least one more additional datastore in your vCenter infrastructure, otherwise restore will fail.
  - The datastore must not be the Veeam NFS datastore.
- You can assign storage policies to FCDs only if you redirect redo logs to another datastore.
- The encryption storage policies are not supported. If you select this type of policy, Veeam Backup & Replication will not be able to apply this policy to registered FCDs.
- You must have at least 10 GB of free disk space on the datastore where write cache folder is located. This disk space is required to store virtual disk updates for the restored VM.

  By default, Veeam Backup & Replication writes virtual disk updates to the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`.

Step 1. Launch Instant Disk Recovery Wizard

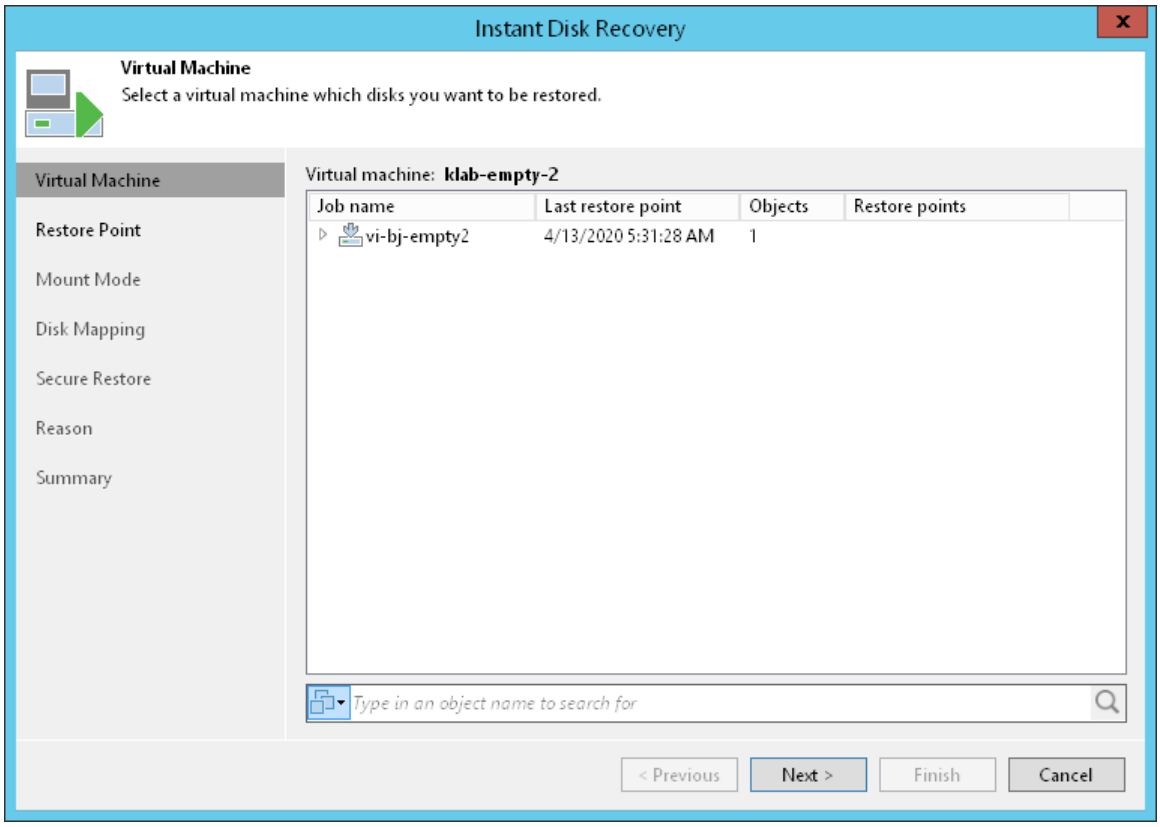
To launch the **Instant Disk Recovery** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **Instant disk recovery**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand a backup, select a VM whose disks you want to restore and click **Instant Disk Recovery** on the ribbon. Alternatively, you can right-click a VM whose disks you want to restore and select **Instant disk recovery**.



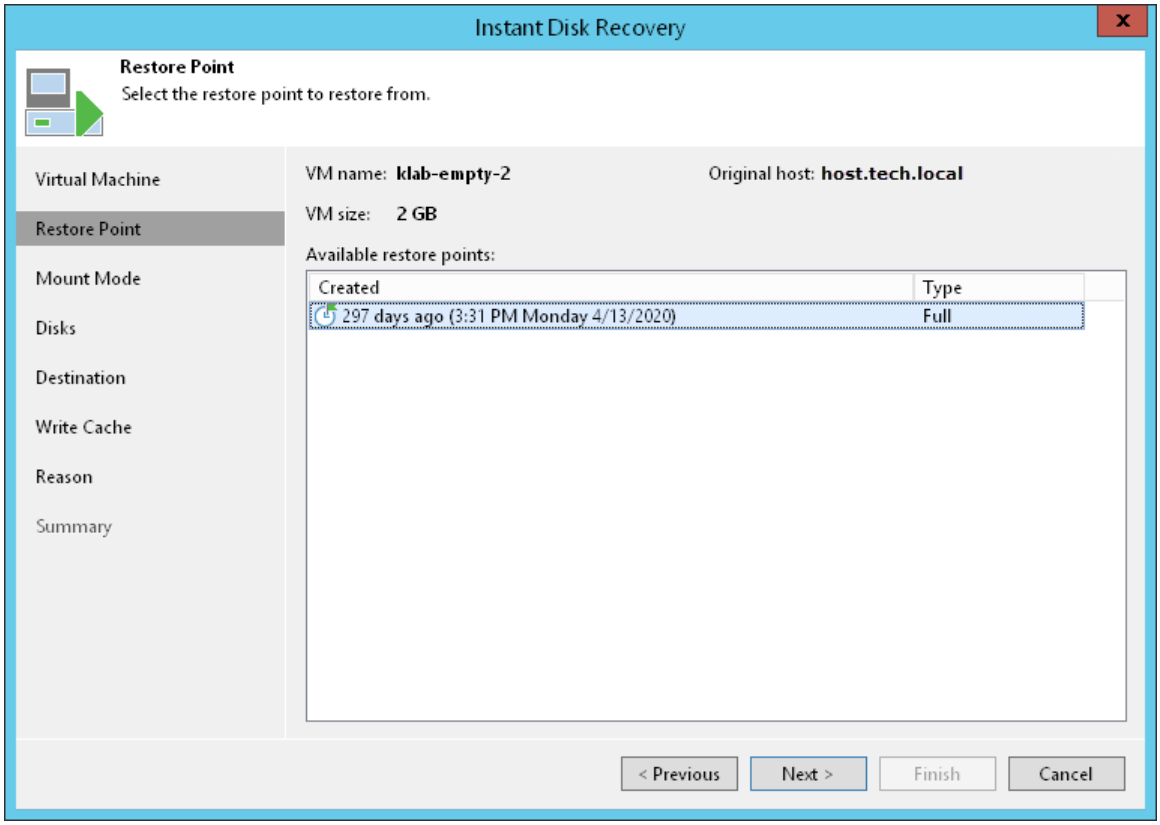
Step 2. Select Source VM

At the **Virtual Machine** step of the wizard, select a VM whose disks you want to register as FCDs.



Step 3. Select Restore Point

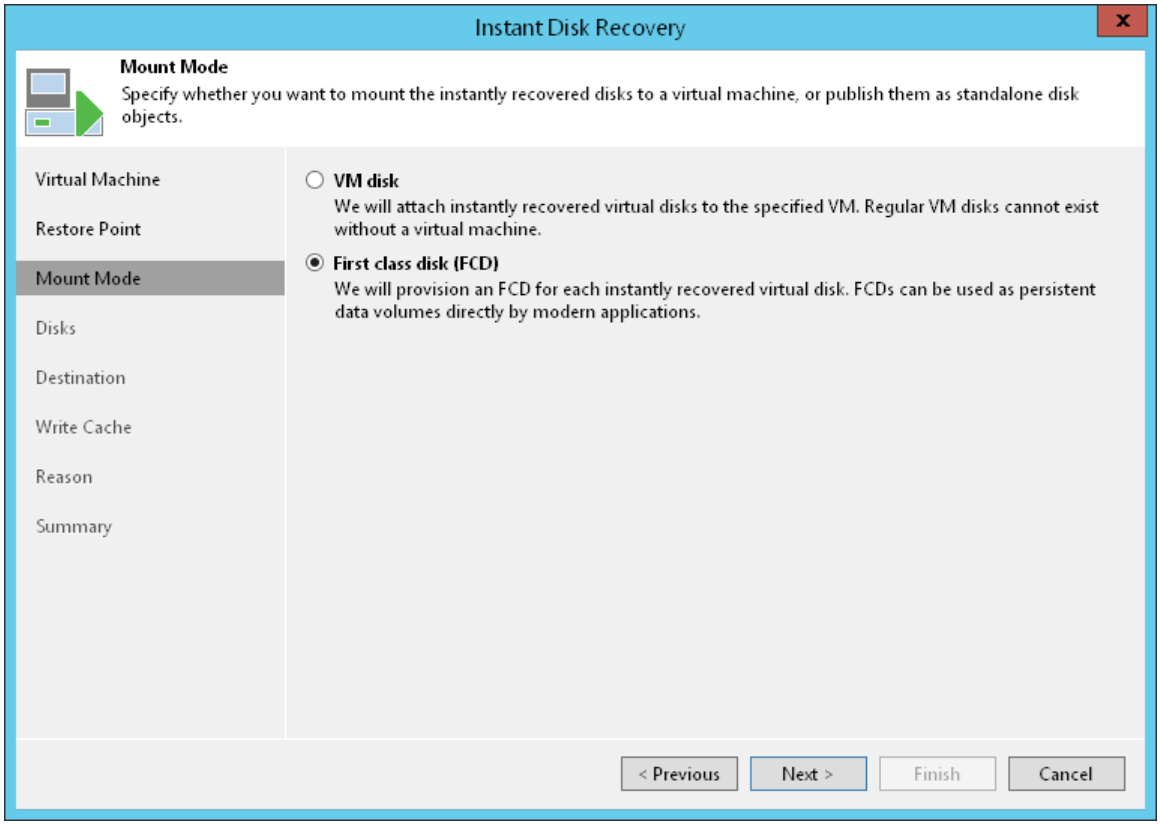
At the **Restore Point** step of the wizard, select a restore point for the VM whose virtual disks you want to register as FCDs.



Step 4. Select Mount Mode

At the **Mount Mode** step of the wizard, select the **First class disk (FCD)** option to register virtual disks on a cluster as FCDs.

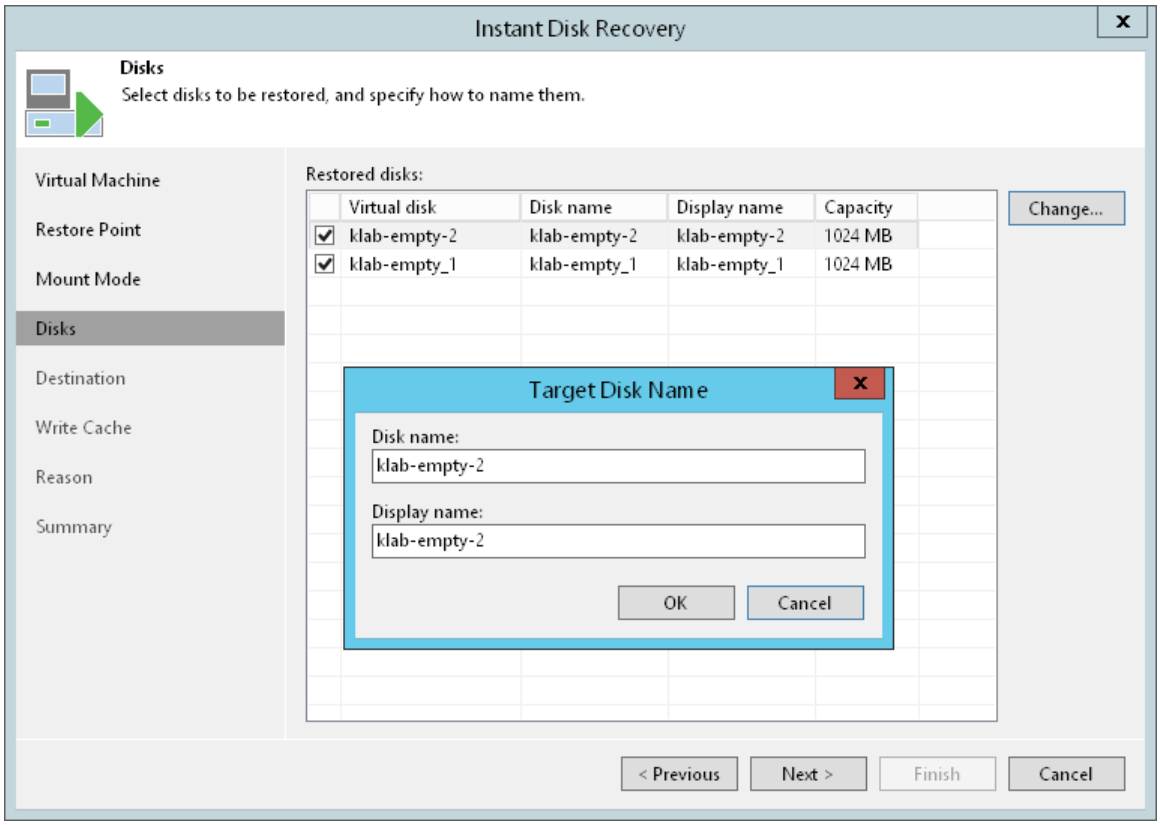
If you want to register virtual disks on a VM added to an ESXi host, select the **VM disk** option. In this case, steps of the wizard differ and Veeam Backup & Replication performs the instant disks recovery as described in Performing Instant Disk Recovery.



Step 5. Select Virtual Disks to Restore

At the **Disks** step of the wizard, select virtual disks that you register as FCDs and define names for the recovered disks and FCDs.

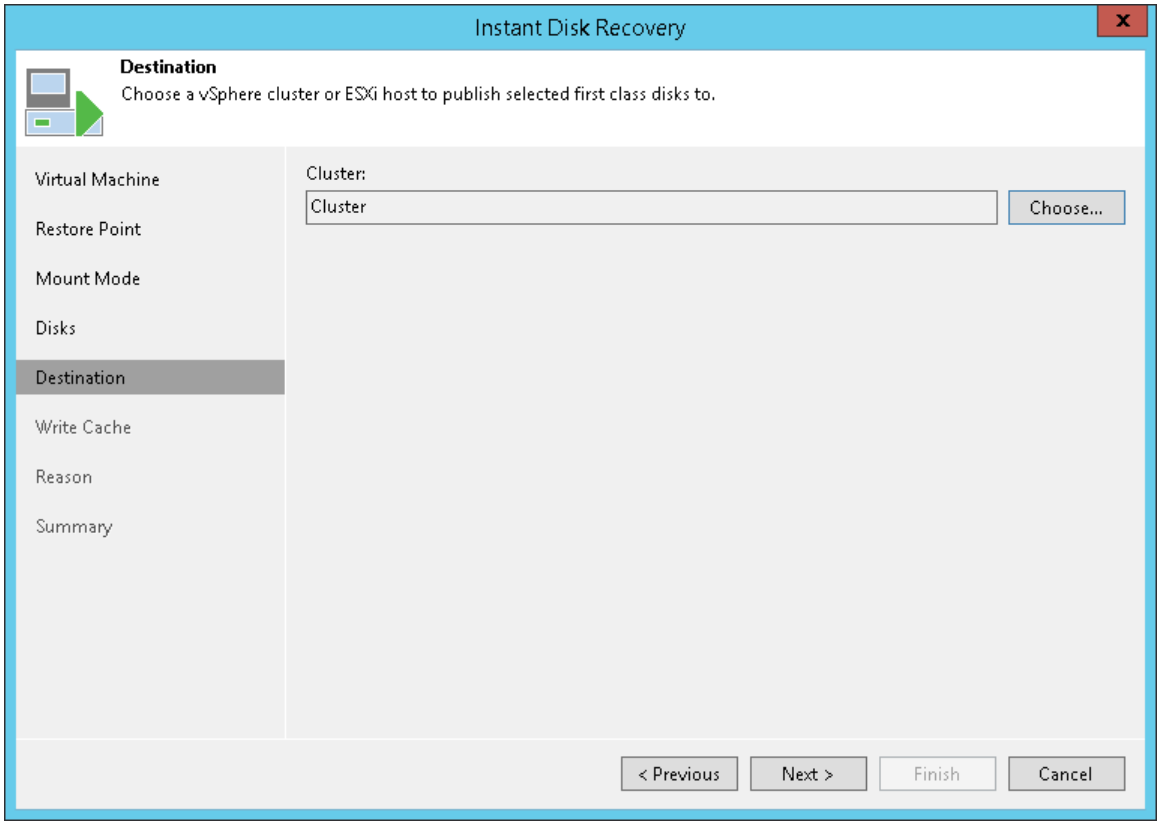
1. Select check boxes next to virtual disks that you want to register as FCDs.
2. By default, Veeam Backup & Replication recovers disks and registers FCDs under the original disk names. To specify new names, select a disk and click **Change**. In the **Target Disk Name** window, do the following:
   a. In the **Disk name** field, specify a name under which the disk will be recovered.
   b. In the **Display name** field, specify a name under which FCD will be registered on the cluster.



Step 6. Specify Target Cluster

At the **Destination** step of the wizard, specify a cluster where Veeam Backup & Replication will register FCDs.

> **IMPORTANT**
>
> To register FCDs on a cluster, Veeam Backup & Replication mounts to a cluster the vPower NFS datastore that contains virtual disks. Veeam Backup & Replication checks the cluster hosts to verify that they are powered on and Veeam Backup & Replication is able to connect to these hosts. If one of these hosts is not available, Veeam Backup & Replication will not include them to the FCD restore scenario.
> Veeam Backup & Replication will not start the Instant FCD Recovery process if there are no powered on hosts or Veeam Backup & Replication is not able to connect to them.



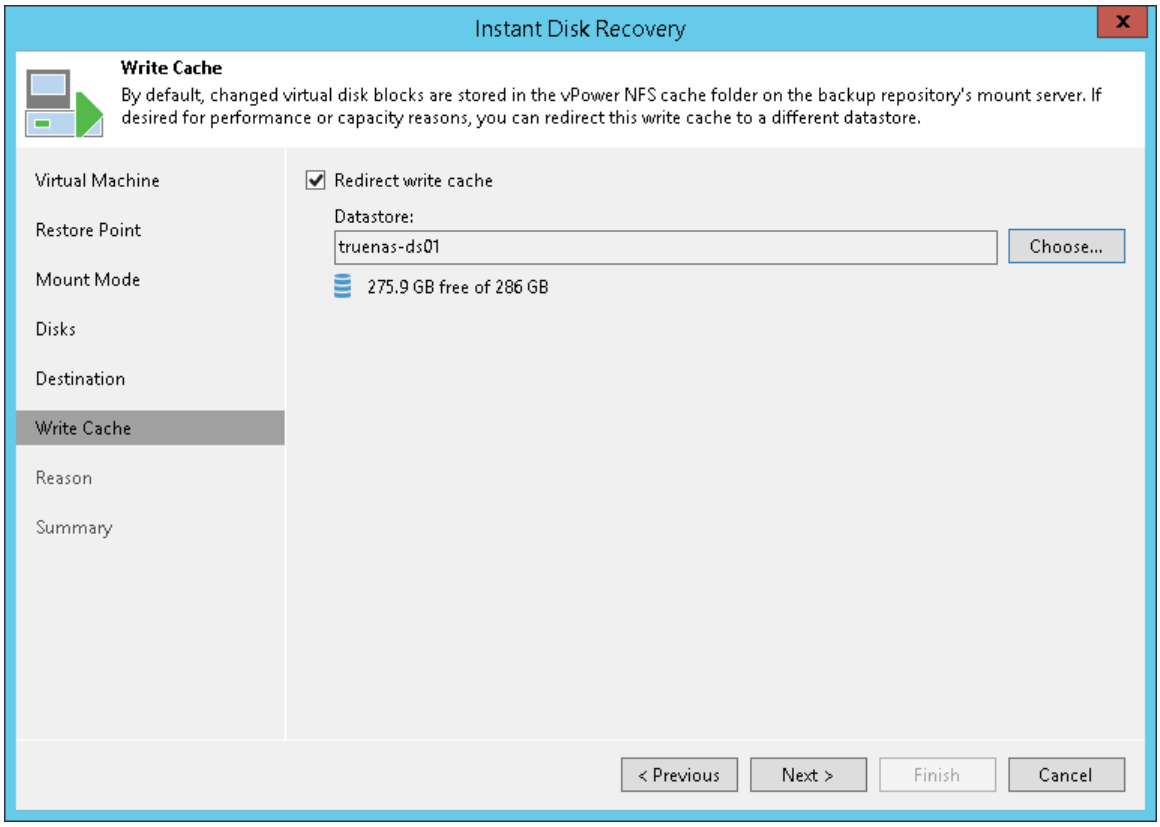
Step 7. Select Destination for FCD Updates

At the **Write Cache** step of the wizard, specify where to store the redo logs. Veeam Backup & Replication uses redo logs to write all changes of the virtual disks that take place while performing Instant FCD Recovery.

By default, redo logs are stored on the vPower NFS server. You can store redo logs on any datastore in the virtual environment. As soon as a recovery verification job completes, Veeam Backup & Replication deletes redo logs.

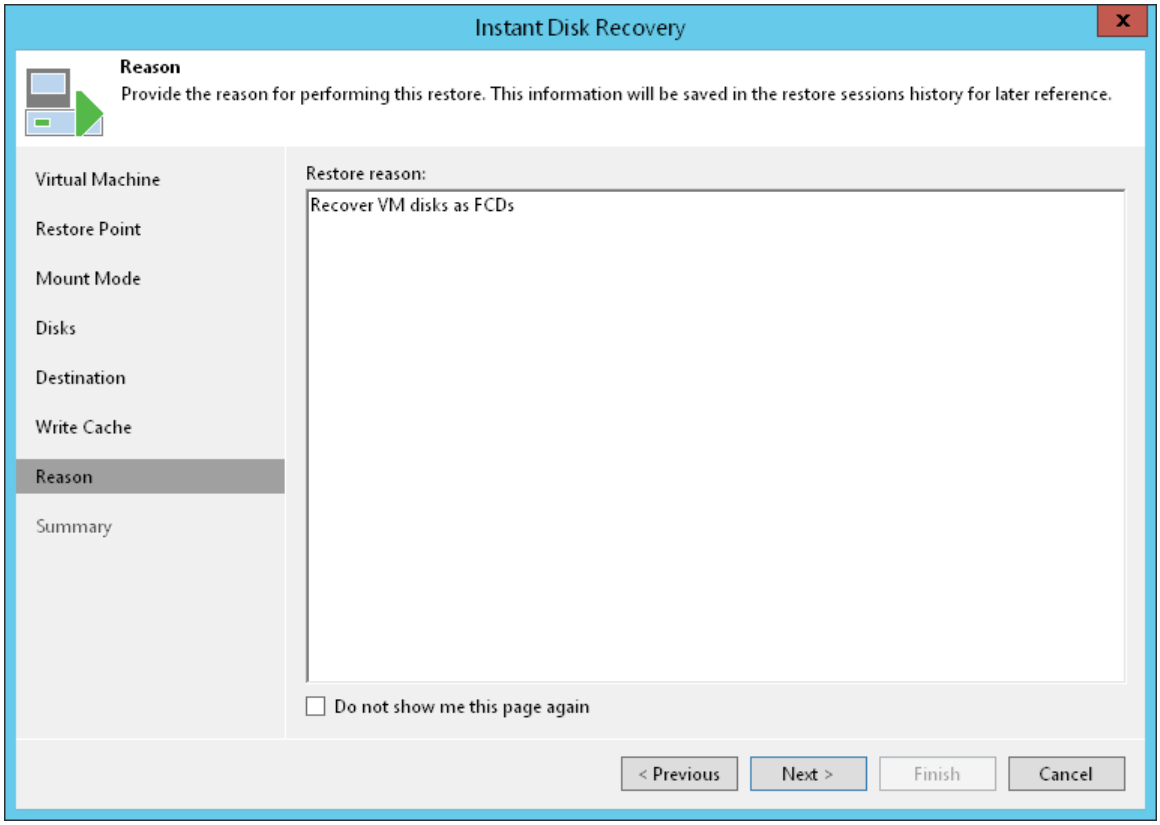
To redirect redo logs:

1. Select the **Redirect write cache** check box.
2. Click **Choose** and select a datastore from the list.



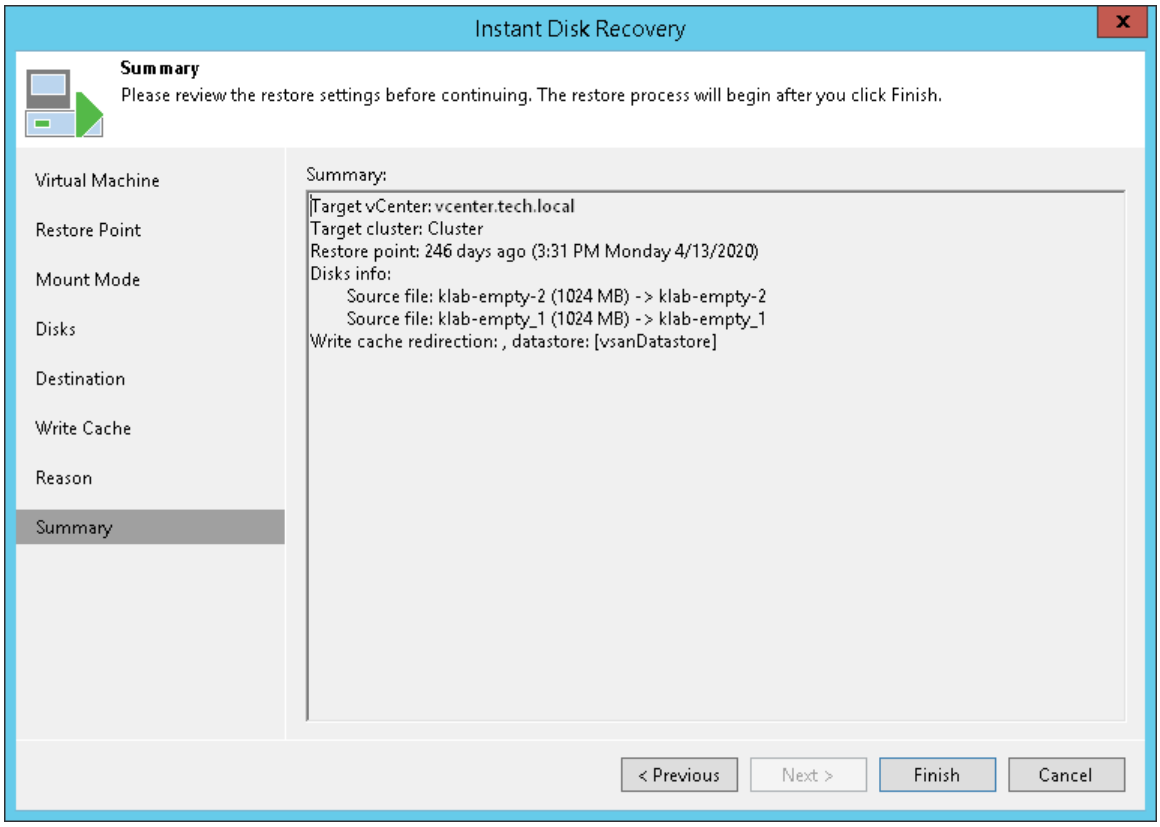
Step 8. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM disks. This information will be saved in the session history so that you can reference it later.



Step 9. Verify Instant FCD Recovery Settings

At the **Summary** step of the wizard, check settings of Instant FCD Recovery and click **Finish**.



Finalizing Instant FCD Recovery

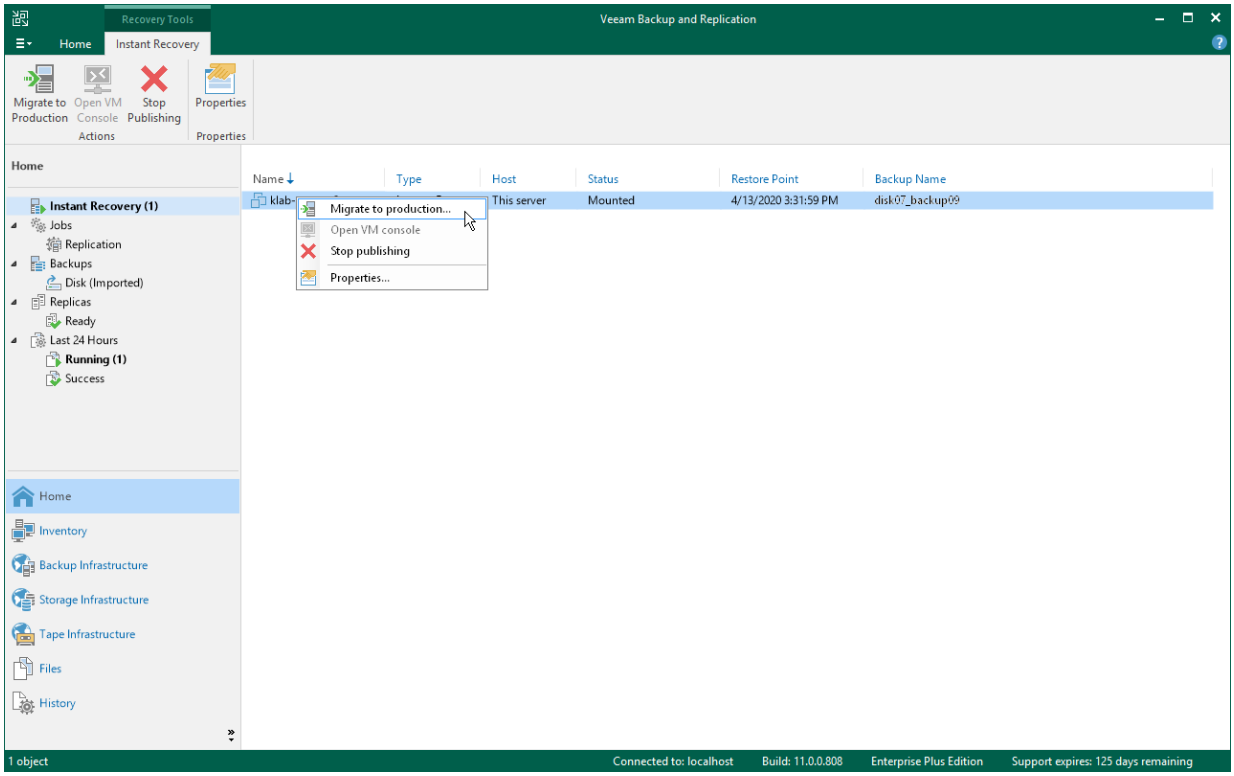
After FCDs have been successfully registered, you must finalize the process and decide whether to migrate the disks to the production environment or stop publishing them.

Migrating FCDs

To migrate FCDs to the production environment:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM whose disks you 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.



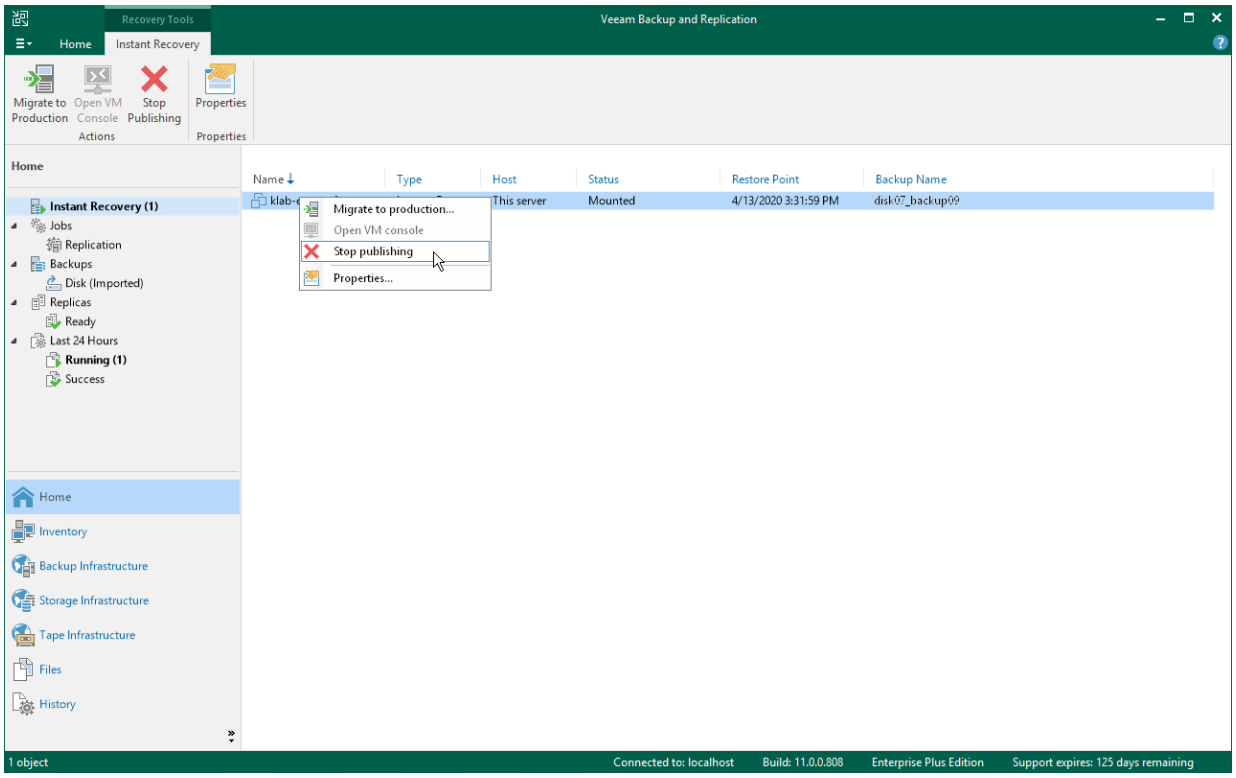
Stop Publishing FCDs

You can stop publishing FCDs. This will remove the registered FCDs from the cluster that you selected as the destination for recovery. Note that all changes made to the FCDs will be lost.

To stop publishing FCDs:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM whose disks you recovered and select **Stop publishing**.



Virtual Disk Restore

Veeam Backup & Replication allows you to recover individual virtual disks of a VM. Recovered virtual disks can be attached to the original VM (for example, if you want to replace a corrupted disk) or mapped to any other VM in the virtual infrastructure. This recovery option can be helpful if a VM disk becomes corrupted.

You can restore VM virtual disk to the latest state or any valid restore point. You can preserve the format of a recovered virtual disk or convert it to the thin or thick provisioned disk format.

> **NOTE**
>
> If a VM has several VM disks, Veeam Backup & Replication restores VM disks in parallel.

Restoring Virtual Disks

To restore virtual disks, use the **Virtual Disk Restore** wizard.

Before You Begin

Before you restore virtual disks, check the following prerequisites:

- You can restore virtual disks from a backup that has at least one successfully created restore point.
- During the virtual disk restore, Veeam Backup & Replication turns off the target VM to reconfigure its settings and attach restored disks. It is recommended that you stop all activities on the target VM for the restore period.
- If you want to scan virtual disk data for viruses, check the secure restore requirements and limitations.

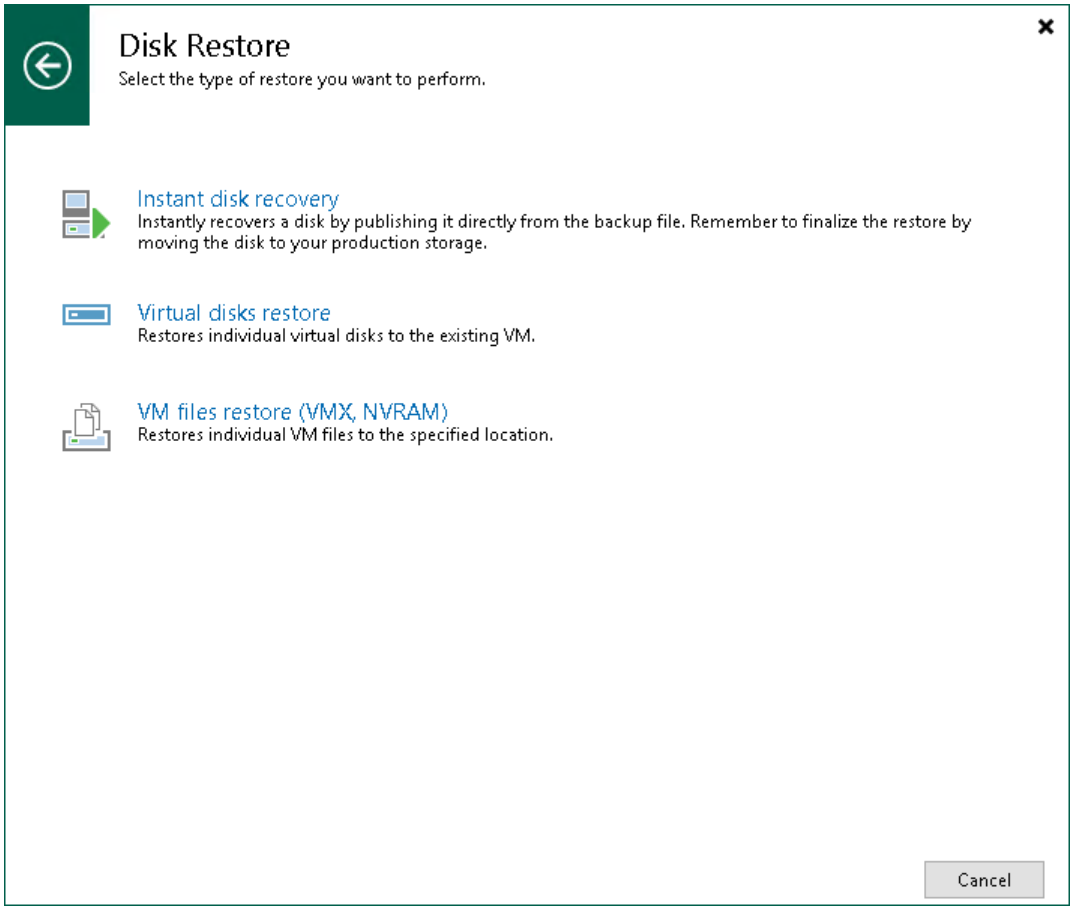
> **NOTE**
>
> If you back up a VM with vRDM disks, Veeam Backup & Replication converts them into VMDK files. Thus, when you restore a vRDM disk, Veeam Backup & Replication will restore it as a VMDK file. If you want to preserve the vRDM format for restored disks, use Quick Rollback. For more information, see Quick Rollback.

Step 1. Launch Virtual Disk Restore Wizard

To launch the **Virtual Disk Restore** wizard, do one of the following:

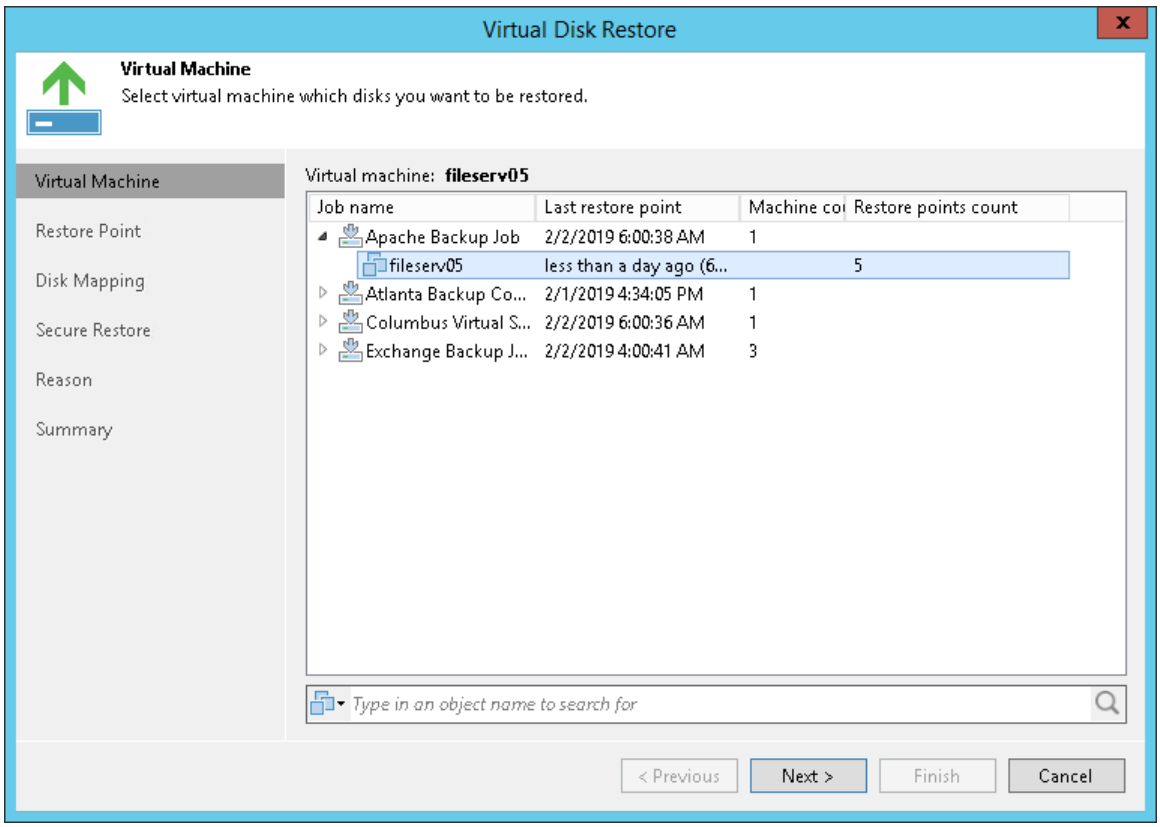
- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **Virtual disks restore**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **Virtual Disks** on the ribbon. Alternatively, you can right-click the VM whose files you want to restore and select **Restore virtual disks**.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Virtual disks**.

  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.



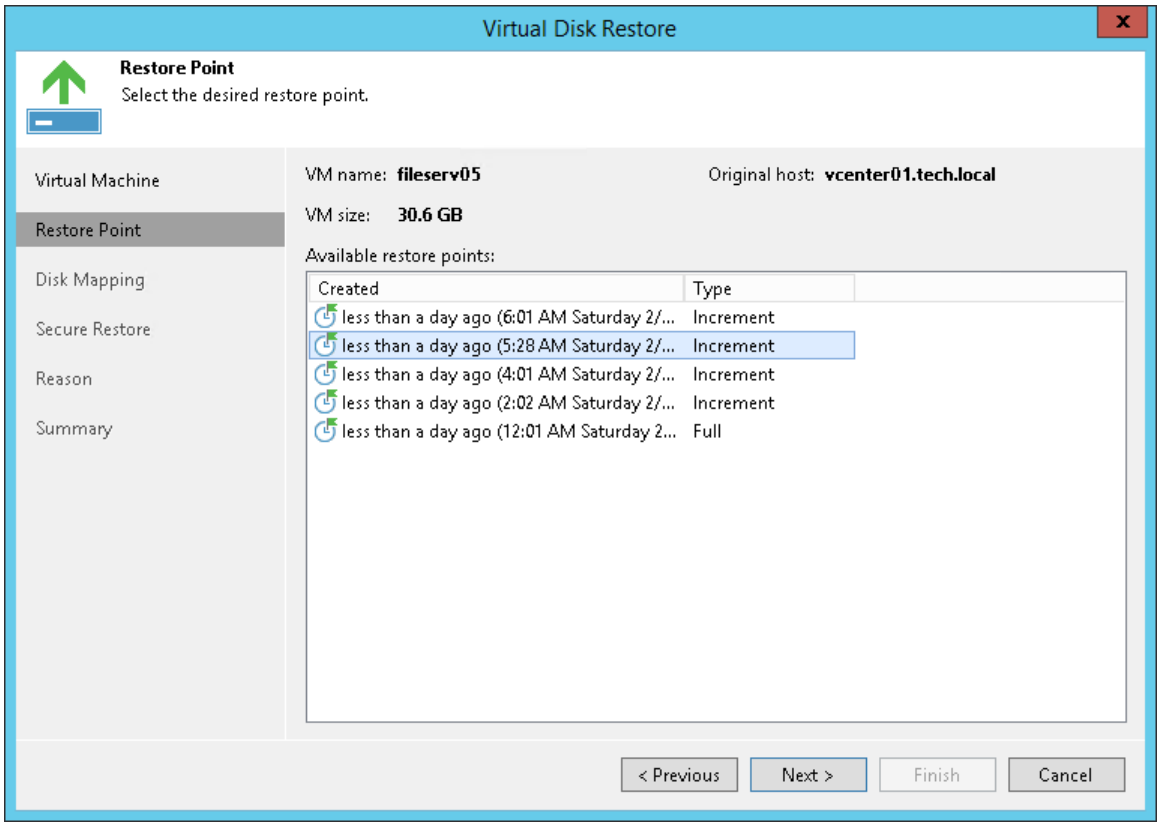
Step 2. Select VM

At the **Virtual Machine** step of the wizard, expand a backup and select a VM whose disks you want to restore.



Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to restore VM disks.



Step 4. Select Virtual Hard Disks to Restore

At the **Disk Mapping** step, select virtual hard disks to restore, choose a VM to which the disks must be attached and define additional restore settings.

1. By default, Veeam Backup & Replication maps restored disks to the original VM. If the original VM was relocated or if you want to attach disks to another VM, you need to select the target VM manually. Click **Choose** and select the necessary VM from the virtual environment.

   > **IMPORTANT**
   >
   > You cannot attach restored disks to a VM that has one or more snapshots.

2. In the **Disk mapping** list, select virtual disks that you want to restore.
3. To configure virtual disk properties, select a disk in the **Disk mapping** list and click **Change**. In the **Virtual Disk Properties** window, do the following:

   a. Click **Choose** to pick a datastore where the restored hard disk must be placed.

      If you use storage policies in the virtual environment, Veeam Backup & Replication displays information about storage policies in the **Select Datastore** window. You can select a datastore associated with the necessary storage policy.

   b. From the **Virtual device node** drop-down list, select a virtual device node. If you want to replace an existing virtual disk, select an occupied virtual node. If you want to attach the restored disk to the VM as a new disk, select a node that is not occupied yet.

   c. Click **OK**.

4. Veeam Backup & Replication preserves the format of the restored virtual hard disks. To change disk format, select the required option from the **Restore disk type** drop-down list. For more information about disk types, see VMware Docs.

   > **NOTE**
   >
   > Disk format change is supported only for VMs with Virtual Hardware version 7 or later.

5. [For disk restore to the original location and with original format] Select the **Quick rollback** check box if you want to use incremental restore for the VM disk. For more information on quick rollback, its requirements and limitations, see Quick Rollback.

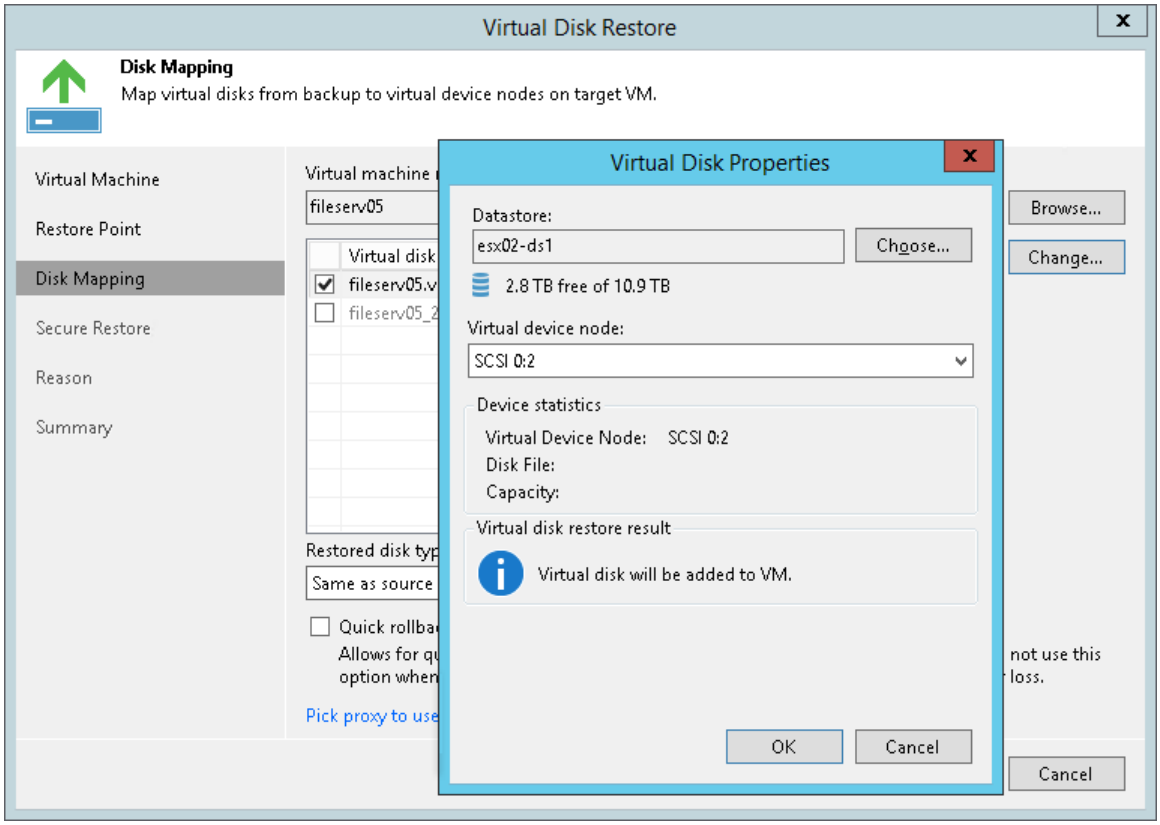
   It is recommended that you enable this option if you restore a VM disk after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.

6. Click the **Pick proxy to use** link to select backup proxies over which VM data must be transported to the target datastore. By default, Veeam Backup & Replication assigns proxies automatically.

   - 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.

     During the restore process, VM hard disks are processed simultaneously. 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 writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VM hard disk processing.

- If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that must be used for restore. It is recommended that you select at least two backup proxies to ensure that VM hard disks are recovered if one of backup proxies fails or loses its connectivity to the target datastore during restore.



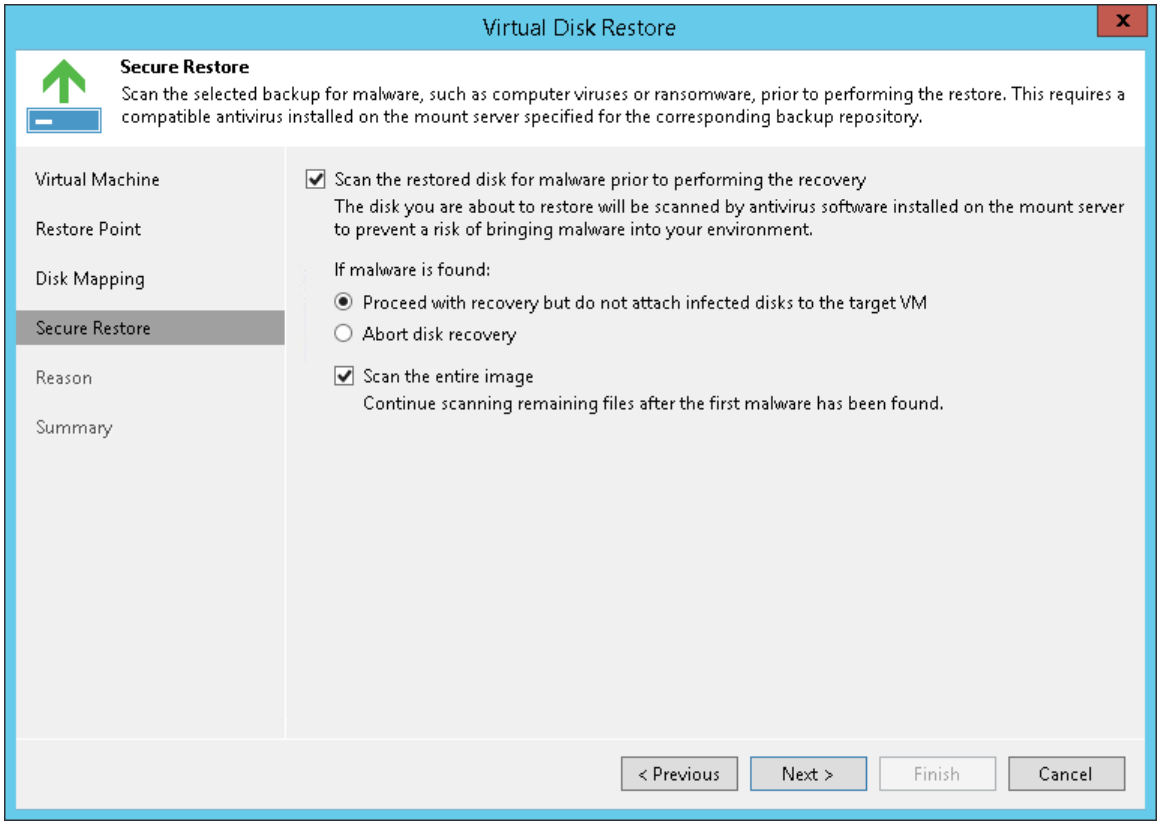
Step 5. Specify Secure Restore Settings

This step is available if you restore disks of Microsoft Windows VMs.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to perform secure restore — scan virtual disk data with antivirus software before restoring the disk. For more information on secure restore, see Secure Restore.

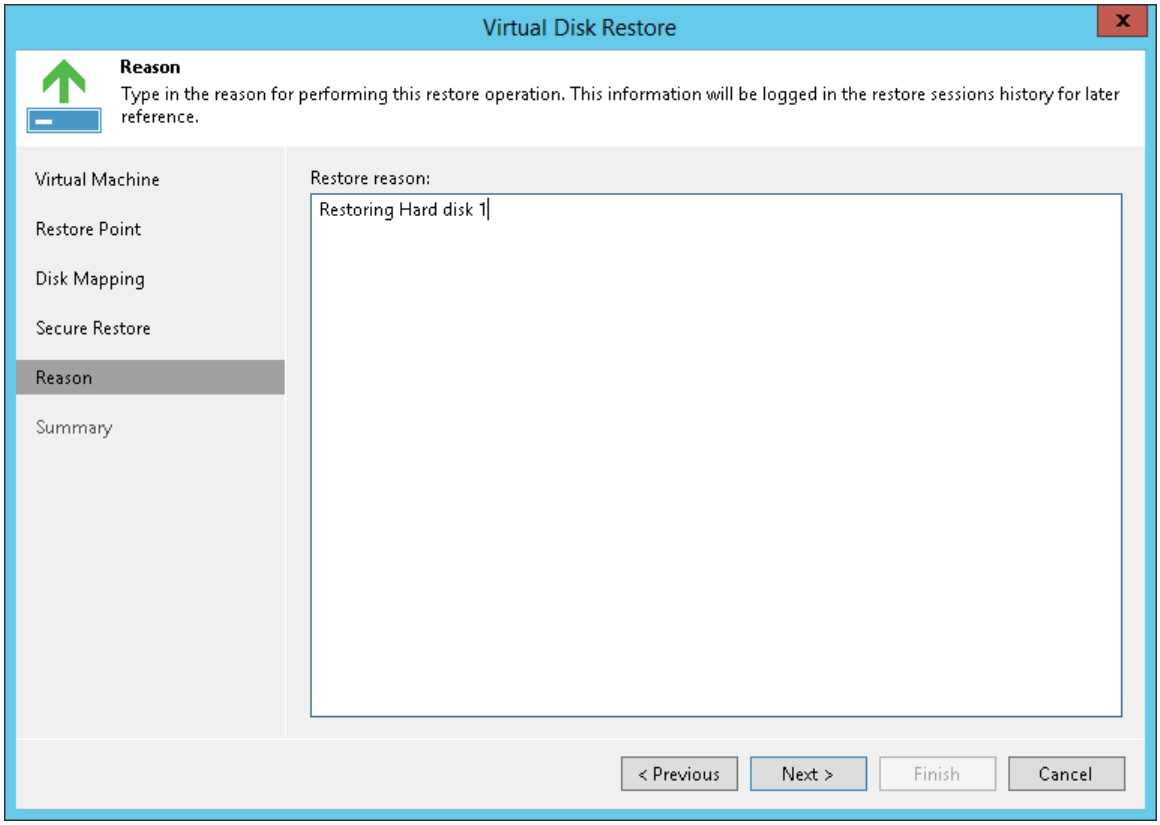
To specify secure restore settings:

1. Select the **Scan the restored disk for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software detects a virus threat:
   - **Proceed with recovery but do not attach infected disks to the target VM**. Select this action if you want to continue the virtual disk restore. In this case, the restored disk will not be attached to the target VM.
   - **Abort disk recovery**. Select this action if you want to cancel the restore session.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue the virtual disk scan after the first virus threat is detected. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



Step 6. Specify Restore Reason

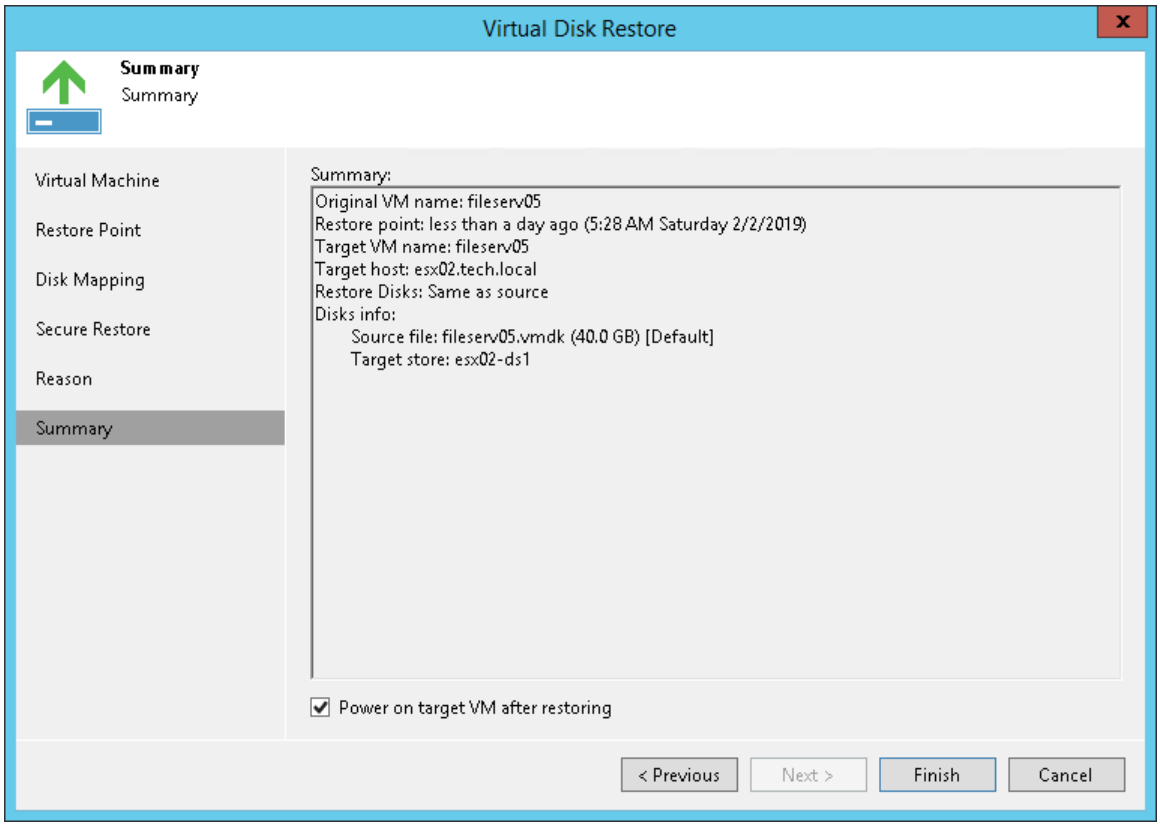
At the **Reason** step of the wizard, enter a reason for restoring VM disks. The information you provide will be saved in the session history and you can reference it later.



Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.

If you want to start the VM to which the disk will be attached right after restore, select the **Power on target VM after restoring** check box.



Disk Export

Veeam Backup & Replication allows you to export disks, that is, restore disks from backups of physical or virtual machines and convert them to the VMDK, VHD or VHDX formats.

During disk export, Veeam Backup & Replication creates the following files that can be used by VMware vSphere and Microsoft Hyper-V VMs:

- When you export a disk in the VMDK format, Veeam Backup & Replication creates a pair of files that make up the VM virtual disk: a descriptor file and file with the virtual disk content.
- When you export a disk in the VHD/VHDX format, Veeam Backup & Replication creates a file of the VHD or VHDX format.

You can save the exported disks to any server added to the backup infrastructure or place disks on a datastore connected to an ESXi host (for VMDK disk format only). VMDK disks can be exported as thin provision and thick disks. VHD/VHDX disks are always exported as dynamically expanding.

Veeam Backup & Replication supports batch disk export. For example, if you choose to export 2 disks, Veeam Backup & Replication will convert them to 2 virtual disks and store these disks in the specified location.

Supported Backup Types

You can restore disks from the following backups:

- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux or Veeam Agent for Mac
- Backups of EC2 instances created by Veeam Backup for AWS
- Backups of Azure VMs created by Veeam Backup for Microsoft Azure
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of Google VM instances created by Veeam Backup for Google Cloud Platform
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups of RHV VMs created by Veeam Backup for RHV
- [Starting from Veeam Backup & Replication cumulative patch P20211211 for v11a] Backups exported by K10 policies

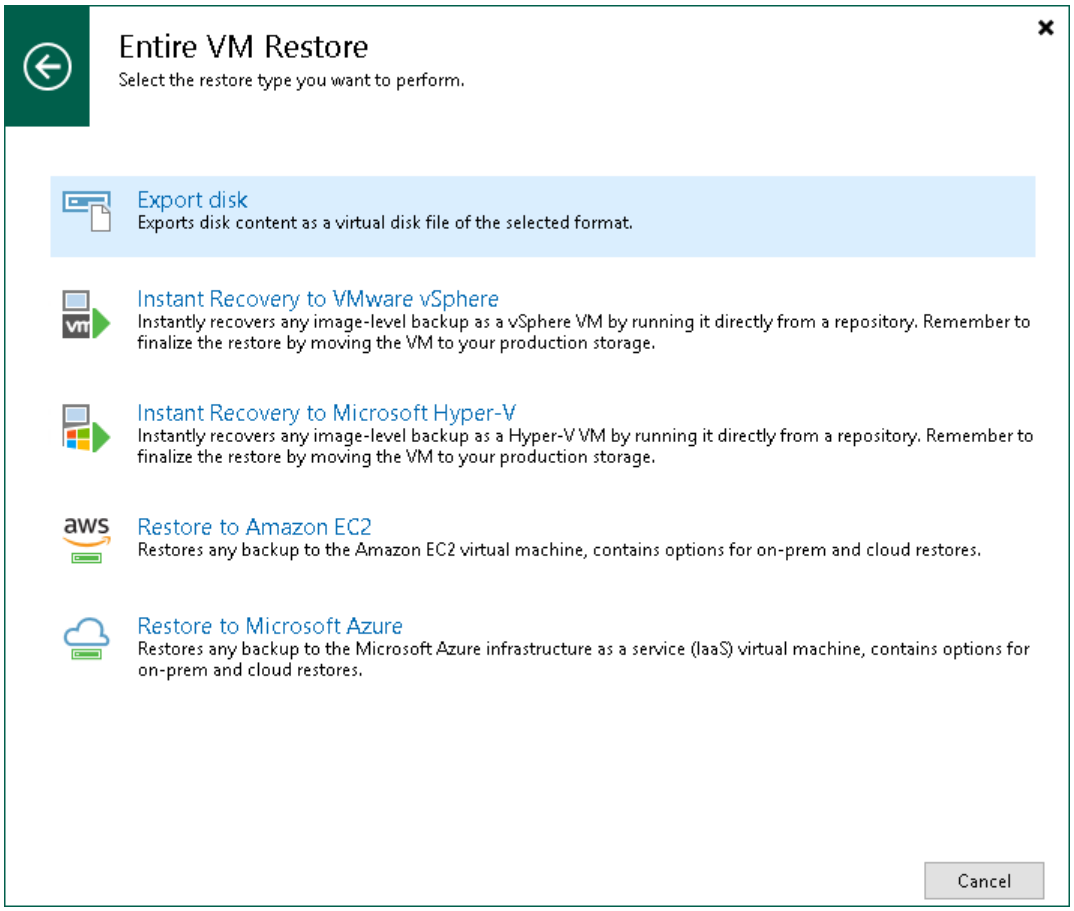
Exporting Disks

To export disks of physical or virtual machines and convert the disks to the VMDK, VHD or VHDX format, use the **Export Disk** wizard.

Step 1. Launch Export Disk Wizard

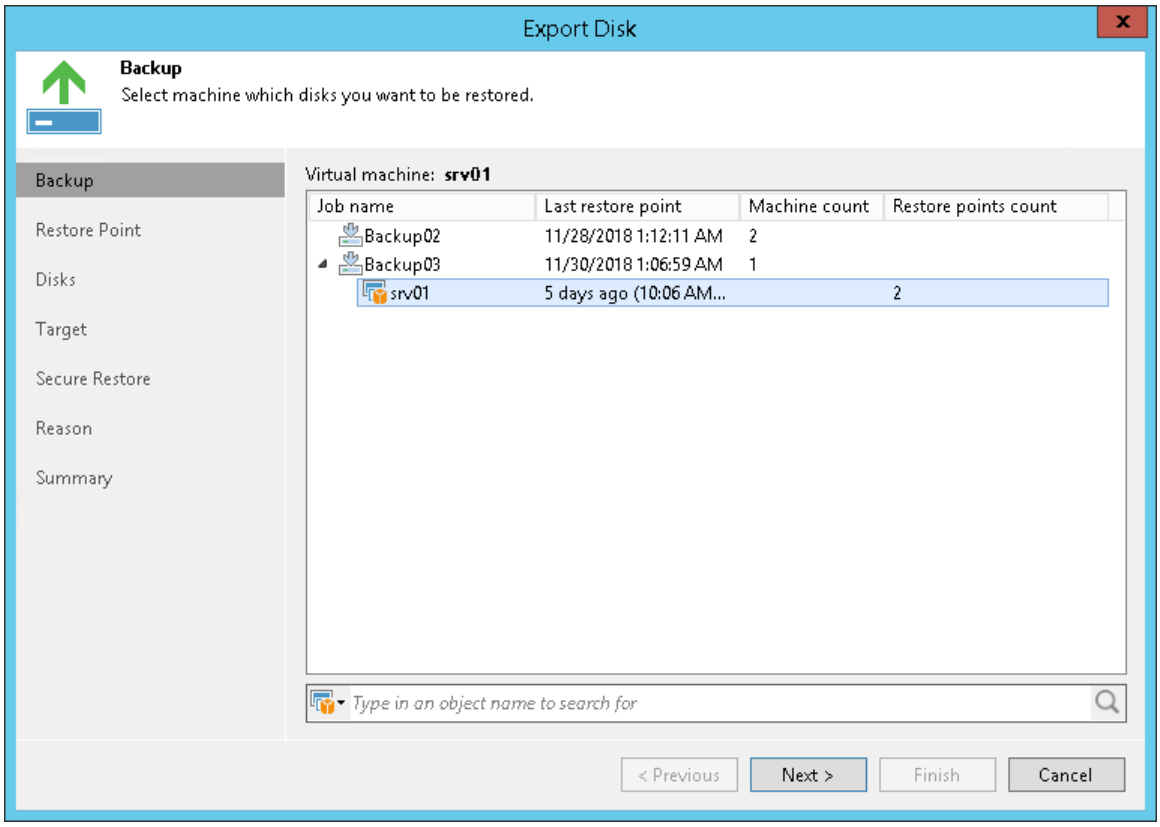
To launch the **Export Disk** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select one of the following:
  - **Agent** > **Disk Restore** > **Export disk** — to export disks of physical machines and virtual machines from backups created with Veeam Agent for Microsoft Windows, Veeam Agent for Linux or Veeam Agent for Mac.
  - **AWS EC2 backup** > **Entire machine restore** > **Export disk** — to export disks of EC2 instances from backups created with Veeam Backup for AWS.
  - **Azure IaaS** > **Entire machine restore** > **Instant Recovery to VMware vSphere** — to export disks of Azure VMs from backups created with Veeam Backup for Microsoft Azure.
  - **GCE backup** > **Entire machine restore** > **Export disk** — to export disks of VM instances from backups created with Veeam Backup for Google Cloud Platform.
  - **Nutanix backup** > **Entire machine restore** > **Export disk** — to export disks of VMs from backups created with Veeam Backup for Nutanix AHV.
  - **RHV Proxy** > **Entire machine restore** > **Export disk** — to export disks of VMs from backups created with Veeam Backup for RHV.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select a workload whose disks you want to export and click **Export Disks** on the ribbon. Alternatively, you can right-click the workload and select **Export content as virtual disks**.



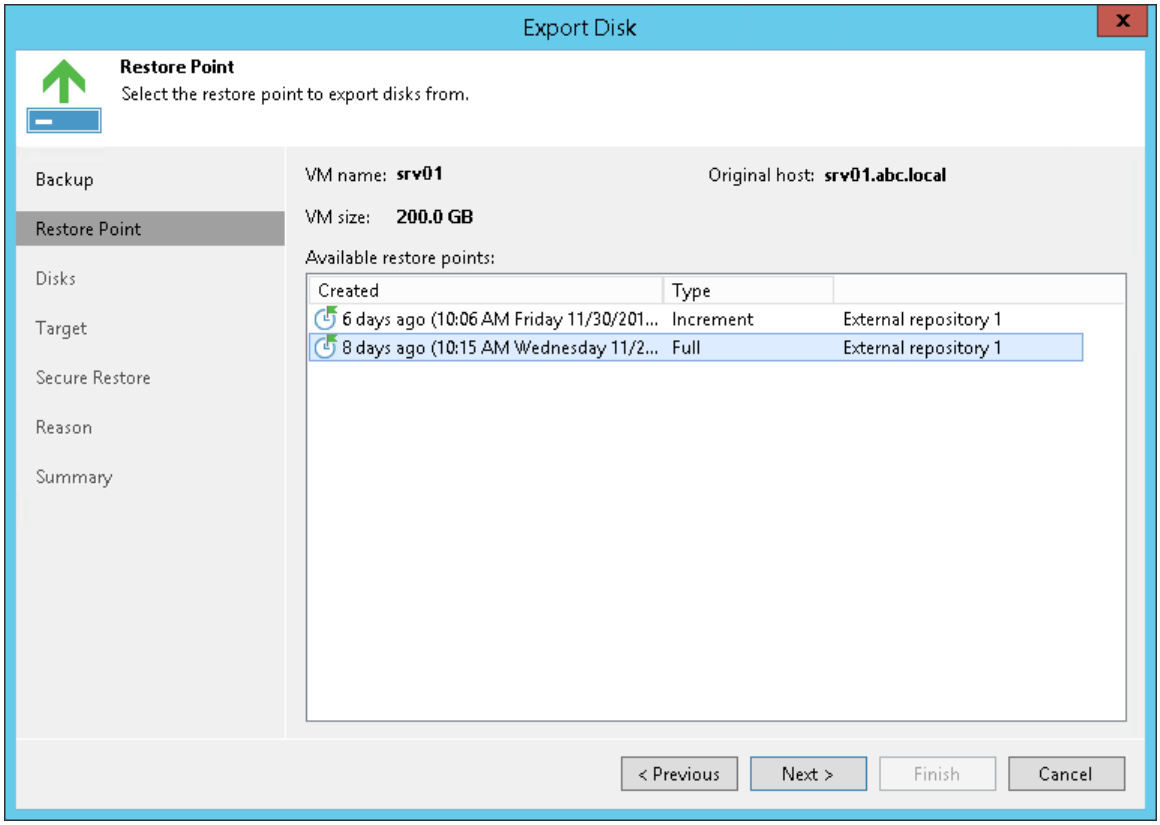
Step 2. Select Backup

At the **Backup** step of the wizard, expand a backup and select the workload whose disks you want to export.



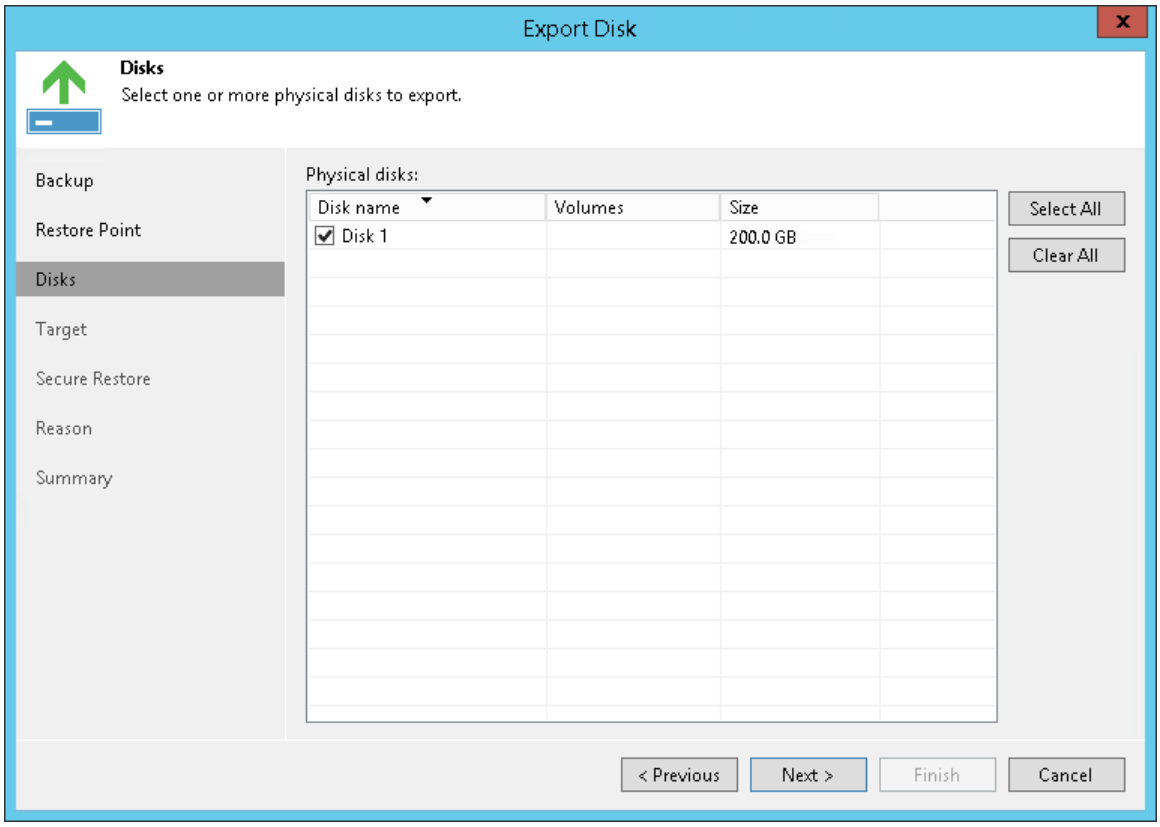
Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select the restore point from which you want to restore disks. By default, Veeam Backup & Replication restores disks to the latest state. However, you can restore disks to an earlier state.



Step 4. Select Disks

At the **Disks** step of the wizard, select check boxes next to the disks that you want to export.



Step 5. Select Destination and Disk Format

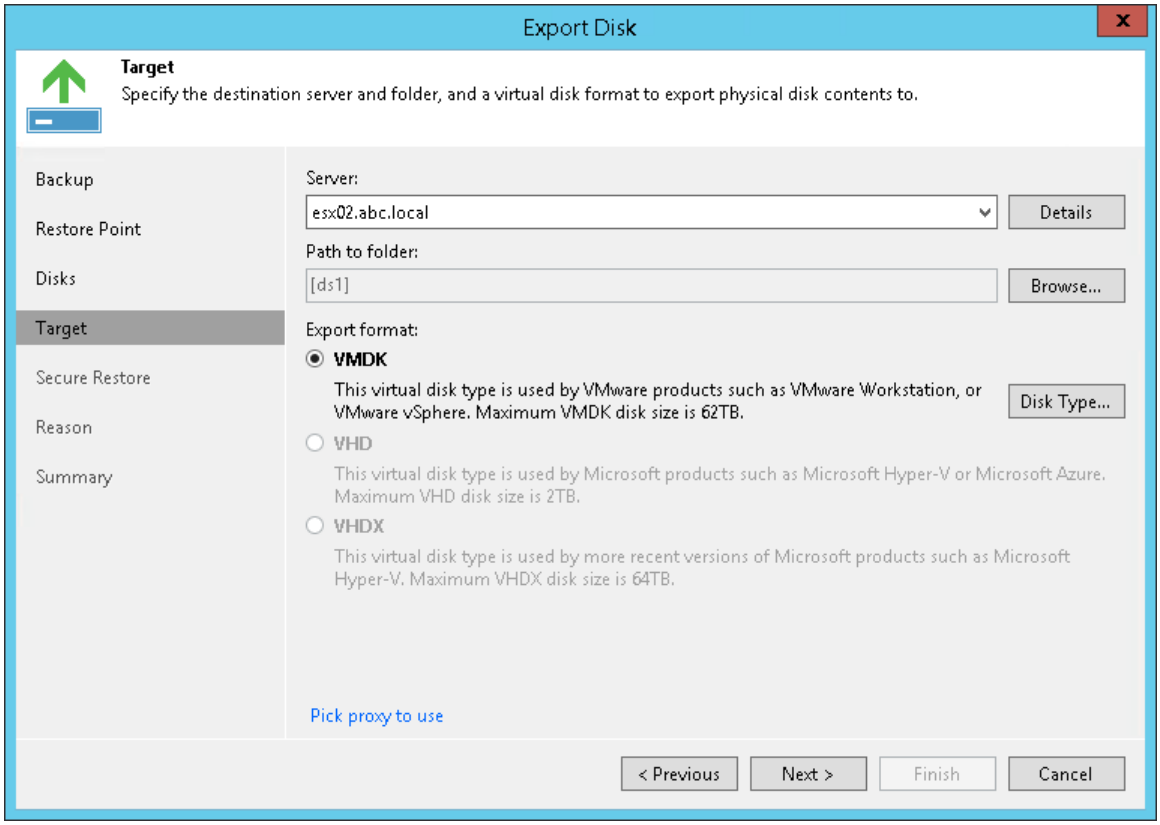
At the **Target** step of the wizard, select the destination for disk export and format in which you want to save the resulting virtual disks:

1. From the **Server** list, select a server on which the resulting virtual disks must be saved. If you plan to save the disks in the VMDK format on a datastore, select an ESXi host to which this datastore is connected.
2. In the **Path** to folder field, specify a datastore or folder on the server where the virtual disks must be placed.
3. Select the export format for the disks:
   - **VMDK** – select this option if you want to save the resulting virtual disk in the VMware VMDK format. This is the only available option if you export disks to a datastore.
   - **VHD** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHD format.
   - **VHDX** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHDX format (supported by Microsoft Windows Server 2012 or later).
4. [For VMDK disks exported to a datastore] Click the **Pick proxy to use** link to select backup proxies over which disk data must be transported to the target datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.

NOTE

Consider the following:

- If you select to export the resulting virtual disk to a datastore, you are able to save the virtual disk in the VMDK format only. Other options are disabled.
- If you export VMDK disks to a datastore, disks are saved in the thin provisioned format.
- If you export VMDK disks to a server, disks are saved in the thick format.
- VHD/VHDX disks are always saved as dynamically expanding.



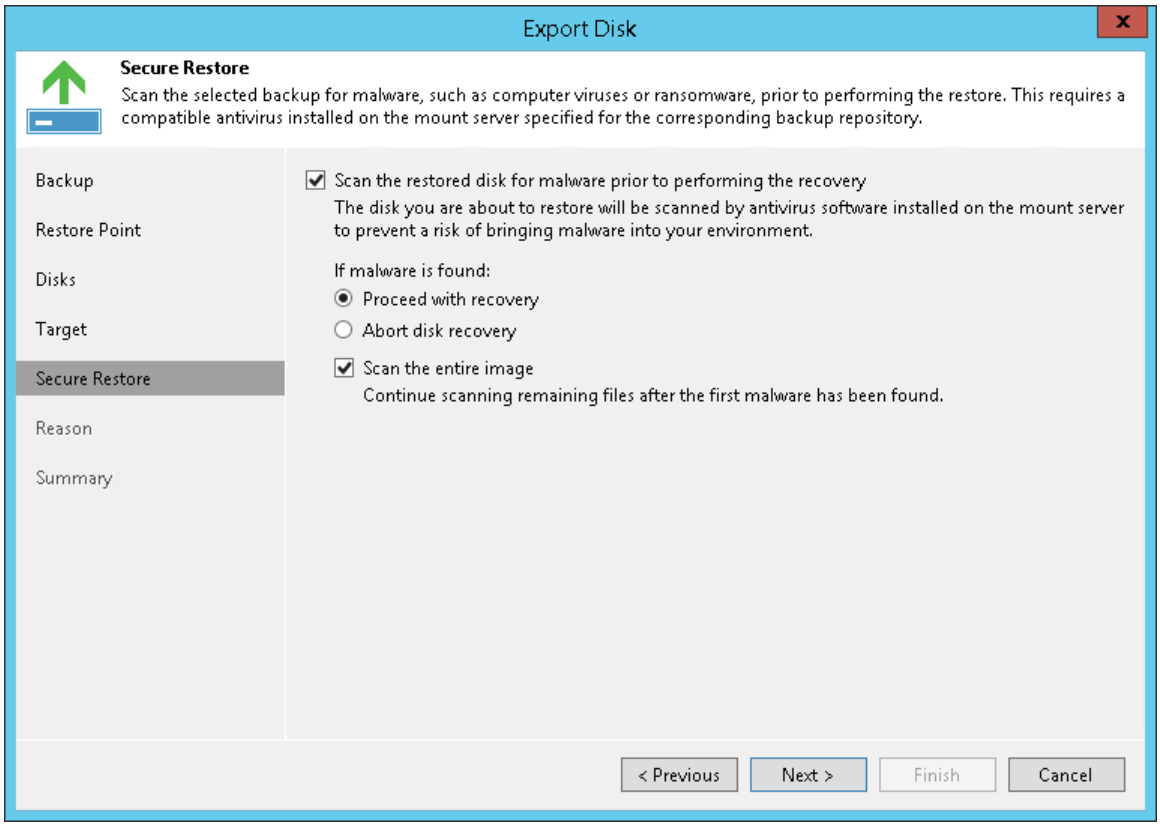
Step 6. Specify Secure Restore Settings

This step is available if you export disks of Microsoft Windows workloads.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to perform secure restore — scan exported disk data with antivirus software before restoring the disk. For more information on secure restore, see Secure Restore.

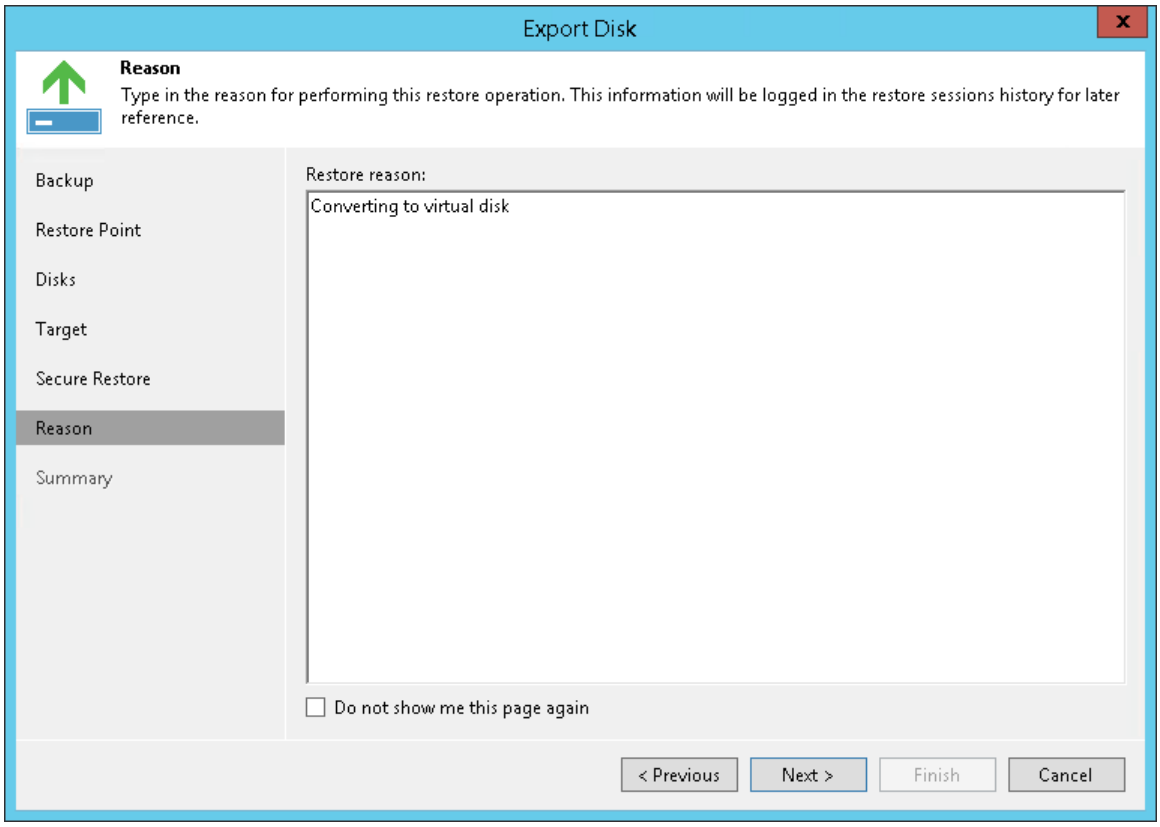
To specify secure restore settings:

1. Select the **Scan the restored disk for malware prior performing the recovery** check box.
2. Specify how the antivirus software must behave after the first malware threat is found: continue or abort disk recovery.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue disk scan after the first malware threat is found. For information on how to view results of the scan, see Viewing Malware Scan Results.



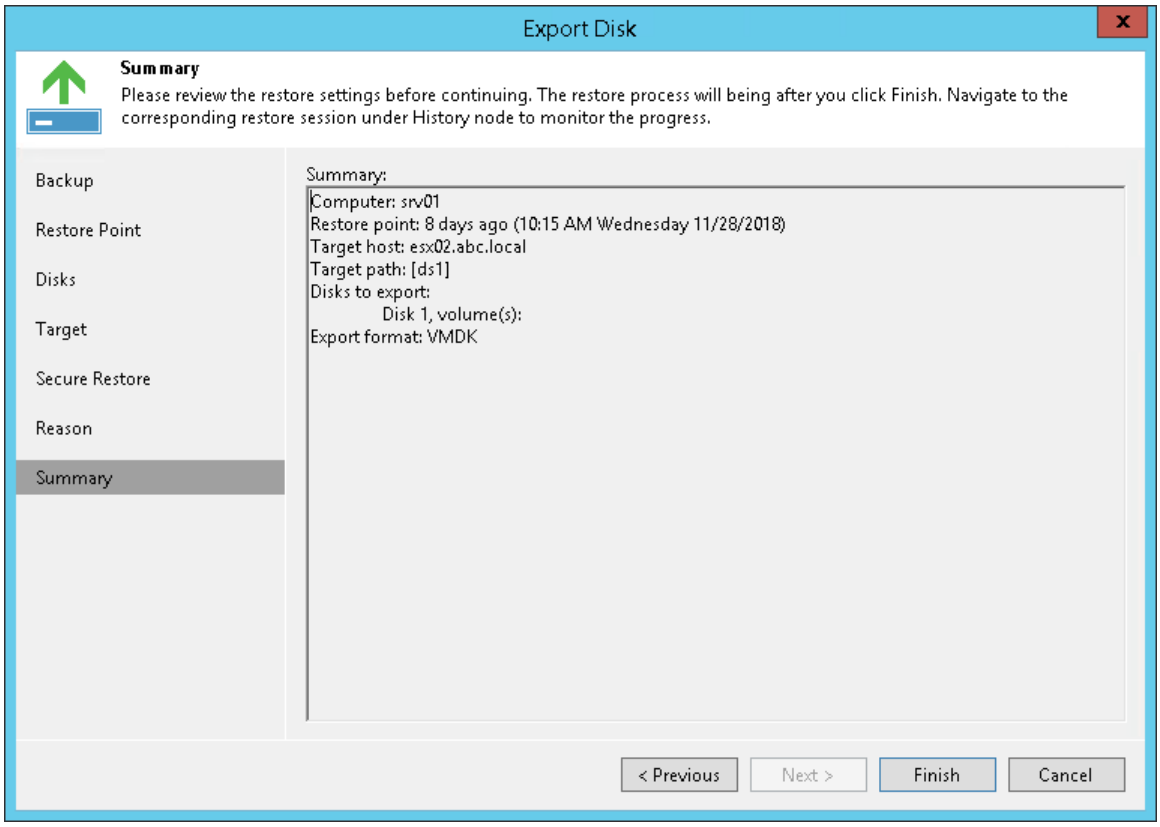
Step 7. Specify Export Reason

At the **Reason** step of the wizard, enter a reason for disk export.



Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.



Item Recovery

Item recovery includes the following methods:

- VM files restore — to restore VM files (.VMX, .VMXF and so on) without restoring the entire VM.
- Guest OS file recovery — to recover individual guest OS files from Windows, Linux, Mac and other guest OS file systems. You can restore files and folders directly from a regular image-level backup or replica.
- Application items restore — to restore items from different applications such as Microsoft Active Directory, Microsoft SQL Server and so on. Application items are recovered directly from VM backups and replicas. To recover application items, Veeam Backup & Replication uses the capabilities of Veeam Backup Explorers.

VM Files Restore

You can restore specific VM files (.VMX, .VMXF, .NVRAM, .VMDK including flat files) if any of these files are deleted or the datastore is corrupted. This option provides a great alternative to entire VM restore, for example, when your VM configuration file is missing and you need to restore it. Instead of restoring the whole VM image to the production storage, you can restore a specific VM file only.

When you perform VM file restore, VM files are restored directly from regular image-level backups, without prior de-staging of VM images from backups. VM files can be restored to the original VM location or to a new location.

Restoring VM Files

To restore VM files from a backup, use the **Virtual Machine Files Restore** wizard.

Before You Begin

Before you restore VM files, check the following prerequisites:

- You can restore VM files from a backup that has at least one successfully created restore point.
- The server on which you plan to save restored VM files must be added to the backup infrastructure.

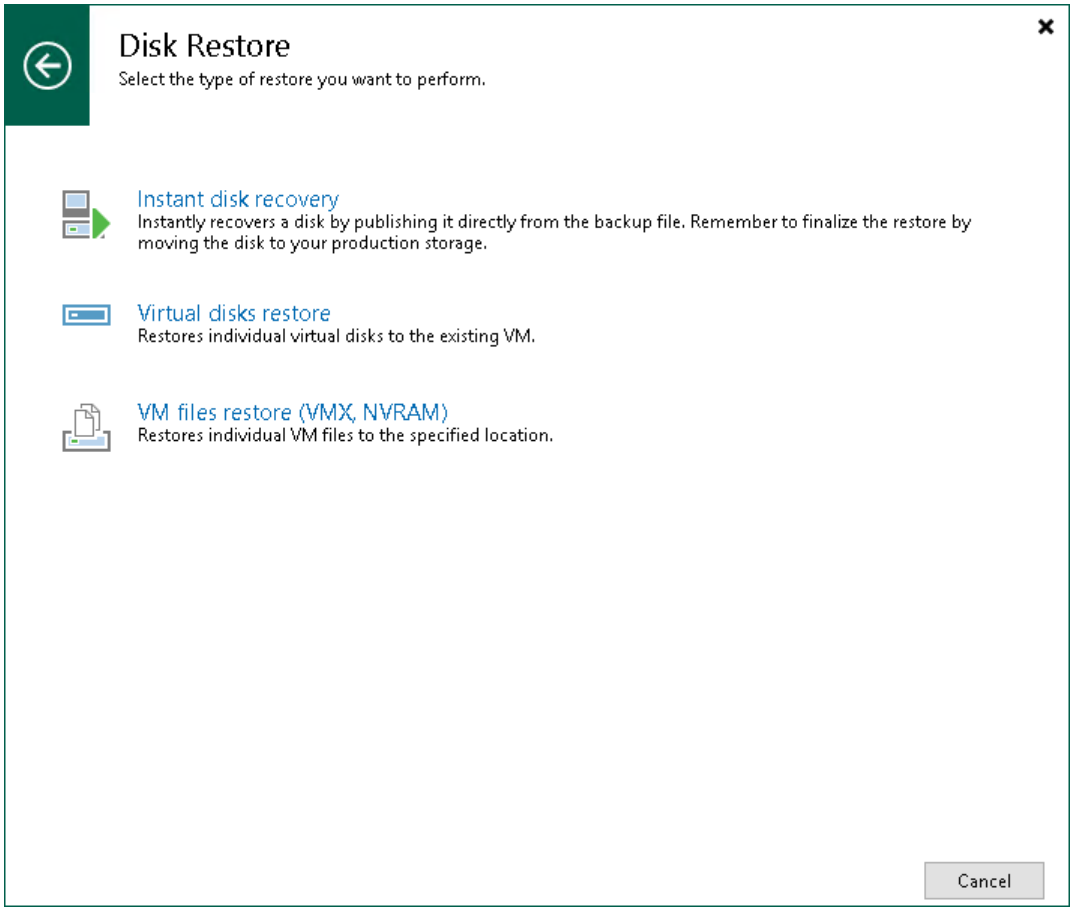
Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **VM files restore (VMX, NVRAM)**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **VM Files** on the ribbon. Alternatively, you can right-click the VM whose files you want to restore and select **Restore VM files**.

  In this case, you will pass to the Restore Point step of the wizard.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **VM files**. In this case, you will pass to the Restore Point step of the wizard.

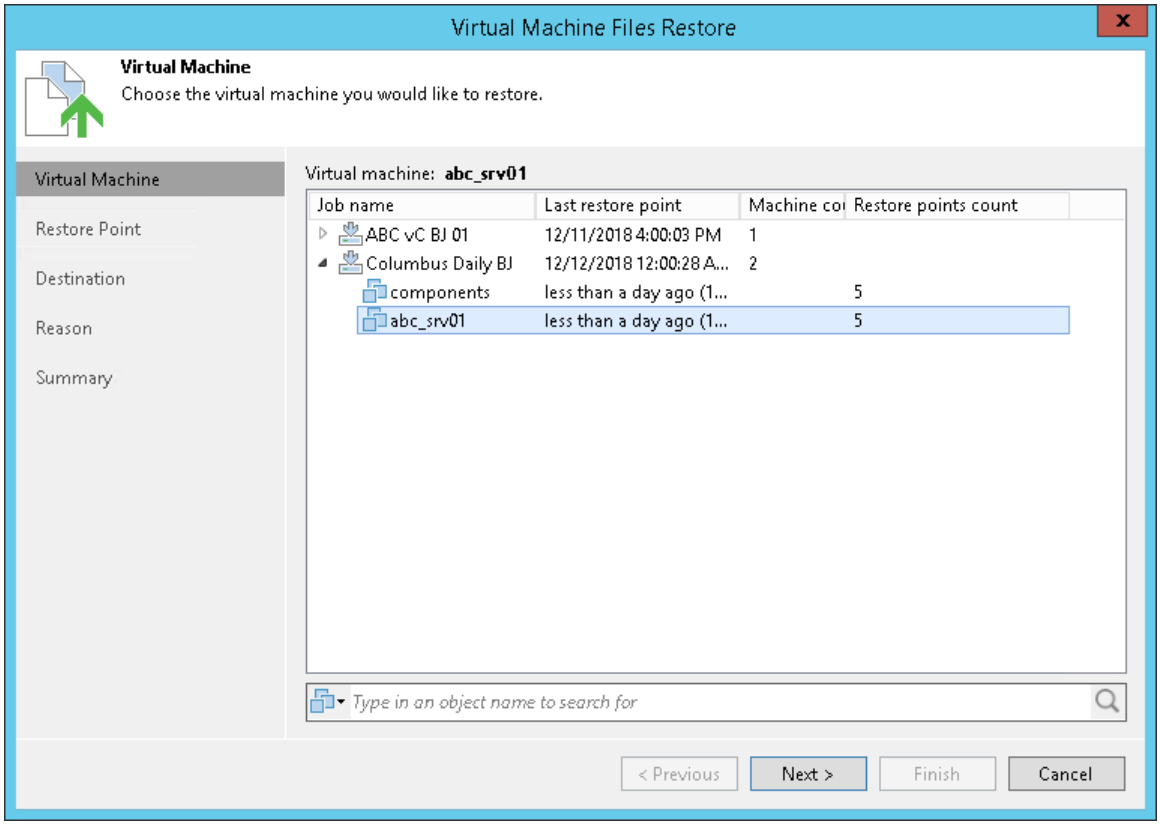
  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.



Step 2. Select VM

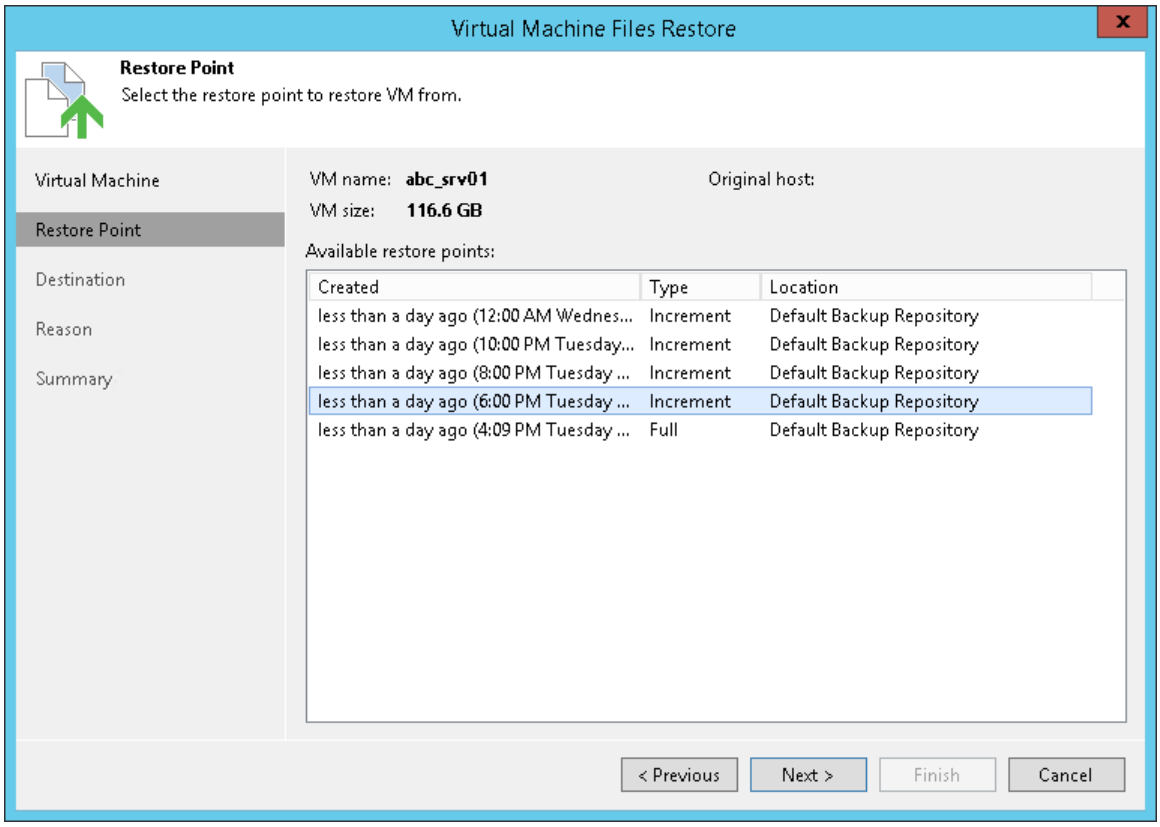
At the **Virtual Machine** step of the wizard, select the VM whose files you want to restore:

1. In the **Virtual machine** list, expand the necessary backup.
2. Select the VM.



Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to restore VM files.

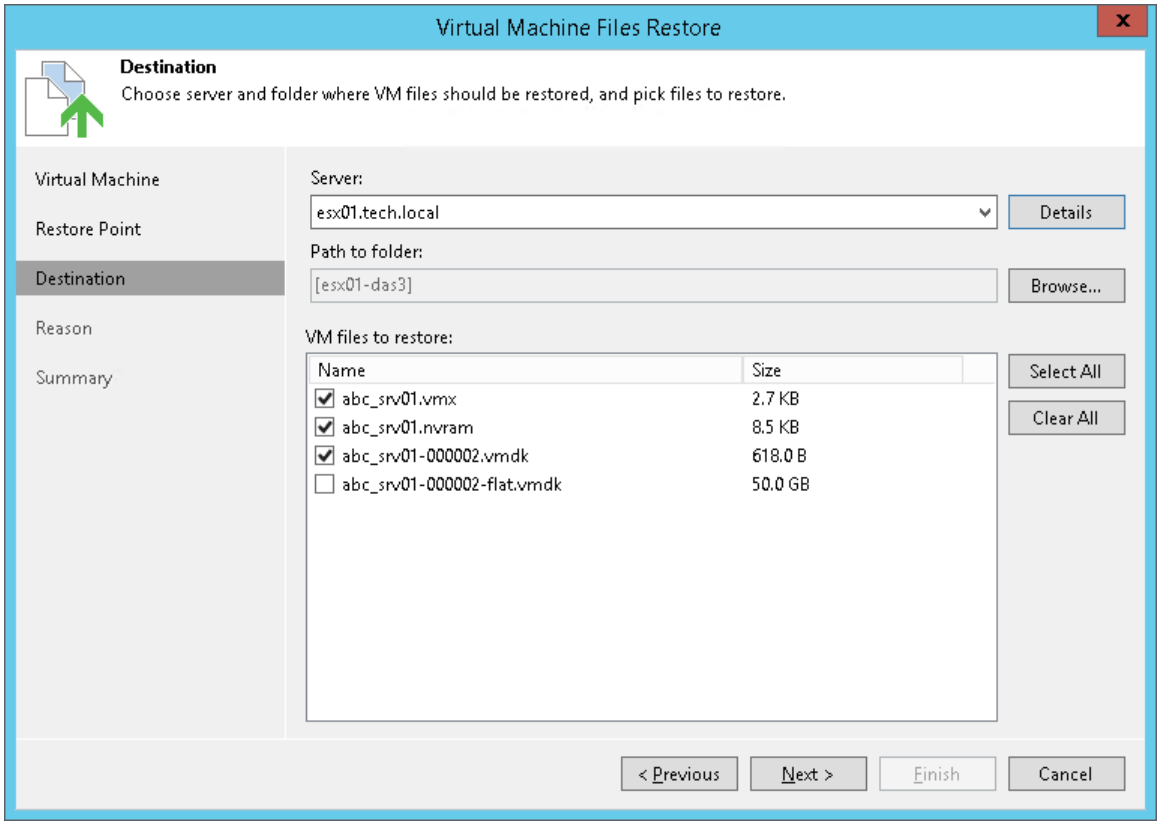


Step 4. Select VM Files and Destination

At the **Destination** step of the wizard, select VM files that you want to restore and destination where the restored files must be stored:

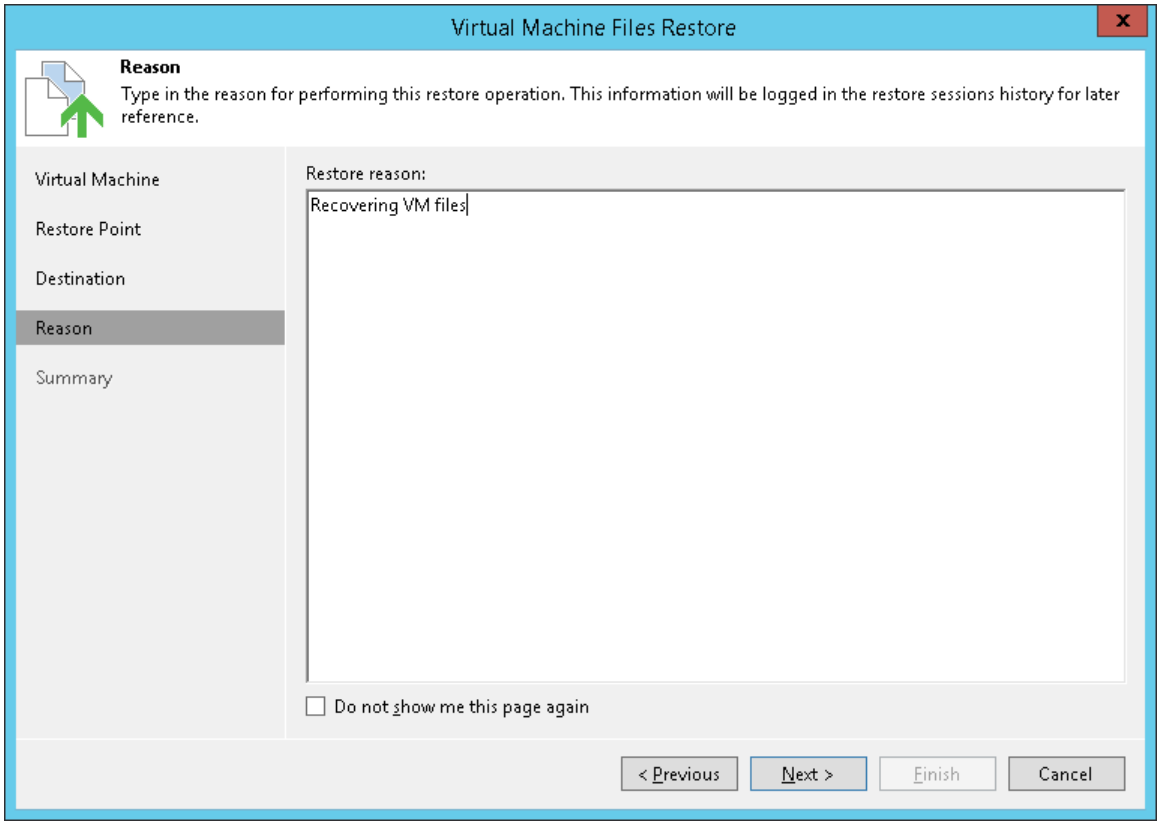
1. From the **Server** list, select where to store VM files: on an ESXi host, on the backup server or on a Microsoft Windows or Linux server added to the backup infrastructure. Use the **Details** button to view or change connection settings of the target host or server.
2. In the **Path to folder** section, specify a path to the folder on the selected host where files must be restored.

   To create a new folder for restored files, click **Browse**. In the **Select Folder** window, select the target location for VM files and click **New Folder** at the bottom of the window.
3. In the **VM files to restore** section, select check boxes next to files that you want to restore.



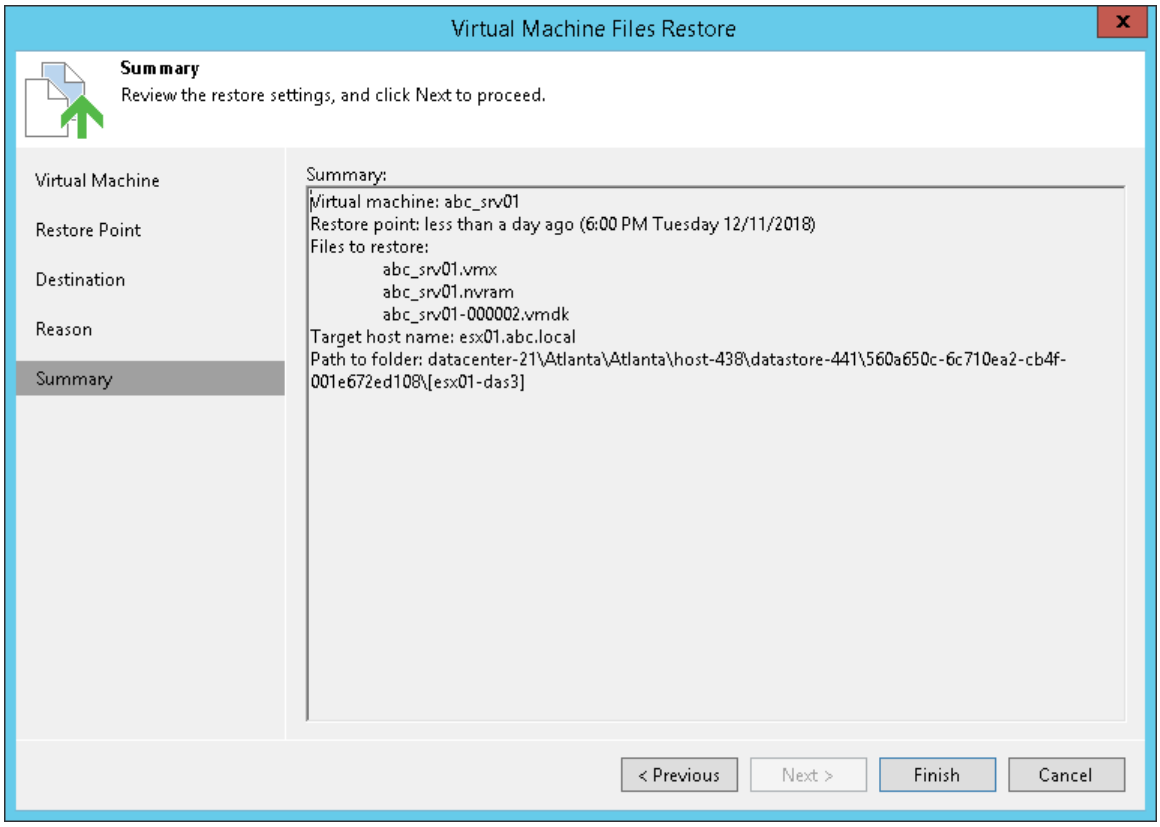
Step 5. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM files. The information you provide will be saved in the session history and you can reference it later.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, check details for the restore task and click **Finish** to start VM files restore.



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. When restoring files or folders, you do not need to extract the VM image to a staging location or start the VM prior to restore. You can restore files and folders directly from a regular image-level backup or replica to the necessary point in time.

Using the following methods, you can restore files from different guest OS file systems:

- Restore from FAT, NTFS or ReFS

  This method that helps you restore files from Microsoft Windows VMs with NTFS, FAT and ReFS file systems.

- Restore from Linux, Unix and Other File Systems (multi-OS file-level restore)

  This method that helps you restore files from Linux, Solaris, BSD, Novell Storage Services, Unix and Mac machines.

  Note that this method supports recovery of files and folders only. Recovery of other file system objects such as pipes is not supported.

- Restore from Other File Systems

  This method that helps you restore files from file systems that other methods from the list do not support.

Restore from FAT, NTFS or ReFS

The restore from FAT, NTFS and ReFS method helps you restore files from Microsoft Windows VMs with NTFS, FAT and ReFS file systems. For the full list of supported file systems, see Platform Support.

You can restore files to the original or new location, work with the restored files using Microsoft Windows File Explorer or launch application item restore for the files. For more information, see Finalize Restore.

How Restore Works

When you perform file-level restore, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a mount point under the `C:\VeeamFLR\<vmname>` folder and mounts VM disks from the backup or replica to it. For more information on which machines Veeam Backup & Replication creates mount points, see File-Level Restore Scenarios and Mount Points.

   To access VM disk content, Veeam Backup & Replication uses a separate program — Virtual Disk Driver (VDK) that is provided with the product. VM disks are not physically extracted from the backup file or VM replica. Veeam Backup & Replication emulates their presence on the backup server or Veeam Backup & Replication console. The backup file or VM replica itself remains in the read-only state.

2. Veeam Backup & Replication launches the Veeam Backup browser on the Veeam Backup & Replication console. The Veeam Backup browser shows the content of disks mounted to the machine where the mount point is created.

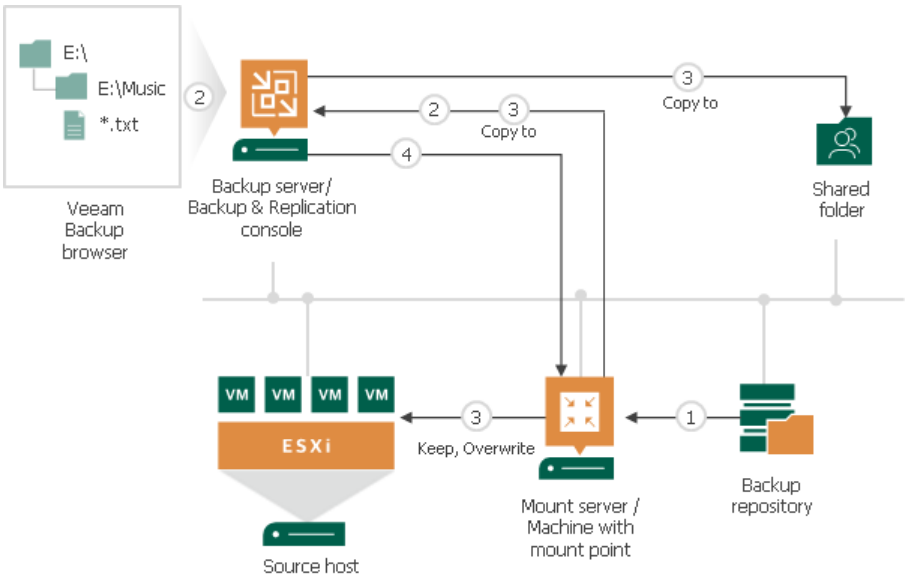
   You can browse the VM guest file system in the Veeam Backup browser and restore files or folders to the original or new location.

3. Depending on whether you restore files or folders to the original or new location, the operations differ:
   - If you restore to the original location, that is, select the **Keep** or **Overwrite** command.

     The machine where the mount point is created connects to the VM over network or VIX API/vSphere Web Services if a connection over the network cannot be established.
   - If you restore files to a new location, that is, select the **Copy to** command.

     The machine where the mount point is created connects to the Veeam Backup & Replication console. The connection is performed over network or VIX API/vSphere Web Services if a connection over the network cannot be established. If you recover files to a shared folder, the Veeam Backup & Replication console connects to this folder.
4. When the restore process is finished or the Veeam Backup browser is closed by timeout, Veeam Backup & Replication removes all the created mount points.

When you perform file-level restore, you can also open files in Microsoft Windows explorer to view file content. In this case, Veeam Backup & Replication creates an additional mount point on the Veeam Backup & Replication console. For more information on how to view file content, see Open Files in Microsoft Windows Explorer.



Mount Points and Restore Scenarios

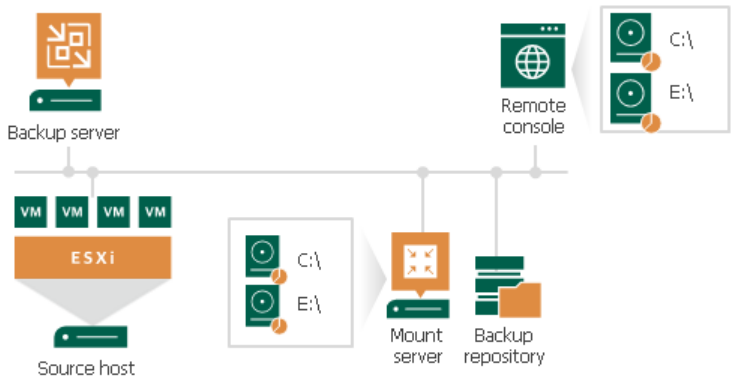
Veeam Backup & Replication creates mount points on different machines depending on a file-level restore scenario you use.

Restoring Files from Backups

When you restore files from backups that reside in the backup repository, Veeam Backup & Replication creates a mount point on the following machines:

1. Mount server associated with the backup repository on which the backup file resides. Veeam Backup & Replication uses this mount point when the restore process starts and allows you to browse the VM file system and restore files.

2. Veeam Backup & Replication console. Veeam Backup & Replication uses this mount point only if you instruct Veeam Backup & Replication to mount VM disks there. For more information, see Open Files in Microsoft Windows Explorer.



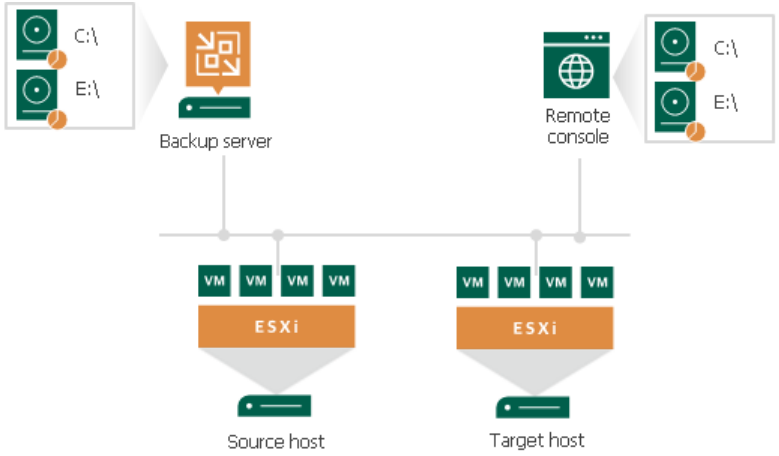
NOTE

Backup files on HPE StoreOnce are locked exclusively by a restore task. For this reason, Veeam Backup & Replication uses only one mount point on the backup server or Veeam Backup & Replication console machine for backups on HPE StoreOnce.

Restoring Files from Replicas

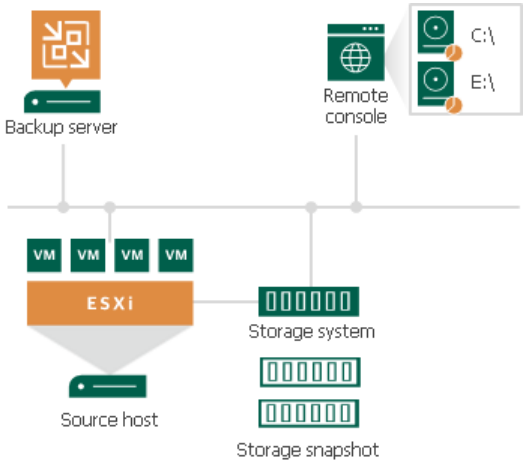
When you restore files from a VM replica, Veeam Backup & Replication creates a mount point on the following machines:

1. Backup server. Veeam Backup & Replication uses this mount point when the restore process starts and allows you to browse the VM file system and restore files.
2. Veeam Backup & Replication console. Veeam Backup & Replication uses this mount point only if you instruct Veeam Backup & Replication to mount VM disks there. For more information, see Open Files in Microsoft Windows Explorer.



Restoring Files from Storage Snapshots

When you restore files from storage snapshots, Veeam Backup & Replication uses one mount point on the backup server or Veeam Backup & Replication console machine.



Restoring Files for Veeam Explorers

If you have launched file-level restore and then launch application items restore from the Veeam Backup browser, Veeam Backup & Replication uses additional restore points apart from the the mount points described in the subsections above:

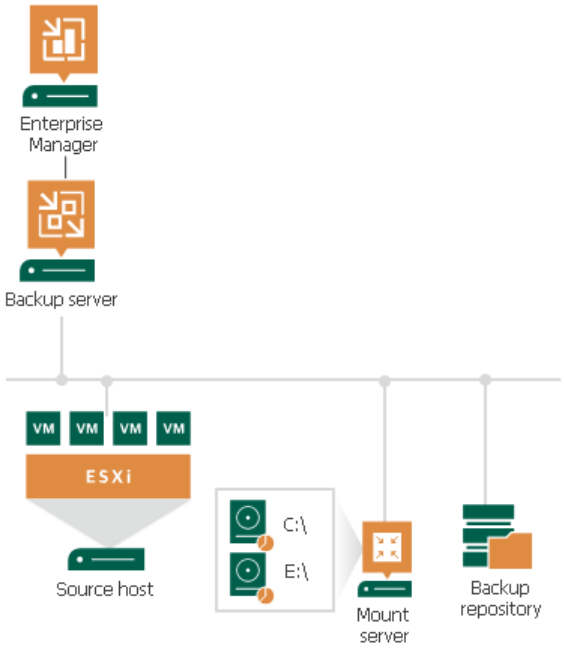
- If you restore Microsoft Active Directory or Microsoft Exchange items, Veeam Backup & Replication creates an additional mount point on the Veeam Backup & Replication console. However, if during file-level restore the mount point was already created on a machine where the Veeam Backup & Replication console is installed, an additional point is not created.
- If you restore Microsoft SharePoint, Veeam Backup & Replication creates an additional mount point on a staging Microsoft SQL Server.
- If you restore Microsoft SQL Server items, Veeam Backup & Replication creates two additional mount points: one on a staging Microsoft SQL Server and the other one on the target VM to which you restore the application items.
- If you restore Oracle server items, Veeam Backup & Replication creates two additional mount points: one on a staging Oracle server and the other one on the target VM to which you restore the application items.

To create mount points on Linux and Windows machines, Veeam Backup & Replication uses different methods:

- To create a mount point on Microsoft Windows machines, Veeam Backup & Replication uses the iSCSI protocol. The remote machine or staging server acts as an iSCSI initiator. The machine on which the Veeam Explorer runs acts as an iSCSI target. The iSCSI mount point is non-persistent — it is created only for duration of the restore process.
- To create a mount point on Linux VMs (for Oracle running on Linux), Veeam Backup & Replication uses FUSE.

Restoring Files from Veeam Backup Enterprise Manager

When you restore files from Veeam Backup Enterprise Manager, Veeam Backup & Replication creates a mount point on the mount server associated with the backup repository on which the backup file resides. Veeam Backup & Replication uses this mount point when the restore process starts and allows you to browse the VM file system and then restore files.



Restoring Files from Nutanix AHV Snapshots

To learn how Veeam Backup & Replication performs guest OS file restore from backup snapshots, user snapshots and persistent disk snapshots created by AHV Backup Proxy and which mount points are used, see the Restoring Guest OS Files section of the Veeam Backup for Nutanix AHV User Guide.

Restoring Files from Amazon EC2 Instances

To perform a file-level restore for EC2 instances, Veeam Backup & Replication utilizes the functionality of AWS Plug-in for Veeam Backup & Replication. To learn more how file-level restore is performed, see the File-Level Restore section of the Veeam Backup for AWS User Guide.

Restoring Files from Microsoft Azure VMs

To perform a file-level restore for Microsoft Azure VMs, Veeam Backup & Replication utilizes the functionality of Microsoft Azure Plug-in for Veeam Backup & Replication. To learn how file-level restore is performed, see the File-Level Restore section of the Veeam Backup for Microsoft Azure User Guide.

Requirements and Limitations

Before you restore VM guest OS files, mind the following requirements and limitations.

Requirements

Consider the following general requirements:

- You can restore files from the file systems listed in the Platform Support section.
- The account that you use to start the Veeam Backup & Replication console and to connect to the backup server must have permissions and privileges described in Veeam Backup & Replication Console Permissions.
- You can restore VM guest OS files from a backup or replica that has at least one successfully created restore point.
- [For restore to the original location] The mount server must have access to the VM guest OS (if file-level restore is performed over the network) or vCenter Server and ESXi host where the target VM runs (if file-level restore is performed over VIX API/vSphere Web Services).
- [For restore to the original location] VMware Tools must be installed on the target VM.

Requirements for ReFS

If you plan to restore files from a VM running Microsoft Windows ReFS, consider the following requirements for the Veeam Backup & Replication components involved in the restore process:

- The machine on which a mount point is created (for example, the mount server) must run Microsoft Windows Server 2012 or later.
- [ReFS 3.x] If you plan to restore files from a VM running Microsoft Windows ReFS 3.x, the machine on which mount point is created must run Microsoft Windows Server 2016 or later and the ReFS version must be supported on it.
- The machine on which a mount point is created must run Microsoft Windows Server of the same version or newer than the guest OS of a VM from which you plan to restore files.

To learn in which scenarios on which machines mount points are created, see Mount Points and Restore Scenarios.

Requirements for Data Deduplication

If you plan to restore files from a VM running Microsoft Windows Server 2012 or later and data deduplication is enabled for some VM volumes, consider the following for the Veeam Backup & Replication components involved in the restore process:

- The machine on which a mount point is created (for example, the mount server) must run Microsoft Windows Server 2012 or later.
- The machine on which a mount point is created must run Microsoft Windows Server of the same version or newer than the guest OS of a VM from which you plan to restore files.
- Data deduplication must be enabled on the machine on which a mount point is created.

To learn in which scenarios on which machines mount points can be created, see Mount Points and Restore Scenarios.

Limitations

Consider the following limitations:

- You cannot restore pipes and other file system objects. File-level restore supports recovery of files and folders only.
- [For Microsoft Windows workloads] You can restore files from basic disks and dynamic disks (including simple, mirrored and striped volumes).
- [For Linux workloads] You can restore files from basic disks, Linux LVM (Logical Volume Manager) and ZFS pools. Encrypted LVM volumes are not supported.
- [For backups of BSD, Mac and Solaris VMs] You cannot restore files directly to the original location. Use the **Copy to** option instead.
- Processing of reparse points is supported only for NTFS. Note that reparse points with reparse tag values other than IO_REPARSE_TAG_MOUNT_POINT, IO_REPARSE_TAG_SYMLINK and IO_REPARSE_TAG_DEDUP may be processed and restored incorrectly.
- You cannot restore files from a backup created in the reverse incremental mode if the backup job is being performed. If the backup is created in the incremental backup mode and the backup job is being performed, you can restore files from any available restore point.
- You cannot restore VM guest OS files to the original location if you have excluded the system disk from the VM backup used for restore.
- You cannot restore VM guest OS files from a running replica or if the replication job with the necessary VM is being performed.

Restoring VM Guest OS Files (FAT, NTFS or ReFS)

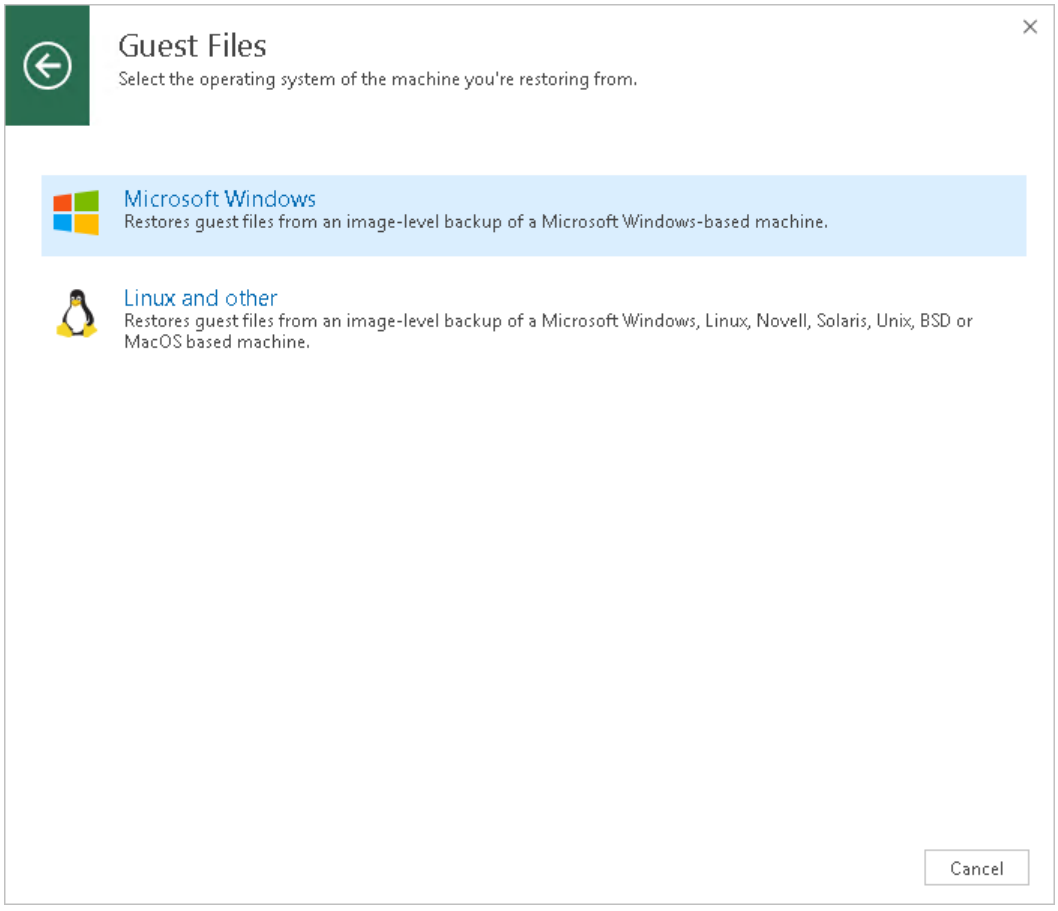
To restore VM guest OS files and folders, use the **File Level Restore** wizard.

Step 1. Launch File Level Restore Wizard

To launch the **File Level Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Guest files restore** > **Microsoft Windows** or **Restore from replica** > **Guest files restore** > **Microsoft Windows**.
- Open the **Home** view. In the inventory pane, select **Backups** or **Replicas**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **Guest files** > **Microsoft Windows** on the ribbon. Alternatively, right-click the VM whose files you want to restore and select **Restore guest files** > **Microsoft Windows**.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Guest files (Microsoft Windows)**.

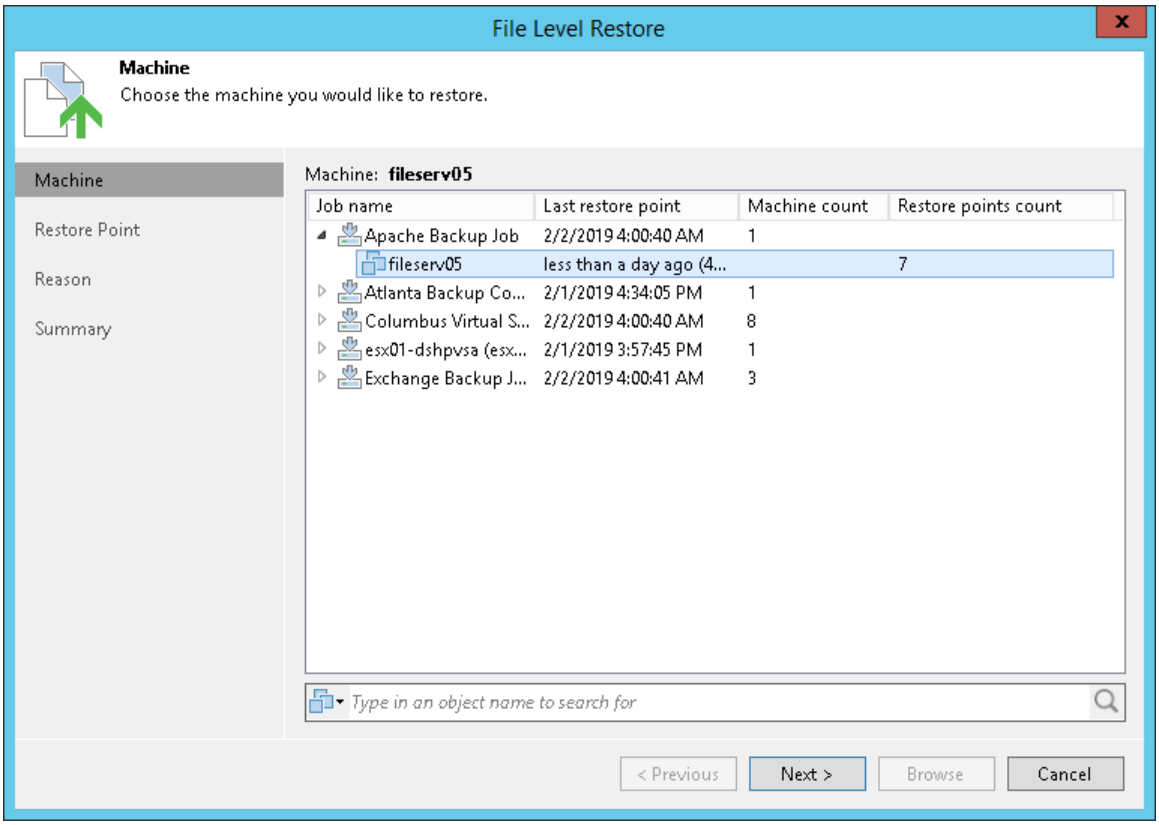
  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.



Step 2. Select VM

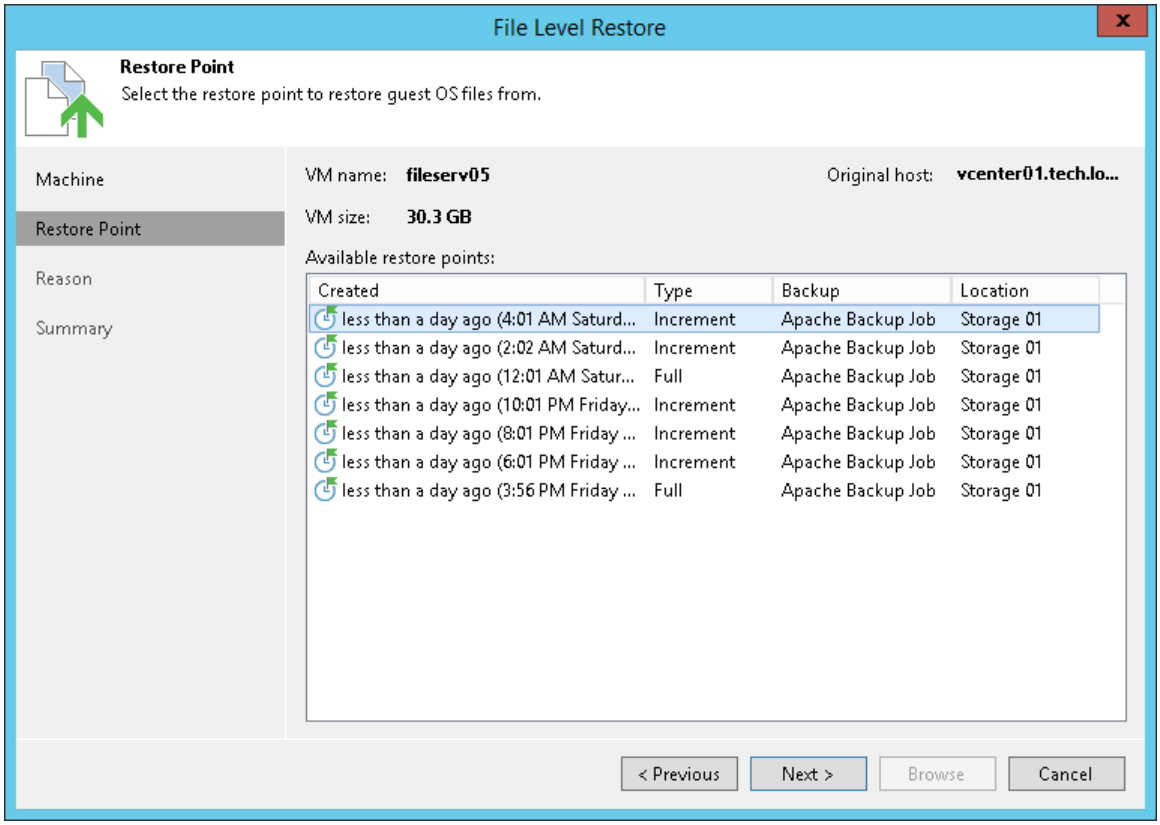
At the **Machine** step of the wizard, select the VM whose guest OS files you want to restore:

1. In the **Machine** list, expand the necessary backup.
2. Select the VM.



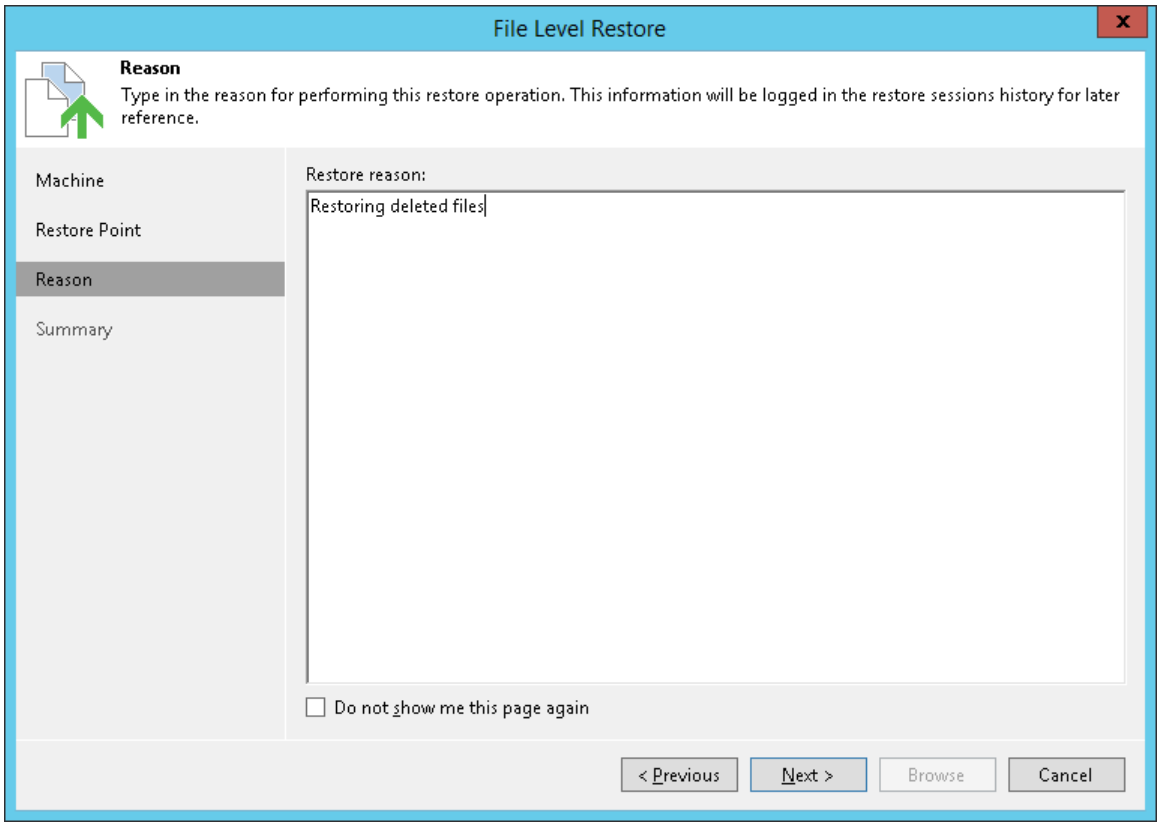
Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to restore VM guest OS files.



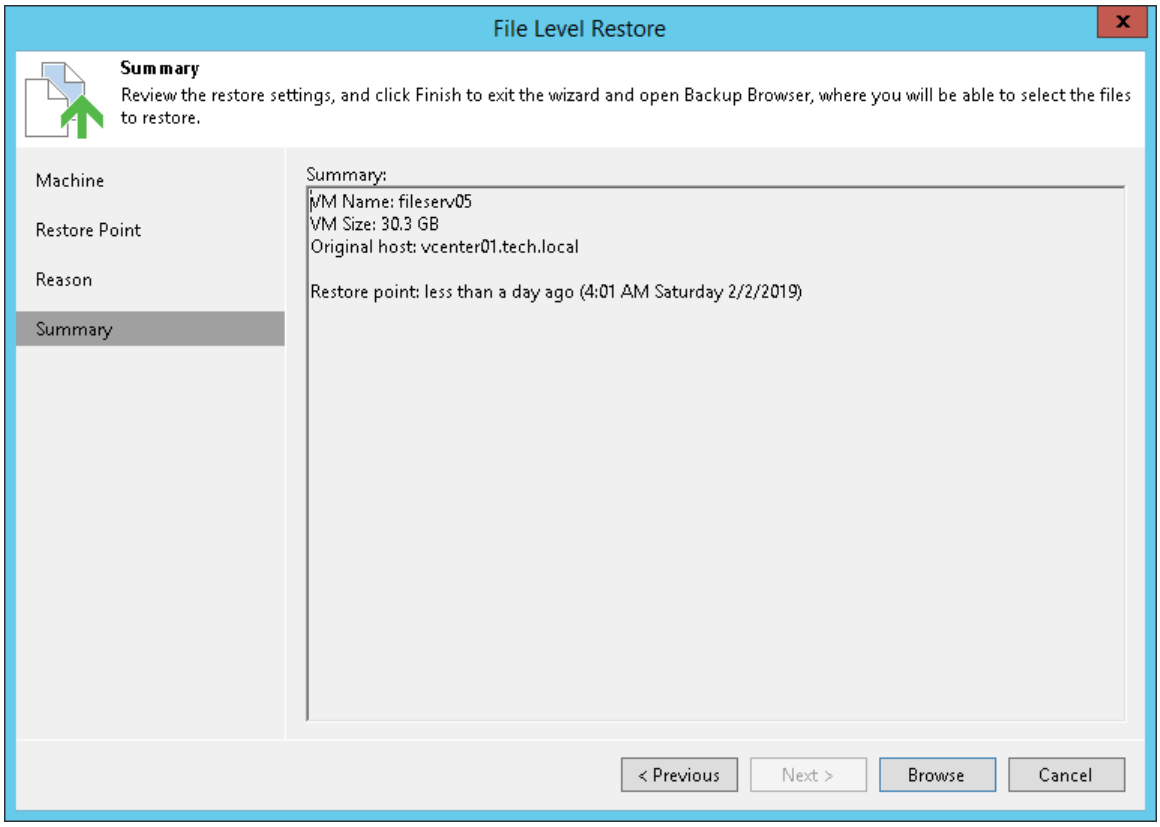
Step 4. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM guest OS files. The information you provide will be saved in the session history and you can reference it later.



Step 5. Verify Restore Settings

At the **Summary** step of the wizard, check details of the restore task and click **Browse** to close the wizard and open the Veeam Backup browser.



Step 6. Finalize Restore

After you close the wizard, Veeam Backup & Replication opens the Veeam Backup browser with the file system tree of the restored VM. Note that names of the restored VM hard disks may differ from the original ones.

You can perform the following operations in the Veeam Backup browser:

- Restore files and folders to the original location
- Save files and folders to a folder on the backup server or network shared folder
- Launch application item restore
- Open files in Microsoft Windows File Explorer

After you finish restoring files, close the Veeam Backup browser.

Restoring Files to Original Location

To restore a file or folder to its original location, in the Veeam Backup browser, right-click the file or folder and select one of the following commands:

- To overwrite the original file on the VM guest OS with the file restored from the backup, select **Restore** > **Overwrite**.
- To save the file restored from the backup next to the original file, select **Restore** > **Keep**.

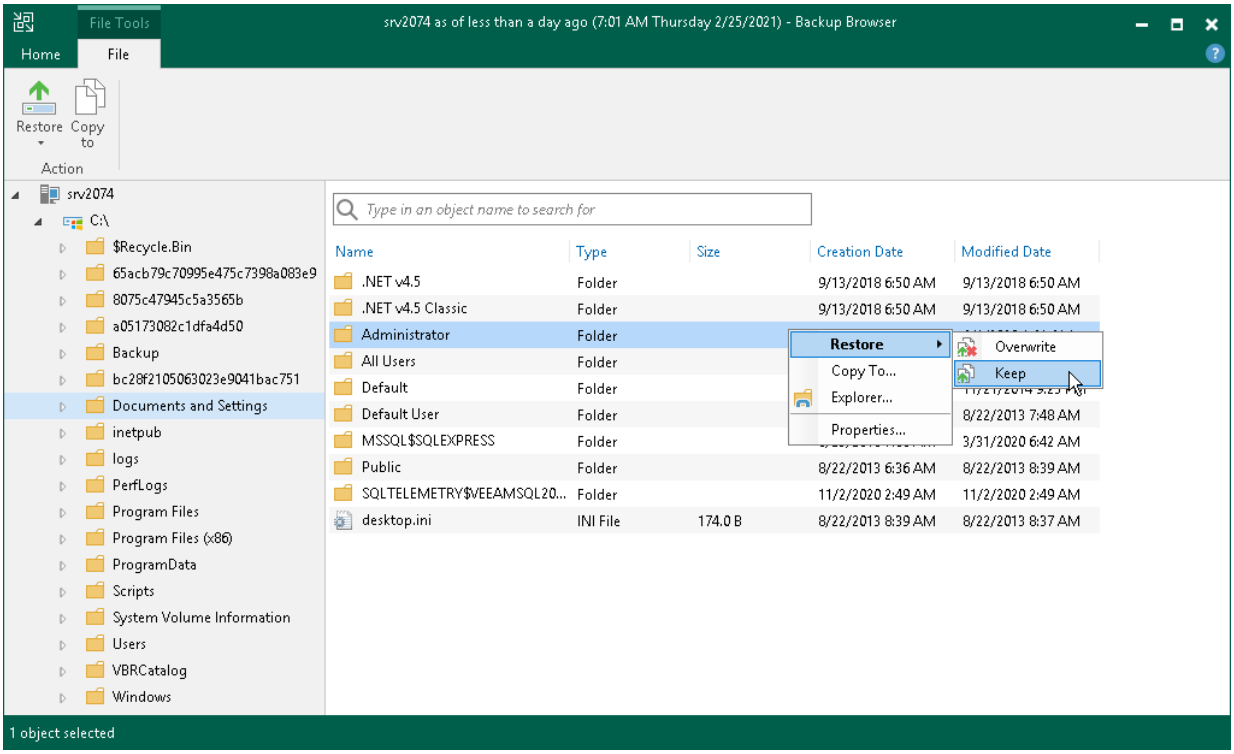
  Veeam Backup & Replication will add the *RESTORED-* prefix to the original file name and store the restored file in the same folder where the original file resides. If the file with the *RESTORED-* prefix already exists in the original location, Veeam Backup & Replication will name the restored file in the following format: *RESTORED-<filename>_YYYYMMDD_HHMMSS*.

IMPORTANT

Restore to the original location may fail for the following reasons:

- VMware Tools are not installed on the target VM.
- Application-aware processing is not supported for the Microsoft Windows OS of the original VM.

To restore guest OS files in such situation, you can use 1-click file-level restore or copy files to the selected folder and then move them to their original location.



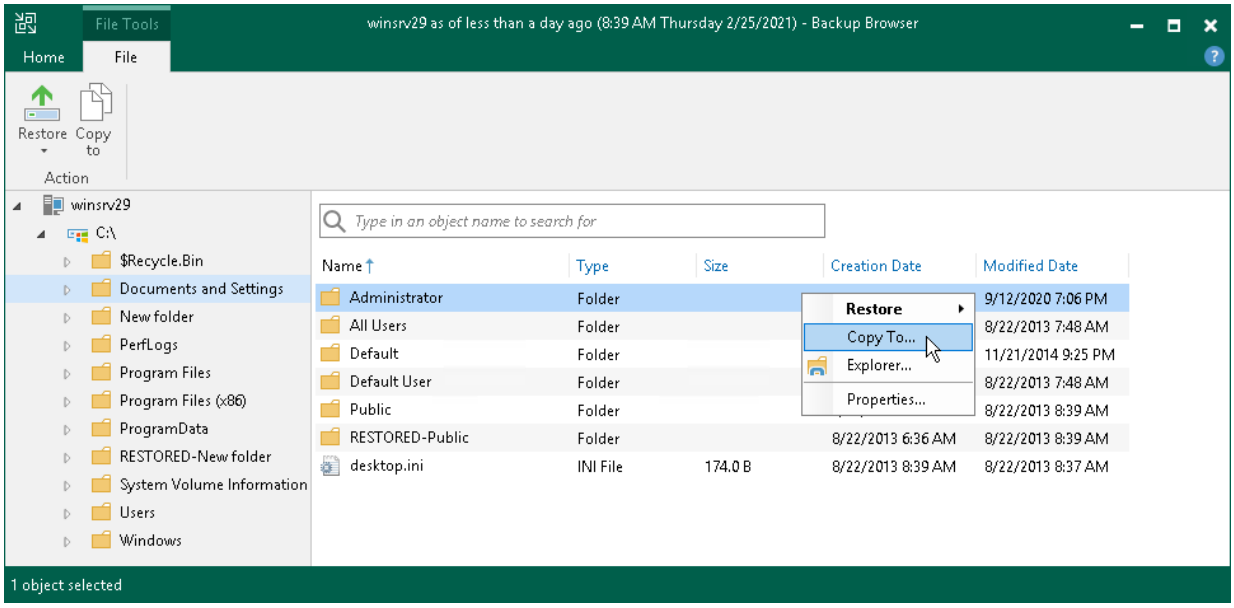
Saving Files to a New Location

To save restored files or folders on the machine where Veeam Backup & Replication console is installed or in a network shared folder:

1. Right-click the necessary file or folder in the file system tree or in the details pane on the right and select **Copy To**.
2. In the **Choose Folder** window, select the necessary destination:
   - To recover files to a folder on the server where the Veeam Backup & Replication console is installed, click **Browse** to find the necessary folder.
   - To recover files to a shared folder, enter a path to the destination folder in the **Choose folder** field.
3. In the **Choose Folder** window, choose whether to preserve original NTFS permissions or not:
   - Select the **Preserve permissions and ownership** check box to keep the original ownership and security permissions for restored objects. Veeam Backup & Replication will copy selected files and folders along with associated Access Control Lists, preserving granular access settings.

- Leave the **Preserve permissions and ownership** check box not selected if you do not want to preserve the original ownership and access settings for restored objects. Veeam Backup & Replication will change security settings: the user who launched the Veeam Backup & Replication console will be set as the owner of the restored object, while access permissions will be inherited from the folder to which the restored object is copied.

3. If prompted, in the **Credentials** window, specify settings of the user account to access the destination location.



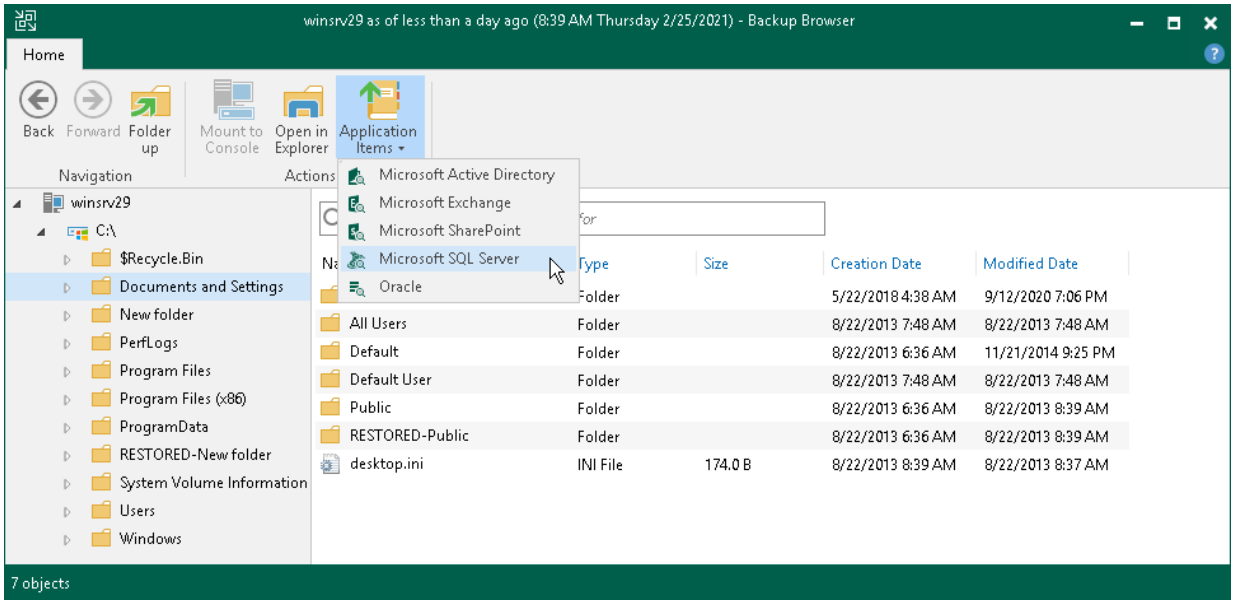
Launching Application Item Restore

If you are restoring guest OS files from VMs where the supported applications (Microsoft Active Directory, Oracle and so on) are installed, you can also launch application item restore directly from the Veeam Backup browser. To restore application items, Veeam Backup & Replication uses special tools called Veeam Explorers.

To launch application item restore, do the following:

1. On the **Home** tab, click **Application Items**.
2. Select the required application:
   - Microsoft Active Directory
   - Microsoft Exchange
   - Microsoft Share Point
   - Microsoft SQL Server
   - Oracle

3. In the opened Veeam Explorer, perform the necessary operations. For more information on Veeam Explorers, see the Veeam Explorers User Guide.



Working with Microsoft Windows File Explorer

You can use Microsoft Windows File Explorer to work with restored files and folders:

1. On the ribbon of the Veeam Backup browser, click **Mount to Console** to mount the VM disks to the Veeam Backup & Replication console.

   The VM disks are mounted to the `C:\VeeamFLR\<vmname>\<volume n>` folder.

2. To open Microsoft Windows File Explorer, do the following:

   - Click **Open in Explorer** on the Veeam Backup browser ribbon or right-click the necessary folder and select **Explorer**.
   - Click **File Explorer** in the **Start** menu of the machine where Veeam Backup & Replication console is installed. Browse to the folder where the VM disks are mounted and find the necessary files.

It is recommended that you use Microsoft Windows File Explorer only to view file content, not to restore files. For guest OS file restore, use Veeam Backup browser. This browser has the following advantages:

1. You can browse the VM guest OS file system ignoring the file system ACL settings.
2. You can preserve permissions and ownership during file-level restore.

If you open the VM file system in the Microsoft Windows Explorer, these capabilities are not available. For more information, see Microsoft Docs.

Closing Veeam Backup Browser

You can browse to VM guest OS files only while the Veeam Backup browser is open. After the Veeam Backup browser is closed, Veeam Backup & Replication unmounts VM disks from the machine where the Veeam Backup & Replication console is installed and mount server (if you have restored VM guest OS files to the original location).

It is recommended that you close the Veeam Backup browser after you have finished restoring VM guest OS files. When the Veeam Backup browser is open, the backup file whose VM guest OS file system is displayed in the browser is locked in the backup repository. As a result, some scheduled operations that use this backup file may fail.

Veeam Backup & Replication checks if there is any activity in the Veeam Backup browser with an interval of 5 minutes. If the user or Veeam Backup & Replication components and services do not perform any actions for 30 minutes, Veeam Backup & Replication displays a warning that the Veeam Backup browser is to be closed in 5 minutes.

After the warning is displayed, you can perform one of the following actions:

- You can close the Veeam Backup browser manually.
- You can click **Cancel** to postpone the close operation. In this case, the Veeam Backup browser will remain open for 5 minutes. After this period expires, Veeam Backup & Replication will display the warning again.
- You can perform no action at all. In this case, the Veeam Backup browser will be automatically closed in 5 minutes.

Restore from Linux, Unix and Other File Systems

The restore from Linux, Unix and other file systems method helps you restore files from Linux, Solaris, BSD, Novell Storage Services, Unix and Mac machines. For the list of supported OSes and file systems, see Platform Support.

You can restore files to the original or new location, access restored files using FTP or launch application item restore for the files. For more information, see Finalize Restore.

How Restore Works

When you perform guest OS file restore, Veeam Backup & Replication provides the following options for mounting VM disks from the backup or replica:

- Mounting disks to a helper host. As a helper host, you can select the target host where you want to restore files from the backup or any other Linux server. We recommend you to specify the same server to which you want to restore the files. This will improve the performance.
- Mounting disks to a helper appliance. The helper appliance is a helper VM running a stripped down Linux kernel that has minimal set of components. The appliance is quite small — around 50 MB. It requires 1024 MB RAM and takes around 10 seconds to boot.

When you perform guest OS file restore, Veeam Backup & Replication performs the following operations:

1. [If you have selected to mount disks to a helper appliance] Veeam Backup & Replication deploys the helper appliance on the ESXi host in the virtual infrastructure.
2. Veeam Backup & Replication mounts disks of a VM from the backup or replica to the host selected as a helper host or helper appliance. The backup file or VM replica itself remains in the read-only state in the backup repository or datastore.
3. Veeam Backup & Replication launches the Veeam Backup browser where mounted VM disks are displayed. You can browse the VM guest file system in the Veeam Backup browser and restore files or folders to the original VM or to another location. Also, you can enable an FTP server on the virtual appliance and allow VM owners to restore files themselves.
4. Veeam Backup & Replication connects to the VM and deploys the agent which performs the restore process. When you restore files or folders, the helper host or helper appliance connects to the VM over network or VIX API/vSphere Web Services, if a connection over the network cannot be established.

5. When you close the Veeam Backup browser or it is closed by timeout, Veeam Backup & Replication unmounts the content of the backup file or replica from the helper appliance or helper host.
6. [If you have selected to mount disks to a helper appliance] Veeam Backup & Replication unregisters the helper appliance on the ESXi host.

Requirements and Limitations

Before you restore VM guest OS files, mind the following requirements and limitations.

General Limitations

Mind the following general limitations:

- Check the supported file systems. For details, see the File-Level Restore section.
- If you plan to restore VM guest OS files to their original location, VMware Tools must be installed on the target VM.
- You cannot restore VM guest OS files from a backup or replica that does not have any successfully created restore points.
- You cannot restore files from a backup created in the reverse incremental mode if the backup job is being performed. If the backup is created in the incremental backup mode and the backup job is being performed, you can restore files from any available restore point.
- You cannot restore VM guest OS files from a running replica or if the replication job with the necessary VM is being performed.
- Veeam Backup & Replication must have access over a network to the guest OS of the target VM or direct access to the vCenter or ESXi host where the target VM resides to deploy a coordination process. The coordination process performs a number of administrative actions on the target VM guest OS, for example, collects information about mount points.
- The mount server, whether it is a helper host or a helper appliance, must have access over a network to a VM whose files you restore or direct access to vCenter or ESXi host where the VM resides. If a connection over the network cannot be established, the mount server connects to the VM over VIX API/vSphere Web Services. In this case, you must use a root account for a target VM, otherwise the restore process will fail.

  If you use the FQDN of the ESXi host in the helper appliance configuration window, the helper appliance must be able to resolve the FQDN of the ESXi host.

  If your DHCP does not provide the DNS configuration, assign the DNS server address manually using the registry `[HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication\DnsServersList]` and restart the restore process.
-
- Veeam Backup & Replication uses the ICMP `ping` command to define whether a VM is available over network. If the VM must be available over the network, check that ICMP protocol is enabled on the VM.
- The multi-OS file-level restore wizard does not support restore of deduplicated volumes (for example, Microsoft Windows volumes with data deduplication enabled).
- You can restore from Novell Storage Services (NSS) only if you use the helper appliance option.

Requirements and Limitations for Helper Host

If you select to mount disks to a helper host, mind the following:

- You can restore from ZFS using the helper host option if zpool is configured on the specified helper host.
- The helper host OS kernel must support the file system that you plan to mount on this host. Otherwise, the mount will be refused and in rare cases it may cause kernel panic.
- For the helper host option, mounting of LVM snapshots is not supported. Thus, LVM snapshots are skipped from processing.
- If you want to restore files from a Btrfs disk and select the original host as a helper host, the mount of the Btrfs disk will fail. The issue occurs due to restriction of mounting 2 Btrfs disks with identical IDs to the same machine. To avoid this issue, use the helper appliance option.
- Hardened repositories cannot be selected as helper hosts.

Requirements and Limitations for Target Linux VM

If you want to save files to a new location, mind the following for the target Linux VM:

- If you want to restore files over network, make sure that the SSH daemon is configured and SCP utility is available on the target VM.
- Veeam Backup & Replication can restore ACL for recovered VM guest OS files. To let Veeam Backup & Replication detect the target Linux system architecture and kernel version, the following utilities must be present in the minimal configuration of the system: *arch* and *uname*.

Linux Firewalls in Helper Host and Target Linux VM

If you want to use the helper host option or if you want to restore files to a new location on another Linux VM, mind the following.

Veeam Backup & Replication automatically opens ports used for the restore process on the helper host and on the target Linux VM. Generally, Veeam Backup & Replication automatically opens ports for most 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.

Restoring VM Guest OS Files (Multi-OS)

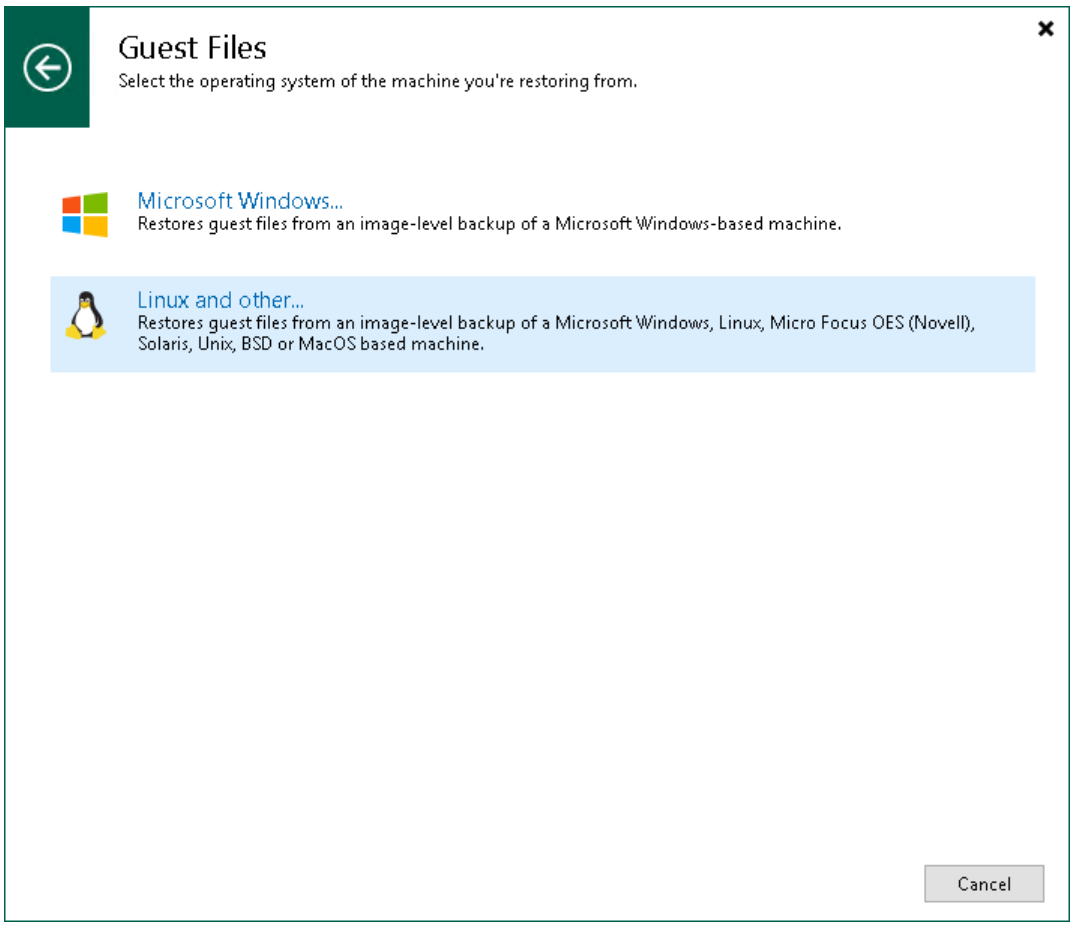
To restore VM guest OS files and folders from Linux or Unix-based machines, use the **Guest File Restore** wizard.

Step 1. Launch Guest File Restore Wizard

To launch the **Guest File Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Guest files restore** > **Linux and other**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **Guest files** > **Linux and other** on the ribbon. Alternatively, right-click the VM whose files you want to restore and select **Restore guest files** > **Linux and other**.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Guest files (Linux and other)**.

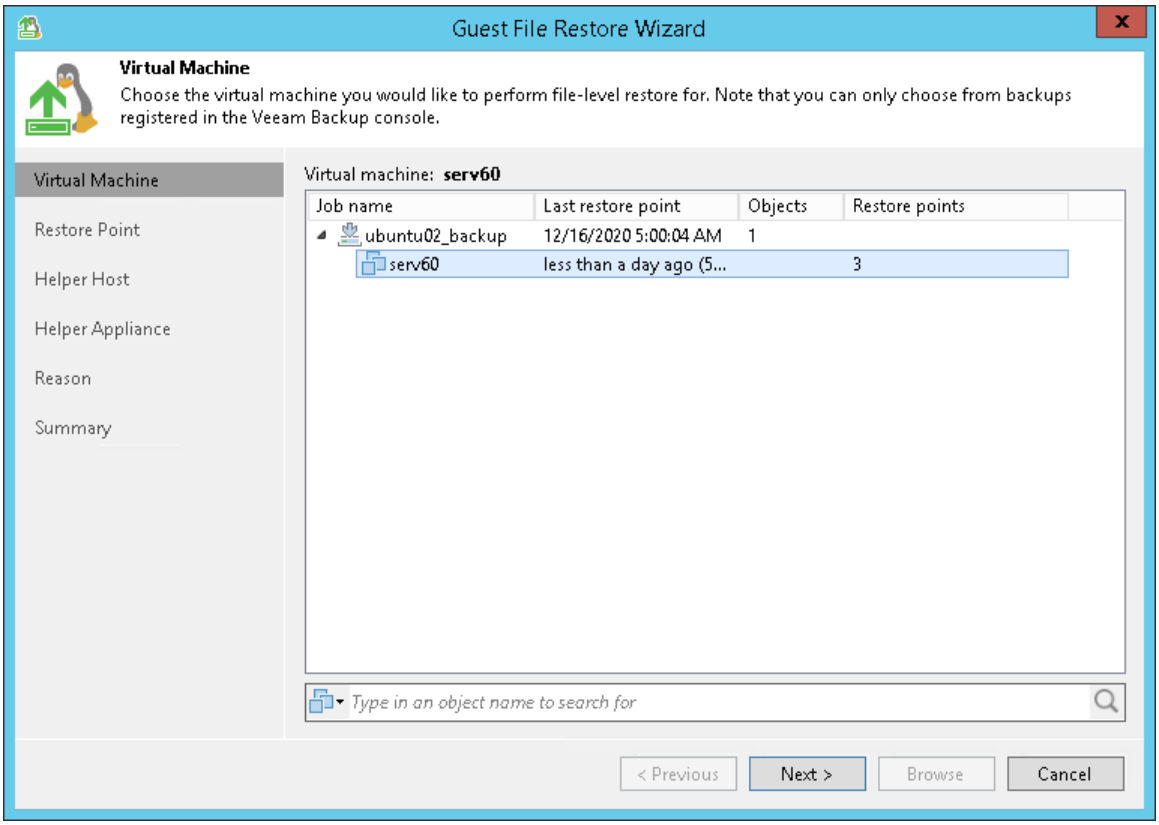
  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.



Step 2. Select VM

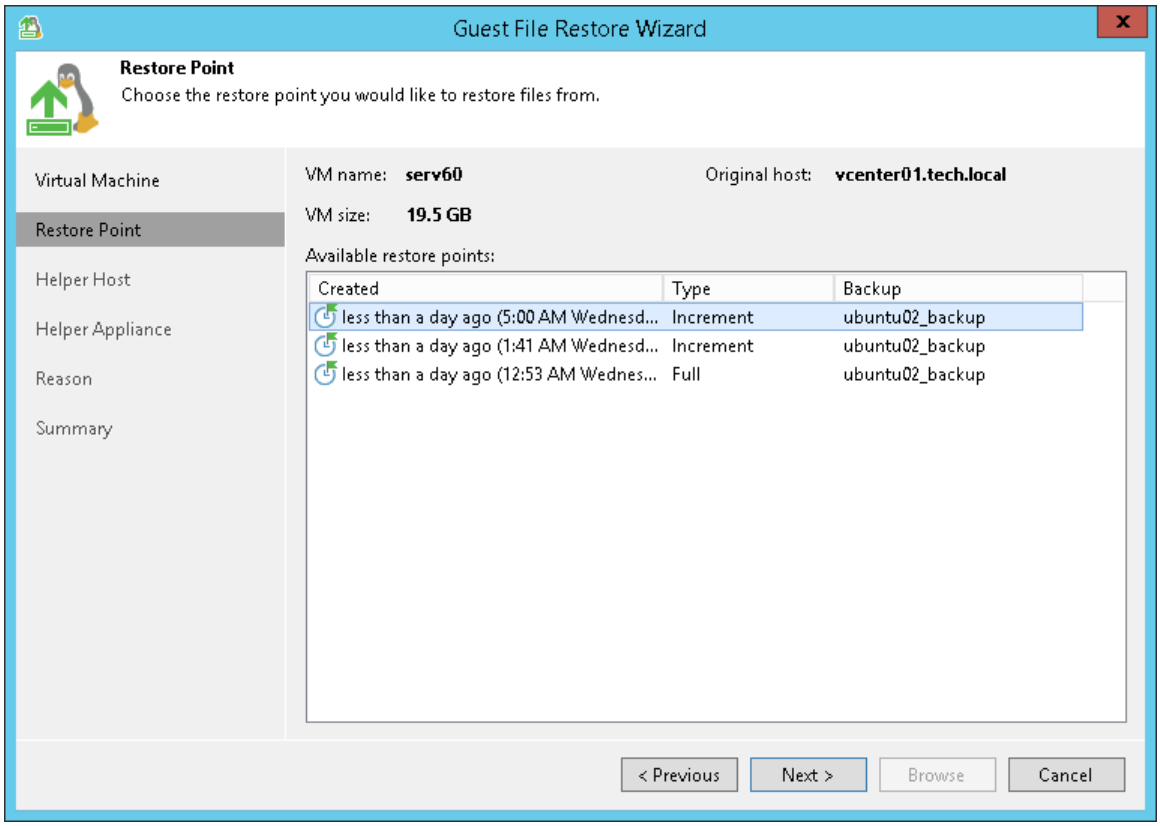
At the **Virtual Machine** step of the wizard, select the VM whose guest OS files you want to restore:

1. In the **Virtual machine** list, expand the necessary backup.
2. Select the VM.



Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select the restore point from which you want to restore VM guest OS files.



Step 4. Specify Helper Host

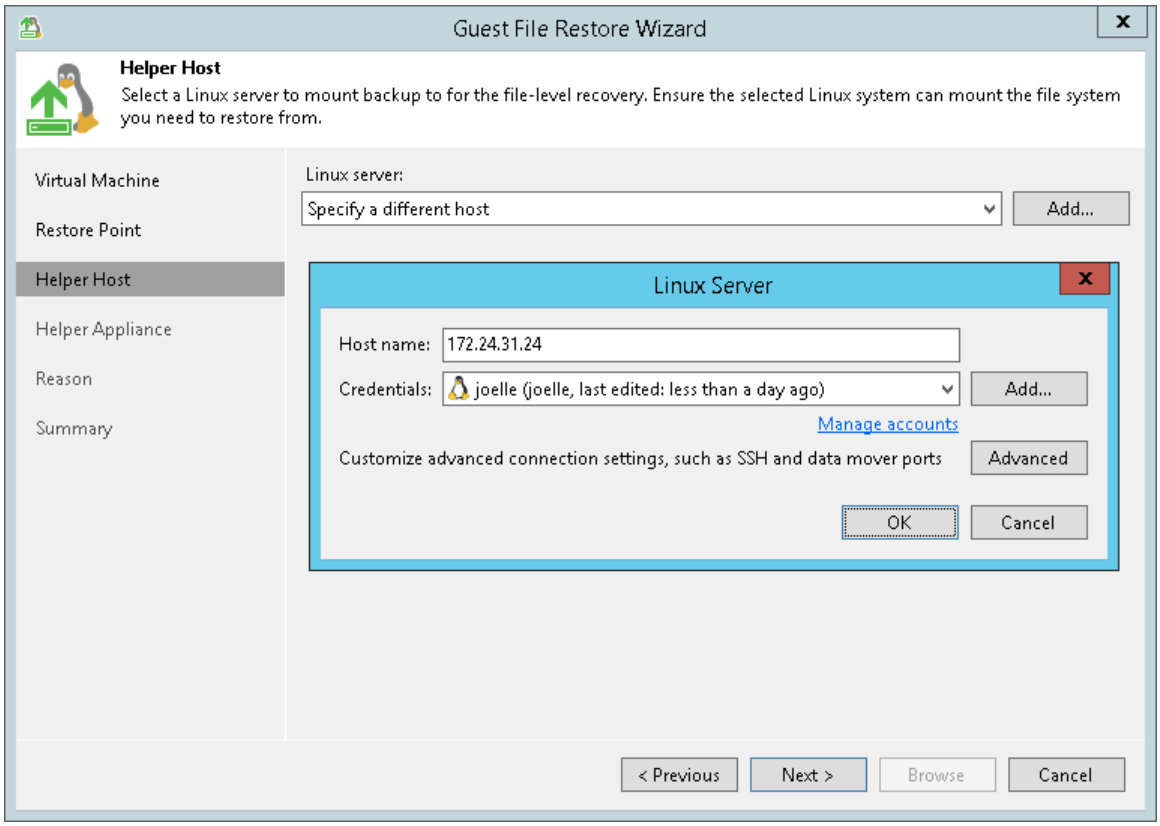
At the **Helper Host** step of the wizard, you can select a machine that will be used as a helper host. Veeam Backup & Replication will mount disks of a machine from the backup or replica to the selected helper host. Alternatively, you can choose to configure a new helper appliance where Veeam Backup & Replication will mount VM disks. For details on helper hosts and appliances, see Restore from Linux, Unix and Other File Systems.

To specify a helper host, do the following:

1. From the **Linux server** list, select one of the following options:
   - *<Hostname>* — to specify a Linux server as the helper host. The wizard displays all Linux servers added the Veeam Backup & Replication infrastructure.
   - **Specify a different host** — to specify a Linux server that is not added to the Veeam Backup & Replication infrastructure.
   - **Use a temporary helper appliance** — to use a temporary helper appliance. If you select this option, the wizard will include a step for configuring the helper appliance.
2. [For the **Specify different host** option] In the **Linux Server** window, specify the helper host name and connection setting:
   a. In the **Host name** field, specify the IP address or the host name of the Linux server which will be used as a helper host.
   b. From the **Credentials** list, select an account that will be used to connect to the helper host.

      If you have not added the account before, click the **Add** button on the right of the **Credentials** list and follow the instructions described in Linux Accounts.
   c. If you want to change the default SSH settings, click **Advanced** and configure SSH settings as described in Specify Credentials and SSH Settings.

d. Click **OK** to save the helper host settings.



Step 5. Specify Location for Helper Appliance

The **Helper Appliance** step of the wizard is available only if you have selected the **Use a temporary helper appliance** option at the Specify Helper Host step.

Select an ESXi host on which you want to place the helper appliance:

1. At the bottom of the window, click **Customize**.
2. In the **Host** field, specify the IP address or FQDN of the ESXi host on which the helper appliance must be registered.
3. In the **Resource pool** field, specify a resource pool to which the helper appliance must be placed.
4. Select a network for the helper appliance:
   a. On the right of the **Network** field, click **Choose**.

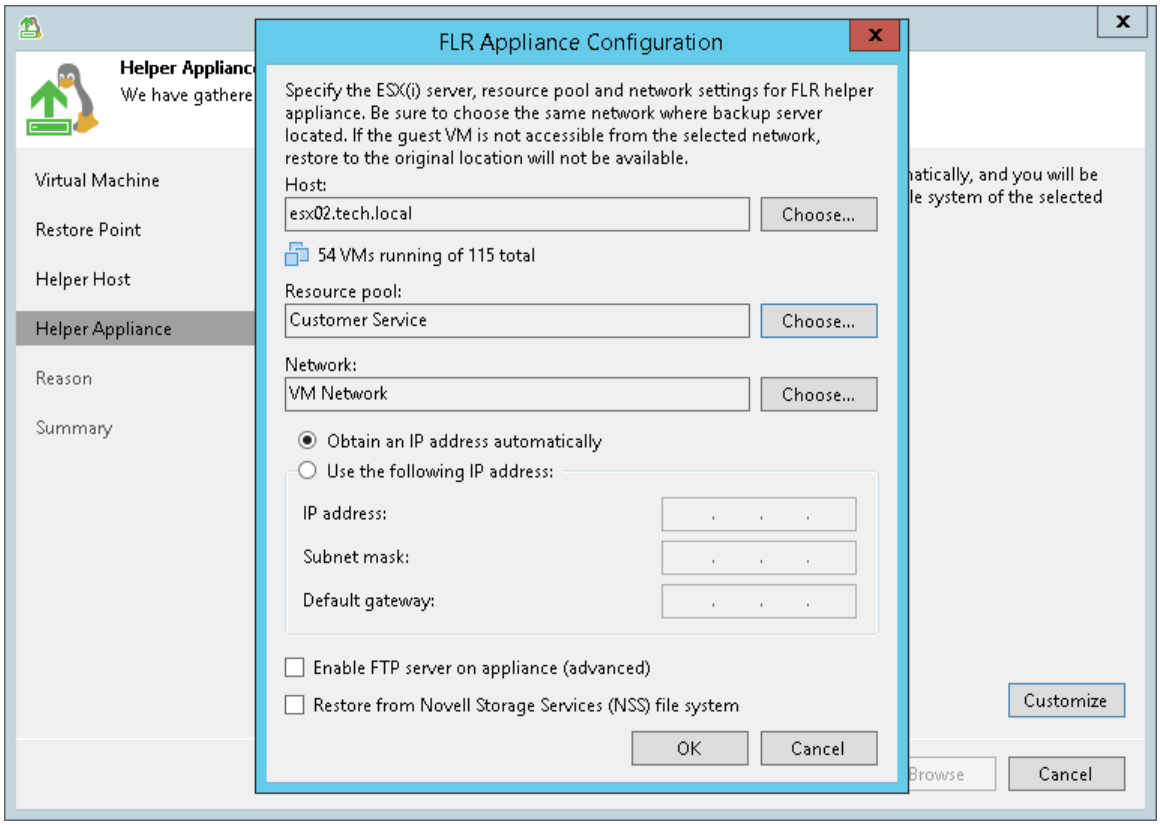
      In the **Select Network** window, Veeam Backup & Replication will display a list of networks to which the specified host is connected.

   b. From the **Networks** list, select a network to which the helper appliance must be connected and click **OK**.

      Mind that the backup server and the mount server must have access to the helper appliance over the network.
5. Specify IP addressing settings for the helper appliance:
   - If you use a DHCP server in the network, leave the **Obtain an IP address automatically** option selected.
   - To manually assign the specific IP address to the helper appliance, select the **Use the following IP address** option and specify the IP address, subnet mask and default gateway address.
6. To enable FTP access to the restored file system, select the **Enable FTP server on appliance (advanced)** check box. As a result, users will be able to access the helper appliance over FTP, browse the file system of the restored VM and download necessary files on their own.
7. If you are performing restore of a VM with the Novell Storage Services file system, select the **Restore Novell Storage Services file system** check box. Veeam Backup & Replication will deploy a specific appliance that supports the Novell Storage Services file system.

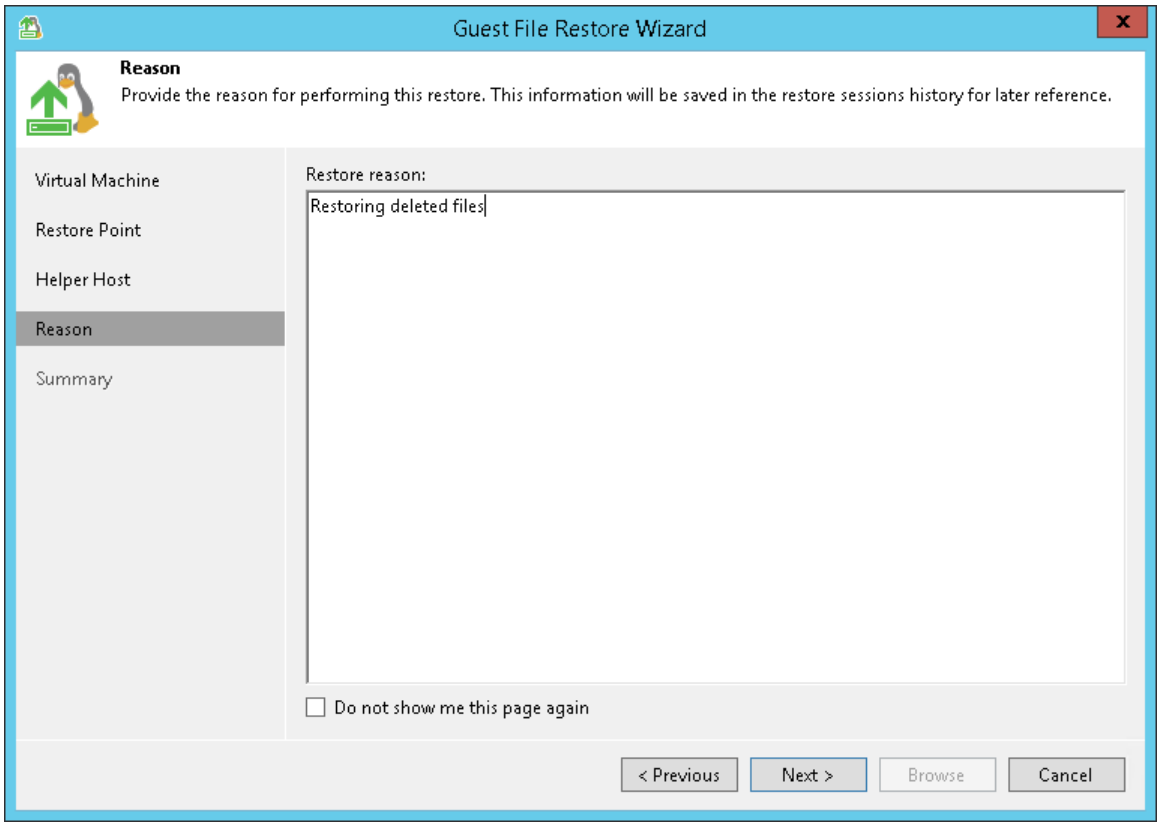
IMPORTANT

When choosing an ESXi host for the helper appliance used for file-level restore from the Novell Storage Services file system, make sure that it allows running VMs with 64-bit guest OSes.



Step 6. Specify Restore Reason

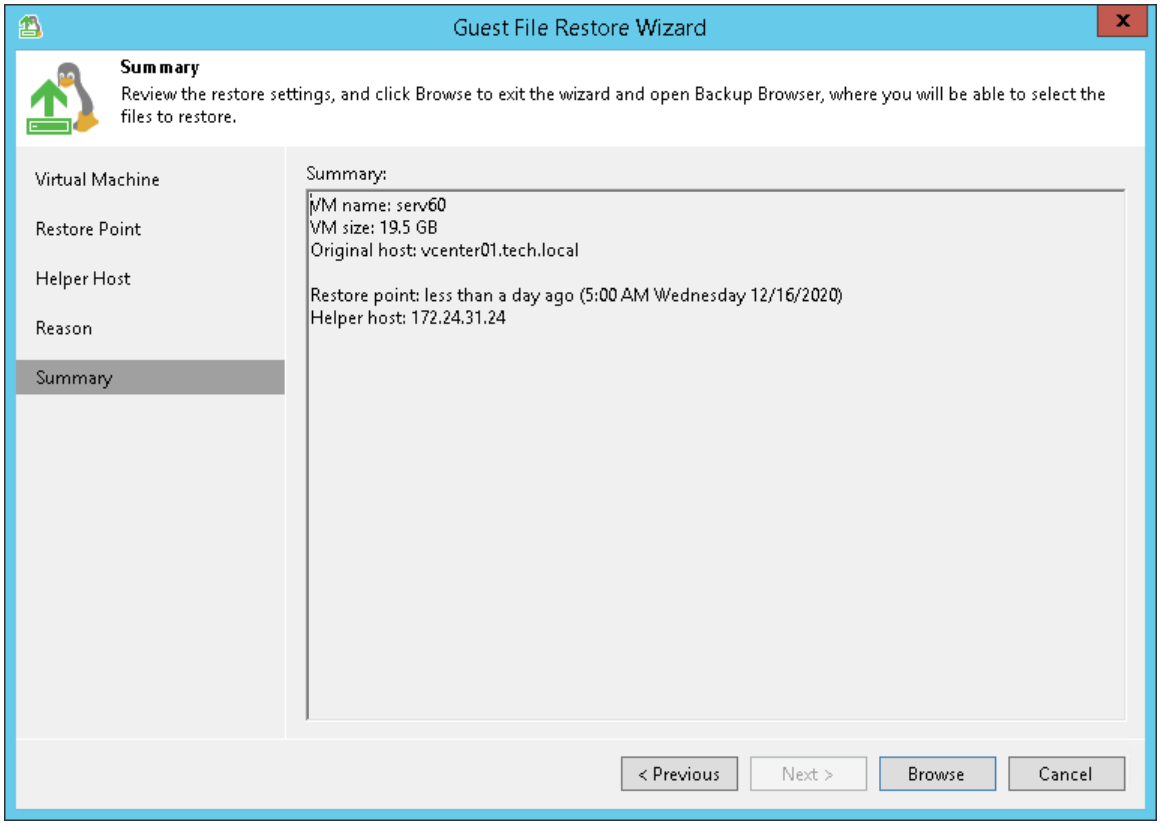
At the **Reason** step of the wizard, enter a reason for restoring VM guest OS files. The information you provide will be saved in the session history and you can reference it later.



Step 7. Verify Restore Settings

At the **Summary** step of the wizard, review the restore settings and click **Browse** to open the Veeam Backup browser.

If you have selected to mount disks to a helper appliance, it may take about 10-40 seconds to boot the helper appliance and open the browser.



Step 8. Finalize Restore

After the wizard is closed, Veeam Backup & Replication opens the Veeam Backup browser displaying the file system tree of the restored VM.

In the Veeam Backup browser, you can perform the following operations:

- Restore files to the original location
- Restore files to a new location
- Access files over FTP
- Access helper appliance logs

> **NOTE**
>
> You can browse the VM guest OS files and access restored files on the FTP only while the Veeam Backup browser with the restored files is open. After the Veeam Backup browser is closed, Veeam Backup & Replication unmounts the VM disks from the helper appliance and removes helper appliance from the ESXi host.

Restoring Files to Original Location

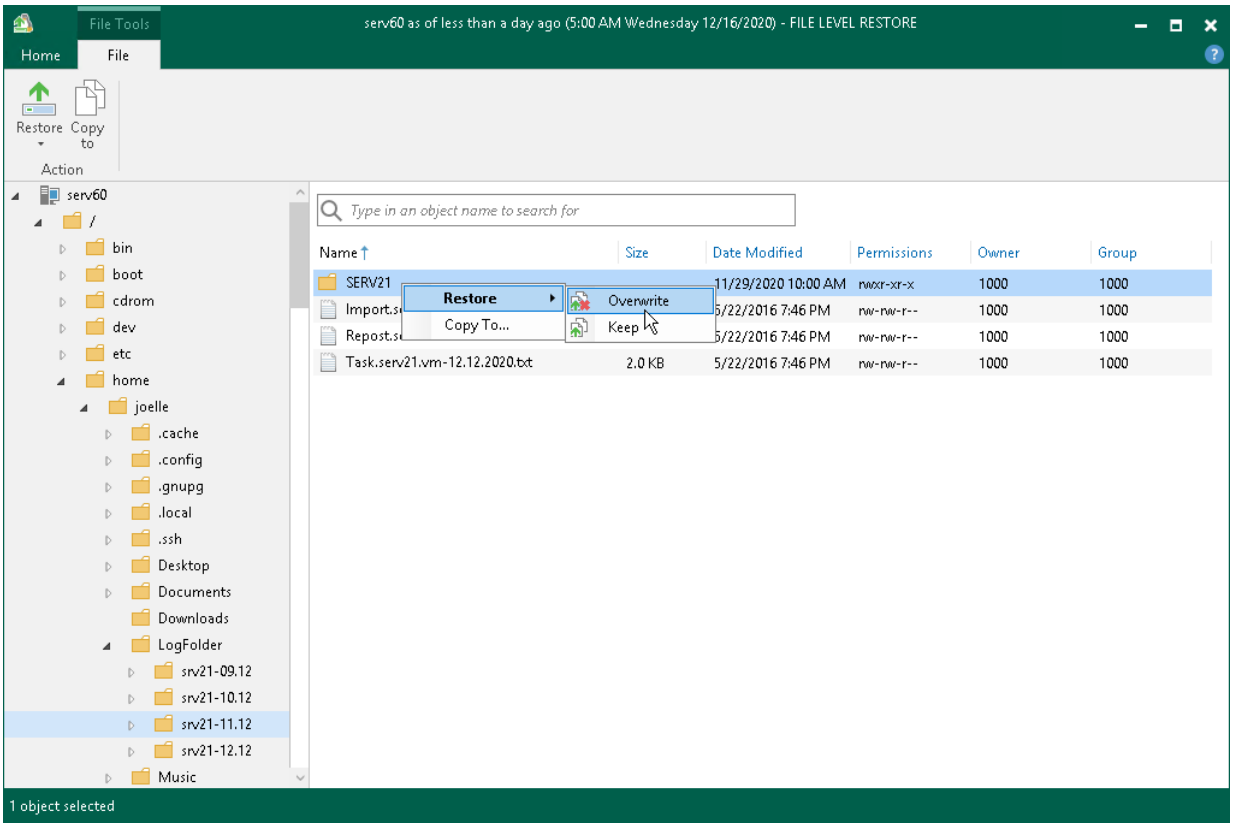
To restore files and folders to the original location, right-click the necessary file or folder in the file system tree or in the details pane on the right and select one of the following commands:

- To overwrite the original file on the VM guest OS with the file restored from the backup, select **Restore** > **Overwrite**.
- To save the file restored from the backup next to the original file, select **Restore** > **Keep**.

  Veeam Backup & Replication will add the *.RESTORED-YYYYMMDDHHMMSS* suffix to the original file name and store the restored file in the same folder where the original file resides.

To restore files to the original location, Veeam Backup & Replication uses the account for VM guest OS access specified in the backup job settings. If this account does not have sufficient rights to access the target VM, you will be prompted to enter credentials. In the **Credentials window**, specify a user account to access the destination location (server or shared folder). For more information on adding credentials, see Credentials Manager.

In some cases, you may remove the original VM and restore it from the backup by the time of file-level restore. If you then attempt to restore VM guest OS files to the original location, Veeam Backup & Replication will not be able to find the original VM by its reference ID, and display a warning. Click **OK** and browse to the target VM in the virtual infrastructure to which you want to restore VM guest OS files.



Saving Files to New Location

To save files and folders to a new location:

1. Right-click the necessary file or folder and select **Copy to**.
2. In the **Select Destination** window, select the necessary destination:
   - To recover files to a server already added to the Veeam Backup & Replication infrastructure, select the server from the **Server** drop-down list and then specify the path to a folder where files will be recovered.
   - To recover files to a Linux server that is not added to the Veeam Backup & Replication infrastructure, select **Specify a different host** from the **Server** drop-down list and follow the steps of the wizard to add a Linux server that will be used as a target host. The server will be added ad-hoc. Then specify the path to a folder where files will be recovered.

     The server you add ad-hoc will not appear in the list of managed hosts in Veeam Backup & Replication: its purpose is to host the files that you recover. It will only remain visible in the Veeam Backup browser until all currently active file-level restore sessions are completed.

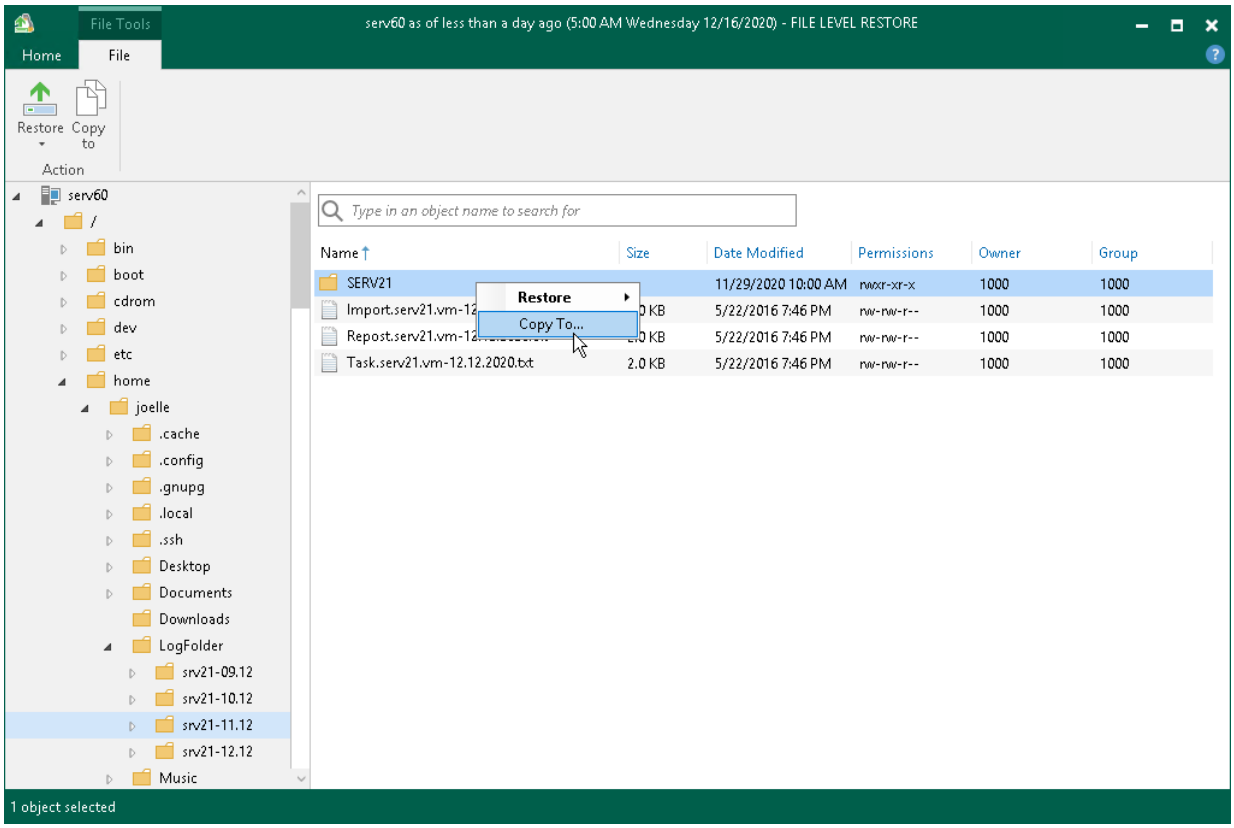
     > **IMPORTANT**
     >
     > Only Linux servers can be added ad-hoc.

   - To recover files to a shared folder, specify a path to the destination folder in the **Path to folder** field.

3. If you want to preserve original permissions and ownership for recovered files, select the **Preserve permissions and ownership** check box in the **Select Destination** window.

> **IMPORTANT**
>
> To restore original permissions and ownership settings, the user account you have specified must have privileges to change the owner on the selected server or shared folder.

4. If prompted, in the **Credentials** window, specify settings of the user account to access the destination location.



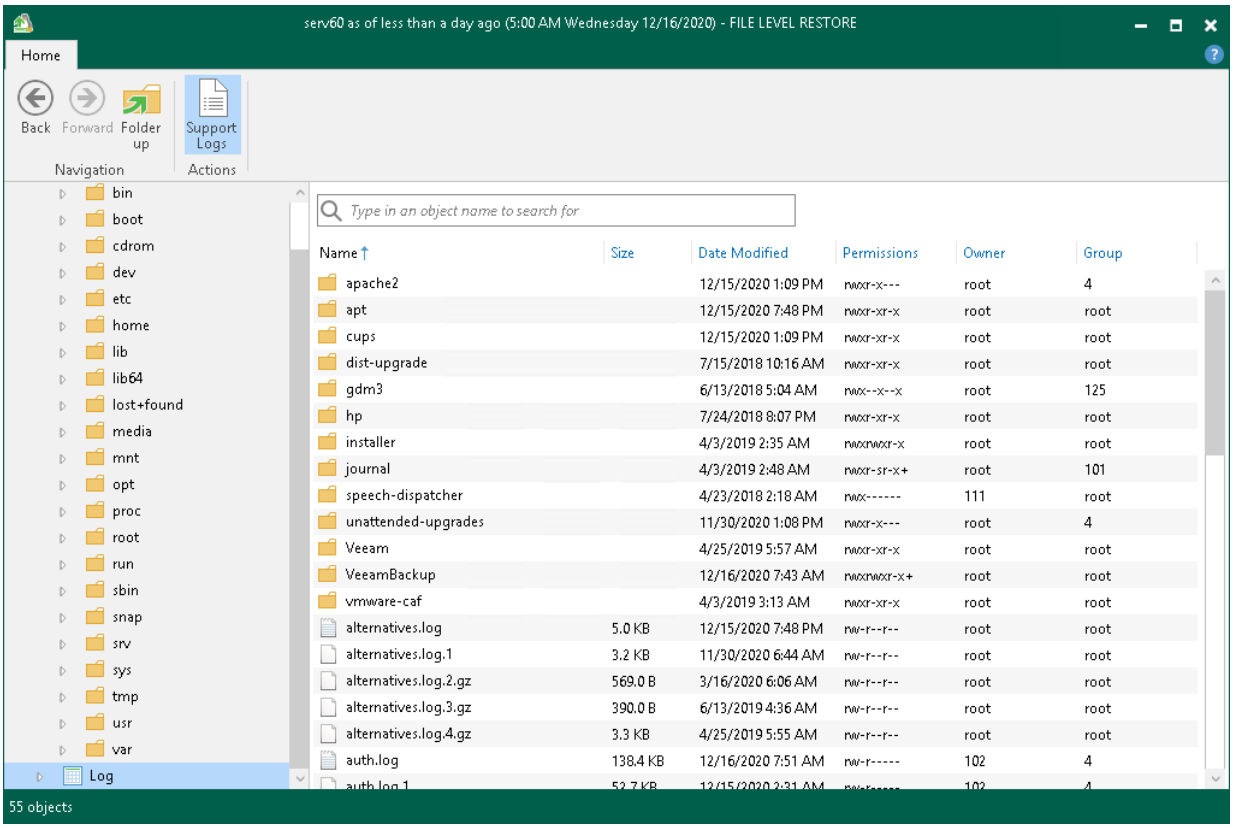
Accessing Files over FTP

If you have chosen to enable FTP server on the helper appliance, the restored file system will also be available over FTP at *ftp://<FLR_appliance_IP_address>*. Other users in the same network can access the helper appliance to restore the files they need.

Accessing the appliance over FTP requires credentials. Use the Guest OS helper appliance credentials specified in managed credentials. If the password has not been updated, refer to the following knowledge base article: this Veeam KB article.

Accessing Helper Appliance Logs

If you need to access logs of the helper appliance, click **Support Logs** on the Veeam Backup browser ribbon. Veeam Backup & Replication will show the **Log** node under the file system tree. To hide this node, click **Support Logs** once again.



Restore from Other File Systems

You can restore file systems other than those listed in the File-Level Restore section. For this purpose, use the Instant Disk Recovery technology.

To restore files and folders, do the following:

1. In the virtual infrastructure, find a VM that can read the file system of the original VM.
2. Use Instant Disk Recovery to mount a disk of the original VM to the VM that can read the file system.
3. Open the VM console, bring the disk to the online mode.
4. Copy the required files.
5. If you no longer need the disk, stop publishing it.

Alternatively, you can mount the VM disks to a Microsoft Windows VM and use file management tools.

Viewing File Restore Session Statistics

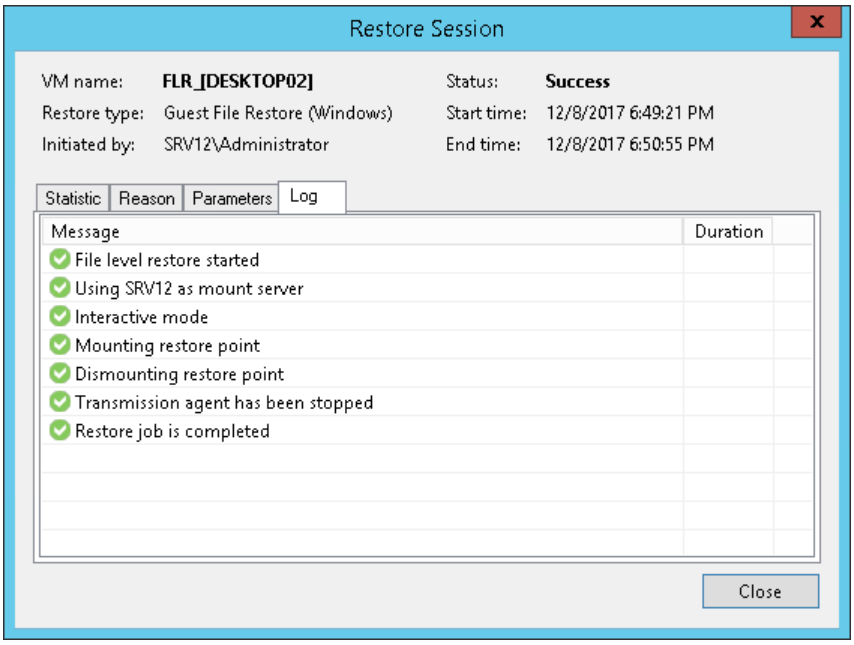
You can view statistics about performed guest OS file restore sessions.

To view the 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**.

The file restore statistics provides detailed data on file restore sessions:

- At the top of the **Restore Session** window, Veeam Backup & Replication shows general session statistics: a name of the machine whose guest OS files are restored during the session, a user name of the account under which the session was started, session status and duration details.
- The **Statistics** tab shows detailed information about the files restored during the session.
- The **Reason** tab shows the reason for the guest OS file restore that was specified at the **Reason** step of the **File Level Restore** wizard.
- The **Parameters** tab shows information about the restore point selected for the guest OS file restore at the **Restore Point** step of the **File Level Restore** wizard.
- The **Log** tab shows a list of operations performed during the session.



Application Item Restore

Veeam Backup & Replication provides auxiliary tools called Veeam Explorers that allow you to restore application items directly from backups or replicas. These backups and replicas must be created with the enabled application-aware processing option.

You can restore items of the following applications:

- Microsoft Active Directory
- Microsoft SQL Server
- Oracle
- Microsoft Exchange
- Microsoft SharePoint
- Microsoft OneDrive for Business
- Microsoft Teams

Veeam Explorers are distributed as part of Veeam Backup & Replication, and you do not need to install them separately. You also do not need to purchase any additional license to use Veeam Explorers. For more general information, see the Veeam Explorers Overview section of the Veeam Explorers User Guide.

If you want to restore application items from storage snapshots, see the Application Item Restore from Storage Snapshots section in Veeam Backup & Replication User Guide.

Restoring Microsoft Active Directory Items

To restore Microsoft Active Directory items from a backup or replica, Veeam Backup & Replication uses Veeam Explorer for Microsoft Active Directory.

> **IMPORTANT**
>
> Before you restore Microsoft Active Directory items, check requirements and limitations listed in the Planning and Preparation section in the Veeam Backup Explorers User Guide. One of the main requirements is that the backup or replica must be created with application-aware processing enabled.

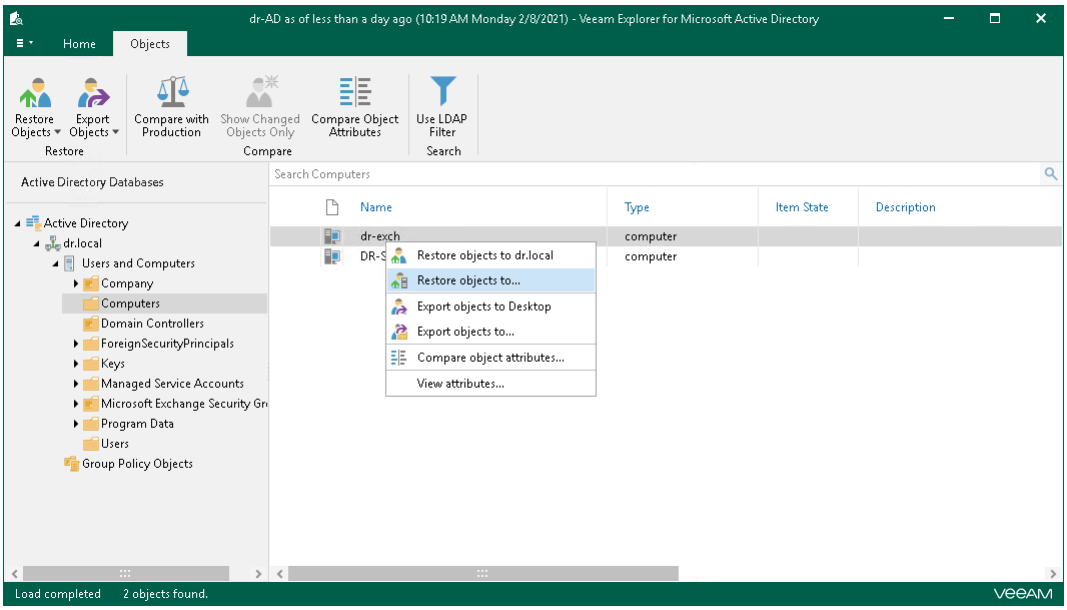
To launch Veeam Explorer for Microsoft Active Directory from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** or **Replicas** node.
3. In the working area, select the necessary machine in the backup or VM replica and click **Application Items > Microsoft Active Directory** on the ribbon. Alternatively, right-click the machine or VM replica and select **Restore application items** > **Microsoft Active Directory objects**.
4. In the opened **Microsoft Active Directory Object Restore** wizard, select a restore point from which you want to restore application items.

Then Veeam Backup & Replication opens Veeam Explorer for Microsoft Active Directory. In this explorer, you can browse, search, restore application items and so on. The detailed information is provided in the Veeam Explorer for Microsoft Active Directory main section of the Veeam Backup Explorers User Guide. For example:

- Information about viewing and searching objects in the backup file is in the Browsing, Searching and Viewing Items section.

- Information about adding a database manually is in the Adding Standalone Databases section.
- Information about restoring process is in the Data Restore section.



Restoring Microsoft SQL Server Items

To restore databases from a backup or replica of a Microsoft SQL Server, Veeam Backup & Replication uses Veeam Explorer for Microsoft SQL Server.

> **IMPORTANT**
>
> Before you restore Microsoft SQL databases, check requirements and limitations listed in the Planning and Preparation section in the Veeam Backup Explorers User Guide. One of the main requirements is that the backup or replica must be created with application-aware processing enabled.

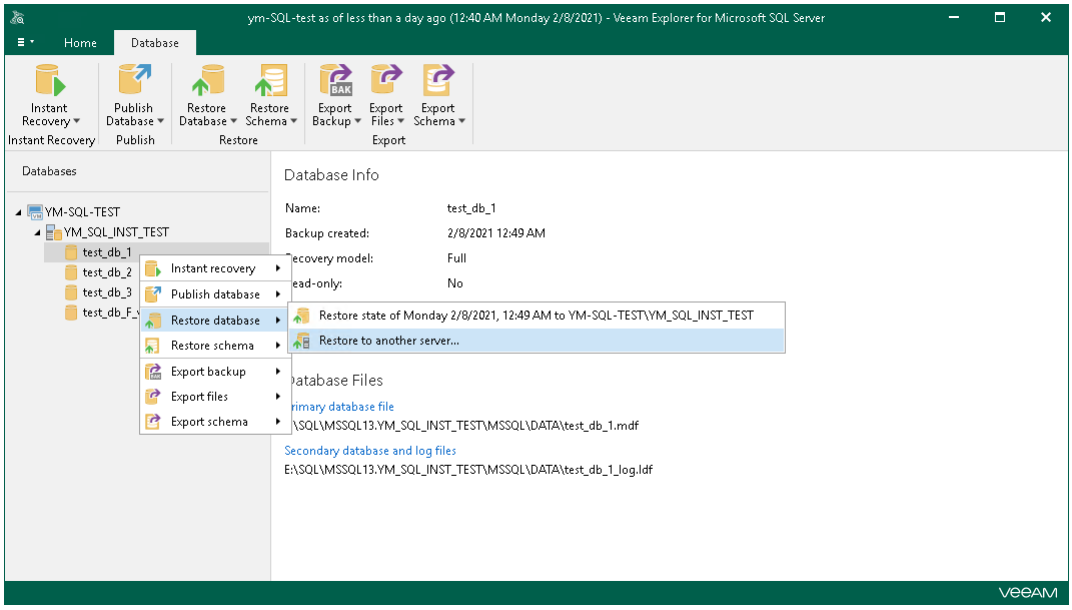
To launch Veeam Explorer for Microsoft SQL Server from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** or **Replicas** node.
3. In the working area, select the necessary machine in the backup or VM replica and click **Application Items** > **Microsoft SQL Server** on the ribbon. Alternatively, right-click the machine or VM replica and select **Restore application items** > **Microsoft SQL Server databases**.
4. In the opened **Microsoft SQL Server Database Restore** wizard, select a restore point from which you want to restore databases.

Then Veeam Backup & Replication opens Veeam Explorer for Microsoft SQL Server. In this explorer, you can view, restore databases and so on. The detailed information is provided in the Veeam Explorer for Microsoft SQL Server main section of the Veeam Backup Explorers User Guide. For example:

- Information about database is in the Viewing Database Information section.
- Information about restoring process is in the Data Restore section.

- Information about adding a database manually is in the Adding Standalone Databases section.



Restoring Oracle Items

To restore databases from a backup or replica of an Oracle system, Veeam Backup & Replication uses Veeam Explorer for Oracle.

> **IMPORTANT**
>
> Before you restore Oracle databases, check requirements and limitations listed in the Planning and Preparation section in the Veeam Backup Explorers User Guide. One of the main requirements is that the backup or replica must be created with application-aware processing enabled.

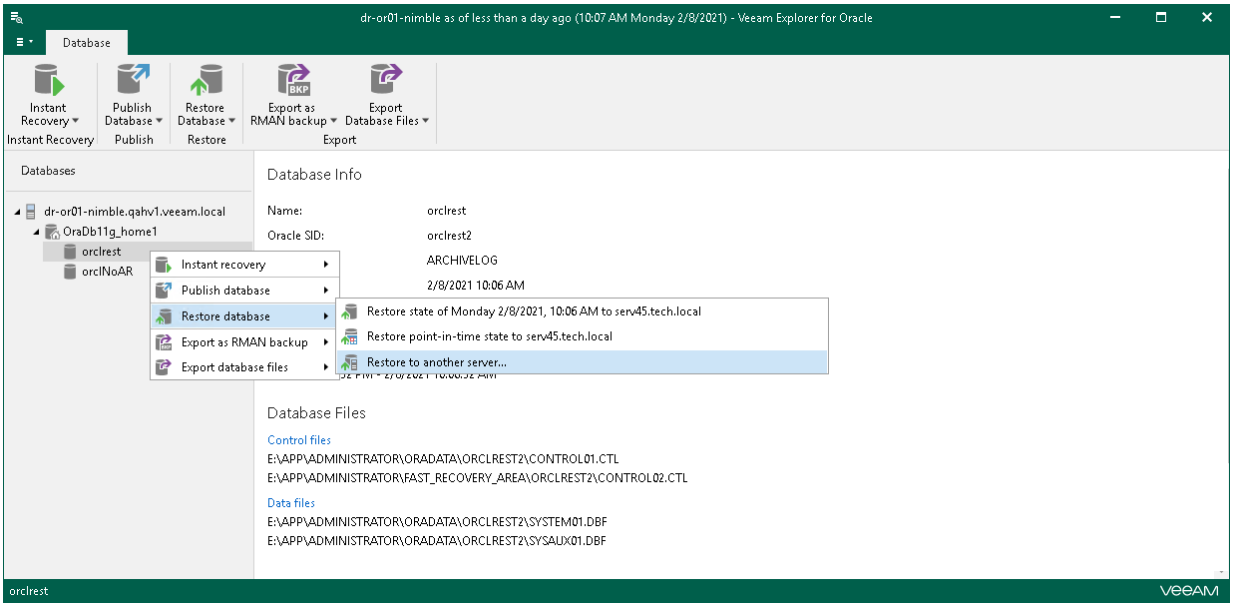
To launch Veeam Explorer for Oracle from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** or **Replicas** node.
3. In the working area, select the necessary machine in the backup or VM replica and click **Application Items** > **Oracle** on the ribbon. Alternatively, right-click the machine or VM replica and select **Restore application items** > **Oracle databases**.
4. In the opened **Oracle Database Restore** wizard, select a restore point from which you want to restore databases.

Then Veeam Backup & Replication opens Veeam Explorer for Oracle. In this explorer, you can view, restore databases and so on. The detailed information is provided in the Veeam Explorer for Oracle main section of the Veeam Backup Explorers User Guide. For example:

- Information about database is in the Viewing Database Information section. Oracle databases cannot be added manually.

- Information about restoring process is in the Data Restore section.



Restoring Microsoft Exchange Items

To restore Microsoft Exchange items from a backup or replica of a Microsoft Exchange Server, Veeam Backup & Replication uses Veeam Explorer for Microsoft Exchange.

> **IMPORTANT**
>
> Before you restore Microsoft Exchange items, check requirements and limitations listed in the Planning and Preparation section in the Veeam Backup Explorers User Guide. One of the main requirements is that the backup or replica must be created with application-aware processing enabled.

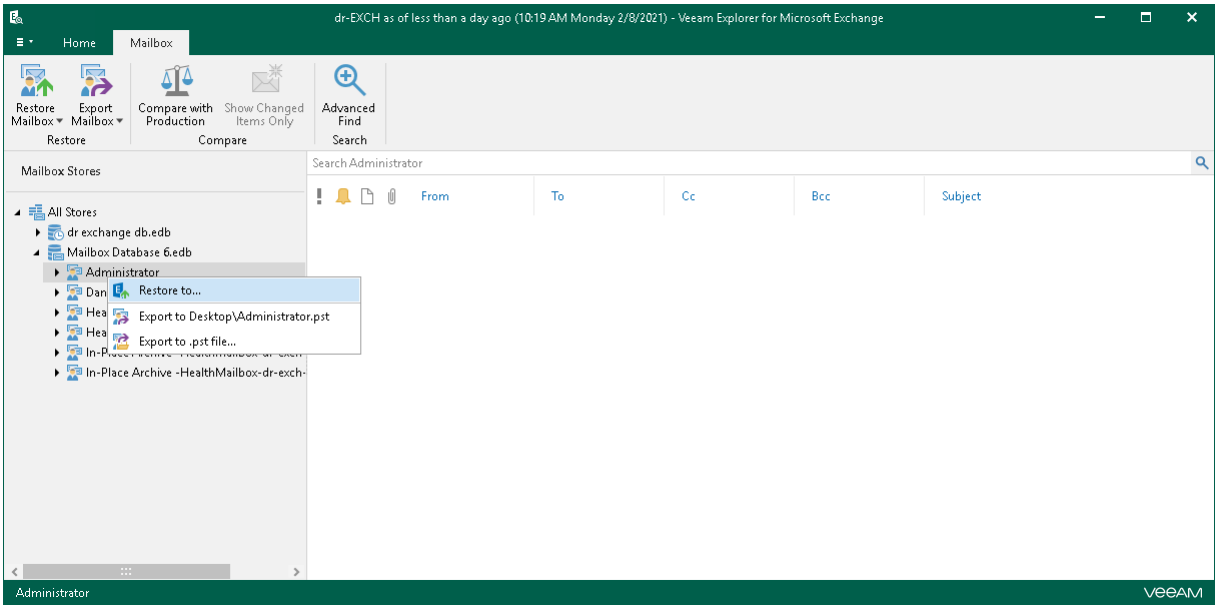
To launch Veeam Explorer for Microsoft Exchange from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** or **Replicas** node.
3. In the working area, select the necessary machine in the backup or VM replica and click **Application Items** > **Microsoft Exchange** on the ribbon. Alternatively, right-click the machine or VM replica and select **Restore application items** > **Microsoft Exchange mailbox items**.
4. In the opened **Microsoft Exchange Item Level Restore** wizard, select a restore point from which you want to restore application items.

Then Veeam Backup & Replication opens Veeam Explorer for Microsoft Exchange. In this explorer, you can view, browse, restore Microsoft Exchange items and so on. The detailed information is provided in the Veeam Explorer for Microsoft Exchange main section of the Veeam Backup Explorers User Guide. For example:

- Information about viewing and searching objects in the backup file is in the Browsing, Searching and Viewing Items section.
- Information about adding a database manually is in the Standalone Databases Management section.

- Information about restoring process is in the Data Restore section.



Restoring Microsoft SharePoint Items

To restore Microsoft SharePoint items from a backup or replica of a Microsoft SharePoint Server, Veeam Backup & Replication uses Veeam Explorer for Microsoft SharePoint.

> **IMPORTANT**
>
> Before you restore Microsoft SharePoint items, check requirements and limitations listed in the Planning and Preparation section in the Veeam Backup Explorers User Guide. One of the main requirements is that the backup or replica must be created with application-aware processing enabled.

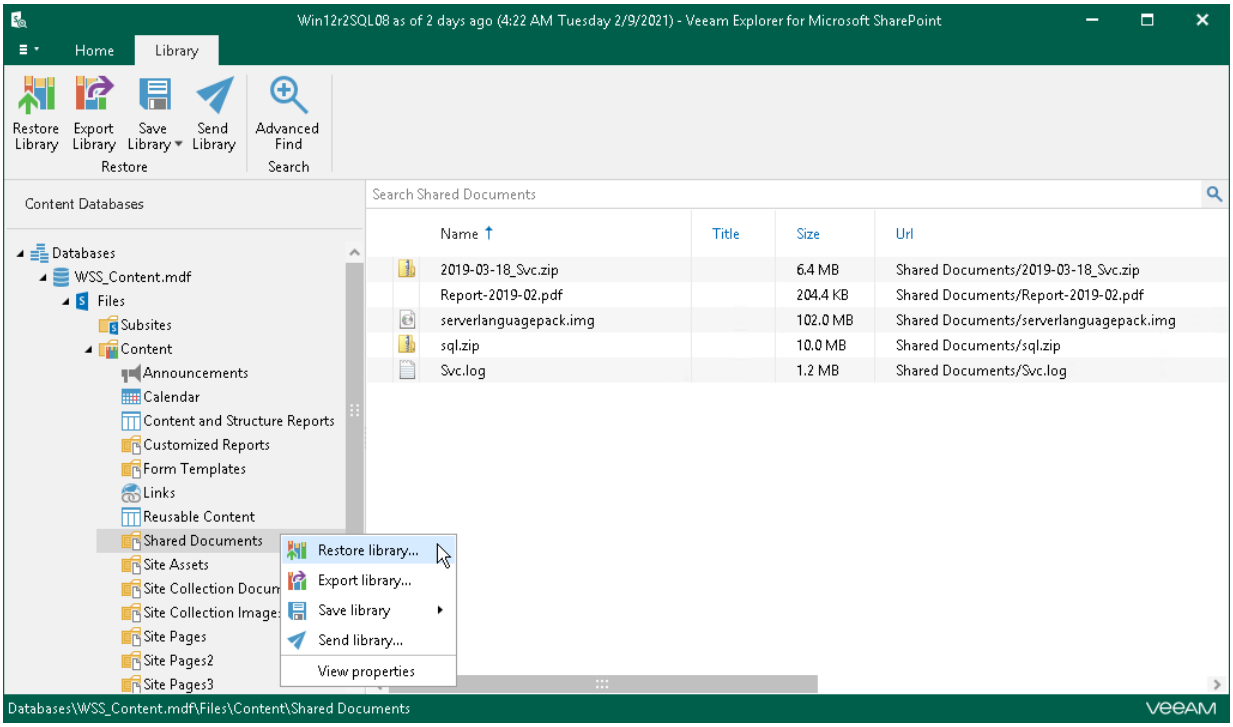
To launch Veeam Explorer for Microsoft SharePoint from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** or **Replicas** node.
3. In the working area, select the necessary machine in the backup or VM replica and click **Application Items** > **Microsoft SharePoint** on the ribbon. Alternatively, right-click the machine or VM replica and select **Restore application items** > **Microsoft SharePoint content**.
4. In the opened **Microsoft SharePoint Item Restore** wizard, select a restore point from which you want to restore application items.

Then Veeam Backup & Replication opens Veeam Explorer for Microsoft SharePoint. In this explorer, you can view, search, restore Microsoft SharePoint items and so on. The detailed information is provided in the Veeam Explorer for Microsoft SharePoint main section of the Veeam Backup Explorers User Guide. For example:

- Information about viewing and searching objects in the backup file is in the Browsing, Searching and Viewing Items section.
- Information about adding a database manually is in the Standalone Databases Management section.

- Information about restoring process is in the Data Restore section.



Restoring Microsoft OneDrive for Business Items

To restore Microsoft OneDrive for Business data from a backup or replica of a Veeam Backup for Microsoft Office 365 server, Veeam Backup & Replication uses Veeam Explorer for Microsoft OneDrive for Business.

> **IMPORTANT**
>
> Before you restore Microsoft OneDrive for Business data, check requirements and limitations listed in the Planning and Preparation section in the Veeam Backup Explorers User Guide. One of the main requirements is that the backup or replica must be created with application-aware processing enabled.

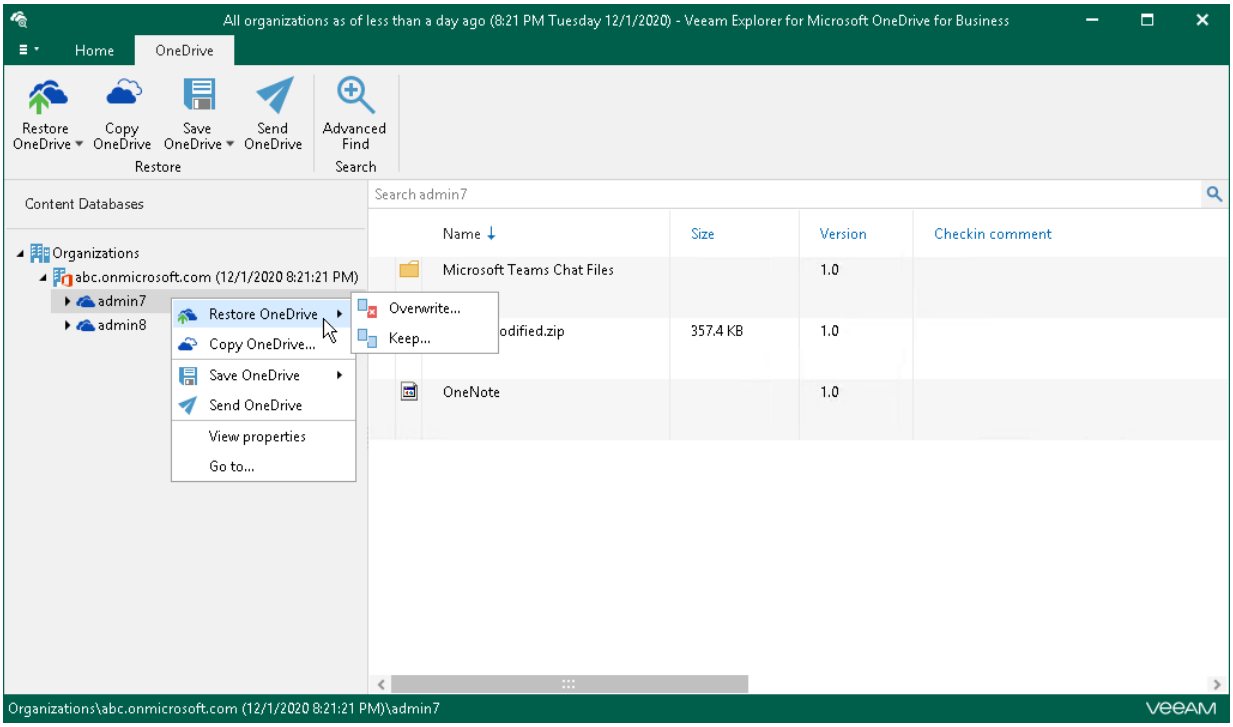
To launch Veeam Explorer for Microsoft OneDrive for Business from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** or **Replicas** node.
3. In the working area, select the necessary machine in the backup or VM replica and click **Application Items** > **Microsoft OneDrive for Business** on the ribbon. Alternatively, right-click the machine or VM replica and select **Restore application items** > **Microsoft OneDrive for Business files**.
4. In the opened **Microsoft OneDrive for Business Files** wizard, select a restore point from which you want to restore application items.

Then Veeam Backup & Replication opens Veeam Explorer for Microsoft OneDrive for Business. In this explorer, you can browse, view, restore Microsoft OneDrive for Business data and so on. The detailed information is provided in the Veeam Explorer for Microsoft OneDrive for Business main section of the Veeam Backup Explorers User Guide. For example:

- Information about viewing and searching objects in the backup file is in the Browsing, Searching and Viewing Items section.
- Information about adding a database manually is in the Standalone Databases Management section.

- Information about restoring process is in the Restoring Microsoft OneDrive Data section.



Restoring Microsoft Teams Items

To restore Microsoft Teams data from a backup or replica of a Veeam Backup for Microsoft Office 365 server, Veeam Backup & Replication uses Veeam Explorer for Microsoft Teams.

> **IMPORTANT**
>
> Before you restore Microsoft Teams data, check requirements and limitations listed in the Planning and Preparation section in the Veeam Backup Explorers User Guide. One of the main requirements is that the backup or replica must be created with application-aware processing enabled.

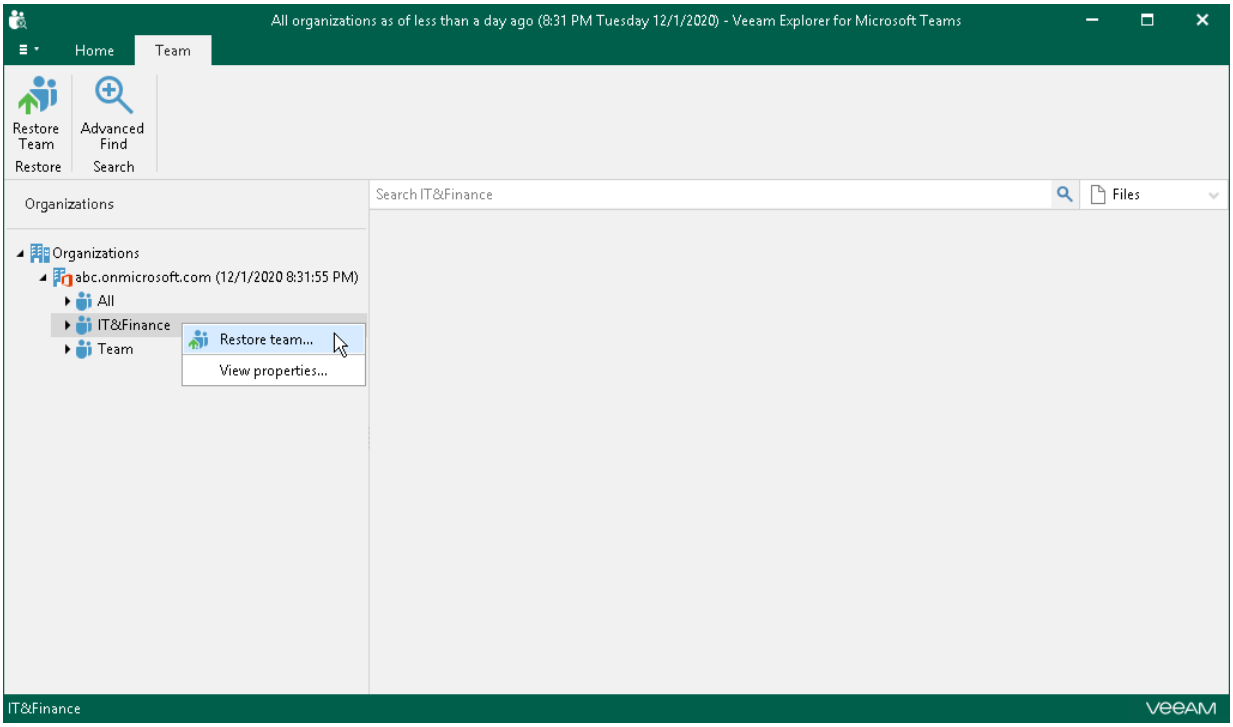
To launch Veeam Explorer for Microsoft Teams from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** or **Replicas** node.
3. In the working area, select the necessary machine in the backup or VM replica and click **Application Items** > **Microsoft Teams** on the ribbon. Alternatively, right-click the machine or VM replica and select **Restore application items** > **Microsoft Teams**.
4. In the opened **Microsoft Teams Files Restore** wizard, select a restore point from which you want to restore data.

Then Veeam Backup & Replication opens Veeam Explorer for Microsoft Teams. In this explorer, you can browse, view, restore Microsoft Teams data and so on. The detailed information is provided in the Veeam Explorer for Microsoft Teams main section of the Veeam Backup Explorers User Guide. For example:

- Information about viewing and searching objects in the backup file is in the Browsing, Searching and Viewing Items section.
- Information about adding a database manually is in the Standalone Databases section.

- Information about restoring process is in the Data Restore section.



Secure Restore

Veeam Backup & Replication allows you to perform secure restore — scan machine data with antivirus software before restoring the machine to the production environment.

During secure restore, Veeam Backup & Replication mounts disks of the machine that you plan to restore to the mount server. On the mount server, Veeam Backup & Replication triggers the antivirus software to scan files from the mounted disks. If during the scan the antivirus software detects malware, Veeam Backup & Replication will either abort the restore process, or restore the machine or its disks with restrictions depending on secure restore settings.

Secure restore is available for the following restore operations:

- Instant Recovery
- Entire VM Restore
- Virtual Disks Restore
- Restore to Microsoft Azure
- Restore to Amazon EC2
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Restore to Google Compute Engine
- Disk Export
- SureBackup

To perform secure restore, you must enable the **Scan the restored machine for malware prior to performing the recovery** option at the **Secure Restore** step of the restore wizard.

> **TIP**
>
> You can also scan machine data for malware regularly within a SureBackup job. For information on how to enable the malware scan for a SureBackup job, see the Settings step of the SureBackup job wizard.

Requirements and Limitations for Secure Restore

Before you perform secure restore, check the following prerequisites:

- You can perform secure restore only for machines that run Microsoft Windows.
- The antivirus software must be installed on the mount server and support the command line interface (CLI).
- The antivirus configuration file must be located on the mount server and must be properly configured. For details, see Antivirus XML Configuration File.
- Veeam Backup & Replication does not perform malware scan for disks or volumes that cannot be mounted to the mount server.

  For example, Storage Spaces disks or ReFS volumes (if ReFS is not supported by the mount server OS) are skipped from the scan and restored in a regular way.

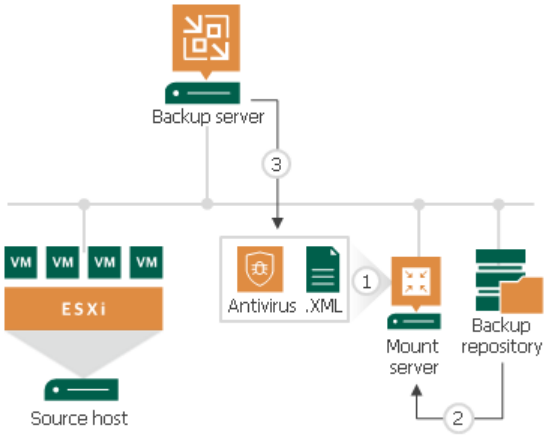
How Secure Restore Works

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. For example, you may want to run the malware scan process on a different server for security reasons. For details on mount server deployment and requirements, see Mount Server.

To run the malware scan, Veeam Backup & Replication performs the following actions:

1. On the mount server, Veeam Backup & Replication runs Veeam Mount Service to check the antivirus configuration file and antivirus software:
   a. Veeam Mount Service verifies if the `AntivirusInfos.xml` configuration file is located in the `%ProgramFiles%\Common Files\Veeam\Backup and Replication\Mount Service` folder.
   b. Veeam Mount Service checks the scan settings in the configuration file and verifies if the antivirus is installed on the mount server.

      Note that if the antivirus is not installed or the configuration file is improperly configured, Veeam Backup & Replication will not start the restore process. In the restore wizard, you will not be able to pass the step with secure restore settings.
2. Veeam Backup & Replication mounts machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
3. Veeam Backup & Replication triggers the antivirus to scan files in the `C:\VeeamFLR\<machinename>` folder.



If during the scan the antivirus does not detect malware, Veeam Backup & Replication restores the machine or its disks to the target location.

If the antivirus detects malware, Veeam Backup & Replication will either abort the restore process, or restore the machine or its disks with restrictions depending on the following secure restore settings:

- Disable the network adapter (NIC) on the restored machine.
- Connect the restored machine to a different Microsoft Azure virtual network.
- Change the AWS security group for the restored machine.
- Disconnect restored virtual disks from the target VM.

You can further access the restored machine or its disks in the isolated environment and clean the infection.

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\Mount Service` 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

If you want to scan machine data with other antivirus software, you must add configuration for this software to the `AntivirusInfos.xml` file. The configuration must contain the same elements and attributes as the antiviruses supported out-of-the-box. Mind that the antivirus software must support the command line interface (CLI).

> **NOTE**
>
> If you made changes to the antivirus configuration file, you do not need to restart Veeam services on the backup server — Veeam Backup & Replication will perform the next malware scan with new settings.

XML File Structure

The XML file describing antivirus settings has the following structure:

```
<Antiviruses>
  <!-- Symantec -->
  <AntivirusInfo Name='Symantec' IsPortableSoftware='false' ExecutableFilePath=
'Veeam.Backup.Antivirus.Scan.exe' CommandLineParameters='/p:%Path%' RegPath='HK
EY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\symcscan' ServiceName='symcs
can' ThreatExistsRegEx='Threat\s+found' IsParallelScanAvailable='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Error' Description='Invalid command line argument'>1</E
xitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>2</ExitCode>
        <ExitCode Type='Error' Description='Antivirus scan was canceled'>4</Exi
tCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>3</Ex
itCode>
     </ExitCodes>
  </AntivirusInfo>
  <!-- Eset -->
  <AntivirusInfo Name='Eset File Security' IsPortableSoftware='true' Executable
FilePath='%ProgramFiles%\ESET\ESET File Security\ecls.exe' CommandLineParameter
s='%Path% /clean-mode=None /no-symlink' RegPath='' ServiceName='' ThreatExistsR
egEx='threat\s*=\s*["&apos;](?!is OK["&apos;])[^"&apos;]+["&apos;]' IsParallelS
canAvailable='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>1</Ex
itCode>
        <ExitCode Type='Warning' Description='Some files were not scanned'>10</
ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>50</E
xitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>100</ExitCode>
     </ExitCodes>
  </AntivirusInfo>
  <AntivirusInfo Name='ESET Antivirus' IsPortableSoftware='true' ExecutableFile
Path='%ProgramFiles%\ESET\ESET Security\ecls.exe' CommandLineParameters='%Path%
/clean-mode=None /no-symlink' RegPath='' ServiceName='' ThreatExistsRegEx='thre
at\s*=\s*["&apos;](?!is OK["&apos;])[^"&apos;]+["&apos;]' IsParallelScanAvailab
le='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>1</Ex
itCode>
        <ExitCode Type='Warning' Description='Some files were not scanned'>10</
ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>50</E
xitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>100</ExitCode>
     </ExitCodes>
  </AntivirusInfo>
  <!-- Windows Defender -->
  <AntivirusInfo Name='Windows Defender' IsPortableSoftware='false' ExecutableF
ilePath='%ProgramFiles%\Windows Defender\mpcmdrun.exe' CommandLineParameters='-
Scan -ScanType 3 -File %Path% -DisableRemediation -BootSectorScan' RegPath='HKE
Y_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WinDefend' ServiceName='WinDe
fend' ThreatExistsRegEx='Threat\s+information' IsParallelScanAvailable='false'>
     <ExitCodes>
```

```
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>2</ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>2</Ex
itCode>
     </ExitCodes>
  </AntivirusInfo>
  <!-- Kaspersky Security -->
  <AntivirusInfo Name='Kaspersky Security' IsPortableSoftware='false' Executabl
eFilePath='kavshell.exe' CommandLineParameters='scan %Path%' RegPath='HKEY_LOCA
L_MACHINE\SYSTEM\CurrentControlSet\services\KAVFS' ServiceName='kavfs' ThreatEx
istsRegEx='' IsParallelScanAvailable='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Warning' Description='There were processing errors for
some files'>-82</ExitCode>
        <ExitCode Type='Warning' Description='Some files were not scanned'>-83<
/ExitCode>
        <ExitCode Type='Warning' Description='Some files were corrupted'>-84</E
xitCode>
        <ExitCode Type='Error' Description='Operation timed out'>2</ExitCode>
        <ExitCode Type='Error' Description='Antivirus scan was canceled'>1</Exi
tCode>
        <ExitCode Type='Error' Description='Service process is not running'>-2<
/ExitCode>
        <ExitCode Type='Error' Description='Access denied'>-3</ExitCode>
        <ExitCode Type='Error' Description='Object not found'>-4</ExitCode>
        <ExitCode Type='Error' Description='Invalid syntax'>-5</ExitCode>
        <ExitCode Type='Error' Description='Invalid operation'>-6</ExitCode>
        <ExitCode Type='Error' Description='Service does not exist'>-7</ExitCod
e>
        <ExitCode Type='Error' Description='Service disabled'>-8</ExitCode>
        <ExitCode Type='Error' Description='Service logon failure'>-9</ExitCode
>
        <ExitCode Type='Error' Description='Unable to create file'>-10</ExitCod
e>
        <ExitCode Type='Error' Description='Invalid command line argument'>-11<
/ExitCode>
        <ExitCode Type='Error' Description='Invalid password'>-12</ExitCode>
        <ExitCode Type='Error' Description='Cannot create report file'>-85</Exi
tCode>
        <ExitCode Type='Error' Description='License is invalid'>-301</ExitCode>
        <ExitCode Type='Error' Description='Antivirus bases are corrupted'>-236
</ExitCode>
        <ExitCode Type='Infected' Description='Infected object was detected'>-8
0</ExitCode>
        <ExitCode Type='Infected' Description='Possibly infected object was det
ected'>-81</ExitCode>
     </ExitCodes>
  </AntivirusInfo>
</Antiviruses>
```

The XML file contains the following elements:

- Antiviruses. The element encapsulates the file with antivirus settings.
- AntivirusInfo. The element describes the antivirus software.
- ExitCodes. The element encapsulates messages that Veeam Backup & Replication displays on scan results.

- ExitCode. The element describes the subject and the body of the message that Veeam Backup & Replication displays on scan results.

AntivirusInfo

The element has the following attributes:

Name	Specifies the antivirus name. Veeam Backup & Replication will display this name in restore session logs.
IsPortableSoftware	Indicates if antivirus software is portable: • If you set this attribute to `True`, Veeam Backup & Replication will treat the antivirus software as portable. Before performing secure restore, Veeam Backup & Replication will verify if the antivirus executable file exists. The path to the file is specified by the ExecutableFilePath attribute. • If you set this attribute to `False`, Veeam Backup & Replication will treat the antivirus software as non-portable. Before performing secure restore, Veeam Backup & Replication will verify if the antivirus registry key exists and if the antivirus service is running. The key is specified by the RegPath attribute. The service name is specified by the ServiceName attribute.
ExecutableFilePath	Specifies the path to the antivirus executable file.
CommandLineParameters	Specifies antivirus commands that you want to execute during the scan. Make sure that the antivirus supports the specified commands. For example, the list of commands for ESET is available in this ESET KB article. **Note:** The `%Path%` variable is required for this attribute. During secure restore, Veeam Backup & Replication substitutes this variable for the path to the folder with mounted disks (`C:\VeeamFLR\<machinename>`).
ServiceName	Specifies the name of the antivirus service. The service must be responsible for data scanning. The attribute value can be an empty string if **IsPortableSoftware** = True and **ExecutableFilePath** is specified.
RegPath	Specifies the registry key of the anitivirus service. The attribute value can be an empty string if **IsPortableSoftware** = True and **ExecutableFilePath** is specified.

ThreatExistsRegEx	Specifies regular expressions. A regular expression is a sequence of characters that form a search pattern. Veeam Backup & Replication will search the antivirus output messages for the specified regular expression. If any of the output messages match the expression, Veeam Backup & Replication will notify you on detected threat. **Note:** You must have a good understanding of the regular expression language to specify this attribute properly. For more information on the regular expression language, see Microsoft Docs.
IsParallelScanAvailable	Indicates if the antivirus will run multiple jobs to scan files on mounted disks simultaneously. If you set this attribute to `True`, Veeam Backup & Replication will lock the antivirus to perform the scan for the current restore session. The antivirus will not be available for other sessions with enabled secure restore until the scan completes. The default value for antivirus lock time-out is 24 hours. If the scan does not complete after this period, Veeam Backup & Replication will finish other restore sessions as specified in the restore wizard: abort restore sessions or restore machines (or its disks) with restrictions. **Note:** You can change the lock time-out using registry keys. For more information, contact Veeam Support. If the antivirus CLI does not support multiple scan jobs, set this attribute to `False`.

ExitCode

The element has the following attributes:

Type	Specifies the subject of the message that Veeam Backup & Replication will display on scan results: • `Success` • `Infected` • `Warning` • `Error`
Description	Specifies the body of the message that Veeam Backup & Replication will display on scan results.

TIP

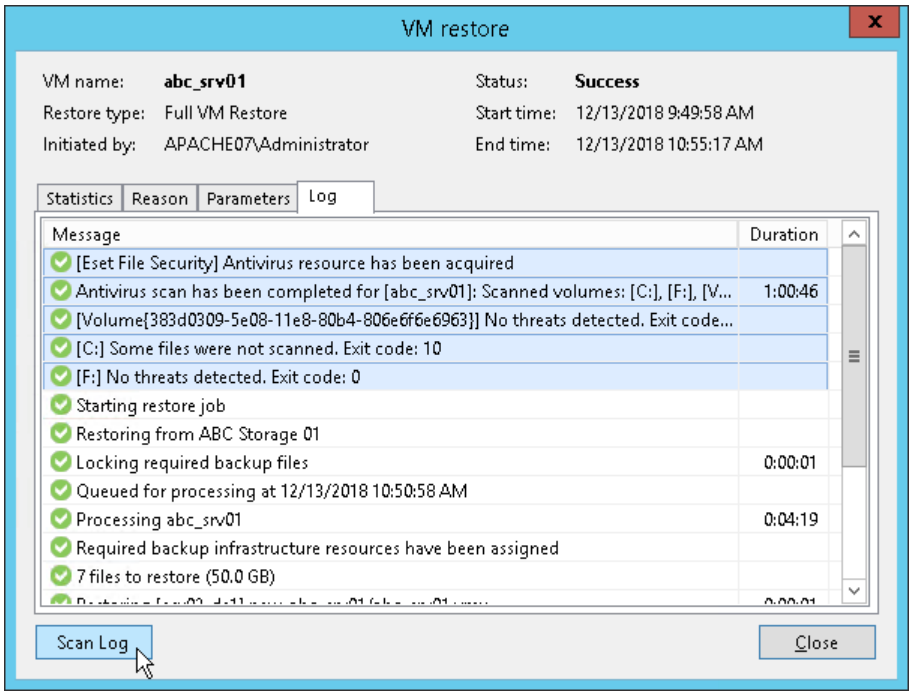
You can distribute the XML configuration file among other mount servers in your backup infrastructure using Veeam PowerShell. For more information, see the Copy-VBRAntivirusConfigurationFile section in the Veeam PowerShell Reference.

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**.



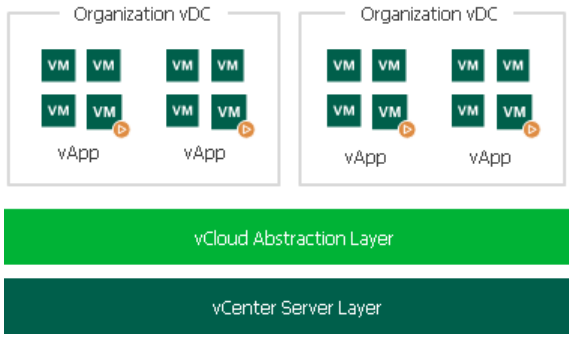
To view the detailed logging of the malware scan, click the **Scan Log** button at the bottom of the window with restore session statistics. Veeam Backup & Replication will display the most recent logs in a file of 1 MB in size.

Full logs of the scan are stored on the mount server in the following folder: `C:\ProgramData\Veeam\Backup\FLRSessions\Windows\FLR__<machinename>_\Antivirus`.

vCloud Director Support

Veeam Backup & Replication provides support for vCloud Director. Veeam Backup & Replication uses vCloud Director API to back up vApps and VMs and restore them directly to the vCloud Director hierarchy.

The main entity with which Veeam Backup & Replication works during backup is a vApp. A vApp is a virtual system that contains one or more individual VMs along with parameters that define operational details — vApp metadata. When Veeam Backup & Replication performs backup of VMs, it captures not only data of VMs being a part of vApps, but also vApp metadata. As a result, you can restore vCloud Director objects back to the vCloud Director hierarchy and do not need to perform any additional actions on import and VM configuration.



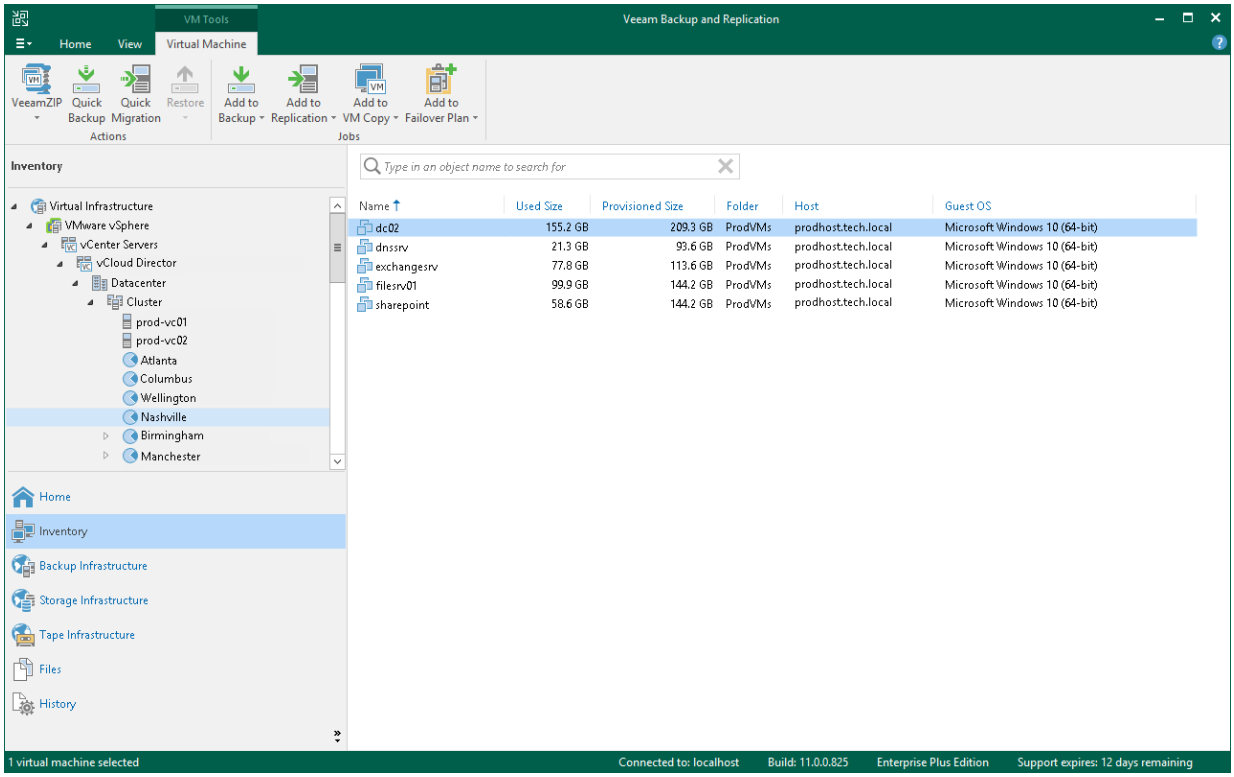
Viewing vCloud Director VMs

After you add the vCloud Director server to the backup infrastructure, you can view the vCloud Director hierarchy in Veeam Backup & Replication and work with vCloud Director VMs.

To open the vCloud Director hierarchy:

1. Open the **Inventory** view.
2. Click the **View** tab on the ribbon.
3. On the **View** tab, click **vCloud View**.

The hierarchy of the vCloud Director server will become available in the inventory pane. VMs managed by vCloud Director will be displayed in the working area. You can work with these VMs just as if you worked with VMs managed by vCenter Servers or registered on ESXi hosts in your backup infrastructure.



Backup and Restore of vApps

Veeam Backup & Replication lets you back up vCloud Director vApps and restore them back to the vCloud Director hierarchy.

In terms of vCloud Director, a vApp is a coherent system that includes one or more VMs. Every vApp is described with a set of operational details — vApp metadata. vApp metadata contains the following information:

- vApp owner settings
- Access rights settings
- vApp network settings: information about organization networks to which the vApp is connected
- Lease settings and so on

When Veeam Backup & Replication performs backup of a vApp, it backs up all VMs being a part of this vApp along with the vApp metadata. Backup of the vApp is performed with the vCD backup job. The vCD backup job may contain one or several vApps. If necessary, you can exclude specific VMs and VM disks from the backup when configuring a vCD backup job.

Veeam Backup & Replication offers the following restore options for backed-up vApps:

- Restoring vApps to vCloud Director
- Restore of separate VMs being a part of the vApp to vCloud Director

> **NOTE**
>
> Veeam Backup & Replication shows restore points of a VM that is no longer processed, for example excluded from the job or deleted, as long as incremental backup files with this VM exist in the active backup chain. After Veeam Backup & Replication deletes the files according to the retention policy, it also hides the VM restore points from the list of restore points created with this backup job. Although the full backup file with the VM may still exist, restore from this file into the vCloud Director is not supported.

Backup of vCloud Director VMs

Veeam Backup & Replication lets you perform backup for vApps and VMs, as well as VM containers in vCloud Director such as organization VDC, organization and even the vCloud Director instance.

When Veeam Backup & Replication performs backup of vApps and VMs, it additionally captures vApp metadata.

vApp metadata includes:

- General information about the vApp where VMs reside, such as: vApp name, description, VMs descriptions
- Information about vApp networks and organization networks to which the vApp is connected
- VMs startup options
- User information
- Lease
- Quota
- Storage template and so on

vApp metadata is stored together with the VM content. Capturing vApp metadata is extremely important for restore: without it, you will not be able to restore vApps and VMs back to vCloud Director.

Data to Back Up

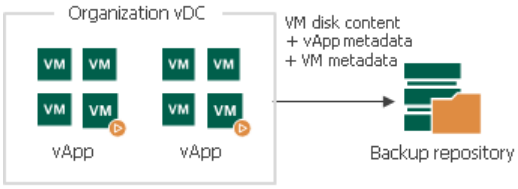
Veeam Backup & Replication lets you back up the following types of vCloud Director VMs:

- Regular VMs that are part of vApps
- Standalone VMs that were created in the vCloud Director tenant portal
- Linked clone VMs that are associated with vApps

Backup of Regular and Standalone VMs

When you back up regular or standalone VMs, Veeam Backup & Replication captures and stores to the backup file the following data:

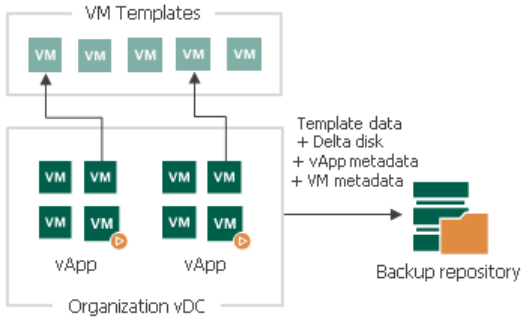
- VM disk content
- vApp metadata
- VM metadata



Backup of Linked Clone VMs

When you back up linked clone VMs, Veeam Backup & Replication captures and stores to the backup file the following data:

- Content of the template to which the VM is linked
- Content of the VM user disk — delta disk
- vApp metadata
- VM metadata



During full backup of linked clone VMs, Veeam Backup & Replication consolidates data of the VM template and delta disk and saves it as a regular VM disk in the backup file. Data merging guarantees proper VM restore: even if a VM template is lost by the time of recovery, you will still be able to restore the linked clone VM from the backup.

During incremental backup, Veeam Backup & Replication saves only changed data of the delta file.

Limitations for Backup of Linked Clone VMs

Before you back up linked clone VMs, consider the following:

- [For vCenter 6.5 or later] If you back up a linked clone VM that has snapshots, Veeam Backup & Replication may fail to produce a valid restore point. To overcome this issue, do one of the following:
  - Disable CBT (Change Block Tracking) in the backup job settings.
  - Ensure that CBT is enabled on the VM template to which the VM is linked.

  For details on how to enable CBT on the VM template, contact Veeam Customer Support.
- Backup of linked clone VMs that were created with services other than vCloud Director may cause snapshot-related problems. To overcome this issue, disable Veeam Snapshot Hunter. For details, see this Veeam KB article.

vCD Backup Jobs

For VMs managed by vCloud Director, Veeam Backup & Replication offers a special type of the backup job — vCD backup job. vCD backup jobs process vCloud Director objects, ensure their proper restore and support of vCloud-specific features.

It is recommended that you use vCD backup jobs to back up VMs managed by vCloud Director. If you back up VMs managed by vCloud Director using a regular backup job, Veeam Backup & Replication will perform backup at the level of the underlying vCenter Server and will not capture vApp metadata. As a result, you will not be able to restore a fully-functioning VM to vCloud Director.

Performing Backup of vCloud Director VMs

The vCD backup is practically the same as a regular VM backup. The vCD backup job aggregates main settings for the backup task and defines when, what, how and where to back up vCloud Director VMs.

You can perform the vCD backup job for single VMs and for VM containers:

- vApp
- Organization VDC
- Organization
- vCloud Director instance

Just like a regular backup job, the vCD backup job can be scheduled or run manually. To create a vCD backup job, do one of the following:

- On the **Home** tab, click **Backup Job** and select **vCloud**.
- Open the **Home** view, in the inventory pane right-click **Jobs** and select **Backup** > **vCloud**.
- Open the **Inventory** view, click the **View** tab and click **vCloud View** on the ribbon. In the inventory pane expand the **vCloud Director** hierarchy, in the working area select one or more VMs, click **Add to Backup** on the ribbon and select **New job**. Alternatively, you can right-click one or several VMs and select **Add to backup job** > **New job**. In this case, the selected VMs will be automatically added to the new vCD backup job. You can add other VMs to the job when passing through the wizard steps.

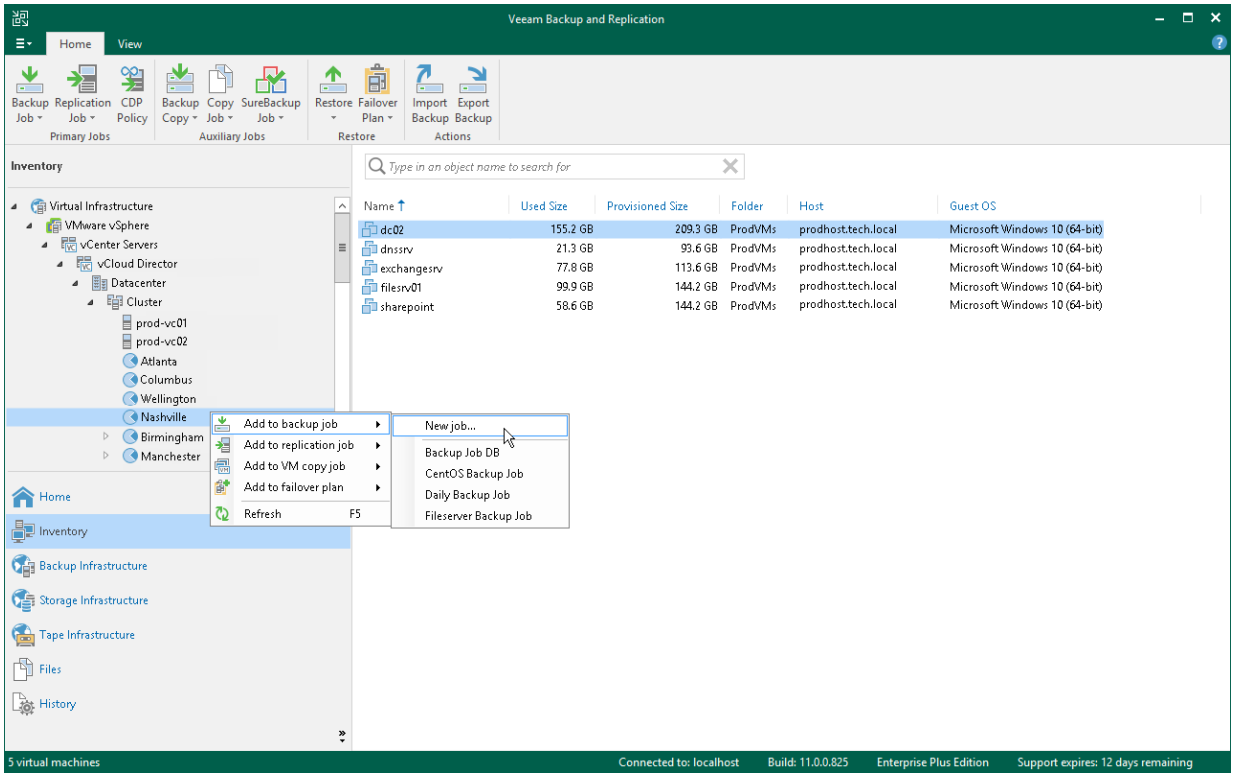
You can quickly include VMs to already existing vCD backup jobs. To do this, in the **Inventory** view, in the working area right-click necessary VMs and select **Add to backup job** > *name of the job*.

The **New vCD Backup Job** wizard offers the same options as a **New Backup Job** wizard. For more information, see Creating Backup Jobs.

IMPORTANT

Mind the following:

- Limitation of concurrent tasks is not supported for vCD backup jobs. Fore more information, see Limitation of Concurrent Tasks.
- If you run a vCD backup job for the vApp, the job is considered to finish with the *Success* status and the complete restore point for the vApp is created only if all VMs in the vApp are successfully backed up. If any VM in the job fails, the restore point for the vApp will be in the incomplete status, and you will not be able to restore the whole vApp from such restore point. However, you will be able to perform restore to the original vApp for VMs that were partially or successfully backed up and whose data is available in the incomplete restore point. Veeam Backup & Replication will restore all data that is available for such VMs in the backup.

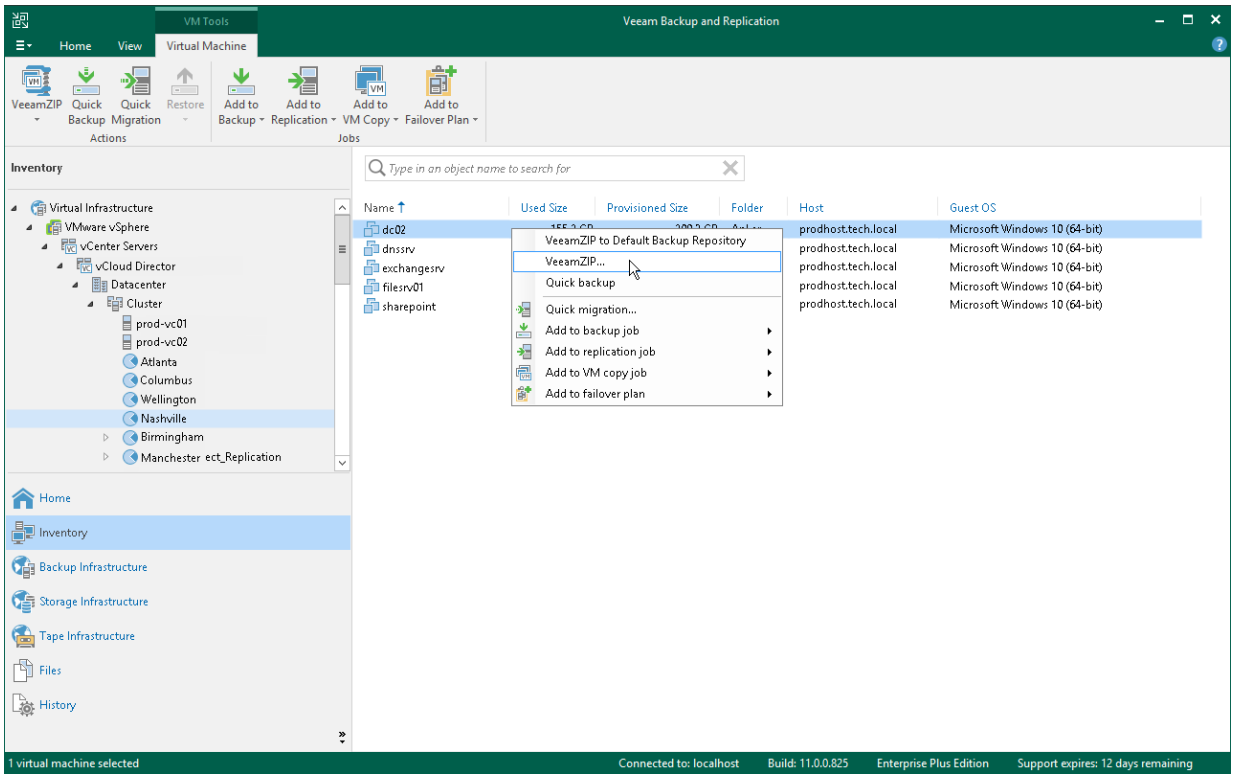


Creating VeeamZIP Files for vCloud Director VMs

You can create a VeeamZIP file for one or more vCloud Director VMs.

When Veeam Backup & Replication creates a VeeamZIP file for a vCloud Director VM, it backs up a VM as separate object. Veeam Backup & Replication does not capture metadata of the vApp to which the VM belongs. When you restore a vCloud Director VM from the VeeamZIP file, Veeam Backup & Replication registers the VM on the underlying ESXi host and does not register the VM in vCloud Director.

The process of VeeamZIP files creation for vCloud Director VMs does not differ from that for regular VMware VMs. For more information, see Creating VeeamZIP Files.



vCloud Director Replication

vCloud Director replication is a technology that allows you to replicate VM containers between production and disaster recovery vCD environments, to the original organization VDC or to another organization VDC added to a different vCloud Director. Using the vCloud Director replication you can replicate the following VM containers:

- vApps
- vCloud organizations
- vCloud organization VDC
- vCloud Director instances

The vCloud Director replication technology utilizes the same mechanisms as the VM replica and follows the same recovery scenarios. The main differences from VM replica are the following:

- The vCD replica target is an organization VDC and it must be set up beforehand in your vCloud Director infrastructure.
- A minimal unit of a vCD replication is a vApp, you cannot replicate a single VM that is added to a vApp.
- The vCD replica snapshots are created on the VMs that are added to the target vApp.
- A single restore point of the vCD replica contains snapshots of all VMs added to a vApp.
- Veeam Backup & Replication assigns a new CLOUD.UUID to the replicated VMs when a vCD replication job runs.

The vCloud Director replication technology supports the same functionality as a normal VM replication: you can create vCD replication jobs, perform failover and failback operations with replicated VM containers and manage these VM containers.

Related Topics

- Creating vCD Replication Job
- Managing vCD Replication Jobs
- Managing vCD Replicas
- Performing Failover and Failback

Creating vCD Replication Job

To replicate a vApp or another VM container, you must configure a vCD replication job. The vCD replication job defines a scope of VM containers to replicate, where to store replicated VM containers and how often to replicate VM containers. Every time the vCD replication job runs, replica snapshots are created on the VMs that are added to the target vApp. If a disaster strikes, you can fail over to the necessary restore point of your VM container. During failover, Veeam Backup & Replication shifts from the original VM container to their replicas. As a result, you have fully functional vApps within a couple of seconds, and your users can access services and applications with minimum disruption.

> **IMPORTANT**
>
> vCD replication job does not support replication of a single VM that is added to a vApp. You can replicate only vApps or VM containers.

To create a vCD replication job, do the following:

1. Check prerequisites
2. Launch the New Replication Job wizard
3. Specify a job name and advanced settings
4. Select vApps to replicate
5. Exclude objects from a replication job
6. Specify the vApp processing order
7. Select a destination for replicas
8. Configure network mapping
9. Specify replication job settings
10. Specify advanced replica settings
11. Specify data transfer and replica settings
12. Configure seeding and mapping settings
13. Specify guest processing settings
14. Define a job schedule
15. Finish working with the wizard

Before You Begin

Before you create a vCD replication job, check the following prerequisites:

- The supported vCloud Director versions are 9.7 - 10.2.
- You must add vCloud Director server to the backup infrastructure. The backup server must be able to resolve short names and connect to source and target virtualization hosts.
- You must set up the organization VDC that will keep the vApp replicas.
- You must disable the VM discovery option in the general system settings of VMware Cloud Director. For more information, see VMware Docs.

- vCloud Director replication does not use and support continuous data protection (CDP).
- vCloud Director replication does not support replication of a vApp to a Cloud Provider.

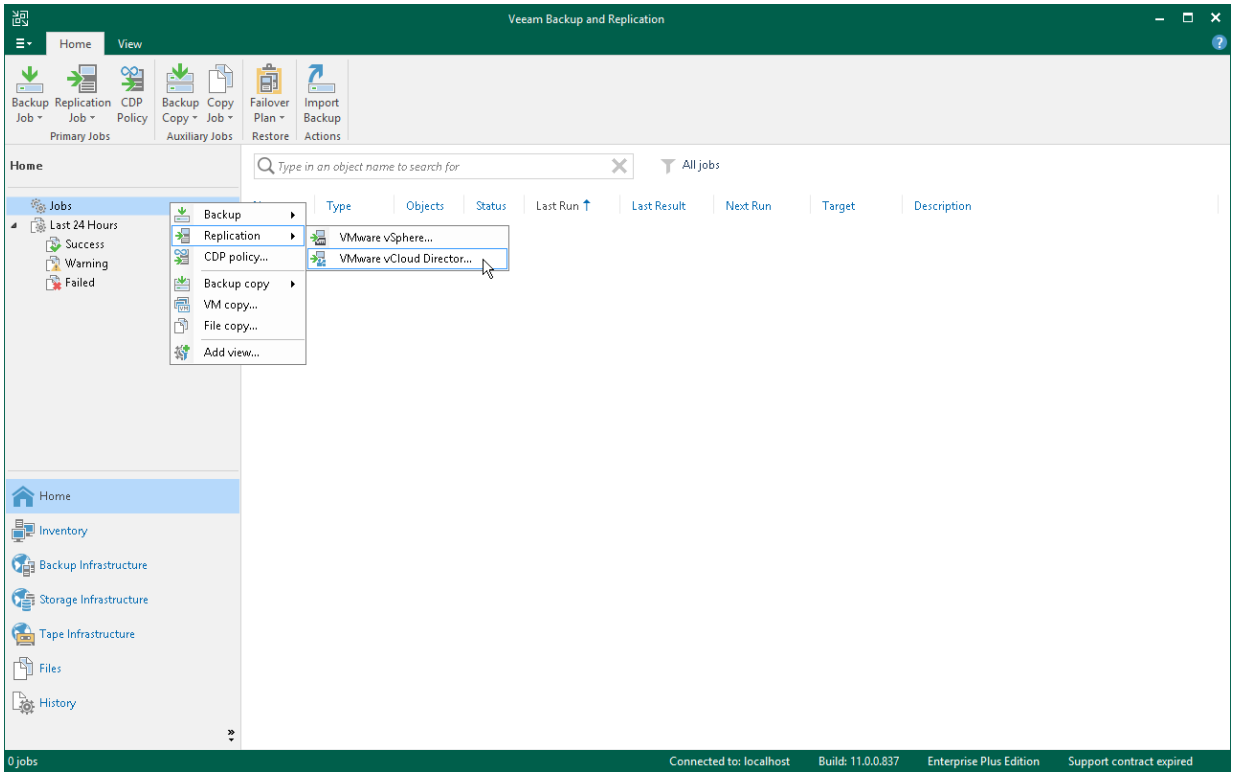
  > **NOTE**
  >
  > It is possible to replicate vSphere VMs to your Cloud Provider vCloud. For more information, see Creating Replication Jobs.

- If you plan to replicate VM containers 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 VM containers using WAN accelerators, it is recommended that you pre-populate global cache on the target WAN accelerator before you start the vCD 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 use pre-job and post-job scripts and/or pre-freeze and post-thaw scripts, you must create scripts before you configure the vCD replication job. Veeam Backup & Replication supports script files in the following formats: EXE, BAT, CMD, JS, VBS, WSF, PS1, SH. For more information, see Pre-Freeze and Post-Thaw Scripts.
- You must check limitations for replication. For more information, see Requirements and Limitations.
- Due to VMware vSphere limitations, if you change the size of VM disks on the source VM, all VM snapshots that were created previously, will not be available for failover, failback and other subsequent operations. For more information, see this VMware KB article.

Step 1. Launch New Replication Job Wizard

To launch the **New Replication Job** wizard, do one of the following:

- On the **Home** tab, click **Replication Job** > **VMware vCloud Director**.
- Open the **Home** view, in the inventory pane right-click **Jobs** and select **Replication** > **VMware vCloud Director**.



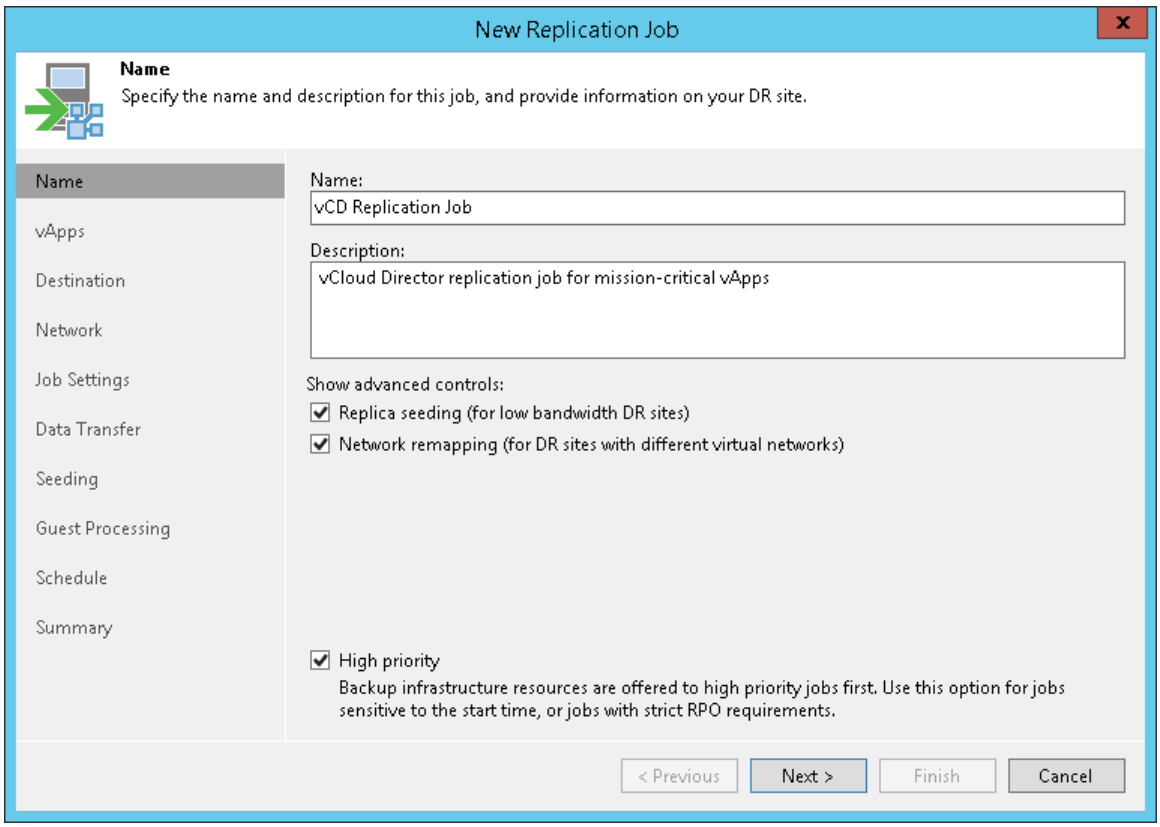
Step 2. Specify Job Name and Advanced Settings

At the **Name** step of the wizard, specify a job name and description, and configure advanced settings for the vCD replication job.

1. In the **Name** field, enter a name of the replication job.
2. In the **Description** field, provide a description for future reference.
3. Depending on your DR site configuration, you can select the following advanced settings for the job:
   - Select the **Replica seeding (for low bandwidth DR sites)** check box to enable the Seeding step in the wizard. Replica seeding can be used if you plan to replicate vApps and want to reduce the amount of traffic sent over the network during the first run of the replication job.
   - Select the **Network remapping (for DR sites with different virtual networks)** check box to enable the Network step in the wizard. If the network in the DR site does not match the production network, you can resolve this mismatch by creating a network mapping table.

     > **NOTE**
     >
     > vCD replication jobs do not support network mapping of the vApp networks. You can configure a mapping table for organization VDC network only.

4. Select the **High priority** check box if you want Veeam Backup & Replication to prioritize this job higher than other similar jobs. Veeam Backup & Replication will allocate resources to this job in the first place. For more information on job priorities, see Job Priorities.



Step 3. Select vApps to Replicate

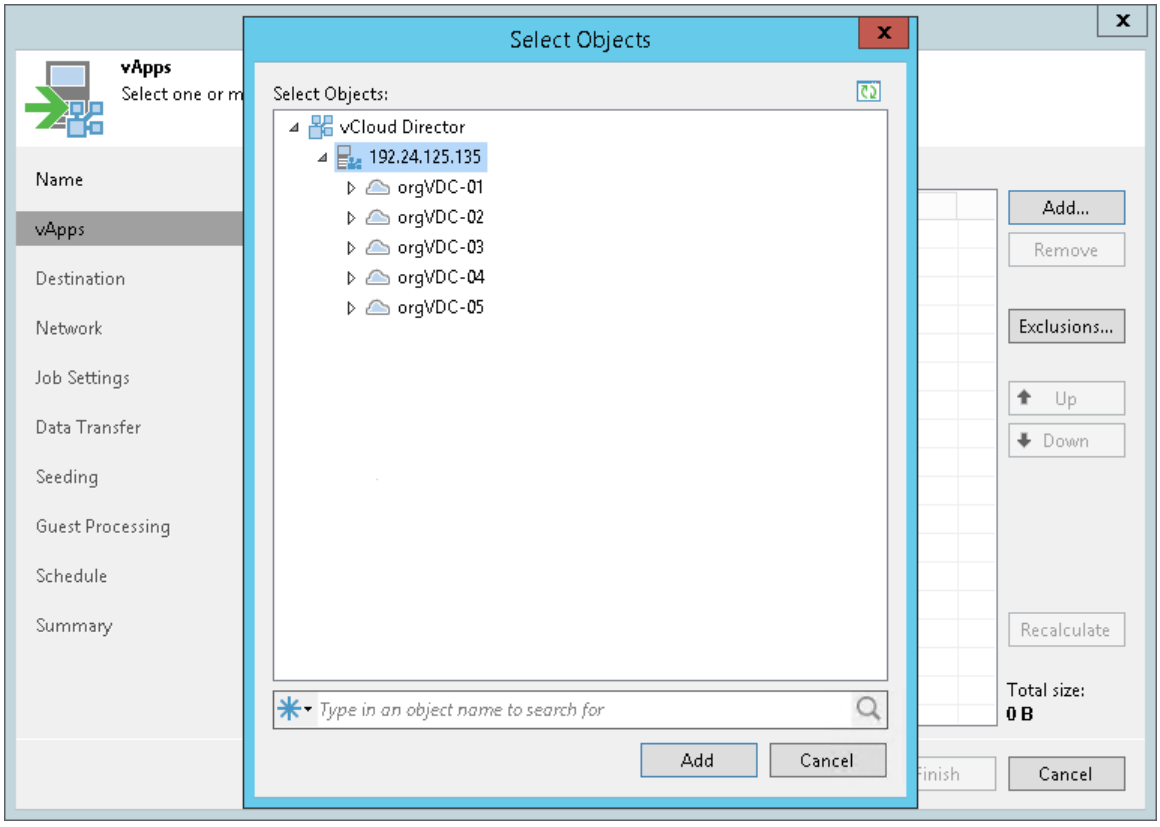
At the **vApps** step of the wizard, select vApps or VM containers which you want to replicate.

1. Click **Add**.
2. In the **Selects Objects** window, select the necessary vApps or VM containers. Click **Add**.

   If you select vApps containers (for example, organization VDCs) and add new vApps to this container in future, Veeam Backup & Replication will update replication policy settings automatically to include these vApps.

   > **IMPORTANT**
   >
   > vCD replication job does not support a replication of a single VM that is added to a vApp. You can replicate only vApps or VM containers.

   To quickly find necessary vApps, you can use the search field at the bottom of the Selects Objects window. If you want to switch between types of VM containers you want to search through, use the button to the left of the search field.



Step 4. Exclude Objects

After you have added vApps or VM containers to the replication job, at the **vApps** step of the wizard you can specify which objects you want to exclude from the replication job.

> **NOTE**
>
> You can exclude only child objects that are added to a VM container. For example, if you add an organization VDC to a replication job, you can exclude only vApps that are available in this organization VDC.

You can exclude the following types of objects:

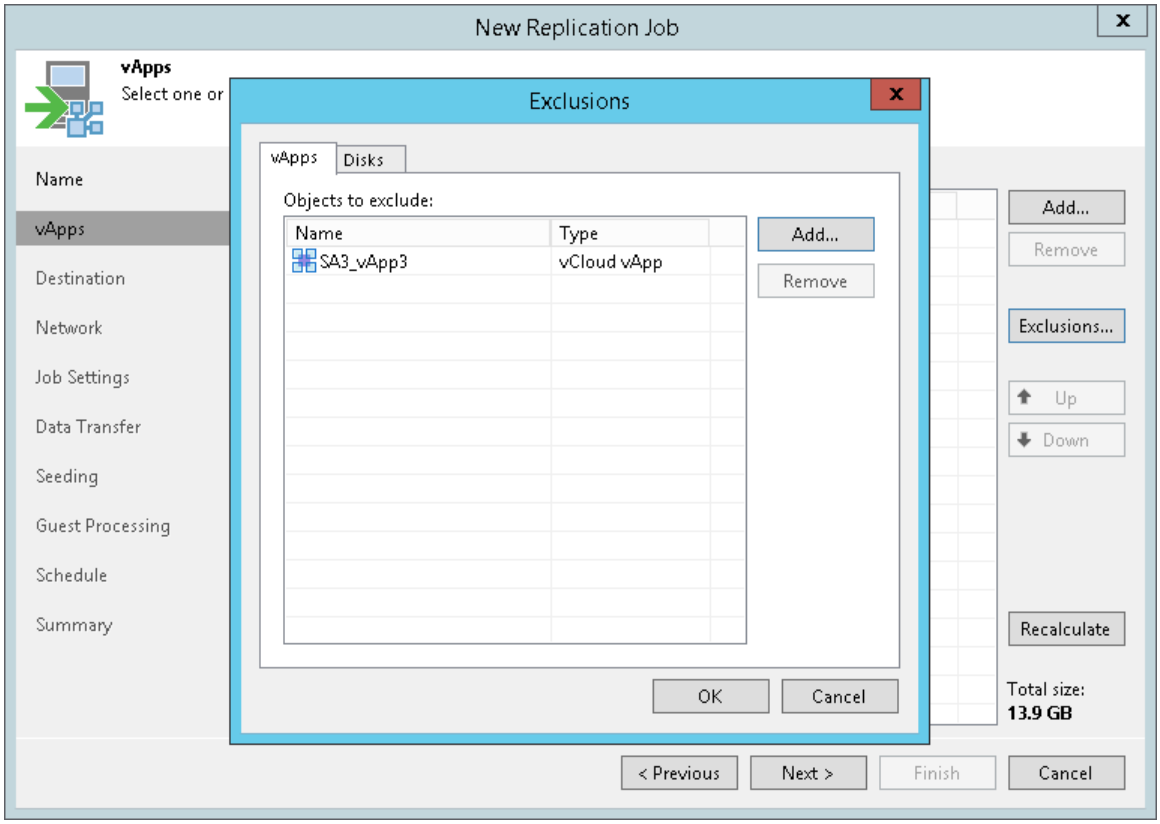
- vApps and VM containers
- Disks

Excluding vApps and VM Containers

To exclude vApps or VM containers:

1. At the **vApps** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, check that the **vApps** tab is selected. Click **Add**.
3. In the **Select Objects** window, select vApps or VM containers that you want to exclude from replication jobs. Click **OK**.

   Select the **Show full hierarchy** check box to display the hierarchy of all vCloud Director servers added to the backup infrastructure.



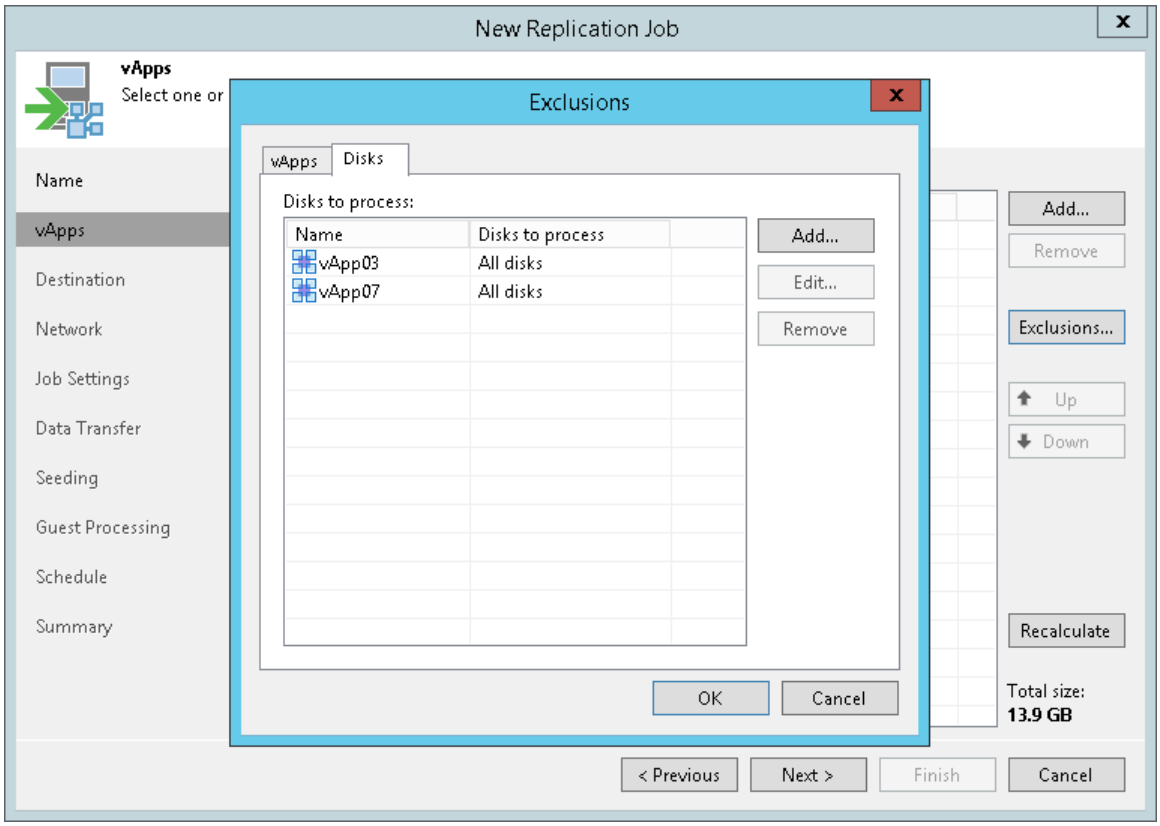
Specifying Disks

This option allows you to specify VM disks that you want to replicate and the VM disks that you want to exclude from replication job.

1. At the **vApps** step of the wizard, click **Exclusions**.
2. Switch to the **Disks** tab.
3. Click **Add**.
4. In the **Select Objects** window, select vApps or VM containers which disks you want to replicate.

   > **TIP**
   >
   > You can specify disks of a single VM.

5. To replicate a specific type of disks, select the necessary vApps or VM containers in the list and click **Edit**.
6. 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.
7. Click **OK**.

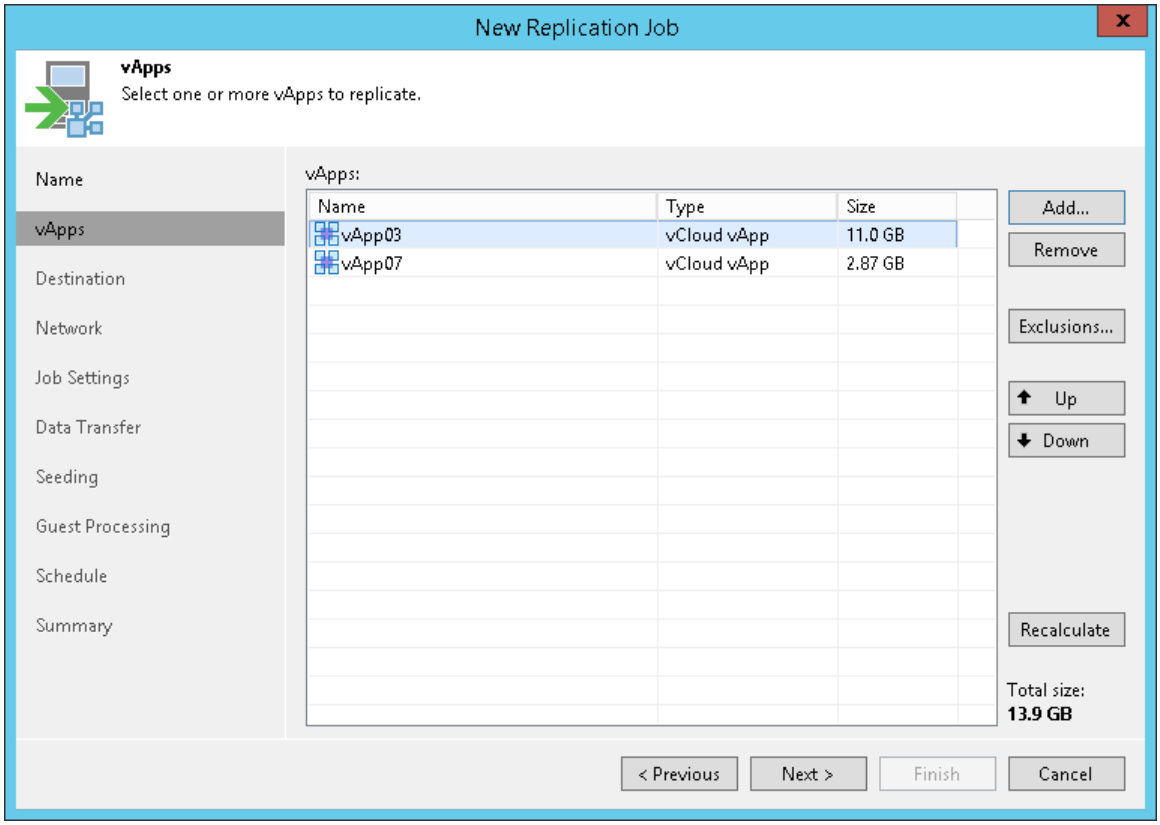


Step 5. Define vApp Processing Order

You can define the order in which the replication job must process vApps. Setting vApps order can be helpful, for example, if you have added some mission-critical vApps to the job and want the job to process them first. You can set these vApps first in the list to ensure that their processing fits the backup window.

At the **vApps** step of the wizard, click **Up** and **Down** to change the processing order. The vApps that are listed at the top have a higher priority and will be processed first.

> **NOTE**
>
> Mind the following:
>
> - VMs inside vApps or a VM container are processed at random.
> - The processing order may differ from the order that you have specified. For example, if resources of a vApp that is higher in the priority are not available, and resources of a vApp that is lower in the priority are available, Veeam Backup & Replication will process the vApp with the lower priority first.



Step 6. Select Replica Destination

At the **Destination** step of the wizard, select an organization VDC to which you want to restore the vApp replicas and the storage policy that you want to apply to the vApp replicas. You can assign a storage policy to all objects of your vCloud Director infrastructure and also specify a storage policy for a particular vApp or a VM.

1. Next to the **Organization VDC** field, click **Choose** and select an organization VDC where replicas must be registered.
2. Next to the **Storage policy VDC** field, click **Choose** and select the storage policy.
   Veeam Backup & Replication will apply this storage policy to the vApp replicas.

   If you want to assign a specific storage policy to individual vApp replicas:

   a. Click the **Select storage policy for individual replicas** link.

   b. In the **Storage policy** window, click **Add** on the right.

   c. In the **Select vApp(s)** window, select necessary vApps.

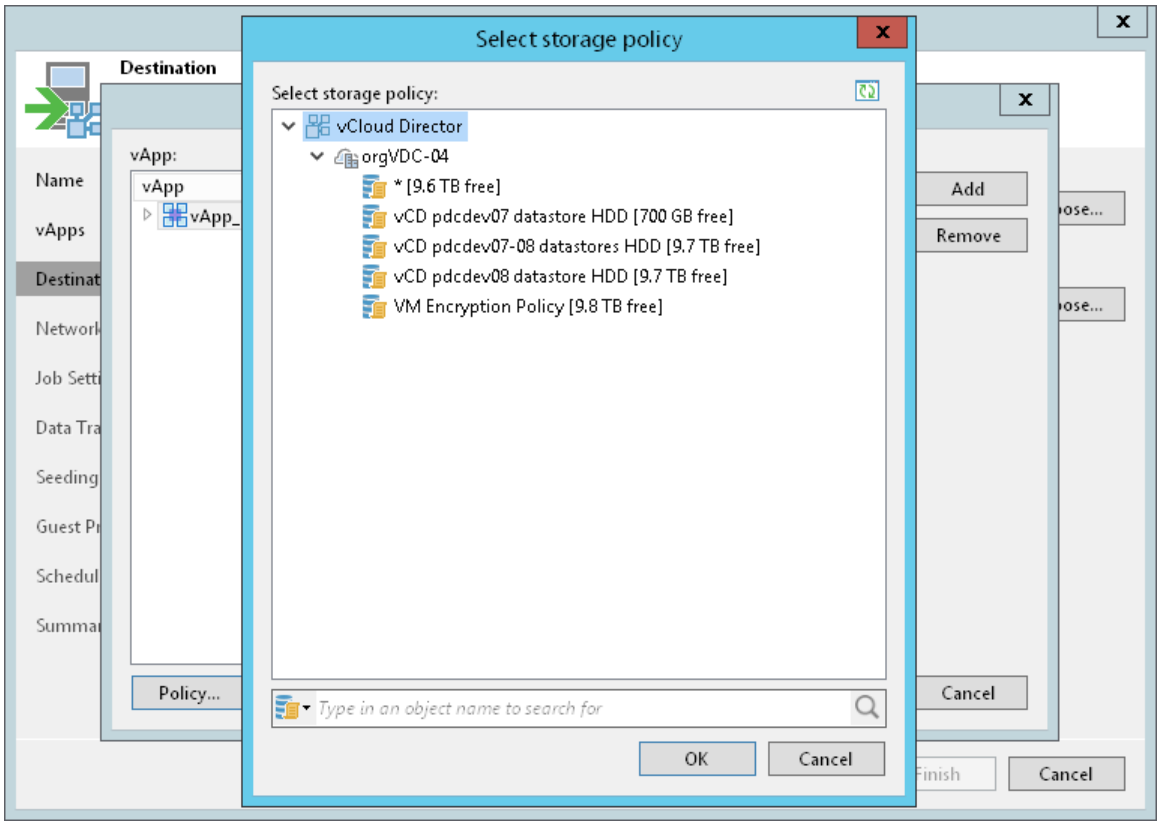
   d. Select the added vApp in the **Storage policy** list and click **Policy** at the bottom of the window.

   e. From the list of available storage policies, select the necessary storage policy.

      You can select a storage policy for a single VM added to a VM container.

TIP

When selecting the necessary object in the virtual infrastructure, you can use the search field at the bottom of the corresponding window. Click the button on the left of the field to select the necessary type of object, enter the object name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.



Step 7. Configure Network Mapping

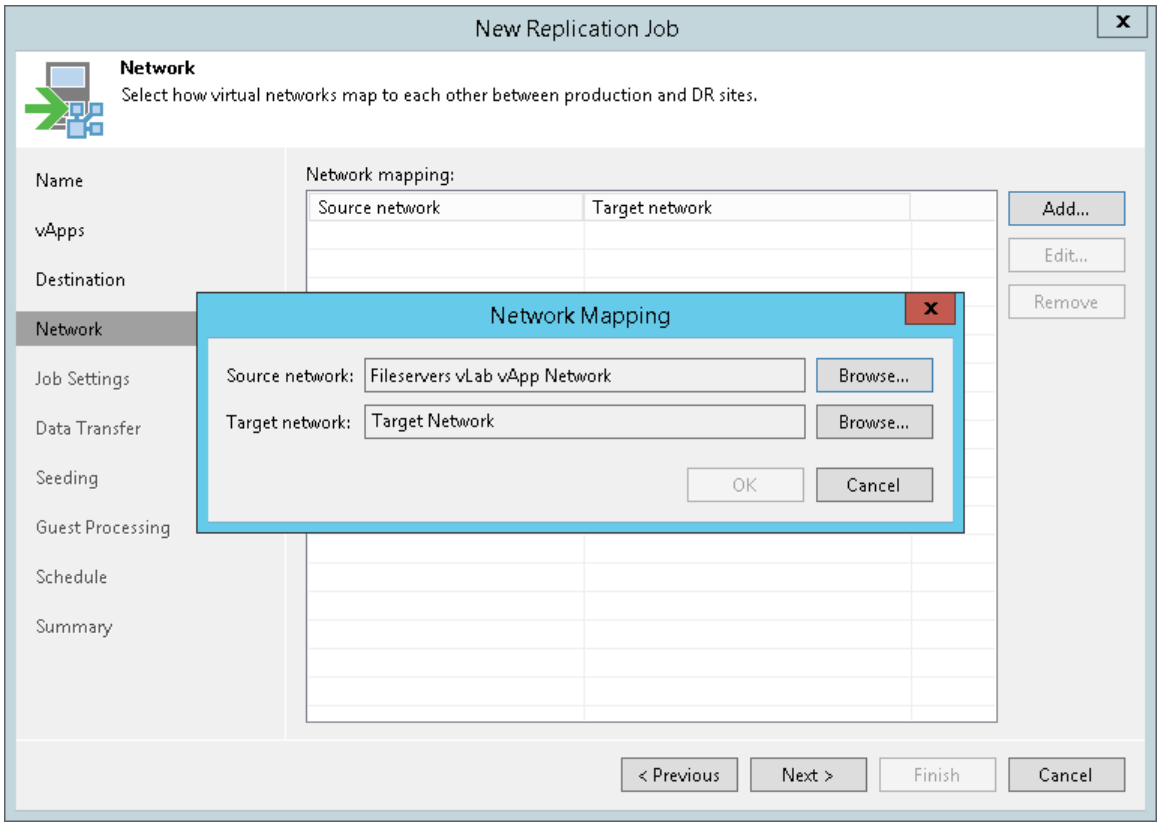
The **Network** step is available if you have chosen to use network mapping at the **Name** step of the wizard.

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. During every vCD replica job session, Veeam Backup & Replication will check the network mapping table and update the replica configuration file to replace the production network with the specified network in the DR site. As a result, you will not have to re-configure network settings manually, a replicated vApp will be connected to a vDC network according to this mapping.

> **NOTE**
>
> Mind the following:
>
> - vCD replication jobs do not support network mapping of the vApp networks. You can configure a mapping table for organization VDC network only.
> - If your vCD replica has a configured network remapping and IP addresses are set as a static pool on the target, Veeam Backup & Replication will assign these IP addresses to VMs NICs on a "first come, first served" basis: the first available IP address will be assigned to the first NIC, the second available IP address will be assigned to the second NIC and so on.

To add a row to a network mapping table:

1. Click **Add**.
2. Click **Browse** next to the **Source network** field and select the production network. The vApps that you want to replicate must be connected to this network.
3. Click **Browse** next to the **Target network** field and select the network in the DR site. Veeam Backup & Replication will connect the replicated vApps to this network.



Step 8. Specify Replication Job Settings

At the **Job Settings** step of the wizard, define vCD replication job settings.

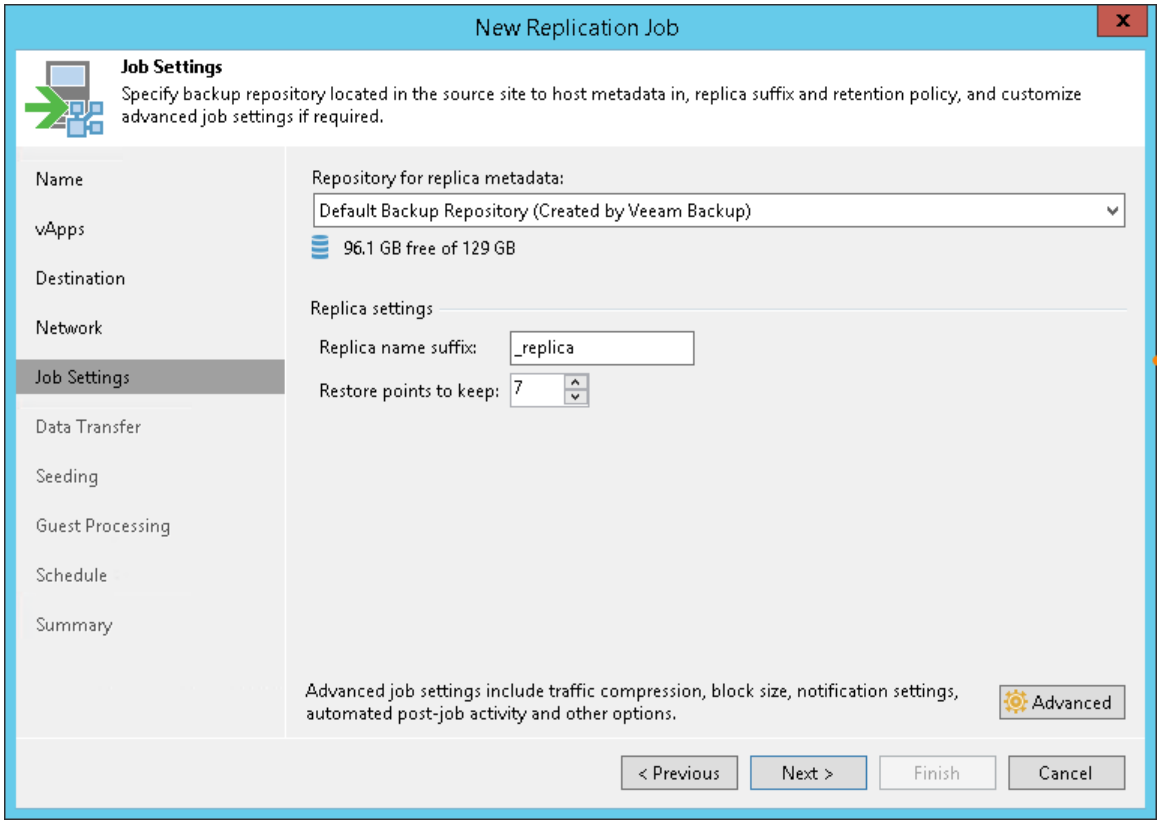
1. From the **Repository for replica metadata** list, select a backup repository. Veeam Backup & Replication uses this backup repository to keep metadata for vApp replicas. Metadata is required when you perform failback with the Quick Rollback option. For more information on metadata, see Changed Block Tracking.
2. In the **Replica name suffix** field, enter a suffix for the name of vApp replicas. Veeam Backup & Replication will add this suffix to the name of the target vApps and will register vApp replicas with this suffix.
3. In the **Restore points to keep** field, specify the number of restore points the replication job must maintain. If this number is exceeded, Veeam Backup & Replication removes the earliest restore point.

   When you specify the retention policy settings for the replication job, consider available space on the target datastore. Due to VMware restrictions on the number of VM snapshots, the maximum number of restore points for vApp replicas is limited to 28. An excessive amount of restore points may overfill the target datastore.

IMPORTANT

Mind the following:

- You cannot store vApp 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 vApp replica metadata in a scale-out backup repository or NFS backup repository.
- At least one successful run of a vCD replication job is required to apply a retention policy to a vCD replica and to delete previous restore points.



Step 9. Specify Advanced Replica Settings

At the **Job settings** step of the wizard, you can specify the following settings for the vCD 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.

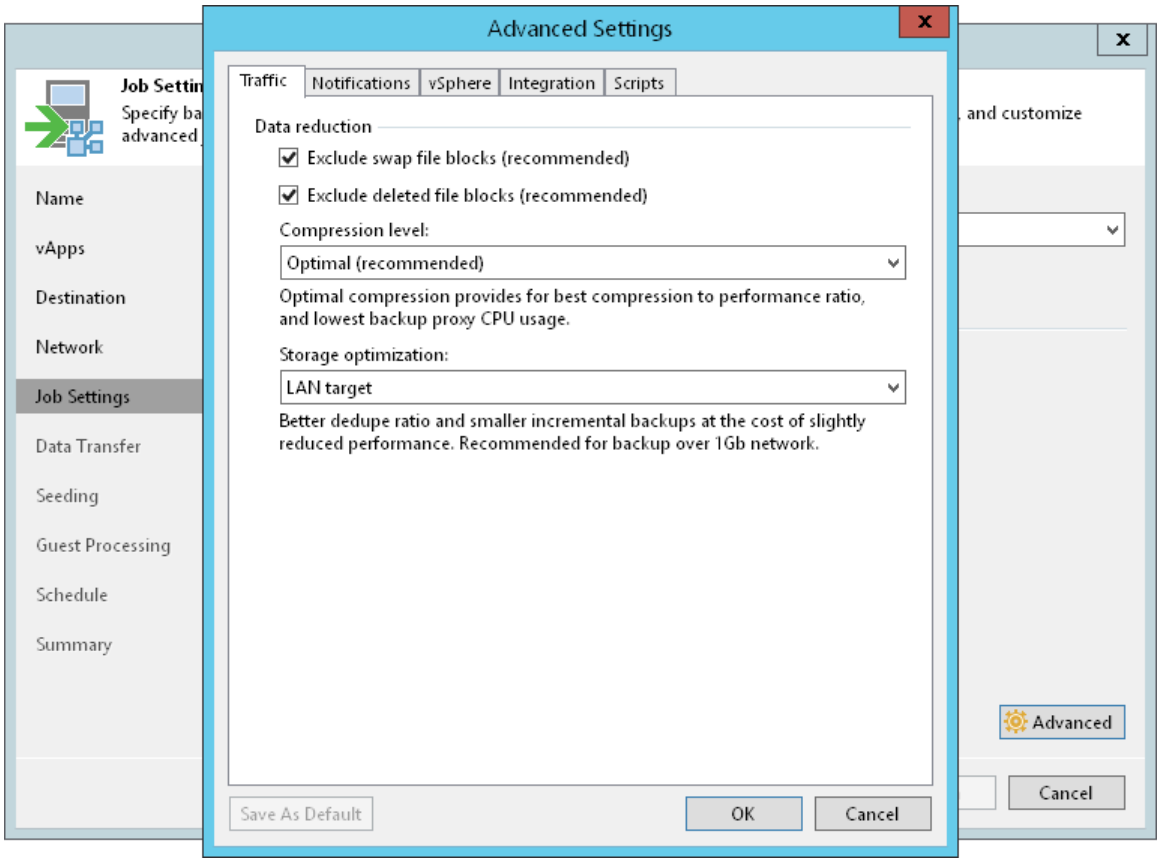
Traffic Settings

To specify traffic settings for a vCD replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Verify that the **Traffic** tab is selected.
3. 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 replication 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 data.

   If you want to include data blocks of the `hiberfil.sys` file and `pagefile.sys` file to the replica, clear the **Exclude swap file blocks** check box. For more information, see Swap Files.
4. 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 the replica and increase the job performance. If you want to copy dirty blocks, clear the **Exclude deleted file blocks** check box. For more information, see Deleted File Blocks.
5. From the **Compression level** list, select a compression level for the created replica: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*. Compression is applicable only if data is transferred between two backup proxies. If one backup proxy acts as the source and target backup proxy, data is not compressed at all.

6. 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.



Notification Settings

To specify notification settings for a vCD replication job:

1. At the **Job Settings** 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 configure 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 by email if a job completes with a failure, warning or success. In the field below, specify a recipient email address. You can enter several addresses separated by a semicolon.

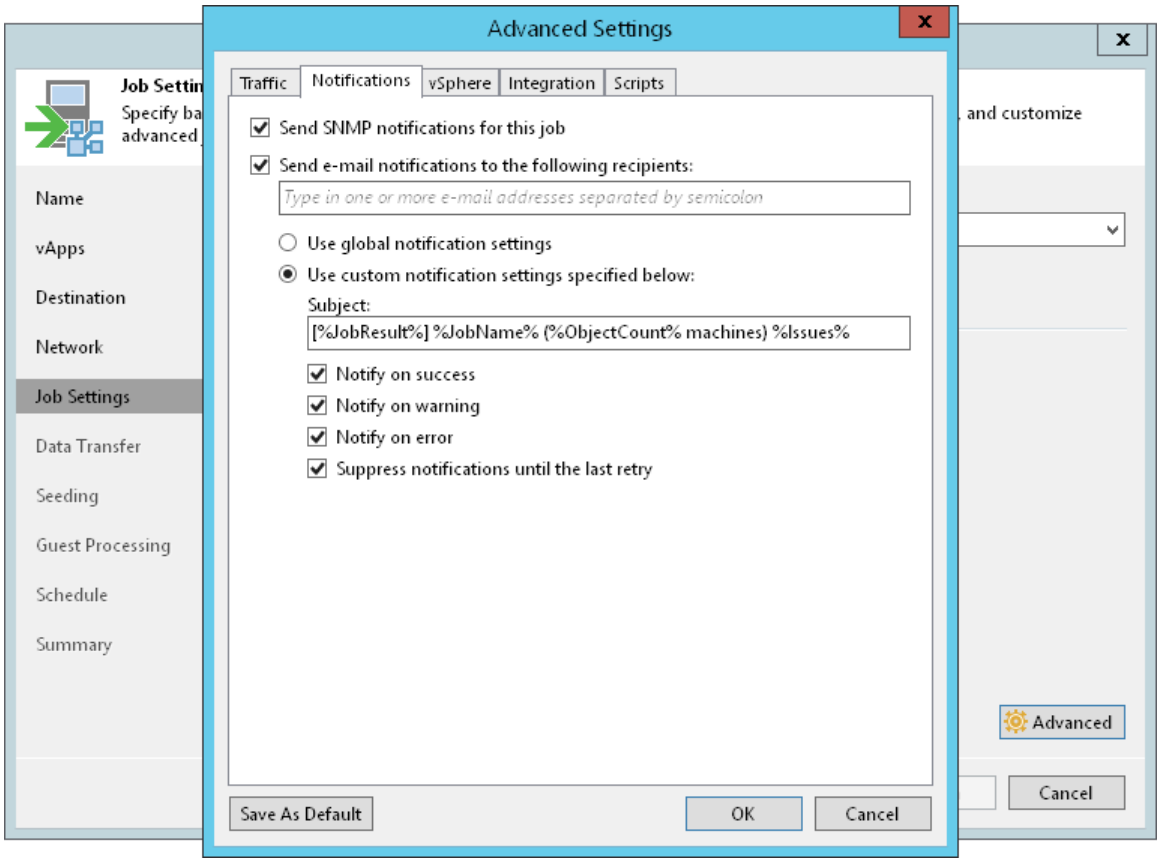
   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.
5. You can choose to use global notification settings or specify custom notification settings.
   - To receive 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 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%*, *%JobResult%*, *%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 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.



vSphere Settings

To specify vSphere settings for a vCD 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 VMware FileSystem Sync Driver (vmsync) 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 replication.
4. In the **Changed block tracking** section, configure VMware vSphere Changed Block Tracking 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.

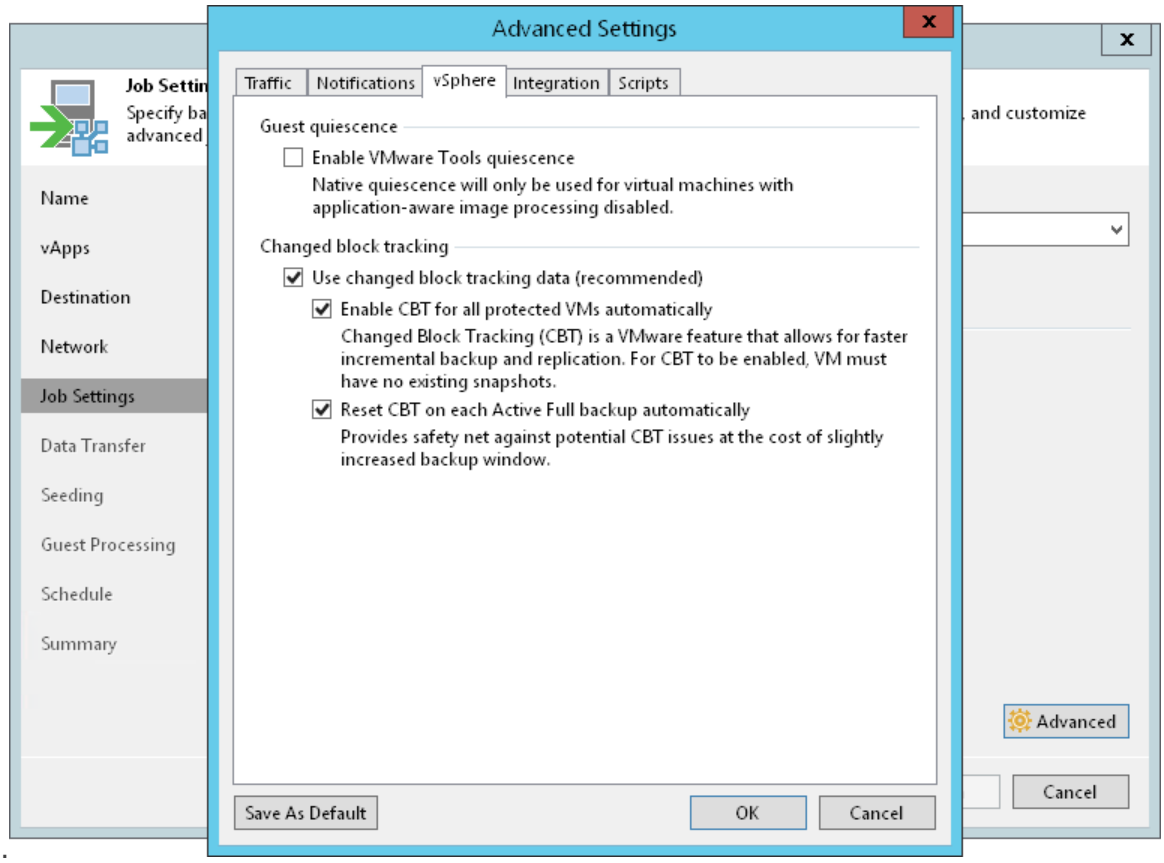
c. Make sure that the **Reset CBT on each Active Full backup automatically** check box is selected if you want to reset CBT after the replication job starts for the first time.

   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 vCD 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 Backup from Storage Snapshots section.

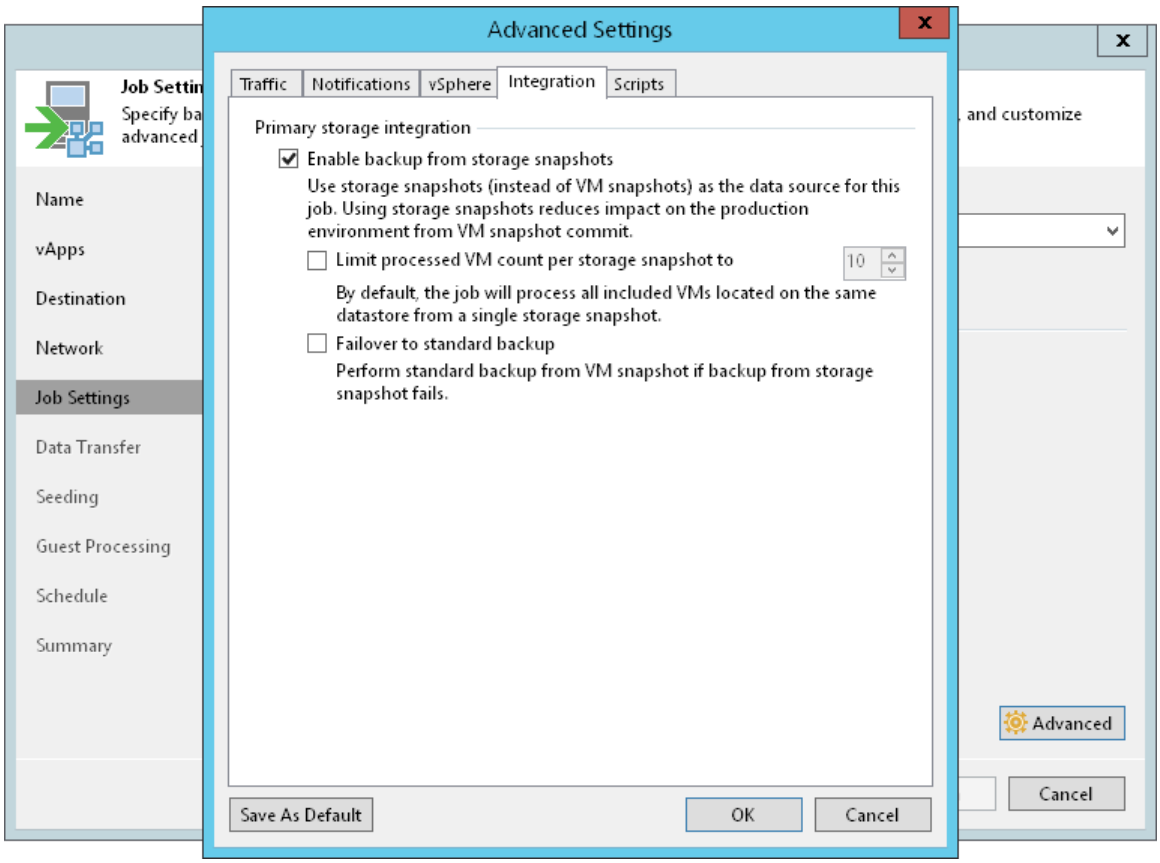
To specify storage integration settings for the replication job:

1. At the **Job Settings** 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.

4. If you want to replicate vApps with multiple 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 example, you add to the job vApps with 30 VMs whose disks are located on the same volume and set the **Limit processed VM count per storage snapshot to** option to 10. Veeam Backup & Replication will divide all VMs into 3 groups and create 3 storage snapshots from which it will read VM data.

5. If the backup infrastructure is configured incorrectly, for example, the backup proxy does not meet the necessary requirements, Backup from Storage Snapshots will fail and VMs residing on the storage systems will not be processed by the job at all. To fail over to the regular VM processing mode and process such VMs in any case, 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. Click the **Scripts** tab.

3. If you want to execute custom scripts before and/or after the replication 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.

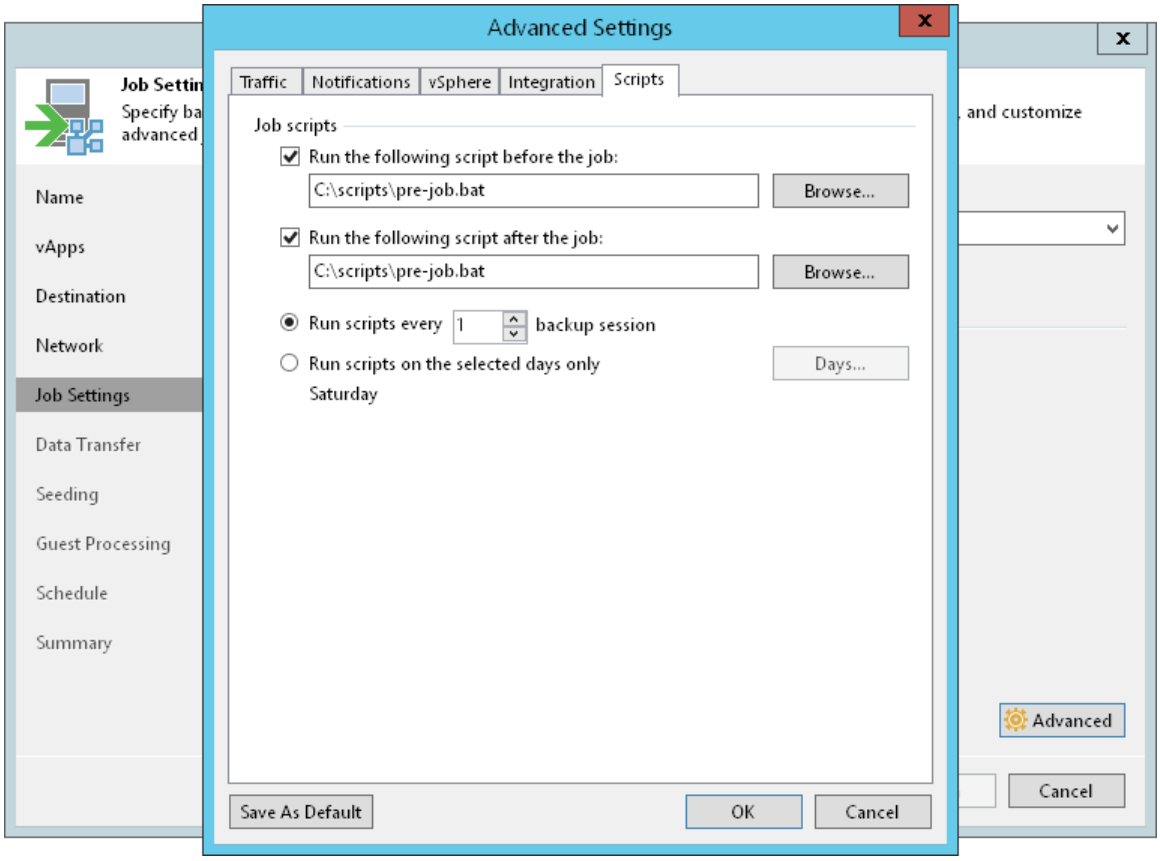
   You can select to execute pre- and post-replication actions after a number of job sessions or on specific week days.

   - If you select the **Run scripts every... backup session** option, specify the number of the replication job sessions after which scripts must be executed.
   - If you 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 replication 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 10. Specify Data Transfer Settings

At the **Data Transfer** step of the wizard, select backup infrastructure components that Veeam Backup & Replication will use for the replication process and choose a path for vApp data transfer.

1. If you plan to replicate vApp data within one site, the same proxy can act as the source and target proxy. For off-site replication, you must deploy at least one proxy in each site to establish a stable connection for vApp data transfer across sites.

   Click **Choose** next to the **Source proxy** and **Target proxy** fields to select backup proxies for the job. In the **Backup Proxy** window, you can choose automatic proxy selection or assign backup proxies explicitly.

   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source and target datastores and automatically assign optimal proxy resources for processing vApp data.

     Veeam Backup & Replication assigns resources to vApp included in the replication job one by one. Before processing a new vApp from the list, Veeam Backup & Replication checks available backup proxies. If more than one 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 proxy for vApp processing.

   - If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that the job can use. It is recommended that you 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.

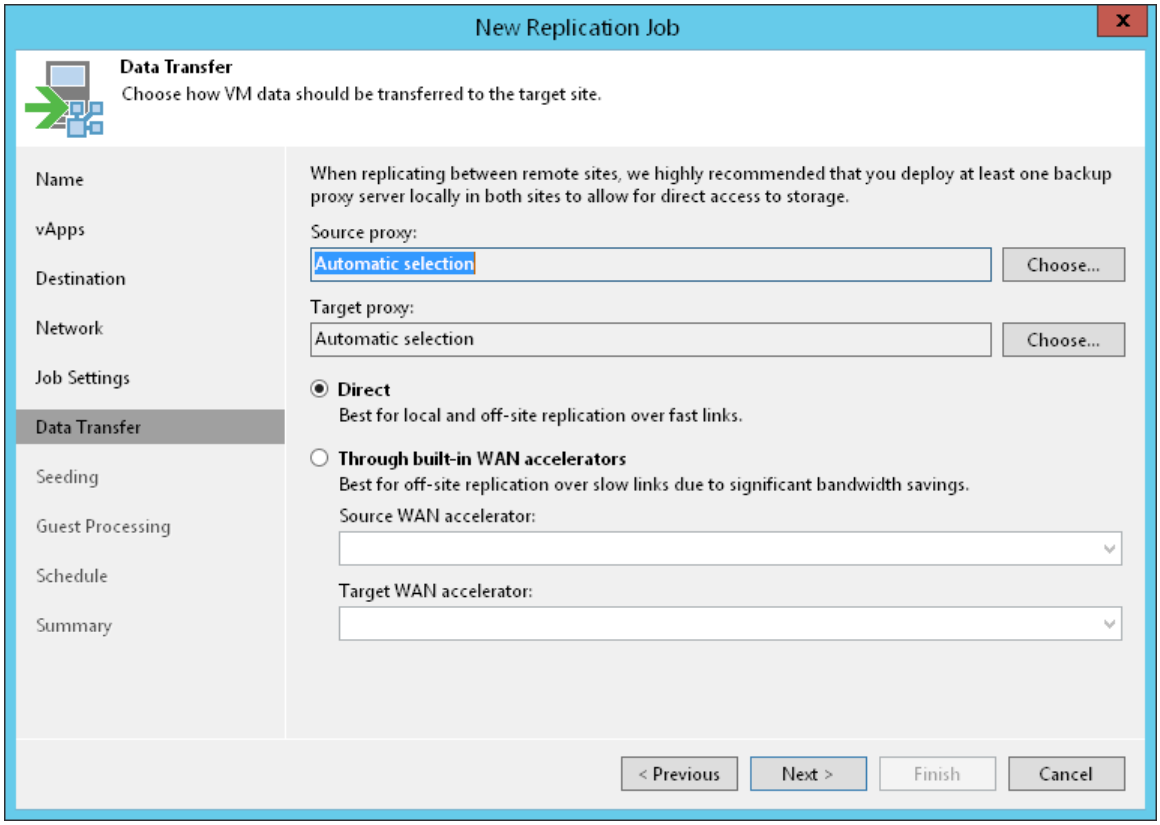
2. Select a path for vApp data transfer:

   - To transport vApp data directly using backup proxies to the target datastore, select **Direct**.

   - To transport vApp 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.

     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. Alternatively, you can create one replication job for all vApps 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 11. Configure Seeding and Mapping Settings

The **Seeding** step is available if you have selected the **Replica seeding** option at the **Name** step wizard. You can use this step to configure replica seeding and mapping for a vCD replication job.

If you use replica seeding or mapping, make sure that you select the following backup infrastructure components for the job properly: source-side backup repository to keep metadata and proxies. It is recommended that you explicitly assign proxies in the production site and disaster recovery (DR) site. For more information, see Specify Data Transfer Settings.

Configuring Replica Seeding

If you plan to replicate to a remote DR site over WAN or low-bandwidth network, you can use replica seeding. Replica seeding helps reduce the amount of VM data transferred over the network.

Replica seeding can be used if you have a backup for the replicated vApp in the backup repository located in the DR site. In this case, you can point the replication job to this backup. During the first session of the replication job, Veeam Backup & Replication uses this backup file as a "seed". Veeam Backup & Replication restores the image of VMs that are added to vApps from the backup and registers the vApp replica on the target host. After that, Veeam Backup & Replication synchronizes the vApp replica with the source vApp. All subsequent incremental replication runs are performed in the regular manner.

Before you start a replication job that uses replica seeding, you must perform a number of preparatory tasks:

1. Create a backup of the vApp that you plan to replicate. This backup will act as a seed for a replication job.

   To create a backup, configure a backup job that points to an on-site backup repository. Run the job to create a full backup.

   If you have previously created a backup containing all necessary vApps, there is no need to configure and run a new backup job.

   For seeding, you can use any existing backup 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 vApp 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. If you do not have a backup repository in the DR site, you will need to create it.

   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.

3. After the backup is copied to the backup repository in the DR site, perform rescan of this backup repository. Otherwise, Veeam Backup & Replication will not be able to detect the copied backup.

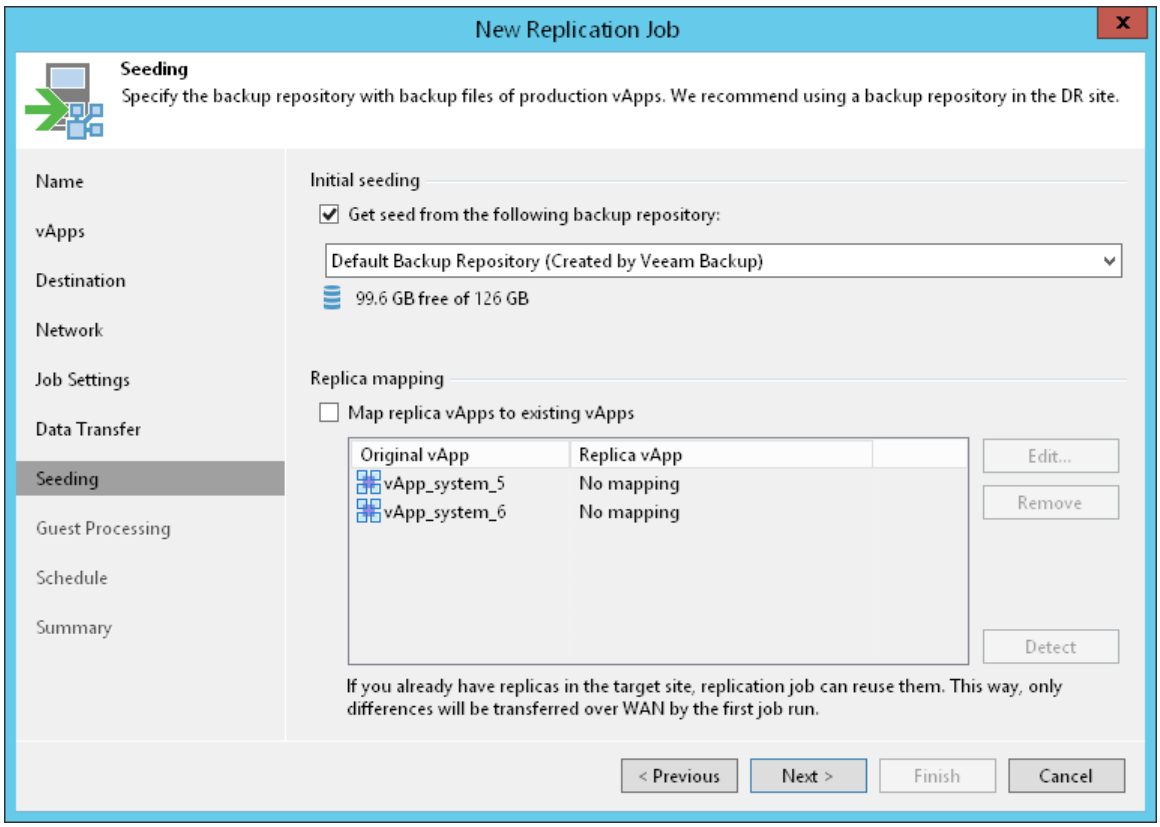
> **IMPORTANT**
>
> When you start the replication job, Veeam Backup & Replication will attempt to restore all VMs added to the job from the vApp seed that you have specified. If a VM is not found in the vApp seed, the VM will be skipped from replication.

When you complete the preliminary steps, you can configure replica seeding settings for the job.

1. In the **Initial seeding** section, select the **Get seed from the following backup repository** check box.
2. From the list of available backup repositories, select the repository where your replica seed is stored.

IMPORTANT

You cannot use a backup located in the scale-out backup repository as a seed for a replication job.



Configuring Replica Mapping

If a replica for the vApp that you plan to replicate already exists on the target host in the DR site, you can use replica mapping. Replica mapping helps reduce the amount of vApp data transferred over the network.

To use replica mapping, you must point the replication job to a vApp replica on the host in the DR site. During the first session of the replication job, Veeam Backup & Replication will calculate the difference between the source vApp and vApp replica and copy necessary data blocks to synchronize the vApp replica to the latest state of the source vApp. All subsequent incremental replication sessions will be performed in the regular manner.

TIP

If there is no existing vApp replica in the DR site, you can restore a vApp from the backup and map it to the original vApp.

To set up replica mapping:

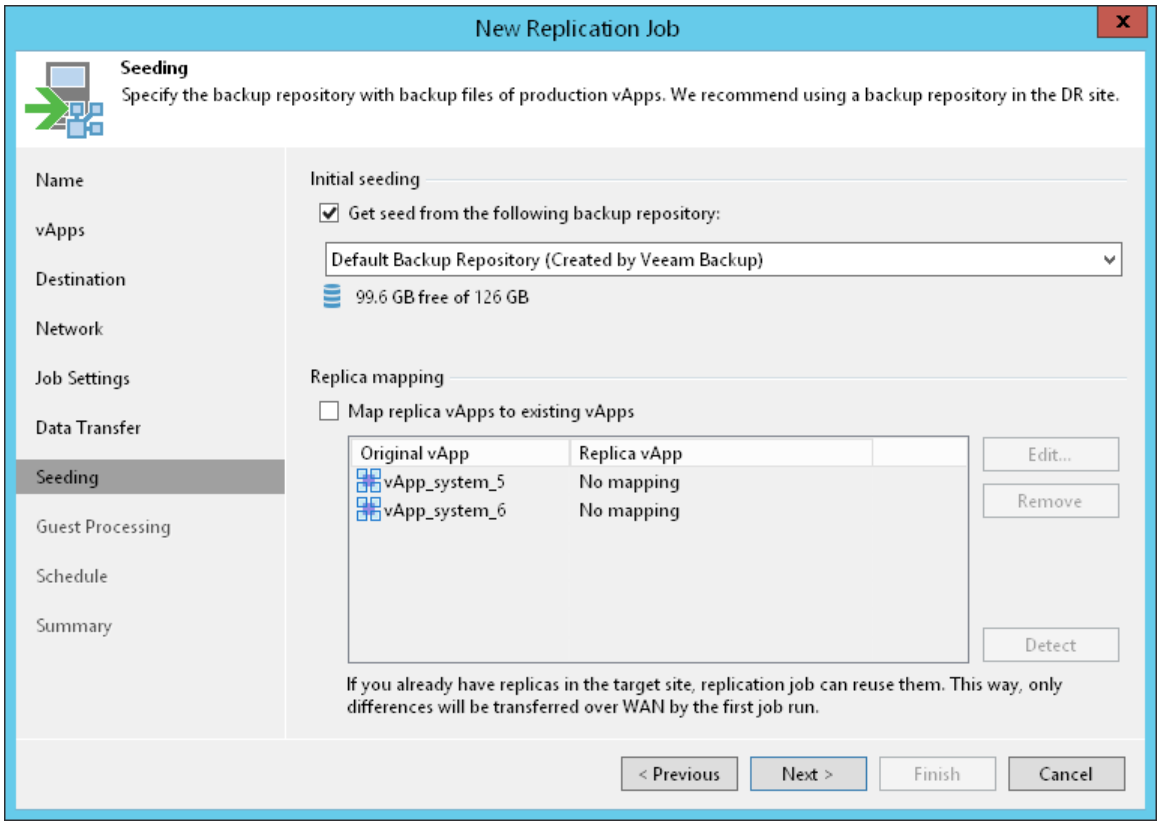
1. Select the **Map replica vApps to existing vApps** check box.
2. Click **Detect**. Veeam Backup & Replication will scan the destination location to detect existing vApp replicas. If any matches are found, Veeam Backup & Replication will populate the mapping table.

   If Veeam Backup & Replication does not find a match, you can map a vApp to its vApp replica manually. To do this, select a production vApp from the list, click **Edit** and choose an existing vApp replica. To facilitate selection, use the search field at the bottom of the window.

To break a mapping association, select the vApp in the list and click **Remove**.

IMPORTANT

The mapping list does not display VMs added to the list of exclusions. For more information, see Exclude objects from replication job.



Step 12. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, you can enable the following settings for guest OS processing of VMs added to vApps you want to replicate:

- Application-aware processing
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- File exclusion settings
- 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 replication job session or after the replication job completes).

> **TIP**
>
> You can specify guest processing settings for vApps and for specific VMs.

You must specify a user account that will be used to connect to the vApp guest OS and deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components:

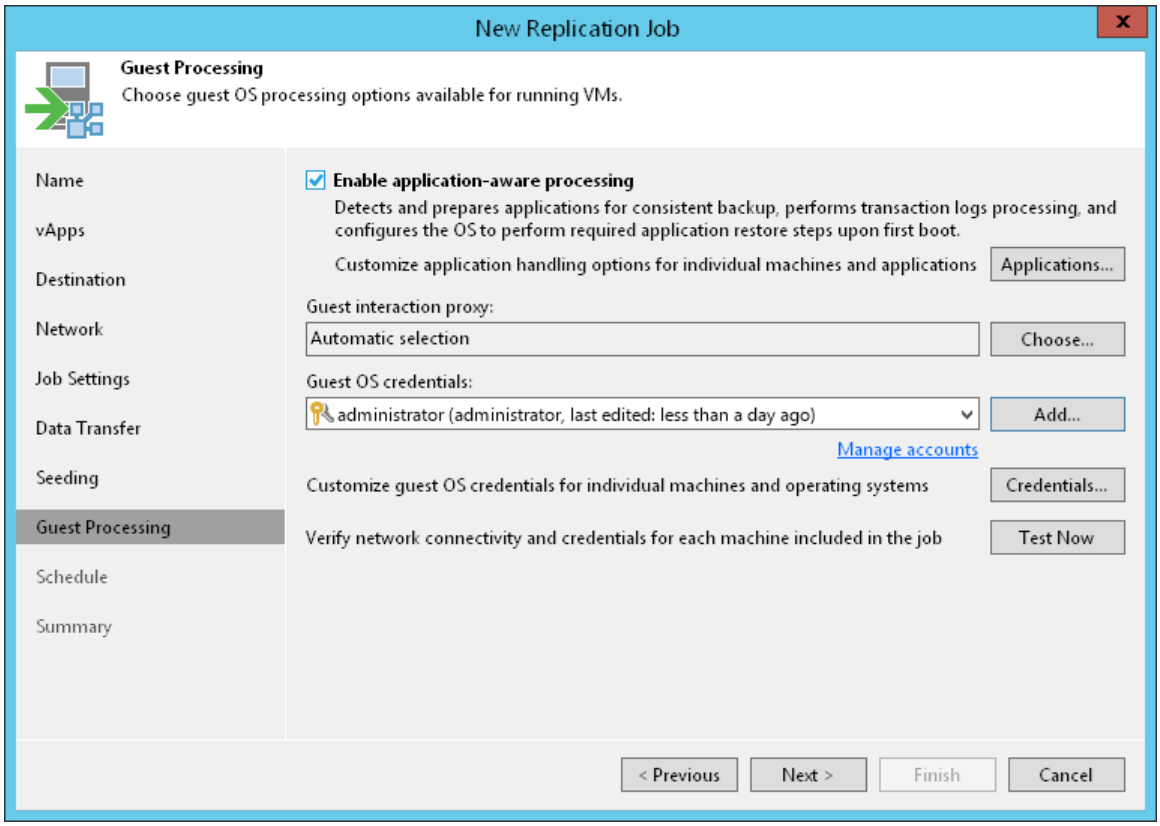
1. From the **Guest OS credentials** list, select a user account with local Administrator privileges on the VM guest OS. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials. For more information, see Guest Processing.

   > **NOTE**
   >
   > Mind the following:
   >
   > - Local accounts do not support Kerberos authentication. To authenticate with Microsoft Windows guest OS using Kerberos, specify an Active Directory account.
   > - If you plan to use Kerberos authentication, check limitations and requirements listed in the Guest Processing 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.
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.

> **NOTE**
>
> The guest interaction proxy functionality is available in the Enterprise and Enterprise Plus editions of Veeam Backup & Replication.

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.



Application-Aware Processing

If you add to the replication job vApps with VMs running VSS-aware applications, you can enable application-aware processing to create transactionally consistent replicas. The transactionally consistent replica 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 vApp 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 replication process if any VSS errors occur.

- Select **Try application processing, but ignore failures** if you want to continue the replication process even if VSS errors occur. This option is recommended to guarantee completion of the job. The created VM replica 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 replicas 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 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.

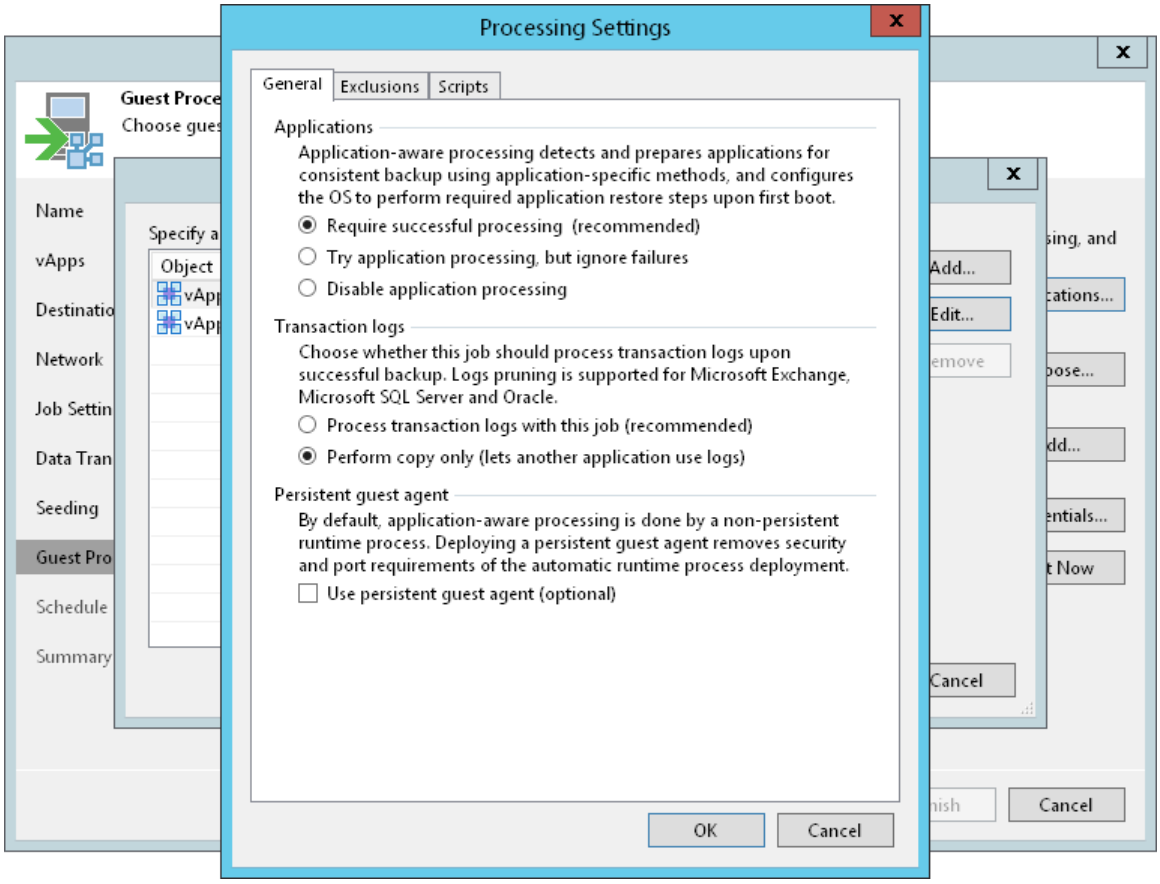
      [For Microsoft SQL Server VMs and Oracle VMs] You will have to specify settings for transaction log handling on the **SQL** and **Oracle** tabs of the **VM Processing Settings** window. For more information, see Microsoft SQL Server Transaction Log Settings and Oracle Archived Log Settings.

   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 replica for the selected VM. The copy only replica preserves the chain of full/differential backup files and transaction logs on the VM. For more information on copy-only option, 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 replication job starts, and removes the runtime components as soon as the replication 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.

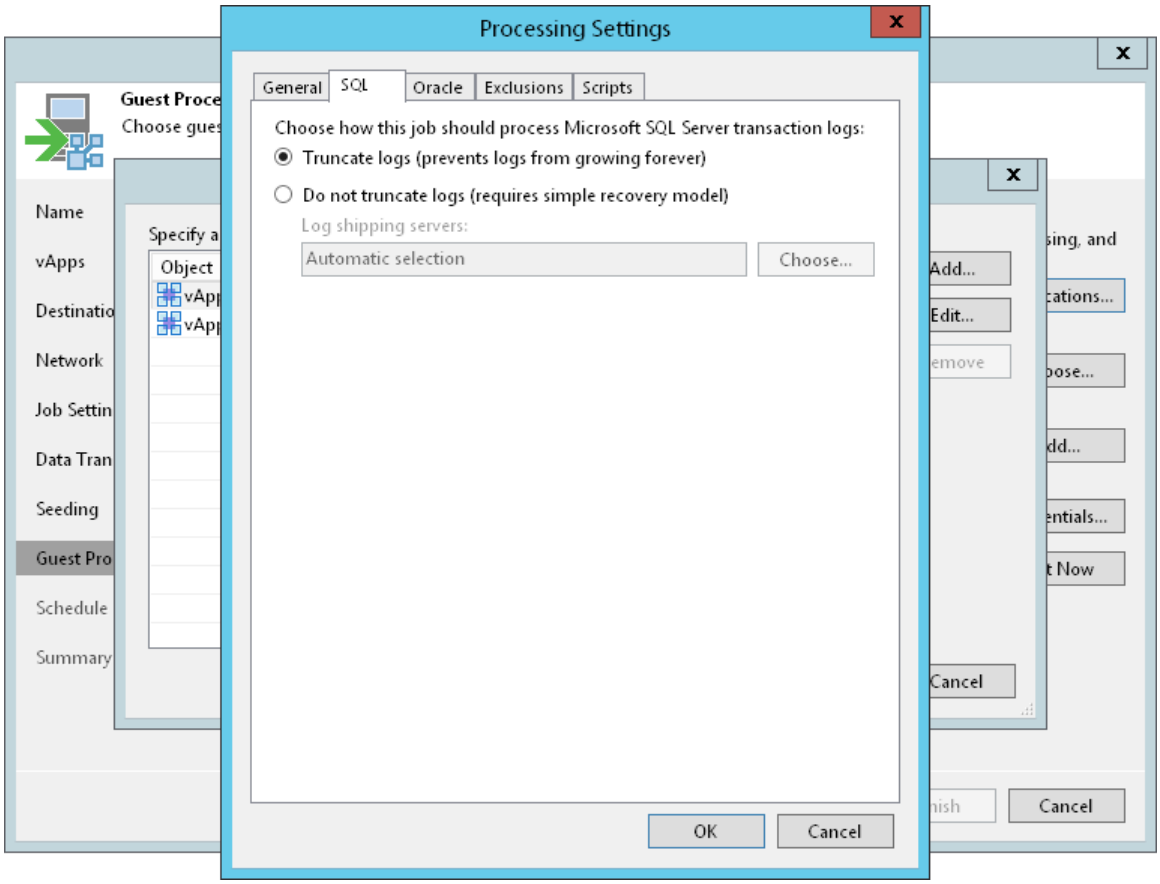


Microsoft SQL Server Transaction Log Settings

If you replicate 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 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.

- 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 replica, you will be able to apply transaction logs to get the database system to the necessary point in time between replication job sessions.



Oracle Archived Log Settings

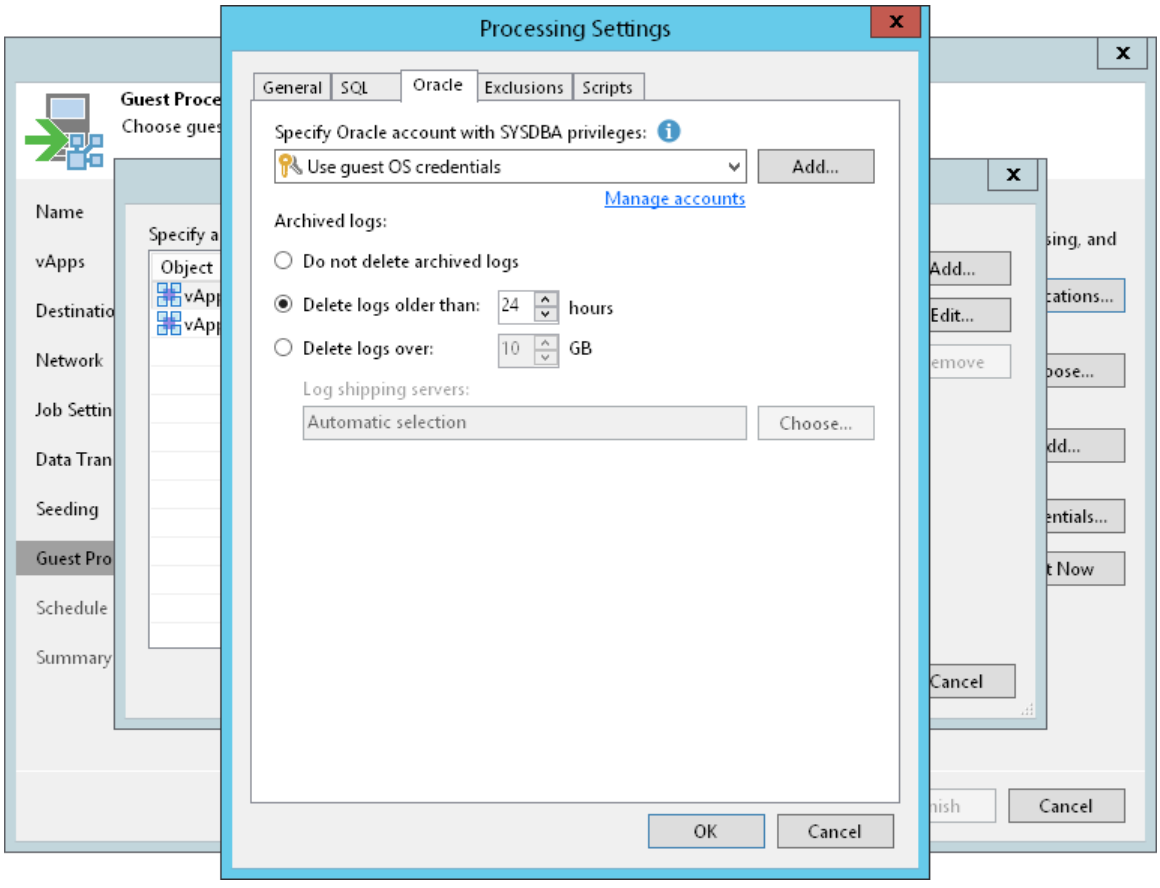
If you replicate an Oracle VM, you can specify how Veeam Backup & Replication must process transaction logs:

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 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 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 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.



Guest OS File Exclusion Settings

If you do not want to replicate specific files and folders on the VM guest OS Windows, you can exclude them from the VM replica.

To define files and folders that must be excluded:

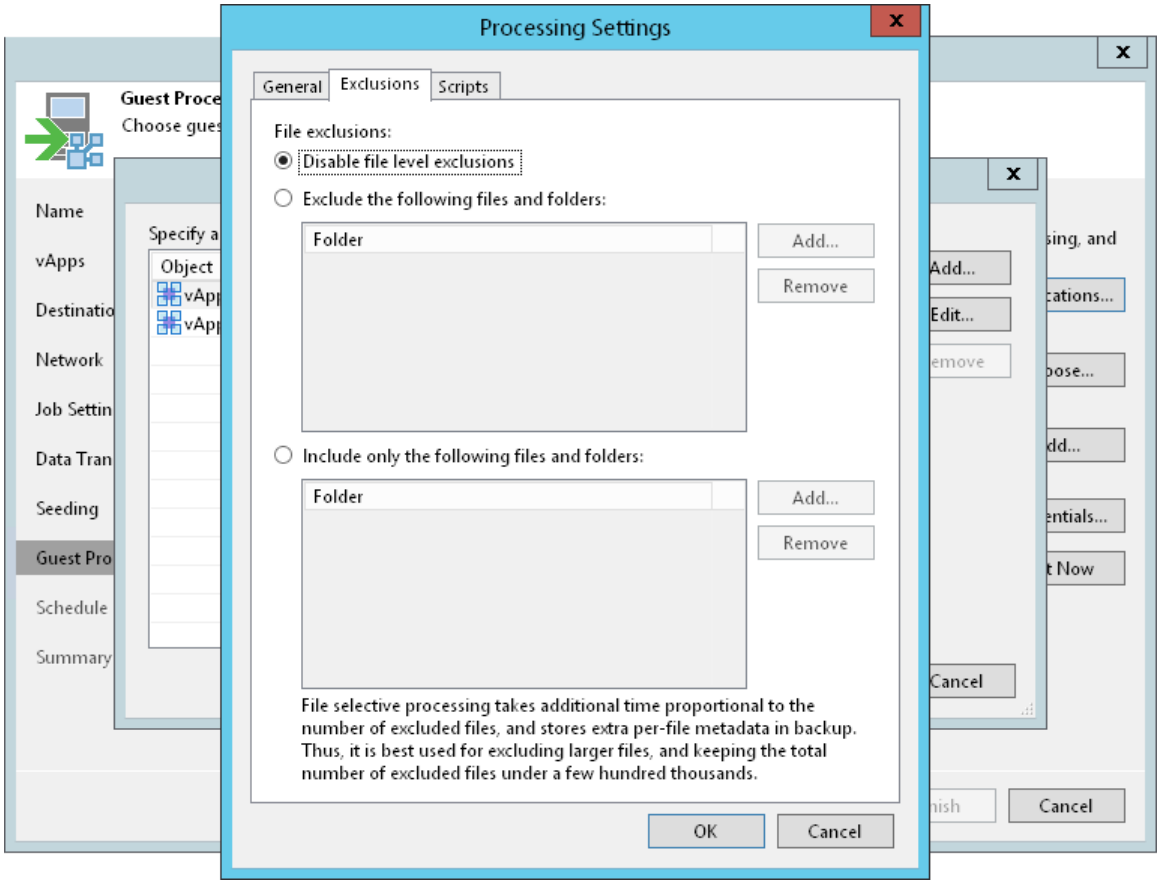
1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.

3. In the displayed list, select the VM and click **Edit**.

> **NOTE**
>
> 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 VM replica:
   - Select **Exclude the following files and folders** to remove the individual files and folders from the VM replica.
   - Select **Include only the following files and folders** to leave only the specified files and folders in the VM replica.
5. Click **Add** and specify what files and folders you want to include or exclude. Click **OK**.

   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.



Pre-Freeze and Post-Thaw Scripts

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.

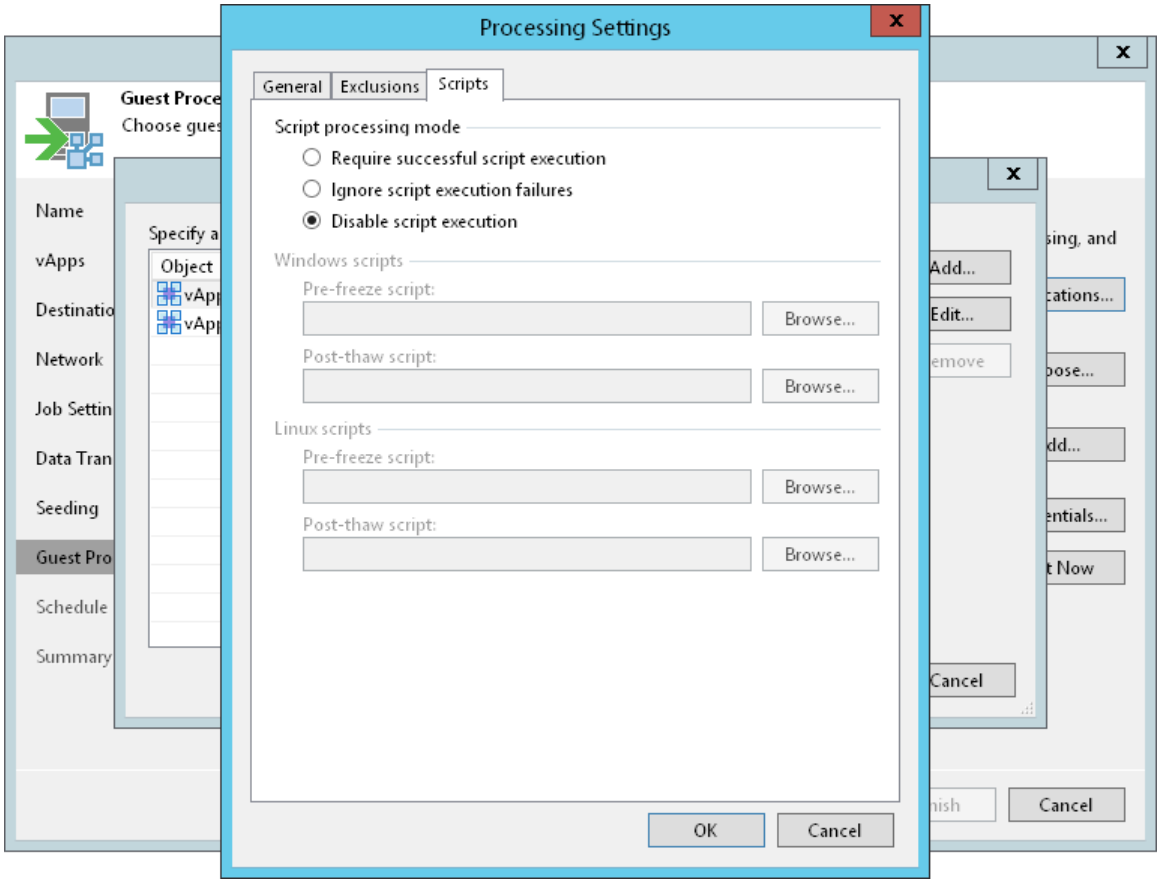
To specify pre-freeze and post-thaw scripts for the job:

1. 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 replication process if the script fails.
   - 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 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.

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 13. Define Job Schedule

At the **Schedule** step of the wizard, select to run the vCD replication job manually or schedule the job to run on a regular basis.

To define a job schedule:

1. Select the **Run the job automatically** check box. If you do not select this check box, you will have to start the job manually to perform vCD 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 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.

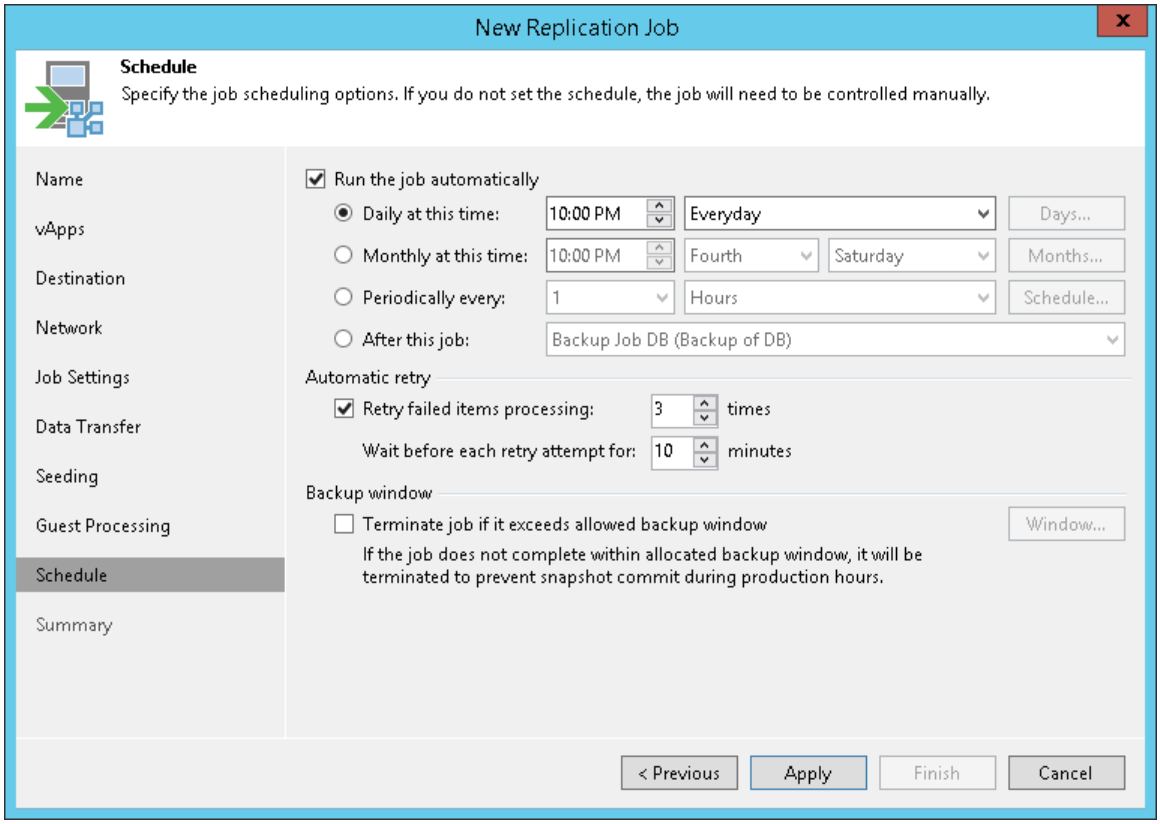
     > **NOTE**
     >
     > The **After this job** function will only start a job if the first job in the chain is started automatically by schedule. If the first job is started manually, jobs chained to it will not be started.

3. In the **Automatic retry** section, define whether Veeam Backup & Replication should attempt to run the job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed vApps 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:

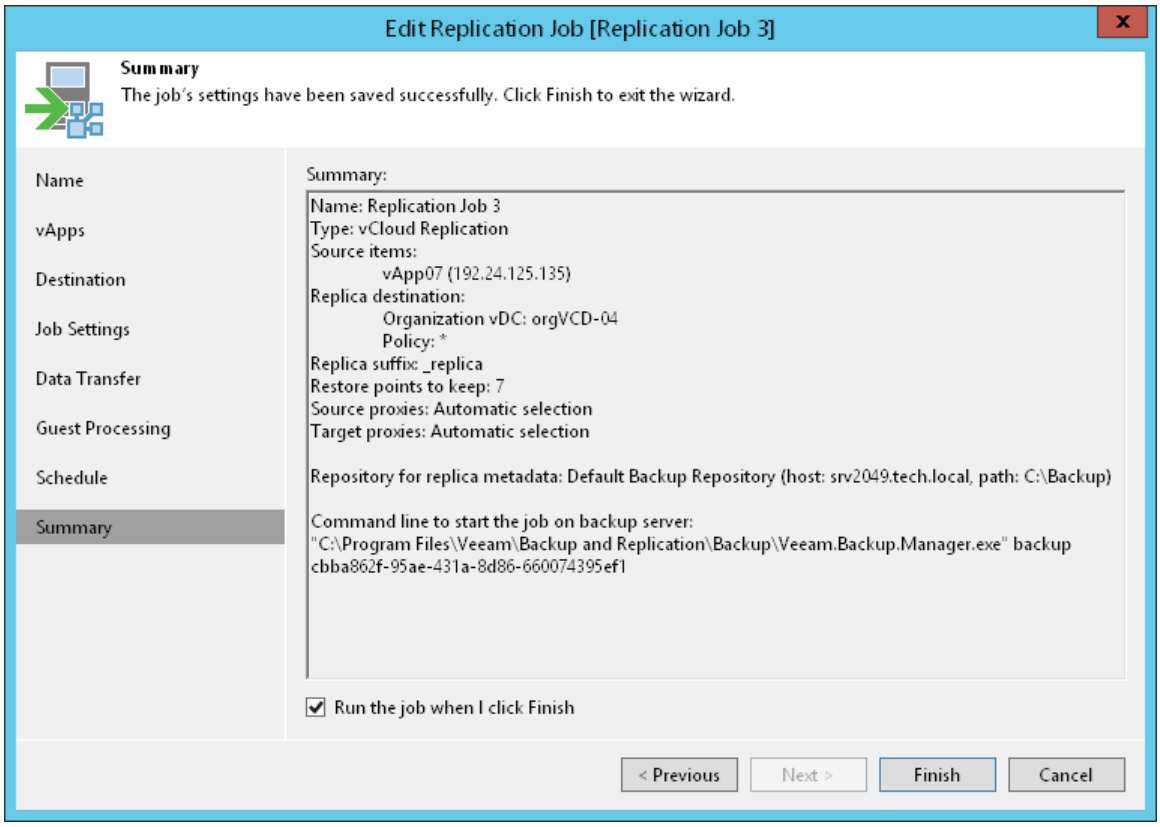
   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** section, define the allowed hours and prohibited hours for vApp replication. If the job exceeds the allowed window, it will be automatically terminated.



Step 14. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the vCD replication job. If you want to start the job right after you close the wizard, leave the **Run the job when I click Finish** check box selected, otherwise clear the check box. Then click **Finish** to close the wizard.



Managing vCD Replication Jobs

After you create vCD replication jobs, you can edit, disable and delete them.

To view all created vCD replication jobs, open the **Home** view and navigate to the **Jobs** > **Replication** node. The working area displays the full list of the created replication jobs. Here, you can manage the jobs.

You can manage vCD replication jobs with the following options:

- Retrying vCD Replication Jobs
- Editing vCD Replication Jobs
- Disabling vCD Replication Jobs
- Deleting vCD Replication Jobs

Retrying vCD Replication Jobs

The retry option is necessary when a vCD job fails and you want to retry this operation again. When you perform a retry, Veeam Backup & Replication restarts the same procedure only for the failed VMs added to the vApp and will not process VMs that have been processed successfully. As a result, the replication job will take less time and will not consume as many resources as when processing a whole vApp.

To a retry a vCD replication job that is failed:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary vCD replication job and select **Retry** on the ribbon. Alternatively, you can right-click the necessary vCD replication job and select **Retry**.

Editing vCD Replication Jobs

To edit a vCD replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary vCD replication job and select **Edit** on the ribbon. Alternatively, you can right-click the necessary vCD replication job and select **Edit**.
4. Follow the instructions provided in Creating vCD Replication Job.

Disabling vCD Replication Jobs

After you disable the job, Veeam Backup & Replication will pause it and will not process this job according the specified schedule. You can enable a disabled job at any time.

To disable a vCD replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.

3. In the working area, select the necessary policy and select **Disable** on the ribbon. Alternatively, you can right-click the necessary policy and select **Disable**.

> **TIP**
>
> To enable a disabled vCD replication job, select it and click **Disable** once again.

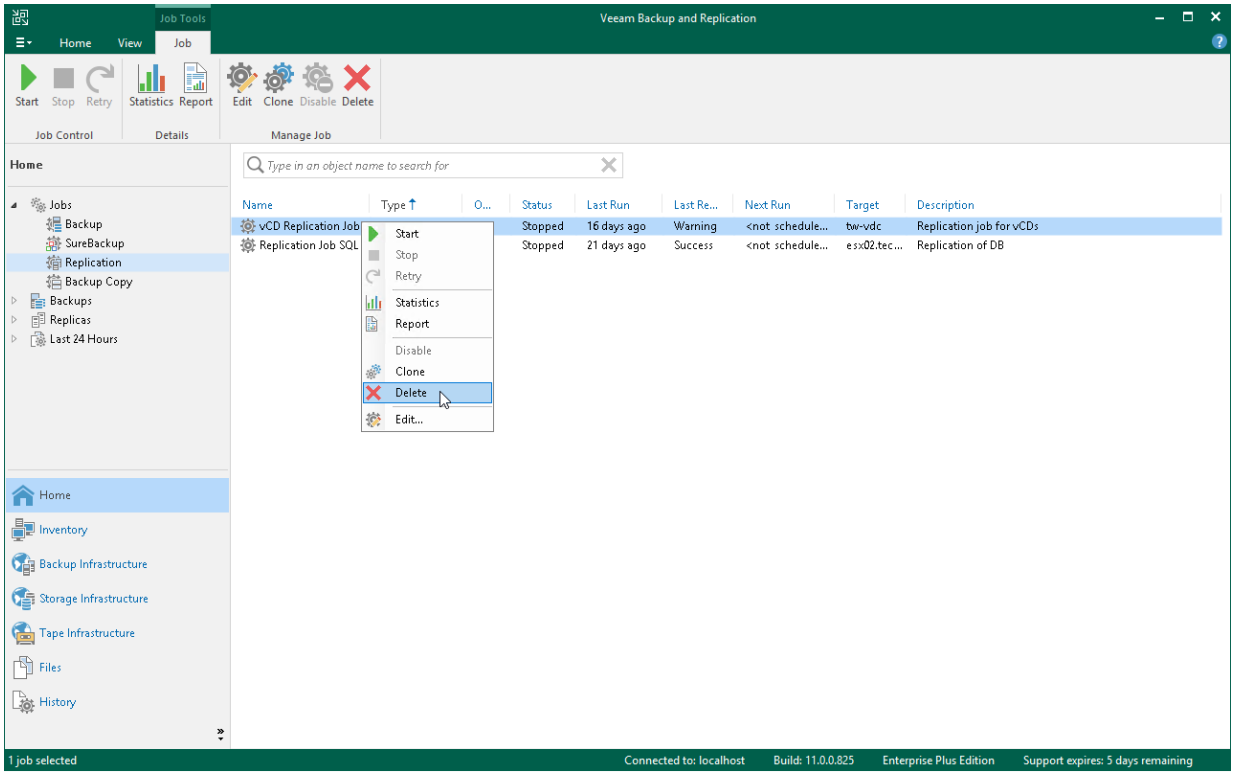
Deleting vCD Replication Jobs

After you delete the job, the replicas created by this job are displayed under the **Replicas** node.
Veeam Backup & Replication allows you to delete only disabled vCD replication jobs. You can also permanently delete a job from Veeam Backup & Replication and from the configuration database.

Veeam Backup & Replication allows you to delete only disabled vCD replication jobs.

To delete a vCD replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary vCD replication job and select **Delete** on the ribbon. Alternatively, you can right-click the necessary vCD replication job and select **Delete**.



Managing vCD Replicas

You can perform the following operations with vCD replicas:

- View vApp replica properties
- Rescan vApp replicas
- Remove a vApp replica from configuration
- Delete a vApp replica from disk

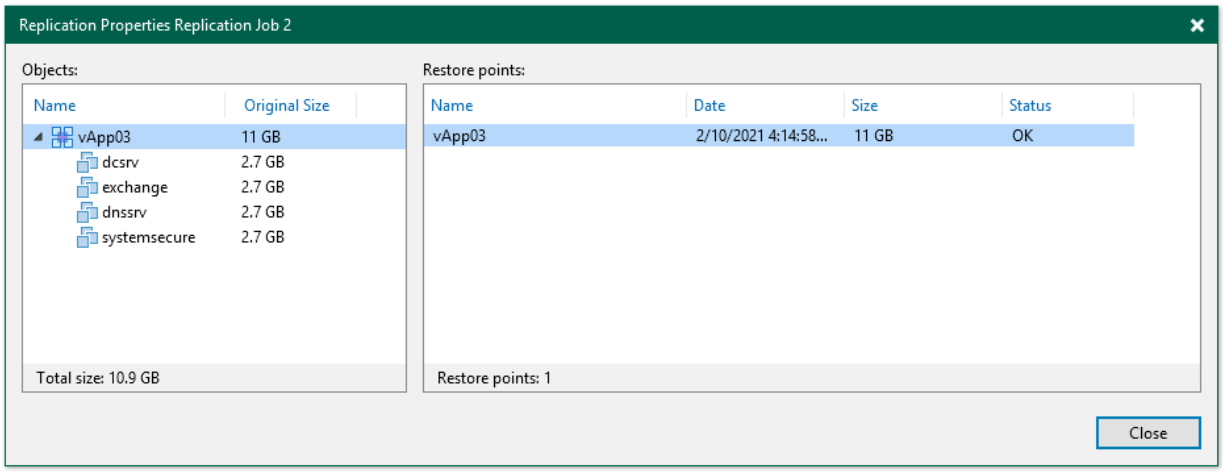
Viewing Replica Properties

You can view summary information about 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, select the necessary vCD replica and select **Properties**. Alternatively, you can right-click the necessary vCD replica and select **Properties**.



Rescanning Replicas

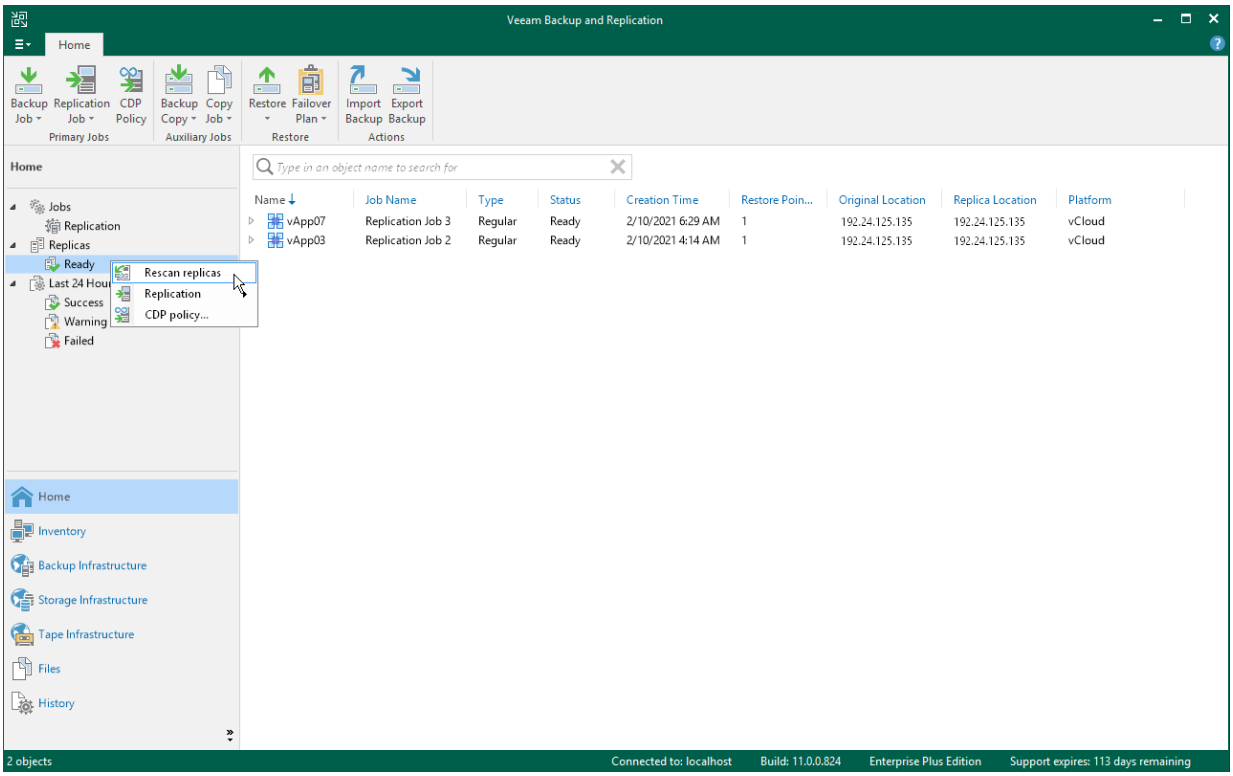
Rescan is a process that allows Veeam Backup & Replication to verify that the data written in the configuration database about the disaster recovery (DR) site, its components and VM containers that are added to the DR site is up to date. The replica rescan process is performed the following way:

1. Veeam Backup & Replication gathers information on replicas that are currently available in the DR site.
2. Veeam Backup & Replication compares this information with information stored in the configuration database about replicas from this DR site.

3. If information about replicas from the DR site differs from information stored in the configuration database about these replicas, Veeam Backup & Replication updates 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 from Configuration

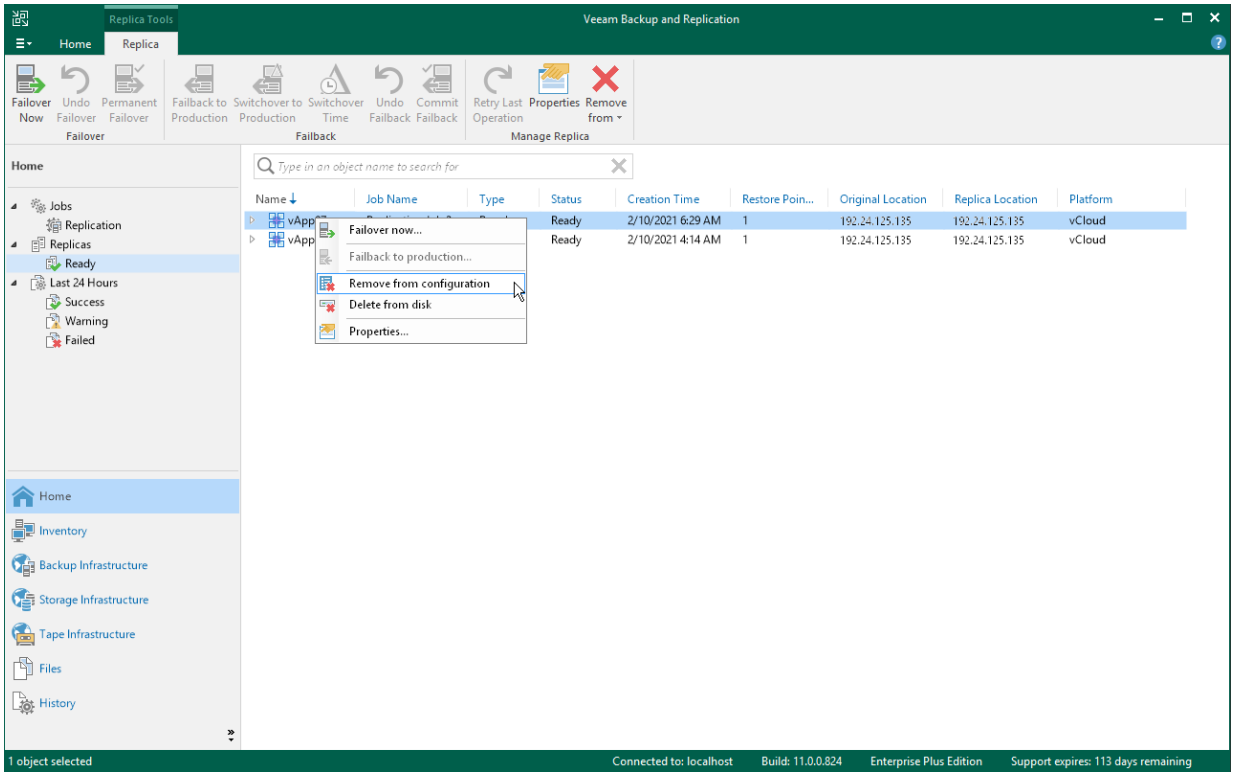
The **Remove from configuration** operation is necessary, if you want to remove records about vApp replicas from the backup infrastructure and configuration database. After you remove vApp replicas, the replicated vApp remain on target hosts. If necessary, you can start them manually after you perform the **Remove from configuration** operation.

> **NOTE**
>
> - The **Remove from configuration** operation can be performed only for vApp replicas in the *Ready* state. If the vApp replica is in the *Failover* or *Failback* state, this option is disabled.
> - When you perform the **Remove from configuration** operation for a vApp that is replicated as a standalone object, Veeam Backup & Replication removes this vApp from the initial replication job. When you perform the **Remove from configuration** operation for a vApp that is replicated as part of a VM container, Veeam Backup & Replication adds this vApp to the list of exclusions in the initial replication job. For more information, see Step 5. Exclude Objects from Replication Job

To remove records about vApp replicas from the backup infrastructure and configuration database:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.

3. In the working area, select the vApp replica and click **Remove from** > **Configuration** on the ribbon.Alternatively, right-click the replica and select **Remove from configuration**.



Deleting from Disk

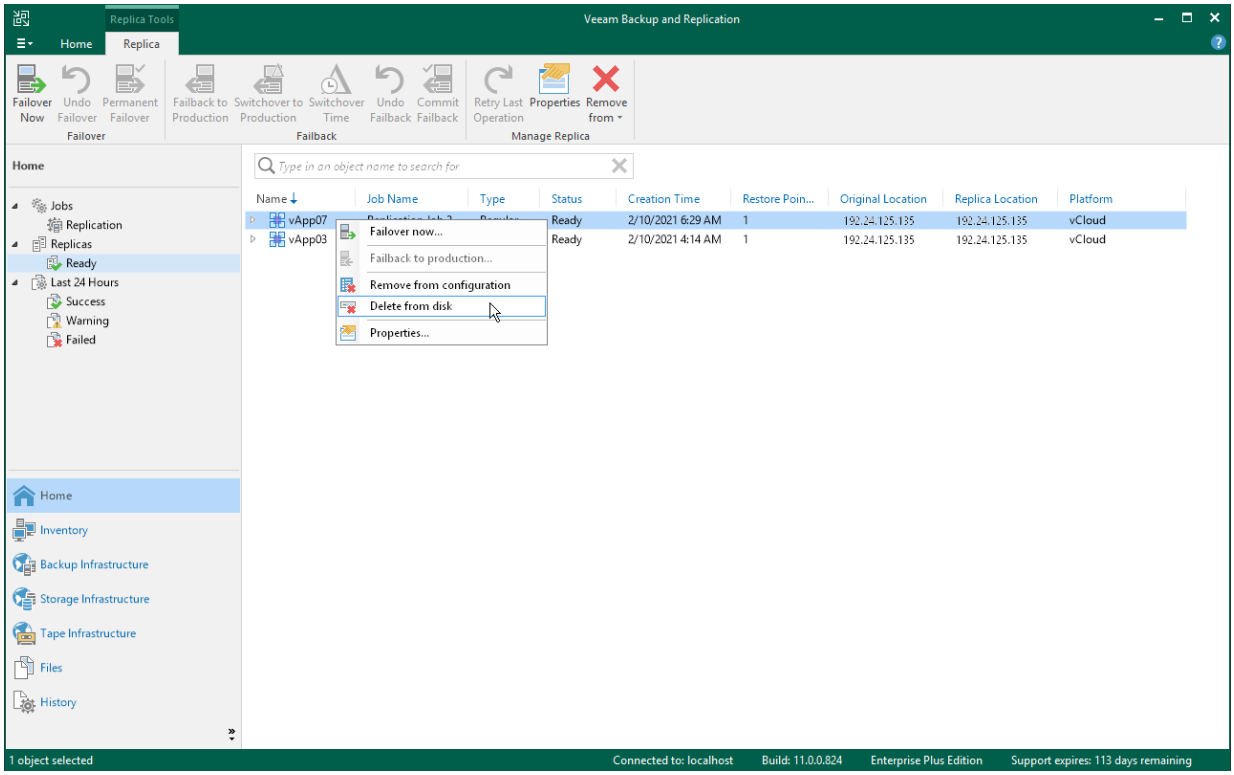
To delete vApp replicas files from the destination storage, you can use the **Delete from disk** operation.

> **NOTE**
>
> - 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 vApp replicas are in the *Ready* state. If the vApp replicas are 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 vApp replica and click **Remove from** > **Disk** on the ribbon. Alternatively, you can right-click the VM replica and select **Delete from disk**.



Failover and Failback for vCD

Failover and failback are operations that allow you to manage your production and disaster recovery (DR) sites if a disaster strikes. Failover is a process of switching from the original vApp on the source organization VDC to its vApp replica on a target organization VDC that is set up as the DR site. Failback is a process of returning from the vApp replica to the original vApp or a new vApp.

Veeam Backup & Replication provides the following failover and failback operations:

- **Perform failover**

  When you perform failover, you shift from the original vApp in the production organization VDC to the vApp replica in the disaster recovery organization VDC. Failover is an intermediate step that needs to be finalized: you can perform permanent failover, perform failback or undo failover.

  For more information on how failover is performed, see Failover.

- **Perform permanent failover**

  When you perform permanent failover, you permanently switch from the original source vApp to a replica vApp and use this replica as the production vApp. The original vApp is excluded from vCD replica processing.

  For more information on how permanent failover is performed, see Permanent Failover.

- **Undo failover**

  When you undo failover, you switch back to the original vApp and discard all changes made to the vApp replica while it was running. For example, you can use the undo failover scenario if you have failed over to the vApp replica for testing and troubleshooting purposes, and you do not need to synchronize the original vApp 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 vApp and send to the original vApp all changes that took place while the vApp replica was running. If the source organization VDC is not available, you can recover a vApp with the same configuration as the original vApp 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 or recovered vApp but not published. You must test whether the original or recovered vApp 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 or recovered vApp 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 or recovered vApp is not working as expected and you want to get back to the vApp replica.

    For more information on how failback undo is performed, see Failback Undo.

> **IMPORTANT**
>
> Currently, vCloud Director replication does not support failover plans.

Related Topics

- Replica Failover
- Permanent Failover
- Failover Undo
- Failback
- Failback Undo
- Failback Commit

Replica Failover

Failover is a process of switching from the original vApp in the production organization VDC to its vApp replica in the disaster recovery organization VDC. During failover, Veeam Backup & Replication recovers the vApp replica to the required restore point and switches from the original vApp to its replica. As a result, you have a fully functional vApp within several minutes, and your users can access services and applications with minimum disruption.

> **TIP**
>
> You can fail over to replicas not only when a disaster strikes the production organization VDC, but also to test replicas for recoverability. If the original vApps and vApp replicas are located in the same network, consider temporary disconnecting the original vApps 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 vApp replica is changed from *Ready* to *Processing*. After failover completes successfully, the sate is changed from *Processing* to *Failover*.
3. Veeam Backup & Replication recovers a vApp replica to the required restore point.
4. Veeam Backup & Replication powers on the vApp replica.

   The original vApp still exists and failover does not change its state: if it was powered on when you performed failover, it remains powered on when failover completes; if it was powered off, it will remain in this state.
5. All changes made to the vApp replica while the replica is running in the *Failover* state are written to the delta file of the snapshot and stored in a target host.

Performing Failover

To perform failover, do the following.

Before You Begin

Before you fail over to a vApp replica, check the following prerequisites:

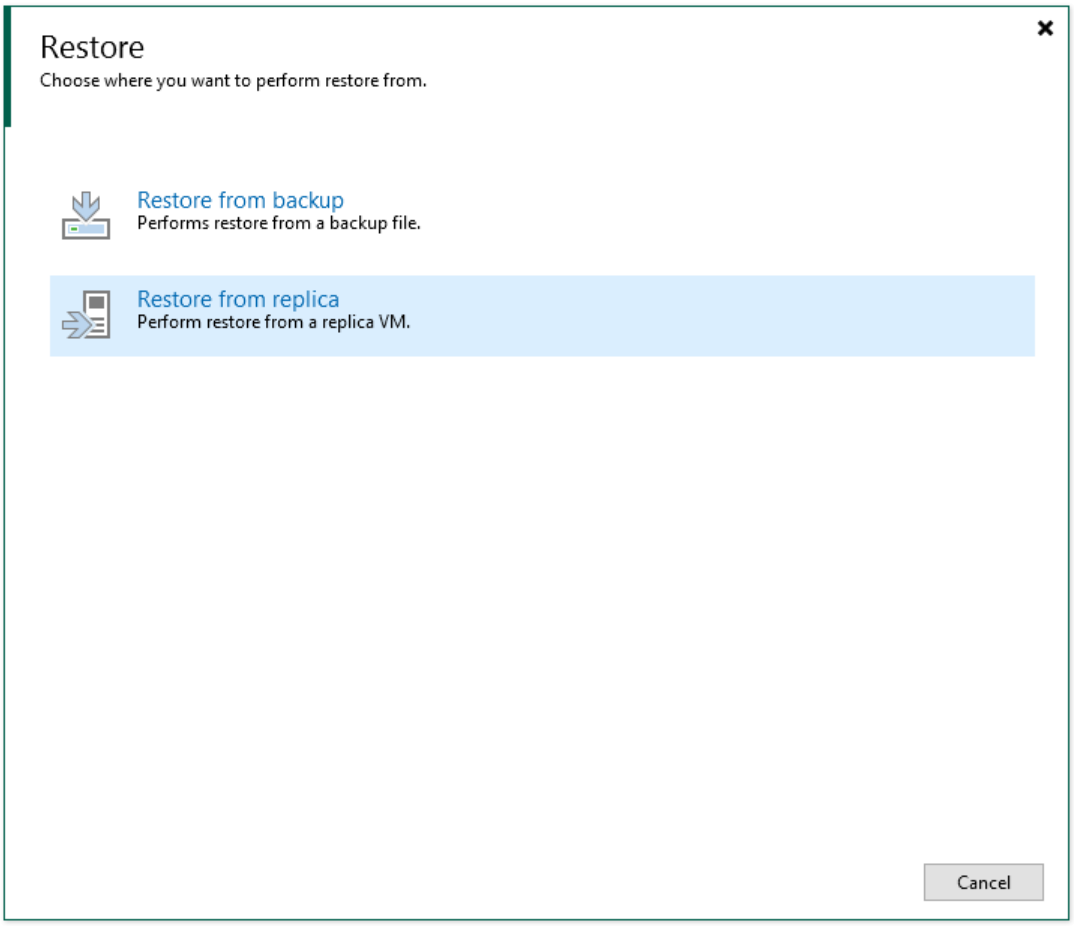
- The failover operation can be performed for vApp that have been successfully replicated at least once.

- vApp replicas must be in the *Ready* state.
- The VMs added to the target vApp must be powered off.

Step 1. Launch Failover Wizard

To launch the **Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vCloud Director** > **Restore from replica** > **Entire vApp** > **Failover vApp to a replica**.
- Open the **Home** view, in the inventory pane select **Replicas**. 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**. In the working area, right-click the necessary replica and select **Failover Now**.



Step 2. Select vApps to Fail Over

At the **vApp** step of the wizard, select one or more vApps for which you want to perform failover.

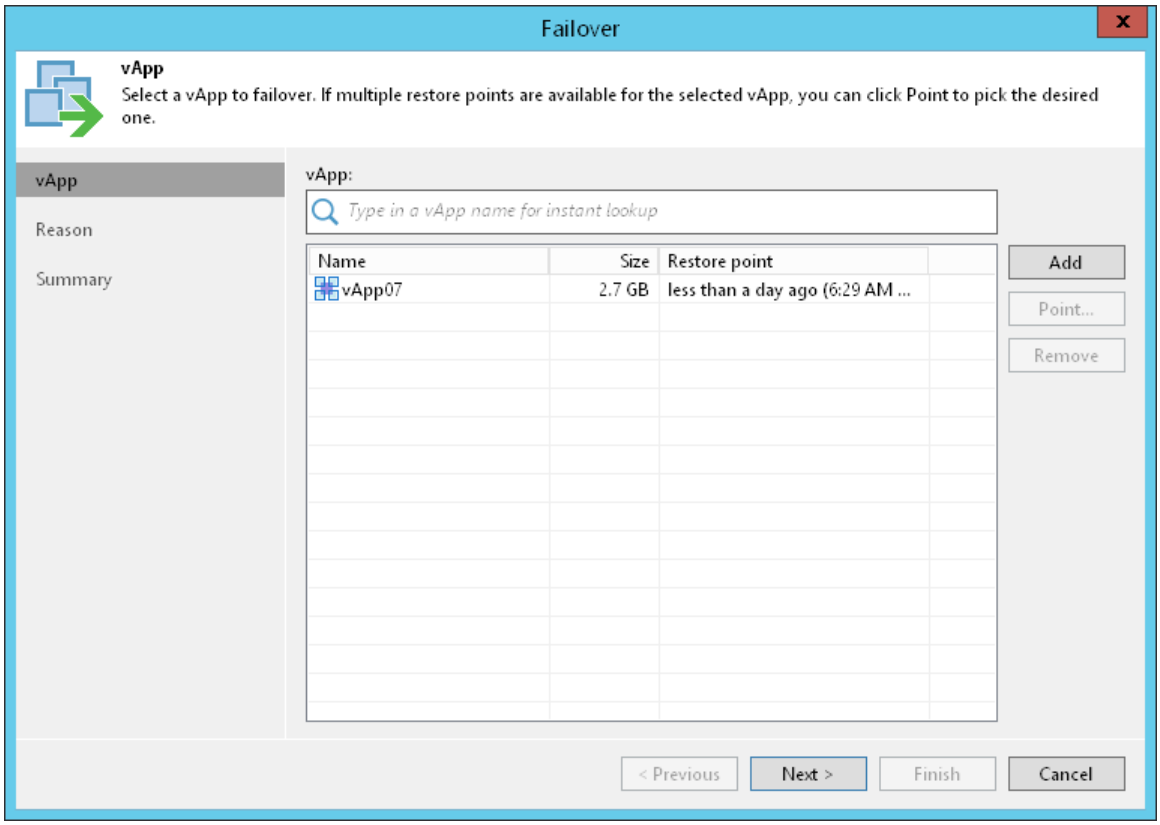
To select vApps:

1. Click **Add**.
2. Select where to browse for vApps:
   - **From infrastructure** – browse the virtual environment and select vApps.
   - **From replicas** – browse existing replicas and select vApps under replication jobs.

To quickly find vApps, you can use the search field at the top of the wizard.

1. Enter a vApp name or a part of it in the search field. Veeam Backup & Replication will display possible matches.
2. If the vApp is not in the list, click the **Show more** link to browse the virtual infrastructure.

> **NOTE**
>
> Make sure that vApps you select from the virtual environment have been successfully replicated at least once.

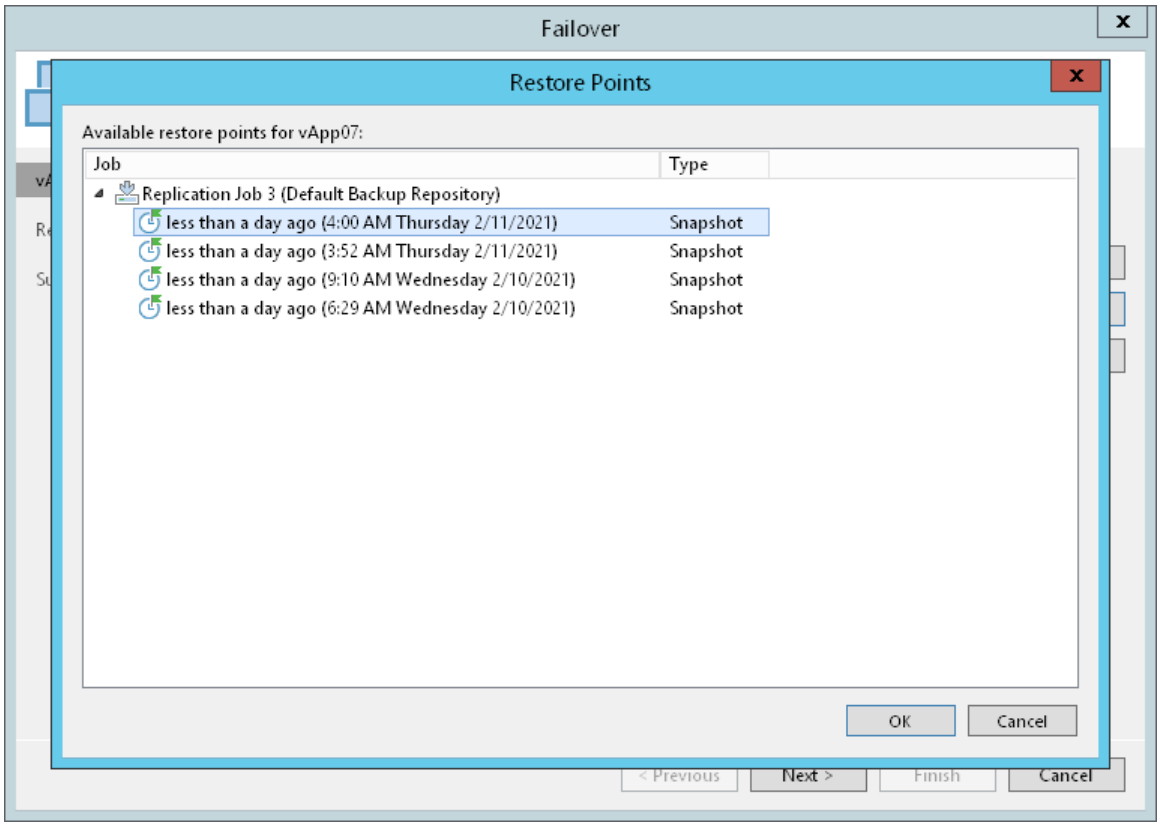


Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore point of the vApps replica. However, you can fail over to an earlier state of the vApps. If you have chosen to perform failover for several vApps, you can select the necessary restore point for every vApp in the list.

To select a restore point for a vApp:

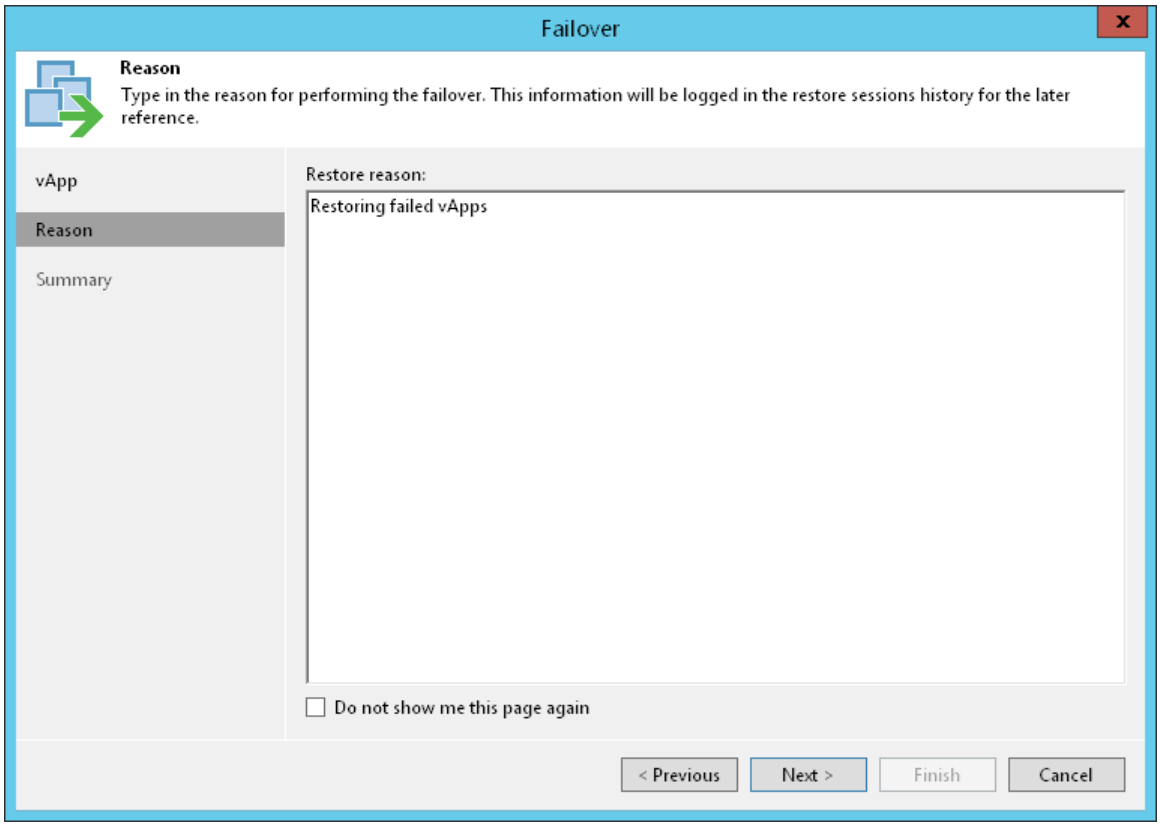
1. In the **vApp** list, select a vApp.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point to which you want to fail over.



Step 4. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the vApp replicas. The information you provide will be saved in the session history and you can reference it later.

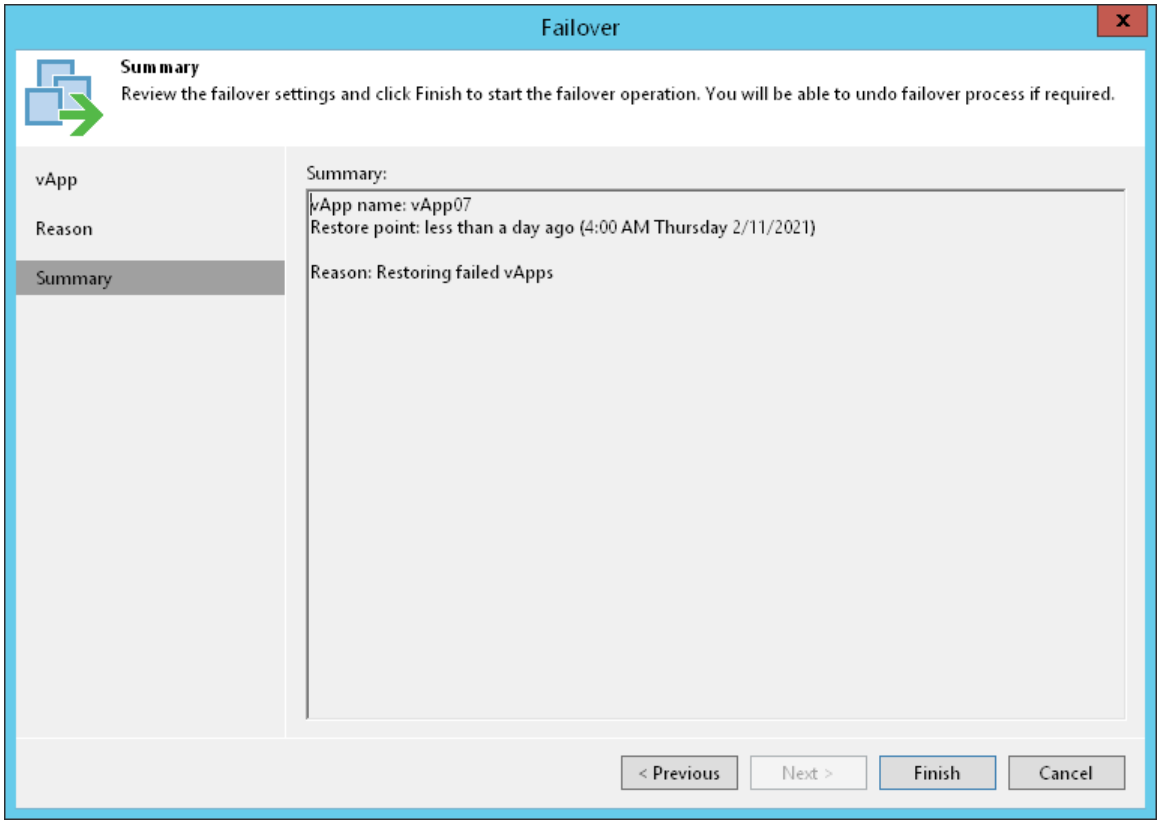
> **TIP**
>
> If you do not want to display the **Reason** step of the wizard in future, select the **Do not show me this page again** check box.



Step 5. Review Summary and Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task. Then click **Finish** to start the failover process.

When the failover process is complete, the vApp replicas will be started on the target hosts.



Failover Retry

The failover retry option is necessary when failover of vApps fails with the *Incomplete* state. When you perform a retry, Veeam Backup & Replication restarts failover only for the failed VMs that are added to vApps and will not process VMs that have been processed successfully. As a result, failover will take less time and will not consume as many resources as when processing a whole vApp.

To retry 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 vApp and select **Retry Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry failover**.

Permanent Failover

Permanent failover is one of the ways to finalize failover. When you perform permanent failover, you permanently switch from the original vApp to it replica. As a result, the vApp replica stops acting as a replica and starts acting as the production vApp.

> **NOTE**
>
> It is recommended to perform permanent failover if the original vApp 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 vApp replica from the snapshot chain and deletes associated files from the datastore.
2. Veeam Backup & Replication removes the vApp replica from the list of replicas in the Veeam Backup & Replication console and the replica vApp becomes a production vApp.
3. To protect from corruption after permanent failover is complete the vApp to which you have failed over, Veeam Backup & Replication reconfigures the replication job and adds the original vApp to the list of exclusions. When the replication job starts, the original vApp is skipped from processing. As a result, no data is written to the vApp to which you failed over.

Performing Permanent Failover

To perform permanent failover, do one of the following:

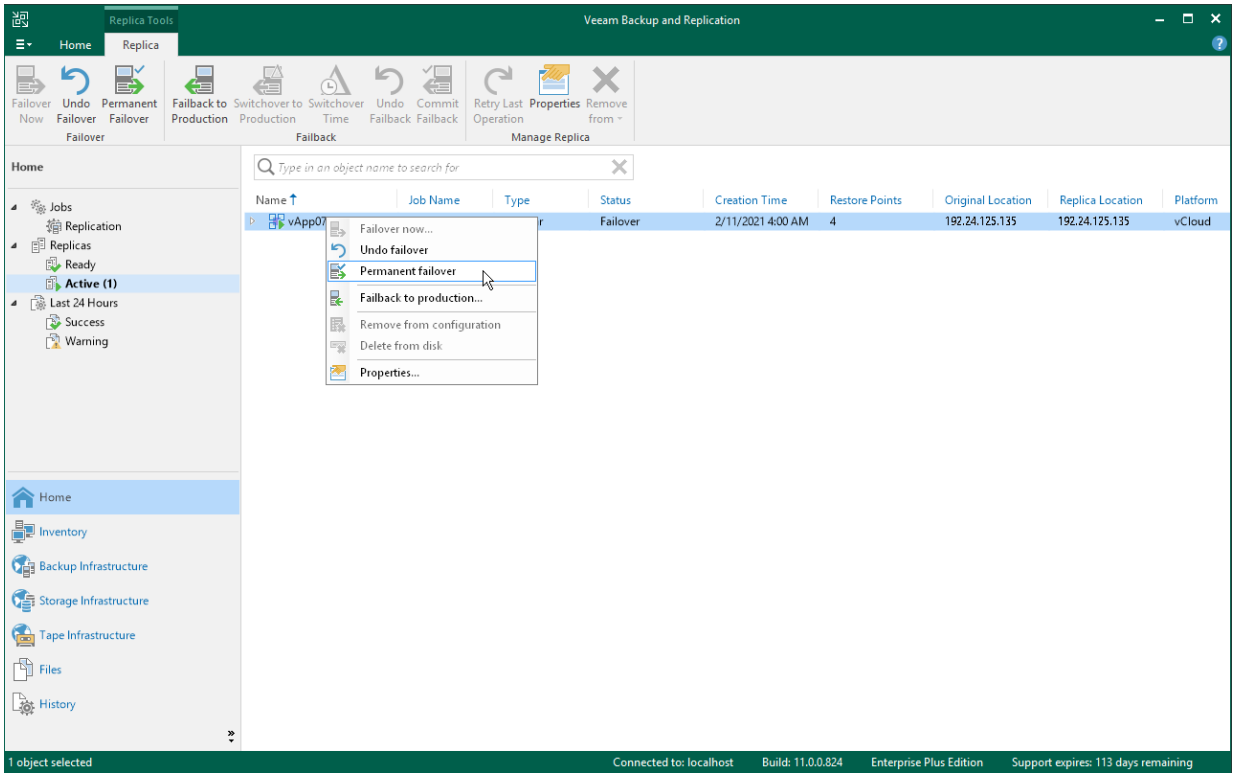
- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, select the necessary vApp and click **Permanent Failover** on the ribbon.
- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, right-click the necessary vApp and select **Permanent Failover**.

If permanent failover failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts the permanent failover only for the failed VMs that are added to vApps and will not process VMs that have been processed successfully. As a result, permanent failover will take less time and will not consume as many resources as when processing a whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.

2. In the working area, select the necessary vApp and select **Retry Permanent Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry permanent failover**.



Failover Undo

Failover undo is one of the ways to finalize failover. When you undo failover, you switch back from a vApp replica to the original vApp. Veeam Backup & Replication discards all changes made to the vApp replica while it was in the Failover state.

The undo failover operation is performed in the following way:

1. Veeam Backup & Replication reverts the vApp replica to its pre-failover state. To do this, Veeam Backup & Replication powers off the vApp replica and gets it back to the latest restore point in the replication chain.
2. The state of the vApp replica gets back to *Ready*, and Veeam Backup & Replication resumes replication activities for the original vApp on the source host.

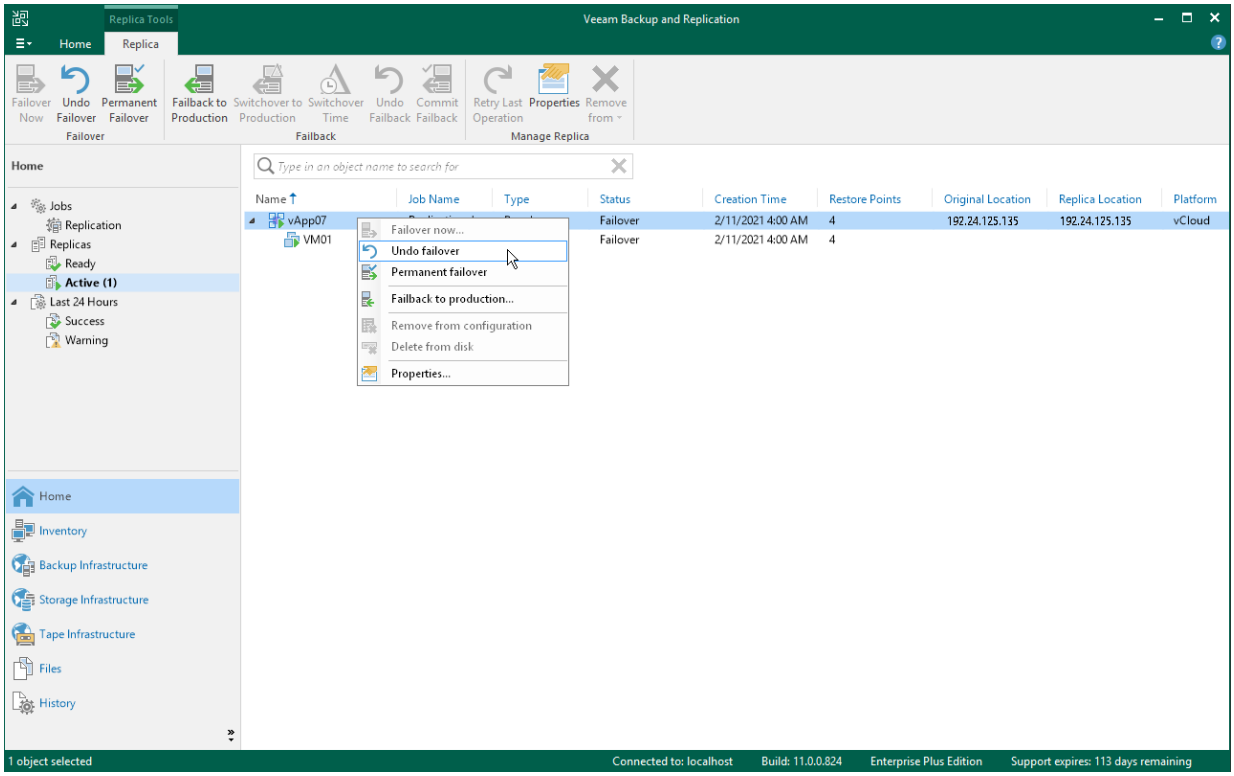
Undoing Failover

With the undo failover operation, you can power off running vApp replicas on target hosts and roll back to initial state of vApp replicas.

To undo failover:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, select the necessary replica and click **Undo Failover** on the ribbon. You can also right-click the necessary replica and select **Undo Failover**.
4. In the displayed window, click **Yes** to confirm the operation.

If the undo failover operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Undo Failover** on the ribbon. Alternatively, right-click the necessary vApp and select **Retry undo failover**.



Failback

Failback is an option that allows you to switch back from the replicated vApp on a disaster recovery organization VDC to the production vApp. When you perform failback, you shift from the disaster recovery site to the production site.

You can perform failback with one of the following options:

- **Failback to the original vApp**: Use this option to fail back to the original vApp in the original location.
- **Failback to the original vApp restored in a different location**: Use this option to fail back to the original vApp recovered to a new location. This vApp must be recovered before you perform failback. For example, you can recover the vApp from a backup.
- **Failback to the specified location**: Use this option to fail back to a vApp recovered from a replica. Veeam Backup & Replication will recover a vApp to a new location with different settings (such as a vApp location, network settings, storage policy and so on). The vApp 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 vApp and vApp replica. For the third option, Veeam Backup & Replication needs to transfer the whole vApp data, including its configuration and virtual disk content. Use the third option if there is no way to use the original vApp or restore it from a replica.

Veeam Backup & Replication performs failback in two phases:

- **First phase**: Veeam Backup & Replication transfers the necessary snapshot of a replicated vApp from the disaster recovery organization VDC to the production vApp. This phase may take a lot of time especially if the vApp contains a lot of VMs. While Veeam Backup & Replication performs the first phase of failback, VMs from the replica vApp are still up and running, users can access these VMs and perform daily routine tasks as normal. All changes that are made to vApps during the first phase of a failback a written to a delta file.
- **Second phase**: Veeam Backup & Replication transfers all changes made to the vApp replica during the first phase of failback to the production vApp, switches all processes from the vApp replica to the production vApp and turns off the replica.

The time when the second phase starts depends on how you want to switch from the replica to the production vApp. You can switch to the production vApp 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 vApp or an original vApp restored in a different location differs from the process of failing back to a specific location:

- How failback to the original vApp or an original vApp restored in a different location works
- How failback to a specific location works

How Failback to Original vApp or Original vApp Restored in Different Location Works

When you fail back to the original vApp or an already recovered vApp, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication calculates the difference between disks of the production vApp and disks of the vApp replica in the *Failover* state. The calculation of the difference helps Veeam Backup & Replication understand what data needs to be transferred to the production vApp and to synchronize its state with the state of the vApp replica.

   If you fail back to the original vApp in the original location and you have enabled the **Quick rollback** option, this calculation can be performed much faster. For more information on the **Quick rollback** option, see Quick Rollback.

2. Veeam Backup & Replication transfers the data that was detected as different to the production vApp. The transferred data is written to the production vApp.
3. Veeam Backup & Replication changes the state of the vApp replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the following operations:

1. The guest OS of the vApp replica is shut down or the vApp replica is powered off.

   If VMware Tools are installed on the VM added to the vApp replica, Veeam Backup & Replication tries to shut down the replica guest OS. If nothing happens in 15 minutes, Veeam Backup & Replication powers off the vApp replica. If VMware Tools are not installed on the VM added to the vApp replica or the vApp is suspended, Veeam Backup & Replication powers off the vApp. The vApp replica remains powered off until you commit failback or undo failback.

2. Veeam Backup & Replication calculates the difference between disks of the production vApp and disks of the vApp replica. The calculation of the difference helps Veeam Backup & Replication understand what data was changed while the vApp replica was in the *Ready to switch* state.
3. Sends data changed on the vApp replica while it was in the *Ready to switch* state to the production vApp.

4. The state of the vApp replica is changed from *Ready to switch* to *Failback*.
5. [If you fail back to a recovered vApp] Veeam Backup & Replication updates the ID of the original vApp in the Veeam Backup & Replication configuration database. The ID of the original vApp is replaced with the ID of the recovered vApp.
6. If you have selected to power on the production vApp after failback, Veeam Backup & Replication powers on the production vApp on the host.

How Failback to Specific Location Works

When you fail back to a vApp recovered from a replica, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication requests VMware vCloud Director to create on the target organization VDC an empty vApp with the same configuration as the vApp replica. VMware vCloud Director server registers the created production vApp.
2. Veeam Backup & Replication transfers data of the vApp replica to the production vApp to update the production vApp state to the vApp replica state.
3. Veeam Backup & Replication changes the state of the vApp 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 vApp or Already Recovered vApp Works except for the step 2.

Failback is an intermediate step that needs to be finalized. If the production vApp works as expected and you want to get back to it, commit failback. If the vApp does not work as expected, undo failback.

Performing Failback

With the **Failback** option, you can switch from a vApp replica back to the original vApp or restore a vApp from a vApp replica in a new location.

To perform failback, do the following:

1. Check prerequisites
2. Launch the Failback wizard
3. Select vApp replicas to fail back
4. Select the failback destination
5. Select the target vApp
6. Specify the organization VDC
7. Specify the storage policy
8. Specify the network
9. Define Failback Mode
10. Review summary and finish working with the wizard
11. [If failback fails] Perform the failback retry option

Before You Begin

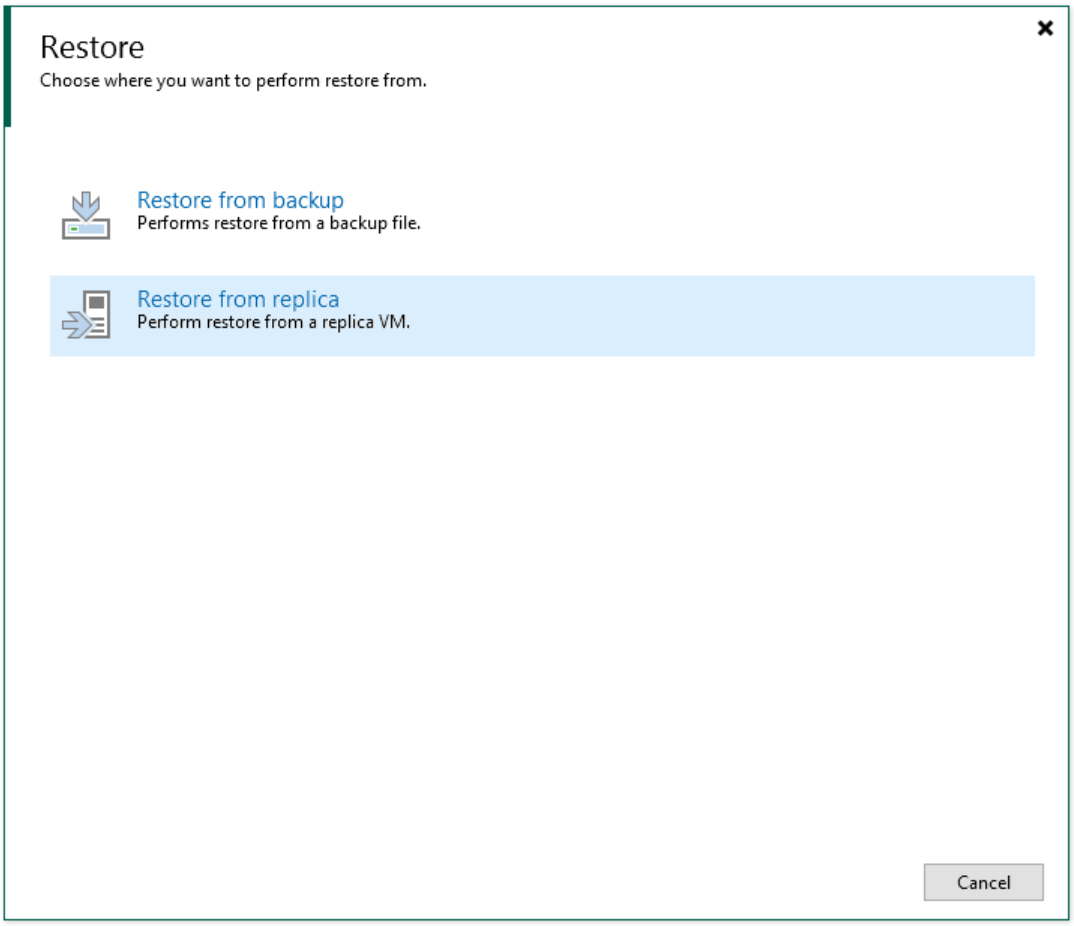
Before you perform failback, check the following prerequisites:

- vApps for which you plan to perform failback must be successfully replicated at least once.
- vApp replicas must be in the *Failover* state.

Step 1. Launch the Failback Wizard

To launch the **Failback** wizard, do one of the following:

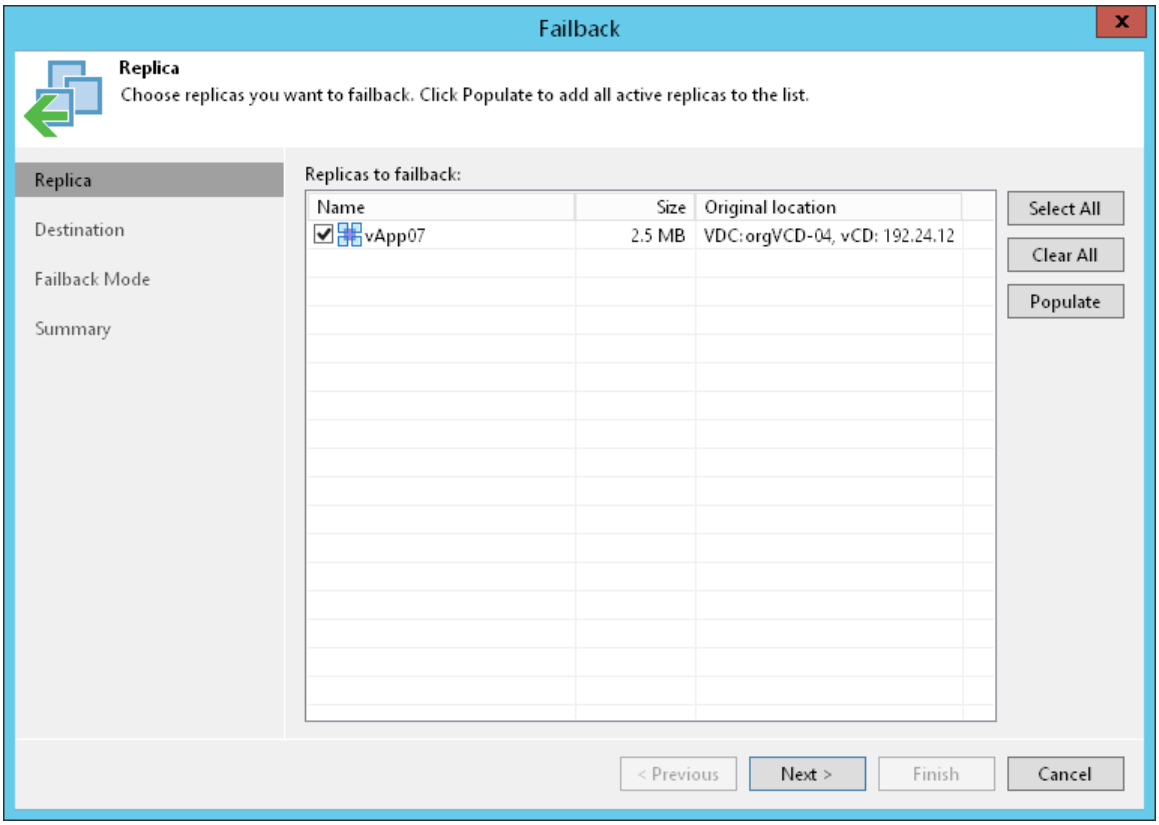
- On the **Home** tab, click **Restore** > **VMware vCloud Director** > **Restore from replica** > **Entire vApp** > **Failback to production**.
- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, select the necessary replica and click **Failback to Production** on the ribbon.
- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, right-click the necessary replica and select **Failback to production**.



Step 2. Select vApps to Fail Back

At the **Replica** step of the wizard, select one or more VM replicas from which you want to fail back.

1. Click **Populate** to display all existing replicas in the *Failover* state.
2. Leave check boxes selected for those VM replicas from which you want to fail back.



Step 3. Select Failback Destination

At the **Destination** step of the wizard, select the failback destination and backup proxies for vApp data transport during failback.

1. Select a destination for failback. Veeam Backup & Replication supports the following options
   - **Failback to the original vApp** — select this option if you want to fail back to the original vApp that resides on the source host. Veeam Backup & Replication will synchronize the state of the original vApp with the current state of its replica and will apply all changes occurred to the vApp replica while it was running in the disaster recovery (DR) site.

     If this option is selected, you will proceed to the Failback Mode step of the wizard.
   - **Failback to the original vApp restored in a different location** — select this option if you have recovered the original vApp from a backup in a new location, and you want to switch to it from the replica. Veeam Backup & Replication will synchronize the state of the recovered vApp with the current state of the vApp replica and will apply all changes occurred to the vApp replica while it was running in the disaster recovery (DR) site.

     If this option is selected, you will proceed to the Target vApp step of the wizard.

     > **TIP**
     >
     > You can restore to another vCD.
   - **Failback to the specified location** — select this option if you want to restore the original vApp from a replica. In this case, Veeam Backup & Replication will recover a vApp to a new location with different settings (such as a vApp location, network settings, storage policy and so on). Select this option if there is no way to fail back to the original vApp or an already recovered vApp.

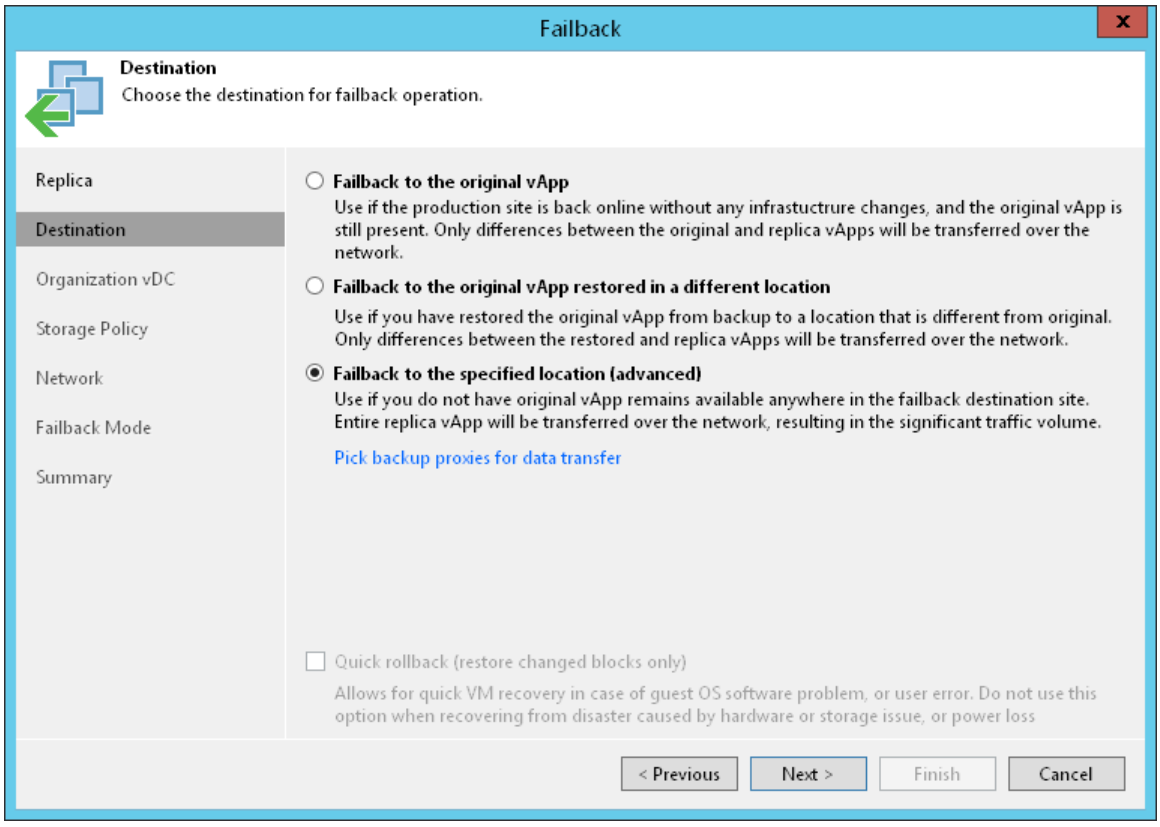
     If this option is selected, you will proceed to the Organization vDC step of the wizard.

     > **NOTE**
     >
     > Mind the following:
     >
     > - If you fail back to the original vApp or the original vApp is restored in a new location, the difference between the virtual disks of the replicated vApp and their state is transferred to the original vApp. Veeam Backup & Replication does not transfer replica configuration changes such as a different IP address or network settings (if replica Re-IP and network mapping were applied), new hardware or virtual disks added while the replica was in the *Failover* state.
     > - If you choose to fail back to the specified location, Veeam Backup & Replication will restore the entire vApp replica, including its configuration and virtual disks content.
2. Click the **Pick backup proxies for data transfer** link to select backup proxies. Veeam Backup & Replication will use them to transfer data during failback. If vApps and their replica are located in different sites, select at least one backup proxy in the production site and one proxy in the disaster recovery site. If vApps and their replicas are located in the same site, you can use the same backup proxy as the source and target one.

3. In the **Choose backup Proxy** section, click **Choose** to assign a backup proxy. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to select backup proxies automatically.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing vApp data.

     vApps selected for failback are processed one by one. Before processing a new vApp in the vApp 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, the current workload on the backup proxies to select the most appropriate resource for vApp processing.
   - If you choose **Use the selected backup proxy servers only**, you can explicitly define backup proxies that must be used for data transfer. It is recommended that you select at least two backup proxies to ensure that failover is performed should one of backup proxies fail or lose its connectivity to the source or target datastore.
4. If you want to perform failback with the quick rollback option, select the **Quick rollback (restore changed blocks only)** check box.

> **IMPORTANT**
>
> It is recommended to use the quick rollback option if you fail back to the original vApp after a problem occurred at the guest OS level of a VM added to a vApp. Do not use quick rollback if the problem has occurred at the VM hardware level, storage level or due to a power loss.



Restoring Storage Policies

If the replicated vApp was associated with the storage policy, in the failback to original location scenario, Veeam Backup & Replication will associate the restored vApp 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.

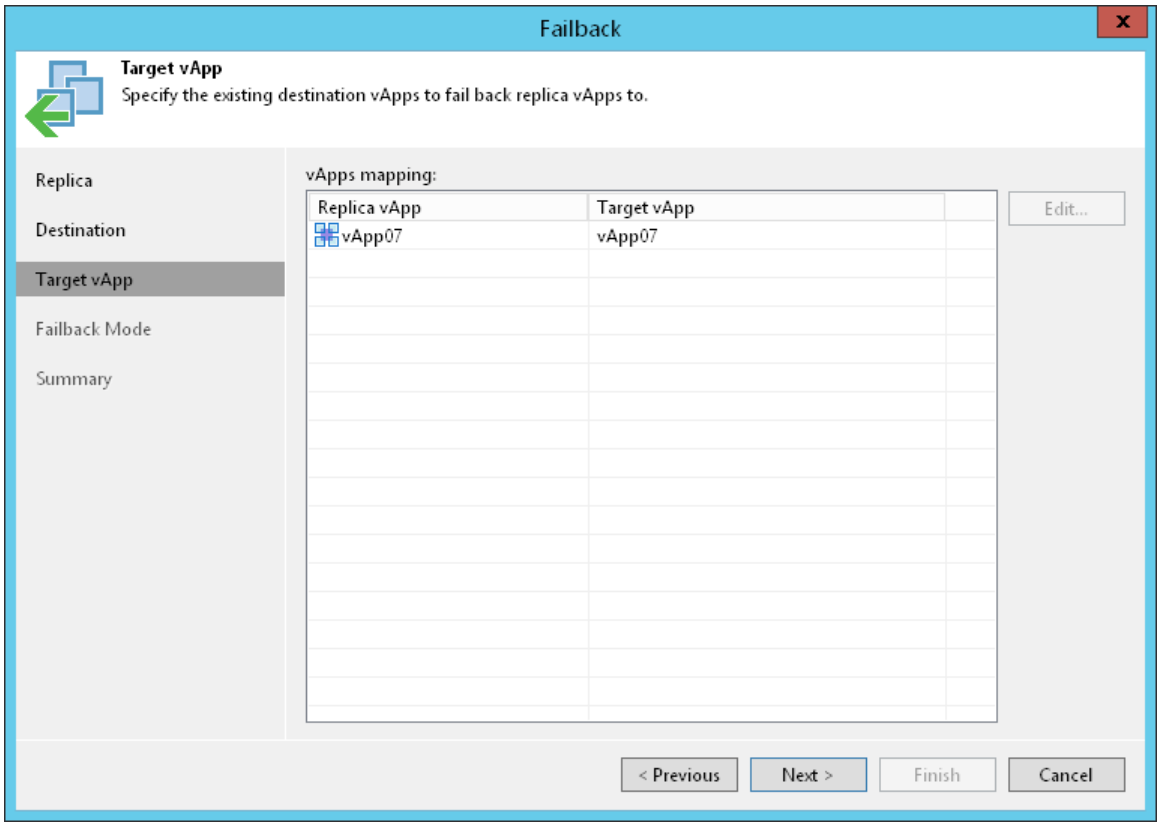
Step 4. Select Target vApp

The **Target vApp** step of the wizard is only available if you have chosen to fail back to the original vApp restored in a different location.

At the **Target vApp** step of the wizard specify the target vApp to which Veeam Backup & Replication will map the replica vApp.

By default, Veeam Backup & Replication will map the replica vApp to the original vApp. If you want to specify the target vApp manually, perform the following steps:

1. Select the vApp that you want to map to the target vApp and click **Edit**.
2. In the **Selects Objects** window, select the necessary vApps or vApps containers.
3. Click **Add**.

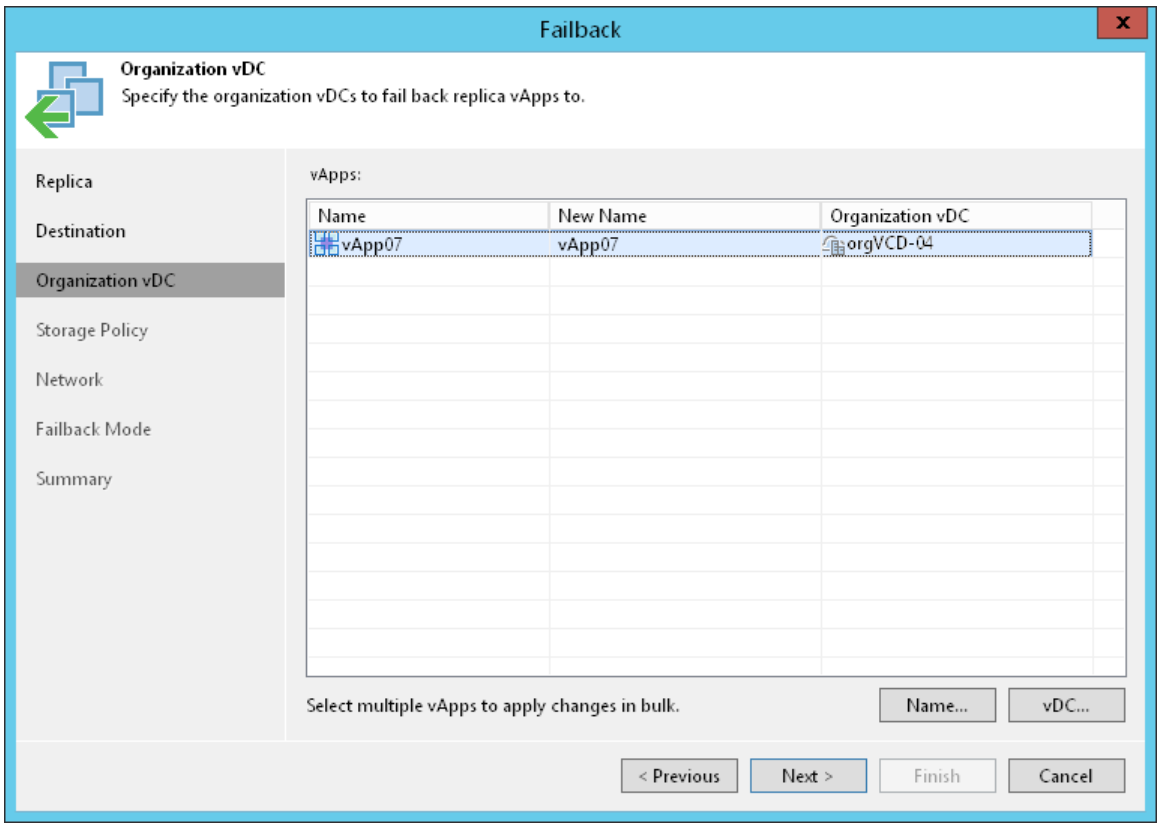


Step 5. Specify Organization VDC

The **Organization vDC** step of the wizard is only available if you have chosen to fail back to the specified location.

At the **Organization vDC** step of the wizard, specify the organization VDC to which Veeam Backup & Replication will add a restored vApp and will assign a name to this vApp. To do this, select one or multiple vApps, and perform the following steps:

1. To specify a name for the restored vApps, click **Name** and in the **Name** window specify the name for vApps.
2. To specify an organization VDC, click **vDC** and in the **Selected Objects** window, select the necessary organization VDC.



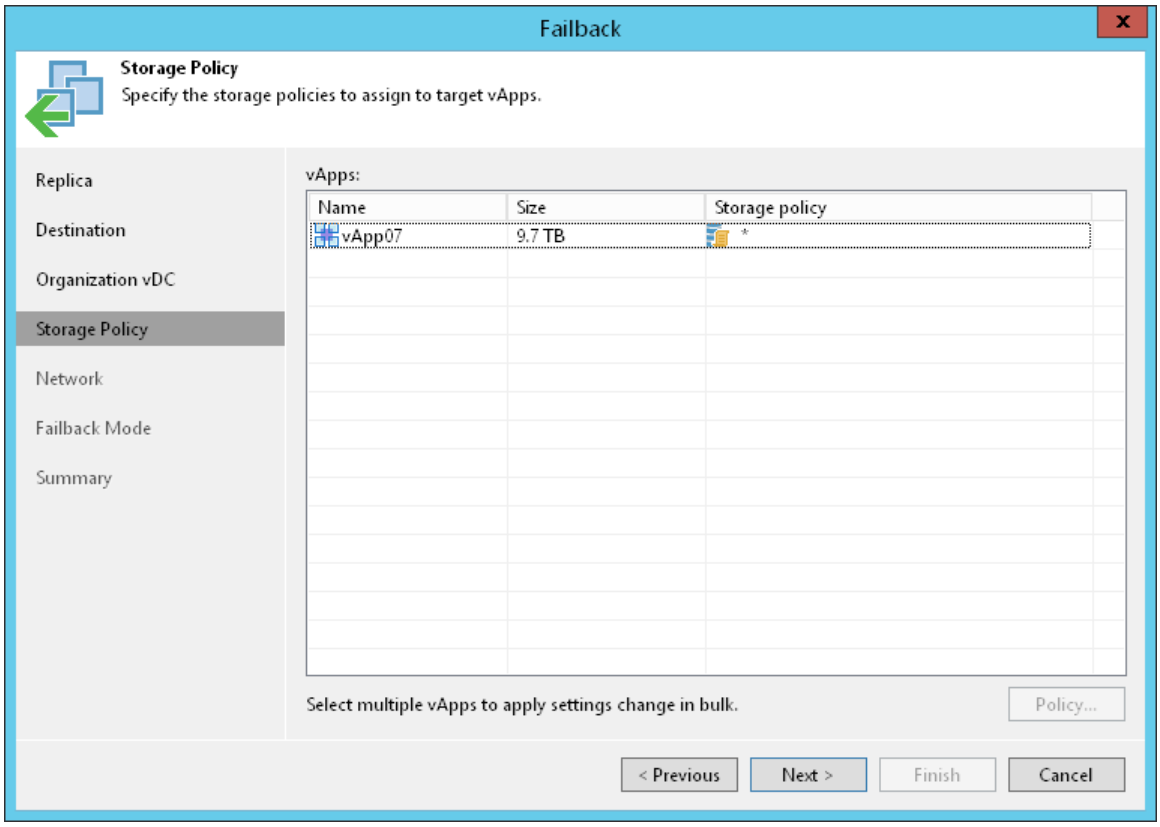
Step 6. Specify Storage Policy

The **Storage Policy** step of the wizard is only available if you have chosen to fail back to the specified location.

At the **Storage Policy** step of the wizard, specify the storage policy setting. Veeam Backup & Replication will apply these settings to a vApp that you want to restore.

To change storage policy setting:

1. In the **vApps location** list, select vApps for which you want to change the policy settings and click **Policy**.
2. In the **Select storage policy** window specify the policy that you want to apply.



Step 7. Specify Network

The **Network** step of the wizard is only available if you have chosen to fail back to the specified location.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the DR site to networks in the site where recovered vApps reside. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs added to vApps on the fly, during the failback process.

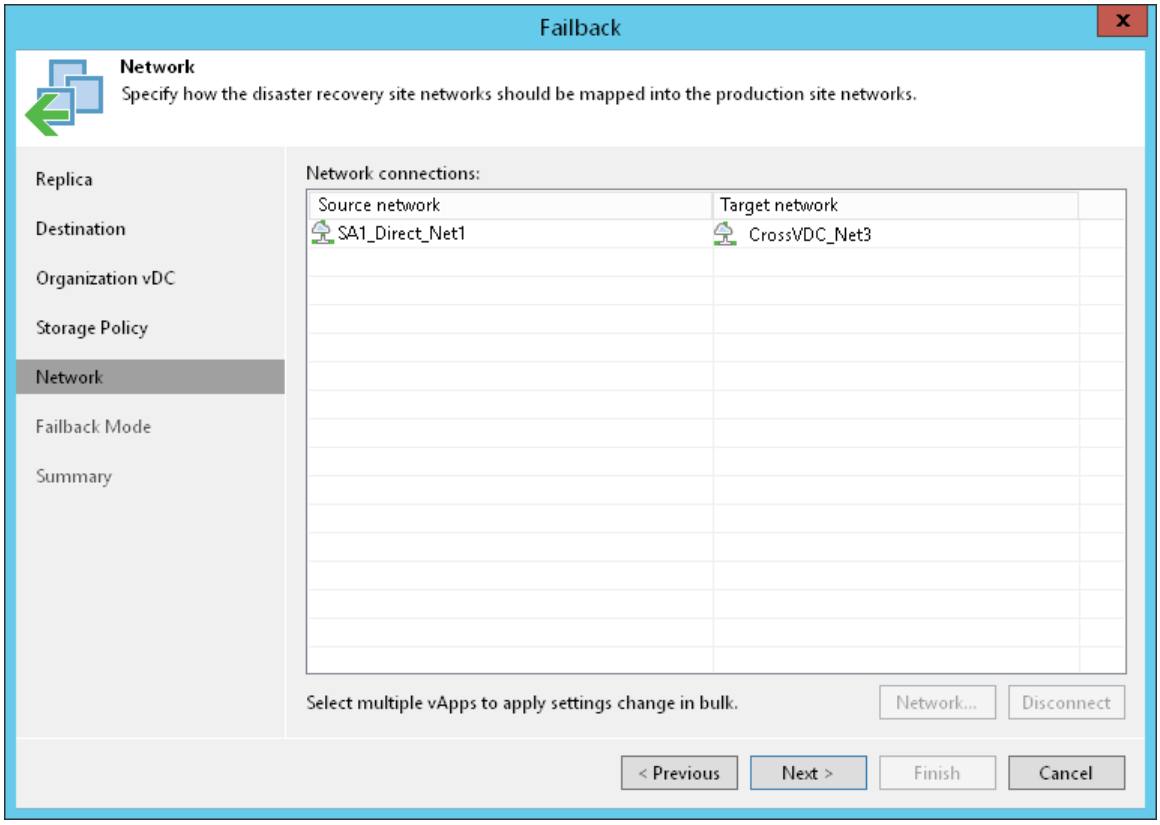
To change networks to which the restored vApps will be connected:

1. In the **Network connections** list, select one or more vApps and click **Network**.

   If vApps are connected to multiple networks, select a necessary network and click **Network**.

2. In the **Select network** window, select networks to which vApps must be connected after restore.

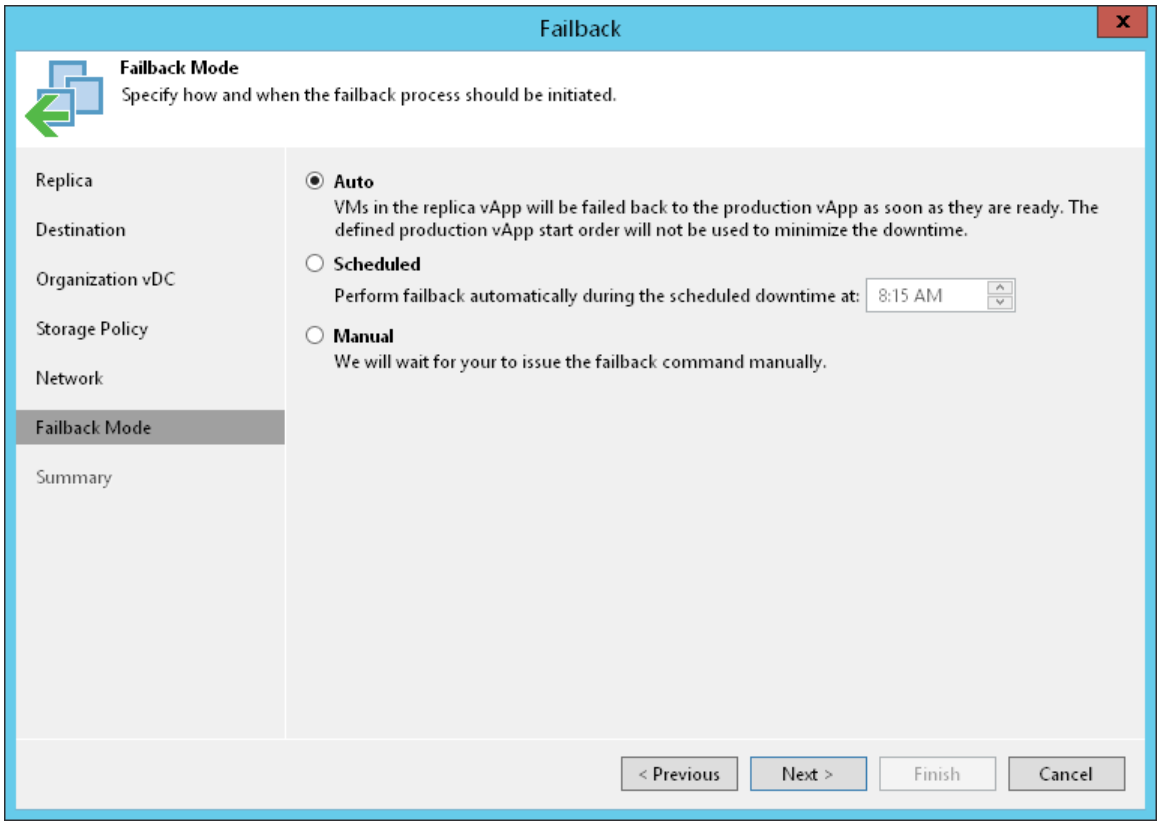
If you do not want to connect restored vApps to any virtual network, select the necessary vApps and click **Disconnect**.



Step 8. Schedule Switch to Production vApps

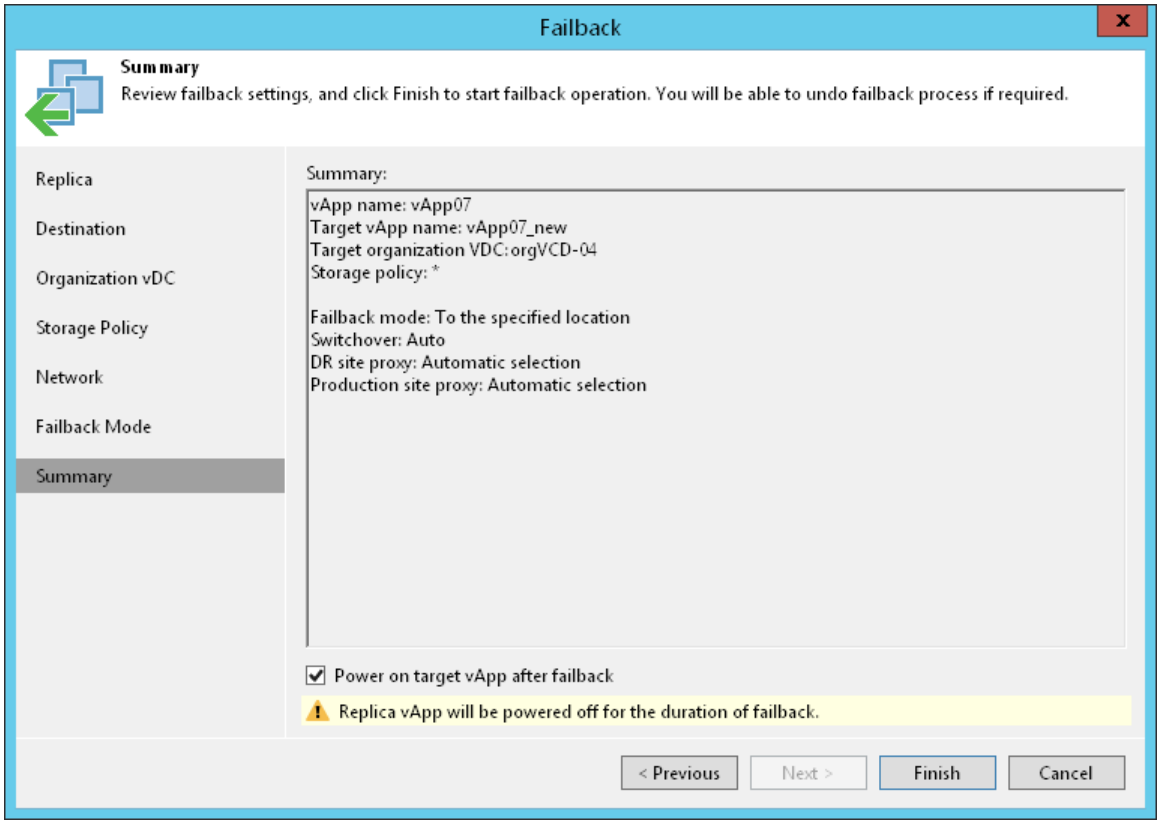
At the **Failback Mode** step of the wizard, specify when switch from replicas to production vApps must be performed:

- Select **Auto** if you want Veeam Backup & Replication to perform the switch automatically right after the state of the production vApps 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 **Manually** if you want to perform the switch manually. For more information on how to perform the switch, see Switching to Production vApps Manually.



Step 9. Review Summary and Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failback. If you want to start the restored vApps right after failback completes, select the **Power on target vApp after failback** check box. Then click **Finish** to close the wizard.



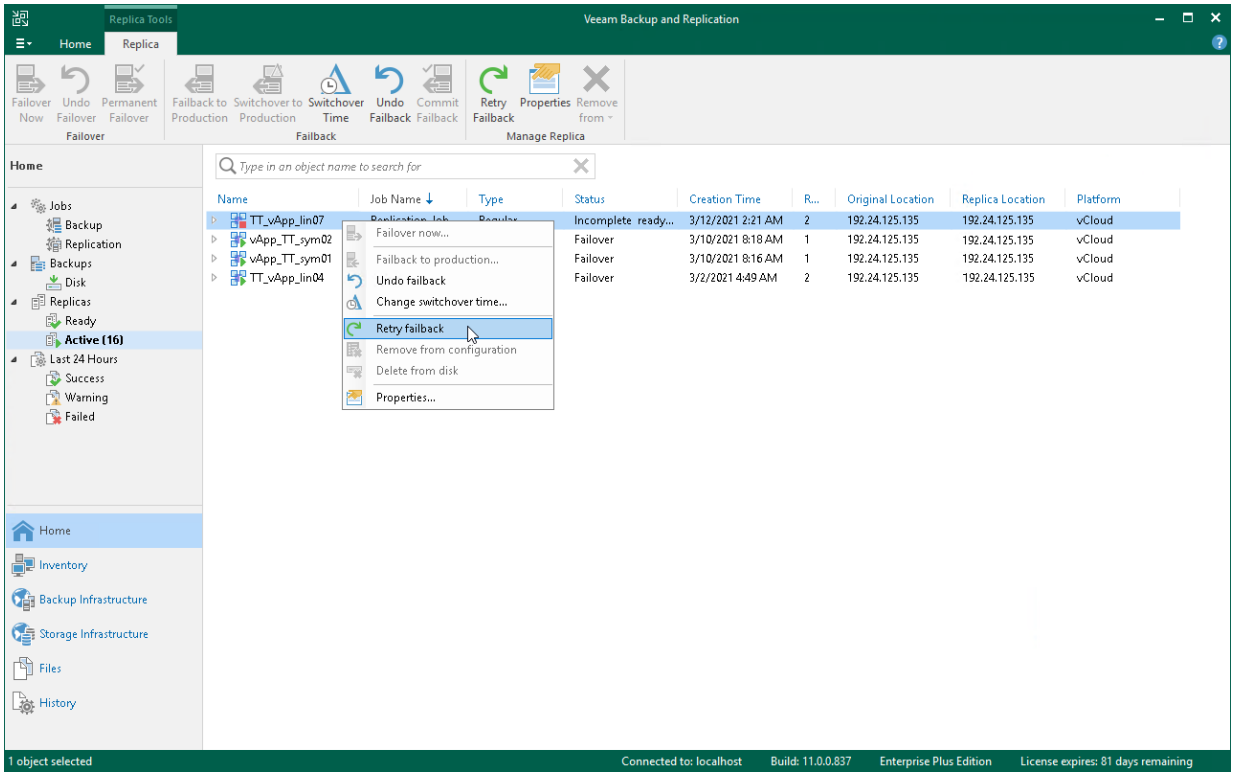
Failback Retry

The failback retry option is necessary when failback of vApps fails with the *Incomplete* state. When you perform a retry, Veeam Backup & Replication restarts failback only for the failed VMs that are added to vApps and will not process VMs that have been processed successfully. As a result, failback will take less time and will not consume as many resources as when processing a whole vApp.

To retry 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 vApp and select **Retry Failback** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry failback**.



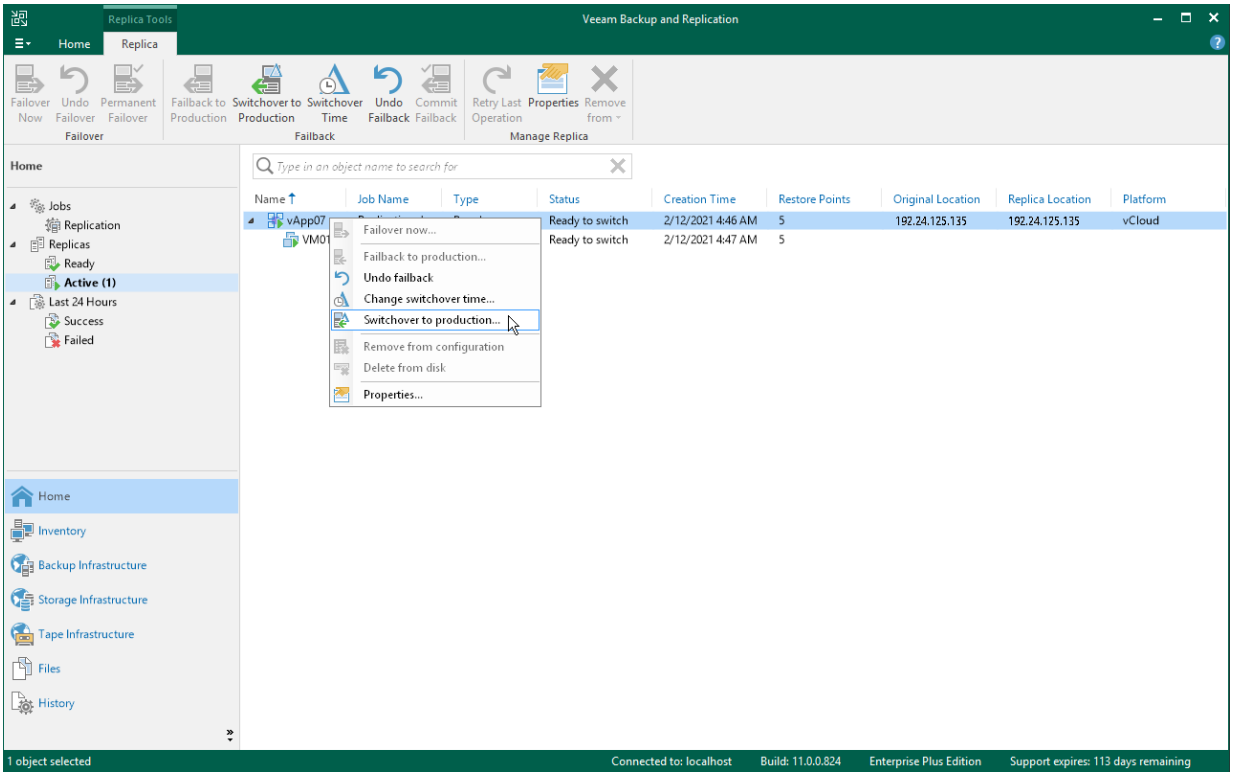
Switching to Production vApps Manually

The following instructions apply if you have selected to switch from replicas to production vApps manually at the Failback Mode step of the **Failback** wizard.

To switch to a production vApp from its replica, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary vApp and select **Switchover to production** on the ribbon. As an alternative, you can right-click the necessary vApp and select **Switchover to production**.

If the switch to production operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Switchover to Production** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry switchover to production**.



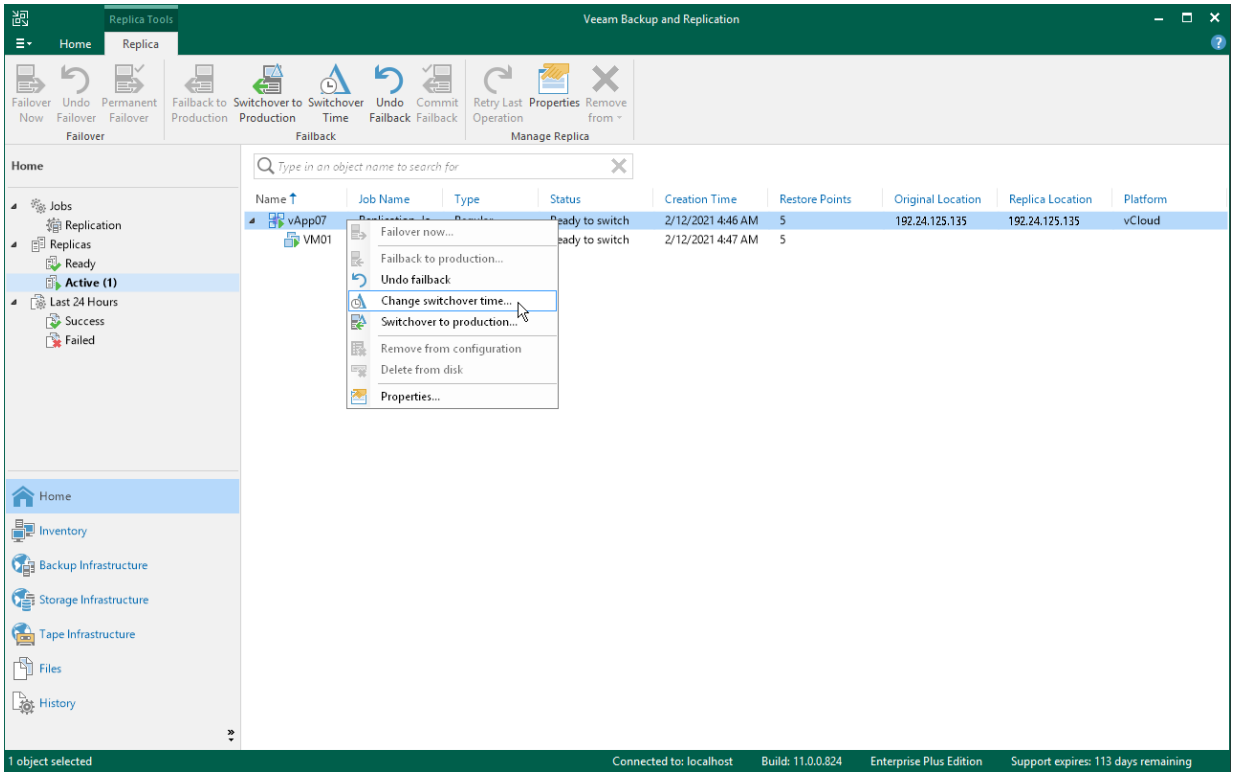
Changing Switching Time

The following instructions apply if you have selected to switch from replicas to production vApps manually or at the scheduled time at at the Failback Mode step of the **Failback** wizard.

To change the time when the switch from replicas to production vApps must be performed:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the vApp in the *Ready to switch* state and select **Switchover Time** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Change switchover time.**

If the switching time operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Switchover Time** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry switchover time.**



Failback Undo

Failback undo is one of the ways to finalize failback. You can use this option if the vApp to which you failed back (the production vApp) works in a wrong way and you want to get back to the replica.

The undo failback operation is performed in the following way:

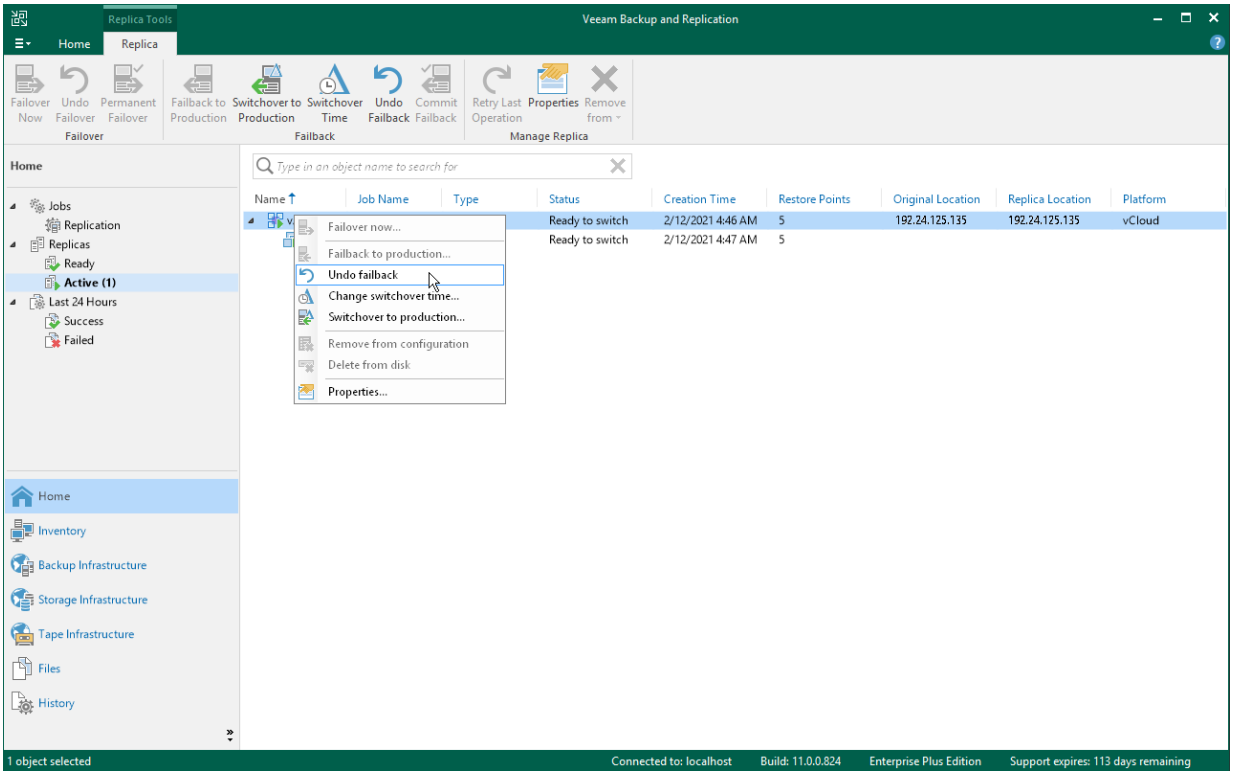
1. Veeam Backup & Replication powers off the production vApp.
2. Veeam Backup & Replication reverts the vApp replica to its pre-failback state.
3. Veeam Backup & Replication powers on the vApp replica and changes the vApp replica state from *Failback* or *Ready to Switch* to *Failover*.

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. Alternatively, you can right-click the necessary replica and select **Undo Failback**.

If the undo failback operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Undo Failback** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry undo failback.**



Failback Commit

To confirm failback and finalize recovery of the original vApp, you need to commit failback.

When you commit failback, Veeam Backup & Replication gets back to the normal operation mode and resumes replication activities for the original vApp to which you failed back.

The commit failback operation is performed in the following way:

1. Depending on whether you have failed back to the original vApp or recovered vApp:
    - If you have failed back to a vApp recovered from a backup or replica, Veeam Backup & Replication will reconfigure all existing jobs where the original vApp is present and adds the original vApp to the list of exclusions. The recovered vApp takes the role of the original vApp and is included into all jobs instead of the excluded vApp. When the vCD replication job starts, Veeam Backup & Replication processes the recovered vApp instead of the former original vApp.
    - If you have failed back to the original vApp, the vCD replication job is not reconfigured. When the vCD replication job starts, Veeam Backup & Replication still processes the original vApp.
2. Veeam Backup & Replication changes the state of the replica from *Failback* to *Ready*.

During failback commit, the failback protective snapshot that saves the pre-failback state of a vApp replica is not deleted. Veeam Backup & Replication uses this snapshot as an additional restore point for vApp 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

The commit failback operation finalizes failback from the vApp replica to the original vApp.

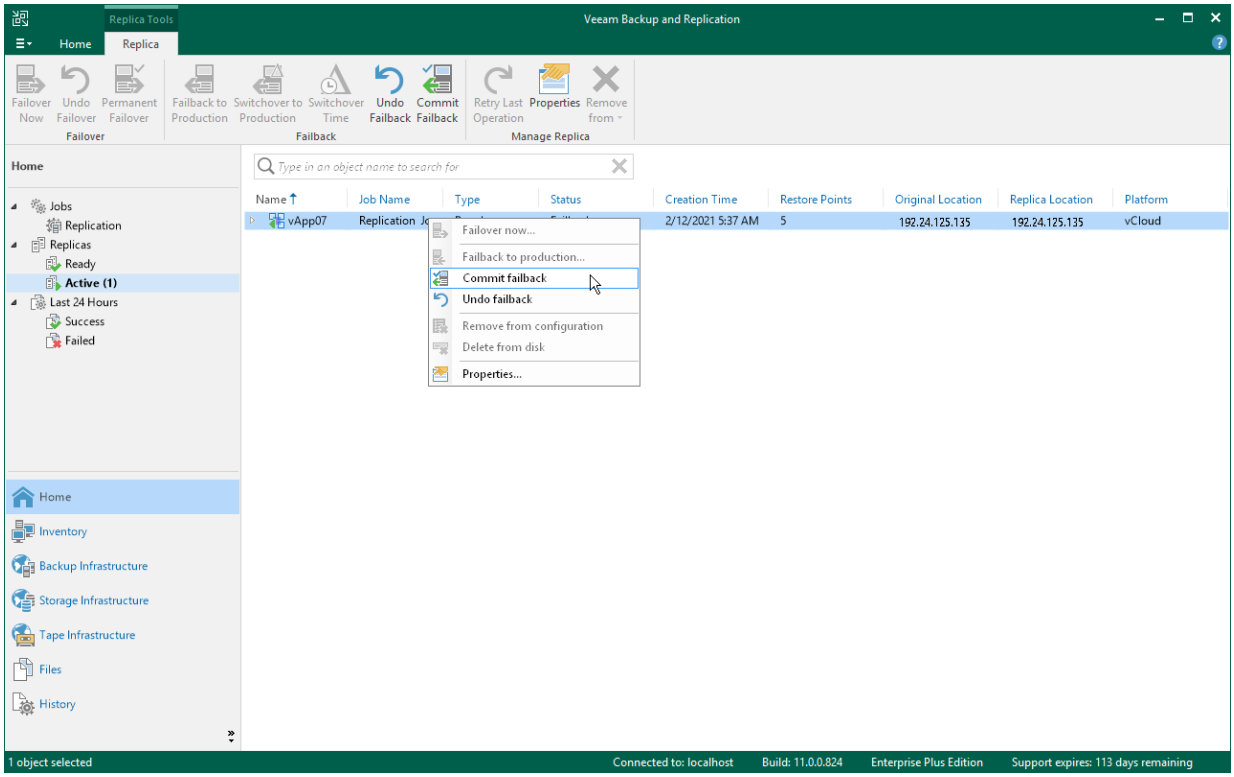
To commit failback, do one of the following:

- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, select the necessary replica and click **Commit Failback** on the ribbon. In the displayed window, click **Yes** to confirm the operation.
- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, right-click the necessary replica and select **Commit Failback**. In the displayed window, click **Yes** to confirm the operation.

If commit failback failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts committing failback only for the failed VMs that are added to vApps and will not process VMs that have been processed successfully. As a result, commit failback will take less time and will not consume as many resources as when processing a whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Commit Failback** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry commit failback**.



Restore of vCloud Director VMs

Veeam Backup & Replication enables full-fledged restore of VMs to vCloud Director. You can restore separate VMs to vApps, as well as VM data.

For restore, Veeam Backup & Replication uses VM metadata saved to a backup file and restores specific VM attributes. As a result, you get a fully-functioning VM in vCloud Director, do not need to import the restored VM to vCloud Director and adjust the settings manually.

Backed-up objects can be restored to the same vCloud Director hierarchy or to a different vCloud Director environment. Restore options include:

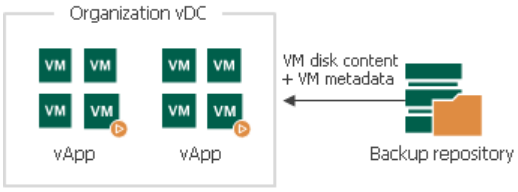
- Instant Recovery
- Full restore for vApps and VMs
- Restore of VM disks
- Restore of VM files
- Guest OS file-level restore for VMs

Restoring Regular and Standalone VMs to vCloud Director

Veeam Backup & Replication lets you restore regular VMs that are part of vApps and standalone VMs that were created in the vCloud Director tenant portal.

When you restore regular or standalone VMs back to the vCloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses the captured vApp metadata to define the vApp settings and VM original location in the vCloud Director hierarchy.
2. Veeam Backup & Replication restores VMs from the backup file to their original location or to a different location. Additionally, Veeam Backup & Replication restores all VM settings.



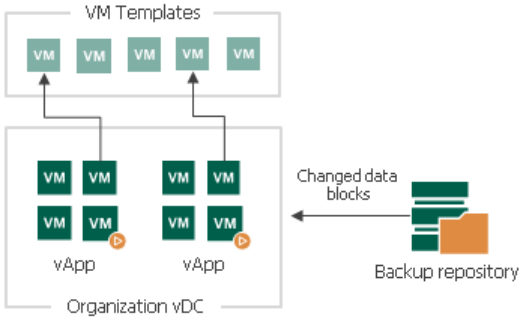
Restoring Linked Clone VMs to vCloud Director

Veeam Backup & Replication lets you restore linked clone VMs – VMs that were deployed from a VM template using the fast provisioning technology. There are several mechanisms for processing linked clone VMs.

Restore of Existing VMs

If you are restoring a vCloud Director linked clone VM that exists in the vCloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses the captured vApp metadata to define the initial settings of the VM.
2. Veeam Backup & Replication calculates a signature for the consolidated VM disk in the backup file (containing the VM template data and data of the delta file) and the signature for the VM existing in vCloud Director. Veeam Backup & Replication then compares the disk signatures to define what data blocks have changed.
3. Veeam Backup & Replication restores only changed data blocks from the backup file and writes them to the user delta file.

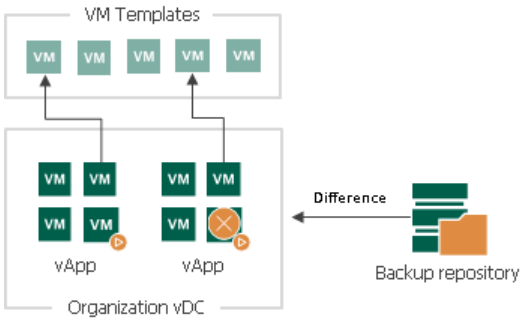


Restore of Deleted VMs

If you are restoring a VM that no longer exists in vCloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses vCloud Director to create a new linked clone VM from the VM template that the user selects. The new VM has a blank user delta file.
2. Veeam Backup & Replication calculates a signature for the consolidated VM disk in the backup file (containing the VM template data and data of the delta file) and the signature for the created VM in vCloud Director. Veeam Backup & Replication then compares the disk signatures to define what data blocks need to be restored.
3. Veeam Backup & Replication restores only those data blocks that need to be restored from the backup file and writes them to the blank user delta file.

By default, Veeam Backup & Replication links the VM to the same VM template that was used by the original VM. During restore, Veeam Backup & Replication checks the settings of the VM template to which the restored VM is linked: verifies connection settings, makes sure the disk size coincide and so on.

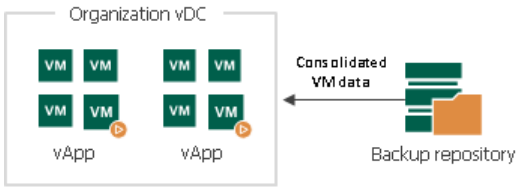


Restore of Linked Clone VMs as Regular VMs

In some cases, Veeam Backup & Replication can restore a VM from a backup file as a regular VM. This type of restore is accomplished in the following situations:

- You have intentionally chosen to restore a linked clone VM as a regular VM.
- You are restoring a VM to the organization VDC which has the fast provisioning option disabled.
- A VM template to which the restored VM should be linked is not accessible in the location to which the VM is restored.

In this case, Veeam Backup & Replication uses the same algorithm as for restore of full VMs in the virtual environment. It retrieves the data of the consolidated VM disk from the backup file and restores the VM in the vCloud Director hierarchy.



Performing Instant Recovery for VMs

Veeam Backup & Replication provides two options for Instant Recovery of vCloud Director VMs:

- You can instantly recover a VM to a vApp in vCloud Director.
- You can instantly recover a VM to the virtual infrastructure. In this case, the VM will be restored at the level of the underlying vCenter Server, and the Instant Recovery process will be the same as for regular VMware VMs.

When you instantly recover a VM to vCloud Director, Veeam Backup & Replication uses the vPower NFS datastore, just as with other VMware VMs. To import the VM to the vApp, Veeam Backup & Replication needs to associate the vPower NFS datastore with some storage policy. To do this, Veeam Backup & Replication creates for the underlying vCenter Server an auxiliary storage policy — *Veeam-InstantVMRecovery*, and displays it in vCloud Director.

The created storage policy is added to the Provider vDC and organization VDC hosting the vApp to which the VM is restored. When the vPower NFS datastore is mounted to the ESXi host, the vPower NFS datastore is associated with the *Veeam-InstantVMRecovery* storage policy. After that, the VM is instantly restored in a regular manner and imported to the selected vApp.

When an Instant Recovery session is finished, the storage policy is not deleted from the Provider vDC, it remains on vCenter Server. This helps speed up all subsequent Instant Recovery operations. However, the storage policy is deleted from the organization VDC as organization VDC settings can be accessed only by organization administrators.

Restoring VMs with Instant Recovery into vCloud vApp

With Instant Recovery, you can immediately start a VM from a backup file stored in the backup repository. Instant Recovery accelerates the restore process, allows you to improve RTOs and decrease downtime of production VMs.

Before you start Instant Recovery, check prerequisites. Then use the **vCloud Instant Recovery** wizard to recover the necessary VM.

Before You Begin

Before you perform Instant Recovery, check the following prerequisites:

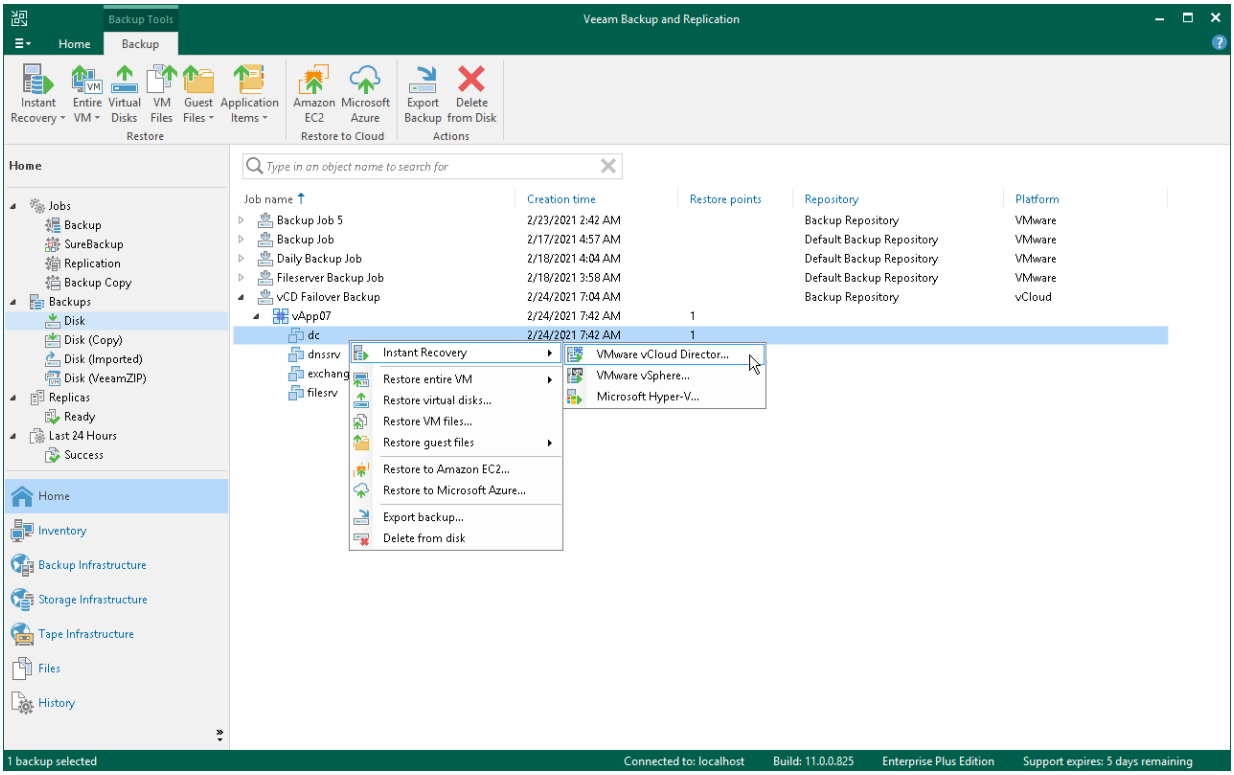
- You can perform Instant Recovery for a VM that has been successfully backed up at least once.
- You must have at least 10 GB of free disk space on the datastore where write cache folder is located. This disk space is required to store virtual disk updates for the restored VM.

  By default, Veeam Backup & Replication writes virtual disk updates to the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`. The write cache is not used when you select to redirect virtual disk updates to a VMware vSphere datastore when configuring the job.
- If you are recovering a VM to the production network, make sure that the original VM is powered off to avoid conflicts.
- If you want to scan VM data for viruses, check the secure restore requirements and limitations.
- Mind that during Instant Recovery, Veeam Backup & Replication restores standalone VMs as regular vCloud Director VMs.

Step 1. Launch vCloud Instant VM Recovery Wizard

To launch the **vCloud Instant VM Recovery** wizard, do one of the following:

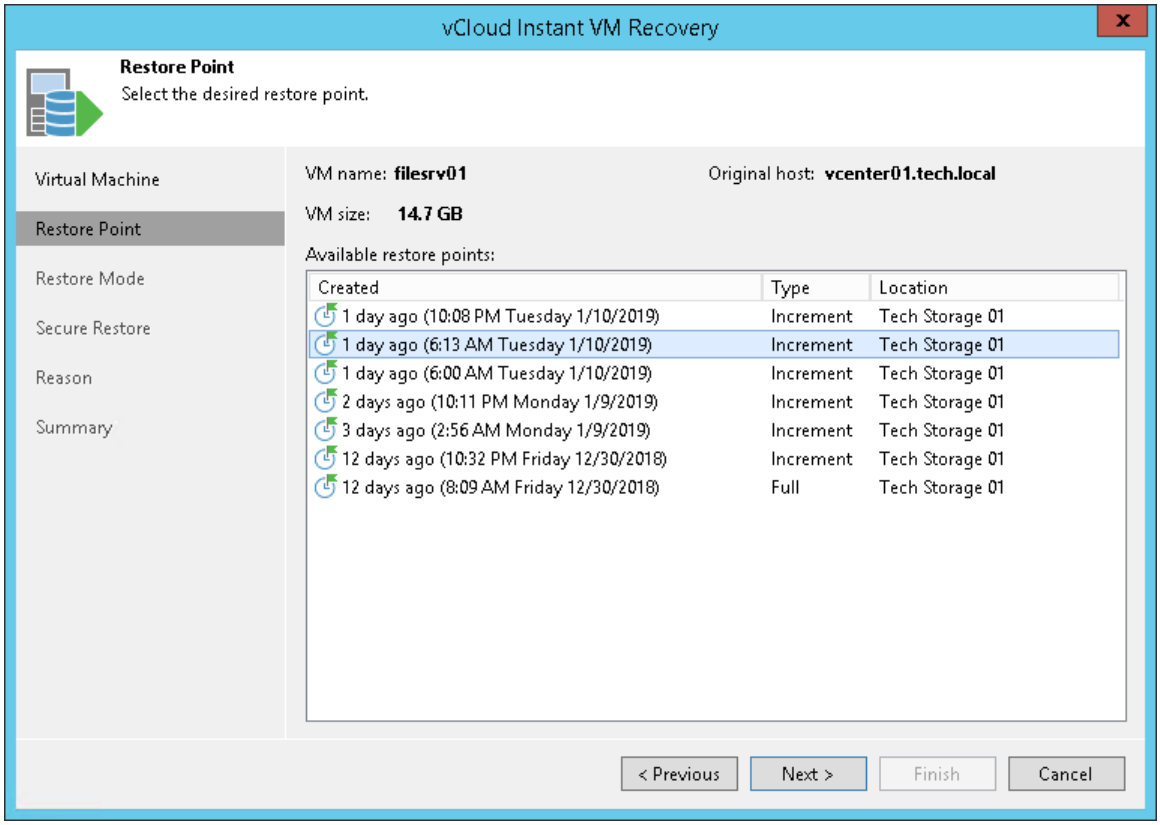
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Select the machine that you want to restore and click **Instant Recovery** > **VMware vCloud Director** on the ribbon.
  - Right-click the machine that you want to restore and select **Instant recovery** > **VMware vCloud Director**.
- Open the **Inventory** view. On the **View** tab, click **vCloud View**. In the inventory pane, expand the **vCloud Director** hierarchy. In the working area, right-click the VM you want to restore and select **Restore > Instant recovery** > **VMware vCloud Director**.



Step 2. Select Restore Point

At the **Restore Point** step of the wizard, select the restore point for the VM.

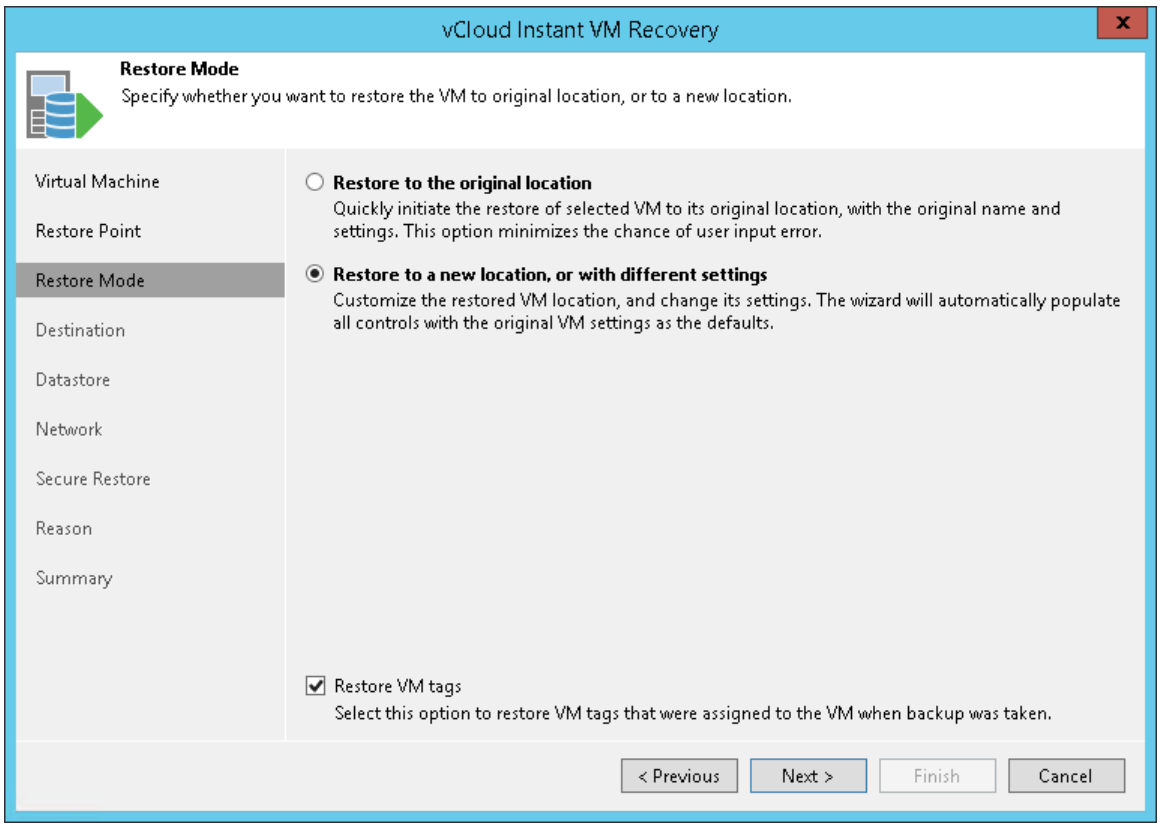
In the **Location** column, you can view a name of a backup repository where a restore point resides.



Step 3. Select Restore Mode

At the **Restore Mode** of the wizard, choose the necessary restore mode.

1. Choose a restore mode:
   - Select **Restore to the original location** if you want to restore the VM with its initial settings to its original location. If this option is selected, you will pass directly to the Reason step of the wizard.
   - Select **Restore to a new location, or with different settings** if you want to restore the VM to a different location and/or with different settings (such as vApp, VM name, network settings and so on). If this option is selected, the **Instant Recovery** wizard will include additional steps for customizing VM settings.
2. Select the **Restore VM tags** check box if you want to restore tags that were assigned to the original VM, and assign them to the restored VM. Veeam Backup & Replication will restore the VM with original tags if the following conditions are met:
   - The VM is restored to its original location.
   - The original VM tag is still available on the source vCenter Server.



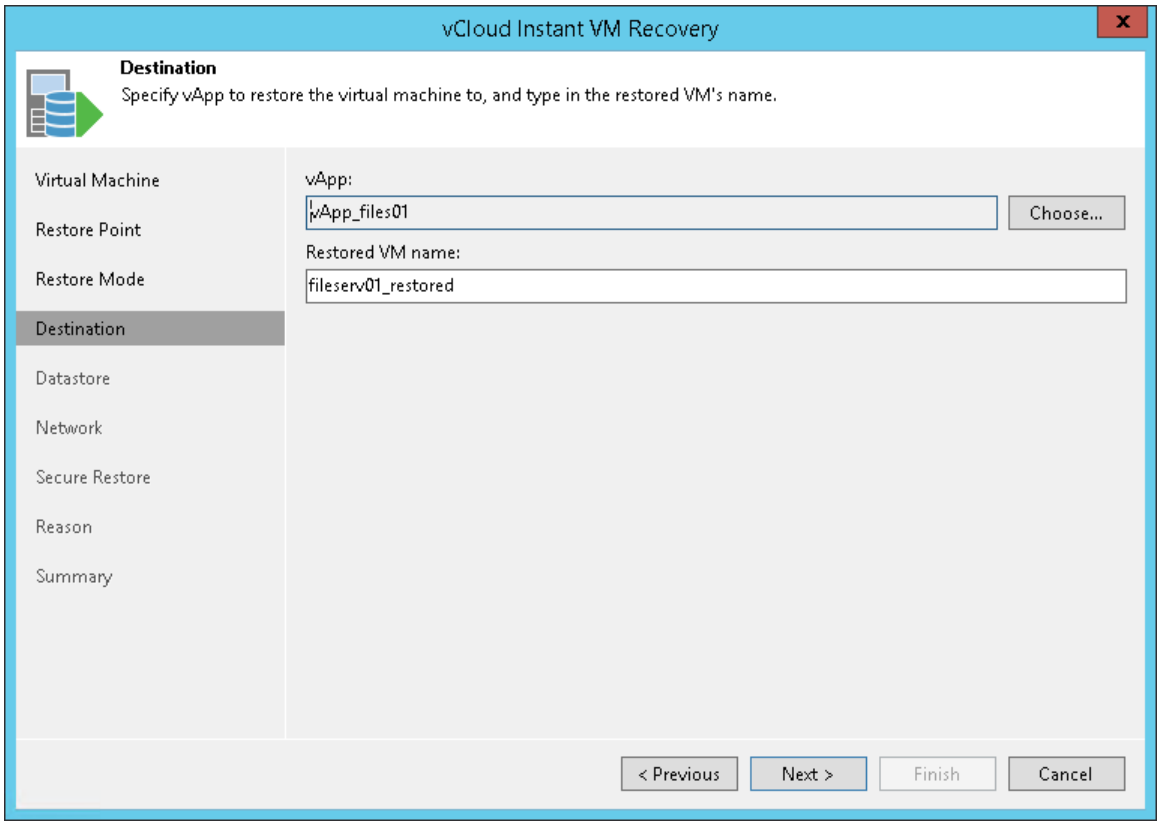
Step 4. Select Destination for Restored VM

The **Destination** step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select a destination and specify a name for the restored VM:

1. In the **vApp** field, specify a vApp to which the VM must be restored. By default, Veeam Backup & Replication restores the VM to its original vApp.
2. In the **Restored VM name** field, enter a name under which the VM must be restored and registered. By default, Veeam Backup & Replication uses the original name of the VM. If you are restoring the VM to the same vApp where the original VM is registered and the original VM still resides there, it is recommended that you change the VM name to avoid conflicts.

> **NOTE**
>
> Veeam Backup & Replication checks the lease term for the vApp to which the VM is restored. In case the lease period has expired, the lease will be automatically updated.



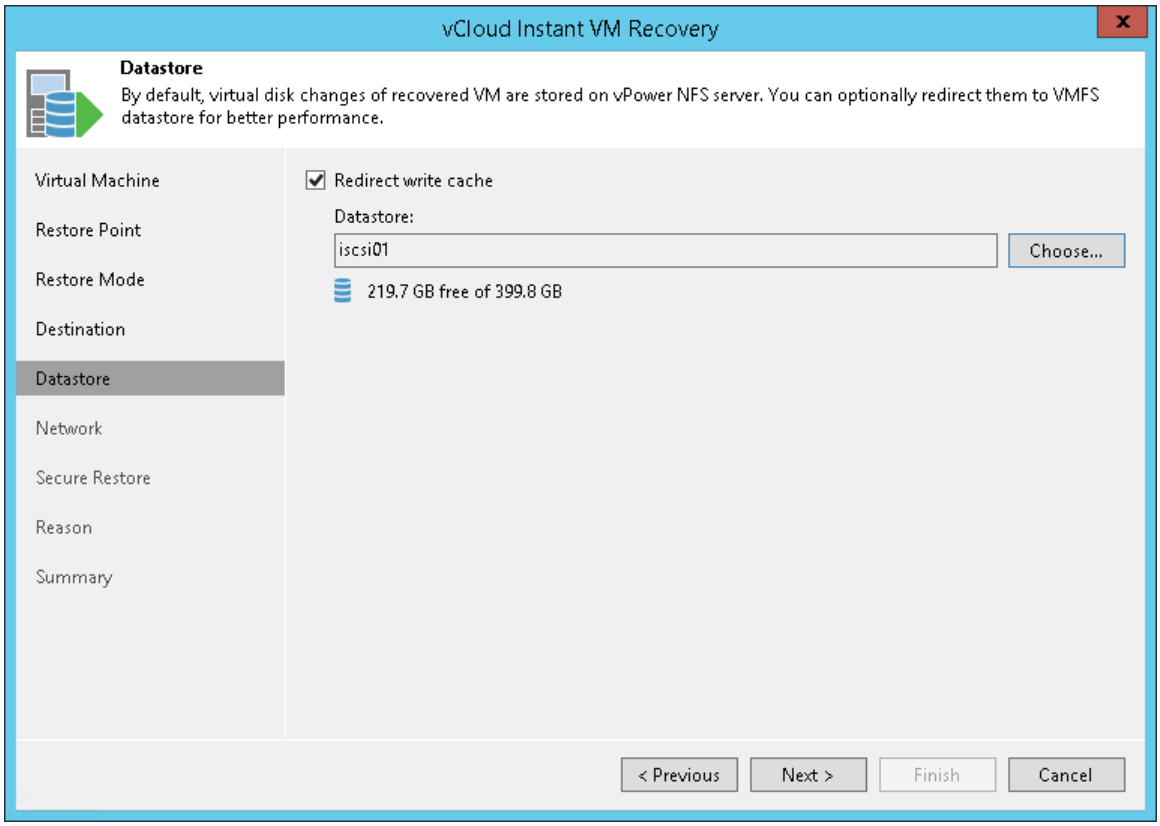
Step 5. Select Destination for Virtual Disk Updates

The **Datastore** step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select the location for holding the VM disk changes when the VM is restored. By default, disk changes are stored directly on the vPower NFS server. However, you can store disk changes on any datastore in your virtual environment.

To select a datastore:

1. Select the **Redirect virtual disk updates** check box.
2. From the **Datastore** list, choose the necessary datastore. You can select only a datastore that is available in the organization VDC hosting the vApp to which the VM is restored.



Step 6. Select Destination Network

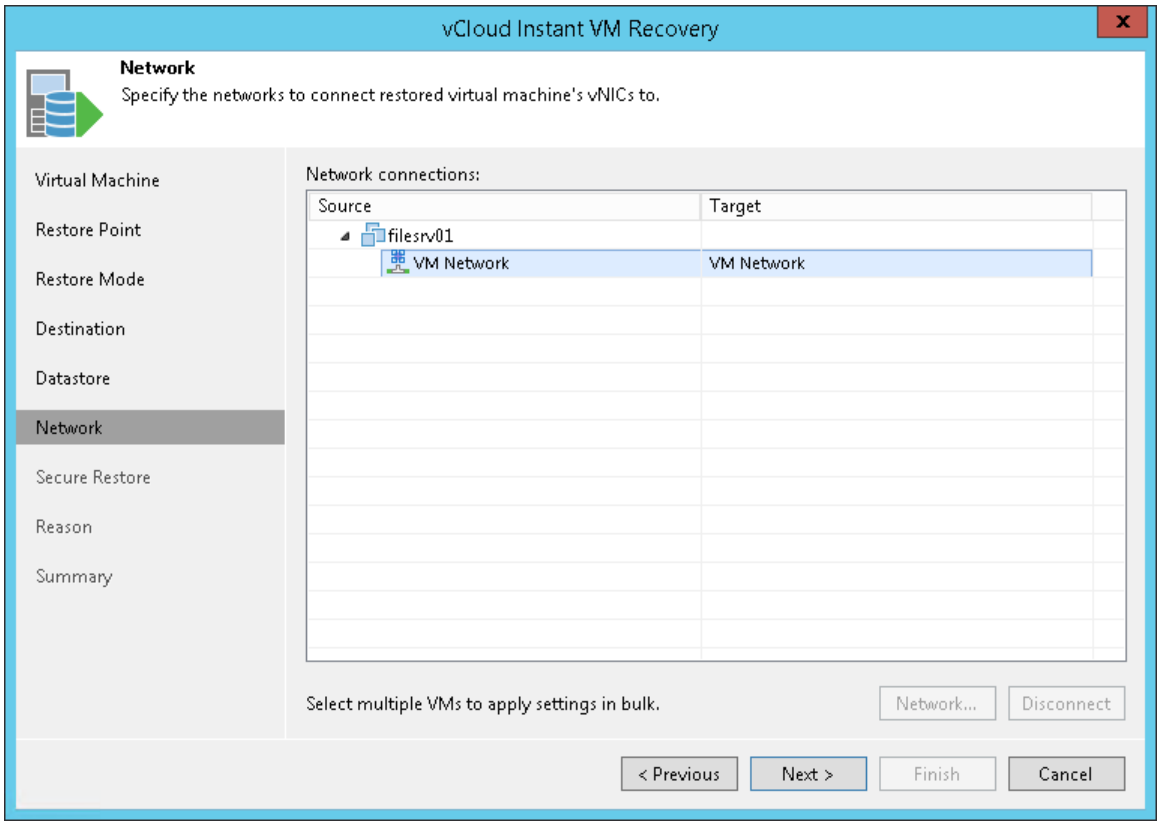
The **Network** step of the wizard is available if you have chosen to change the location and settings of the restored VM.

To select networks to which restored VMs must be connected:

1. Select a VM in the list and click **Network**.
2. The **Select Network** window displays all networks that are configured for the destination vApp. From the list of available networks, choose a network to which selected VM should have access upon restore.

   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.
3. To prevent the restored VM from accessing any network, select it in the list and click **Disconnect**.

Veeam Backup & Replication maps the network settings you define and network settings of the original VM. If necessary, Veeam Backup & Replication makes changes to the network settings of the recovered VM. For example, if the original VM was connected to the network using the static IP mode and you have selected to connect a recovered VM to a network using the dynamic IP mode, Veeam Backup & Replication will change the network settings to the dynamic mode.



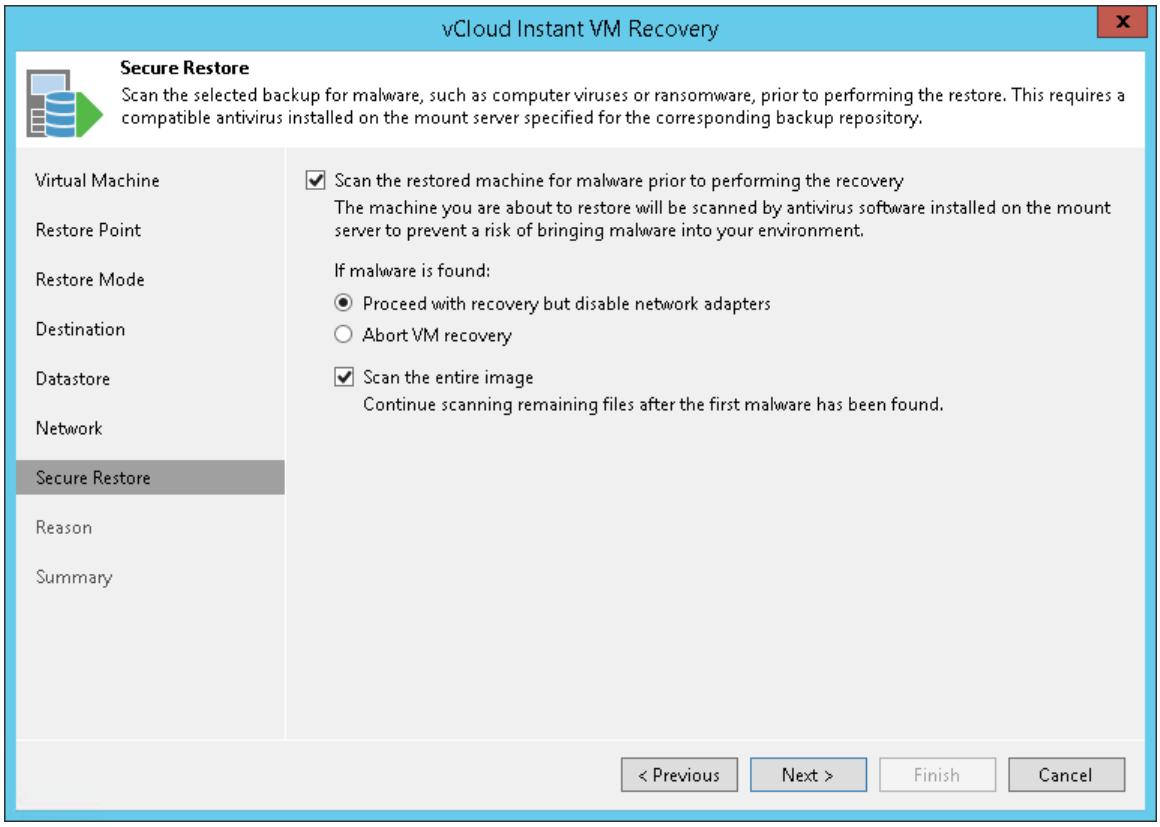
Step 7. Specify Secure Restore Settings

This step is available if you restore Microsoft Windows VMs and restore them to a new location or with different settings.

You can instruct Veeam Backup & Replication to perform secure restore — scan VM data with antivirus software before restoring the VM to the production environment. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

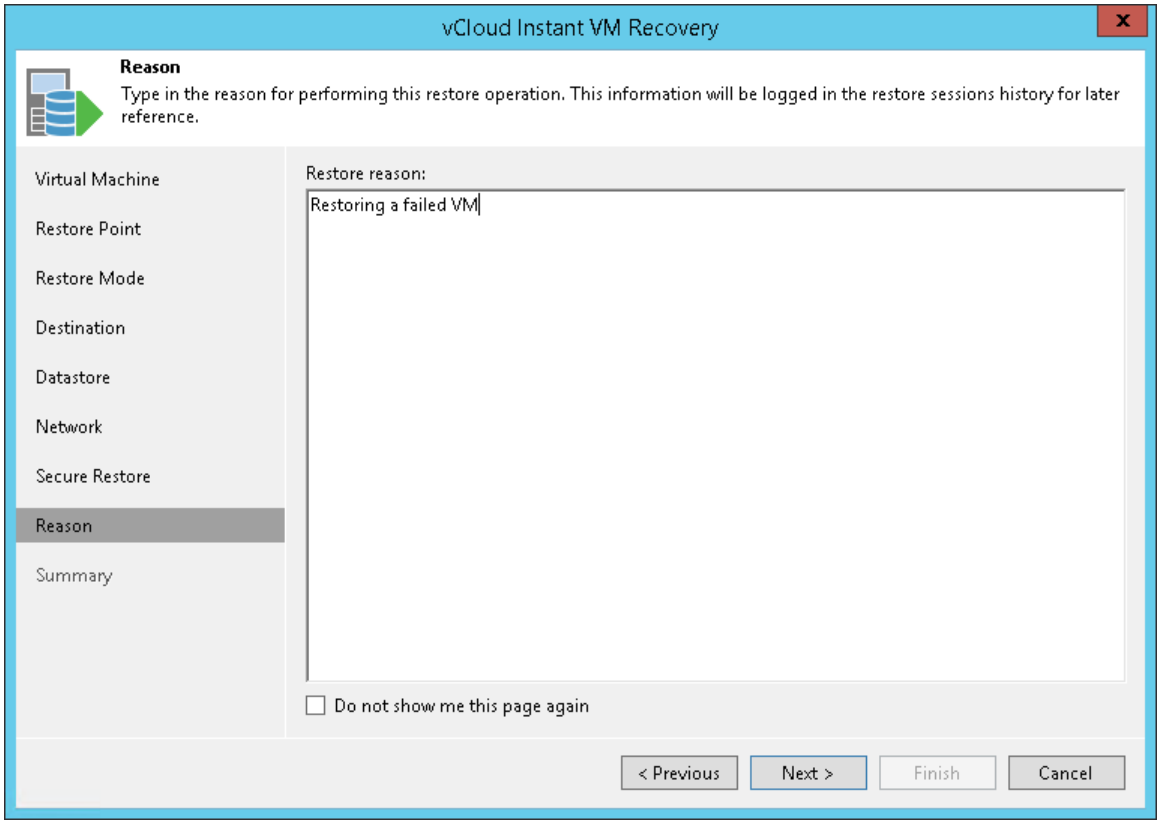
1. At the **Secure Restore** step of the wizard, select the **Scan the restored machine for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software finds a virus threat:
   - **Proceed with recovery but disable network adapters.** Select this action if you want to restore the VM with disabled network adapters (NICs).
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue the VM data scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



Step 8. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for performing Instant Recovery of the VM. The information you provide will be saved in the session history and you can reference it later.

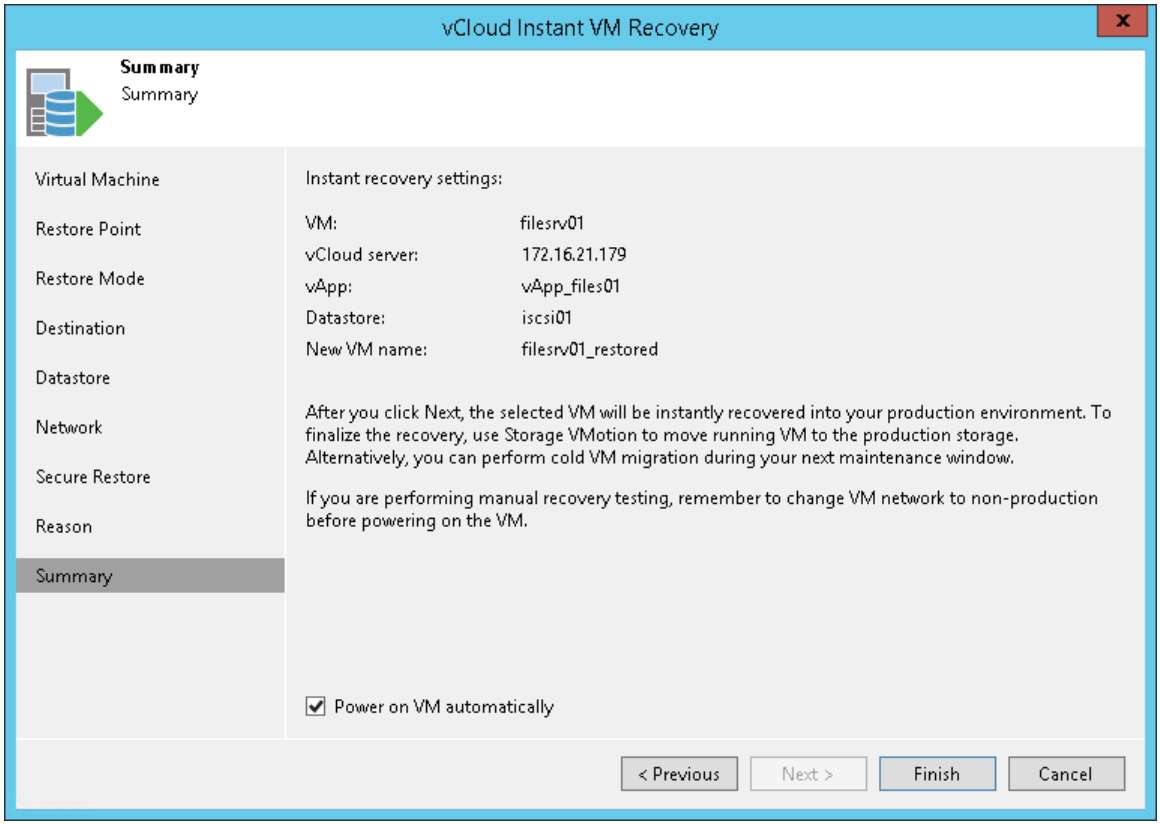
> **TIP**
>
> If you do not want to display the **Reason** step of the wizard in future, select the **Do not show me this page again** check box.



Step 9. Verify Instant Recovery Settings

At the **Summary** step of the wizard, specify additional settings for Instant Recovery:

1. If you want to start the recovered VM, select the **Power on VM automatically** check box.
2. Check the specified settings of Instant Recovery and click **Finish**. Veeam Backup & Replication will recover the selected VM in the specified destination.



Step 10. Finalize Instant Recovery

All VMs restored with Instant Recovery are displayed in the **Home** view, under the **Instant Recovery** node.

To check the progress of Instant Recovery and view session details:

1. Open the **Home** view.
2. In the inventory pane, click the **Instant Recovery** node.
3. Right-click the VM in the working area and select **Properties**.

Alternatively, you can open the **History** view, select the **Instant Recovery** node under **Restore** in the inventory pane and double-click the necessary instant VM restore session.

After the VM has been successfully recovered, you can finalize Instant Recovery in one of two ways:

- Migrate the recovered VM to the production environment
- Unpublish the recovered VM

Migrating Recovered VM

To migrate the restored VM to production:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM and select **Migrate to production**. Veeam Backup & Replication will launch the Quick Migration wizard.

   During migration, Veeam Backup & Replication will restore the VM from the backup file and additionally move all changes that were made while the VM was running from the backup in the Instant Recovery mode.

> **TIP**
>
> When you pass though the **Quick Migration** wizard, enable the **Delete source VM files upon successful migration** option. Veeam Backup & Replication will restore the VM to production and automatically stop the Instant Recovery session. If you do not enable this option, the Instant Recovery session will still be running, and you will need to unpublish the recovered VM manually.

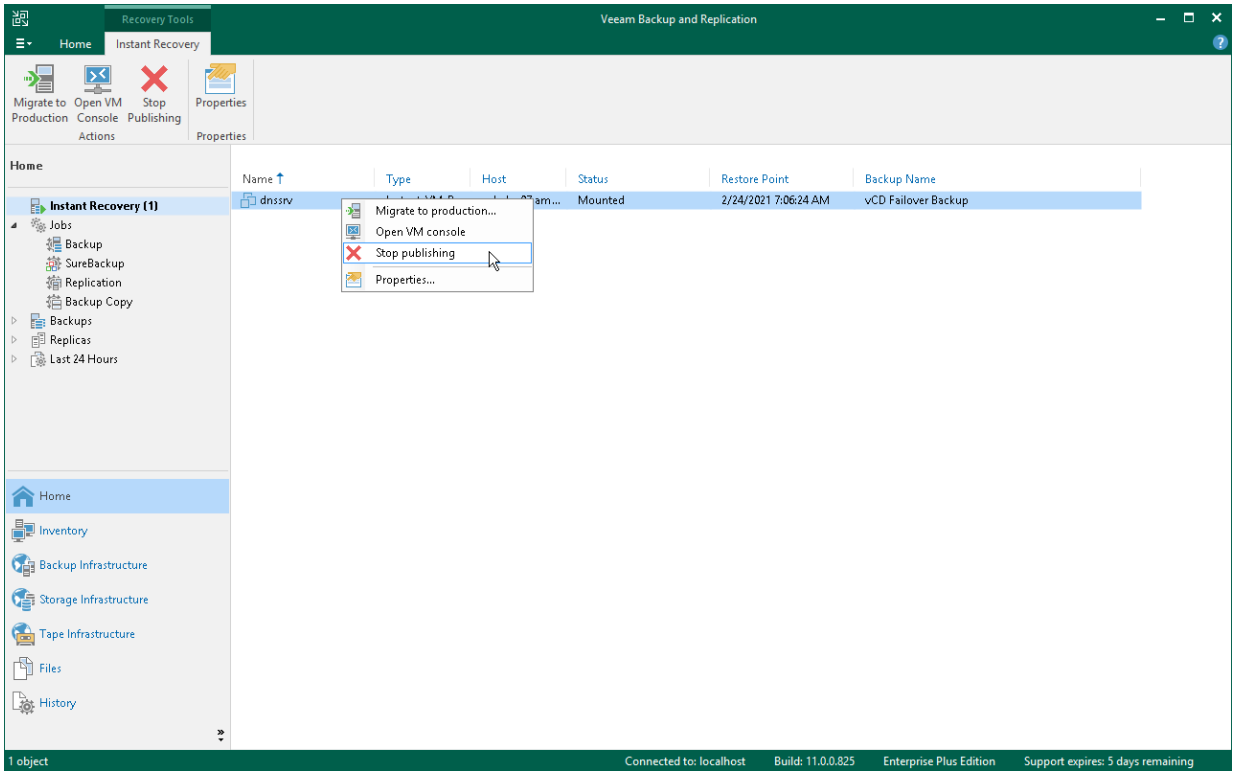
Unpublish Recovered VM

If you have disabled the **Delete source VM files upon successful migration** option in the Quick Migration settings, you must unpublish the VM manually. After you unpublish the VM, the Instant Recovery session will end and the recovered VM will be unmounted from the vPower NFS server. The migrated VM will remain on the production environment.

To unpublish a recovered VM:

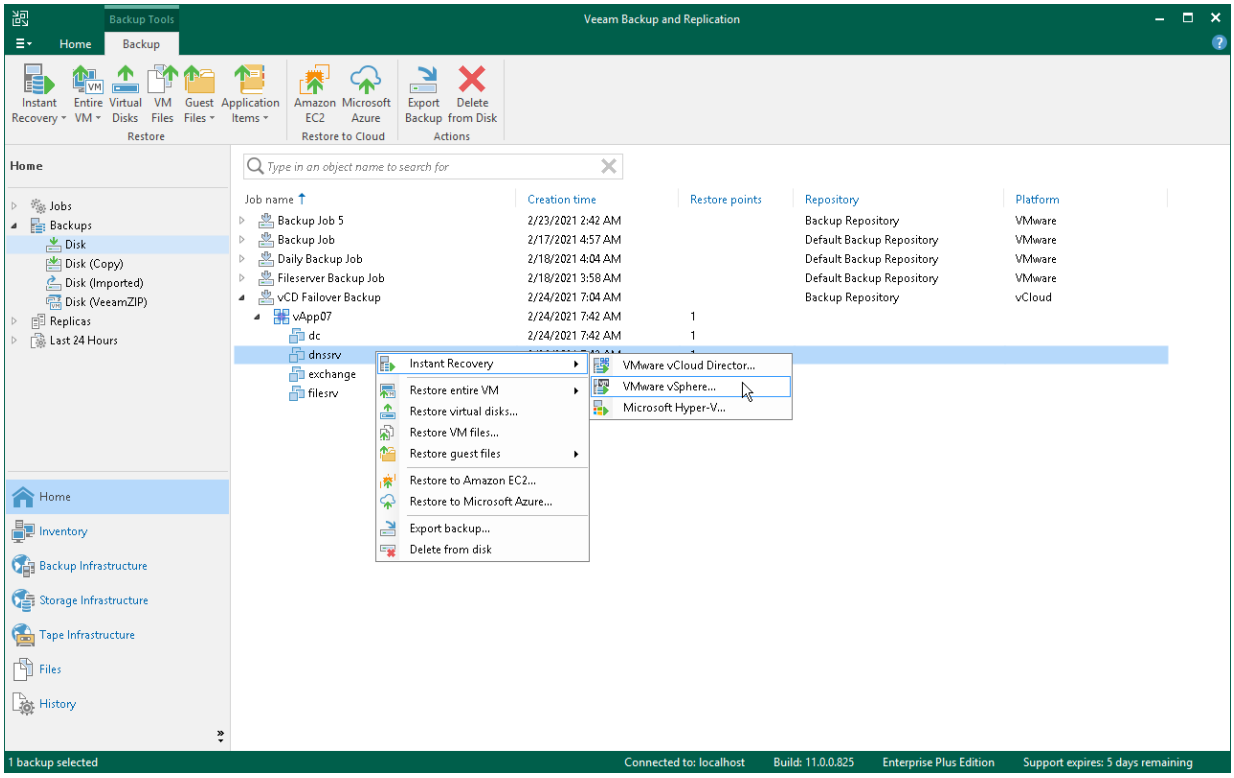
1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM and select **Stop publishing**.

> **TIP**
>
> After the VM has been published from the backup, you can open the VM console directly from Veeam Backup & Replication. To do this, in the working area right-click the VM and select **Open VM Console**.



Restoring VMs with Instant Recovery into vSphere infrastructure

The process of Instant Recovery for vCloud Director VMs does not differ from the regular Instant Recovery process. For more information, see Performing Instant Recovery to VMware vSphere.



Restoring vCloud vApps

You can restore the whole vApp from the backup to vCloud Director.

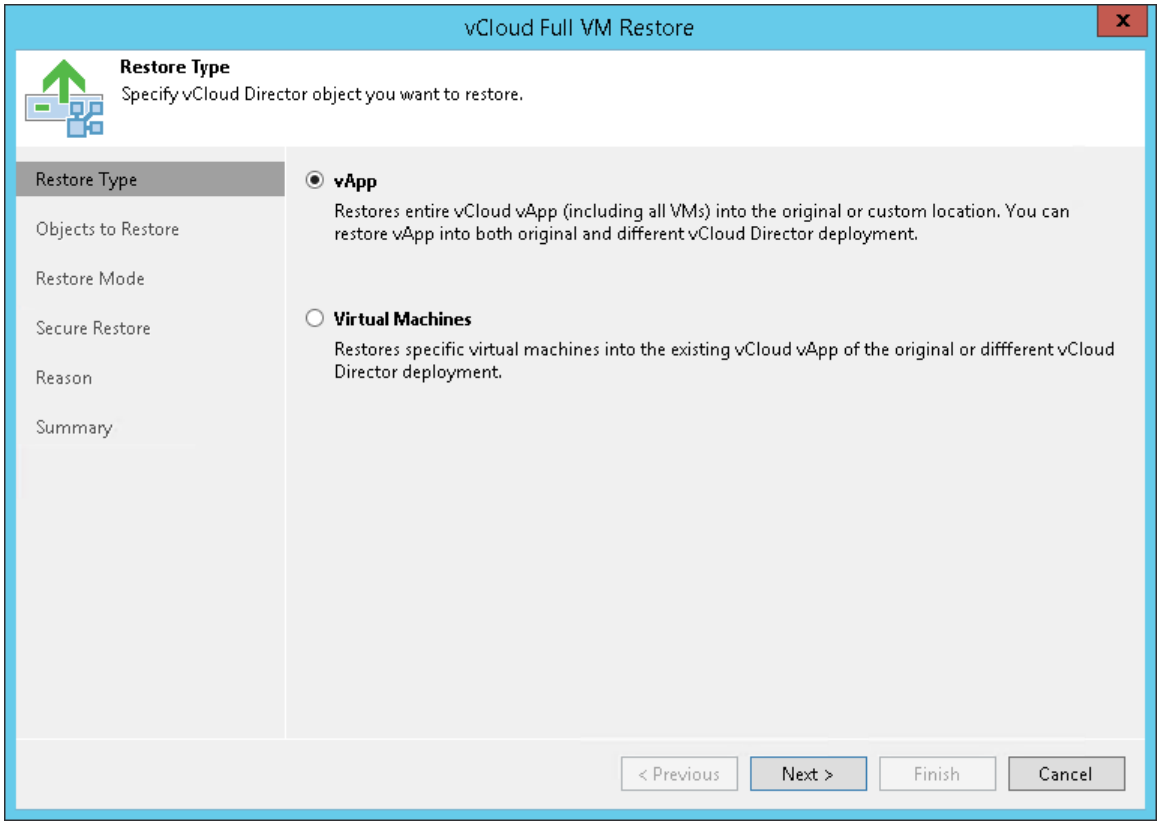
vApps can be restored to their organization VDC or to any other organization VDC. You can restore the vApp that already exists, for example, in case the vApp is corrupted or you want to revert to an earlier state of the vApp, or the vApp that no longer exists, for example, if it was deleted by mistake. If you restore a vApp that already exists, the vApp is overwritten with that from the vCD backup.

To restore a vApp to vCloud Director, use the **vCloud Full vApp Restore** wizard.

Step 1. Launch Full vApp Restore Wizard

To launch the **Full vApp Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware vCloud Director**. At the **Restore Type** step of the wizard, select the object you would like to restore: *vApp*.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Select the vApp and click **Restore vCloud vApp** on the ribbon.
  - Right-click the vApp and select **Restore vCloud vApp**.



Step 2. Select vApp to Restore

At the **Objects to Restore** step of the wizard, select the vApp you want to restore.

To add a vApp, click **Add vApp** and select where to browse for vApps:

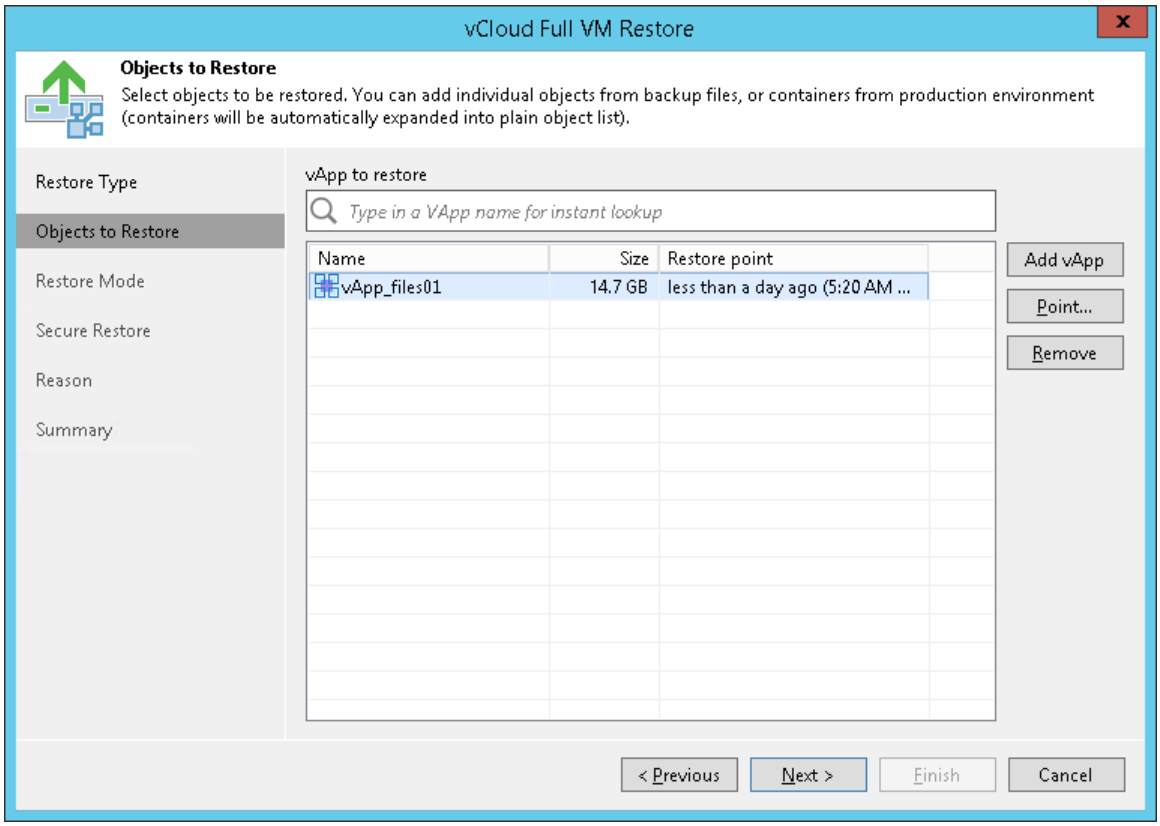
- **From infrastructure** — browse the vCloud Director hierarchy and select a vApp to restore. Note that the vApp you select from the vCloud Director hierarchy must be successfully backed up at least once.
- **From backup** — browse existing backups and select the vApp under backup jobs.

To facilitate selection, use the search field at the bottom of the **Select VMs** window: enter an object's name or a part of it in the search field and click the **Start search** button on the right or press [Enter] on the keyboard.

You can also use the search field at the top of the window to add vApps:

1. Enter a vApp name or a part of it in the search field. Veeam Backup & Replication will search existing backups for the specified vApps and display matching results.
2. To add the vApp to the list, double-click it in the list of search results.
3. If the necessary vApp is not found, click the **Show more** link to browse existing backups and choose the necessary vApp.

To remove a vApp from the list, select it and click **Remove** on the right.



Step 3. Select Restore Point

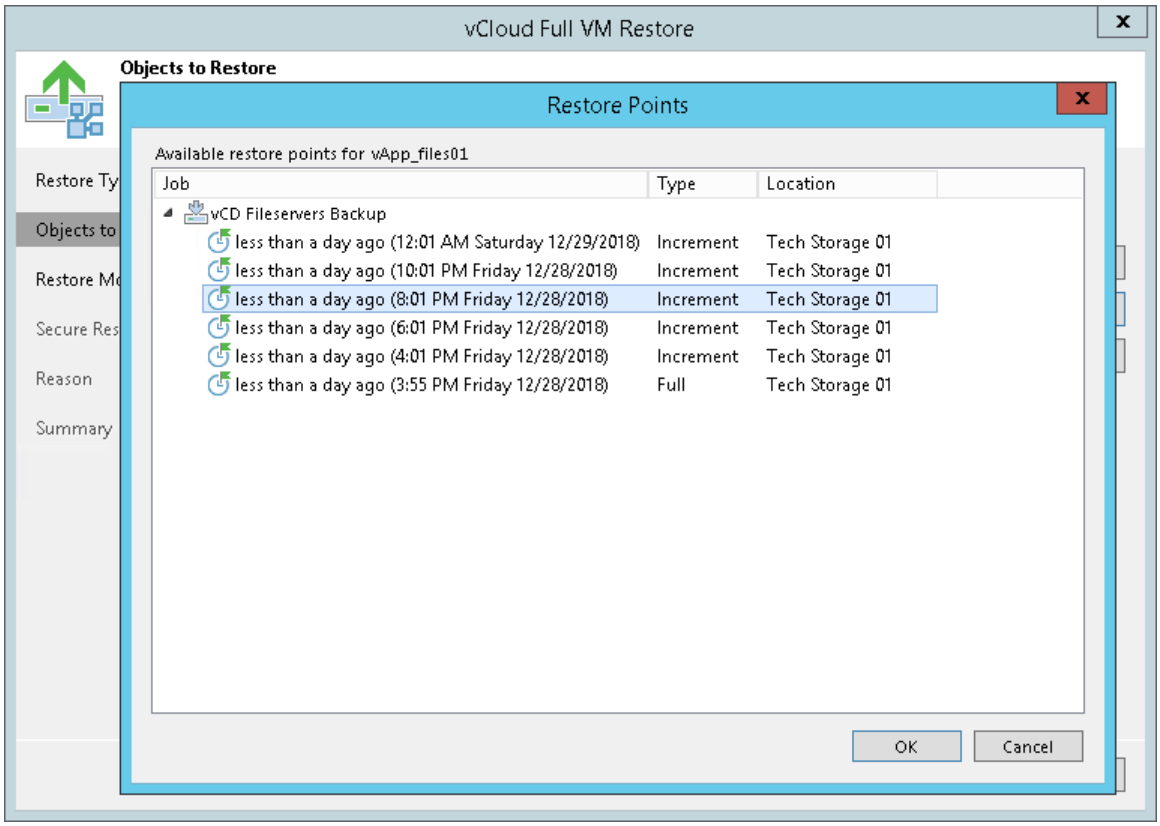
You can select the restore point for the vApp.

By default, Veeam Backup & Replication uses the latest valid restore point to recover the vApp. However, you can restore the vApp to an earlier state.

To select a restore point for the vApp:

1. Select the vApp in the list and click **Point** on the right.
2. In the **Restore Points** window, select a restore point that must be used to recover the vApp.

   In the **Location** column, you can view a name of a backup repository where a restore point resides.



Step 4. Select Restore Mode

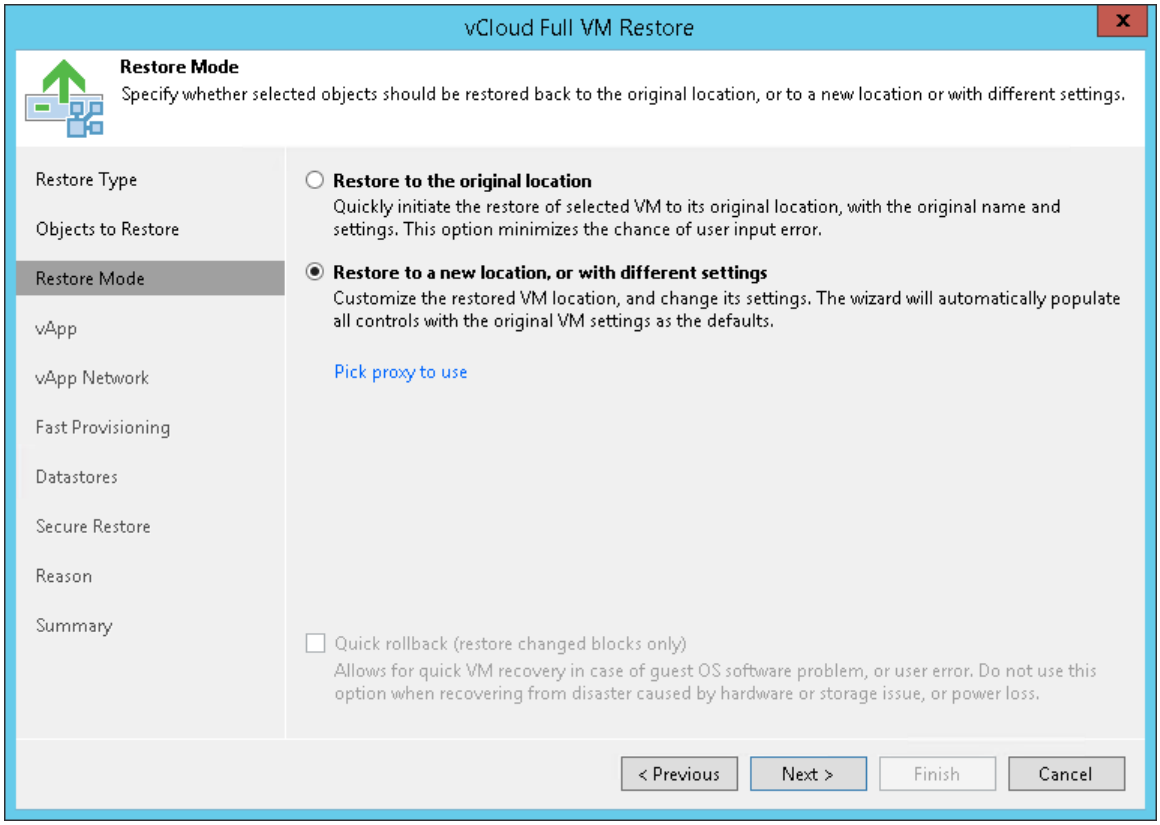
At the **Restore Mode** step of the wizard, choose the necessary restore mode and backup proxy for VM data transport:

1. Choose a restore mode:
   - Select **Restore to original location** if you want to restore the vApp with its initial settings and to its original location. If this option is selected, you will immediately pass to the Summary step of the wizard.
   - Select **Restore to a new location, or with different settings** if you want to restore the vApp to a different location and/or with different settings (such as organization VDC, network settings, fast provisioning settings and so on). If this option is selected, the **vCloud Full vApp Restore** wizard will include additional steps for customizing vApp settings.
2. [For vApp restore to the original location] Select the **Quick rollback** check box if you want to use incremental restore for the vApp. Veeam Backup & Replication will query CBT to get data blocks that are necessary to revert the vApp to an earlier point in time, and will restore only these data blocks from the backup. Quick rollback significantly reduces the restore time and has little impact on the production environment.

   It is recommended that you enable this option if you restore VMs in the vApp after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.
3. Click the **Pick proxy to use** link to select backup proxies over which vApp data must be transported to the source datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing vApp data.

     During the restore process, VMs in the vApp are processed simultaneously.
     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 writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VMs processing.

- If you choose **Use the selected backup proxy serves only**, you can explicitly select backup proxies that will be used for restore. It is recommended that you select at least two proxies to ensure that VMs are recovered should one of backup proxies fail or lose its connectivity to the source datastore during restore.



Step 5. Select vApp Location

The **vApp** step of the wizard is available if you have chosen to change the location and settings of the restored vApp.

By default, Veeam Backup & Replication restores the vApp to its original location with its original name.

To restore the vApp to a different location:

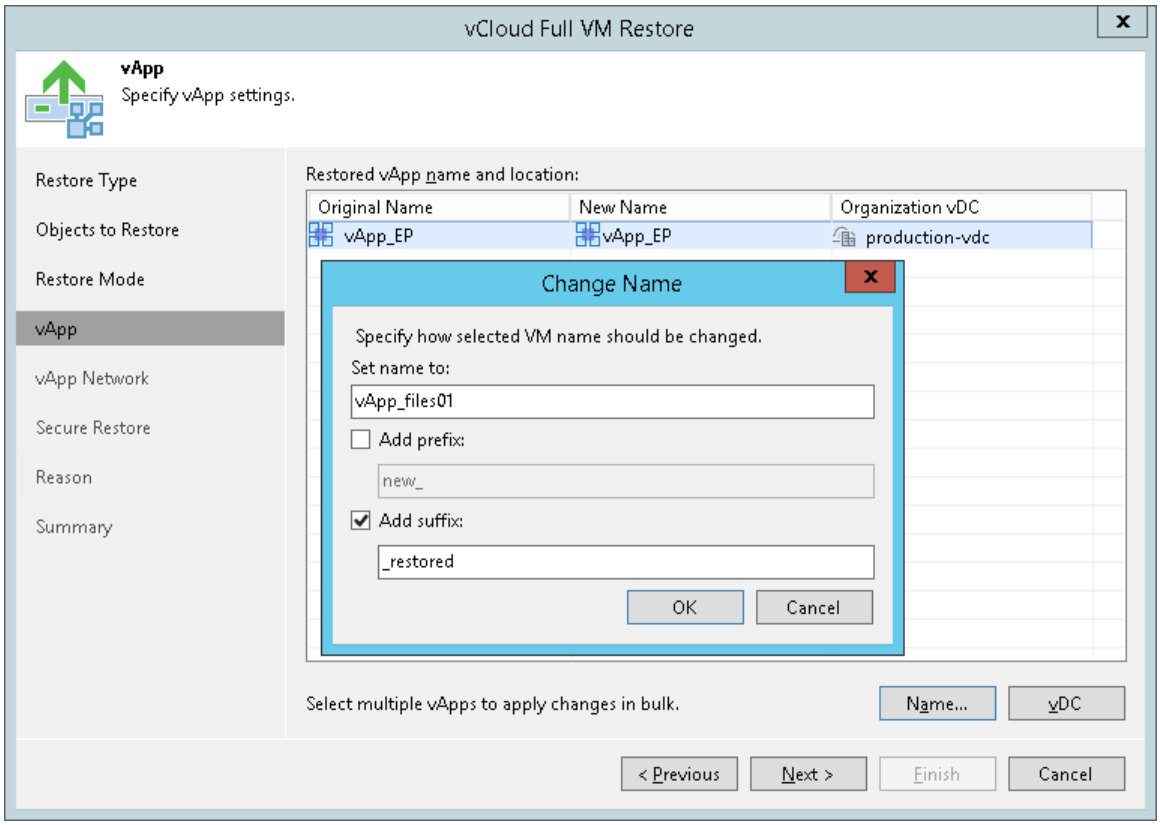
1. Select the App in the list and click **vDC**.
2. From the vCloud Director hierarchy, choose an organization VDC where the selected vApp must be registered.

   To facilitate selection, use the search field at the bottom of the window: enter an object's name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

> **NOTE**
> If a vApp contains a standalone VM, Veeam Backup & Replication restores the standalone VM in such vApp as a regular vCloud Director VM.

To change the vApp name:

1. Select the vApp in the list and click **Name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule by adding a prefix and/or suffix to the original vApp name.
3. You can also change the vApp name directly in the list: select a vApp, click the **New Name** field and enter the name to be assigned to the recovered vApp.



Step 6. Select Destination Network

The **vApp Network** step of the wizard is available if you have chosen to change the location and settings of the restored vApp.

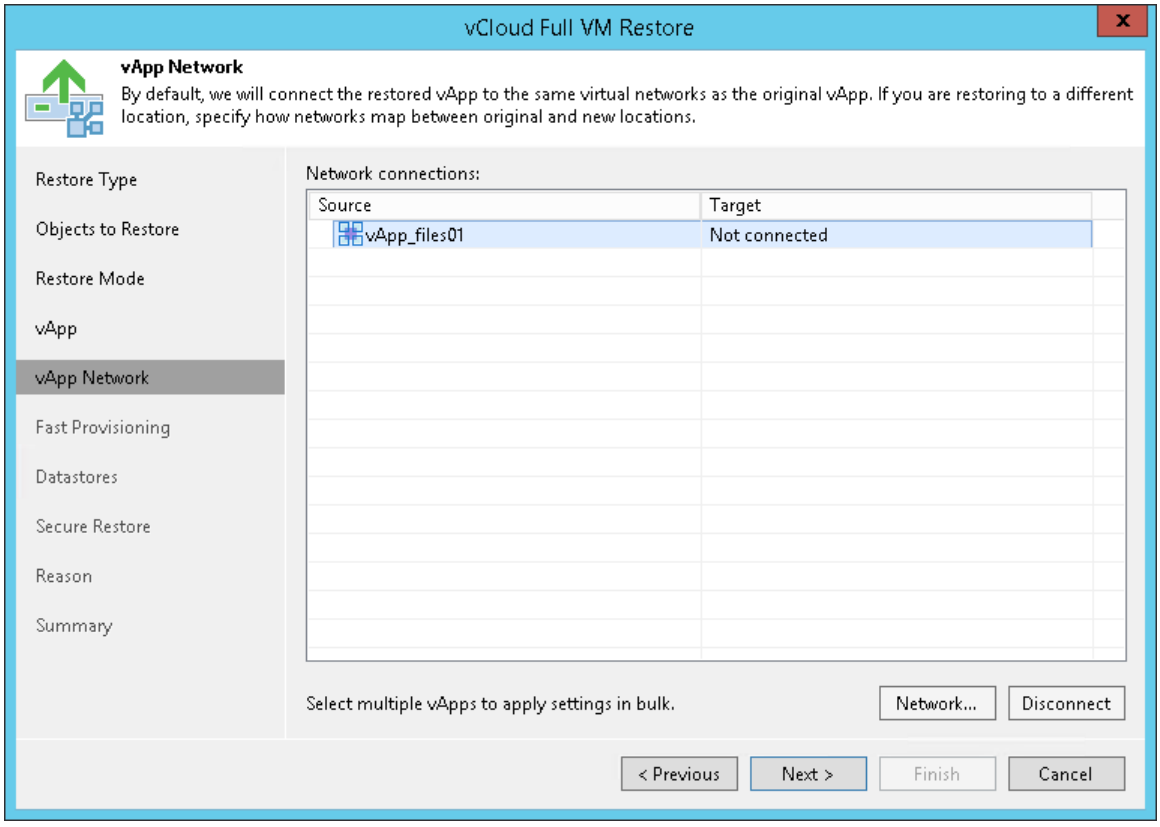
To select networks to which the restored vApp must connected:

1. Select the vApp in the list and click **Network**.
2. The **Select Network** window displays all networks that are configured for the destination organization VDC. From the list of available networks, choose a network to which selected vApp must have access upon restore.

   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

To prevent the restored vApp from accessing any network, select it in the list and click **Disconnect**.

> **NOTE**
>
> When Veeam Backup & Replication backs up a vApp, along with vApp networks data it saves information about organization networks to which the vApp is connected. If you restore the vApp to the original vDC organization and do not change the organization network settings, Veeam Backup & Replication attempts to connect the vApp to all source organization networks to which the vApp was connected at the moment of backup. If one or several source organization networks are not detected, for example, if they have been changed or removed by the time of restore, Veeam Backup & Replication will not be able to restore the vApp to its original organization. In this case, you will need to change the network settings for the restored vApp: map the vApp to other organization networks or disconnect it from organization networks at all.



Step 7. Select Template to Link

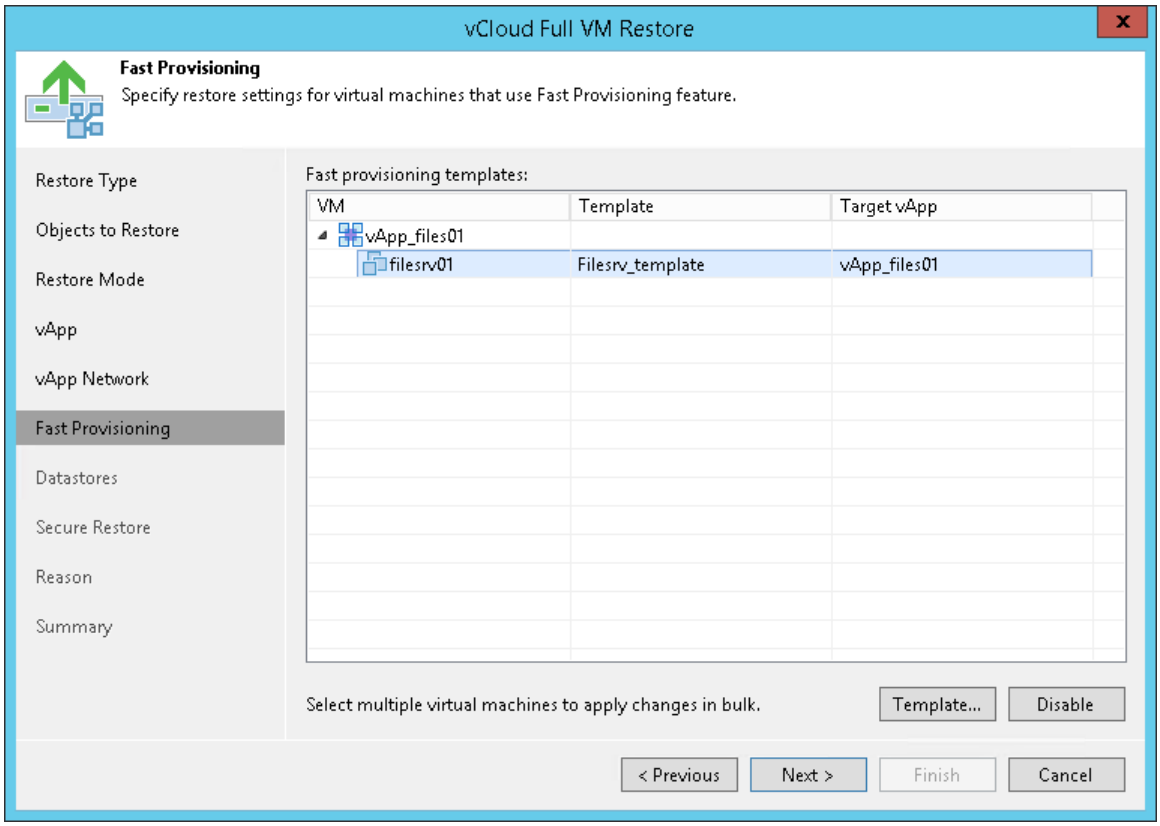
The **Fast Provisioning** step of the wizard is available if you have chosen to change settings of the restored vApp, for example, its name or location.

To select a VM template:

1. Select a VM in the list and click **Set Template**.
2. From the vCloud Director hierarchy, choose a template to which the VMs from the restored vApp must be linked.

   To facilitate selection, use the search field at the bottom of the window: enter a VM template name or a part of it and click the Start search button on the right or press [Enter] on the keyboard.

If you want to disable fast provisioning for the VM and restore it as a regular VM, select the VM in the list and click **Disable**.



Step 8. Select Storage Policy and Datastores

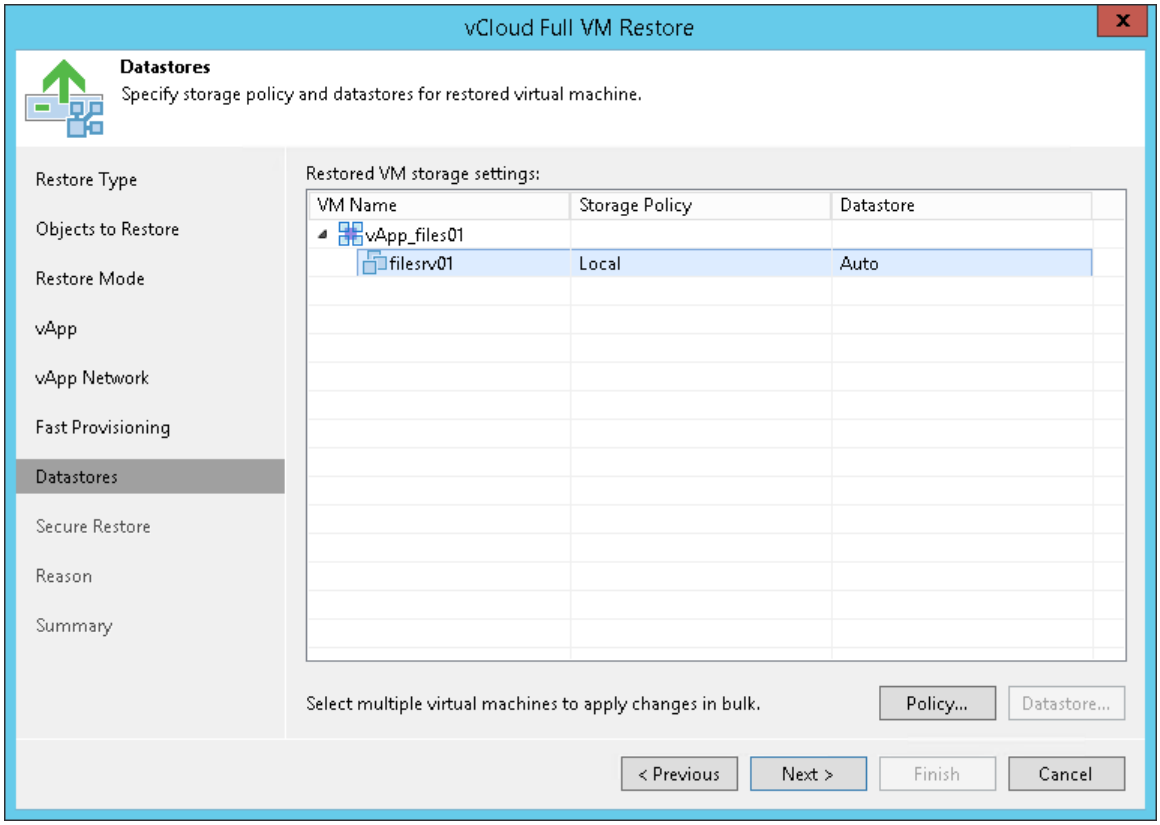
The **Datastores** step of the wizard is available if you have chosen to change settings of the restored vApp, for example, its name or location.

To select a storage policy for the vApp:

1. Select the vApp in the list and click **Policy**.
2. In the displayed window, select the necessary policy for the vApp.

If you have selected to disable fast provisioning at the previous step of the wizard, you must select a datastore on which the disks of restored VMs will be placed. To do this:

1. Select VM or vApp in the list and click **Datastore**.
2. In the displayed window, select the datastore on which the disks of the VM must be placed.

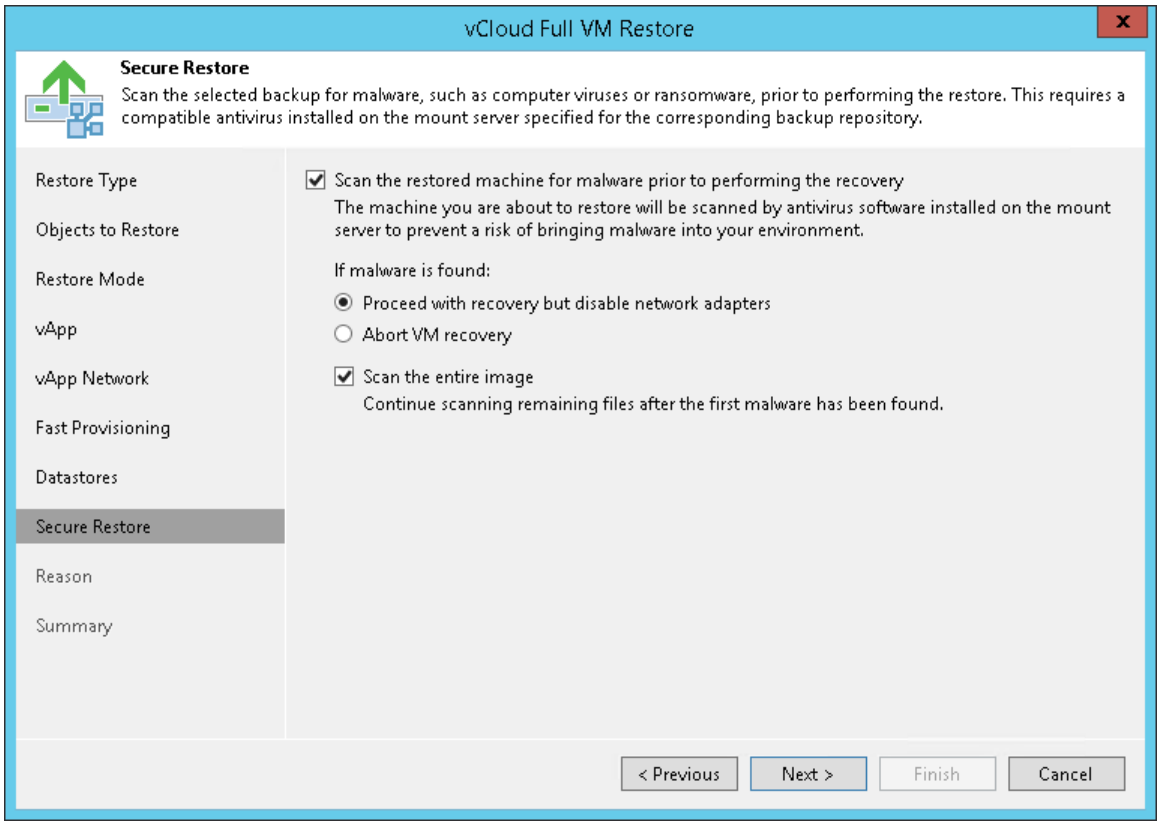


Step 9. Specify Secure Restore Settings

You can instruct Veeam Backup & Replication to perform secure restore — scan vApp data with antivirus software before restoring the vApp. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

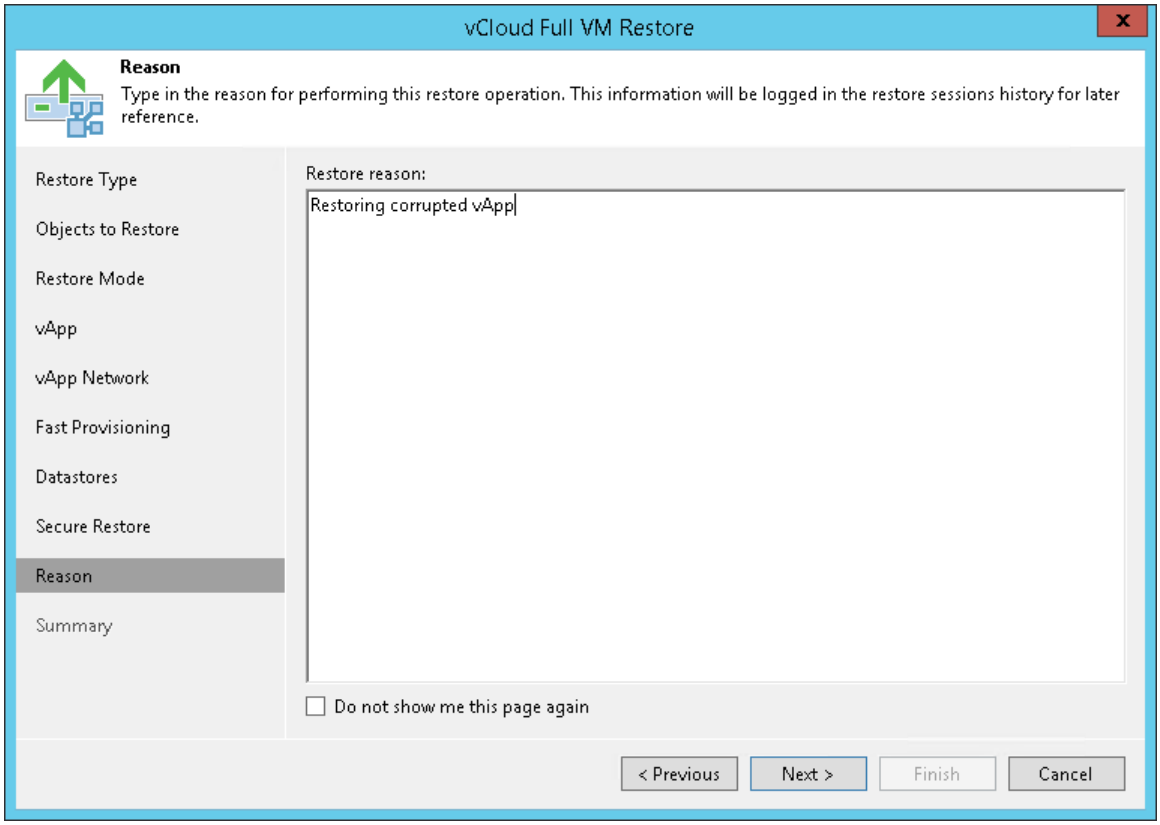
1. At the **Secure Restore** step of the wizard, select the **Scan the restored machine for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software finds a virus threat:
   - **Proceed with recovery but disable network adapters.** Select this action if you want to restore the vApp VMs with disabled network adapters (NICs).
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue the vApp data scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



Step 10. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the selected vApp. 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.

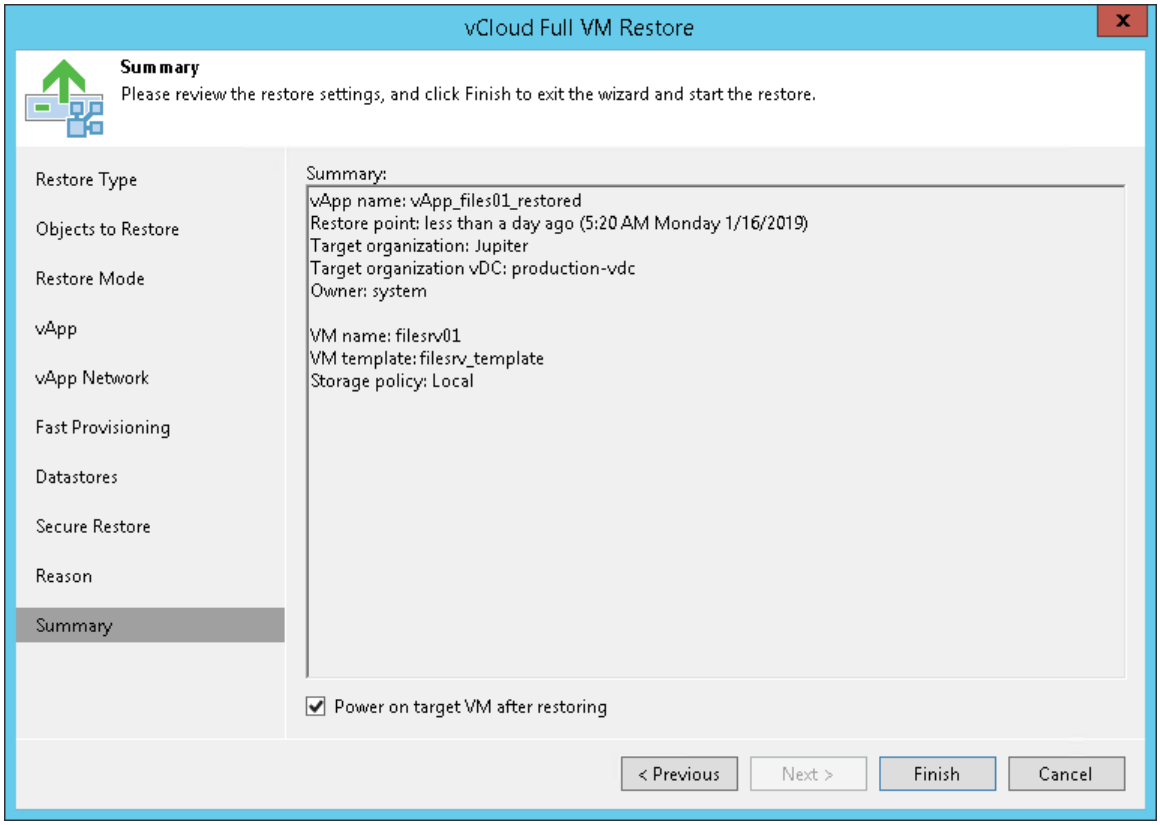


Step 11. Verify Recovery Settings and Finish Working with Wizard

At the **Summary** step of the wizard, specify additional settings for vApp restore:

1. If you want to start VMs in the restored vApp, select the **Power on VM after restoring** check box.
2. Check the settings for vApp restore and click **Finish**. Veeam Backup & Replication will recover the vApp in the specified destination.

> **NOTE**
>
> Veeam Backup & Replication checks the lease term for the restored vApp. If the lease period has expired, the lease will be automatically updated.



Restoring VMs into vCloud vApp

You can restore one or several VMs from vCD backups back to vCloud Director.

The vCloud Director VM can be restored to its original location — to a vApp in which the VM is already registered, or to a different location. You can restore a VM that already exists, for example, if the original VM is corrupted or you want to revert to an earlier state of the VM, or a VM that no longer exists, for example, if the VM was deleted by mistake. If you restore a VM that already exists, the original VM is overwritten with that from the vCD backup.

When restoring VMs to the vCloud Director hierarchy, make sure that you select the **Restore into vCloud vApp** option. If you select the **Restore into vSphere infrastructure** option, the VM will be restored at the level of the underlying vCenter Server. To get a fully functional VM managed by vCloud Director, you will need to manually import the restored VM to the vCloud Director hierarchy.

Before you restore a VM to the vCloud Director hierarchy, check prerequisites. Then use the **vCloud Full VM Restore** wizard to restore the necessary VM.

Before You Begin

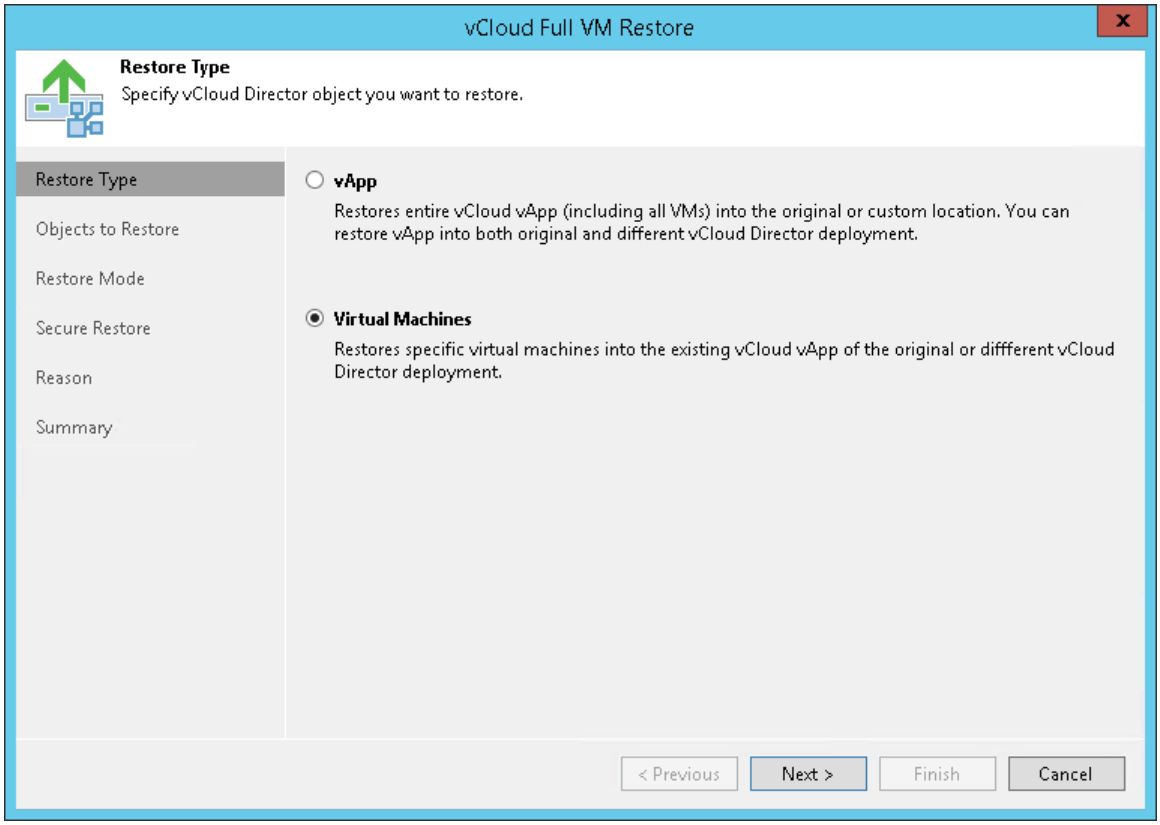
Before you restore vCloud Director VMs to a vApp, consider the following:

- Veeam Backup & Replication restores standalone VMs as regular vCloud Director VMs in the following cases:
  - If you selected to restore a standalone VM to a different location.
  - If the original standalone VM no longer exists in the vCloud DIrector hierarchy.
- It is not recommended to restore vCloud Director VMs to a vApp that already contains a standalone VM. After the restore process is complete, the vApp may not work as expected.
- If you restore linked clone VMs to a different location, make sure that fast provisioning is enabled at the level of the target organization VDC. Otherwise, Veeam Backup & Replication will restore the linked clone VM to a selected vApp as a regular VM.
- If you want to scan VM data for viruses, check the secure restore requirements and limitations.

Step 1. Launch vCloud Full VM Restore Wizard

To launch the **vCloud Full VM Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware vCloud Director**. At the **Restore Type** step of the wizard, select the object you would like to restore: *Virtual Machines*.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Select the VM you want to restore and click **Entire VM** > **Into vCloud vApp** on the ribbon.
  - Right-click the VM you want to restore and select **Restore entire VM** > **Into vCloud vApp**.
- Open the **Inventory** view. On the **View** tab, click **vCloud View**. In the inventory pane, expand the vCloud Director hierarchy. In the working area, right-click the VM you want to restore and select **Restore** > **Restore Entire VM** > **Into vCloud vApp**.



Step 2. Select VMs to Restore

At the **Objects to Restore** step of the wizard, select one or several VMs to restore.

To add a VM, click **Add VM** and select where to browse for VMs:

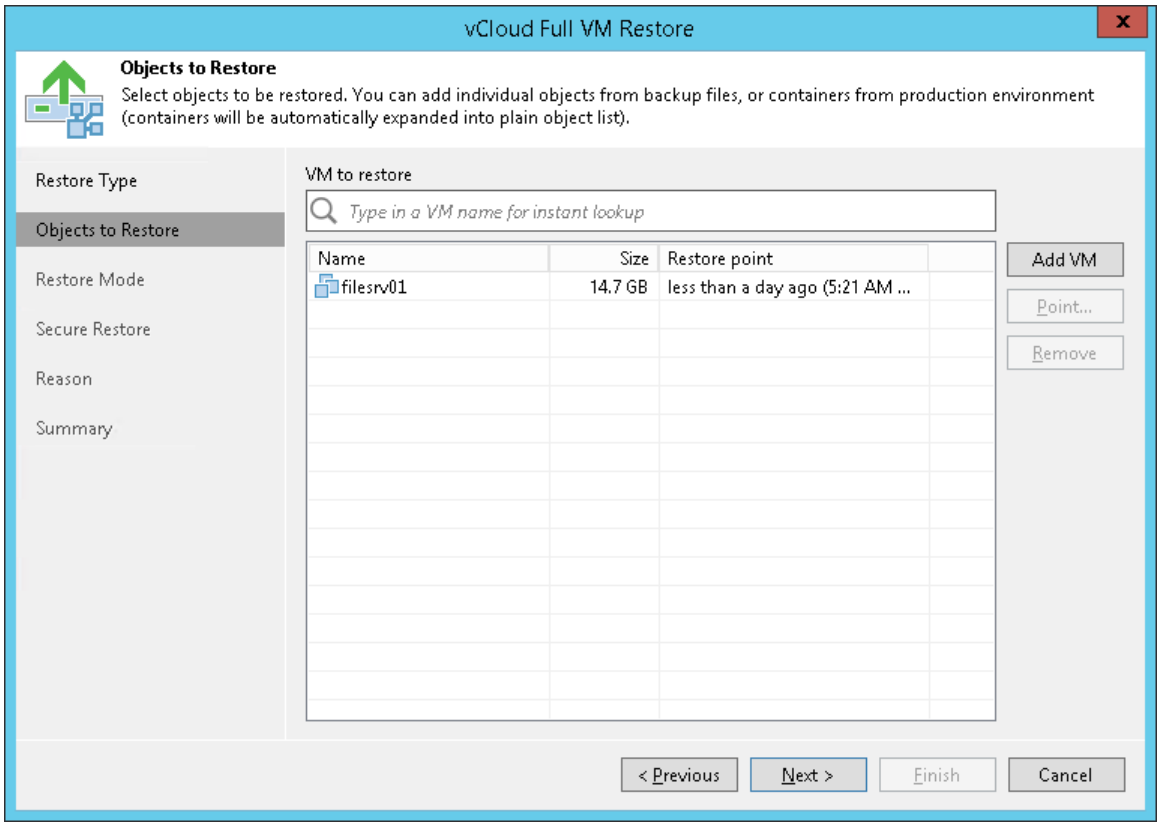
- **From infrastructure** – browse the vCloud Director hierarchy and select VMs to restore. Note that the VM you select from the vCloud Director hierarchy must be successfully backed up at least once.
- **From backup** – browse existing backups and select VMs under backup jobs.

To facilitate selection, use the search field at the bottom of the **Select VMs** window: enter an object's name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

To add VMs to the list, you can also use the search field at the top of the window:

1. Enter a VM name or a part of it in the search field and Veeam Backup & Replication will search existing backups for the specified VM and display matching results.
2. To add the VM to the list, double-click it in the list of search results.
3. If the necessary VM is not found, click the **Show more** link to browse existing backups and choose the necessary VM.

To remove a VM from the list, select it and click **Remove** on the right.



Step 3. Select Restore Point

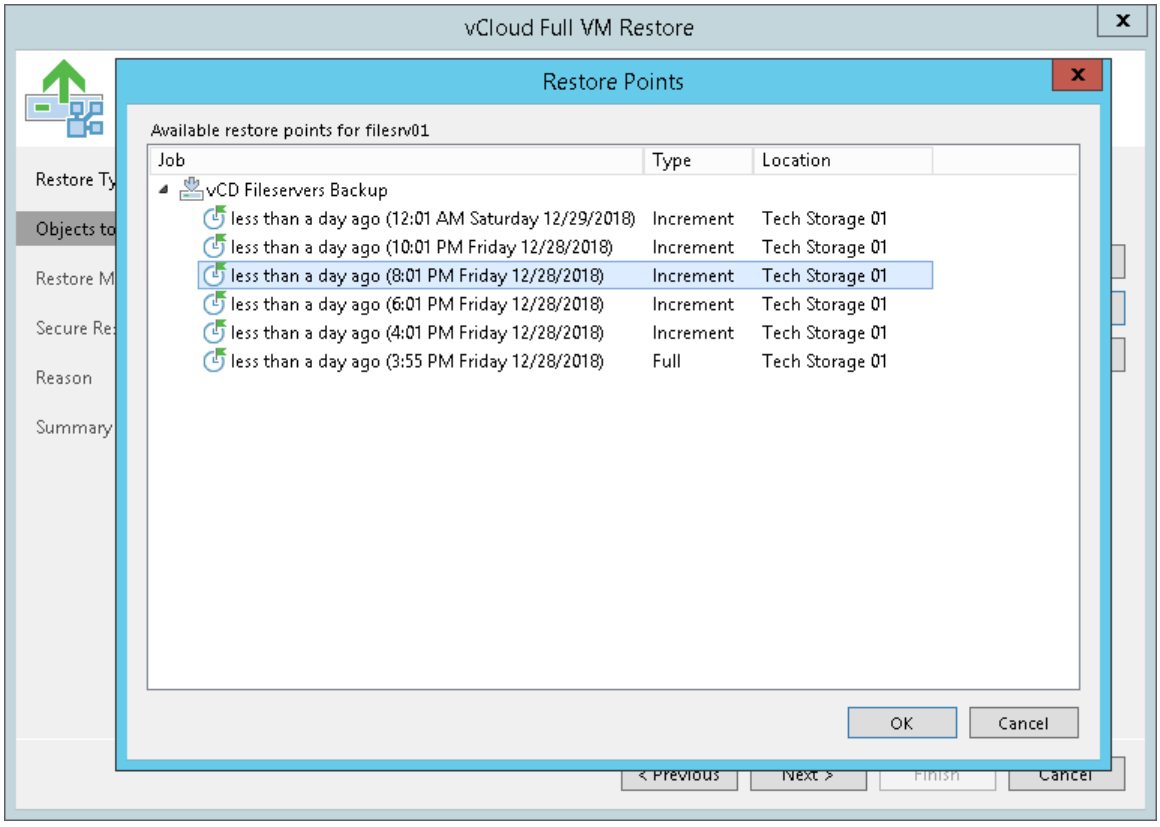
You can select the restore point for the VM.

By default, Veeam Backup & Replication uses the latest valid restore point to recover a VM. However, you can restore a VM to an earlier state. If you have chosen to restore multiple VMs, you can select a different restore point for every VM specifically.

To select a restore point for a VM:

1. Select a VM in the list and click **Point** on the right.
2. In the **Restore Points** window, select the restore point that must be used to recover the VM.

   In the **Location** column, you can view a name of a backup repository where a restore point resides.



Step 4. Select Restore Mode

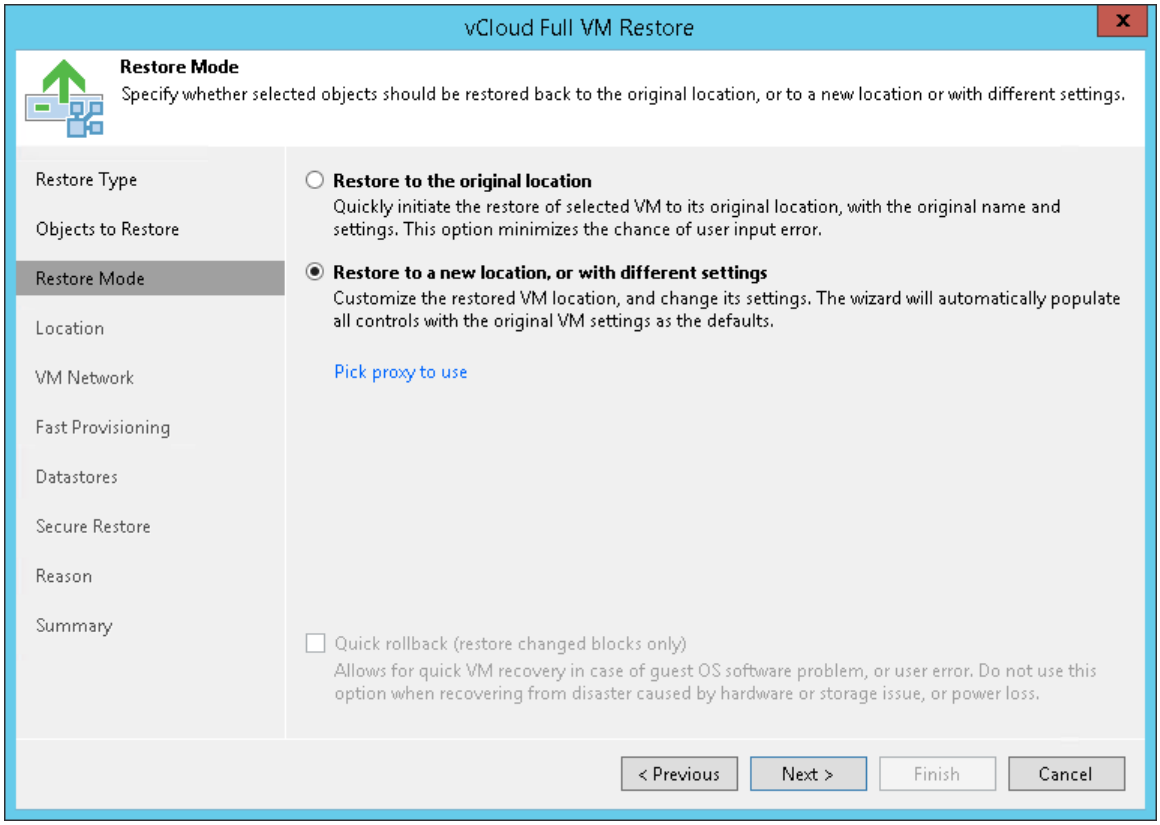
At the **Restore Mode** step of the wizard, choose the necessary restore mode and backup proxy for VM data transport:

1. Choose a restore mode:
   - Select **Restore to original location** if you want to restore the VMs with initial settings and to the original location. If this option is selected, you will immediately pass to the Summary step of the wizard.
   - Select **Restore to a new location, or with different settings** if you want to restore the VMs to a different location and/or with different settings (such as VM location, network settings, fast provisioning settings and so on). If this option is selected, the **vCloud Full VM Restore** wizard will include additional steps for customizing VM settings.
2. [For VM restore to the original location] Select the **Quick rollback** check box if you want to use incremental restore for the VMs. Veeam Backup & Replication will use CBT to get data blocks that are necessary to revert the VMs to an earlier point in time, and will restore only these data blocks from the backup. Quick rollback significantly reduces the restore time and has little impact on the production environment.

   It is recommended that you enable this option if you restore the VMs after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.
3. Click the **Pick proxy to use** link to select backup proxies over which VM data must be transported to the source datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the source datastore and will automatically assign optimal proxy resources for processing VM data.

     During the restore process, VMs are processed simultaneously. 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 writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VMs processing.

- If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that will be used for restore. It is recommended that you select at least two proxies to ensure that VMs are recovered should one of backup proxies fail or lose its connectivity to the source datastore during restore.



Step 5. Select VM Location

The **Location** step of the wizard is available if you have chosen to change the location and settings for the restored VMs.

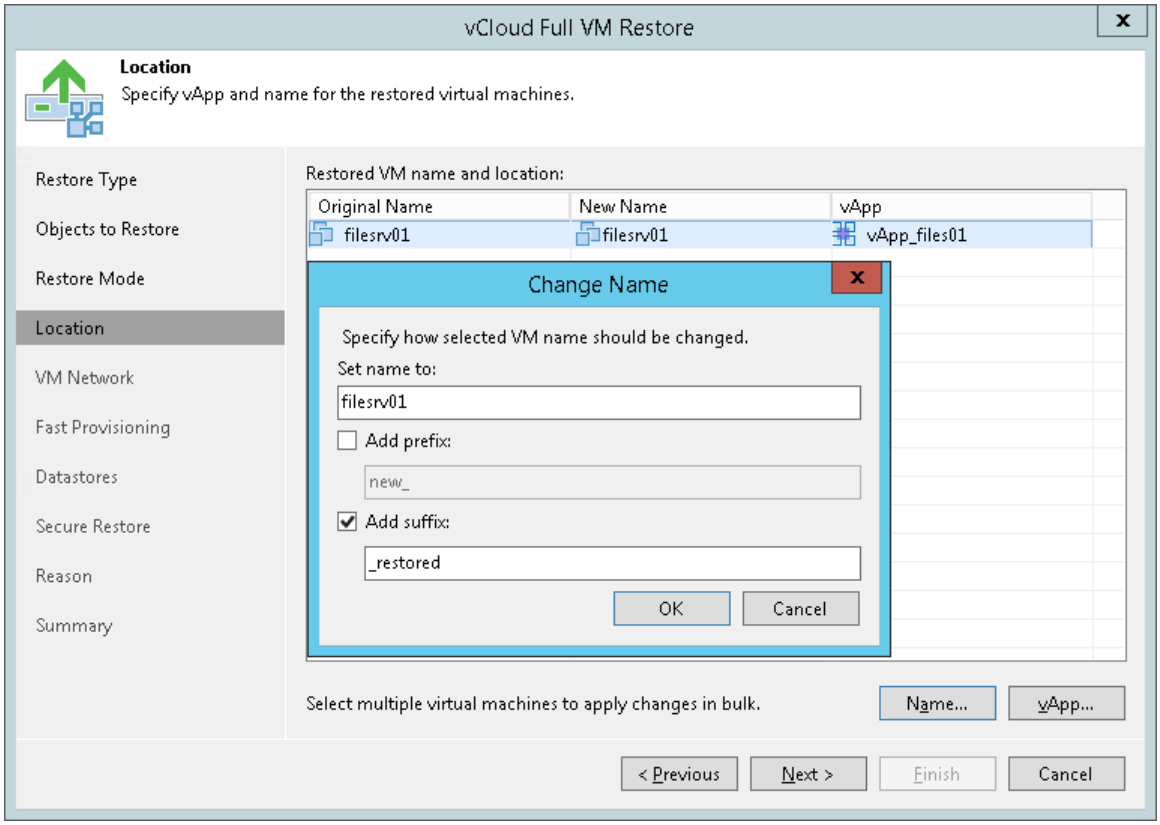
By default, Veeam Backup & Replication restores the VM to its original location. To restore the VM to a different location:

1. Select the VM in the list and click **vApp**.
2. From the vCloud Director hierarchy, choose a vApp in which the restored VM must be registered.

   To facilitate selection, use the search field at the bottom of the window: enter the vApp name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

To change the VM name:

1. Select a VM in the list and click **Name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule by adding a prefix and/or suffix to the original VM name.
3. You can also change VM names directly in the list: select a VM, click the **New Name** field and enter the name to be assigned to the recovered VM.



Step 6. Select Destination Network

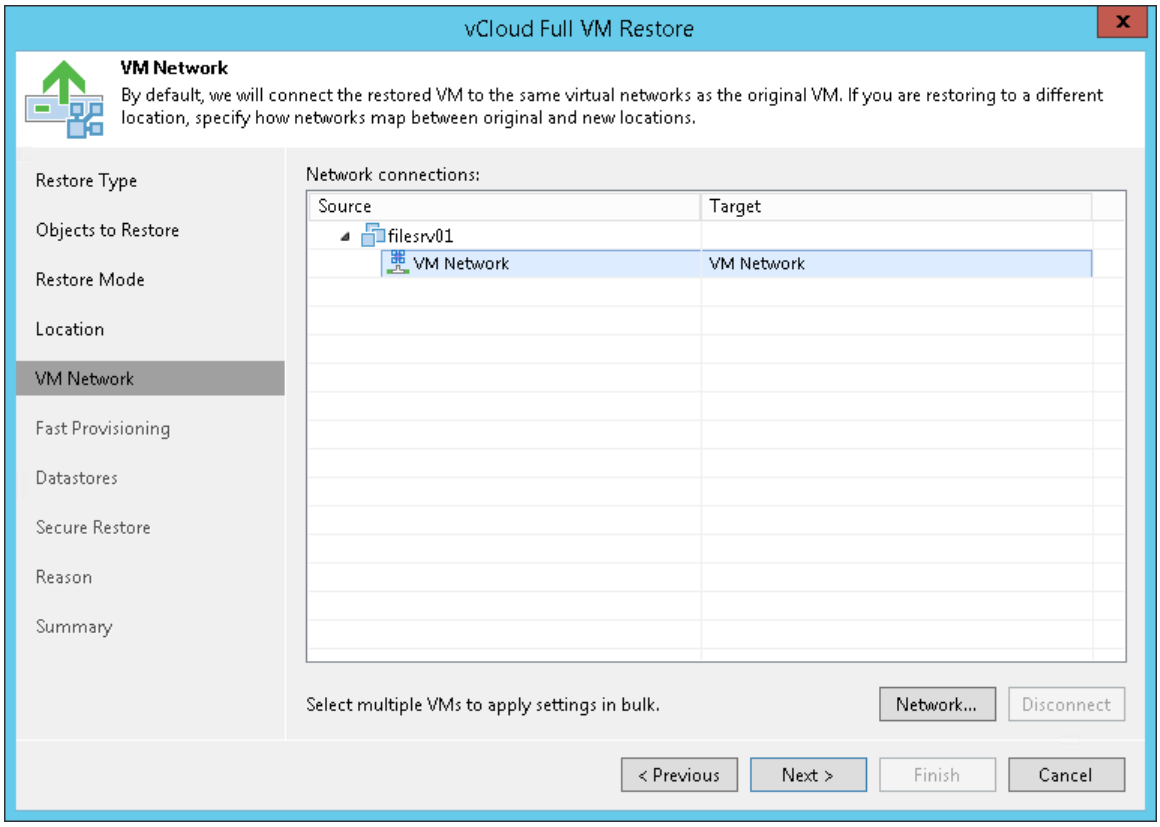
The **VM Network** step of the wizard is available if you have chosen to change the location and settings of the restored VMs.

To select networks to which the restored VM must be connected:

1. Select a VM in the list and click **Networks**.
2. The **Select Network** window displays all networks that are configured for the destination vApp. From the list of available networks, choose a network to which the restored VM must have access upon restore.

   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

To prevent the restored VM from accessing any network, select it in the list and click **Disconnected**.



Step 7. Select Template to Link

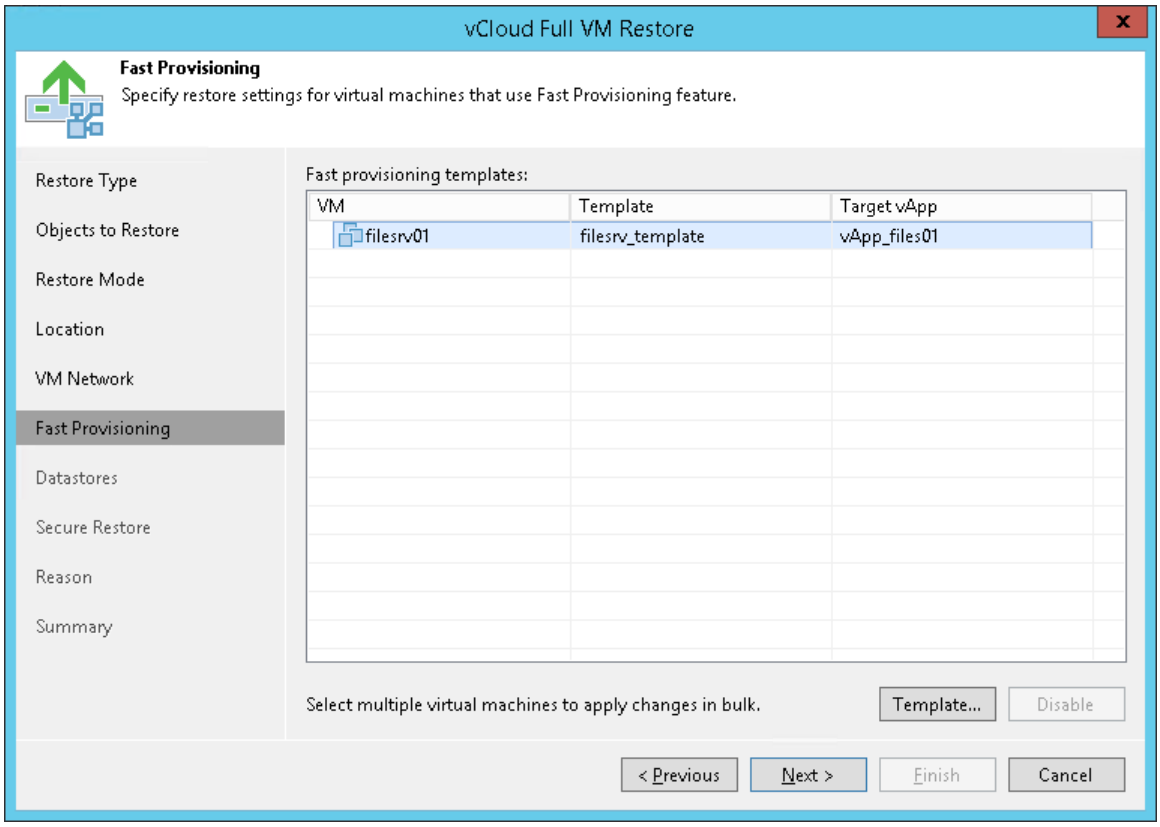
The **Fast Provisioning** step of the wizard is available if you have chosen to change the settings of the restored VMs.

To select a VM template:

1. Select a VM in the list and click **Set Template**.
2. From the vCloud Director hierarchy, choose a template to which the restored VM must be linked.

   To facilitate selection, use the search field at the bottom of the window: enter a VM template name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

If you want to disable fast provisioning for the VM and restore it as a regular VM, select the VM in the list and click **Disable**.



Step 8. Select Storage Policy and Datastores

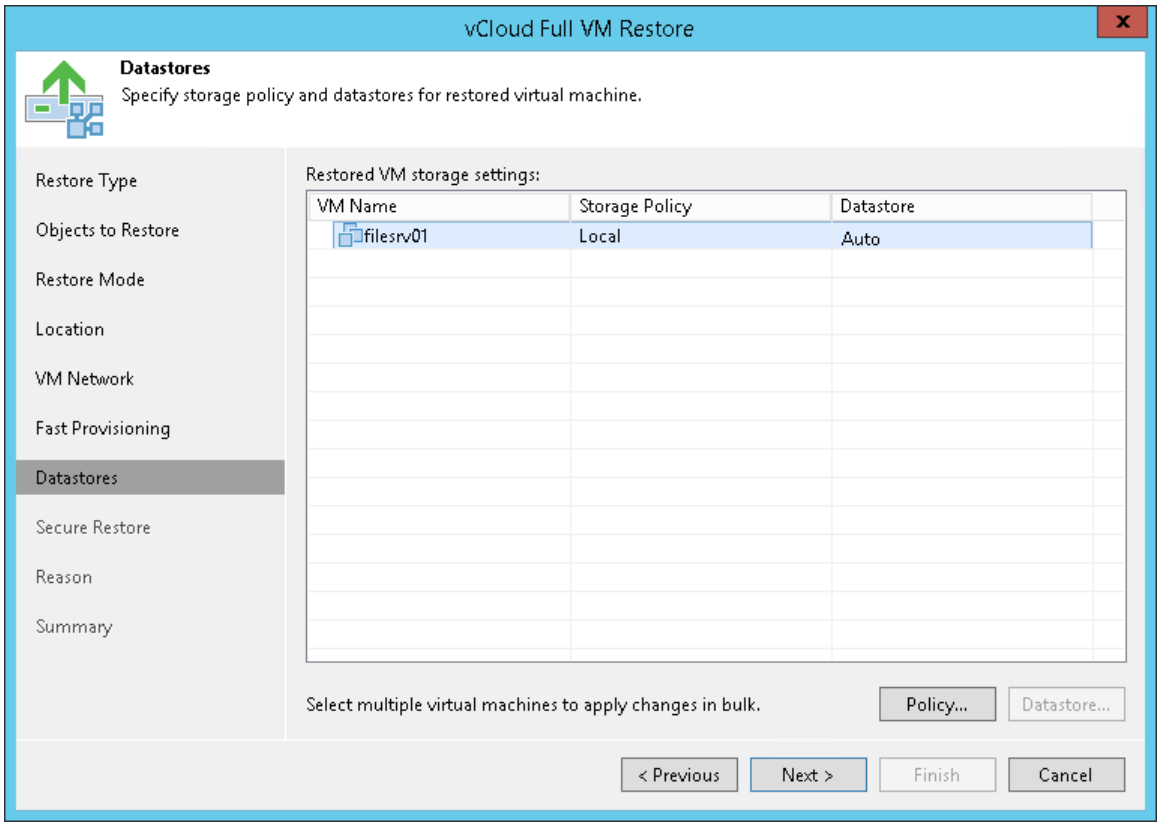
The **Datastores** step of the wizard is available if you have chosen to change the settings of the restored VMs.

To select a storage policy for the restored VM:

1. Select a VM in the list and click **Policy**.
2. In the displayed window, select the necessary policy for the VM.

If you have selected to disable fast provisioning at the previous step of the wizard, you must select a datastore on which disks of the restored VM will be placed.

1. Select a VM in the list and click **Datastore**.
2. In the displayed window, select the datastore on which the VM disks must be located.

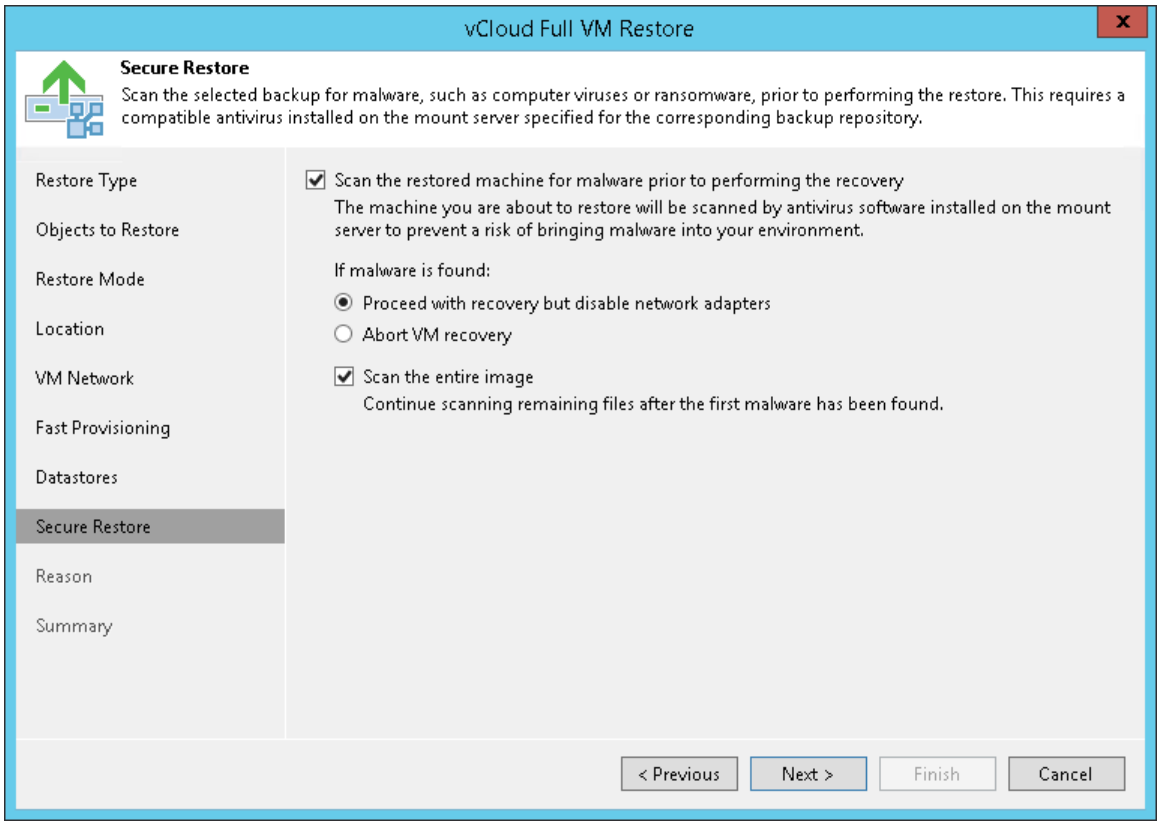


Step 9. Specify Secure Restore Settings

You can instruct Veeam Backup & Replication to perform secure restore — scan VM data with antivirus software before restoring the VM into a vApp. For more information on secure restore, see Secure Restore.

To specify secure restore settings:

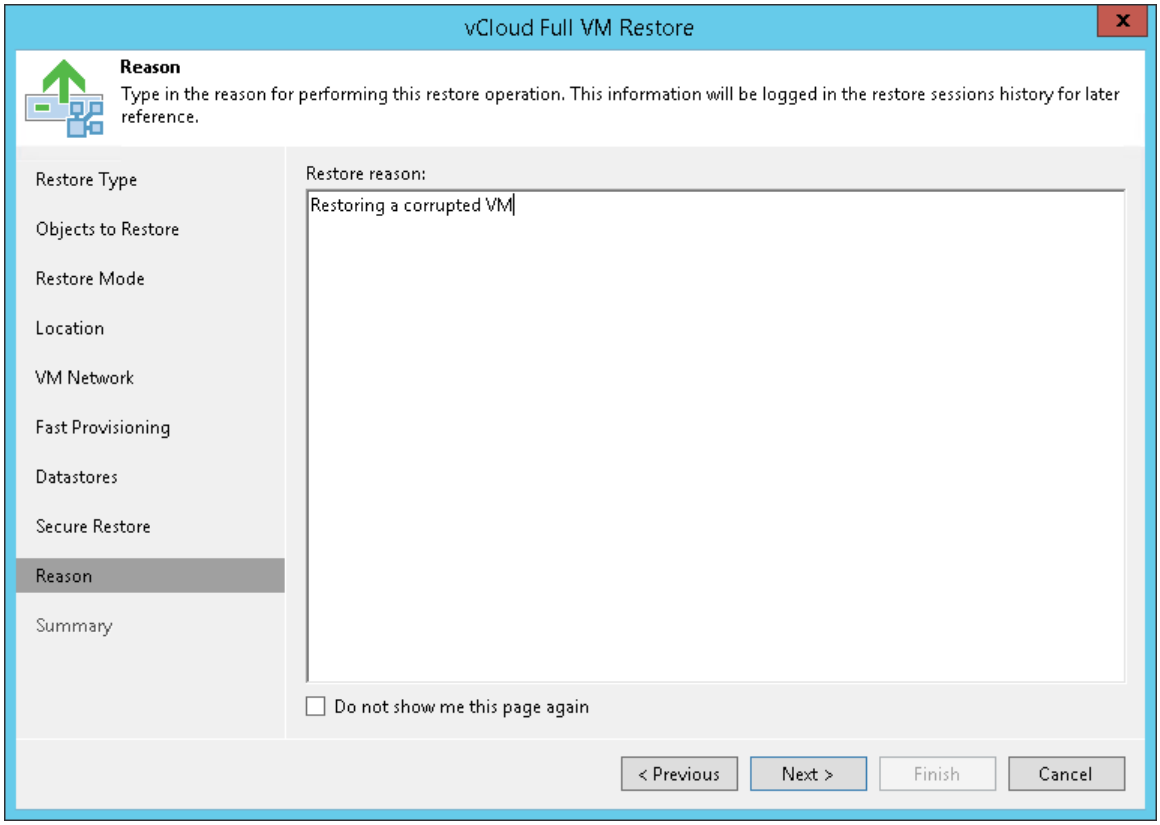
1. At the **Secure Restore** step of the wizard, select the **Scan the restored machine for malware prior to performing the recovery** check box.
2. Select which action Veeam Backup & Replication will take if the antivirus software finds a virus threat:
   - **Proceed with recovery but disable network adapters.** Select this action if you want to restore the VM with disabled network adapters (NICs).
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
3. Select the **Scan the entire image** check box if you want the antivirus software to continue the VM data scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.



Step 10. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the selected VMs. 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.

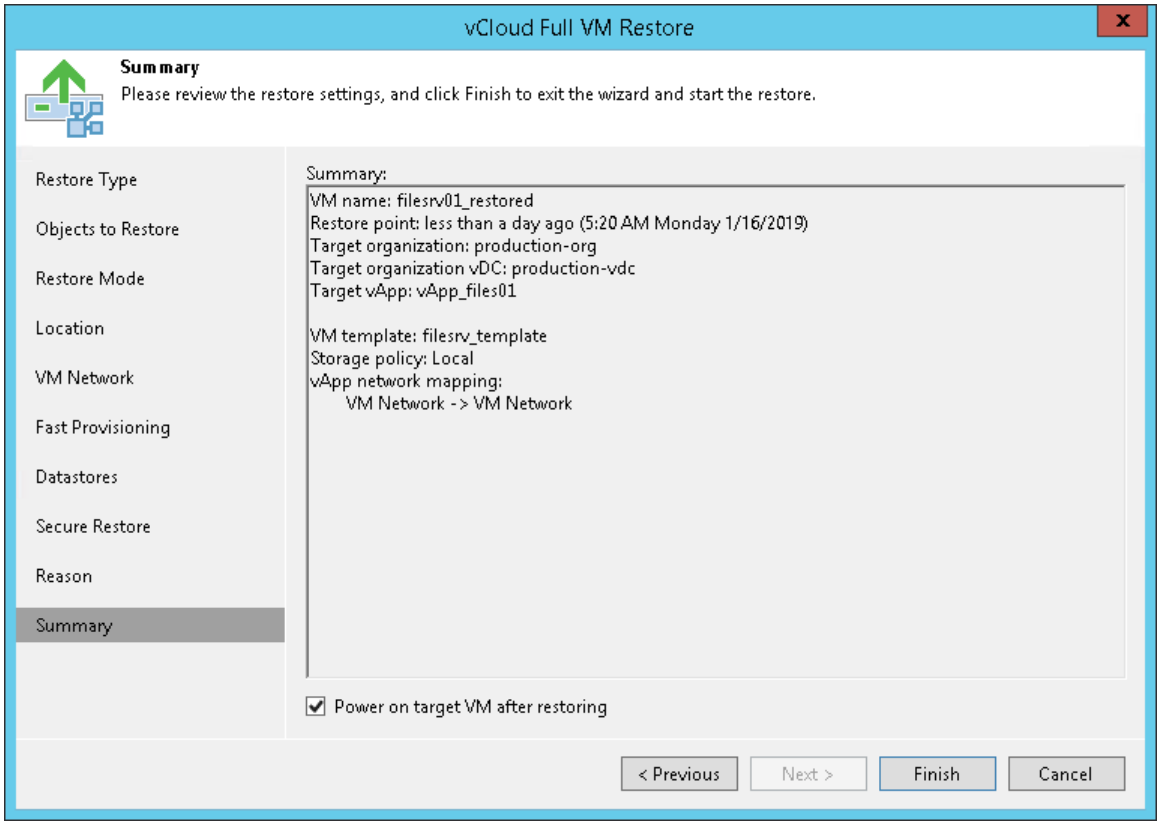


Step 11. Verify Recovery Settings and Finish Working with Wizard

At the **Summary** step of the wizard, specify additional settings for VMs restore:

1. If you want to start the restored VMs, select the **Power on VM after restoring** check box.
2. Check the settings for VMs restore and click **Finish**. Veeam Backup & Replication will recover the VMs in the specified destination.

> **NOTE**
>
> Veeam Backup & Replication checks the lease term for the restored VMs. In case the lease period has expired, the lease will be automatically updated.



Restoring Entire VMs into vSphere Infrastructure

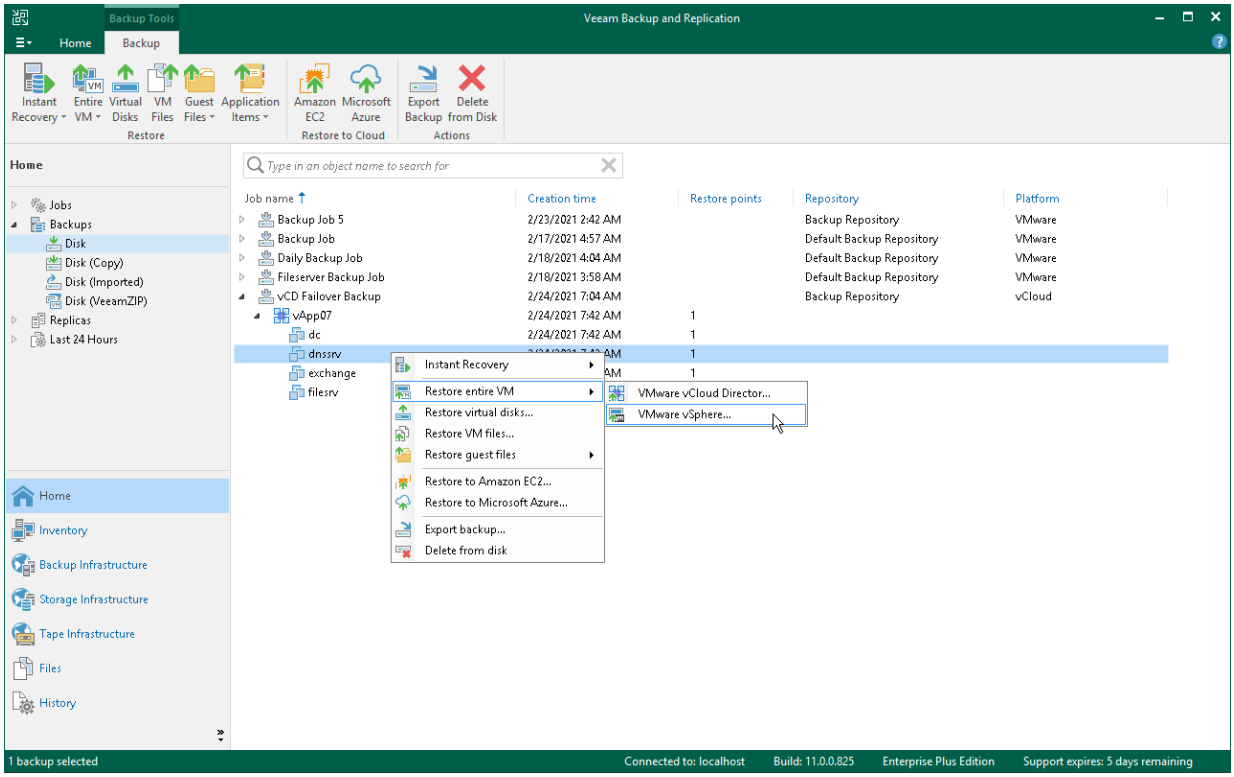
You can restore vCloud Director VMs from the backup to the VMware vSphere infrastructure.

During restore, Veeam Backup & Replication neglects the vApp metadata saved to the backup file and performs a regular entire VM restore process. The VM is restored to the vCenter Server or ESXi host and is not registered in vCloud Director. vCloud-specific features such as fast provisioning are not supported for such type of restore.

To launch the **Full VM Restore** wizard, do one of the following:

- Open the **Home** view, in the inventory pane select **Backups**. In the working area, expand the necessary backup, select the VMs you want to restore and click **Entire VM** > **VMware vSphere** on the ribbon.
- Open the **Inventory** view. In the inventory pane, expand the vCloud Director hierarchy and select the vCenter Server. In the working area, right-click the VM you want to restore and select **Restore entire VM** > **VMware vSphere**.

Entire VM restore of vCloud Director VMs does not differ from entire VM restore of regular VMware VMs. For more information, see Performing Entire VM Restore.



Restoring VM Files

The process of VM files restore for vCloud Director VMs does not differ from that for regular VMware VMs. For more information, see Restoring VM Files.

Restoring VM Hard Disks

The process of VM hard disks restore for vCloud Director VMs does not differ from that for regular VMware VMs. For more information, see Restoring Virtual Disks.

Restoring VM Guest OS Files

The process of VM guest files restore for vCloud Director VMs does not differ from that for regular VMware VMs. For more information, see Restoring VM Guest OS Files (FAT, NTFS or ReFS) and Restoring VM Guest OS Files (Multi-OS).

VMware Cloud on AWS Support

Veeam provides support for VMware Cloud on AWS. With Veeam Backup & Replication, you can administer backup, replication and restore operations in VMware Cloud on AWS environments.

Deployment

To perform data protection and disaster recovery tasks in VMware Cloud on AWS, consider the following recommendations and requirements on the backup infrastructure deployment:

- **Backup Server:** it is recommended to deploy Veeam backup server in the VMware Cloud on AWS environment. The machine must run Microsoft Windows.
- **Backup Proxy:** you must deploy a backup proxy in the VMware Cloud on AWS environment. The machine must run Microsoft Windows. You can assign the role of the backup proxy to a dedicated VM or to the backup server.

  To provide sufficient resources, deploy at least one backup proxy per each SDDC cluster in the VMware Cloud on AWS. This is required for VMware Cloud on AWS specific Hot-Add processing.
- **Backup Repository:** it is recommended to use a backup repository created outside of the VMware Cloud on AWS environment, for example, on the Amazon EC2 server. This type of deployment allows for efficient data transfer over the fast ENI connection used by VMware to communicate with AWS.

  Alternatively, you can store backups in a Veeam backup repository in the on-premises VMware environment or use Veeam Cloud Connect to transfer backups to the cloud. Note that in this scenario you may be charged additional fee for the traffic from VMware Cloud on AWS to the Internet.

To add VMware Cloud on AWS to the backup infrastructure, follow the same steps as described in the Adding VMware vSphere Servers section.

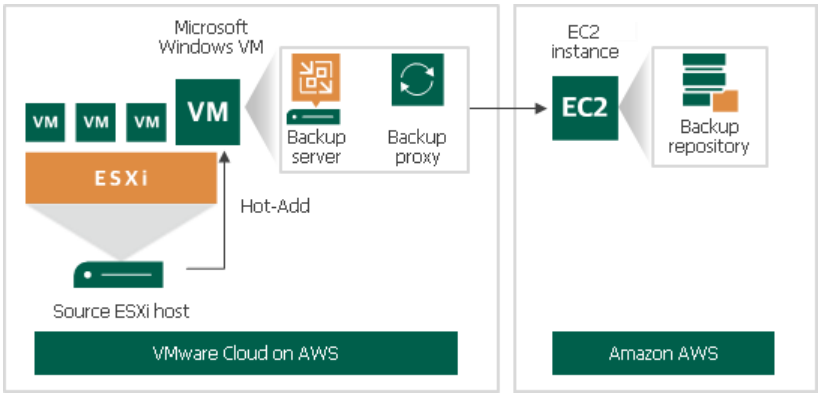
Simple Deployment

Simple deployment is preferable for VMware Cloud on AWS environments with low traffic load. Per this deployment type, you can install the backup server and the backup proxy on the same VM.

In a simple VMware Cloud on AWS deployment the backup infrastructure includes the following components:

- Source ESXi host
- Veeam backup server
- Veeam backup proxy

- Veeam backup repository: an EC2 instance in AWS

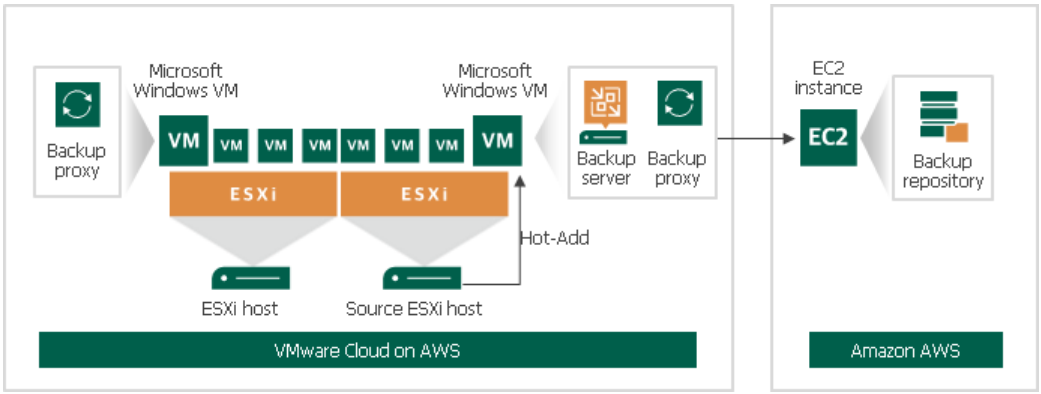


Advanced Deployment

Advanced deployment is intended for large-scale VMware Cloud on AWS environments with a large number of backup and replication jobs. Per this deployment type, it is recommended to install several backup proxies on dedicated VMs to move the workload from the backup server.

In an advanced VMware Cloud on AWS deployment the backup infrastructure includes the following components:

- Source ESXi host
- Veeam backup server
- Several Veeam backup proxies for better performance and workload distribution
- Veeam backup repository: an EC2 instance in AWS



To increase scalability and optimize performance in an advanced deployment, follow the recommendations below:

- Deploy additional backup proxies.
- Scale accordingly CPU and RAM resources of the EC2 instance used as a backup repository. Make sure it has enough free space for storing backups.

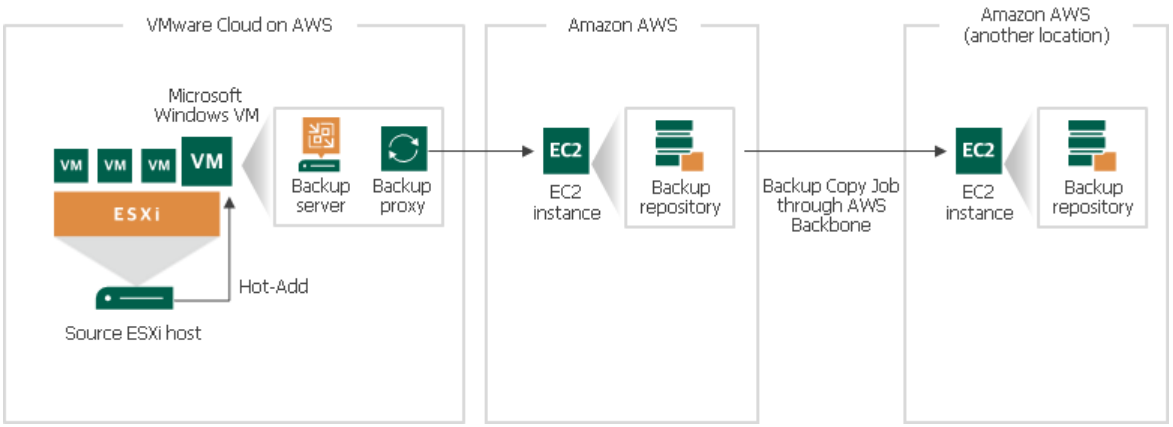
Deployment Scenarios for Off-Site Backup

To keep up with the 3-2-1 backup rule, it is recommended that you have a copy of your backups in an off-site location. To transfer your backups off-site, you can leverage Veeam backup copy.

Mind that transferring backups over the Internet may require paying additional fees. As a cost-effective alternative, you can store backups in a different AWS geographical location. In this case, backup copies are transferred through the AWS backbone. Using such AWS network solution provides data transfer at lower latency and cost when compared to the public Internet.

To perform backup copy to a different AWS location, the backup infrastructure must contain the following components:

- Source ESXi host
- Veeam backup server
- Veeam backup proxy
- Veeam backup repository: an EC2 instance in AWS
- Veeam backup repository for backup copy: an EC2 instance in another AWS location



TIP

As an off-site backup solution, you can copy backups to virtual tapes and store them in Amazon S3/Glacier cloud storage. In this case, AWS Storage Gateway performs the role of a Virtual Tape Library (VTL).

Deployment Scenarios for Capacity Tier

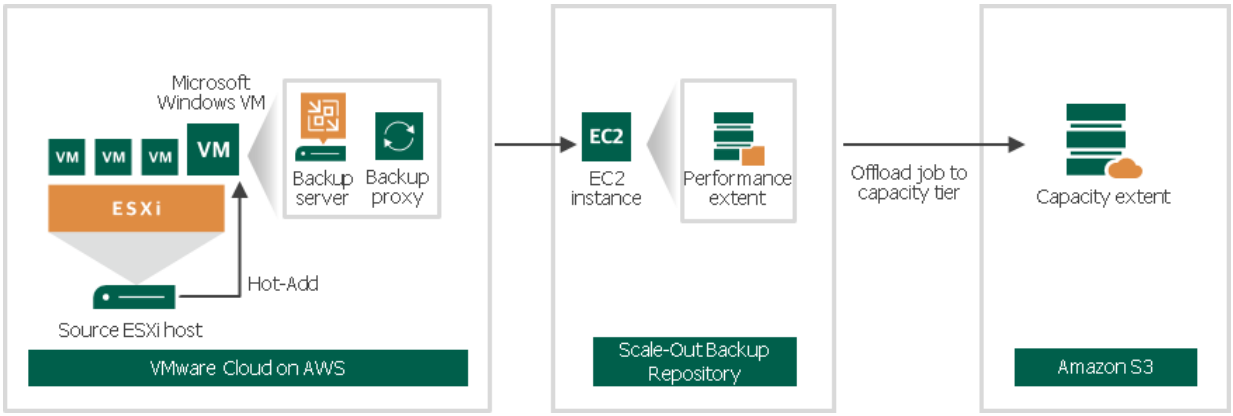
If you have a scale-out backup repository with a capacity tier option configured, you can transfer your backups to the capacity tier for long-term storage. To do it, you can leverage Veeam capacity tier copy mode.

Note that capacity tier is available only as part of scale-out backup repository. For more information on capacity tier, see Capacity Tier.

To transfer backup files to the capacity tier, the backup infrastructure must contain the following components:

- Source ESXi host
- Veeam backup server
- Veeam backup proxy
- Veeam backup repository: an EC2 instance in AWS

- A configured scale-out backup repository with an object storage added as a capacity extent



Considerations, Limitations and Troubleshooting

Some of VMware features and permissions are not granted by default at the start of VMware Cloud on AWS. For details, see this Veeam KB article.

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.

> **NOTE**
>
> WAN acceleration is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise or Enterprise Plus editions of Veeam Backup & Replication are required.

High Bandwidth Mode

Starting with Veeam Backup & Replication 10, Veeam WAN acceleration introduces the **High bandwidth mode**. We recommend to use it for WAN connections faster than 100 Mbps. If compared with the **Low bandwidth mode**, it does not leverage the global cache, but utilizes a faster compression method, optimized digests processing and an alternative deduplication mechanism. As a result, the new mode provides better performance and higher speed of data transfer.

Mind that to use the **High bandwidth mode**, you must enable this option for WAN accelerators at both sites of the data transfer: the source and the target. If the target WAN accelerator has the **High bandwidth mode** enabled, different source accelerators can parallelly interact with it in different modes, depending on the mode selected for each source WAN accelerator.

Global Data Deduplication

The goal of WAN acceleration is to send less data over the network. To reduce the amount of data going over WAN, Veeam Backup & Replication uses the global data deduplication mechanism.

1. When you first run a job to the remote location, Veeam Backup & Replication analyzes data blocks going over WAN.
2. With every new cycle of a job to the remote location, Veeam Backup & Replication uses the data redundancy algorithm to find duplicate data blocks in copied files. Veeam Backup & Replication analyzes data blocks in files on the source side and compares them with those that have been previously transferred over WAN. If an identical data block is found, Veeam Backup & Replication deduplicates it.

As a result, only unique data blocks go over WAN. Data blocks that have already been sent are not sent. This way, Veeam Backup & Replication eliminates transfer of redundant data over WAN.

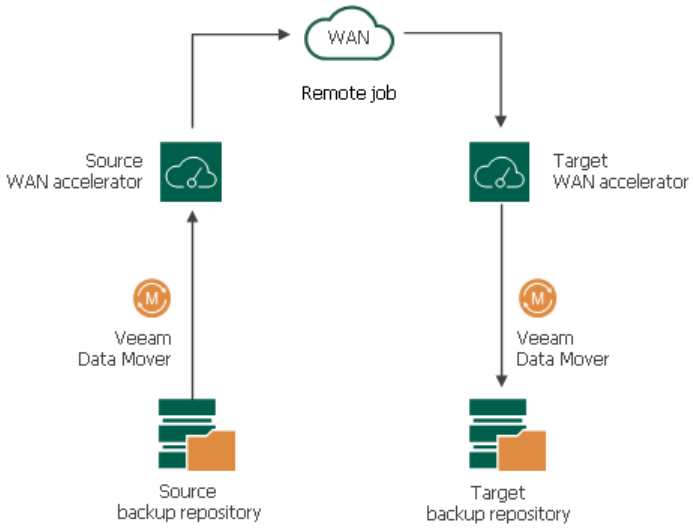
Veeam Backup & Replication uses three sources for data deduplication:

- VM disks. Veeam Backup & Replication analyses data blocks within the same VM disk. If identical blocks are found, duplicates are eliminated.
  For example, in case of a virtualized Microsoft Exchange server, the same email is typically stored in sender's Outbox folder and recipient's Inbox folder, which results in duplicate data blocks. When a job to the remote location runs, Veeam Backup & Replication detects such VM data blocks and performs deduplication.
- Previous restore points for the processed VM on the target repository. Veeam Backup & Replication analyses data in the restore point that is about to be copied and the restore points that are already stored on the target side. If an identical block is found on the target side, Veeam Backup & Replication eliminates the redundant data block in the copied restore point.
- Global cache. Veeam Backup & Replication creates a global cache holding data blocks that repeatedly go over WAN. In a new job session, Veeam Backup & Replication analyzes data blocks to be sent and compares them with data blocks stored in the global cache. If an identical data block is already available in the global cache, its duplicate on the source side is eliminated and not sent over WAN.

> **NOTE**
>
> Mind the following:
>
> - Veeam Backup & Replication deduplicates data blocks within one VM disk and in restore points for one VM only. Deduplication between VM disks and restore points of different VMs is performed indirectly, using the global cache. For more information, see WAN Global Cache.
> - Global data deduplication and deduplication within the same VM disk are not used if both WAN accelerators in the pair (the source one and the target one) operate in the **High bandwidth mode**.

WAN Accelerators

For WAN acceleration, Veeam Backup & Replication uses dedicated components — WAN accelerators. WAN accelerators are responsible for global data caching and data deduplication. Technically, WAN accelerators add a new layer in the backup infrastructure — between Veeam Data Movers on the source side and the Veeam Data Mover on the target side.



WAN Accelerators Deployment

To enable WAN acceleration and data deduplication technologies, you must deploy a pair of WAN accelerators in your backup infrastructure.

- One WAN accelerator is deployed on the source site, closer to the source backup repository or source host.
- The other WAN accelerator is deployed on the target site, closer to the target backup repository or target host.

On each WAN accelerator Veeam Backup & Replication creates the `VeeamWAN` folder containing the following data:

- The `VeeamWAN` folder on the source WAN accelerator stores files with digests required for deduplication. For more information, see How WAN Acceleration Works.
- The `VeeamWAN` folder on the target WAN accelerator stores global cache data.

> **NOTE**
>
> Global cache is not used if both WAN accelerators in the pair (the source one and the target one) operate in the **High bandwidth mode**.

To learn how to add a WAN accelerator to the Veeam Backup & Replication infrastructure, see Adding WAN Accelerators.

Recommendations for WAN Accelerators

You should not assign one source WAN accelerator to several jobs to remote locations that you plan to run simultaneously. The source WAN accelerator requires a lot of CPU and RAM resources, and does not process multiple tasks in parallel. As an alternative, you can create one job to the remote location for all VMs you plan to process over one source WAN accelerator.

The target WAN accelerator, however, can be assigned to several jobs to remote locations.

How WAN Acceleration Works

When you create a job to the remote location, you can select to use WAN acceleration in its properties.

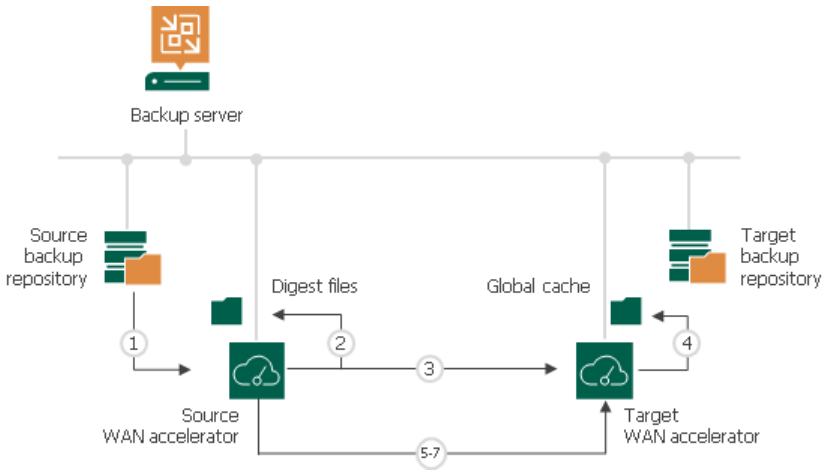
The procedure of data copying with WAN acceleration enabled is performed in the following way:

1. [For backup copy job] Veeam Backup & Replication decompresses the backup file to analyze its content.
2. The Veeam WAN Accelerator Service on the source WAN accelerator analyzes data blocks of the file to be transferred and creates a file with digests for these data blocks. The created file with digests is stored to the `VeeamWAN` folder on the source WAN accelerator.
3. Veeam Backup & Replication compresses the backup file data (for backup copy) or VM data (for replication) and copies it to the target side.

   At this point, Veeam Backup & Replication can perform deduplication within the VM itself — that is, deduplicate identical data blocks in every VM disk.
4. During the data transfer process, the Veeam WAN Accelerator Service on the target WAN accelerator populates the global cache storage with data blocks from the copied file.
5. During the next job cycle, the Veeam WAN Accelerator Service on the source WAN accelerator analyzes data blocks in the file that must be transferred this time and creates digests for these data blocks.
6. The Veeam WAN Accelerator Service compares the created digests with the digests that have been previously stored to the `VeeamWAN` folder on the source WAN accelerator. If duplicate data blocks are found, the actual data block in the backup file is not copied over WAN. Instead, it is taken from the global cache and written to the restore point in the backup copy folder or on the target data volume.
7. Additionally, Veeam Backup & Replication analyzes restore points that have been previously copied to the target side. If duplicates are found, Veeam Backup & Replication does not copy such blocks over WAN but takes them from the previously copied restore points.

As a result, Veeam Backup & Replication copies only new data blocks to the target side and uses data blocks that are already stored in the global cache or in restore points in the target backup repository.

If the target WAN accelerator is used by several jobs, the target backup repository may already contain data blocks of the necessary VM type. In this situation, Veeam Backup & Replication will copy the required data blocks to the global cache before the copying process starts and use these data blocks further on. For more information, see Many to One WAN Acceleration.



If WAN acceleration is performed in the **High bandwidth mode**, the procedure of data transfer with WAN acceleration has the following peculiarities:

- Global cache is not used. Thus, the target WAN accelerator does not need extra disk space to store the global cache folder.
- Deduplication is performed only by using previous restore points for the processed VM on the target repository. Therefore, Veeam Backup & Replication performs less deduplication operations and saves resources and time for data processing.
- The data chunk size used by the Changed Block Tracking mechanism during deduplication is smaller if compared with the **Low bandwidth mode**. This reduces the size of redundant data to transfer.

WAN Global Cache

From the technical point of view, the global cache is a folder on the target WAN accelerator. By default, global cache data is stored in the `VeeamWAN` folder on the disk with the most amount of space available. However, you can define any folder of your choice when you configure the target WAN accelerator.

> **NOTE**
>
> Global cache is not used if both WAN accelerators in the pair (the source one and the target one) operate in the **High bandwidth mode**.

By default, the size of the global cache is 100 GB. You can increase the size or decrease it if necessary. The more space you allocate, the more repeating data blocks will be written to the global cache and the more efficient WAN acceleration will be. It is recommended that you allocate at least 40 GB to the global cache storage.

The global cache size is specified per source WAN accelerator. That is, if you plan to use one target WAN accelerator with several source WAN accelerators, the specified amount of space will be allocated for every source WAN accelerator that will be working with the target WAN accelerator and the size of the global cache will increase proportionally. For more information, see WAN Accelerator Sizing.

The WAN global cache is a "library" that holds data blocks repeatedly going from the source side to the target side. The global cache is populated at the first cycle of a job to the remote location. The priority is given to data blocks of Windows-based OSes, other OSes like Linux/Unix, and Microsoft Exchange Server.

Veeam Backup & Replication constantly maintains the global cache in the actual state. To do that, it continuously monitors data blocks going over WAN and data blocks in the global cache.

- If some new data block is constantly sent over WAN, it is added to the global cache.
- If some data block in the global cache is not sent over WAN and are not re-used for some period of time, it is removed from the global cache to make room for new data blocks.

Veeam Backup & Replication also performs periodic consistency checks. If some data block in the global cache gets corrupted, Veeam Backup & Replication removes it from the global cache.

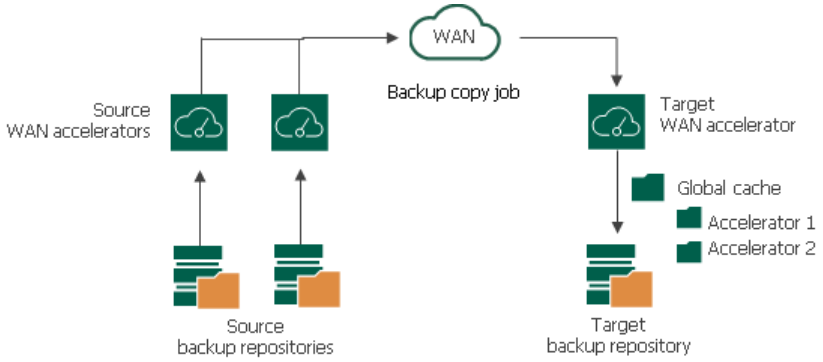
The efficiency of the WAN acceleration increases with every new backup copy interval in the backup copy job. During the first backup copy interval in the backup copy job, the WAN acceleration level is minimal. Veeam Backup & Replication populates the global cache. With every new job cycle, Veeam Backup & Replication updates the global cache to include the most "popular" data blocks and the WAN acceleration efficiency increases.

> **NOTE**
>
> You can populate the global cache before you run the job to the remote location for the first time. In this case, Veeam Backup & Replication will use the global cache starting from the first session of the job to the remote location, and the WAN traffic will be minimal. For more information, see Population of Global Cache.

Many to One WAN Acceleration

The WAN global cache can be used by several source WAN accelerators simultaneously. For example, if you have several remote/branch offices, you can configure several source WAN accelerators in remote sites and one target WAN accelerator in the head office.

In this case, the global cache will hold cache data for separate source WAN accelerators. The cache data for every source WAN accelerator will be stored in a dedicated subfolder in the global cache folder.



When one target WAN accelerator is used by several source WAN accelerators, Veeam Backup & Replication can copy data blocks between global cache of these WAN accelerators. This mechanism works if there are no matching backups of VMs in the target backup repository, but matching data is available in cache of other WAN accelerators.

For example, you have two backup copy jobs: *Job 1* and *Job* 2. The *Job 1* uses the source WAN accelerator *Source 1* and the target WAN accelerator *Target 3*. The *Job 2* uses the source WAN accelerator *Source 2* and the same target WAN accelerator *Target 3*. In the global cache folder, Veeam Backup & Replication will store data for 2 WAN accelerators: *Source 1* and *Source 2*.

- *Job* 1 processes a VM running Microsoft Windows Server 2008 R2, and it has been running for some time. In the global cache, there is already data for this type of OS.
- *Job 2* also processes a VM running Microsoft Windows Server 2008 R2. When you start *Job 2* for the first time, there is no data for this type of OS in the global cache for *Source 2* WAN accelerator. In such situation, Veeam Backup & Replication will copy the necessary data block from the *Source 1* cache to the *Source 2* cache and will not transport this data block over WAN.

> **NOTE**
>
> Beside using global cache of other WAN accelerator, Veeam Backup & Replication also utilizes backup files residing in the backup repository. For example, if the backup repository contains a backup file created with a backup job and the backup copy job starts copying a backup of a VM of the same type, Veeam Backup & Replication will populate global cache on the WAN accelerator from the backup file not to transfer redundant data over WAN.

Manual Population of Global Cache

You can manually pre-populate the global cache to avoid the situation where the cache remains empty. As a result, by the time a job to the remote location starts, the global cache will contain data blocks that can be used for data deduplication.

Manual population of the global cache can be helpful in the following scenarios:

- First run of a job to the remote location. When you run a first session of a job to the remote location, the global cache is empty, and the whole amount of VM data needs to be transferred over WAN. It is recommended that you manually populate the global cache before you start a job to the remote location for the first time.

- Global cache corruption. If the global cache gets corrupted for some reason, Veeam Backup & Replication needs to perform at least one session of the job to the remote location to replace corrupted data blocks with valid data blocks. In this situation, you can clean the global cache and manually populate it with valid data before the job to the remote location begins.

> **IMPORTANT**
>
> Veeam Backup & Replication does not use encrypted backups for manual global cache population.

Limitations for Manual Population of Global Cache

The manually performed global cache population task has the following limitations:

- Veeam Backup & Replication does not use encrypted backups for global cache population.
- Veeam Backup & Replication writes only data blocks for Windows-based OSes to the default cache. Data blocks for other OSes like Linux/Unix and application data blocks are not written to the cache.
- You can start the global cache population task for the target WAN accelerator that is not currently used by any job to the remote location.
- If the global cache population task is currently running, the corresponding target WAN accelerator is locked. You cannot start any job to the remote location by using this target WAN accelerator.
- [For global cache corruption scenario] You must clean the global cache before you populate it with valid data.
- [Veeam Cloud Connect] Veeam Backup & Replication does not use tenant backups to populate global cache on the service provider side.

How Manual Population of Global Cache Works

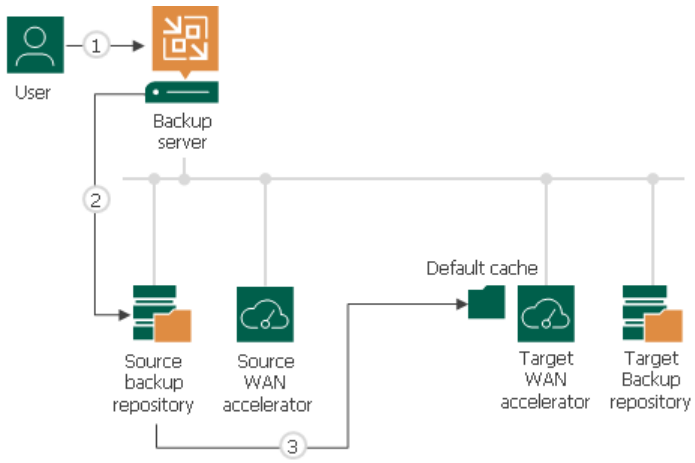
Global cache population is a manual operation performed by the user. When you run the global cache population task, Veeam Backup & Replication creates a 'default cache' on the target WAN accelerator. The default cache is used as a basic, universal cache for every new job to the remote location.

To populate the default cache, Veeam Backup & Replication uses backup files stored in backup repositories as a source of data.

The procedure of global cache population includes the following steps:

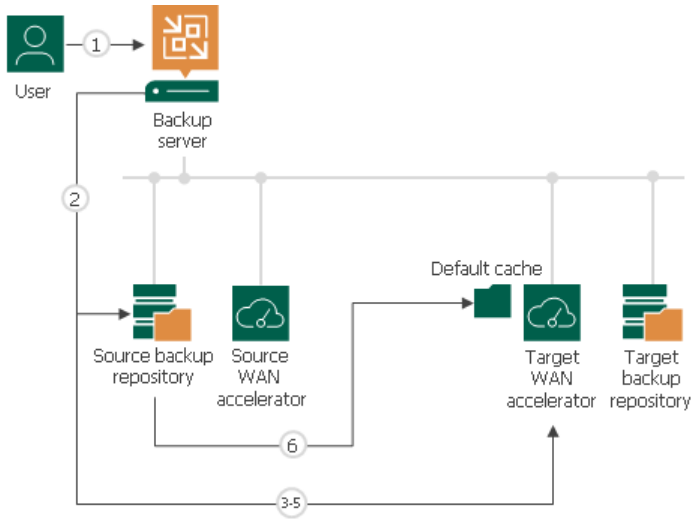
1. The user manually starts the global cache population tasks and selects backup repositories from which data blocks should be retrieved.
2. Veeam Backup & Replication scans backup repositories and makes up a list of OSes whose data blocks are available in backup files.
3. Veeam Backup & Replication copies data blocks from backup repositories and populates the default cache with these data blocks.

When a job to the remote location starts, Veeam Backup & Replication renames the default cache to a folder for the source WAN accelerator engaged in the job. As a result, Veeam Backup & Replication can use data blocks in this folder for deduplication starting from the very first session of the job to the remote location.



In many-to-one WAN accelerator deployment scenarios, the global cache may have folders for other source WAN accelerators, and these folders may contain data blocks for some OSes. If the global cache already contains some data, the procedure of global cache population includes the following steps:

1. The user manually starts the global cache population tasks and selects backup repositories from which data blocks should be retrieved.
2. Veeam Backup & Replication scans backup repositories and makes up a list of OSes whose data blocks are available in backup files.
3. Veeam Backup & Replication scans folders for other source WAN accelerators in the global cache and makes up a list of OSes whose data blocks are available there.
4. The list of OSes in the global cache is compared to the list of OSes in backup repositories. This way, Veeam Backup & Replication detects data blocks for which OSes are missing in the global cache.
5. In the global cache, Veeam Backup & Replication detects a folder with the maximum amount of data. This folder is used as a basis for the default cache.
6. Veeam Backup & Replication copies data blocks only for missing OSes from backup repositories and populates the default cache with these blocks. Data blocks for OSes available in folders for other source WAN accelerators are not copied to the default cache during the population task.
   Veeam Backup & Replication copies these data blocks on the fly, when a job to the remote location runs. For more information, see Many to One WAN Acceleration.



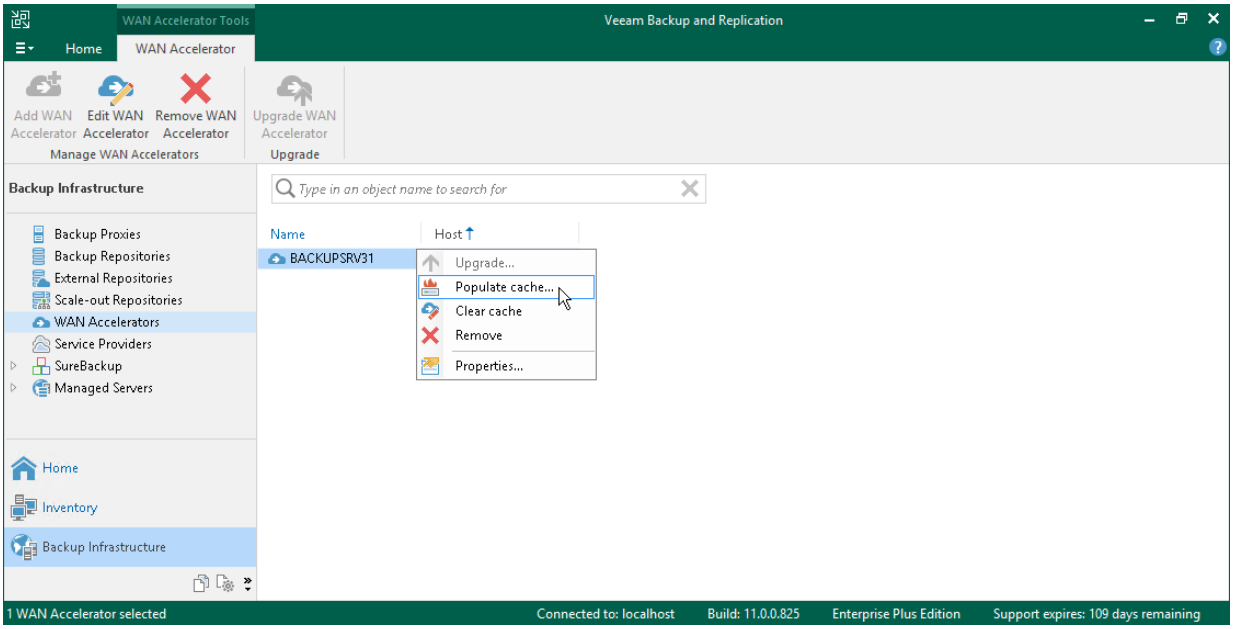
Manually Populating Global Cache

To manually populate the global cache:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **WAN Accelerators** node.
3. In the working area, right-click the target WAN accelerator and select **Populate cache**.

   If the selected WAN accelerator is not assigned as a target WAN accelerator to any job to the remote location, Veeam Backup & Replication will display a warning.
4. In the **Source Backup Repositories** window, select backup repositories from which OS data blocks must be retrieved.

   It is strongly recommended that you select backup repositories on the same site where the target WAN accelerator is located. In the opposite case, the traffic will travel between sites, which will increase load on the network.
5. Click **OK**.



Clearing Global Cache

You can clear the global cache on the target WAN accelerator. It is recommended that you clear the global cache in the following situations:

- Global cache is corrupted.
- Global cache contains data that is no longer needed. This situation may occur, for example, if you have decided to fully switch to the **High bandwidth mode** that does not use the global cache data.

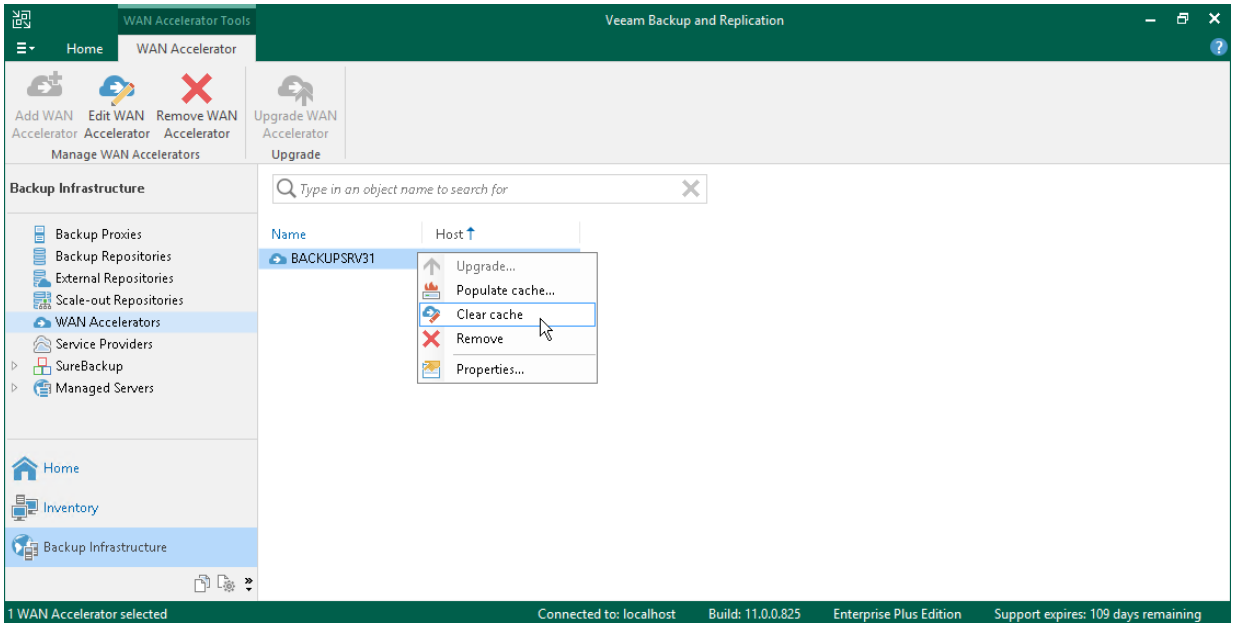
In such cases, it is recommended that you clear the global cache and populate it anew before you start jobs to remote locations processing new types of VMs.

To clear the global cache:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **WAN Accelerators**.

3. In the working area, right-click the target WAN accelerator and select **Clear cache**.

> **IMPORTANT**
>
> Mind the following:
>
> - Before you clear the global cache, make sure that you do not have any running jobs that use this target WAN accelerator. When the global cache is cleared, Veeam Backup & Replication will restart the Veeam WAN Accelerator Service, and running jobs will complete with the Failed status.
> - When you clear the global cache, you also clear all digest data stored at this WAN accelerator.



WAN Accelerator Sizing

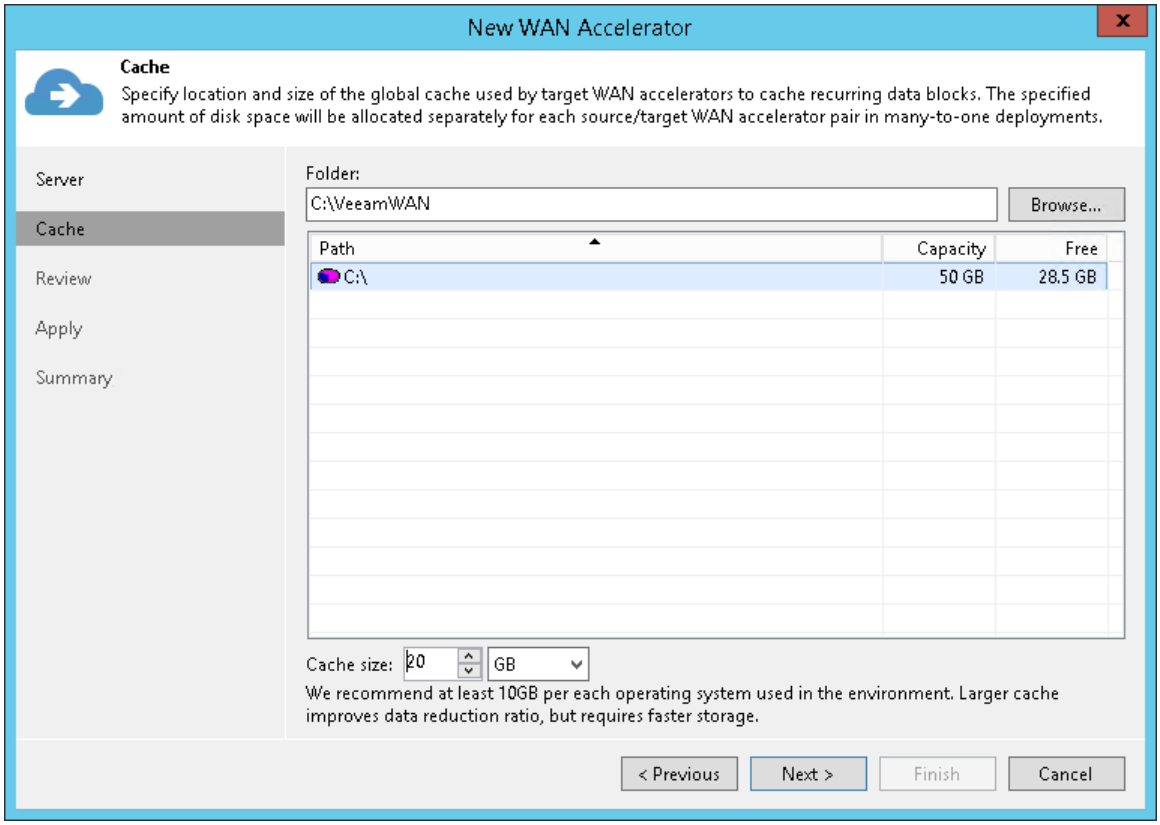
To ensure correct work of remote jobs over WAN accelerators, you must provide enough free space for service data on source and target WAN accelerators.

When configuring WAN accelerators, mind that there can be situations when WAN acceleration switches from the **High bandwidth mode** to the **Low bandwidth mode**: for example, the link to the remote location changes and you decide to disable the **High bandwidth mode** for one of the accelerators in the pair. If you disable the **High bandwidth mode** and start a job which utilizes this WAN accelerator, Veeam Backup & Replication deletes digest data that was used in the **High bandwidth mode** and recreates it for the **Low bandwidth mode**. Besides, Veeam Backup & Replication will also use the global cache at the target WAN accelerator.

To avoid problems caused by the lack of free space when switching from the **High bandwidth mode** to the **Low bandwidth mode**, we recommend that you configure WAN accelerators as if you planned to use them in the **Low bandwidth mode**.

Source WAN Accelerator

When you run a remote job over WAN accelerators, Veeam Backup & Replication analyses data blocks going to target and calculates digests for these data blocks. Digests data is stored on the source WAN accelerator, in the `VeeamWAN` folder on the disk that you select when you configure the WAN accelerator.



You must make sure that there is enough disk space on the source WAN accelerator to store digest data.

The amount of disk space required for a source WAN accelerator operating in the **Low bandwidth mode** is calculated by the following formula:

```
Digest Size = 2% of Provisioned VM Size
```

For example, if you plan to process 10 VMs whose provisioned size is 2 TB, you must allocate 40 GB of disk space for digest data on the source WAN accelerator operating in the **Low bandwidth mode**.

The amount of disk space required for a source WAN accelerator operating in the **High bandwidth mode** is calculated by the following formula:

```
Digest Size = 1% of Provisioned VM Size
```

For example, if you plan to process 10 VMs whose provisioned size is 2 TB, you must allocate 20 GB of disk space for digest data on the source WAN accelerator operating in the **High bandwidth mode**.

Target WAN Accelerator

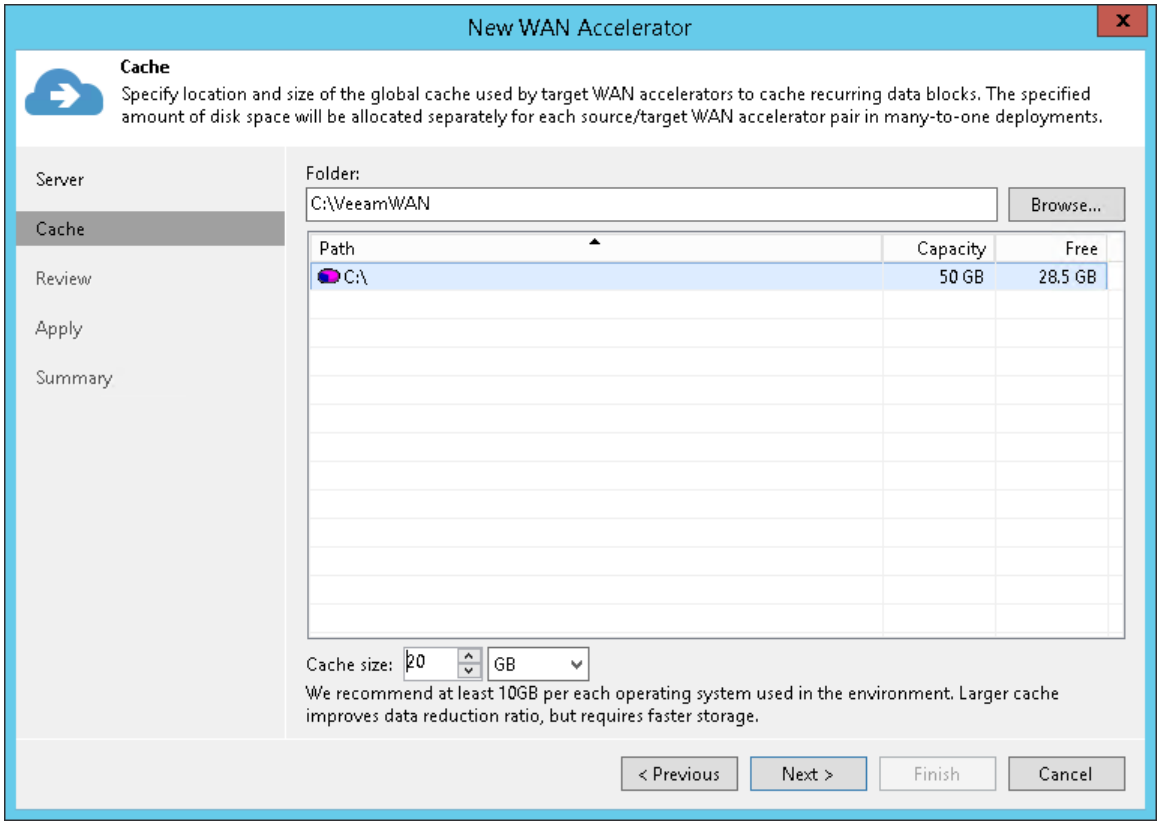
You must make sure that you provide enough free space for the following data on the target WAN accelerator:

- Global cache data
- Digest data

> **NOTE**
>
> For the target WAN accelerator operating in the **High bandwidth mode** only, you must provide enough free space to generate digest data. Global cache data is not used in the **High bandwidth mode**.
>
> When you enable the **High bandwidth mode** for an existing WAN accelerator, Veeam Backup & Replication does not automatically remove the global cache that was previously used for acceleration. If you are planning to use the **High bandwidth mode** for WAN acceleration and you do not need the global cache anymore, you can free the disk space by manually removing the cache. If you are planning to use WAN acceleration in the **Low bandwidth mode** in the future, we recommend that you keep the global cache. You can disable the **High bandwidth mode** and switch back to the **Low bandwidth mode** at any time.

Global Cache Data

Global cache is stored on the target WAN accelerator, in the `VeeamWAN` folder on the disk that you select when you configure the WAN accelerator. The size of global cache is specified in the properties of the target WAN accelerator.



You must provide enough free space for global cache data. It is recommended that you provide 10 GB per every type of OS on VMs that you plan to process. By default, Veeam Backup & Replication allocates 100 GB for the global cache size.

For example, you want to process the following VMs:

- 1 VM that runs Microsoft Windows 7
- 3 VMs that run Microsoft Windows Server 2008 R2
- 2 VMs that run Microsoft Windows Server 2012 R2

There are 3 types of OSes so you must allocate 30 GB for the global cache on the target WAN accelerator.

> **NOTE**
>
> Global cache is stored only on the target WAN accelerator. You do not have to provide space for global cache on the source WAN accelerator.

Digest Data

In some cases, Veeam Backup & Replication may require more space on the target WAN accelerator than specified in the WAN accelerator properties. This can happen if digest data on the source WAN accelerator is missing or cannot be used. For example:

- You have performed the **Clear Cache** operation on the source WAN accelerator and it no longer contains digest data. For more information, see Clearing Global Cache.
- Veeam Backup & Replication has attempted to resume operation of backup data transfer, but the backup file was not prepared for the operation in a proper way. The digest data must be recalculated.

In such situations, the target WAN accelerator calculates digest data on its own, which requires additional space. After the digest data is calculated, the target WAN accelerator transfers it to the source WAN accelerator. After the transfer, the copy of the digest data is removed from the target WAN accelerator.

For safety reasons, it is recommended that you provide the following amount of space for digest data on the target WAN accelerator:

The amount of disk space required for digest data at a target WAN accelerator operating in the **Low bandwidth mode** is calculated by the following formula:

```
Digest Size = 2% of Provisioned VM Size
```

The amount of disk space required for digest data at a target WAN accelerator operating in the **High bandwidth mode** is calculated by the following formula:

```
Digest Size = 1% of Provisioned VM Size
```

This amount of space is required for digest data recalculation. If you do not provide this amount of space and a situation where Veeam Backup & Replication needs to recalculate digest data occurs, the job to the remote location will work in the limited mode. Veeam Backup & Replication will not deduplicate data against the previous restore points copied to target. For more information, see Global Data Deduplication.

> **IMPORTANT**
>
> When you specify the global cache size for a target WAN accelerator, you do not allocate any space for storing digest data. To let Veeam Backup & Replication recalculate digest data, you must make sure that necessary amount of free space is available on the target WAN accelerator (in addition to the space allocated for the global cache).

For example:

- You have allocated 100 GB for global cache on the target WAN accelerator operating in the **Low bandwidth mode**.
- Provisioned size of VMs to be processed is 2 TB.

In this case, the needed amount of free disk space for the global cache on the target WAN accelerator is:

```
100 GB + 40 GB = 140 GB
```

Many-to-One WAN Acceleration Scenario

Global cache size is calculated per 1 source WAN accelerator working with the target WAN accelerator. If you plan to use several source WAN accelerators with 1 target WAN accelerator, you must increase the size of the global cache proportionally. The cache data for every source WAN accelerator will be stored in a dedicated subfolder in the global cache folder of the target WAN accelerator. The global cache size is calculated by the following formula:

```
Total Global Cache Size = (# of Source WAN Accelerators) * (Size of Global Cach
e Configured in Target WAN Accelerator Properties)
Total Free Disk Space to Provide = Total Global Cache Size + Digest Size
```

For example:

- You have 4 source WAN accelerators in the source side working with 1 target WAN accelerator in the disaster recovery (DR) site.
- The global cache size configured in properties of the target WAN accelerator is 100 GB.
- The size of VMs to be processed is 2 TB.

In this case, the needed amount of free disk space for the global cache and digests on the target WAN accelerator is:

```
4*100 GB + 40 GB = 440 GB
```

NOTE

For more information and recommendations on WAN accelerator cache sizing, see the Veeam KB1877 article.

Adding WAN Accelerators

To add a WAN accelerator, you must assign the WAN accelerator role to a Microsoft Windows server added to the backup infrastructure.

You must deploy a pair of WAN accelerators: one WAN accelerator on each side of the WAN link.

Before you add a WAN accelerator, check prerequisites. Then use the **New WAN Accelerator** wizard to add a WAN accelerator.

Before You Begin

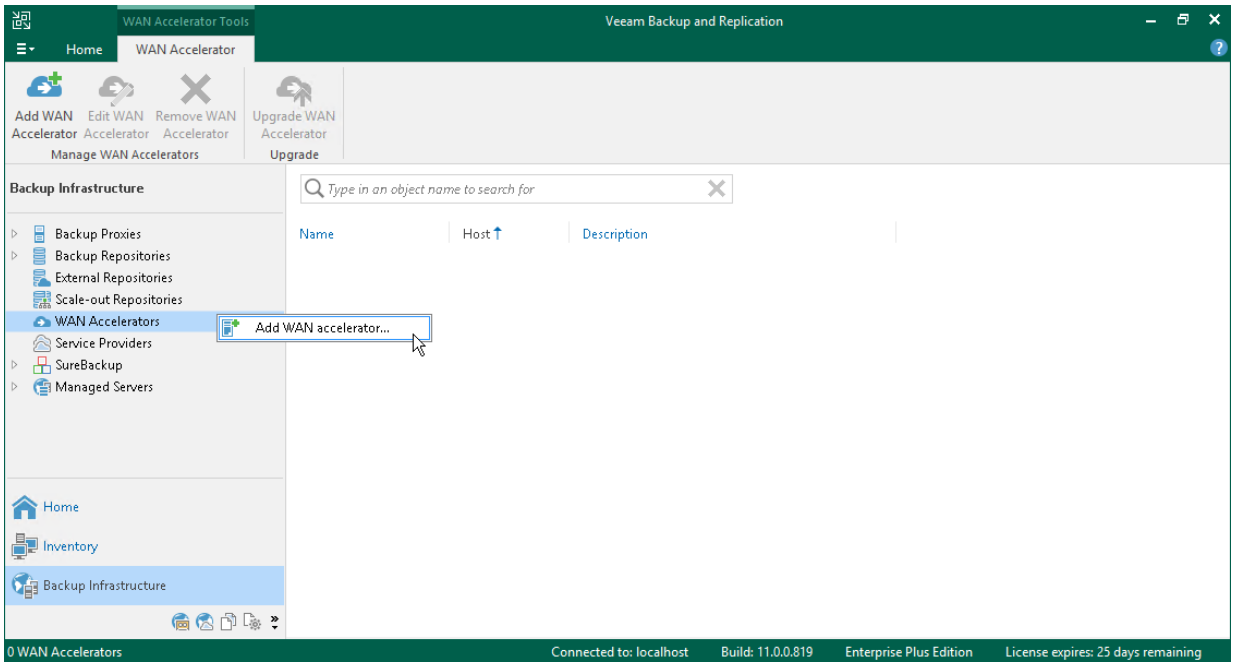
Before you add a WAN accelerator, check the following prerequisites:

- The machine that will operate as a WAN accelerator must meet the system requirements. For more information, see System Requirements.
- You must assign the WAN accelerator role to a Microsoft Windows server (physical or virtual). The WAN accelerator role can be assigned to backup proxies and Microsoft Windows backup repositories already configured in the backup infrastructure.
- You must use 64-bit Microsoft Windows machines as WAN accelerators. Veeam Backup & Replication does not support 32-bit versions of Microsoft Windows used as WAN accelerators.
- WAN acceleration operations are resource-consuming. When assigning the WAN accelerator role, mind available CPU and memory resources of the Microsoft Windows server. It is recommended that you assign the WAN accelerator role to servers with 8 GB RAM and more.
- The machine must have enough free disk space to store digests or global cache data. For more information, see WAN Accelerator Sizing.
- You must add the machine to the Veeam Backup & Replication console as a managed server before adding it as a WAN accelerator.

Step 1. Launch New WAN Accelerator Wizard

To launch the **New WAN Accelerator** wizard, do either of the following:

- Open the **Backup Infrastructure** view, in the inventory pane select **WAN Accelerators** and click **Add WAN Accelerator** on the ribbon.
- Open the **Backup Infrastructure** view, in the inventory pane right-click **WAN Accelerators** and select **Add WAN Accelerator**.



Step 2. Choose Server

At the **Server** step of the wizard, select a Microsoft Windows server that you plan to use as a WAN accelerator and define port and connection settings for this server.

1. From the **Choose server** list, select a Microsoft Windows server 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 Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.

2. In the **Description** field, provide a description for future reference.

   It is recommended that you describe the added WAN accelerator as the source or target one. When you create a job to the remote location, this hint will be displayed in brackets next to the WAN accelerator name, which will help you choose the necessary WAN accelerator to be used on the source or target side.

3. In the **Traffic port** field, specify the number of the port over which WAN accelerators must communicate with each other. By default, port 6165 is used.

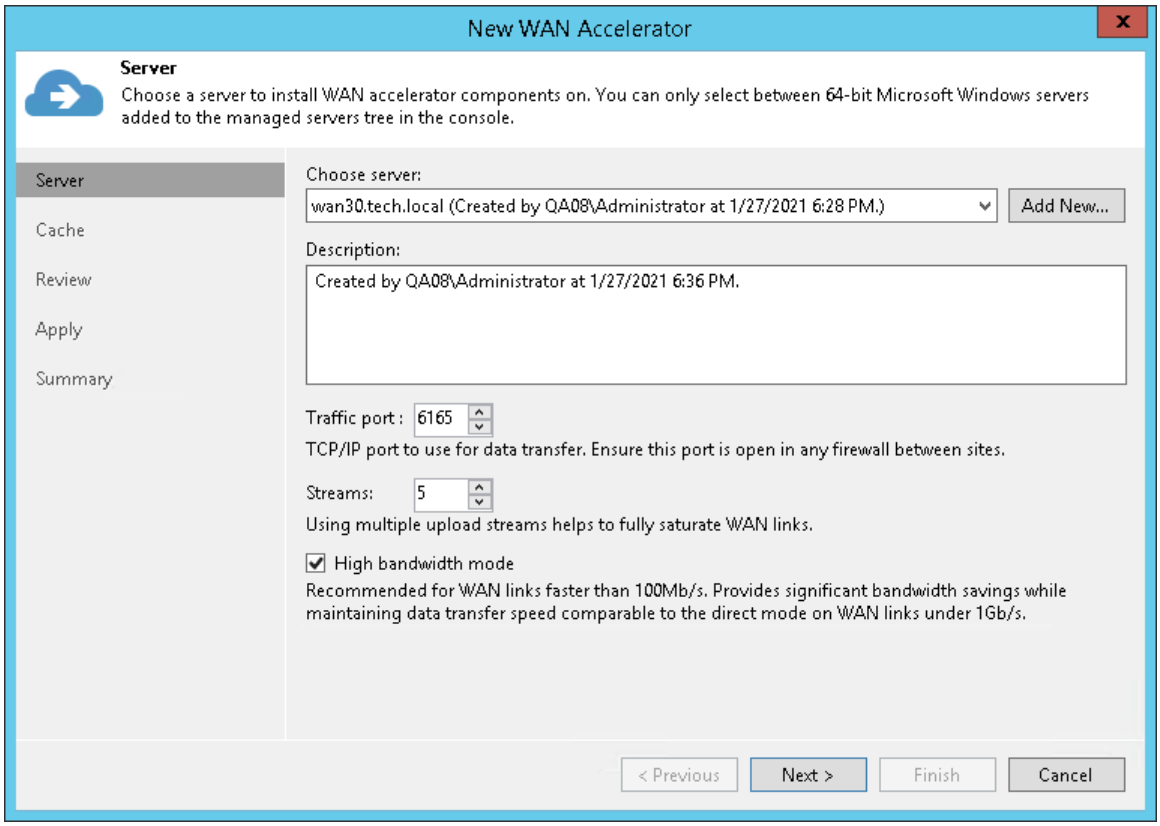
4. In the **Streams** field, specify the number of connections that must be used to transmit data between WAN accelerators. By default, 5 connections are used.

   This setting applies only to the source WAN accelerator. The greater is the number of streams, the more bandwidth resources Veeam Backup & Replication will use. A great number of streams engage more CPU and memory resources of the source WAN accelerator.

5. If your network bandwidth is more than 100 Mbps, we recommend that you use the **High bandwidth mode** option. This mode provides significant bandwidth savings comparable to the direct mode on WAN links under 1 Gbps.

   To use the **High bandwidth mode**, enable the option for WAN accelerators at both sites of the data transfer: the source one and the target one. If the **High bandwidth mode** is enabled for the target WAN accelerator, different source accelerators can parallelly interact with it in different modes, depending on the mode selected for each source WAN accelerator.

When you enable the **High bandwidth mode** for an existing pair of WAN accelerators, Veeam Backup & Replication does not automatically remove the global cache that was previously used for acceleration. If you are planning to use only the **High bandwidth mode** for WAN acceleration and you do not need the global cache anymore, you can free the disk space by manually removing the cache. If you are planning to use WAN acceleration in the **Low bandwidth mode** in the future, we recommend that you keep the global cache. You can disable the **High bandwidth mode** at any time.



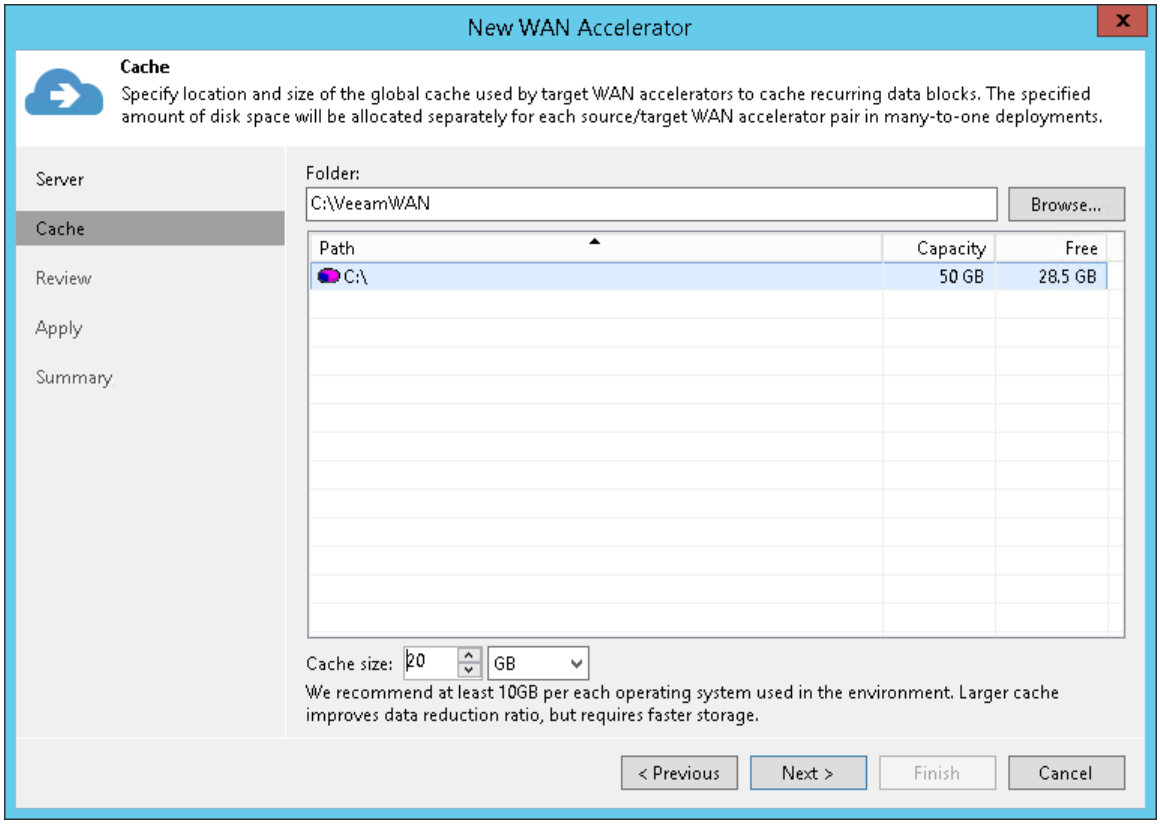
Step 3. Define Cache Location and Size

At the **Cache** step of the wizard, define settings for the folder where service files and global cache data will be stored.

> **NOTE**
>
> If both WAN accelerators (the source one and the target one) work in the **High bandwidth mode**, WAN acceleration does not use the global cache. But bear in mind that you can disable the **High bandwidth mode** and switch back to the **Low bandwidth mode** at any time, which will require to use the global cache. To provide correct operation of WAN accelerators, we recommend that you allocate enough disk space for the global cache folder when adding a new WAN accelerator, no matter if the **High bandwidth mode** is enabled for it or not.

1. In the **Folder** field, specify a path to the folder in which service files (for source and target WAN accelerators) and global cache data (for target WAN accelerator) must be stored. When selecting a folder on the target WAN accelerator, make sure that there is enough space for storing global cache data.
2. [For target WAN accelerator] In the **Cache size** field, specify the size for the global cache. The global cache size is specified per source WAN accelerator. If you plan to use one target WAN accelerator with several source WAN accelerators, the specified amount of space will be allocated to every source WAN accelerator and the size of the global cache will increase proportionally. For more information, see WAN Accelerator Sizing.

> **IMPORTANT**
>
> Do not nest the global cache folder deep in the file tree. During WAN acceleration operations, Veeam Backup & Replication generates service files with long file names. Placing such files to a folder of significant depth may cause problems on the NTFS file system.



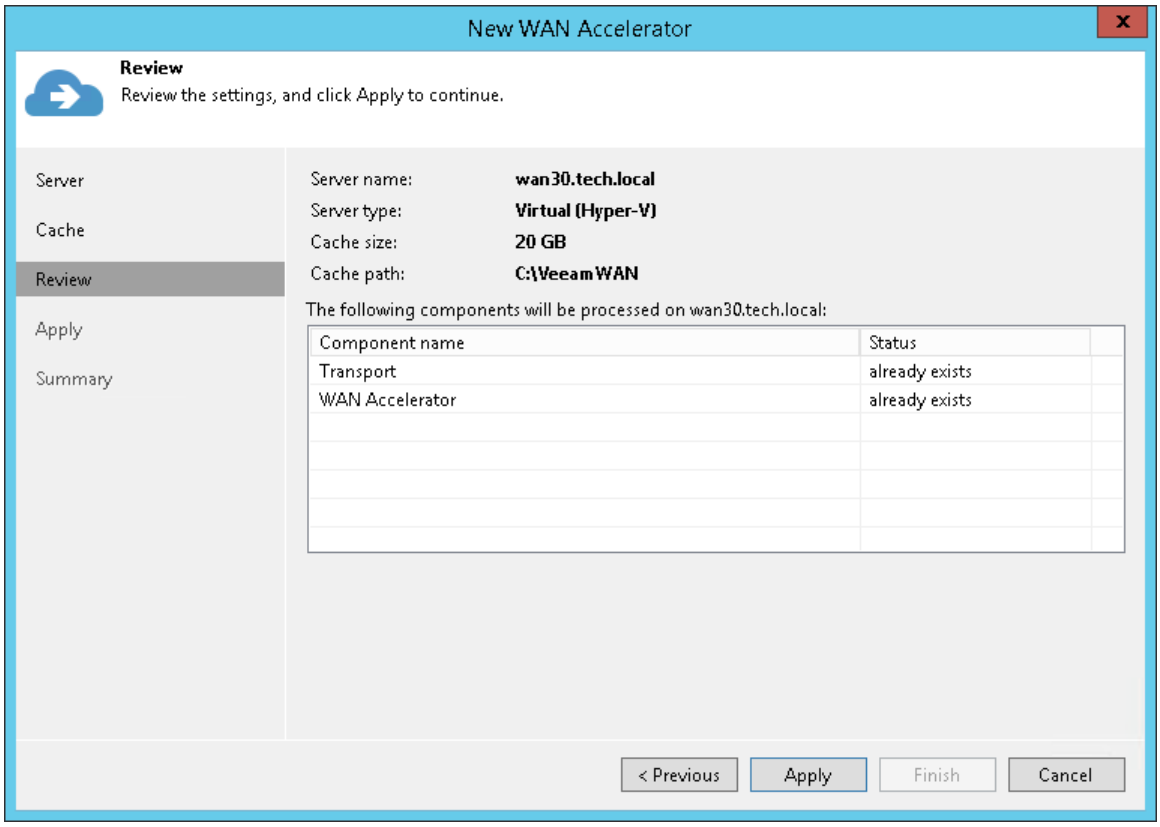
Step 4. Review Components

Veeam Backup & Replication installs the following components on the server used as a WAN accelerator:

- Veeam Data Mover
- Veeam WAN Accelerator Service

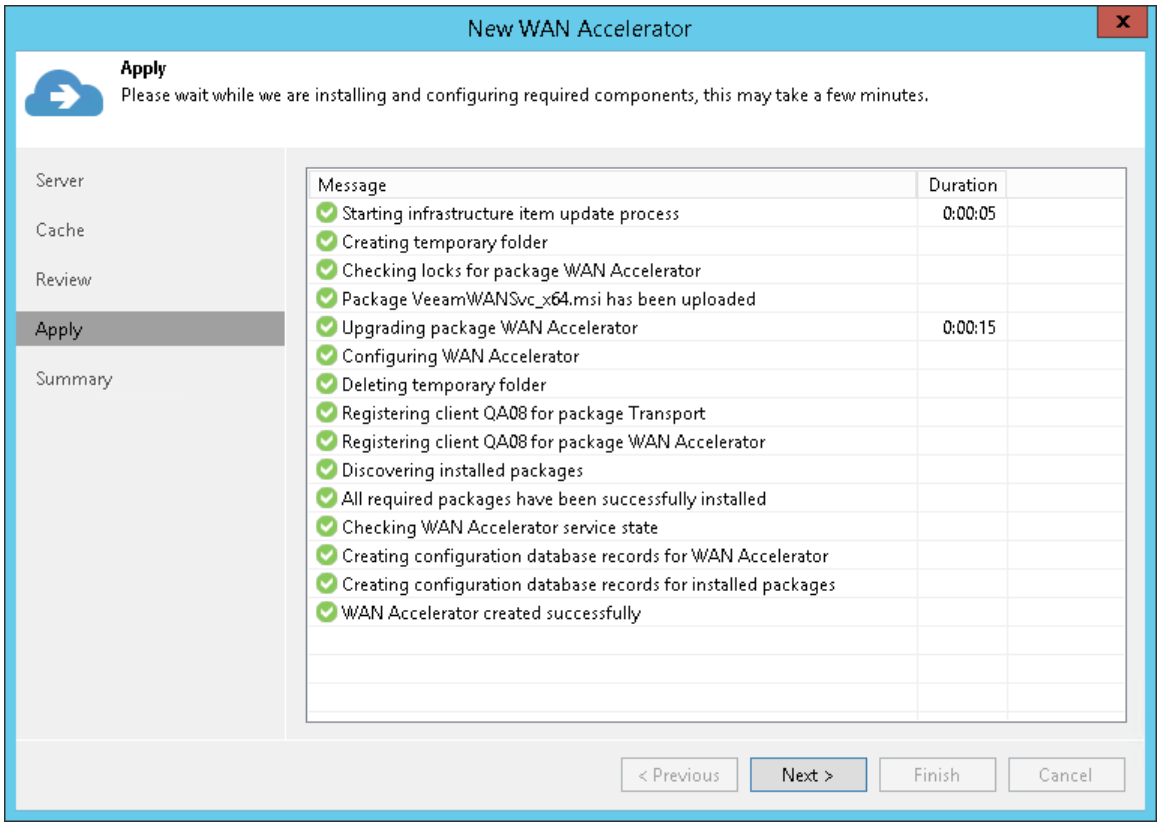
At the **Review** step of the wizard, review what components are already installed on the server and what components will be installed.

1. Review the components.
2. Click **Next** to install the components on the server.



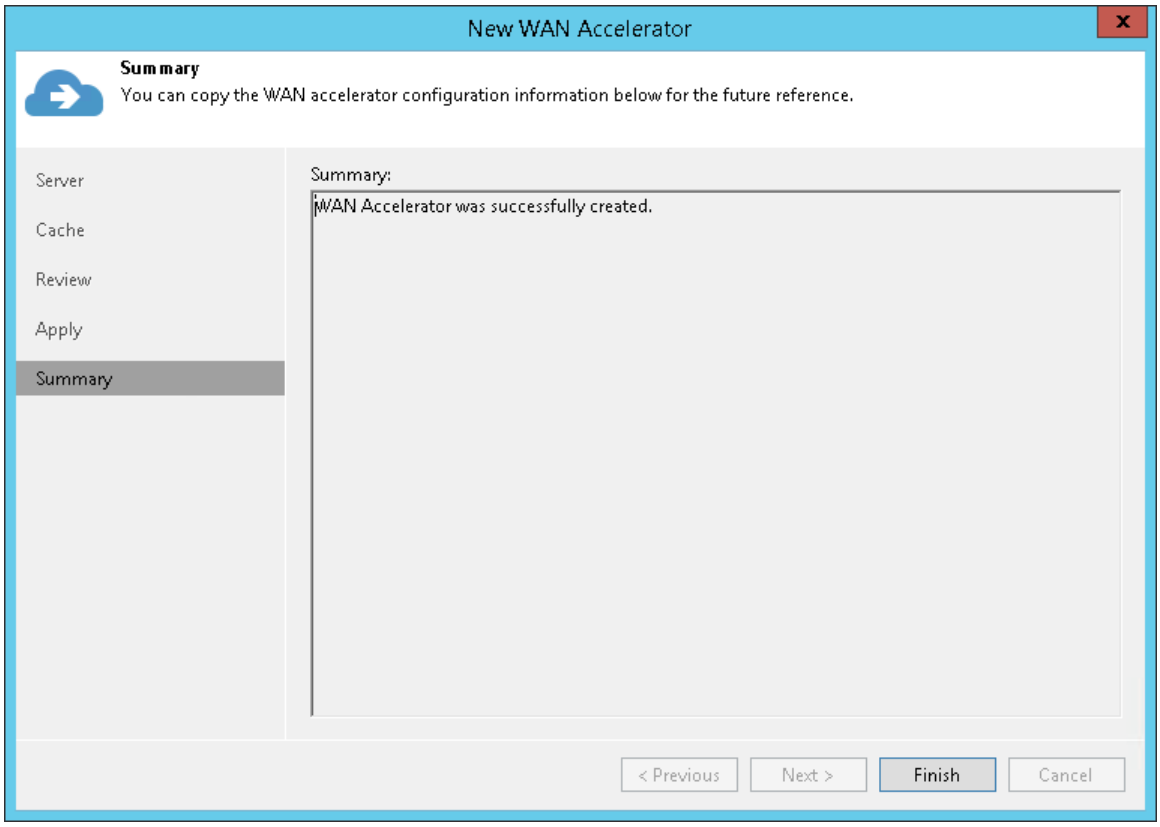
Step 5. Apply WAN Accelerator Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of creating the WAN accelerator and adding it to the backup infrastructure.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added WAN accelerator and click **Finish** to exit the wizard.

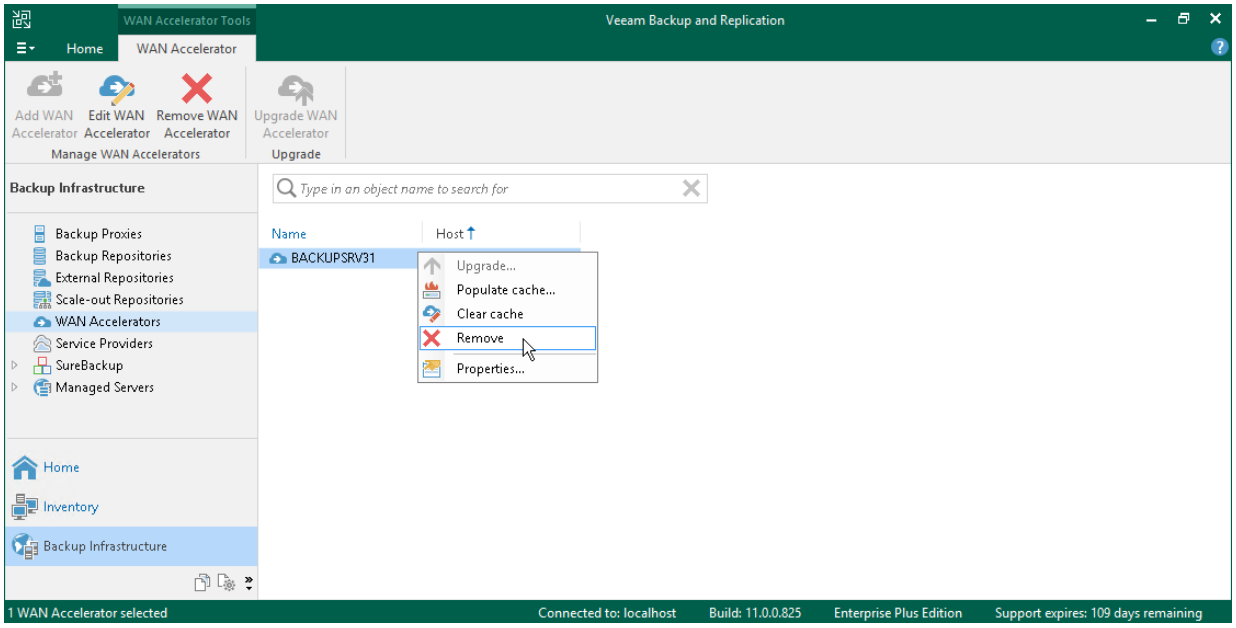


Removing WAN Accelerators

You can permanently remove a WAN accelerator from the backup infrastructure. When you remove a WAN accelerator, Veeam Backup & Replication unassigns the WAN accelerator role from the server, and this server is no longer used as a WAN accelerator. The server itself remains in the backup infrastructure.

To remove a WAN accelerator:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **WAN accelerators**.
3. In the working area, select the WAN accelerator and click **Remove WAN Accelerator** on the ribbon or right-click the WAN accelerator and select **Remove**.



Data Block Verification

During the VM copy process, Veeam Backup & Replication performs a CRC check for the VM traffic going between the source and target WAN accelerators. The CRC check helps ensure that the correct VM data goes to the target side and no corrupted data blocks are written to the global cache or to backup files in the target backup repository.

The check is performed in the following way:

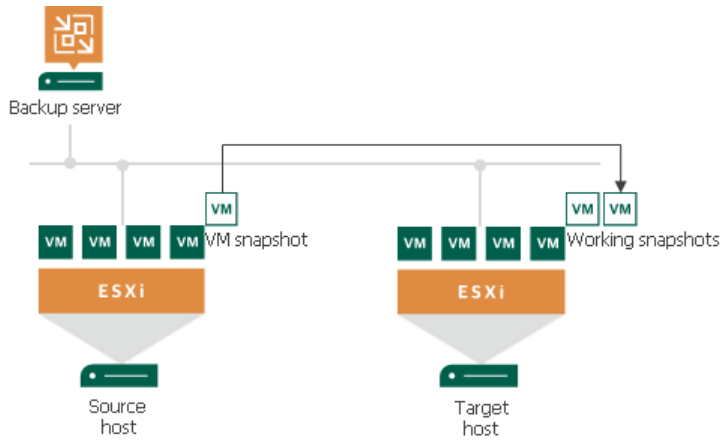
1. Before sending a data block to the target side, Veeam Backup & Replication calculates a checksum for the copied data block.
2. Once the data block is copied over WAN and before it is written to the global cache or to the target backup repository, Veeam Backup & Replication recalculates the checksum for this data block on the target side.
3. The source and target checksums are compared. If the checksums do not coincide, the target WAN accelerator sends a request to the source WAN accelerator for the correct data block. The source WAN accelerator re-sends the necessary data blocks to the target WAN accelerator as is and the re-sent data block is written to the global cache or to the backup file in the target backup repository on the fly.

Data Transport on WAN Disconnect

If you replicate VMs over WAN accelerators, and a WAN connection drops for short periods of time (less than 30 minutes), Veeam Backup & Replication transparently handles disconnect issues. It automatically resumes the data transport process from the point when the connection was lost. The resume on disconnect capability improves the reliability of off-site replication, reduces the backup window and minimizes the load on the WAN link.

If a WAN connection is lost for more than 30 minutes, Veeam Backup & Replication still does not finish the job with a failed status. After a WAN connection is resumed, Veeam Backup & Replication starts a new data transfer cycle. Data transported with every new transport cycle is written to a new working snapshot of a VM replica. As the WAN connection may drop several times, Veeam Backup & Replication can create a number of working snapshots.

Not to keep long snapshot chains, Veeam Backup & Replication merges earlier snapshots and maintains only two working snapshots for the VM replica. When all VM data is transferred to the target host, the two working snapshots are also merged to create one fully functional VM restore point.



If the WAN link is weak and drops constantly, Veeam Backup & Replication may fail to transport VM data by the time a new replication job session starts. In this case, during a new replication job session Veeam Backup & Replication attempts to transfer VM data that have changed since the last replication job session and VM data that were not transferred during the previous replication job session.

Data Encryption

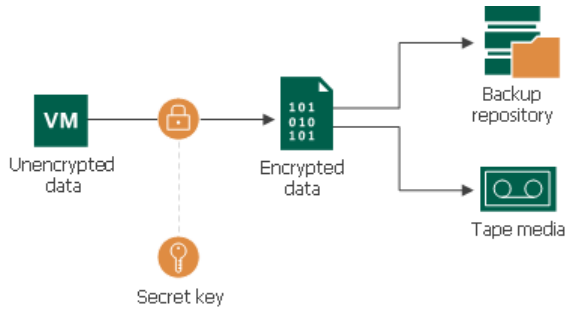
Data security is an important part of the backup strategy. You must protect your information from unauthorized access, especially if you back up sensitive VM data to off-site locations or archive it to tape. To keep your data safe, you can use data encryption.

Data encryption transforms data to an unreadable, scrambled format with the help of a cryptographic algorithm and a secret key. If encrypted data is intercepted, it cannot be unlocked and read by the eavesdropper. Only intended recipients who know the secret key can reverse encrypted information back to a readable format.

In Veeam Backup & Replication, encryption works at the following levels:

- Backup job
- Transaction log backup job
- Backup copy job
- VeeamZIP
- Tapes in media pools

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. Data decryption is also performed on the source side: Veeam Backup & Replication transfers encrypted data back to the source side and decrypts it there.



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

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).

Data Encryption and Deduplication

Data encryption has a negative effect on the deduplication ratio if you use a deduplicating storage appliance as a target. Veeam Backup & Replication uses different encryption keys for every job session. For this reason, encrypted data blocks sent to the deduplicating storage appliances appear as different though they may contain duplicate data. If you want to achieve a higher deduplication ratio, you can disable data encryption. If you still want to use encryption, you can enable the encryption feature on the deduplicating storage appliance itself.

Data Encryption and Compression

If data compression and data encryption are enabled for a job, Veeam Backup & Replication compresses VM data first and after that encodes the compressed data blocks. Both operations are performed at the source side.

Note, however, that if the **Decompress backup data blocks before storing** check box is selected in the backup repository settings, Veeam Backup & Replication does not compress VM data before encryption. 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.

Encryption Standards

Veeam Backup & Replication uses the following industry-standard data encryption algorithms:

Data Encryption

- To encrypt data blocks in backup files and files archived to tape, Veeam Backup & Replication uses the 256-bit AES with a 256-bit key length in the CBC-mode. For more information, see Advanced Encryption Standard (AES).
- To generate a key based on a password, Veeam Backup & Replication uses the Password-Based Key Derivation Function, PKCS #5 version 2.0. Veeam Backup & Replication uses 10,000 HMAC-SHA1 iterations and a 512-bit salt. For more information, see Recommendation for Password-Based Key Derivation.

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.

For more information, see RSA Cryptography Specifications.

Hashing Algorithms

Veeam Backup & Replication uses the following hashing algorithms:

- For digital signature generation: SHA-256
- For backward compatibility and certificate thumbprint generation: SHA-1
- For HMAC generation: SHA-1
- For random number generation: OpenSSL, cryptographic libraries provided by the operating system

Encryption Libraries

For Linux-based components and services, Veeam Backup & Replication uses Veeam Cryptographic Module.

For Veeam Data Movers installed on Microsoft Windows-based machines, Veeam Backup & Replication also uses Veeam Cryptographic Module. For other Microsoft Windows-based components and services, Veeam Backup & Replication uses Microsoft Crypto API.

Veeam Backup & Replication uses the following cryptographic service providers:

- Microsoft Base Cryptographic Provider. For more information, see Microsoft Docs.
- Microsoft Enhanced RSA and AES Cryptographic Provider. For more information, see Microsoft Docs.
- Microsoft Enhanced Cryptographic Provider. For more information, see Microsoft Docs.

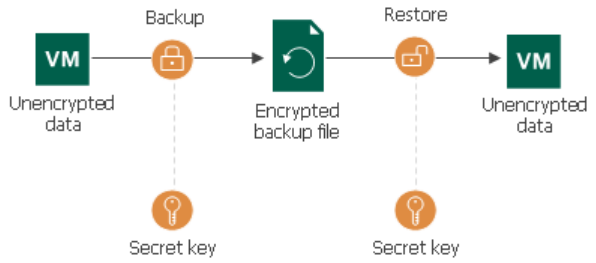
If you need Veeam Cryptographic Module and Microsoft Crypto API to be compliant with the Federal Information Processing Standards (FIPS 140), enable FIPS compliance as described in FIPS Compliance.

Veeam Backup & Replication encrypts stored credentials using the Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs.

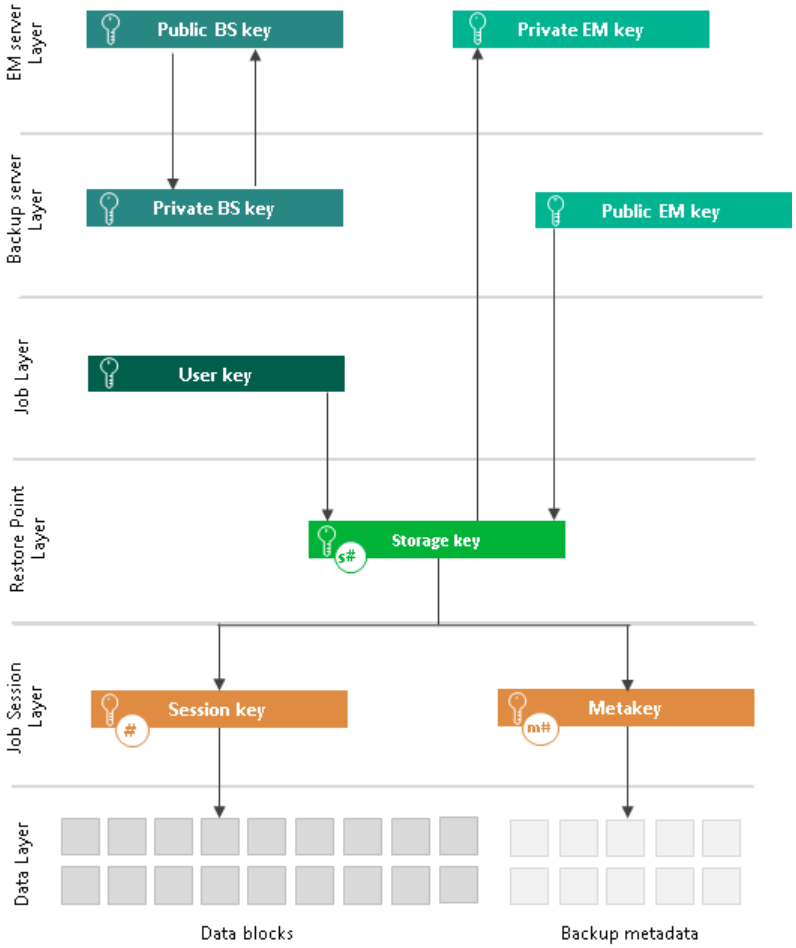
Encryption Algorithms

To encrypt data in backups and files, Veeam Backup & Replication employs a symmetric key encryption algorithm.

The symmetric, or single-key encryption algorithm, uses a single, common secret key to encrypt and decrypt data. Before data is sent to target side, it is encoded with a secret key. To restore encrypted data, you must have the same secret key. Users who do not have the secret key cannot decrypt data and get access to it.



Veeam Backup & Replication relies on a hierarchical encryption scheme. Each layer in the hierarchy encrypts the layer below with a key of specific type.



Encryption Keys

An encryption key is a string of random characters that is used to bring data to a scrambled format and back to unscrambled. Encryption keys encode and decode initial data blocks or underlying keys in the key hierarchy.

Veeam Backup & Replication uses the following types of keys:

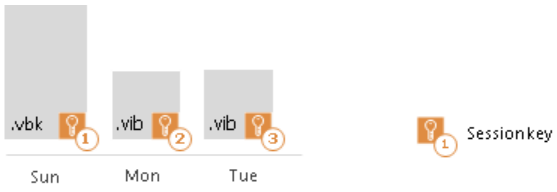
- Service keys generated by Veeam Backup & Replication:
  - Session key
  - Metakey
  - Storage key
- User key — a key generated using a user password.
- Enterprise Manager keys — a pair of keys used for data recovery without a password.
- Backup server keys — a pair of keys used for identity verification of the backup server.

Session Keys and Metakeys

The session key is the lowest layer in the encryption key hierarchy. When Veeam Backup & Replication encrypts data, it first encodes every data block in a file with a session key. For session keys, Veeam Backup & Replication uses the AES algorithm with a 256-bit key length in the CBC-mode.

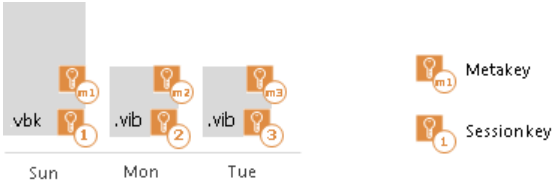
Veeam Backup & Replication generates a new session key for every job session. For example, if you have created an encrypted backup job and run 3 job sessions, Veeam Backup & Replication will produce 3 backup files that will be encrypted with 3 different session keys:

- Full backup file encrypted with session key 1
- Incremental backup file encrypted with session key 2
- Incremental backup file encrypted with session key 3



The session key is used to encrypt only data blocks in backup files or files archived to tape. To encrypt backup metadata, Veeam Backup & Replication applies a separate key — metakey. Use of a metakey for metadata raises the security level of encrypted backups.

For every job session, Veeam Backup & Replication generates a new metakey. For example, if you have run 3 job sessions, Veeam Backup & Replication will encrypt metadata with 3 metakeys.

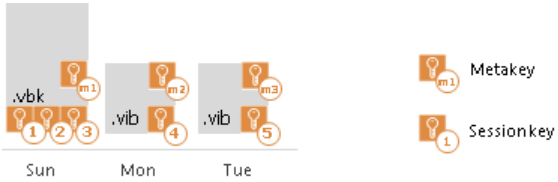


In the encryption process, session keys and metakeys are encrypted with keys of a higher layer — storage keys. Cryptograms of session keys and metakeys are stored to the resulting file next to encrypted data blocks. Metakeys are additionally kept in the configuration database.

Storage Keys

Backup files in the backup chain often need to be transformed, for example, in case you create a reverse incremental backup chain. When Veeam Backup & Replication transforms a full backup file, it writes data blocks from several restore points to the full backup file. As a result, the full backup file contains data blocks that are encrypted in different job sessions with different session keys.

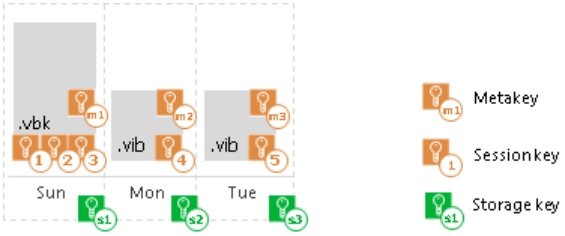
To restore data from such "composed" backup file, Veeam Backup & Replication would require a bunch of session keys. For example, if the backup chain contains restore points for 2 months, Veeam Backup & Replication would have to keep session keys for a 2-month period.



In such situation, storing and handling session keys would be resource consuming and complicated. To facilitate the encryption process, Veeam Backup & Replication introduces another type of service key — a storage key.

For storage keys, Veeam Backup & Replication uses the AES algorithm. A storage key is directly associated with one restore point in the backup chain. The storage key is used to encrypt the following keys in the encryption hierarchy:

- All session keys for all data blocks in one restore point
- Metakey encrypting backup metadata



During the restore process, Veeam Backup & Replication uses one storage key to decrypt all session keys for one restore point, no matter how many session keys were used to encrypt data blocks in this restore point. As a result, Veeam Backup & Replication does not need to keep the session keys history in the configuration database. Instead, it requires only one storage key to restore data from one file.

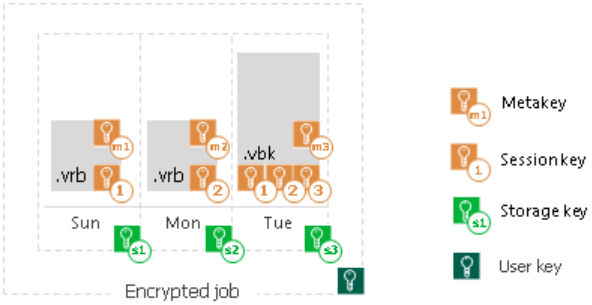
In the encryption process, storage keys are encrypted with keys of a higher layer — user keys and optionally a public Enterprise Manager key. Cryptograms of storage keys are stored to the resulting file next to encrypted data blocks, and cryptograms of session keys and metakeys.

Storage keys are also kept in the configuration database. To maintain a set of valid storage keys in the database, Veeam Backup & Replication uses retention policy settings specified for the job. When some restore point is removed from the backup chain by retention, the storage key corresponding to this restore point is also removed from the configuration database.

User Keys

When you enable encryption for a job, you must define a password to protect data processed by this job, and define a hint for the password. The password and the hint are saved in the job settings. Based on this password, Veeam Backup & Replication generates a user key.

The user key protects data at the job level. In the encryption hierarchy, the user key encrypts storage keys for all restore points in the backup chain.



During the encryption process, Veeam Backup & Replication saves a hint for the password to the encrypted file. When you decrypt a file, Veeam Backup & Replication displays a hint for the password that you must provide. After you enter a password, Veeam Backup & Replication derives a user key from the password and uses it to unlock the storage key for the encrypted file.

According to the security best practices, you must change passwords for encrypted jobs regularly. When you change a password for the job, Veeam Backup & Replication creates a new user key and uses it to encrypt new restore points in the backup chain.

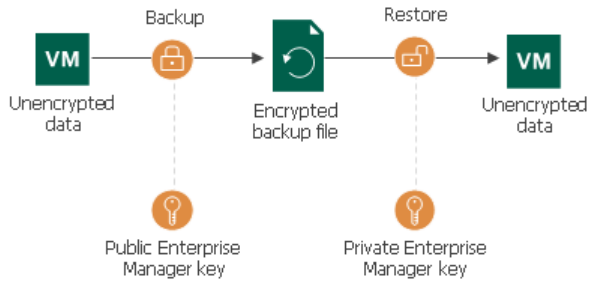
> **IMPORTANT**
>
> You must always remember passwords set for jobs or save these passwords in a safe place. If you lose or forget the password, you can restore data from a backup file by issuing a request to Veeam Backup Enterprise Manager. For more information, see How Decryption Without Password Works.

Enterprise Manager Keys

In some cases, a password required for data decryption may be lost or forgotten, or a user who knows the password may leave your organization. As a result, you cannot recover data from backups or tapes encrypted with this password, and encrypted data becomes unusable.

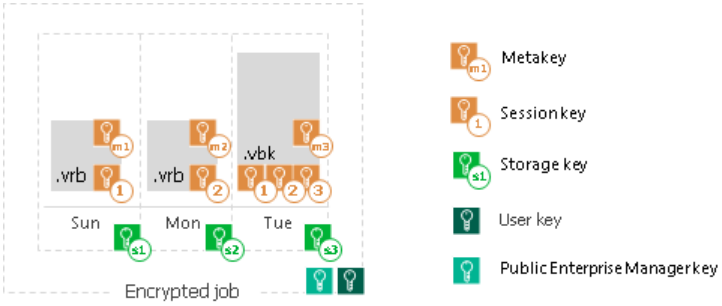
Veeam Backup & Replication offers you a way to restore encrypted data even if you do not have a password. For this purpose, Veeam Backup & Replication employs an additional pair of keys in the encryption process — Enterprise Manager keys.

Enterprise Manager keys is a pair of matching RSA keys: a public key and a private key. The public Enterprise Manager key is used to encrypt data, while the private Enterprise Manager key is used to decrypt data encrypted with the public key.



In the encryption process, Enterprise Manager keys perform a role similar to the user key: the public Enterprise Manager key encrypts storage keys and the private Enterprise Manager key decrypts them. Technically, Enterprise Manager keys offer an alternative to the user key. When you create an encrypted backup file or archive encrypted data to tape, Veeam Backup & Replication encrypts storage keys with two types of keys simultaneously:

- User key
- Public Enterprise Manager key



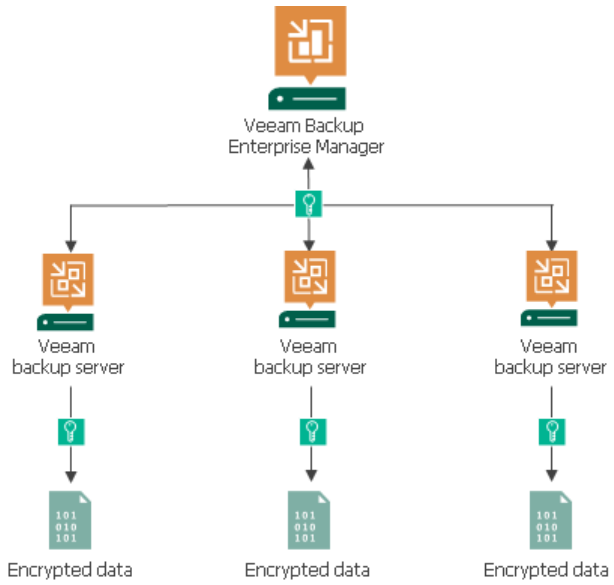
When you decrypt a file and the password is lost, Veeam Backup & Replication cannot derive the user key from the password. In this situation, you can send a request to Veeam Backup Enterprise Manager. Veeam Backup Enterprise Manager will employ the private Enterprise Manager key instead of the user key to unlock storage keys and decrypt the file content. For more information, see How Decryption Without Password Works.

Enterprise Manager keys take part in the encryption process if the following two conditions are met:

1. If you are using a legacy socket-based license, Enterprise or higher edition is required. Note that Enterprise Manager keys functionality is included in the Veeam Universal License.
2. You have Veeam Backup Enterprise Manager installed and your backup servers are connected to Veeam Backup Enterprise Manager.

Enterprise Manager keys make up a pair of matching keys — a keyset. Enterprise Manager keysets are created and managed on the Veeam Backup Enterprise Manager server. During installation of Veeam Backup Enterprise Manager, the setup automatically generates a new keyset containing a public Enterprise Manager key and a private Enterprise Manager key. You can use Veeam Backup Enterprise Manager to create new Enterp rise Manager keysets, activate them, import and export keysets and specify retention for their lifetime.

The public Enterprise Manager key is made publicly available to backup servers. When you connect backup servers to Veeam Backup Enterprise Manager, the public Enterprise Manager key is automatically propagated to these backup servers.



Veeam Backup Enterprise Manager acts as a manager for public Enterprise Manager keys but does not store these keys. After the public Enterprise Manager key is propagated to the backup server, it is kept in the configuration database.

Private Enterprise Manager keys, on the contrary, are not distributed anywhere: they are kept only on Veeam Backup Enterprise Manager.

Backup Server Keys

Eavesdroppers may potentially use Veeam Backup Enterprise Manager to unlock files encrypted with Veeam Backup & Replication. If eavesdroppers intercept an encrypted file, they may generate a request for file unlocking and send such request to Veeam Backup Enterprise Manager Administrators. Having received a response from Veeam Backup Enterprise Manager, eavesdroppers will be able unlock the encrypted file without a password.

To protect you against the "man-in-the-middle" attack, Veeam Backup & Replication uses backup server keys. Backup server keys are a pair of RSA keys, public and private, that are generated on the backup server.

- The public backup server key is sent to Veeam Backup Enterprise Manager to which the backup server is connected, and saved in the Veeam Backup Enterprise Manager configuration database.
- The private backup server key is kept on the backup server in the Veeam Backup & Replication configuration database.

Backup server keys are used to authenticate the identity of the request sender. When the backup server generates a request to unlock a file, it adds a signature encrypted with the private backup server key to this request.

When Veeam Backup Enterprise Manager processes the request, it uses the public backup server key to decrypt the signature and identify the request sender. If the backup server used for request generation is not added to Veeam Backup Enterprise Manager, Veeam Backup Enterprise Manager will not find a matching public key in its database. As a result, Veeam Backup Enterprise Manager will not be able to identify the sender and the storage key decryption process will fail.

How Data Encryption Works

Data encryption is performed as part of backup, backup copy or archiving to tape processes. Encryption works at the source side, before data is transported to the target. Encryption keys are not passed to the target side, unless you run a backup copy job over WAN accelerators or perform health check for the encrypted backup files.

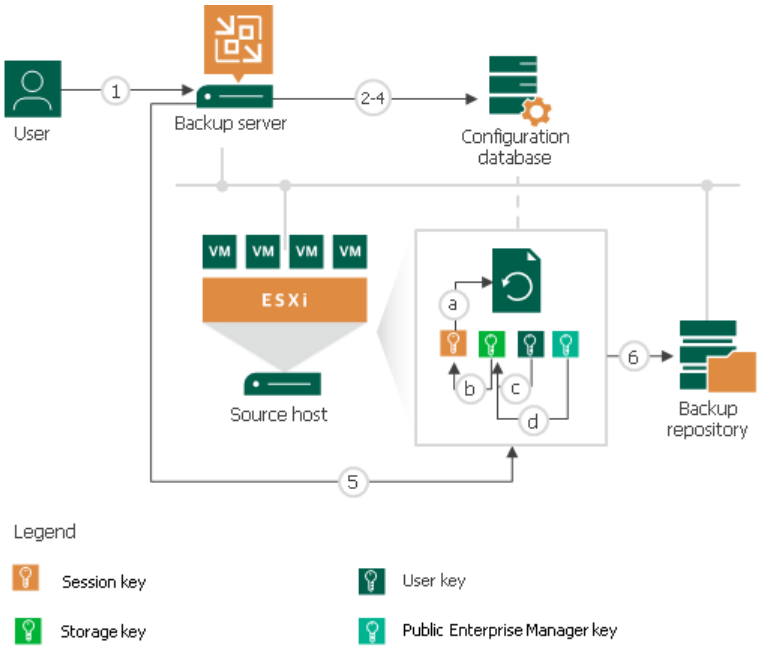
> **NOTE**
>
> The procedure below describes the encryption process for backup, backup copy jobs and VeeamZIP tasks. For more information about encrypting data on tapes, see Tape Encryption.

The encryption process includes the following steps:

1. When you create a new job, you enable the encryption option for the job and enter a password to protect data at the job level.
2. Veeam Backup & Replication generates a user key based on the entered password.
3. When you start an encrypted job, Veeam Backup & Replication creates a storage key and stores this key to the configuration database.
4. Veeam Backup & Replication creates a session key and a metakey. The metakey is stored to the configuration database.
5. Veeam Backup & Replication processes job data in the following way:
    a. The session key encrypts data blocks in the backup file. The metakey encrypts backup metadata.
    b. The storage key encrypts the session key and the metakey.
    c. The user key encrypts the storage key.
    d. If you use the Veeam Universal License, (or, for legacy-based license, Enterprise or higher edition), and the backup server is connected to Veeam Backup Enterprise Manager, the Enterprise Manager key also encrypts the storage key.

6. Encrypted data blocks are passed to the target. The cryptograms of the public Enterprise Manager key (if used), user key, storage key, session key and metakey are stored to the resulting file next to encrypted data blocks.

   If you use the Enterprise or Enterprise Plus edition of Veeam Backup & Replication and the backup server is connected to Veeam Backup Enterprise Manager, Veeam Backup & Replication saves two cryptograms of the storage key to the resulting file: one encrypted with the user key (c) and one encrypted with the Enterprise Manager key (d). Saving the cryptogram twice helps Veeam Backup & Replication decrypt the file even if a password is lost or forgotten. For more information, see How Decryption Without Password Works.



How Data Decryption Works

When you restore data from an encrypted backup file, Veeam Backup & Replication performs data decryption automatically in the background or requires you to provide a password.

- If encryption keys required to unlock the backup file are available in the Veeam Backup & Replication configuration database, you do not need to enter the password. Veeam Backup & Replication uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore does not differ from that from an unencrypted one.

  Automatic data decryption is performed if the following conditions are met:

  a. You encrypt and decrypt the backup file on the same backup server using the same Veeam Backup & Replication configuration database.

  b. [For backup file] The backup is not removed from the Veeam Backup & Replication console.

- If encryption keys are not available in the Veeam Backup & Replication configuration database, you need to provide a password to unlock the encrypted file.

Data decryption is performed at the source side, after data is transported back from the target side. As a result, encryption keys are not passed to the target side, which helps avoid data interception.

> **NOTE**
>
> The procedure below describes the decryption process for backup, backup copy jobs and VeeamZIP tasks. For more information about decrypting tape data, see Tape Encryption.

The decryption process includes the following steps. Note that steps 1 and 2 are required only if you decrypt the file on the backup server other than the backup server where the file was encrypted.

1. You import the file to the backup server. Veeam Backup & Replication notifies you that the imported file is encrypted and requires a password.
2. You specify a password for the imported file. If the password has changed once or several times, you need to specify the password 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 full backup file for import, you must specify the whole set of passwords that were used to encrypt files in the backup chain.
3. Veeam Backup & Replication reads the entered password and generates the user key based on this password. With the user key available, Veeam Backup & Replication performs decryption in the following way:

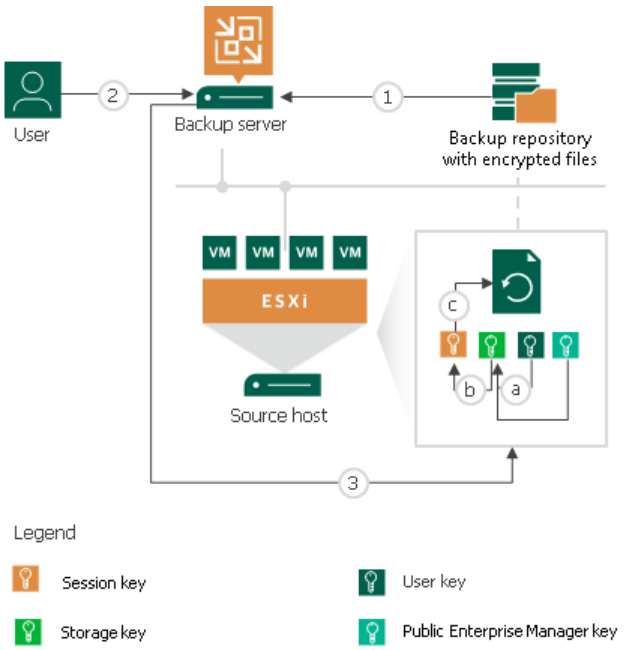
   a. Veeam Backup & Replication applies the user key to decrypt the storage key.

   b. The storage key, in its turn, unlocks underlying session keys and a metakey.

   c. Session keys decrypt data blocks in the encrypted file.

After the encrypted file is unlocked, you can work with it as usual.

If you have lost or forgotten a password for an encrypted file, you can issue a request to Veeam Backup Enterprise Manager and restore data from an encrypted file using Enterprise Manager keys. For more information, see Enterprise Manager Keys and How Decryption Without Password Works.



How Decryption Without Password Works

When you import an encrypted backup file or tape media to the backup server, you need to enter a password to decrypt data. In some cases, however, a password can be lost or forgotten. Veeam Backup & Replication offers a way to restore data from encrypted backups or tapes even if a password is not available.

You can restore of data from encrypted backups or tapes without a password only if your backup infrastructure meets the following conditions:

1. You use Veeam Universal License (or a legacy socket-based license, Enterprise or higher edition).
2. The backup servers on which you encrypted data is added to Veeam Backup Enterprise Manager.
3. The backup server on which you generate a request for data decryption is added to Veeam Backup Enterprise Manager.

If the backup server on which you encrypt data is added to Veeam Backup Enterprise Manager, Veeam Backup & Replication employs the public Enterprise Manager key in the encryption process. To decrypt backups or tapes encrypted with the public Enterprise Manager key, you can apply a matching private Enterprise Manager key, instead of a password. The private Enterprise Manager key unlocks the underlying storage keys and lets you access the content of an encrypted file.

The restore process is accomplished with the help of two wizards that run on two servers:

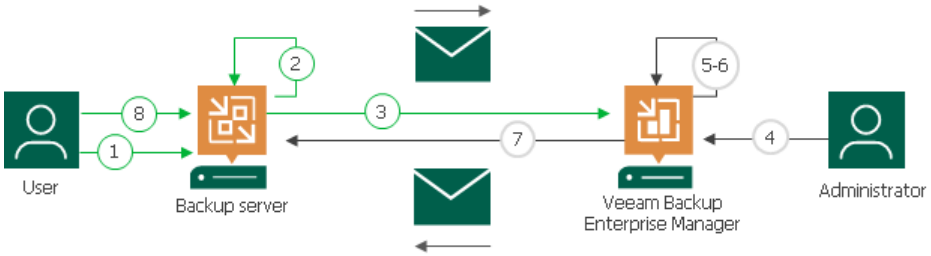
1. The **Encryption Key Restore** wizard on the backup server.
2. The **Password Recovery** wizard on the Veeam Backup Enterprise Manager server.

The restore process includes the next steps:

1. You start the **Encryption Key Restore** wizard on the backup server to issue a request for data recovery.
2. The **Encryption Key Restore** wizard generates a request to Veeam Backup Enterprise Manager. The request has the format of a text document and contains cryptograms of storage keys that must be decrypted, together with information about the public Enterprise Manager key that was used to encrypt data. At the end of the request, the backup server adds a signature encrypted with a private backup server key.
3. You send the request to the Veeam Backup Enterprise Manager Administrator, for example, using email.
4. The Veeam Backup Enterprise Manager Administrator starts the **Password Recovery** wizard on Veeam Backup Enterprise Manager and inserts the text of the request to the wizard.
5. Veeam Backup Enterprise Manager finds a matching public backup server key in Veeam Backup Enterprise Manager configuration database and decrypts the signature with this key.
6. Veeam Backup Enterprise Manager decrypts storage keys with the private Enterprise Manager key available on Veeam Backup Enterprise Manager, and generates a response in the **Password Recovery** wizard. The response has the format of a text document and contains decrypted storage keys.
7. The Veeam Backup Enterprise Manager Administrator sends the response to you, for example, using email.
8. You input the request to the **Encryption Key Restore** wizard. Veeam Backup & Replication processes the response, retrieves the decrypted storage keys and uses them to unlock encrypted backups or tapes and retrieve their content.

IMPORTANT

You can recover data only if Veeam Backup Enterprise Manager has a private Enterprise Manager key matching the public Enterprise Manager key that was used for data encryption. If a matching private Enterprise Manager key is not found in the Veeam Backup Enterprise Manager configuration database, the **Password Recovery** wizard will fail. In such situation, you can import a necessary private Enterprise Manager key using the import procedure. For more information, see *Exporting and Importing Enterprise Manager Keys* in Veeam Backup Enterprise Manager User Guide.



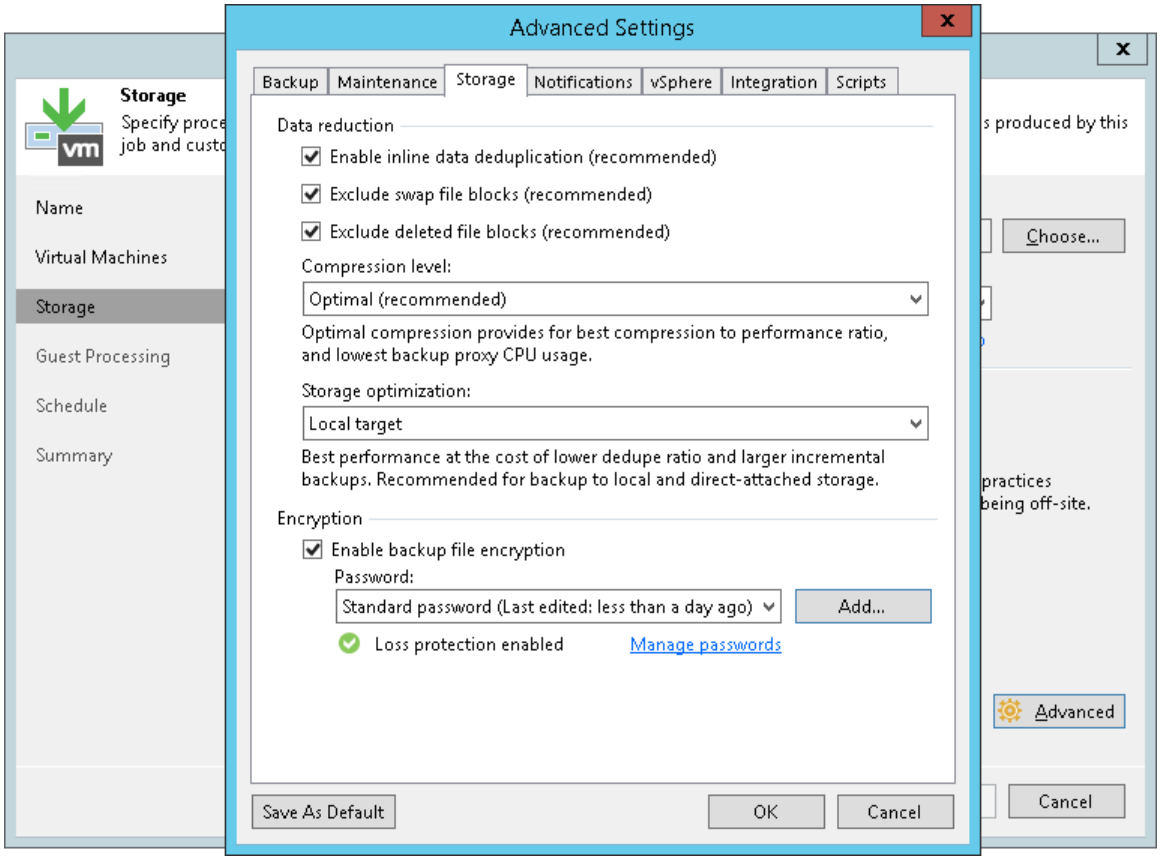
Encrypted Objects

The encryption algorithm works at the job level and media pool level. You can enable encryption for the following objects:

- Backup job
- Backup copy job
- VeeamZIP
- Tape

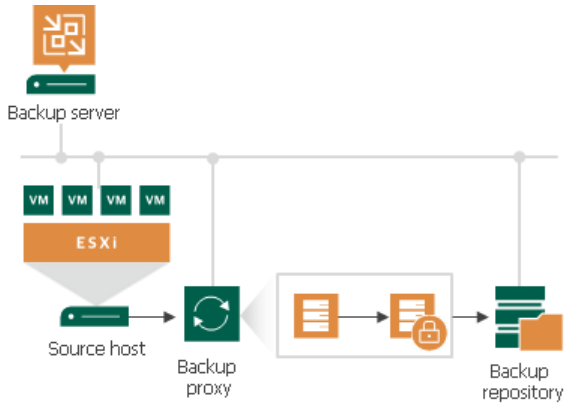
Backup Job Encryption

Encryption for a backup job is configured in the advanced job settings. You should enable the encryption option and specify a password to protect data in backup files produced by the backup job.



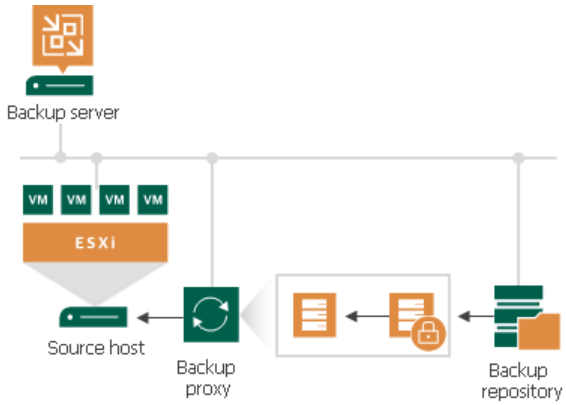
The backup job processing with encryption enabled includes the following steps:

1. You enable encryption for a backup job and specify a password.
2. Veeam Backup & Replication generates the necessary keys to protect backup data.
3. Veeam Backup & Replication encrypts data blocks in the backup proxy, either the dedicated or default one, and transfers them to the backup repository already encrypted.
4. On the backup repository, encrypted data blocks are stored to a resulting backup file.



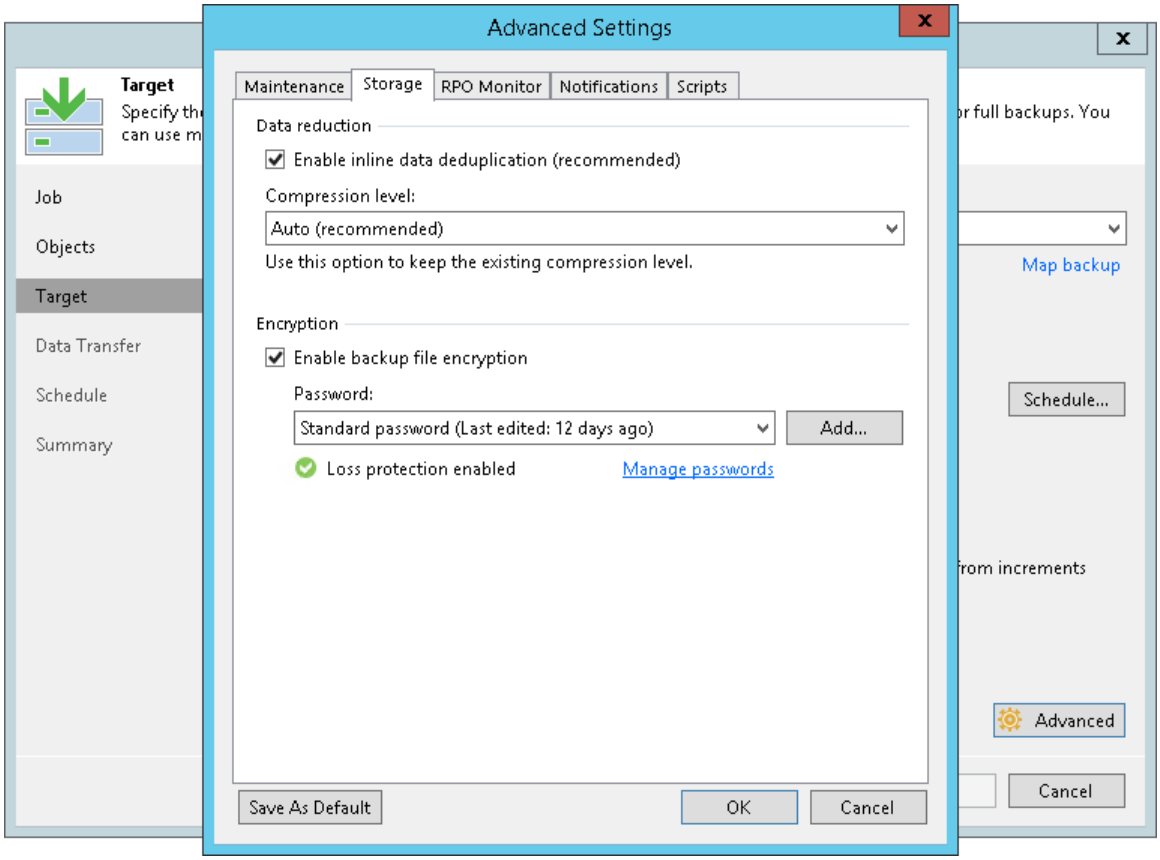
Restore of an encrypted backup file includes the following steps:

1. You import a backup file and define a password to decrypt the backup file. If the password has changed once or several times, you need to specify the password in the following manner:
   - If you select a metadata file (VBM) for import, you must specify the latest password that was used to encrypt files in the backup chain.
   - If you select a full backup file (VBK) for import, you must specify the whole set of passwords that were used to encrypt files in the backup chain.
2. Veeam Backup & Replication uses the provided passwords to generate user keys and unlock the subsequent keys for backup file decryption.
3. Veeam Backup & Replication retrieves data blocks from the backup file, sends them to the source side and decrypts them on the backup proxy, either the dedicated or default one.



Backup Copy Job Encryption

Encryption for a backup copy job is configured in the advanced job settings. You should enable the encryption option and specify a password to protect data in backup files produced by the backup copy job.



The workflow of the encrypted backup copy job depends on the path for data transfer:

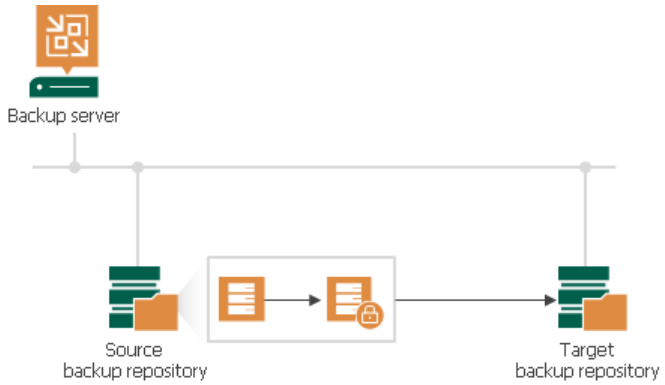
- Direct data path
- Over WAN accelerators

Direct Data Path

If you use a direct data path to transfer backups to the target backup repository, the encrypted backup copy job includes the following steps:

1. You enable encryption for a backup copy job and specify a password.
2. Veeam Backup & Replication generates the necessary keys to protect backup files produced by the backup copy job.
3. Veeam Backup & Replication encrypts data blocks on the source side and transfers them to the target backup repository.

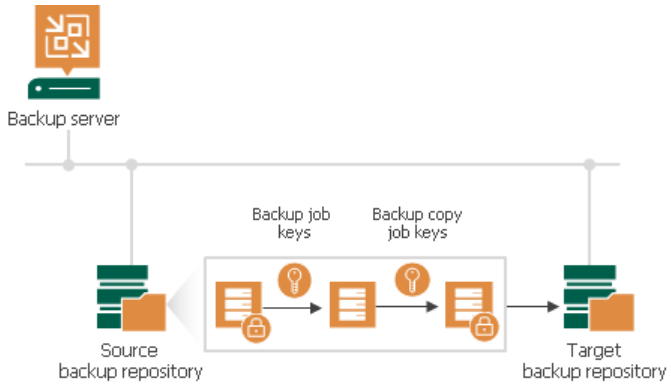
4. On the target backup repository, encrypted data blocks are stored to a resulting backup file.



An encrypted backup copy job may use an encrypted backup file as a source. In this situation, Veeam Backup & Replication does not perform double encryption. The backup copy job includes the following steps:

1. Veeam Backup & Replication decrypts data blocks of the encrypted source backup file. For the decryption process, it uses the storage key and metakeys stored in the configuration database.
2. Veeam Backup & Replication generates the necessary keys to protect backup files produced by the backup copy job.
3. Veeam Backup & Replication encrypts data blocks on the source side using these keys and transfers encrypted data blocks to the target backup repository.
4. On the target backup repository, encrypted data blocks are stored to a resulting backup file.

> **NOTE**
>
> Even if encryption is disabled in the backup copy job, Veeam Backup & Replication will decrypt data blocks of the encrypted source backup files.



The restore process for backups produced by backup copy jobs does not differ from that for backup jobs.

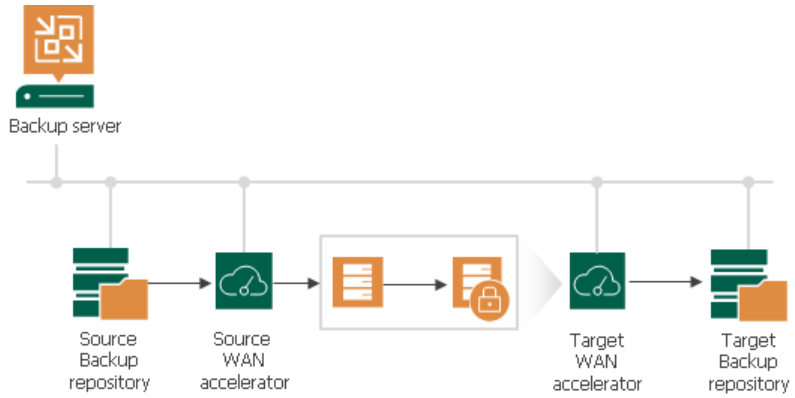
Over WAN Accelerators

WAN accelerators require reading data on the target side to perform such operations as global data deduplication, backup health check and so on. For this reason, if you use WAN accelerators for backup copy jobs, the encryption process is performed on the target side.

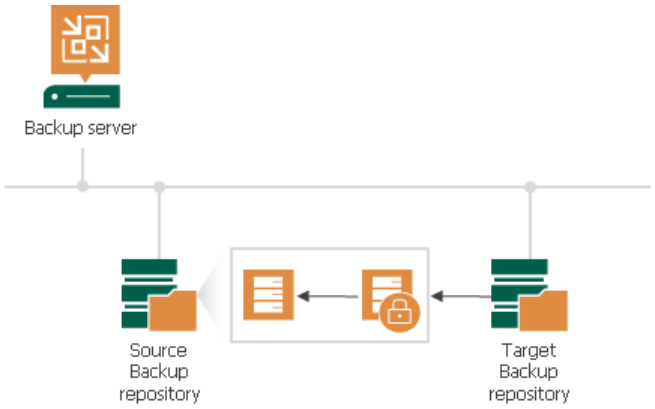
The backup copy job processing over WAN accelerators includes the following steps:

1. You enable encryption for a backup copy job and specify a password.

2. Veeam Backup & Replication generates necessary keys to protect backup files produced by the backup copy job.
3. Data blocks are passed to the target backup repository in the unencrypted format.
4. Received data blocks are encrypted on the target site and stored to a resulting backup file in the target backup repository.



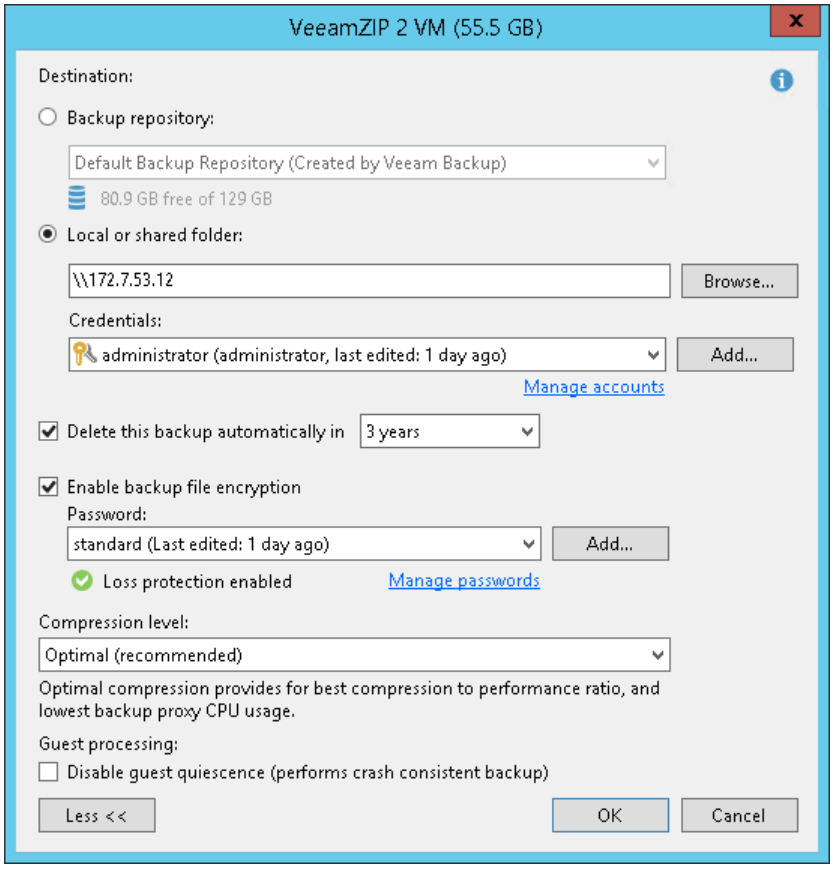
The restore process in this case does not differ from that for backup jobs. Veeam Backup & Replication retrieves data blocks from the backup file in the target backup repository, sends them to the source side and decrypts them on the source side.



When transporting data between WAN accelerators that face external networks, Veeam Backup & Replication encrypts the network traffic by default. For network traffic encryption, Veeam Backup & Replication uses the 256-bit Advanced Encryption Standard (AES). For more information, see Enabling Network Data Encryption.

VeeamZIP Encryption

If you want to create an encrypted VeeamZIP file, you should enable the encryption option and specify a password in VeeamZIP task options.



Data processing during VeeamZIP file creation and restore from a VeeamZIP file does not differ from that of a backup job.

Tape Encryption

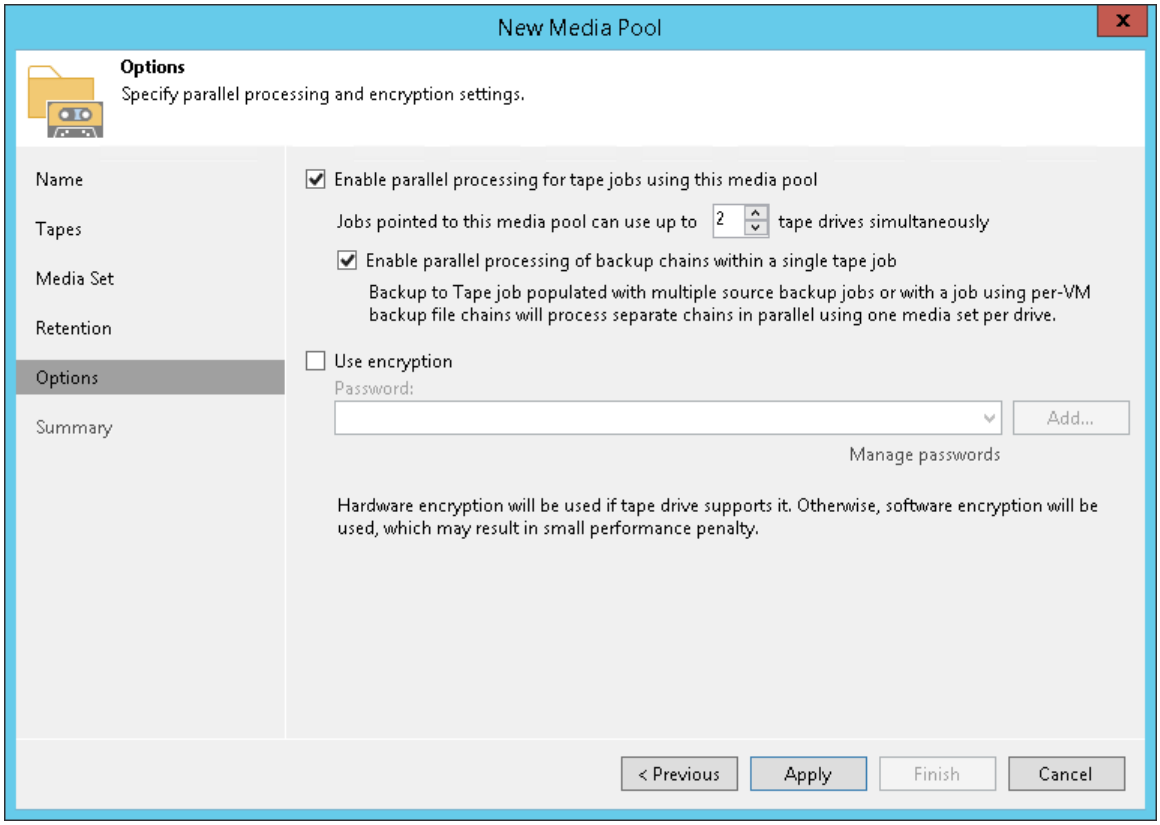
Veeam Backup & Replication supports two types of encryption for tape media:

- Hardware level: library- and driver-managed encryption mechanisms provided by the tape vendor
- Software level: the encryption mechanism provided by Veeam Backup & Replication

Hardware encryption has a higher priority. If hardware encryption is enabled for the tape media, Veeam Backup & Replication automatically disables its software encryption mechanism for such tape libraries. The Veeam encryption mechanism can only be used if hardware encryption is disabled at the tape device level or not supported.

To use the Veeam encryption mechanism, you need to enable encryption at the level of media pool. In this case, Veeam Backup & Replication will encrypt data for all jobs that use tapes from this media pool. Encryption is supported for both types of tape jobs:

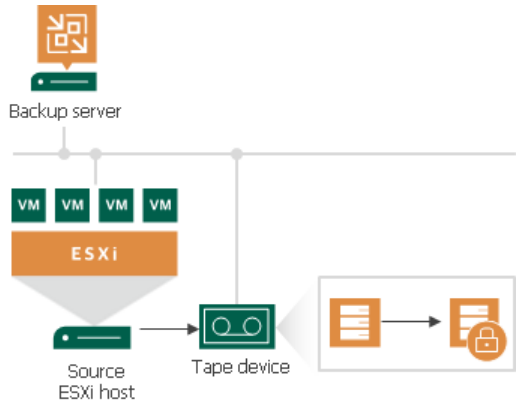
- Backup to tape jobs
- File to tape jobs



Encryption of data on tapes includes the following steps:

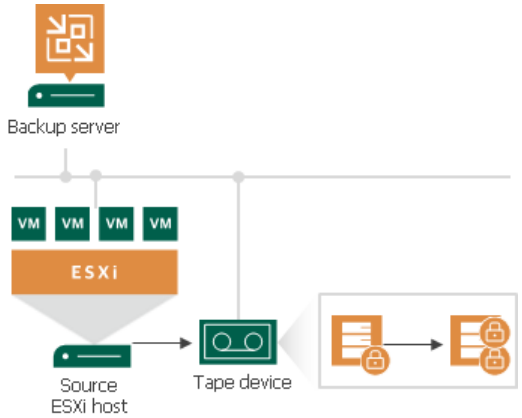
1. You enable encryption for a media pool and specify a password.
2. You select the media pool as a target for a backup to tape or file to tape job.
3. Veeam Backup & Replication generates the necessary keys to protect data archived to tape.

4. During the backup to tape or file to tape job, the key is passed to the target side. In case of hardware encryption, Veeam Backup & Replication passes the key to the tape device, and the tape device uses its mechanism to encrypt data on tapes. In case of software encryption, Veeam Backup & Replication passes the keys to the tape server, and encrypts data when it is archived to tape.



Backup to tape jobs allow double encryption. The backup to tape job uses a backup file as a source of data. If the backup file is encrypted with the initial backup job and the encryption option is enabled for the backup to tape job, too, the resulting backup file will be encrypted twice. To decrypt such backup file, you will need to subsequently enter two passwords:

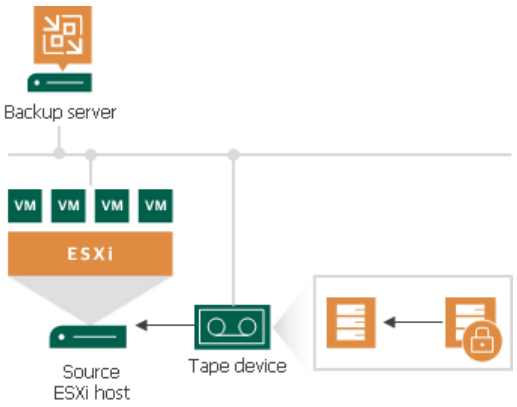
- Password for the initial backup job
- Password for the media pool



Restore of encrypted data from tape includes the following steps:

1. You insert tape with encrypted data into the tape drive and perform tape catalogization. The catalogization operations lets Veeam Backup & Replication understand what data is written to tape.
2. You provide a password to decrypt data archived to tape.

3. Veeam Backup & Replication uses the provided password to generate a user key and unlock the subsequent keys for data Veeam Backup & Replication retrieves data blocks from encrypted files on tapes and decrypts them.



Encryption Best Practices

To guarantee the flawless process of data encryption and decryption, consider the following recommendations.

Password

Mind the following recommendations when you create a password:

1. Use strong passwords that are hard to crack or guess:
   - The password must be at least 8 characters long.
   - The password must contain uppercase and lowercase characters.
   - The password must be a mixture of alphabetic, numeric and punctuation characters.
   - The password must significantly differ from the password you used previously.
   - The password must not contain any real information related to you, for example, date of birth, your pet's name, your logon name and so on.
2. Provide a meaningful hint for the password that will help you recall the password. The hint for the password is displayed when you import an encrypted file or tape to the backup server and attempt to unlock it.
3. Keep passwords in the safe place. If you lose or forget your password, you will not be able to recover data from backups or tapes encrypted with this password, unless you use Enterprise Manager keys in the encryption process.
4. Change passwords for encrypted jobs regularly. Use of different passwords helps increase the encryption security level.

Data Recovery without Password

If you use Veeam Universal License (or a legacy socket-based license, Enterprise or higher edition), connect backup servers to Veeam Backup Enterprise Manager. In this case, Veeam Backup & Replication will employ Enterprise Manager keys in the encryption process, which will let you recover data from encrypted backups and tapes even if the password is lost or forgotten.

Mind the following recommendations for Enterprise Manager keysets:

1. Create and activate new Enterprise Manager keysets regularly. When you activate a keyset, the public Enterprise Manager key is automatically propagated to backup servers connected to Veeam Backup Enterprise Manager and is used for encrypted jobs on these servers.
2. Create backup copies of Enterprise Manager keysets and keep them in a safe place. If your installation of Veeam Backup Enterprise Manager goes down for some reason, you will lose private Enterprise Manager keys. As a result, you will not be able to use the Veeam Backup Enterprise Manager functionality to recover data from backups and tapes without a password.

For more information on data decryption without a password, see Decrypting Data Without Password.

Encryption for Existing Jobs

If you enable encryption for an existing job, except the backup copy job, during the next job session Veeam Backup & Replication will automatically 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.

If you change encryption settings for an existing backup copy job, you will need to create an active full backup manually. For more information, see Creating Active Full Backups.

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 this Veeam KB article.

If you change the password for the already encrypted job, during the next job session Veeam Backup & Replication will create a new incremental backup file. The created backup file and subsequent backup files in the backup chain will be encrypted with the new password.

NOTE

To unlock a backup encrypted with several passwords, you must decrypt it in the following manner:

- If you import a metadata file (VBM), provide the latest password that was used to encrypt files in the backup chain.
- If you import a full backup file (VBK), provide the whole set of passwords that were used to encrypt files in the backup chain.

For more information, see Decrypting Data with Password.

Restoring Data from Encrypted Backups

When you restore data from an encrypted backup, Veeam Backup & Replication performs data decryption automatically in the background or requires you to specify a password.

- If encryption keys required to unlock the backup file are available in the configuration database, you do not need to specify the password. Veeam Backup & Replication uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore from the encrypted backup does not differ from that from an unencrypted one.

  Automatic backup file decryption is performed if the following conditions are met:

  a. You encrypt and decrypt the backup file on the same backup server that uses the same configuration database.

  b. The backup is not removed from the Veeam Backup & Replication console.

- If encryption keys are not available in the configuration database, you can restore data from the encrypted backup with the following methods:

  - You can provide a password or a set of passwords to unlock an encrypted file. For more information, see Decrypting Data with Password.
  - You can use Veeam Backup Enterprise Manager to unlock an encrypted file without a password. For more information, see Decrypting Data Without Password.

Decrypting Data with Password

To unlock an encrypted file, you must specify a password. The password must be the same as the password that was used to encrypt the backup file.

To decrypt a backup file:

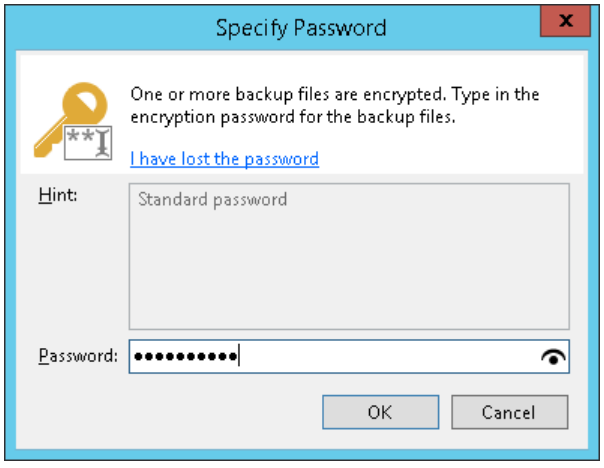
1. Import an encrypted backup file to the Veeam Backup & Replication console. After the import, the encrypted backup will appear under the **Backups** > **Disk (encrypted)** node in the inventory pane.
2. In the inventory pane, select **Disk (encrypted)**.
3. In the working area, select the imported backup and click **Specify Password** on the ribbon or right-click the backup and select **Specify password**.
4. In the **Description** field of the **Specify Password** window, Veeam Backup & Replication displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.
5. 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 metadata file (VBM) for import, you must specify the latest password that was used to encrypt files in the backup chain.
- If you select a full backup file (VBK) 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. You can perform restore operations with the backup file in a regular manner.

> **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.



Decrypting Data Without Password

If you have lost or forgotten a password, you can unlock an encrypted file with the help of Veeam Backup Enterprise Manager.

You can restore data without a password only if the following conditions are met:

1. You use Veeam Universal License (or a legacy socket-based license, Enterprise or higher edition).
2. The backup server on which you encrypted data is connected to Veeam Backup Enterprise Manager.
3. The backup server on which you generate a request for data decryption is connected to Veeam Backup Enterprise Manager.
4. Password loss protection is enabled on Veeam Backup Enterprise Manager (enabled by default). You can check the configured settings as described in the Managing Encryption Keys section in the Veeam Backup Enterprise Manager Guide.

> **IMPORTANT**
>
> Backup servers that you use for data decryption must be connected to the same instance of Veeam Backup Enterprise Manager. If you connect the backup server to several instances of Veeam Backup Enterprise Manager, this may cause unexpected behavior, and the decryption process may fail.

The restore process is accomplished with the help of two wizards that run on two servers:

1. The **Encryption Key Restore** wizard on the backup server.
2. The **Password Recovery** wizard on the Veeam Backup Enterprise Manager server.

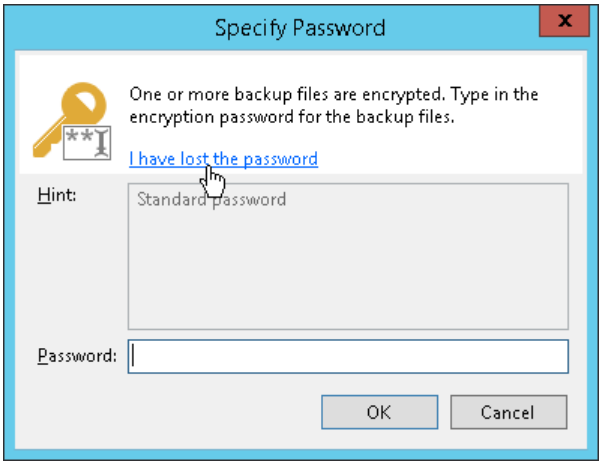
To restore encrypted data without a password:

1. Create a request for data restore
2. Process the request in Veeam Backup Enterprise Manager
3. Complete the key restore process

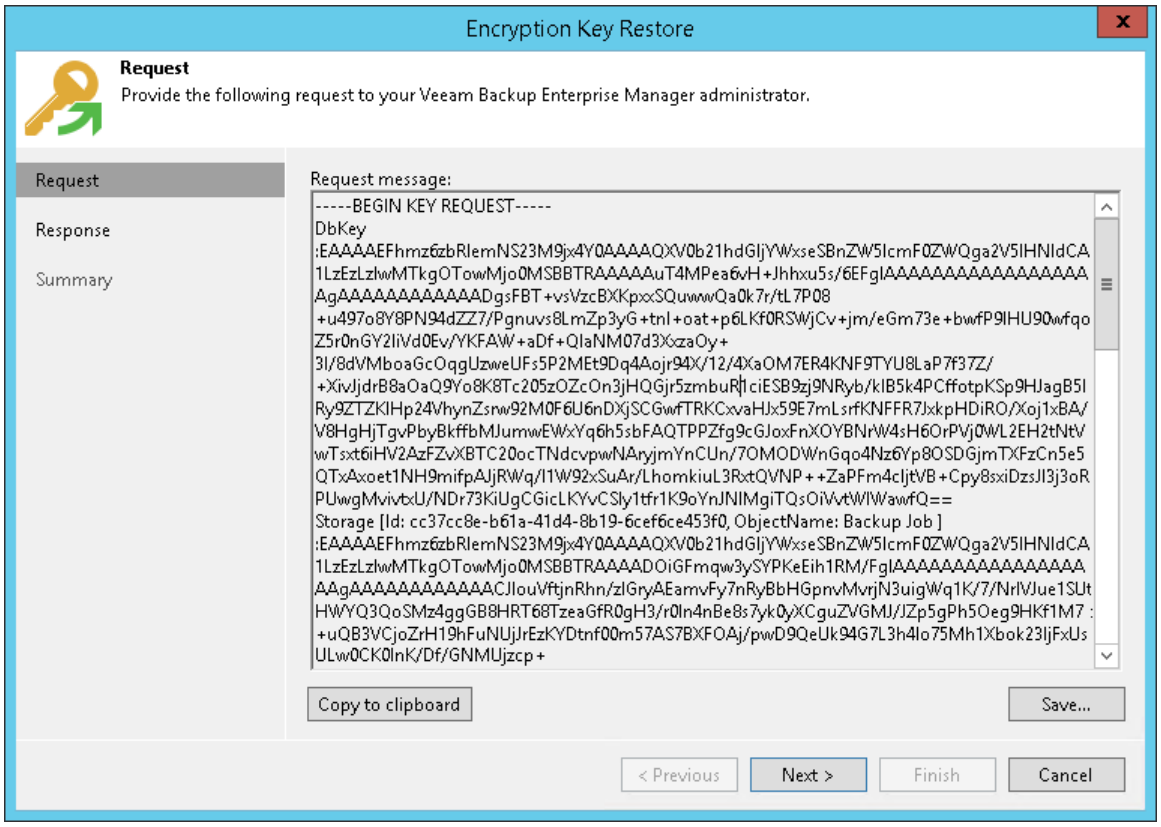
Step 1. Create Request for Data Restore

This procedure is performed by the Veeam Backup Administrator on the backup server.

1. Import encrypted backup to the Veeam Backup & Replication console.
2. Select the imported backup and click **Specify Password** on the ribbon or right-click the backup and select **Specify password**.
3. In the **Specify Password** window, click the **I have lost the password** link.



4. Veeam Backup & Replication will launch the **Encryption Key Restore** wizard. At the **Request** step of the wizard, review the generated request for data recovery. Use buttons at the bottom of the wizard to copy the request to the Clipboard or save the request to a text file.
5. Send the copied request by email or pass it in any other way to the Veeam Backup Enterprise Manager Administrator.



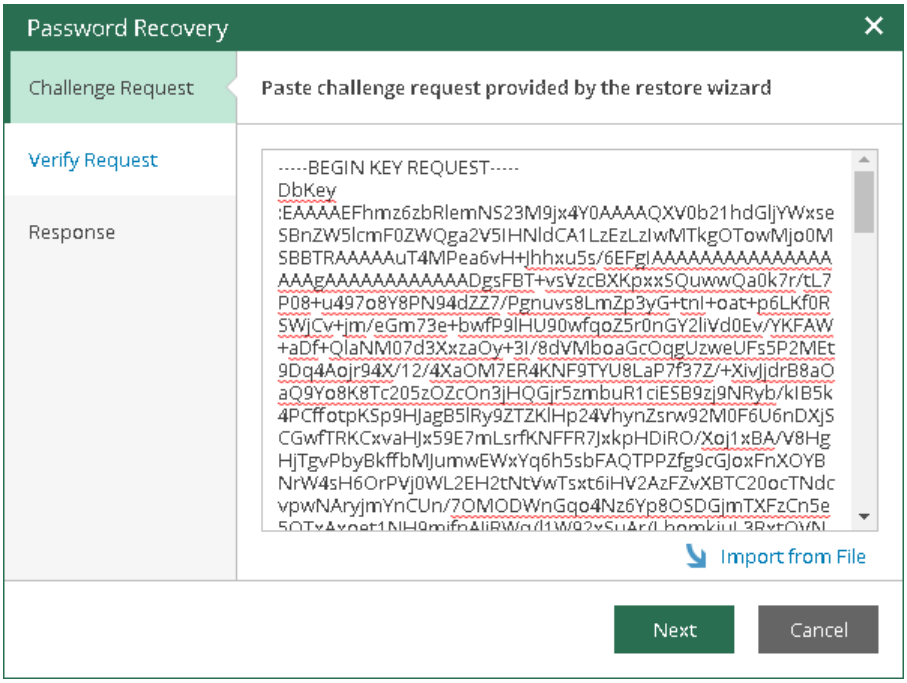
TIP

You can close the **Encryption Key Restore** wizard on the backup server and start it anew when you receive a response from the Veeam Backup Enterprise Manager Administrator.

Step 2. Process Request in Veeam Backup Enterprise Manager

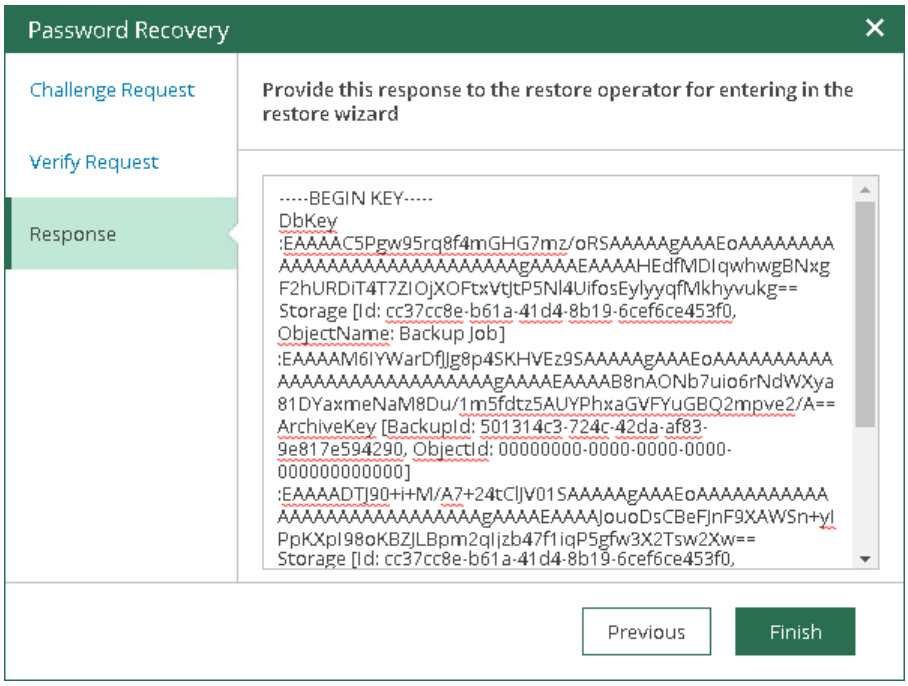
This procedure is performed by the Veeam Backup Enterprise Manager Administrator on the Veeam Backup Enterprise Manager server.

1. Copy the obtained request to the Clipboard.
2. In Veeam Backup Enterprise Manager, go to the **Configuration** > **Key Management** section.
3. Click **Password Recovery** at the top of the section to open the **Password Recovery** wizard.
4. Paste the request that you have received from the Veeam Backup Administrator. You can use the **[CTRL+V]** key combination or click **Paste** at the bottom of the wizard.



5. Follow the next steps of the wizard. At the **Response** step of the wizard, copy the text displayed in the wizard to the Clipboard.

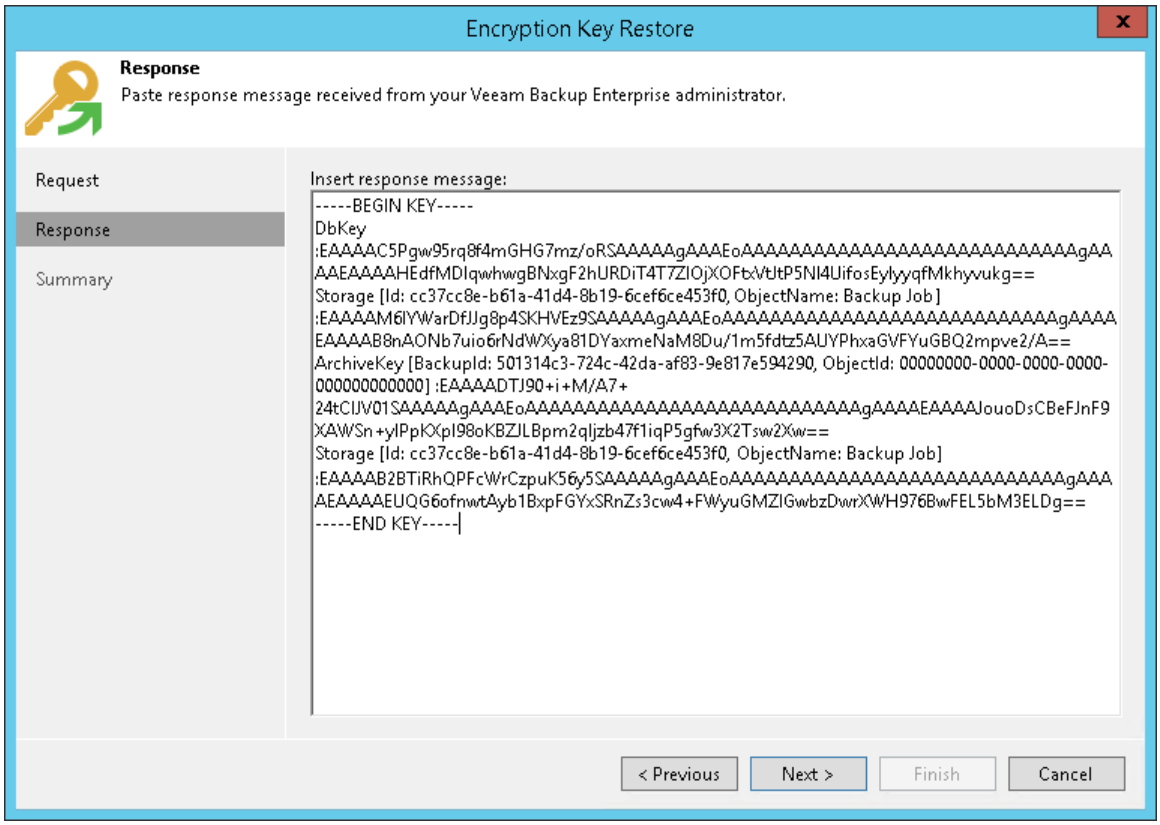
6. Send the copied response by email or pass it in any other way to the Veeam Backup Administrator working on the backup server.



Step 3. Complete Key Restore Process

This procedure is performed by the Veeam Backup Administrator on the backup server.

1. In Veeam Backup & Replication, get back to the **Encryption Key Restore** wizard.
2. Enter the copied response to the text window at the **Response** step of the **Encryption Key Restore** wizard.
3. Follow the next steps of the wizard. At the last step, click **Finish**. Veeam Backup & Replication will retrieve the decrypted storage keys from the response, apply them to the encrypted file and unlock the file content.



Restoring Encrypted Data from Tapes

When you restore data from encrypted tapes, Veeam Backup & Replication performs data decryption automatically in the background or requires you to provide a password.

- If encryption keys required to unlock the tape are available in the Veeam Backup & Replication database, you do not need to enter the password to decrypt the tape. Veeam Backup & Replication uses keys from the database to unlock the encrypted tape. Data decryption is performed in the background and data restore from encrypted tapes does not differ from that from an unencrypted ones.

  Automatic tape decryption is performed if the following conditions are met:

  - You encrypt and decrypt tapes on the same Veeam backup server.
  - The tape is loaded to the tape library and information about this tape is available in the catalog.
  - The password specified in the settings of the media pool to which the tape belongs is the same as the password that was used for tape encryption.
- If encryption keys are not available in the Veeam Backup & Replication database, you can restore data from encrypted tapes with the following methods:
  - You can provide a password or a set of passwords to unlock the encrypted tape. For more information, see Decrypting Tapes with Password.
  - You can use Veeam Backup Enterprise Manager to unlock the encrypted tape without a password. For more information, see Decrypting Data Without Password.

Decrypting Tapes with Password

When you restore encrypted files or backups from tape, you need to specify a password that was used to encrypt data archived to tape.

To unlock encrypted tapes:

1. Insert encrypted tapes into the tape library.
2. Catalog the tapes so that Veeam Backup & Replication can read data archived on tape. After you perform catalogization, encrypted tapes will be displayed under the **Media** > **Encrypted** node in the corresponding tape library. On the cataloged tape, Veeam Backup & Replication displays the key icon to mark it as encrypted.
3. In the inventory pane, select the **Encrypted** node under **Media** node.
4. In the working area, select the imported tape and click **Specify password** on the ribbon or right-click the tape and select **Specify password.**
5. In the **Description** field of the **Specify Password** window, Veeam Backup & Replication displays a hint for the password that was used to encrypt the tape. Use the hint to recall the password.
6. In the **Password** field, enter the password for the tape.
7. If the imported tape is a part of a backup set but is not the last tape in this set, perform catalogization once again.

   When Veeam Backup & Replication creates a backup set, it writes catalog data to the last tape in this set.

   - If the imported group of tapes contains the last tape in the backup set, Veeam Backup & Replication retrieves catalog data from the last tape during the initial catalogization process (see point 2 of this procedure).
   - If the imported group of tapes does not contain the last tape in the backup set, Veeam Backup & Replication needs to additionally catalog files on imported tapes.

If you enter a correct password, Veeam Backup & Replication will decrypt the tape media. The tape will be moved under the corresponding media pool in the inventory pane. You can perform restore operations for data archived to tape as usual.

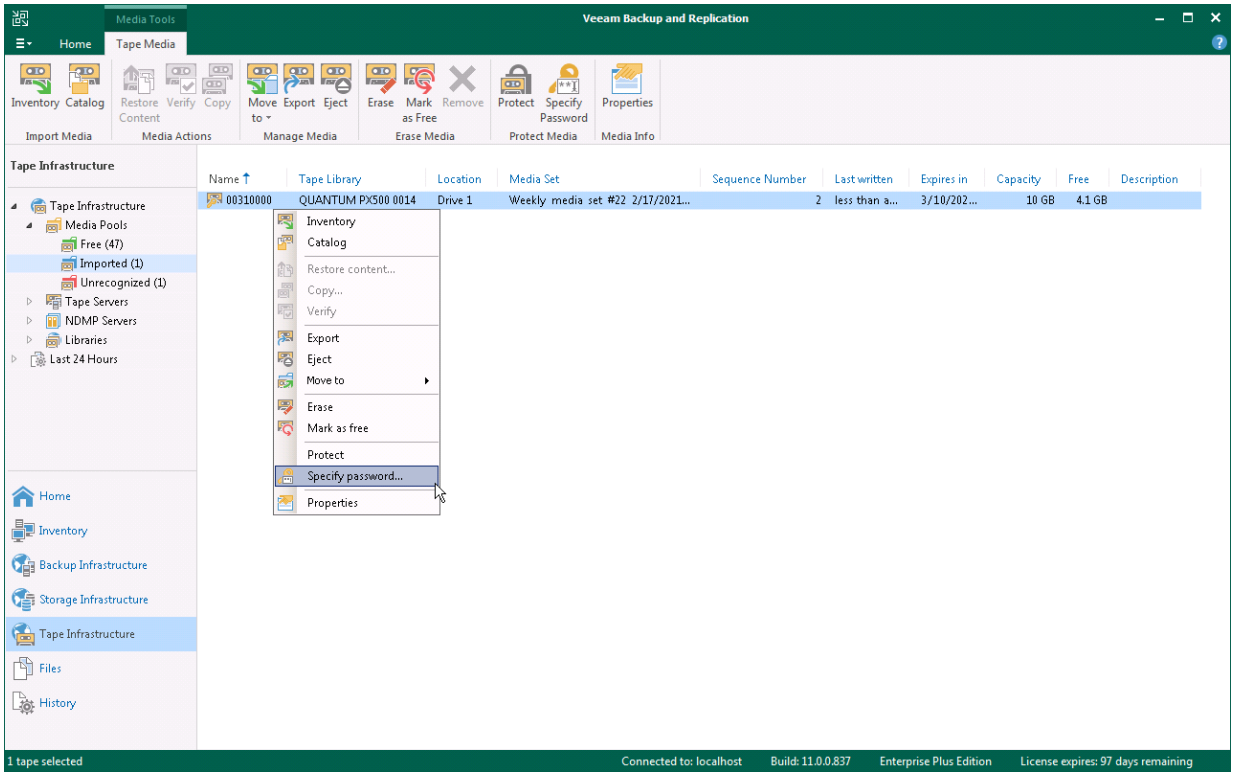
If you import a backup file from tape and the backup file was encrypted twice, with the initial backup job and with the backup to tape job, you must sequentially specify two passwords:

1. Password that was used to encrypt tapes in the media pool.
2. Password for the primary backup job.

After you enter the first password, backups from the tape will be moved under the **Backup** > **Encrypted** node in the inventory pane. You must then enter the second password to decrypt the backup and get access to its content.

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.



Decrypting Tapes Without Password

If you have lost or forgotten a password, you can unlock encrypted tapes with the help of Veeam Backup Enterprise Manager.

You can restore data from tapes without a password only if your backup infrastructure meets the following conditions:

1. You use Veeam Universal License (or a legacy socket-based license, Enterprise or higher edition).
1. Veeam backup server on which you encrypted tapes is added to Veeam Backup Enterprise Manager.
2. Veeam backup server on which you generate a request for data decryption is added to Veeam Backup Enterprise Manager.

The restore process is accomplished with the help of two wizards that run on two servers:

1. The **Encryption Key Restore** wizard on the Veeam backup server.
2. The **Password Recovery** wizard on the Veeam Backup Enterprise Manager server.

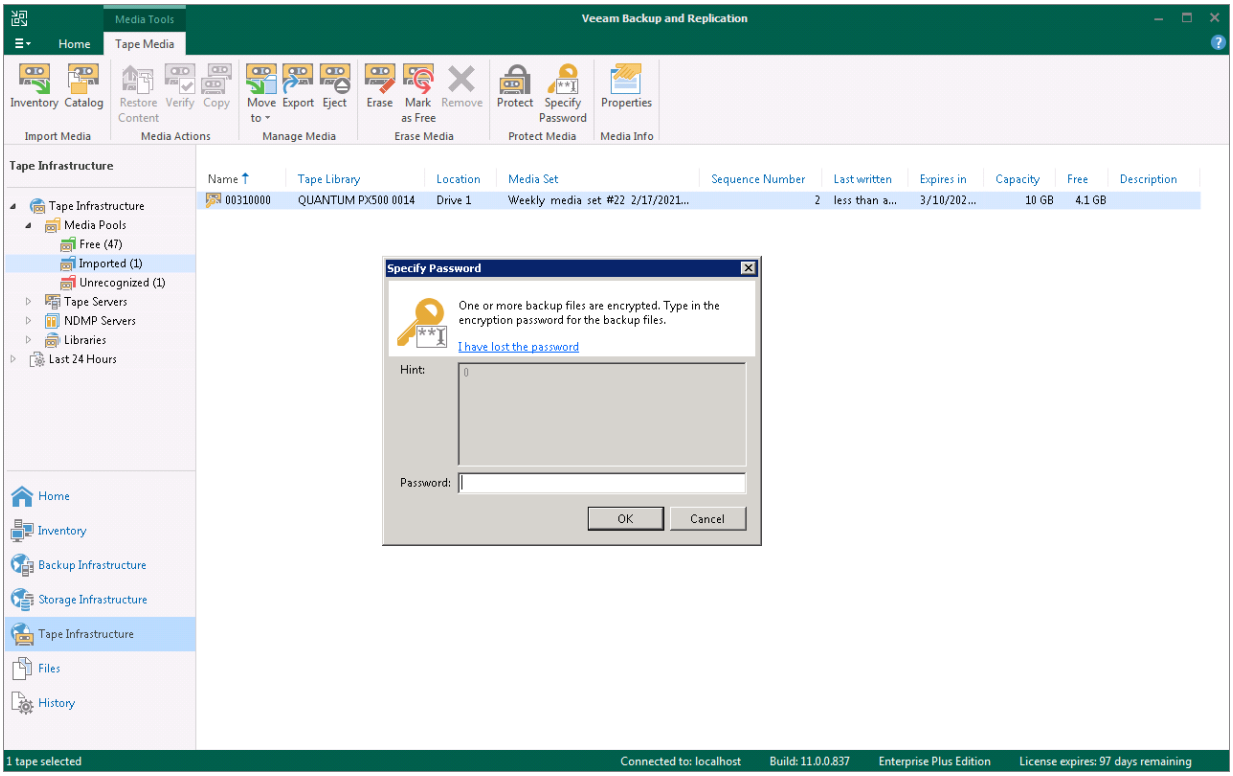
To restore encrypted data from tapes without a password:

1. Create a request for data restore.
2. Process the request in Veeam Backup Enterprise Manager.
3. Complete the key restore process.

Step 1. Create Request for Data Restore

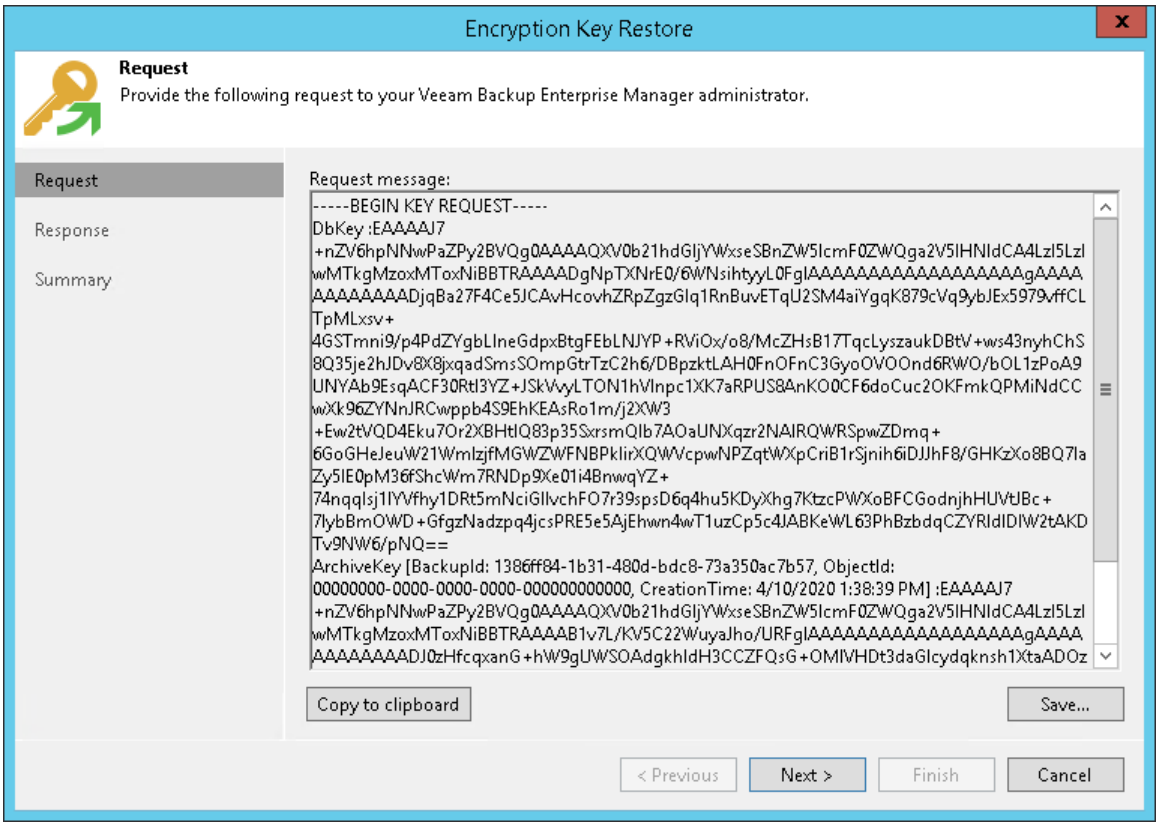
This procedure is performed by the Veeam Backup Administrator on the Veeam backup server.

1. Import encrypted tapes to the Veeam backup server.
2. Select the imported tape and click **Specify Password** on the ribbon or right-click the tape and select **Specify password**.
3. In the **Specify Password** window, click the **I have lost the password** link.



4. Veeam Backup & Replication will launch the **Encryption Key Restore** wizard. At the **Request** step of the wizard, review the generated request for data recovery. Use buttons at the bottom of the wizard to copy the request to the Clipboard or save the request to a text file.

5. Send the copied request by email or pass it in any other way to the Veeam Backup Enterprise Manager Administrator.



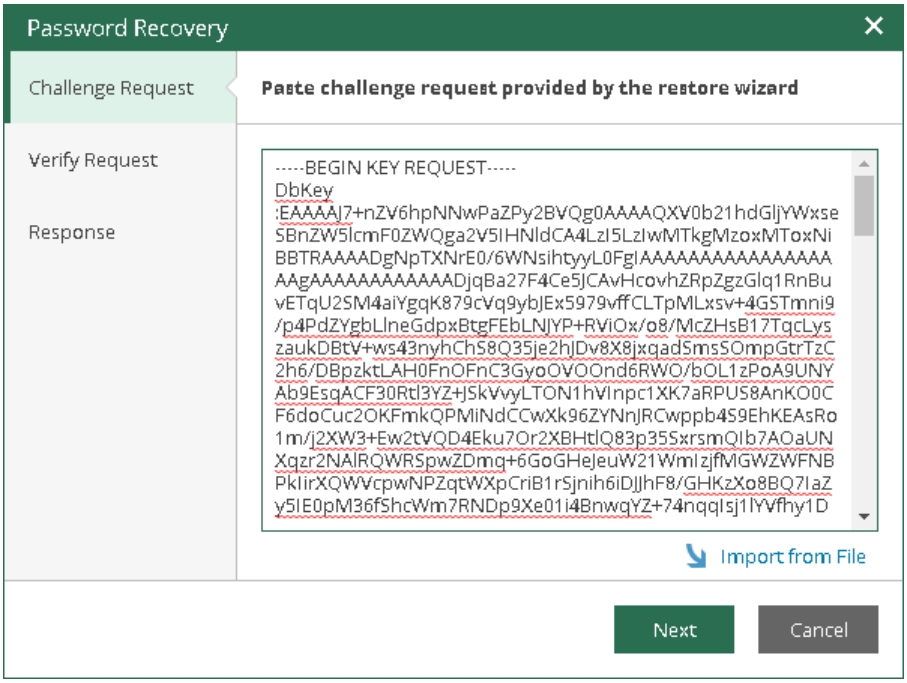
TIP

You can close the **Encryption Key Restore** wizard on the Veeam backup server and start it anew when you receive a response from the Veeam Backup Enterprise Manager Administrator.

Step 2. Process Request in Veeam Backup Enterprise Manager

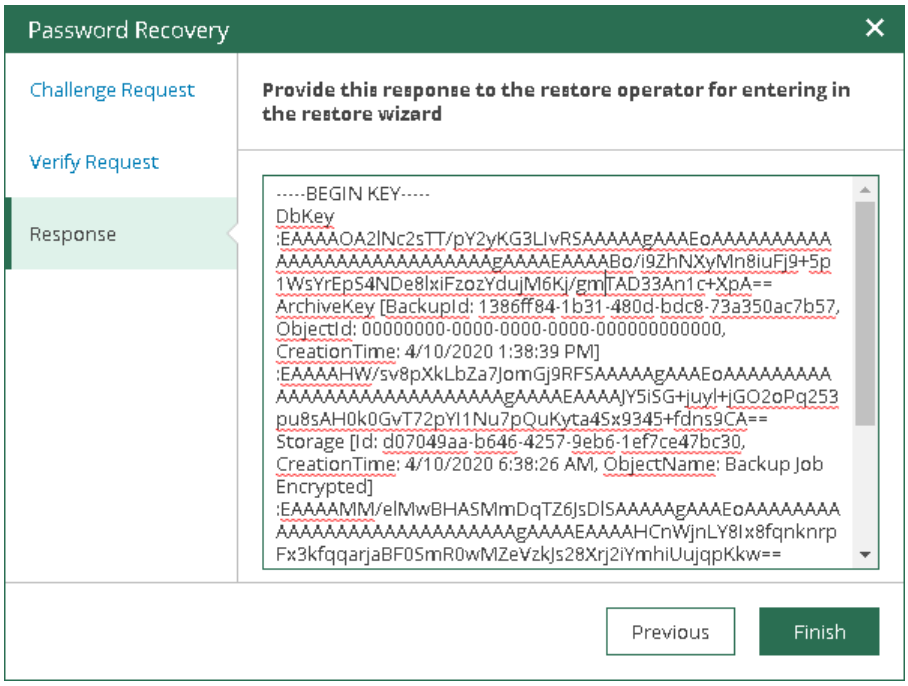
This procedure is performed by the Veeam Backup Enterprise Manager Administrator on the Veeam Backup Enterprise Manager server.

1. Copy the obtained request to the Clipboard.
2. In Veeam Backup Enterprise Manager, go to the **Configuration** > **Settings** > **Key Management** section.
3. Click **Password Recovery** to open the **Password Recovery** wizard.
4. Use the [CTRL+V] key combination to paste the request that you have received from the Veeam Backup Administrator. You can also use the **Import from File** link to import the request from a text file.



5. Follow the next steps of the wizard. At the **Response** step, copy the text displayed in the wizard to the Clipboard.

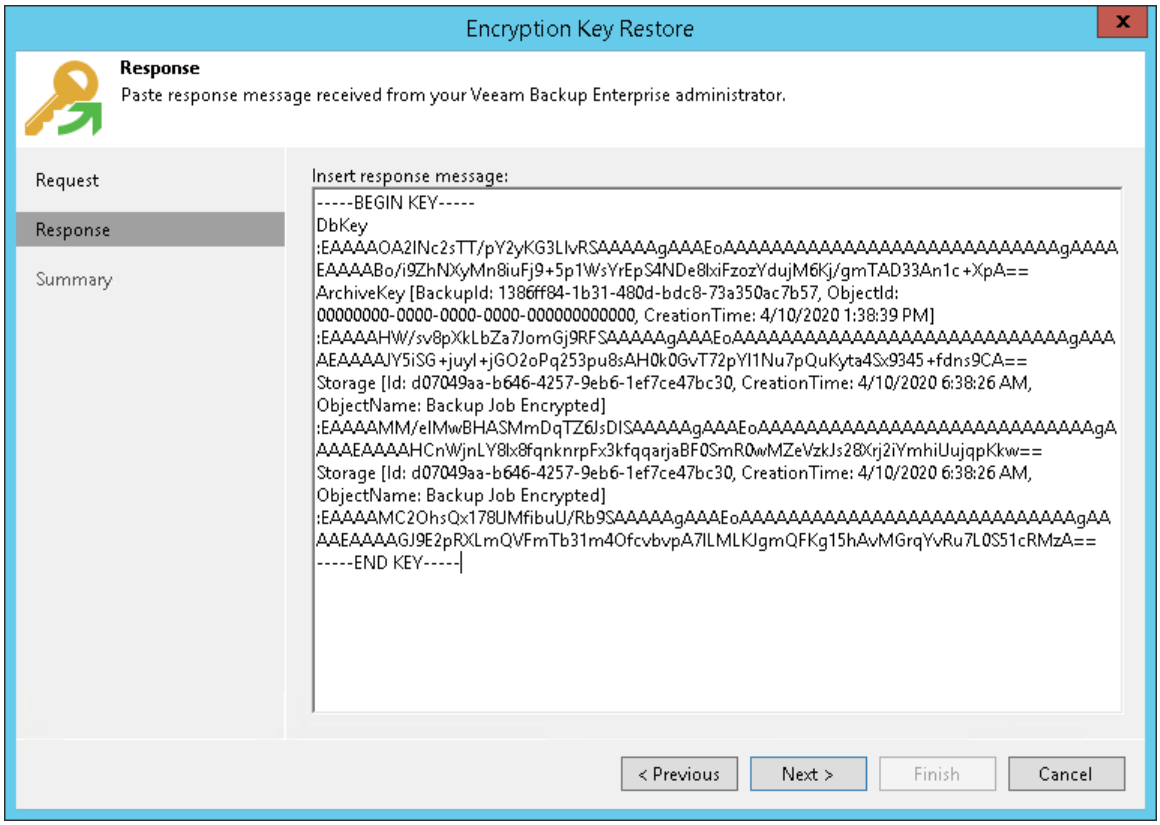
6. Send the copied response by email or pass it in any other way to the Veeam Backup Administrator working on the Veeam backup server.



Step 3. Complete Key Restore Process

This procedure is performed by the Veeam Backup Administrator on the Veeam backup server.

1. In Veeam Backup & Replication, get back to the **Encryption Key Restore** wizard.
2. Enter the copied response to the text window at the **Response** step of the **Encryption Key Restore** wizard.
3. Follow the next steps of the wizard. At the last step, click **Finish**. Veeam Backup & Replication will retrieve the decrypted storage keys from the response, apply them to the encrypted tape and unlock the tape content.



Tape Devices Support

Veeam provides native tape support that is fully integrated into Veeam Backup & Replication. You can administer all operations on tapes from your Veeam console.

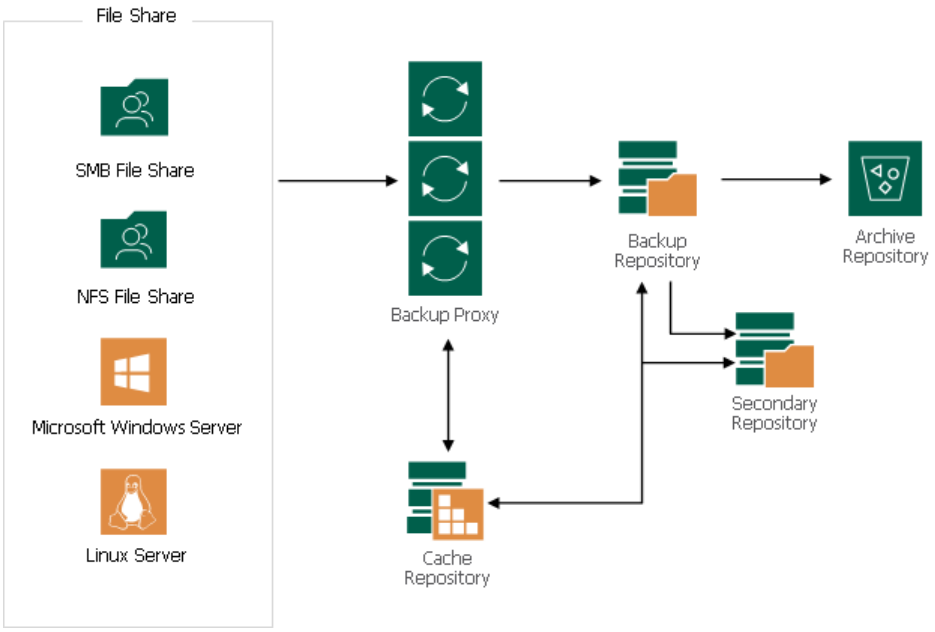
For more information, see the Tape Device Support Guide.

NAS Backup Support

With Veeam Backup & Replication you can easily back up and restore content of various NAS file shares. The solution can be flexibly scaled to reliably protect massive amounts of data, even for the largest of enterprise organizations.

To protect your NAS file shares, you can use your existing Veeam Backup & Replication infrastructure. Just configure the following components:

- File share
- Backup proxy
- Cache repository
- Storage repositories



For system requirements for NAS backup components, see System Requirements.

To learn how NAS backup components interact during file share backup, see How File Share Backup Works.

File Share

A file share is a storage device or data source available to multiple hosts through a computer network.

For supported file shares, see Platform Support.

File share backup jobs in Veeam Backup & Replication can read data from the following sources:

- SMB (CIFS) path
- NFS path
- Path to the storage snapshot folder
- VSS snapshot

NOTE

Mind the following limitations:

- Reading from VSS snapshots on SMB shares is available only under certain conditions, listed in this Veeam KB article.
- Reading from VSS snapshots on DFS shares is not supported.

To learn how to add NAS file shares to the backup infrastructure, see Adding File Share.

Backup Proxy

A backup proxy is an architecture component that sits between the file share and other components of the backup infrastructure. The backup proxy operates as a data mover that transfers data between the source file share and the backup repository. The backup proxy processes jobs and delivers backup and restore traffic.

You can assign the role of a backup proxy to any Windows-managed server added to your Veeam Backup & Replication infrastructure. By default, this role is assigned to the backup server. But this option is sufficient only for small installations where all components are located in the same network segment. For larger installations with larger workload, assign the role of a backup proxy to a dedicated server, as described in Adding Backup Proxy. After that, choose this backup proxy to process the backup traffic from file shares, as described in Adding NFS File Share and Adding SMB File Share.

To optimize performance of several concurrent tasks, you can use several backup proxies. In this case, Veeam Backup & Replication will distribute the backup or restore workload between available backup proxies on per-task basis, taking into account proxy connectivity and their current load. You can deploy backup proxies in the primary site or in remote sites. Mind that using a remote backup proxy may reduce backup performance and is not recommended. To minimize the network load during backup, locate the backup proxy closer to the source file share in the computer network: at the best they should be one hop away from each other.

Cache Repository

A cache repository is a storage location where Veeam Backup & Replication keeps temporary metadata and uses it to reduce the load on the file share during the backup procedure. The cache repository keeps track of all objects that have changed between each backup session. This allows performing incremental backups from the file share super fast and very efficiently.

You can assign the role of a cache repository to a simple backup repository added to the Veeam Backup & Replication infrastructure. To assign this role, select the backup repository as a cache repository, as described in Specify File Share Processing Settings.

To minimize the network load during backup, locate the cache repository closer to the backup proxy in the computer network: at the best, they should be located on one machine.

Storage Repositories

A **backup repository** is a main storage location where Veeam Backup & Replication keeps all versions of backed up files for the configured period and metadata files. Backups stored in the backup repository can be used to quickly restore the entire file share to the state as of a specific restore point.

[Optional] If you want to retain specific files for a longer period of time, you can use cheaper devices for archive purposes. To enable file archiving, configure Veeam Backup & Replication to move backup files and metadata files from the backup repository to an **archive repository**. By default, usage of the archive repository is disabled and, after the retention period for the backup repository is over, backup files are deleted.

[Optional] If you want to store a copy of the file share backup in a different repository, you can configure a **secondary repository** where Veeam Backup & Replication will copy all backups created in the backup repository. The secondary repository can have its own retention policy and encryption settings. By default, no secondary repository is configured.

The table below describes which roles can be assigned to the different storage types.

Storage Type	Backup Repository	Archive Repository	Secondary Repository
Microsoft Windows server	●	●	●
Linux server [1]	●	●	●
SMB (CIFS) share [2]	●	●	●
NFS share	●	●	●
Dell EMC Data Domain	●	●	●
ExaGrid	●	●	●
Quantum DXi	●	●	●
HPE StoreOnce [3]	●	●	●
Scale-out backup repository (SOBR) [4]	●	○	●
Object storage repository [5]	○	●	○
Repository with rotated drives	○	○	○
Veeam Cloud Connect repository	○	○	○

[1] You can store NAS backup files and backup copy files in a hardened repository with immutability, but these files will not be immutable.

[2] If you use a Dell EMC Isilon/PowerScale storage system in the **CIFS Share Access** mode, make sure that you have assigned your service account to the built-in **BackupAdmin** role within Isilon. Otherwise, the access to the share will be denied.

[3] If you plan to use HPE StoreOnce storage appliances, 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.

[4] 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.

[5] The following object storages are not supported for NAS backup: Amazon S3 Glacier, Amazon S3 Snowball Edge, Azure Blob storage Archive Tier, Azure Databox.

You can create two object storage repositories pointing to the same cloud folder/bucket and use these repositories for storing both NAS backups and Capacity Tier backups at the same time: one object storage repository will be used to store NAS backups, the other one – to store virtual and physical machine backups as a capacity tier in a single SOBR. However, these object storage repositories (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 inevitable data loss.

Deployment of Backup, Archive and Secondary Repositories

To use a storage, which is already added to the Veeam Backup & Replication Backup Infrastructure, to store file share backups, define it as a target storage when creating a file share backup job:

- Configure the backup repository and the archive repository as described in Define Target Backup Storage Settings.
- Configure the secondary repository as described in Specify Secondary Repository.

If the required storage is not added as a backup repository in your Veeam Backup & Replication Backup Infrastructure, add it as described in Adding Backup Repositories.

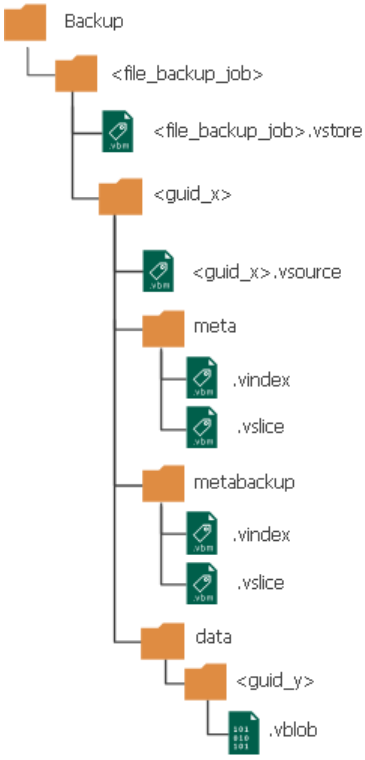
Data Structure in Backup, Archive and Secondary Repositories

Backup, archive and secondary repositories store file share backups as objects in VBLOB files (up to 64 MB each) of unstructured data. They also store metadata about the original files and folders on the file share and their structure.

When you run the file share backup job for the first time, Veeam Backup & Replication creates a full backup for all files and folders of the file share. During subsequent backup job sessions, Veeam Backup & Replication copies only files and folders that have changed since the last backup job session.

Although the file share backup job first creates a full backup and afterwards incremental backups, Veeam Backup & Replication does not create a separate file for each backup job run as it does during VM backup. Instead of this, it consistently creates multiple VBLOB files accompanied by metadata files that track all the changes on the file share.

Veeam Backup & Replication uses the following structure for storing the file share backup data in the backup repository:



Folder/File						Description
Backup						Folder in the repository that is dedicated to store backups.
	<file_backup_job>					Folder that contains all data backed up by a specific file share backup job. Data for each job will be placed to its own directory.
		<file_backup_job>.vstore				XML metadata file that describes the entire backup file for a specific file share backup job.
		<guid_x>				Folder that contains all data for a single source. Data for each source will be placed to its own directory.

Folder/File						Description
			<guid_x>.vsource			XML metadata file that describes the single source object.
			meta			Folder with binary metadata files that describe the content of the backup. If backup files are stored on a scale-out backup repository, this folder on each performance extent contains metadata for data stored on the same extent of the scale-out backup repository. For more information, see Scale-Out Repository as NAS Backup Repository.
				.vindex		Binary metadata that describes backup files (names and versions).
				.vslice		Binary metadata that describes allocation of data in VBLOB backup files.
			metabackup			If backup files are stored on a scale-out backup repository, this folder on each performance extent contains replica of metadata stored on other extents of the scale-out backup repository. For more information, see Scale-Out Repository as NAS Backup Repository.
			data			Folder with binary data.
				<guid_y>		1 Gb basket that stores VBLOB backup files.

Folder/File						Description
					.vblob	By default, a classic 64 Mb file that stores data from the file share backup. If necessary, you can convert NAS backup data files into the single file basket format recommended for storing data on HPE StoreOnce storage appliances. For more information, see description of the `Convert-VBRNASBackupStorageFormat` cmdlet in the Veeam PowerShell Reference.

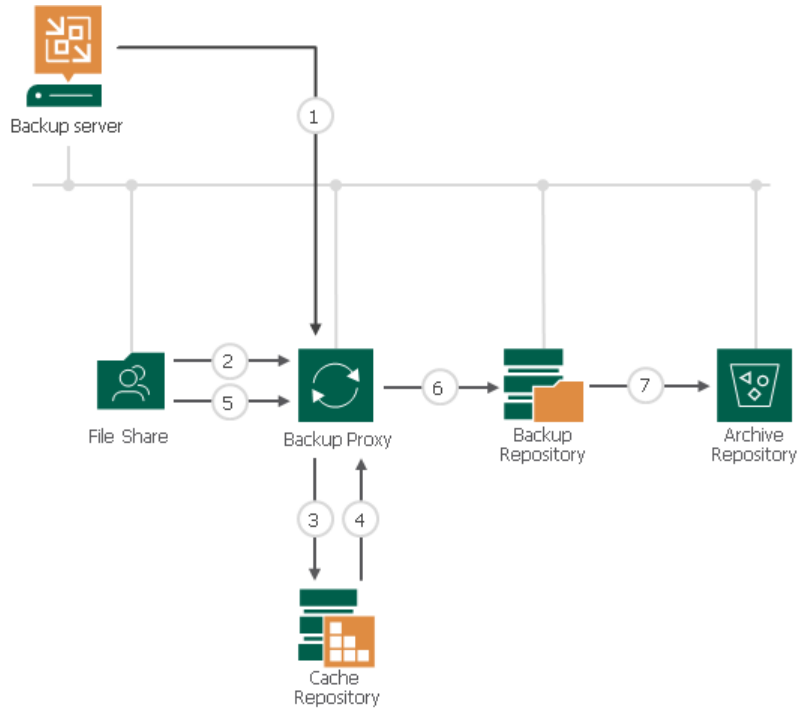
How File Share Backup Works

Veeam Backup & Replication performs file share backup to the backup storage in the following way:

1. When a new backup job session starts, Veeam Backup & Replication assigns a backup proxy to process the file share data.
2. The backup proxy enumerates files and folders on the file share and creates a cyclic redundancy check (CRC) tree.
3. The backup proxy transfers the CRC tree to the cache repository.
4. The cache repository saves the CRC tree.

   When the cache repository receives a new CRC tree structure from the proxy, it compares it with the CRC tree created during the previous run of the backup session. If any files or folders of the file share have changed since the previous backup session run, the cache repository instructs the backup proxy to start reading changed data from the source file share.
5. The backup proxy reads new data from the file share.
6. The backup proxy creates data packages and transfers them to the target backup repository.

   Data packages comprise backup data files (each 64 Mb in size) and metadata files that contain names and versions of backup files and allocation of data in backup files.
7. Veeam Backup & Replication checks file versions in the backup repository against retention settings and moves backup data from the backup repository to the archive repository if necessary.



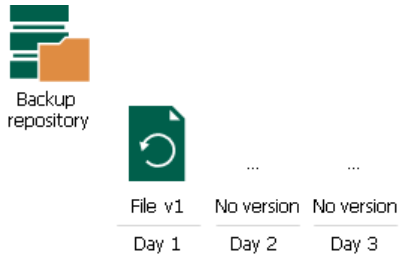
NAS Backup Retention Scenarios

There can be a number of backup retention scenarios depending on the configuration of backup and archive repositories. Below you can find example cases that illustrate NAS backup retention with different settings.

Case 1

Only 1 file version is created. The file does not change.

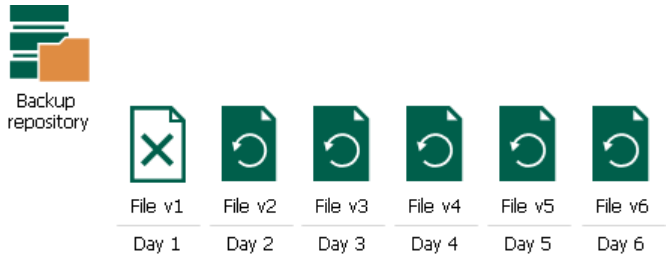
File version 1 always remains in the backup repository and is not moved to the archive repository even if this behavior is enabled and configured in the retention policy settings.



Case 2

Retention for the backup repository is set to 5 days. No archive repository is configured. The file changes once a day. The backup is performed once a day.

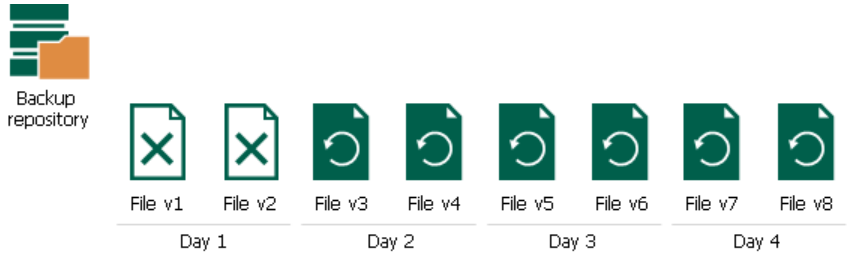
On day 6, file version 6 is added to the backup repository, file version 1 is deleted by retention.



Case 3

Retention for the backup repository is set to 3 days. The file changes every hour. The backup is performed 2 times a day.

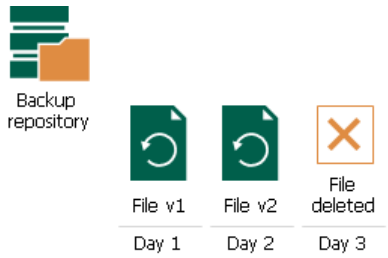
On day 4, versions 7 and 8 are added to the backup repository, file versions 1 and 2 added to the backup repository on day 1 are deleted by retention.



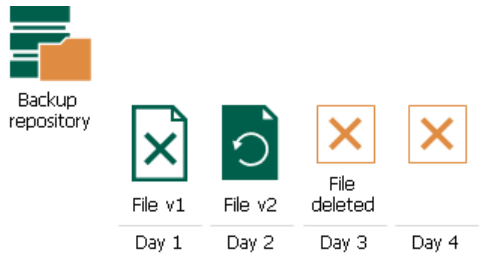
Case 4

Retention for the backup repository is set to 3 days. The file changes once a day.

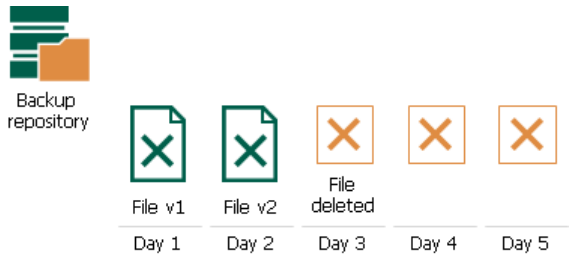
On day 3, the source file is deleted from the source share, the backup repository considers file version created on this day as deleted.



On day 4, the backup repository still detects the file as deleted, file version 1 is deleted from the backup repository by retention.



On day 5, the backup repository still detects the file as deleted, file version 2 is deleted from the backup repository by retention.

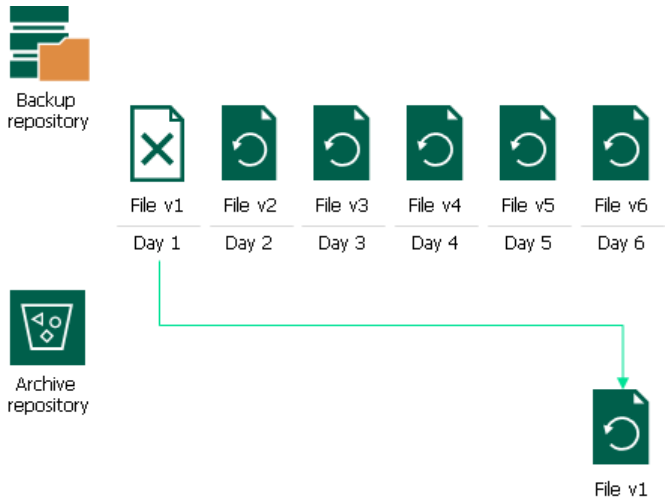


Thus, no file versions are stored in the backup repository for this file any longer.

Case 5

Retention for the backup repository is set to 5 days. The archive repository is enabled with default settings. The file changes every day. The backup is performed once a day.

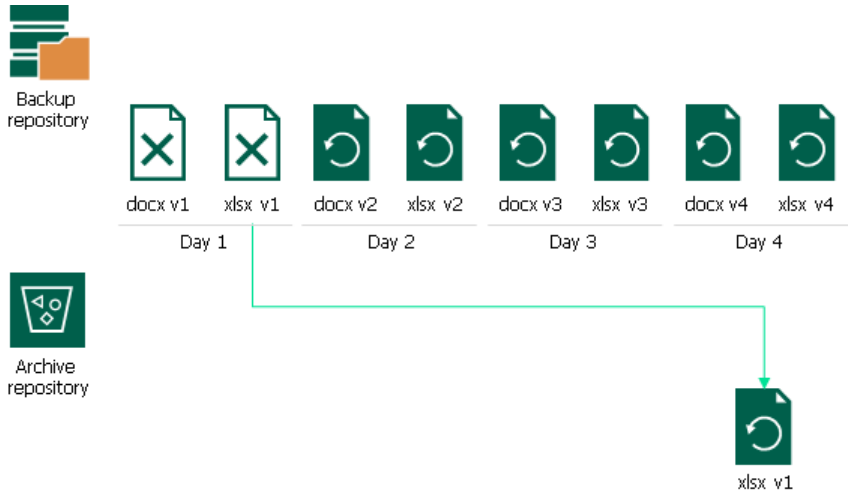
On day 6, file version 6 is added to the backup repository, file version 1 is moved to the archive repository by retention.



Case 6

Retention for the backup repository is set to 3 days. The archive repository is enabled with DOCX files to be excluded from archiving. The files change once a day. The backup is performed once a day.

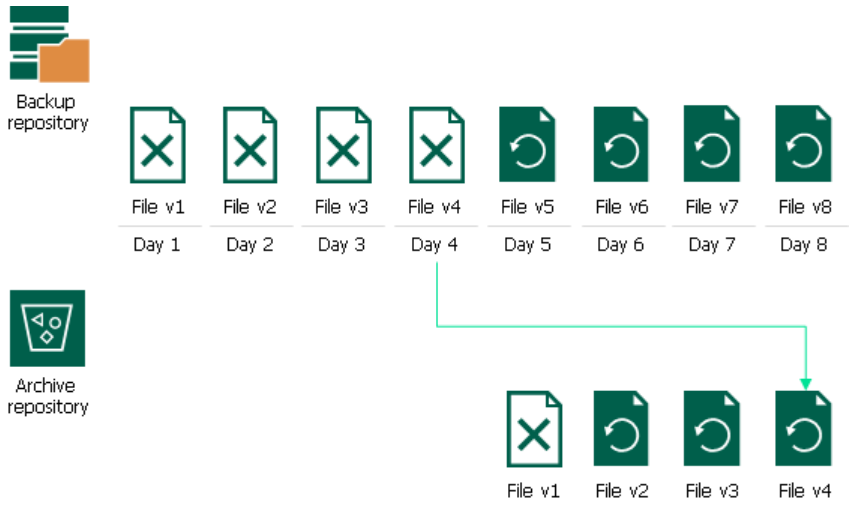
On day 4, file versions created on day 1 are removed from the backup repository. File version 1 for DOCX file is deleted, file version 1 for XLSX file (non-DOCX) is moved to the archive repository.



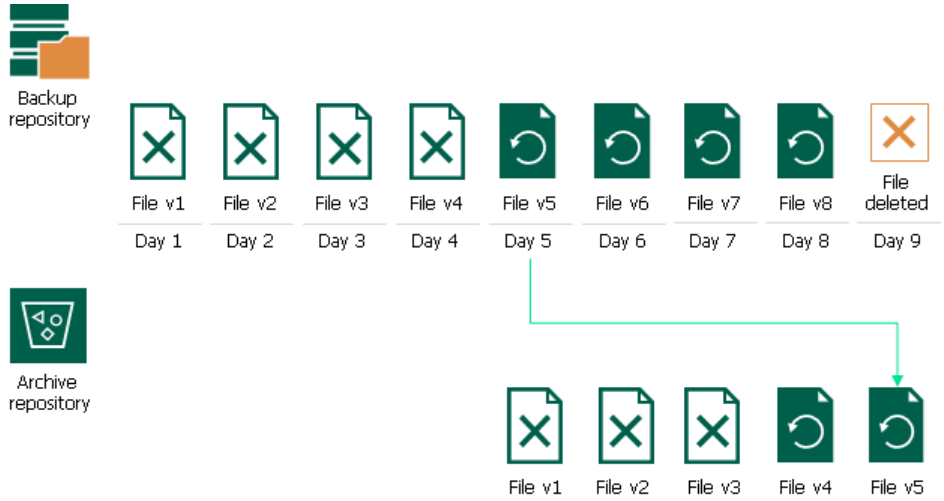
Case 7

Retention for the backup repository is set to 4 days. The archive repository is enabled and configured to keep 3 versions of active files and 2 versions of deleted files.

On day 8, file version 8 is added to the backup repository, file version 4 is moved from the backup repository to the archive repository to keep file versions for 4 days, file version 1 is deleted from the archive repository to keep 3 file versions of the active file (versions 2, 3, 4).

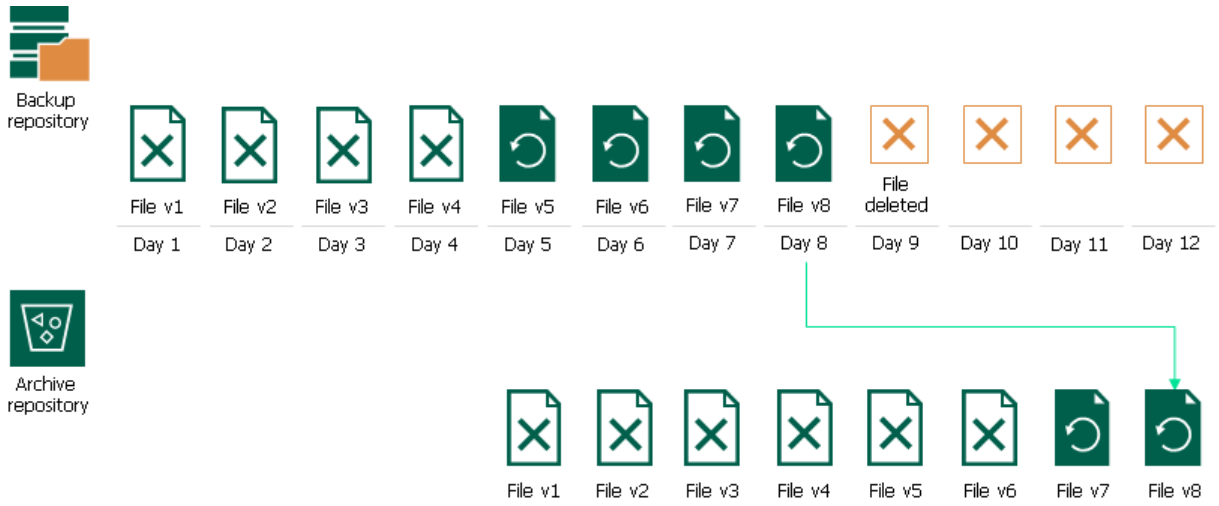


On day 9, the file is removed from the source, file version 9 (denoting the missing file) is added to the backup repository, file version 5 is moved from the backup repository to the archive repository, file versions 2 and 3 are deleted from the archive repository to keep 2 file versions of the deleted file (versions 4 and 5).



On day 10 and 11, file versions 6 and 7 are successively moved from the backup repository to the archive repository. File versions 4 and 5 are deleted from the archive repository.

On day 12, file version 8 (the last file version) is moved from the backup repository to the archive repository, file version 6 is deleted from the archive repository. After that, versions 7 and 8 are stored in the archive repository further on.



NAS Backup Integration with Storage Systems

There are two approaches in backing up file shares residing on enterprise NAS storage systems.

Integration with Storage System as NFS or SMB File Share Server

You can add the storage system as a root folder of the server where NFS or SMB file shares reside. The procedure of configuring the file share protection in this case will look as follows:

1. Add the storage system as an NFS file share to the inventory, as described in Adding NFS File Share, or as an SMB fie share, as described in Adding SMB File Share. As a file share path, specify the root server folder.

   When adding the storage system in this way, you cannot configure what containers, volumes or file shares will be available for further protection. Therefore, to configure them, you must carefully configure inclusion/exclusion settings when creating a file share backup job.

   Mind that servers accessed by NFS (with file shares and folders within them) and servers accessed by SMB (with file shares and folders within them) are added to the inventory separately. For example, if the storage system IP address is 173.25.136.64, add an NFS share for this server by specifying its root folder as `173.25.136.64:/`, and add an SMB share for this server by specifying its root folder as `\\173.25.136.64`.

   > **NOTE**
   >
   > If you add a root server folder as a source for protection, hidden and admin file shares are skipped from processing by default. You can enable their processing with registry keys. For more information, contact Veeam Customer Support.

2. Create a file share backup job, as described in Creating File Share Backup Jobs. As a source to protect, you can select the following entities:

   - whole server
   - file share residing on the server
   - separate folders within the share

   To protect all file shares residing on one server, you must add to the file share backup job both NFS and SMB shares previously added to the inventory.

3. Configure what files and folders must be included in or excluded from processing by the file share backup job. For more information on how to include/exclude files and folders from processing, see Select Files and Folders to Back Up.

Integration with Storage System as NAS Filer

You can add the storage system as a NAS filer. This option is preferable if you want to leverage the backup from storage snapshots technology.

> **NOTE**
>
> If you used to protect NFS and SMB file shares residing on the enterprise storage system and added as file shares in inventory, and now you want to protect them using benefits of NAS filer, you can convert backups created for existing SMB or NFS shares into the format of NAS filer shares. For more information, see Converting Backups from SMB or NFS Shares to NAS Filer Shares.

The procedure of configuring the file share protection in this case will look as follows:

1. Add the storage system to the backup infrastructure, as described in Adding NetApp Data ONTAP, Adding Lenovo ThinkSystem DM Series, or Adding Dell EMC Isilon, depending on the type of the NAS system you use.

   Depending on storage settings, the IP address for accessing the storage system can differ from one used for accessing it as a server where file shares reside. You can also use the DNS name of the server.

   When adding the storage system, make sure that you do not forget to perform the following steps:

   a. Enable the NAS filer role for the added storage system.

   b. Specify what protocols the storage should use as a NAS filer: NFS and/or SMB. Only file shares using the selected protocols will be displayed when you add the storage as a NAS filer and thus available for protection.

   c. Select storage volumes to analyze for the presence of newly added file shares. You can either configure Veeam Backup & Replication to analyze all storage volumes, or exclude some volumes from processing, or specify only certain volumes that will be processed. Only file shares on the selected storage volumes will be displayed when you add the storage as a NAS filer and thus available for protection.

   At this step, you must carefully consider what file shares on what volumes must be protected and via what protocols. Limiting the number of volumes reduces the storage load.

   > **NOTE**
   >
   > Starting from Veeam Backup & Replication 11a (build 11.0.1.1261), hidden and admin file shares on storage systems added as NAS filers are skipped from processing by default. You can enable their processing with a registry key. For more information, contact Veeam Customer Support.

2. Add the configured storage system as a NAS filer to the inventory, as described in Adding Enterprise NAS System Share.

3. Create a file share backup job, as described in Creating File Share Backup Jobs. As a source to protect, you can select the following entities:

   - whole storage
   - container (for Dell EMC Isilon/PowerScale — access zone, for NetApp Data ONTAP — SVM)
   - volume
   - file share

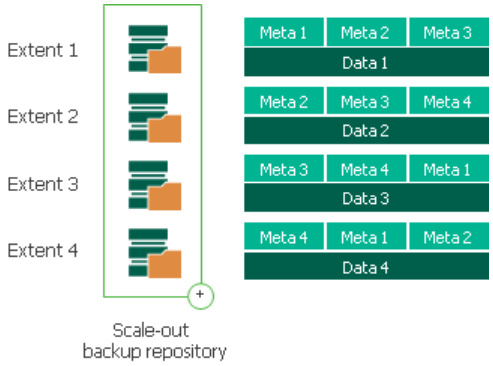
   You cannot specify separate folders within file shares. Therefore, to configure files and folders to be protected, you must properly configure inclusion/exclusion settings.

4. Configure what files and folders must be included in or excluded from processing by the file share backup job. For more information on how to include/exclude files and folders from processing, see Select Files and Folders to Back Up.

Scale-Out Repository as NAS Backup Repository

If you use a scale-out backup repository as a backup repository for storing NAS backups, by default Veeam Backup & Replication evenly distributes backup data among all the performance extents added to this repository. The NAS backup data is accompanied by the metadata files. To provide an opportunity to restore data written to extents remaining available in case one or two of other extents are unavailable, Veeam Backup & Replication triplicates metadata when writing it to extents. Thus, every extent keeps metadata for data stored on this extent (in folder **meta**) and replica of metadata for data stored on adjacent extents (in folder **metabackup**). Even if two of three adjacent extents are lost, you can restore NAS backup data stored on the third extent.

> **NOTE**
>
> Mind the following when storing NAS backups on a scale-out backup repository:
>
> - The metadata redundancy approach protects against two extents out of three being completely lost.
> - The metadata redundancy approach allows protecting metadata, not data.
> - The backup file placement policy set for the scale-out backup repository, as described in Backup File Placement, is ignored.



Scale-Out Repository with Extents in Metadata and Data Roles

When you use a standalone repository for storing NAS backups, it stores both data and metadata. But when you plan to use a scale-out repository, you can configure its performance extents to act as data extents or as metadata extents. Thus, you can store metadata separate from the backup data.

> **NOTE**
>
> We strongly recommend to differentiate metadata and data roles for extents, when you plan to create a scale-out backup repository that will consist of a fast SSD storage and slow deduplicating storage appliances (Dell EMC Data Domain, ExaGrid, HPE StoreOnce, Quantum DXi). In this case, you can set the metadata role to the SSD storage and the data role to deduplicating storage appliances.

Most often, when performing restore, merge, transform operations, Veeam Backup & Replication interacts with metadata rather than with the backup data. Obviously, processing metadata stored on the SSD is faster and more efficient than accessing large data arrays stored on a slow storage.

> **IMPORTANT**
>
> Mind the following when assigning roles to extents in a scale-out backup repository:
>
> - An extent with the metadata role can be used for storing NAS backup jobs metadata only. However, extents with the data role can be used by any job.
> - Make sure that you assign both roles to extents in a scale-out repository. If an extent with one of the roles is missing, Veeam Backup & Replication cannot store backups on this scale-out repository.

> **NOTE**
>
> If the data role is assigned to an extent, Veeam Backup & Replication will also copy replica of metadata to this extent to provide the metadata redundancy. While the original metadata is available, Veeam Backup & Replication does not use the replica of metadata on data extents. If by some reason metadata stored on metadata extents is corrupted or lost, to restore it Veeam Backup & Replication will use the replica of metadata stored on data extents.

To assign the metadata or data role to extents in a scale-out backup repository, use the `Set-VBRRepositoryExtent` cmdlet, as described in the Veeam PowerShell Reference. If previously the role was not assigned to the extents for NAS backup or you changed the assigned role, during the next run of the NAS backup job that writes backups to this scale-out repository, Veeam Backup & Replication will move metadata to the metadata extent, data — to the data extents.

For example, if you already have an existing backup and want to move its metadata to a specific extent that will operate as a metadata only extent, do the following:

1. Make sure that your license allows using a scale-out backup repository with object storage support. For more information, see Viewing License Information and Veeam Backup & Replication Feature Comparison.
2. Create a scale-out backup repository with the following extents: one extent that currently stores the backup (its data and metadata), another extent that will store metadata of the backup. Usually, it is a fast storage, for example, SSD-based. Let us assume that these extents are named “Backup Repository 1” and "NAS Backup Repository on SSD" respectively.

3. Run the `Set-VBRRepositoryExtent` cmdlet to assign the data role to the "Backup Repository 1" extent and the metadata role to the "NAS Backup Repository on SSD" extent.

```
Set-VBRRepositoryExtent –Extent "NAS Backup Repository on SSD" -Metadata
Set-VBRRepositoryExtent –Extent "Backup Repository 1" –Data
```

   For more information, see the Veeam PowerShell Reference.

4. Run the NAS backup job and make sure that the metadata of the backup was moved to the metadata extent: the backup job session displays a line notifying of that.

To view the roles of the extents in a scale-out backup repository, do either of the following:

- Check the role of each extent (the **Role** column) in the list of extents under the certain scale-out repository in the **Backup Infrastructure** view.
- Run the `Get-VBRRepositoryExtent` cmdlet, as described in the Veeam PowerShell Reference.

Adding File Share

You must add to the backup infrastructure file shares that you plan to use as sources for backup.

You can add file shares of the following types:

- Windows-managed or Linux-managed server
- Enterprise NAS system
- NFS file share
- SMB file share

Adding Managed Server File Share

Before you add a Windows- or Linux-managed server as a file share to the inventory of the virtual infrastructure, mind the following:

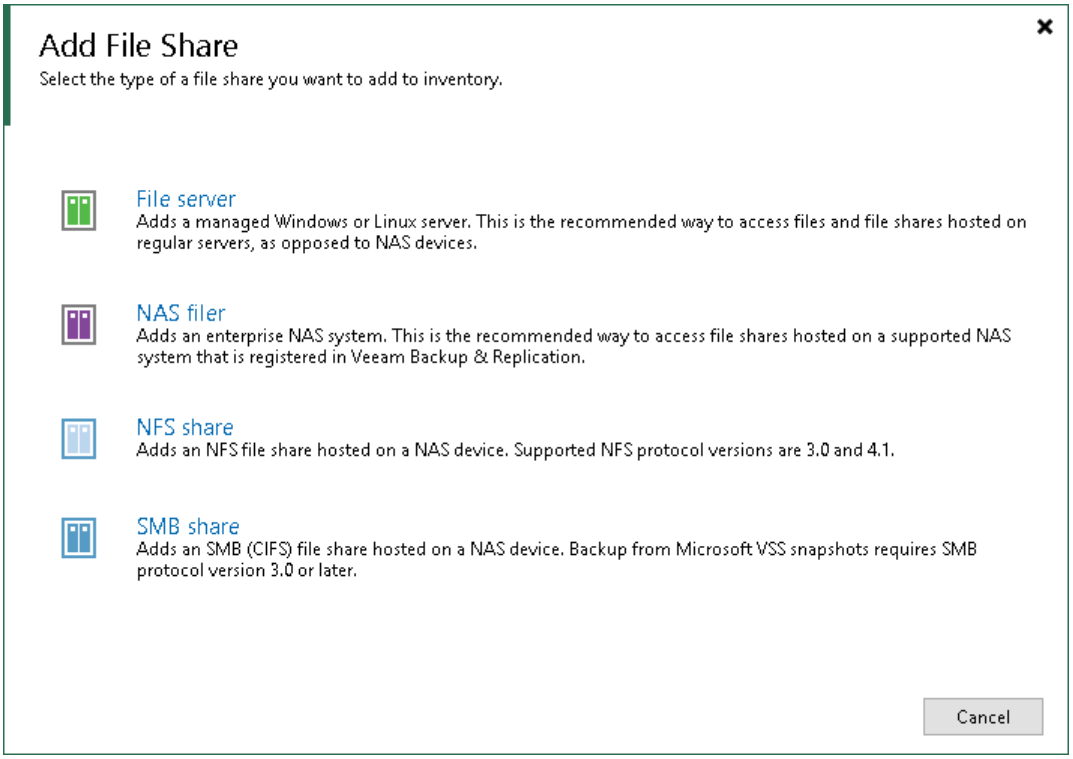
- This server must meet requirements listed in Platform Support.
- You must have this server added in **Backup Infrastructure**.

  For more information on how to add servers, see Adding Microsoft Windows Servers and Adding Linux Servers.
- If you plan to use a devoted cache repository, make sure it is added in **Backup Infrastructure**.
- Data from managed server file shares are transferred directly to the repository without a proxy server.

Step 1. Launch New File Share Wizard

To launch the **New File Share** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **File Shares** node and select **Add File Share**.
   - Select the **File Shares** node and click **Add File Share** on the ribbon.
   - Select the **File Shares** node and click **Add File Share** in the working area.
3. In the **Add File Share** window, click **File server**.



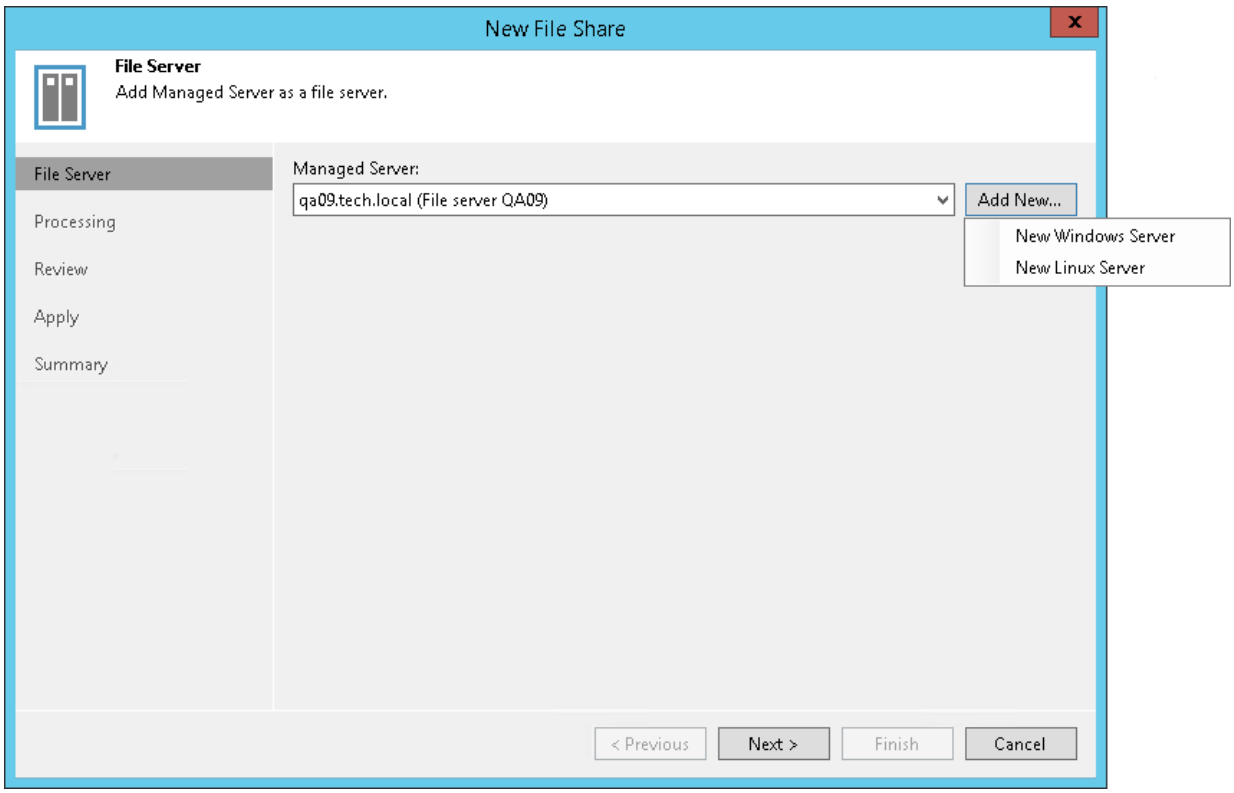
Step 2. Add Managed Server as File Server

At the **File Server** step of the wizard, choose the server, which you want to use as a file share server. Select it from the **Managed Server** drop-down list.

> **NOTE**
>
> If you plan not only to back up the Linux-managed file share, but also to restore files to it, use an account with root access when adding the server to the backup infrastructure.

If the drop-down list does not display the required server, you must add it to the backup infrastructure. To add the server, do the following:

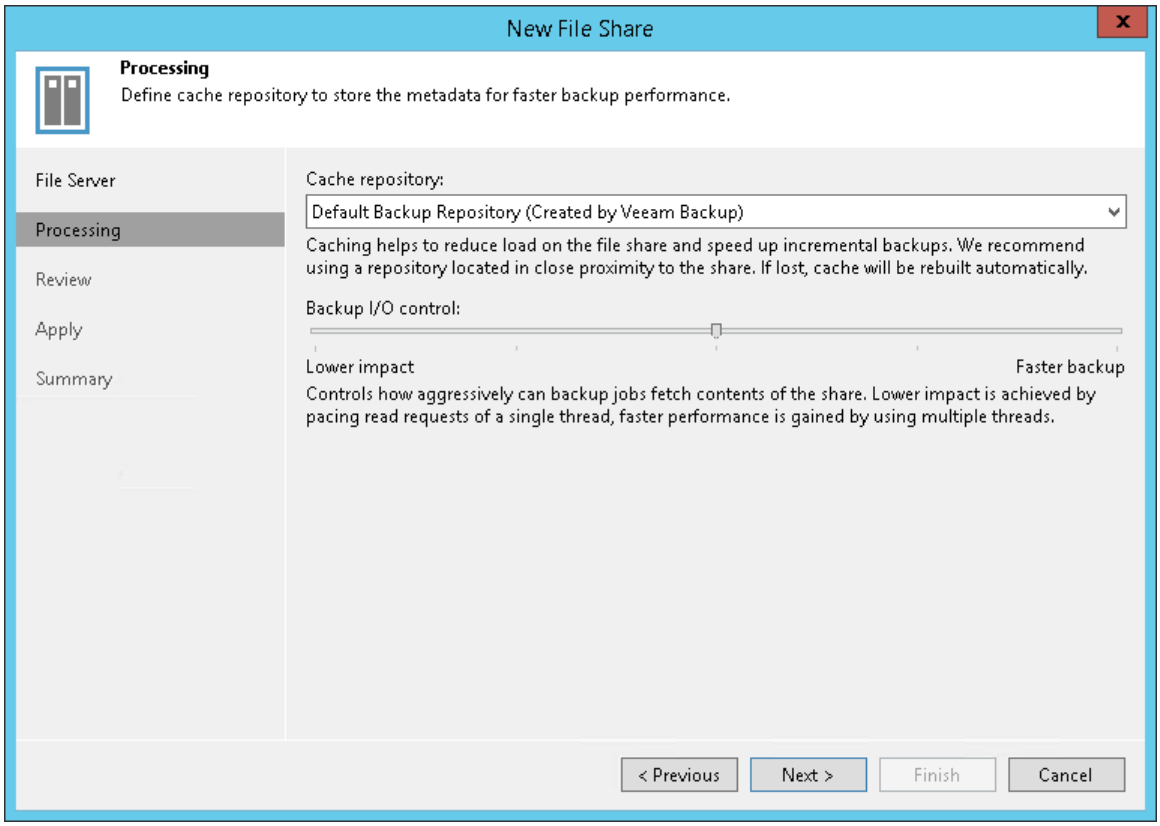
1. Click **Add New**.
2. Select **New Windows Server** or **New Linux Server**.
3. Add a new Windows or Linux server to the backup infrastructure as described in Adding Microsoft Windows Servers and Adding Linux Servers.
4. Select the newly added server from the **Managed Server** drop-down list.



Step 3. Specify File Share Processing Settings

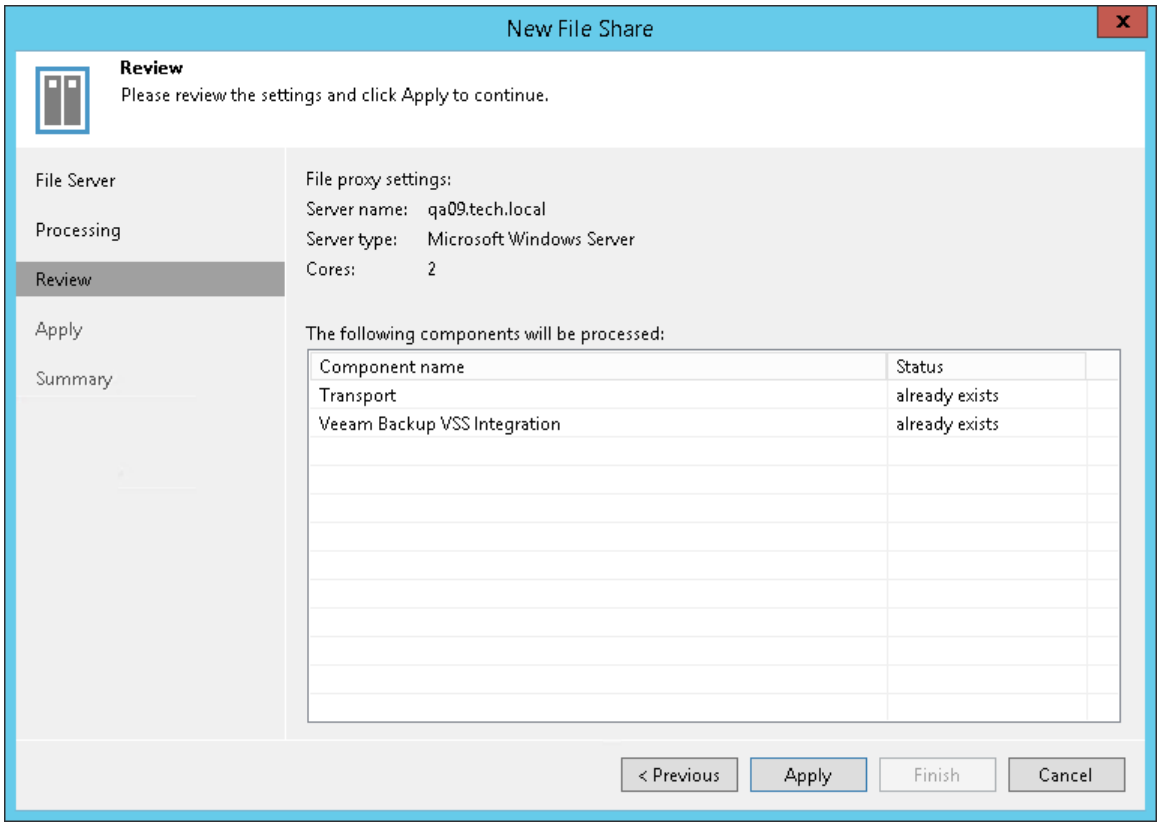
At the **Processing** step of the wizard, define file share processing settings:

1. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located as close to the source file share as possible.
2. Use the **Backup I/O control** slider to define how fast the backup proxy can read data from the source file share. This setting is based on the number of parallel threads that can be used by the proxy configured for processing the file share. If resources of your NAS device are limited, it is recommended that you select the **Lower impact** option. If your NAS device is powerful enough, select the **Faster backup** option.
3. Click **Next** to save the configure settings.



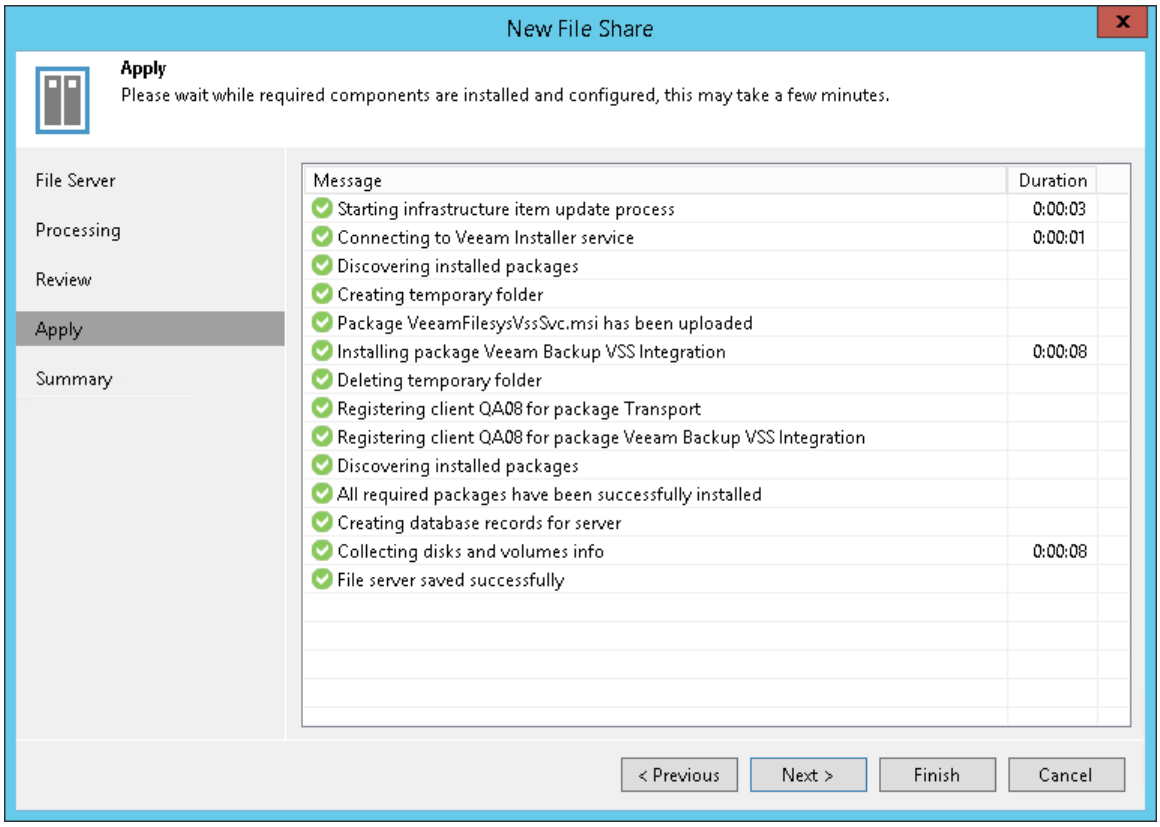
Step 4. Review Components to Install

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and click **Apply** to start installation of missing components.



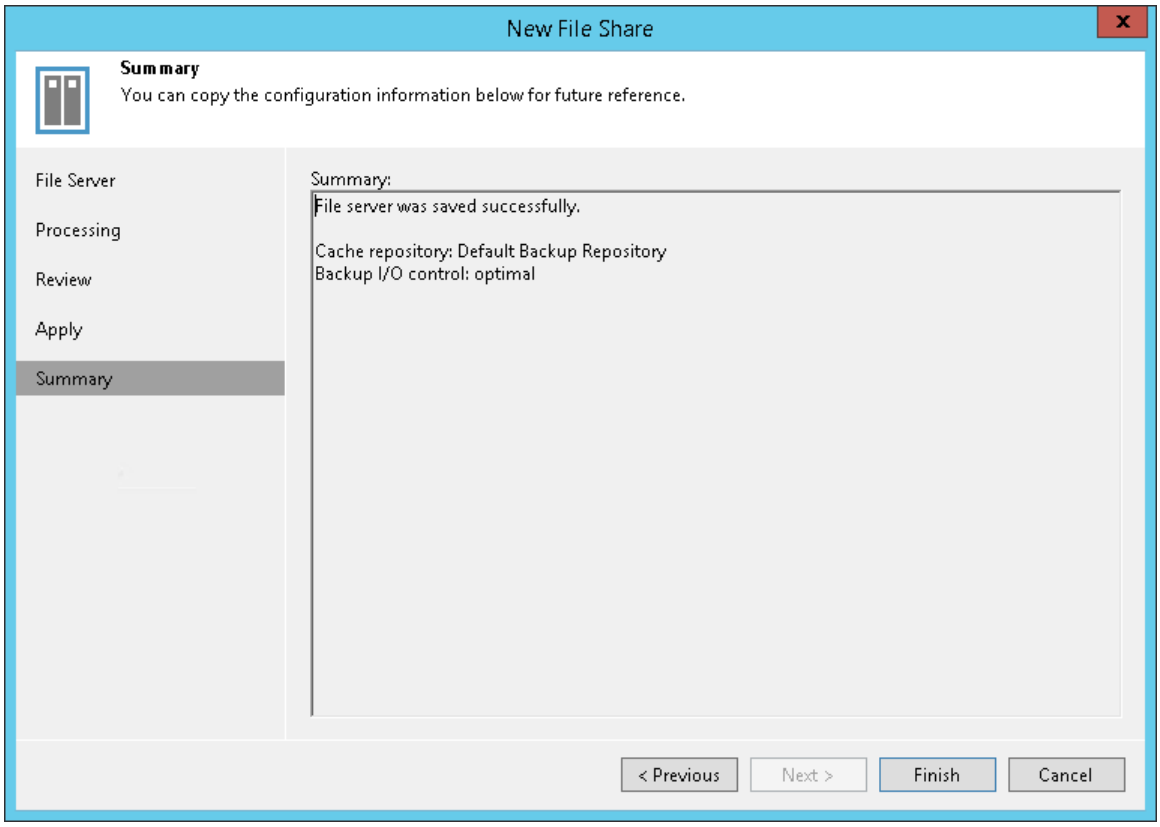
Step 5. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of the file share role assignment to the managed file server.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the file server added as a file share and click **Finish to exit the wizard**.



Adding Enterprise Storage System as NAS Filer

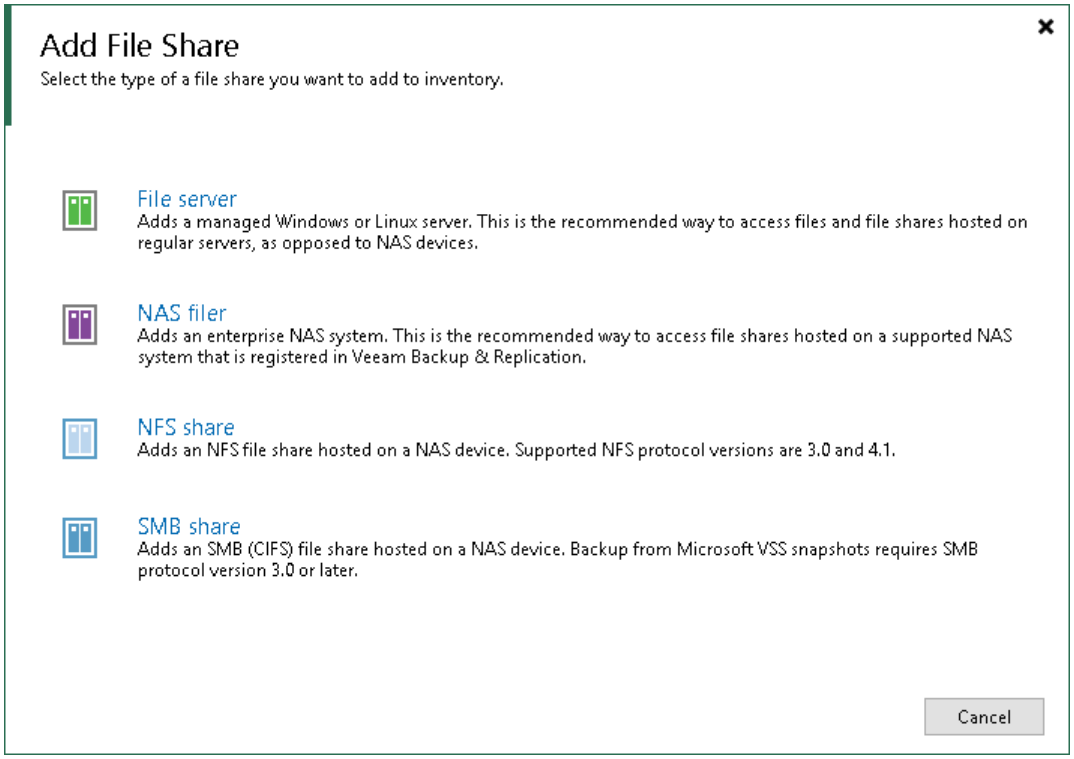
Before you add an enterprise storage system as a NAS filer to the inventory of the virtual infrastructure, check the following prerequisites:

- The NAS device meets requirements listed in Platform Support.
- If you plan to use a devoted backup proxy server or cache repository, make sure these components are added in Backup Infrastructure.

Step 1. Launch New File Share Wizard

To launch the **New File Share** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **File Shares** node and select **Add File Share**.
   - Select the **File Shares** node and click **Add File Share** on the ribbon.
   - Select the **File Shares** node and click **Add File Share** in the working area.
3. In the **Add File Share** window, click **NAS filer**.



Step 2. Select NAS Device

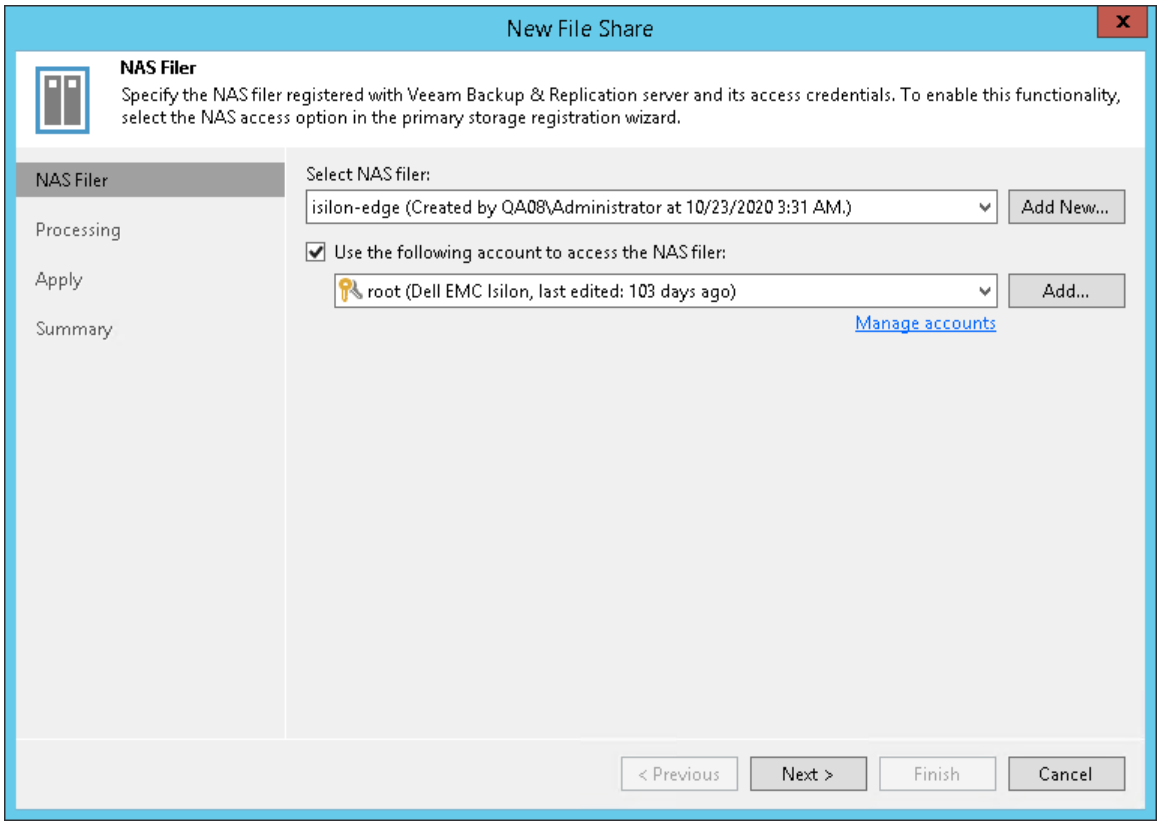
At the **NAS Filer** step of the wizard, choose the NAS device, which you want to use as a NAS filer where protected file shares reside.

1. Select the required NAS device from the **Select NAS filer** drop-down list.

   If the drop-down list does not display the required device, you must add it to the storage infrastructure. To add the NAS device, click **Add New** and follow the instructions described in Adding NetApp Data ONTAP, Adding Lenovo ThinkSystem DM Series, or Adding Dell EMC Isilon, depending on the type of the storage system you use.

2. If you must specify user credentials to access the storage system, select the **Use the following account to access the NAS filer** check box. From the **Credentials** drop-down list, select a credentials record for a user account that has **Full Control** permissions on the storage system.

   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.



Integration with Dell EMC Isilon/PowerScale

To provide proper integration of Veeam Backup & Replication with Dell EMC Isilon/PowerScale to implement the NAS backup functionality, mind the following:

- The Dell EMC Isilon/PowerScale system must be licensed for SnapshotIQ. Otherwise NAS integration will not work.
- The Isilon service account is used to add the Dell EMC Isilon/PowerScale system to Veeam Backup & Replication, as described in the Adding Dell EMC Isilon section of the Integration with Storage Systems Guide. This account is used only for registering the storage system in the backup infrastructure and further storage rescans. You cannot use it for NAS file share backup and restore.

- We recommend that you create a custom **VeeamStorageIntegration** role to group all necessary privileges. This role must be created in the System access zone. For more information on access zones, see the Isilon documentation.

  Add the following privileges to the **VeeamStorageIntegration** role:

Privilege ID	Read Only	Comments
ISI_PRIV_LOGIN_PAPI	True	Used to log in to the Platform API and the web administration interface.
ISI_PRIV_AUTH	True	Used to configure external authentication providers, including root-level accounts.
ISI_PRIV_DEVICES	True	Used to create new roles and assign privileges.
ISI_PRIV_JOB_ENGINE R/W	False	Used for array Change File Tracking report scheduling for changelist API call.
ISI_PRIV_NETWORK	True	Used to configure network interfaces.
ISI_PRIV_NFS R/W	False	Used to update export rules.
ISI_PRIV_SMB	True	Used to configure the SMB server.
ISI_PRIV_SNAPSHOT R/W	False	Used to schedule, take, and view snapshots.

  For more information on privileges, see the Isilon documentation.

  You can use the following command to check privileges added to the **VeeamStorageIntegration** role:

```
isi auth roles view VeeamStorageIntegration
```

  Make sure that you make the service account to be used for NAS integration a member of the **VeeamStorageIntegration** role.

- To enable backup and restore of SMB file shares, additionally do the following:
  - Grant minimum read access to SMB file shares for the account that will be used for NAS integration. That will allow backing up SMB file shares.
  - Grant "run as root" access to SMB file shares for the account that will be used for NAS integration. That will allow restoring SMB file shares.

  Mind that the Isilon BackupAdmin role is not enough for backup and restore. It is mainly used for SMB access and enables backup and restore of files from `/ifs`. It does not propagate down to subfolders and therefore cannot be used to backup anything, but the system access zone.

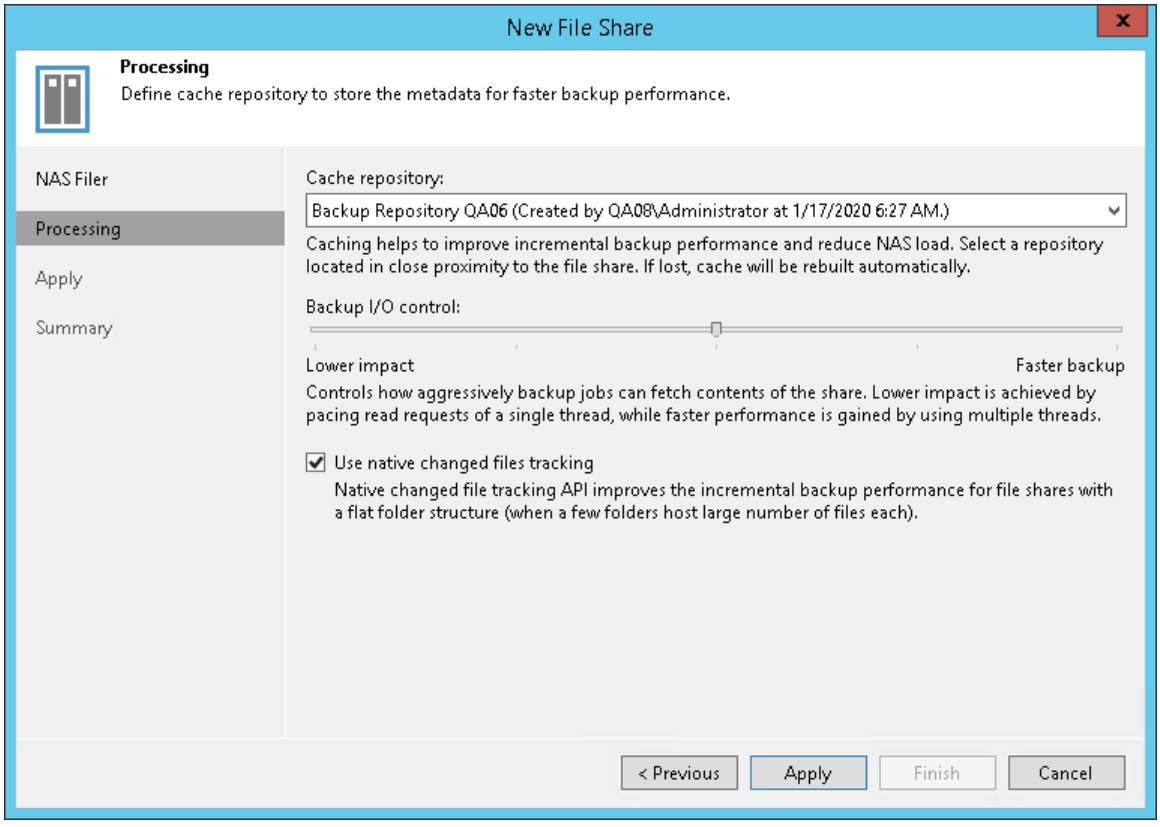
Step 3. Specify File Share Processing Settings

At the **Processing** step of the wizard, define file share processing settings:

1. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located as close to the NAS filer with file shares as possible.

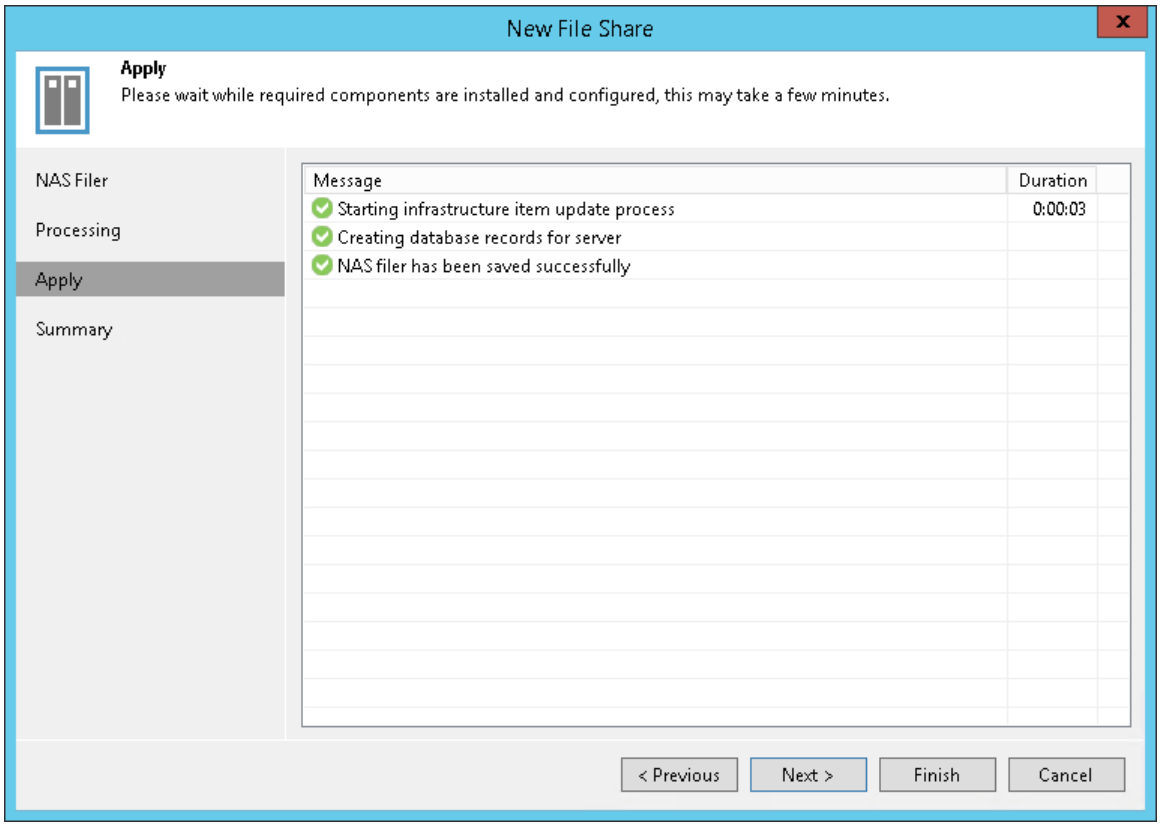
   > **NOTE**
   >
   > Mind that you cannot use a Linux-based server as a cache repository to process content of the protected shares on enterprise storage systems if you enable the **Use native changed files tracking** option.

2. Use the **Backup I/O control** slider to define how fast the backup proxy can read data from the source file share. This setting is based on the number of parallel threads that can be used by the proxy configured for processing the file shares on the NAS filer. If resources of your storage system are limited, it is recommended that you select the **Lower impact** option. If your storage system is powerful enough, select the **Faster backup** option.
3. [For Dell EMC Isilon only] Select the **Use native changed files tracking** check box if you want to use the file change tracking technology provided by the storage system manufacturer.
4. Click **Next** to save the configure settings.



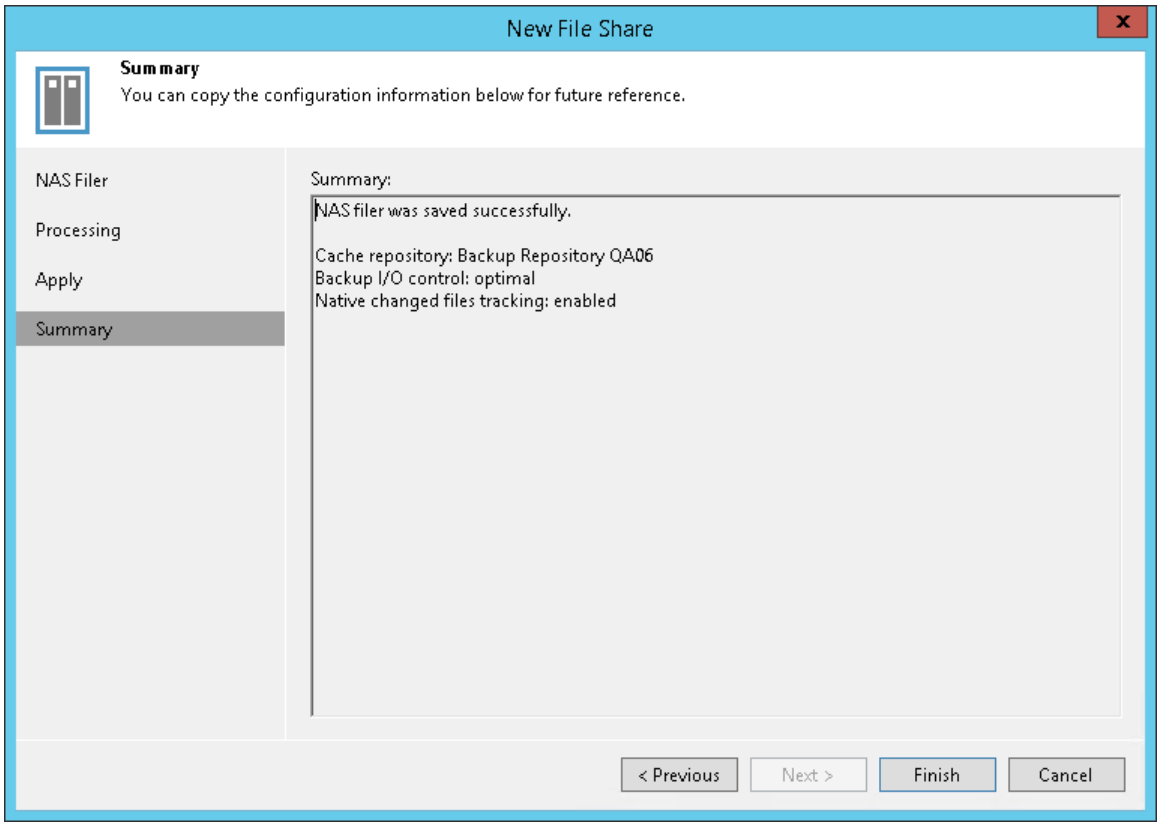
Step 4. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of adding the storage system as a NAS filer.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the storage system added as a NAS filer and click **Finish to exit the wizard**.



Adding NFS File Share

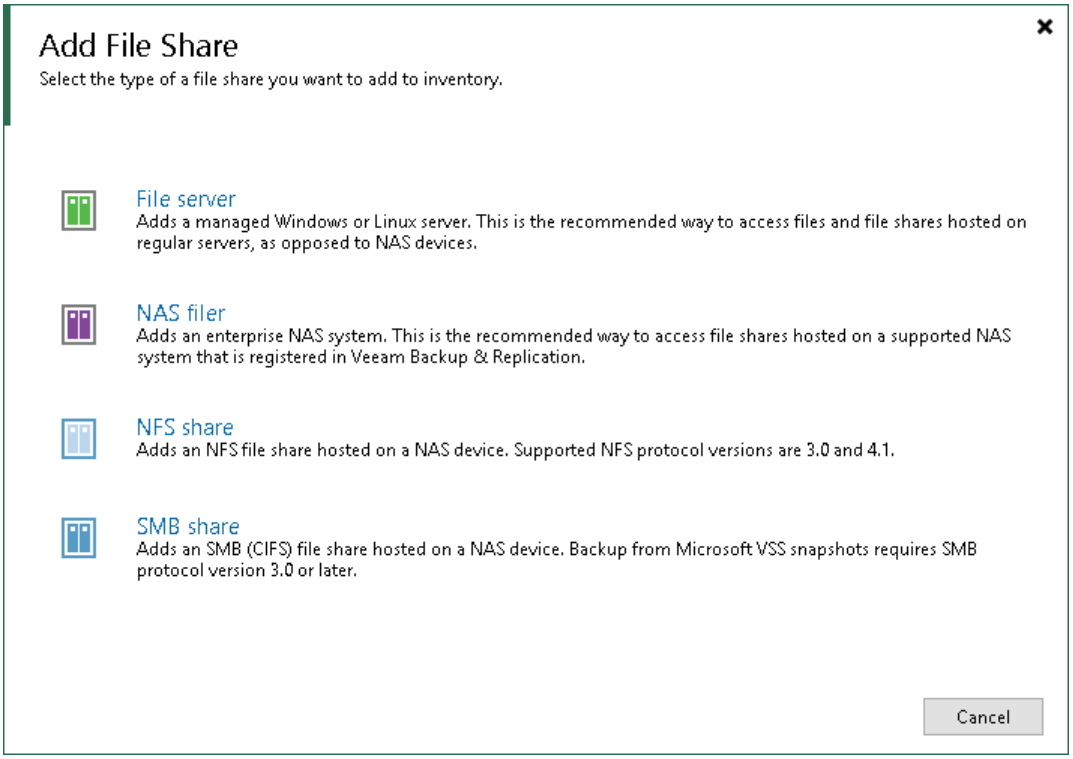
Before you add an NFS file share to the inventory of the virtual infrastructure, check the following prerequisites:

- The file share meets requirements listed in Platform Support.
- If you plan to use a devoted backup proxy server or cache repository, make sure these components are added in Backup Infrastructure.

Step 1. Launch New File Share Wizard

To launch the **New File Share** wizard:

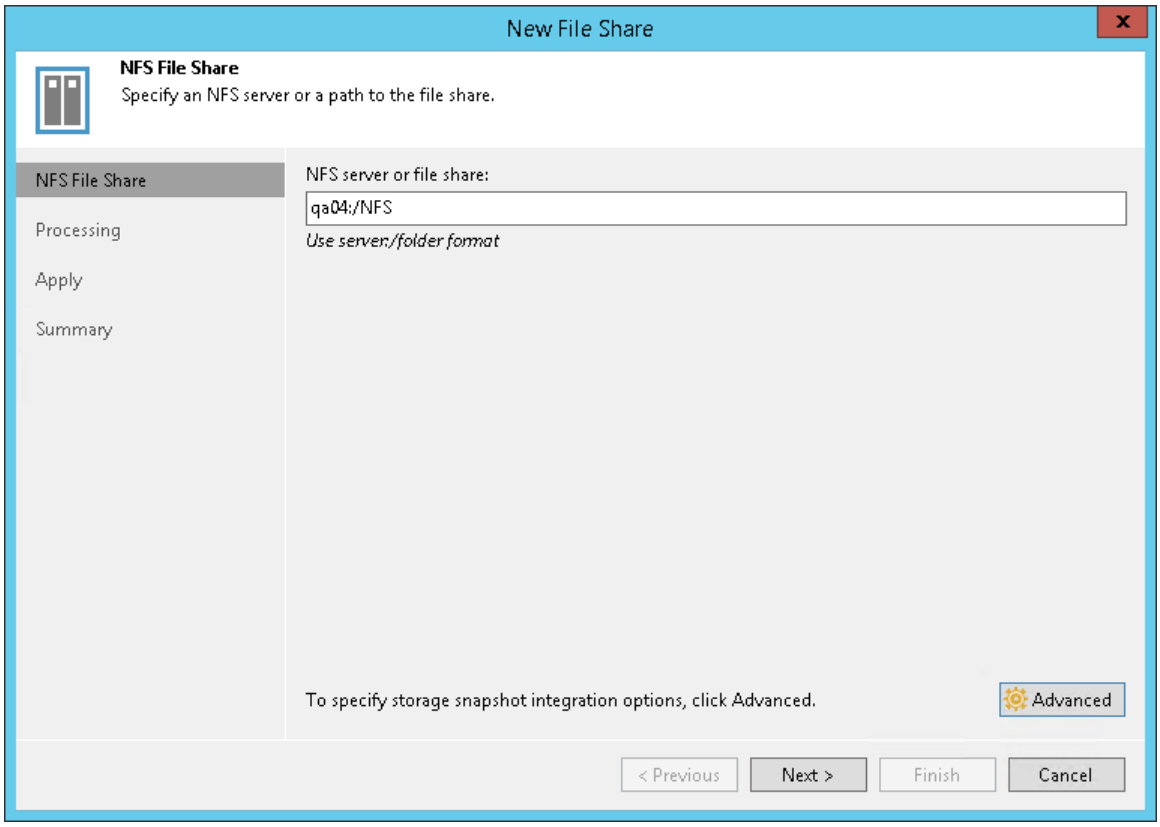
1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **File Shares** node and select **Add File Share**.
   - Select the **File Shares** node and click **Add File Share** on the ribbon.
   - Select the **File Shares** node and click **Add File Share** in the working area.
3. In the **Add File Share** window, click **NFS share**.



Step 2. Specify Path to NFS File Share

At the **NFS File Share** step of the wizard, specify the path to an NFS file share in the *server:/folder* format.

You can add the root server folder in the *server:/* format to protect all NFS file shares residing on this server. After that, create a single file share backup job to protect the added server, as described in Creating File Share Backup Jobs. Then all NFS file shares added on this server will be automatically processed with the file share backup job and protected. If you previously had several separate non-root shared folders residing on the same server and want to switch to using a single root shared folder to cover the same shares, you do not have to run full backups to update data of protected shares. Instead, you can convert existing backups and update existing file share backup jobs to protect single root shared folders comprising all other non-root shared folders residing on the same server. To learn more about the conversion, see Converting Backups from Non-Root to Root Shared Folders. Perform the conversion with extreme caution.



Step 3. Specify Advanced NFS File Share Settings

You can instruct Veeam Backup & Replication to back up data from native storage snapshots. During backup jobs, Veeam Backup & Replication will read data of shared files and folders from snapshots, which speeds up backup operations and improves RPOs.

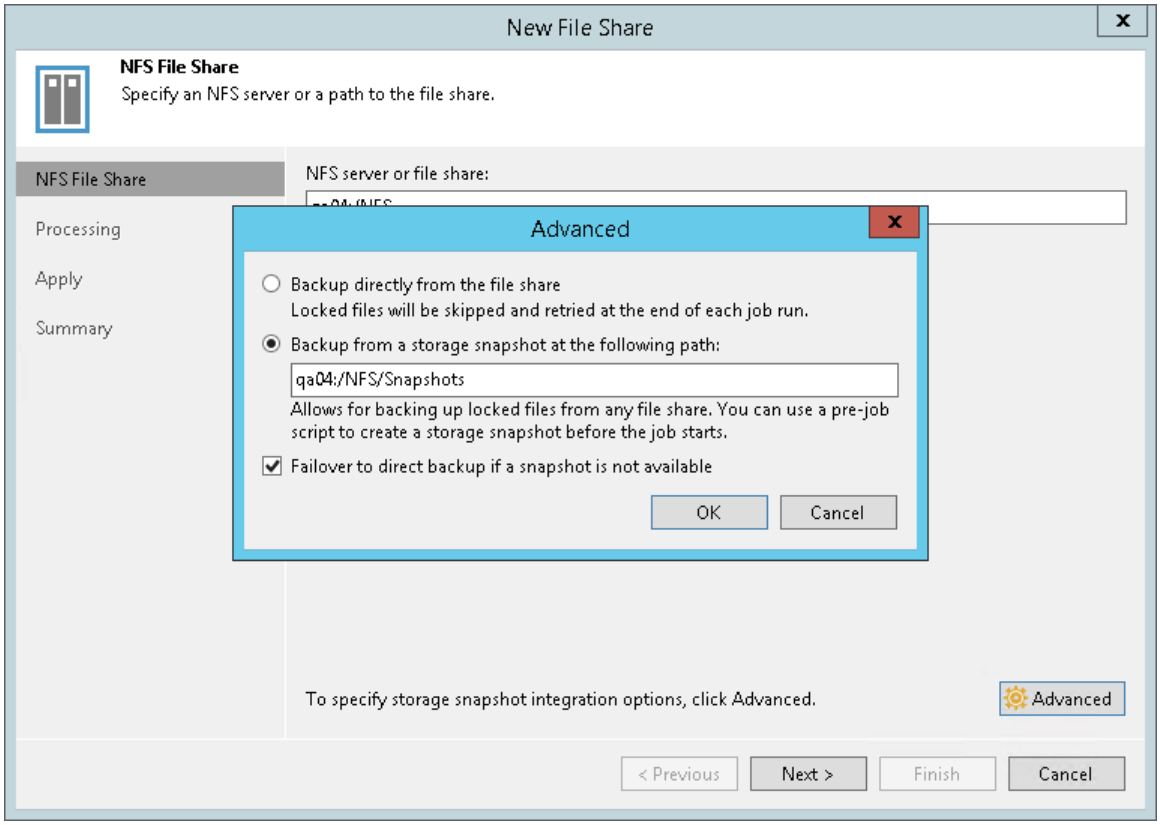
To define if Veeam Backup & Replication will use snapshots for backups:

1. At the **NFS File Share** step of the wizard, click **Advanced**.
2. In the **Advanced** window, select one of the following options:
   - To ignore the snapshot functionality, select **Backup directly from the file share**.
     Veeam Backup & Replication will ignore locked files and folders. When creating a backup job, you can configure notifications to list files and folders that are skipped during the backup procedure. For more information see Notification Settings.
   - To back up files from the native storage snapshot, select **Backup from a native storage snapshot at the following path** and specify the path in the *server:/folder* format to the folder on the NFS file share where the file share snapshot is saved.

     If you select this option, you can additionally use custom scripts written by you, for example, to create a snapshot before the backup and remove it after the backup. You can define these scripts when creating a new file share backup job, as described in Script Settings.

> **NOTE**
>
> Mind that Veeam Backup & Replication does not take snapshots itself, but it can use a snapshot taken by the storage system.
>
> File share backup jobs do not trigger the storage snapshot creation and deletion automatically. You can specify the folder where the storage snapshot is stored. In this case file share backup jobs can access this folder and read data from the storage snapshot.

3. Select **Failover to direct backup if a snapshot is not available** if you want Veeam Backup & Replication to read data for backup directly from the file share when the snapshot is unavailable. If you do not select the option and the snapshot is unavailable, the file share backup job will stop with a failure.



Step 4. Specify File Share Processing Settings

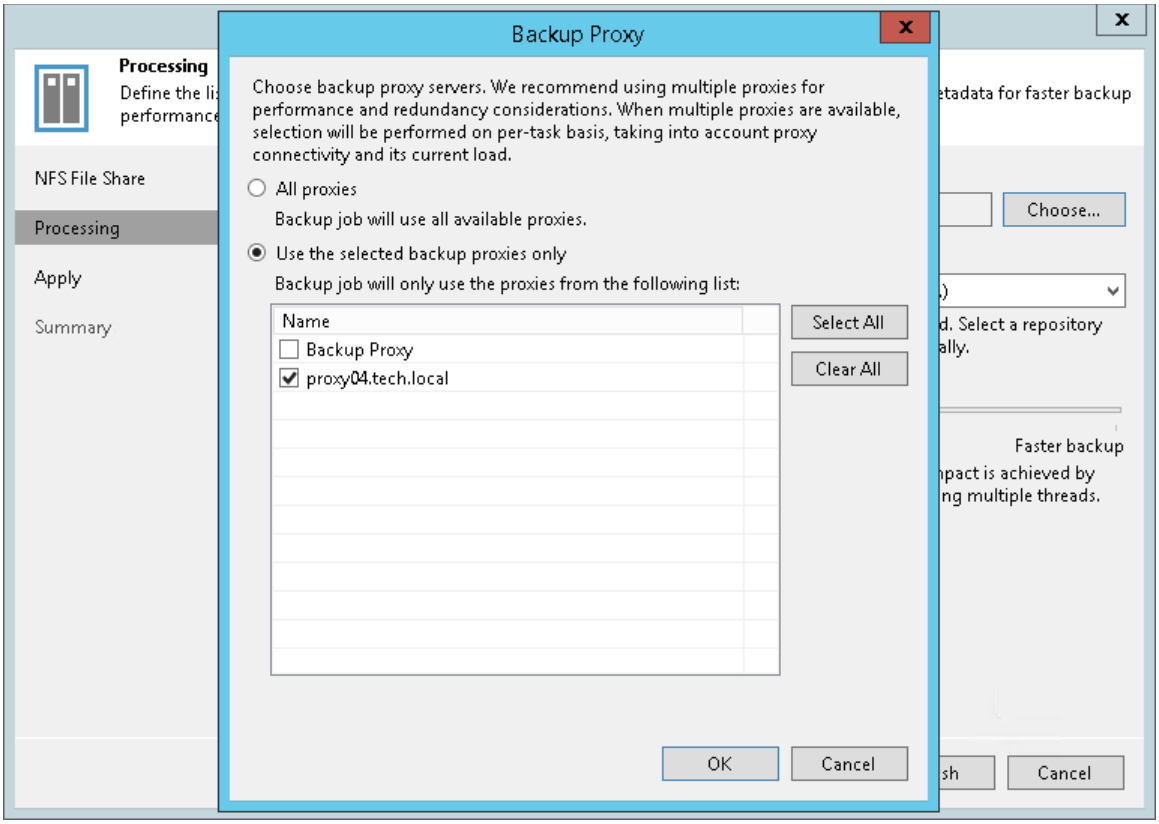
At the **Processing** step of the wizard, do the following:

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
2. In the **Backup Proxy** window, select backup proxy servers:
   - If you select **All proxies**, Veeam Backup & Replication will use all available backup proxies for file backup. The number of proxies in use defines the number of data threads that transfer data from the file share to the backup repository. The more data transfer threads Veeam Backup & Replication uses, the higher is the data transfer speed.

     If the file share is used as a source for a file to tape backup job, the tape server utilized for this job is added as yet another backup proxy when creating a file to tape backup job. This backup proxy has the highest priority over all others and is used by default if it has access rights to the file share. For details on file to tape backup jobs, see the File Backup to Tape section in the Veeam Backup & Replication User Guide.
   - If you select **Use the selected backup proxies only**, you can explicitly specify backup proxies that Veeam Backup & Replication must use for file backup.

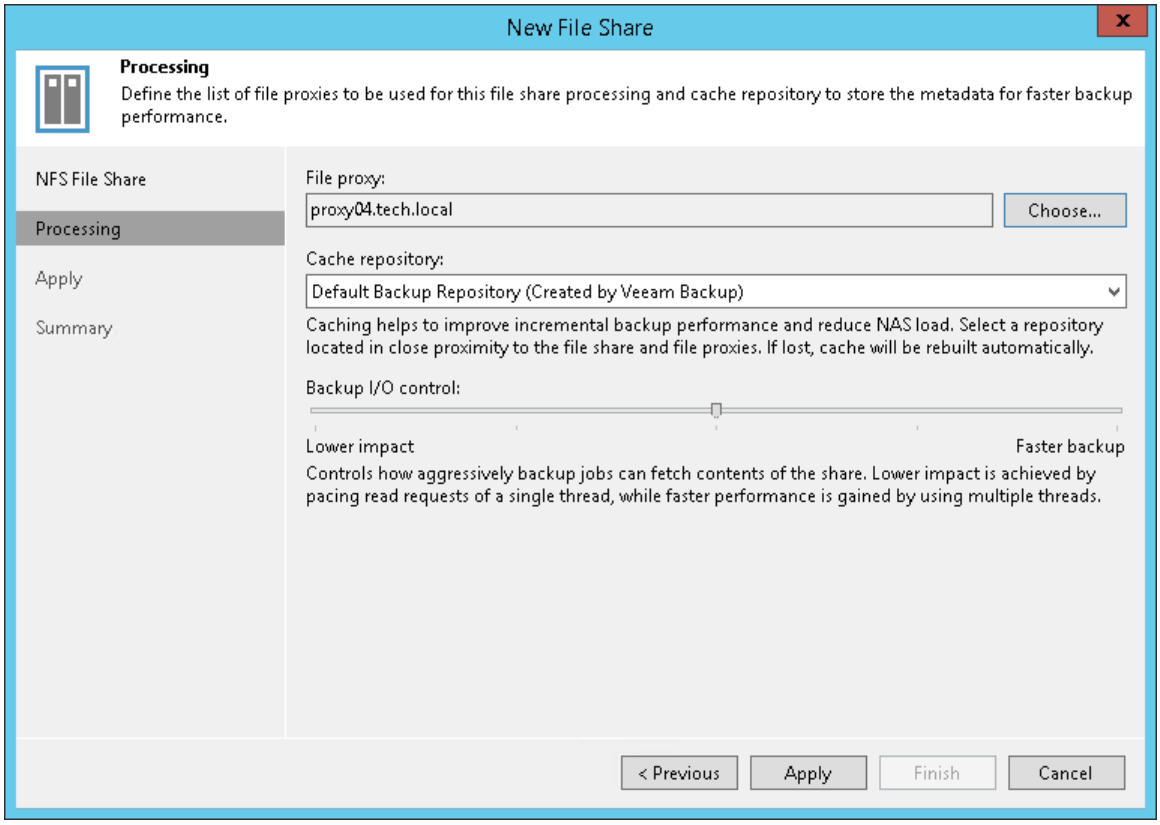
     It is recommended that you select at least two backup proxies to ensure that the backup jobs start even if one of the proxies fails or loses its connectivity to the source file share. The more proxies you select, the more data transfer threads Veeam Backup & Replication will use for backup jobs, thus improving performance.

     Even if the file share is used as a source for file to tape backup jobs, Veeam Backup & Replication will use only proxies selected in the list to process the backup data traffic.



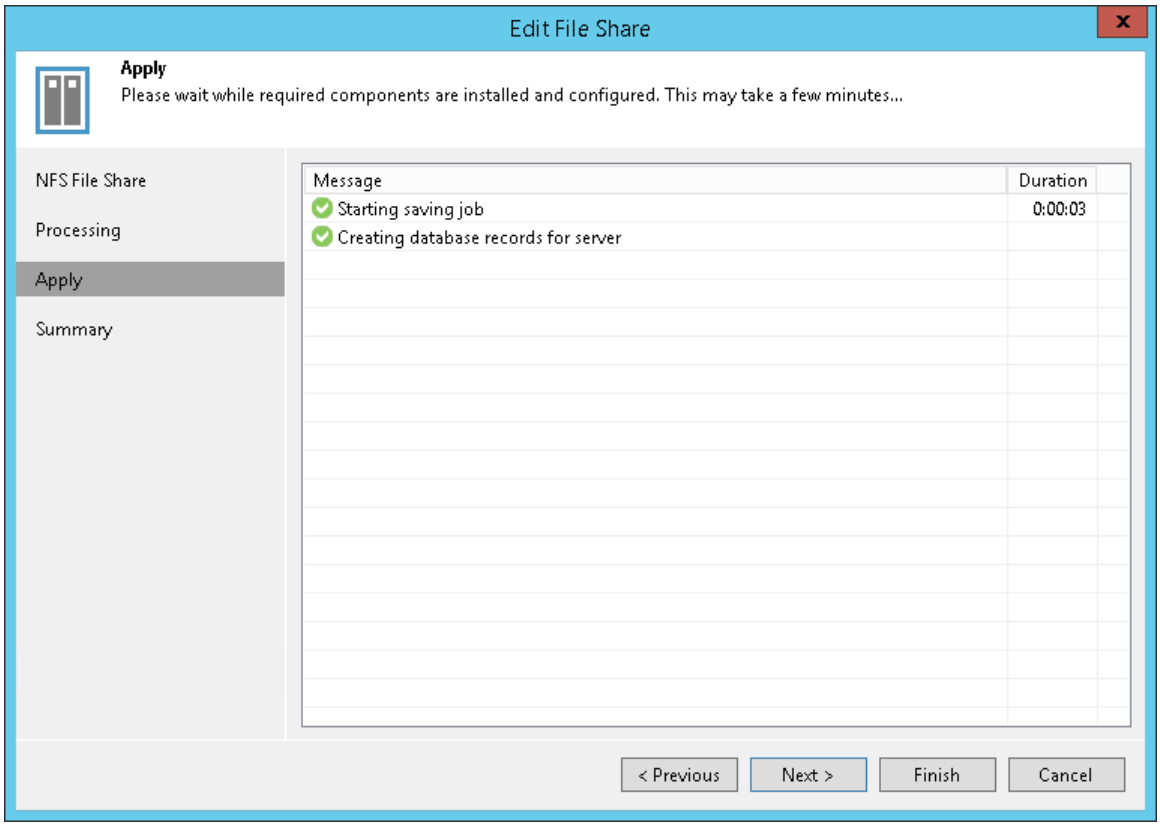
3. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located as close to the source file share as possible.

4. Use the **Backup I/O control** slider to define how fast all proxies can read data from the source file share. This setting is based on the number of parallel threads that can be used by all the proxies configured for processing the file share. If resources of your NAS device are limited, it is recommended that you select the **Lower impact** option. If your NAS device is powerful enough, select the **Faster backup** option.
5. Click **Apply** to save the configured settings.



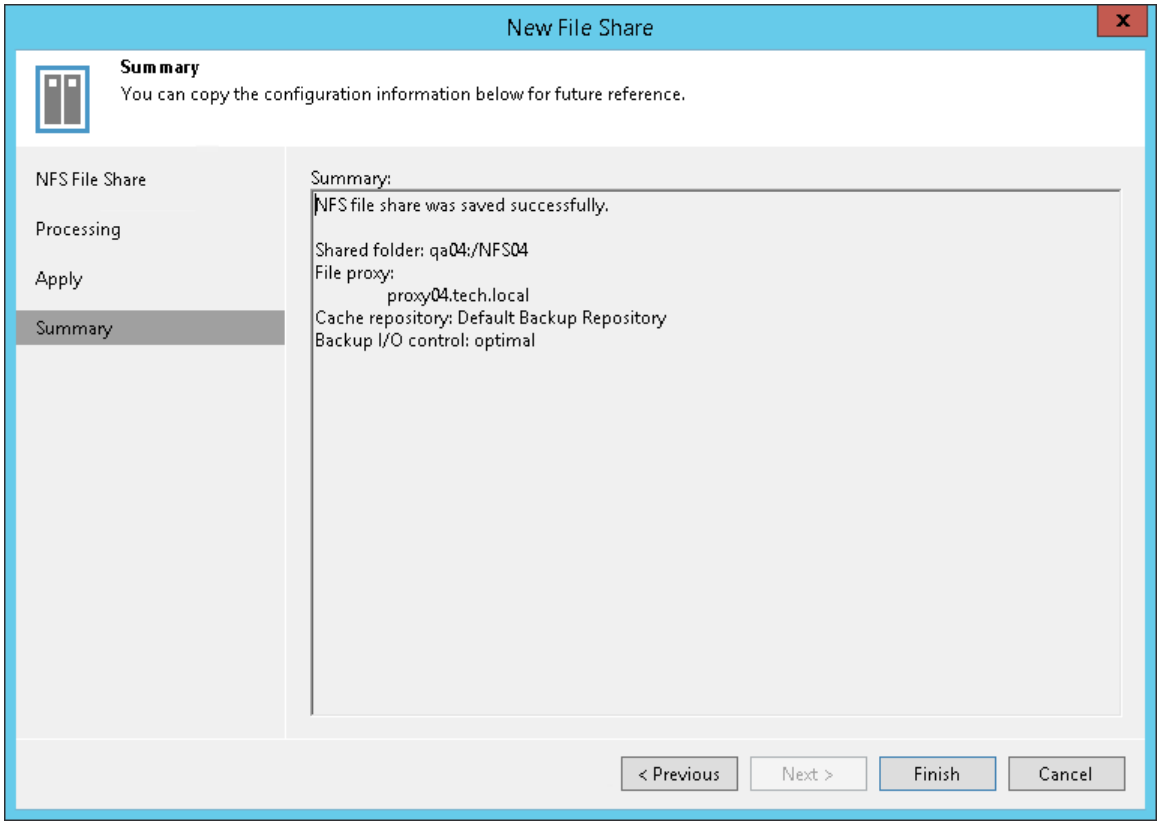
Step 5. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the NFS file share to the backup infrastructure. Click **Next** to proceed.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added NFS share and click **Finish** to exit the wizard.



Adding SMB File Share

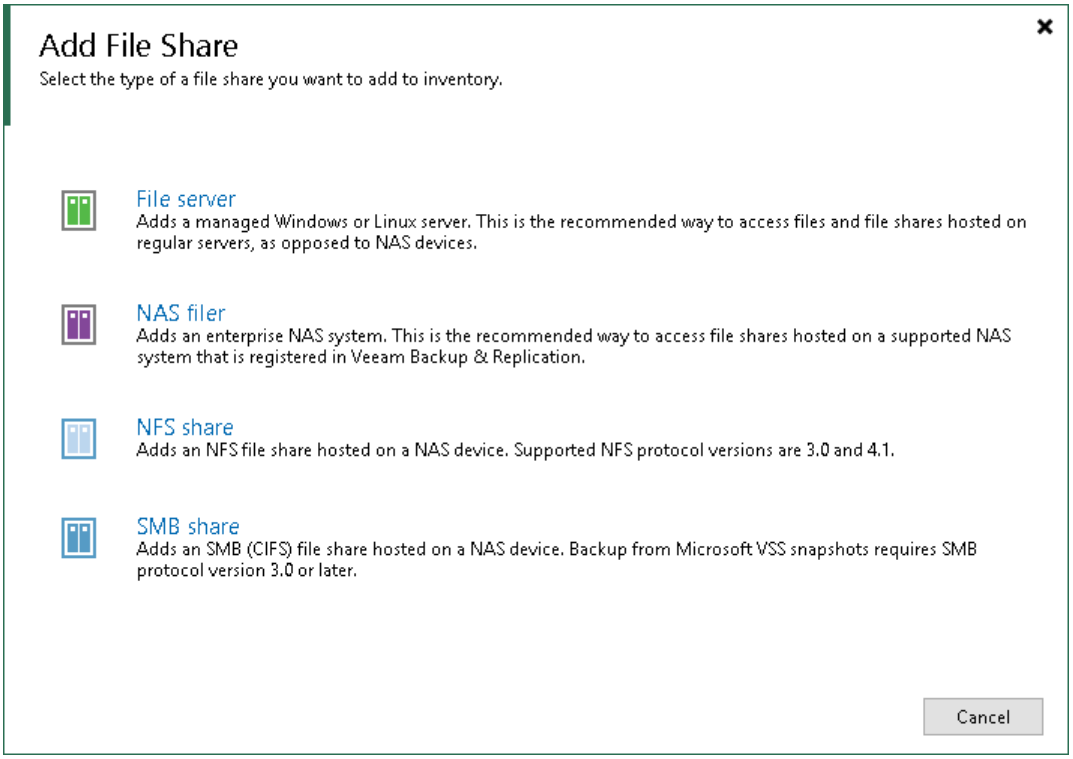
Before you add an SMB file share to the inventory of the virtual infrastructure, check the following prerequisites:

- The file share meets requirements listed in Platform Support.
- If you plan to use a devoted backup proxy server or cache repository, make sure these components are added in **Backup Infrastructure**.

Step 1. Launch New File Share Wizard

To launch the **New File Share** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **File Shares** node and select **Add File Share**.
   - Select the **File Shares** node and click **Add File Share** on the ribbon.
   - Select the **File Shares** node and click **Add File Share** in the working area.
3. In the **Add File Share** window, click **SMB share**.



Step 2. Specify Path to SMB File Share and Access Credentials

At the **SMB File Share** step of the wizard, specify access settings for the SMB file share:

1. In the **SMB server of file share** field, specify the path to an SMB file share in the *||server|folder* format.

   You can also click **Browse** and select the shared folder from the list of available network shares.

   You can add the root server folder in the *||server* format to protect all SMB file shares residing on this server. After that, create a single file share backup job to protect the added server, as described in Creating File Share Backup Jobs. Then all SMB file shares added on this server will be automatically processed with the file share backup job and protected. If you previously had several separate non-root shared folders residing on the same server and want to switch to using a single root shared folder to cover the same shares, you do not have to run full backups to update data of protected shares. Instead, you can convert existing backups and update existing file share backup jobs to protect single root shared folders comprising all other non-root shared folders residing on the same server. To learn more about the conversion, see Converting Backups from Non-Root to Root Shared Folders. Perform the conversion with extreme caution.

2. If you must specify user credentials to access the shared folder, select the **This share requires access credentials** check box. From the **Credentials** drop-down list, select a credentials record for a user account that has **Full Control** permissions on the shared folder.

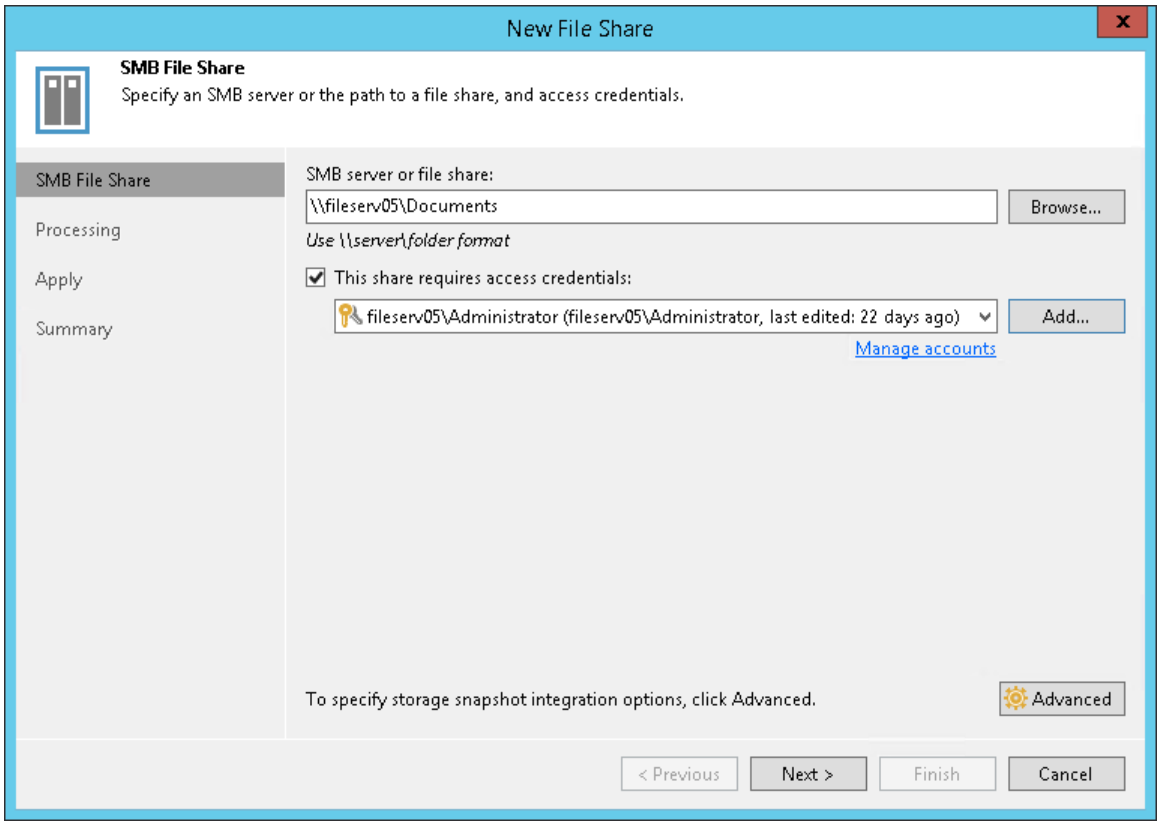
   To access the SMB share, you must use an account that meets either of the following requirements:

   - If you only plan to back up the share, you can use an account with read-only permissions.
   - If you plan not only to back up the share, but also to restore files to it, use an account with read/write permissions.

   > **NOTE**
   >
   > Accessing the SMB file share with credentials in the User Principal Name format (*user@domain.xxx*) is not supported.

   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.

> **NOTE**
>
> If the **This share requires access credentials** check box is not selected, to access the file share Veeam Backup & Replication uses the computer account of the backup proxy server.



Step 3. Specify Advanced SMB File Share Settings

You can instruct Veeam Backup & Replication to back up data from Microsoft VSS snapshots or native storage snapshots. During backup jobs, Veeam Backup & Replication will read data of shared files and folders from snapshots, which speeds up backup operations and improves RPOs.

To define if Veeam Backup & Replication will use snapshots for backups:

1. At the **SMB File Share** step of the wizard, click **Advanced**.
2. In the **Advanced** window, select one of the following options:
   - To ignore the snapshot functionality, select **Backup directly from the file share**.

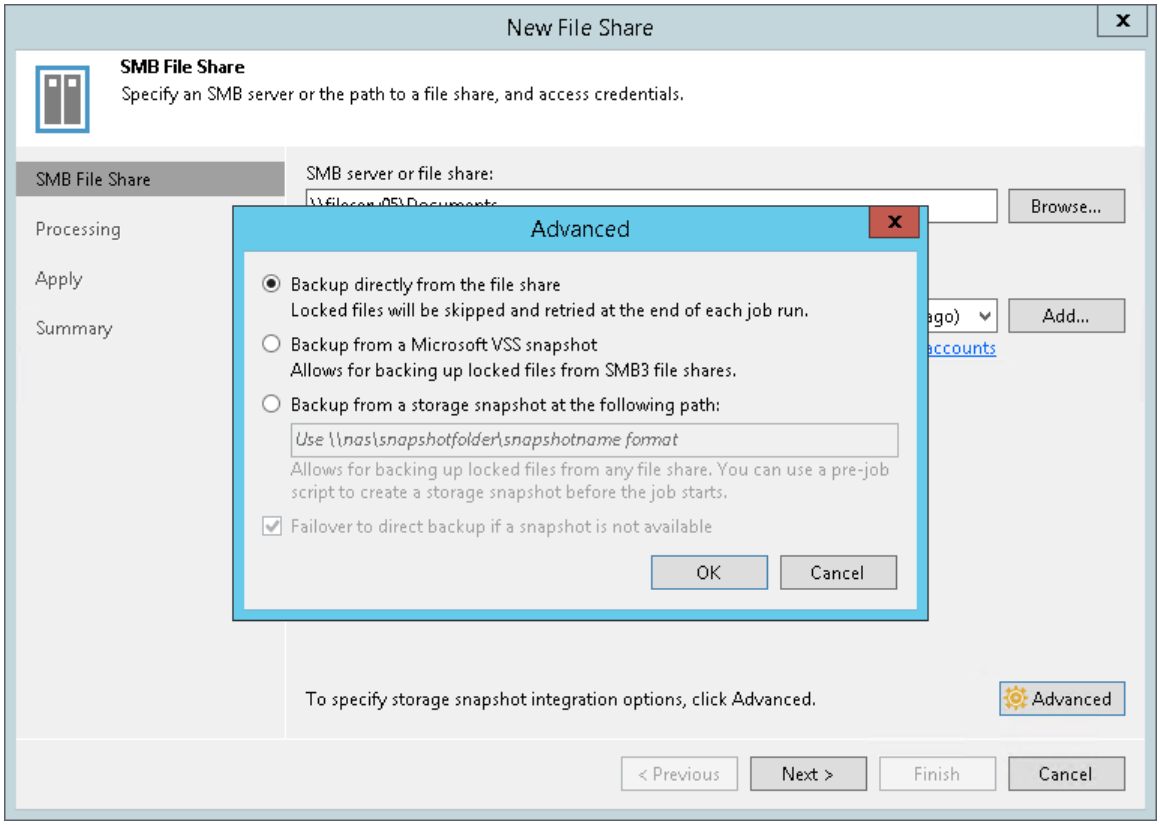
     Veeam Backup & Replication will ignore locked files and folders. When creating a backup job, you can configure notifications to list files and folders that are skipped during the backup procedure. For more information see Notification Settings.
   - To back up files from Microsoft VSS snapshots, select **Backup from Microsoft VSS snapshot**.

     If you select this option, make sure that the file share and the backup proxy used for the file backup job support SMB protocol version 3.0 or later.
   - To back up files from the native storage snapshot, select **Backup from a native storage snapshot at the following path** and specify the path in the *||server|folder* format to the folder on the SMB file share where the file share snapshot is saved.

     If you select this option, you can additionally use custom scripts written by you, for example, to create a snapshot before the backup and remove it after the backup. You can define these scripts when creating a new file share backup job, as described in Script Settings.

> **NOTE**
>
> Mind that Veeam Backup & Replication does not take snapshots itself, but it can use a snapshot taken by the storage system.
>
> File share backup jobs do not trigger the storage snapshot creation and deletion automatically. You can specify the folder where the storage snapshot is stored. In this case file share backup jobs can access this folder and read data from the storage snapshot.

3. Select **Failover to direct backup if snapshot is unavailable** if you want Veeam Backup & Replication to read data for backup directly from the file share when the snapshot is unavailable. If you do not select the option and the snapshot is unavailable, the file share backup job will stop with a failure.



Step 4. Specify File Share Processing Settings

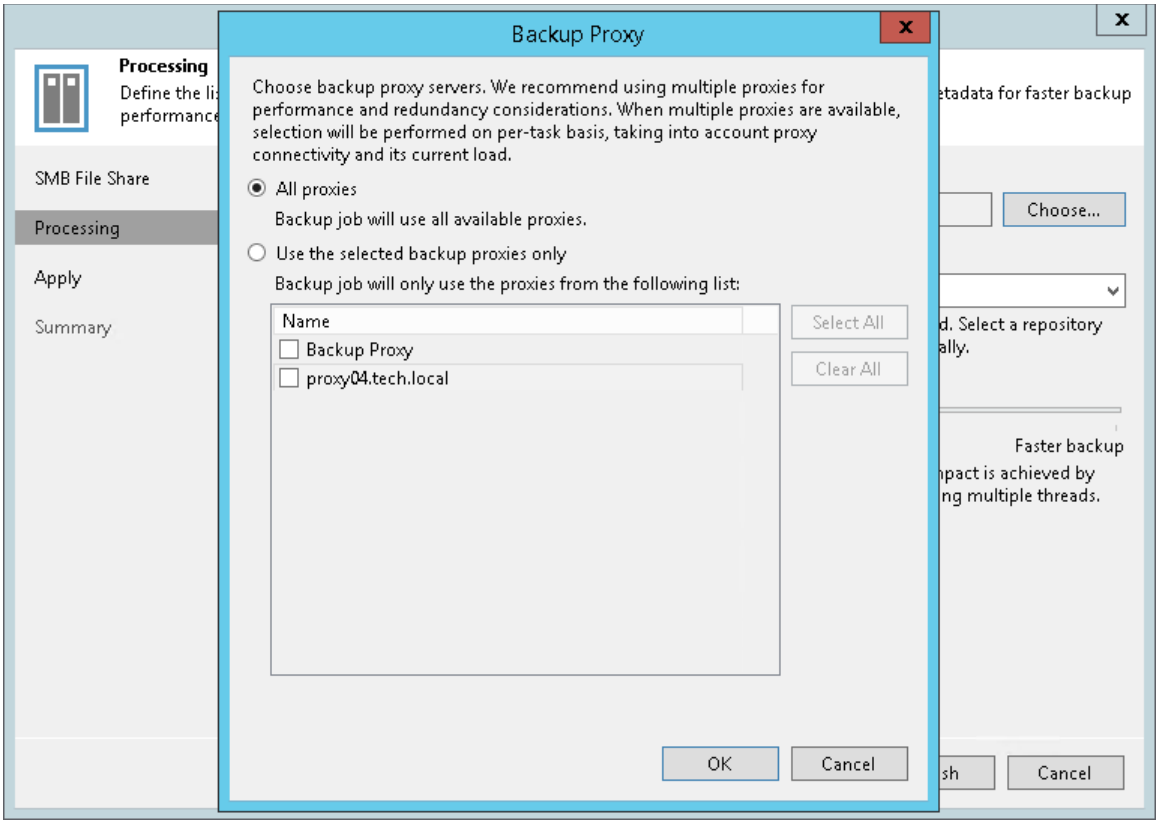
At the **Processing** step of the wizard, do the following:

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
2. In the **Backup Proxy** window, select backup proxy servers:
   - If you select **All proxies**, Veeam Backup & Replication will use all available backup proxies for file backup. The number of proxies in use defines the number of data threads that transfer data from the file share to the backup repository. The more data transfer threads Veeam Backup & Replication uses, the higher is the data transfer speed.

     If the file share is used as a source for a file to tape backup job, the tape server utilized for this job is added as yet another backup proxy when creating a file to tape backup job. This backup proxy has the highest priority over all others and is used by default if it has access rights to the file share. For details on file to tape backup jobs, see the File Backup to Tape section in the Veeam Backup & Replication User Guide.
   - If you select **Use the selected backup proxies only**, you can explicitly specify backup proxies that Veeam Backup & Replication must use for file backup.

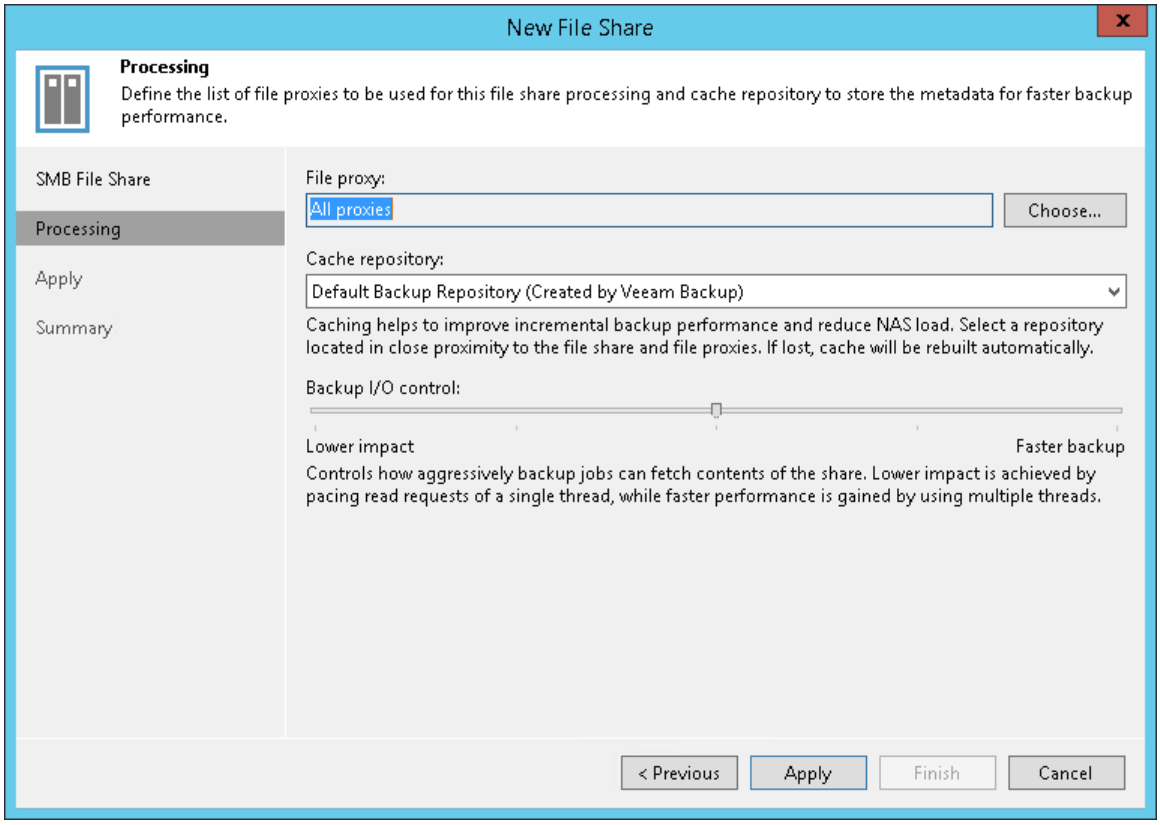
     It is recommended that you select at least two backup proxies to ensure that the backup jobs start even if one of the proxies fails or loses its connectivity to the source file share. The more proxies you select, the more data transfer threads Veeam Backup & Replication will use for backup jobs, thus improving performance.

     Even if the file share is used as a source for file to tape backup jobs, Veeam Backup & Replication will use only proxies selected in the list to process the backup data traffic.



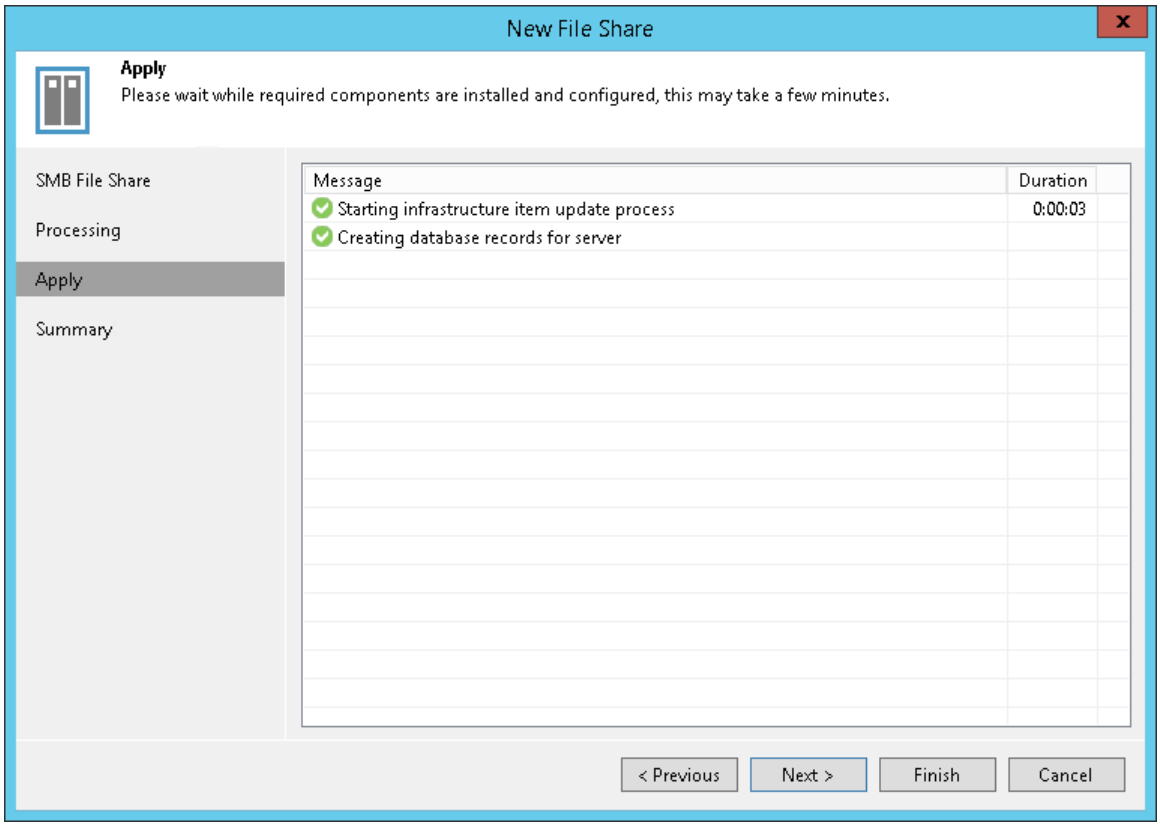
3. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located in the close proximity to the source file share and backup proxies.

4. Use the **Backup I/O control** slider to define how fast all proxies can read data from the source file share. This setting is based on the number of parallel threads that can be used by all the proxies configured for processing the file share. If resources of your NAS device are limited, it is recommended that you select the **Lower impact** option. If your NAS device is powerful enough, select the **Faster backup** option.
5. Click **Apply** to save the configured settings.



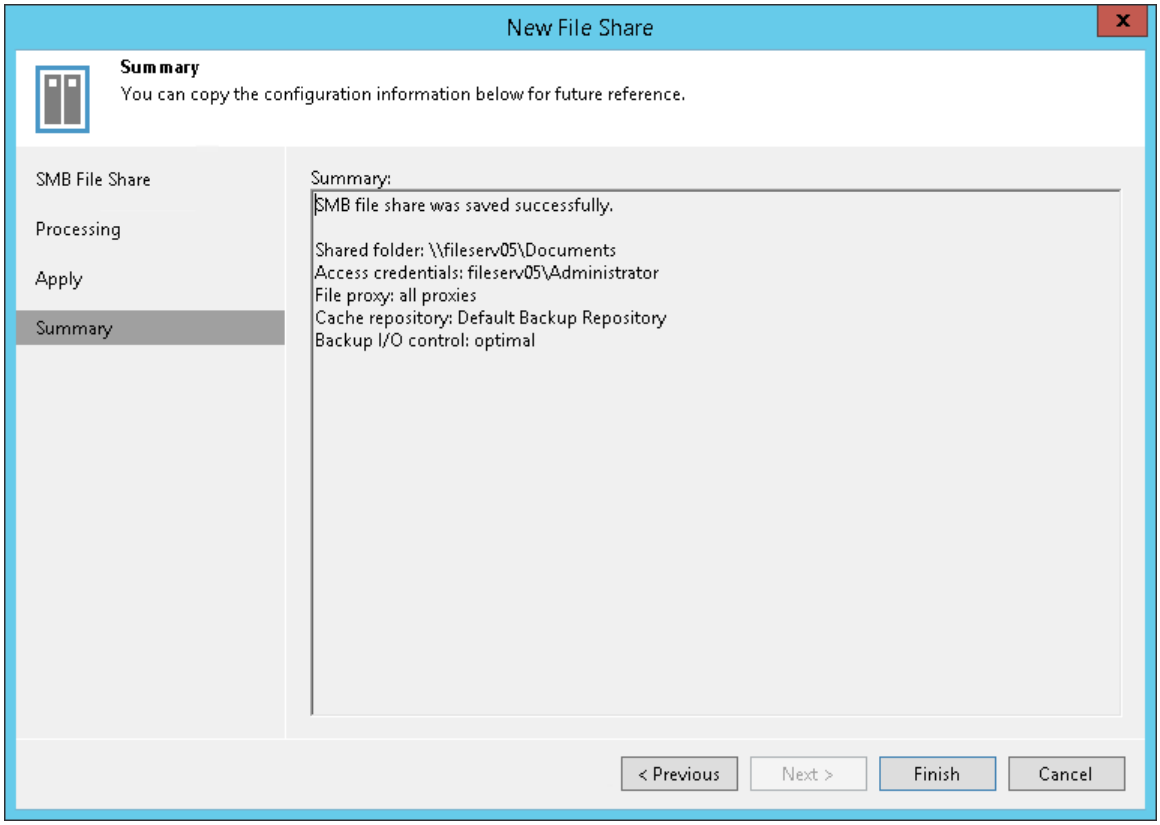
Step 5. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the SMB file share to the backup infrastructure. Click **Next** to proceed.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added SMB share and click **Finish** to exit the wizard.



Adding Backup Proxy

You must add to the backup infrastructure one or more backup proxies that you plan to use for moving backup data from the file share to the backup repository.

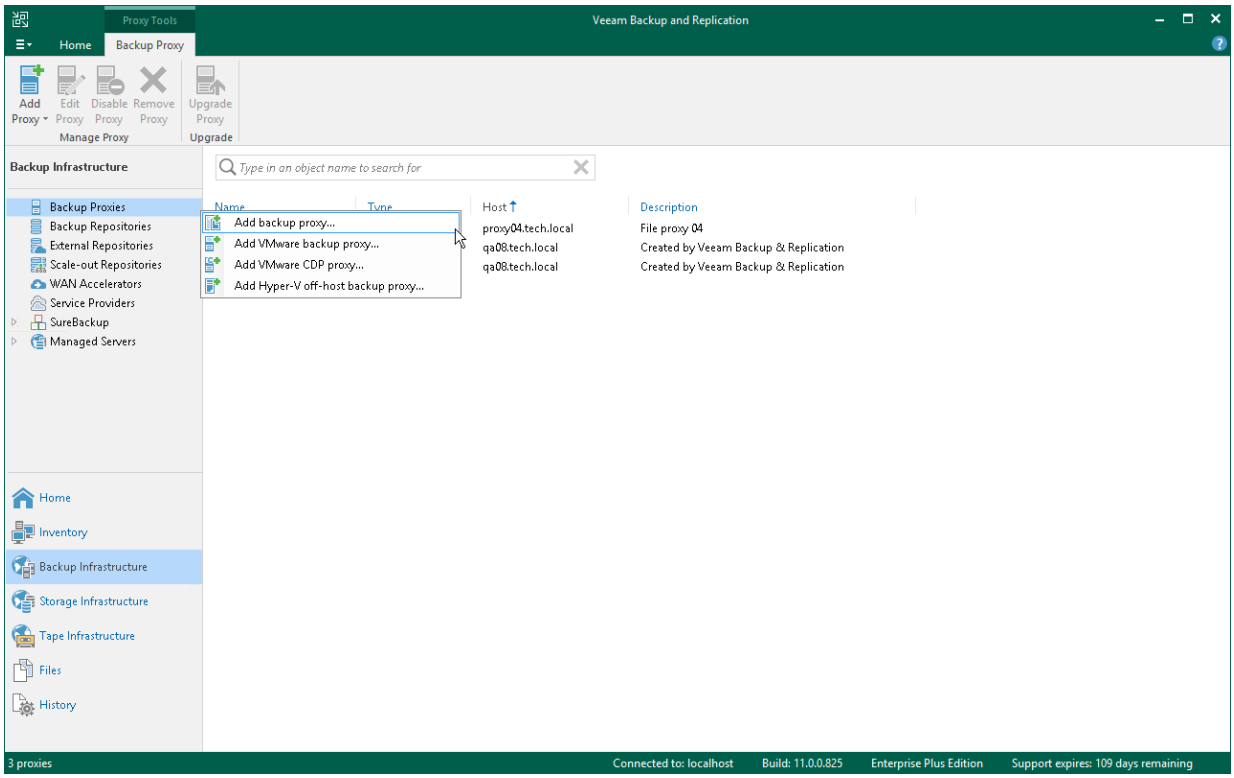
Before you add a backup proxy to the inventory of the virtual infrastructure, check the following prerequisites:

- The backup proxy meets system requirements.
- The backup proxy is a Windows-managed server.

Step 1. Launch New Backup Proxy Wizard

To launch the **New Backup Proxy** wizard, do one 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 **Backup proxy**.
- Open the **Backup Infrastructure** view, in the inventory pane right-click the **Backup Proxies** node and select **Add backup proxy**.



Step 2. Choose Microsoft Windows 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 server that you want to use as a backup proxy.

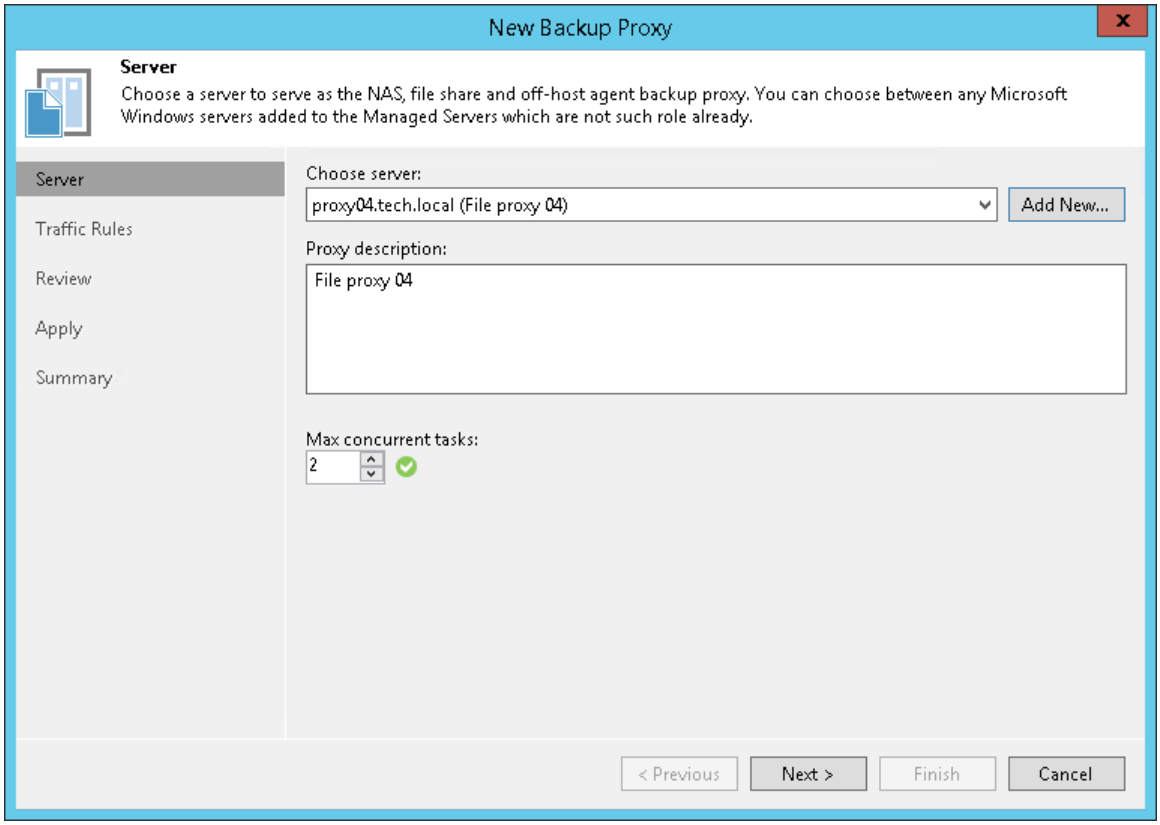
   The list of servers contains only those managed servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** to open the **New Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.

2. In the **Proxy description** field, provide a description for future reference. The default description contains information about the user who added the backup proxy, date and time when the backup proxy was added.

3. In the **Max concurrent tasks** field, specify the number of tasks that the backup proxy can process in parallel.

   If the number of parallel tasks reaches this value, the backup proxy will not start a new task until one of current tasks completes. Veeam Backup & Replication creates one task per every source file share. The recommended number of concurrent tasks is calculated automatically based on the amount of 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 a maximum of 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 infrastructure.



Step 3. 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.

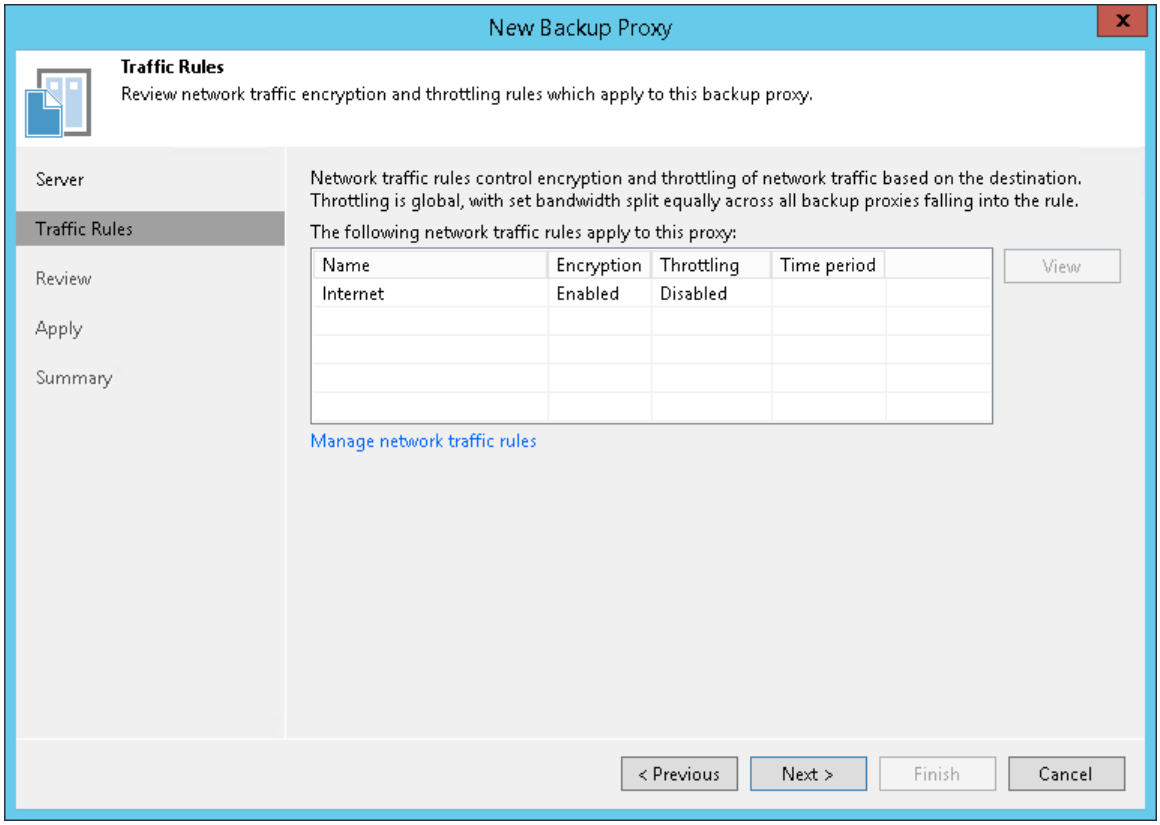
The list of network traffic rules contains only rules applied to the backup proxy: its IP address falls into the IP range configured for the rule.

To view settings configured for the rule:

1. Select the rule in the list.
2. Click **View**. The **View Network Traffic Rule** window will display settings configured for the rule.

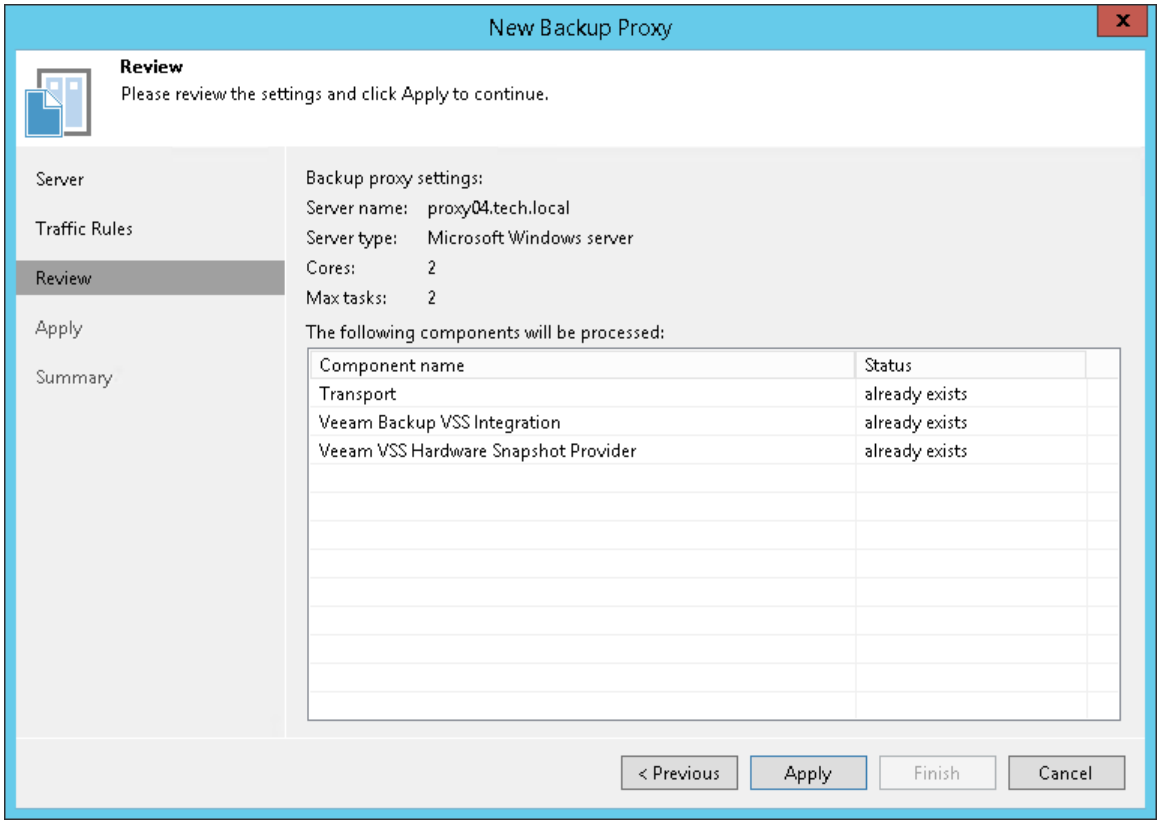
To modify network traffic settings:

1. Click the **Manage network traffic rules** link.
2. The **Global Network Traffic Rules** window will display the full list of all existing global network traffic rules.
3. Select the rule that you want to modify and click **Edit**. For more information on how to configure network traffic rules, see Configuring Network Traffic Rules.



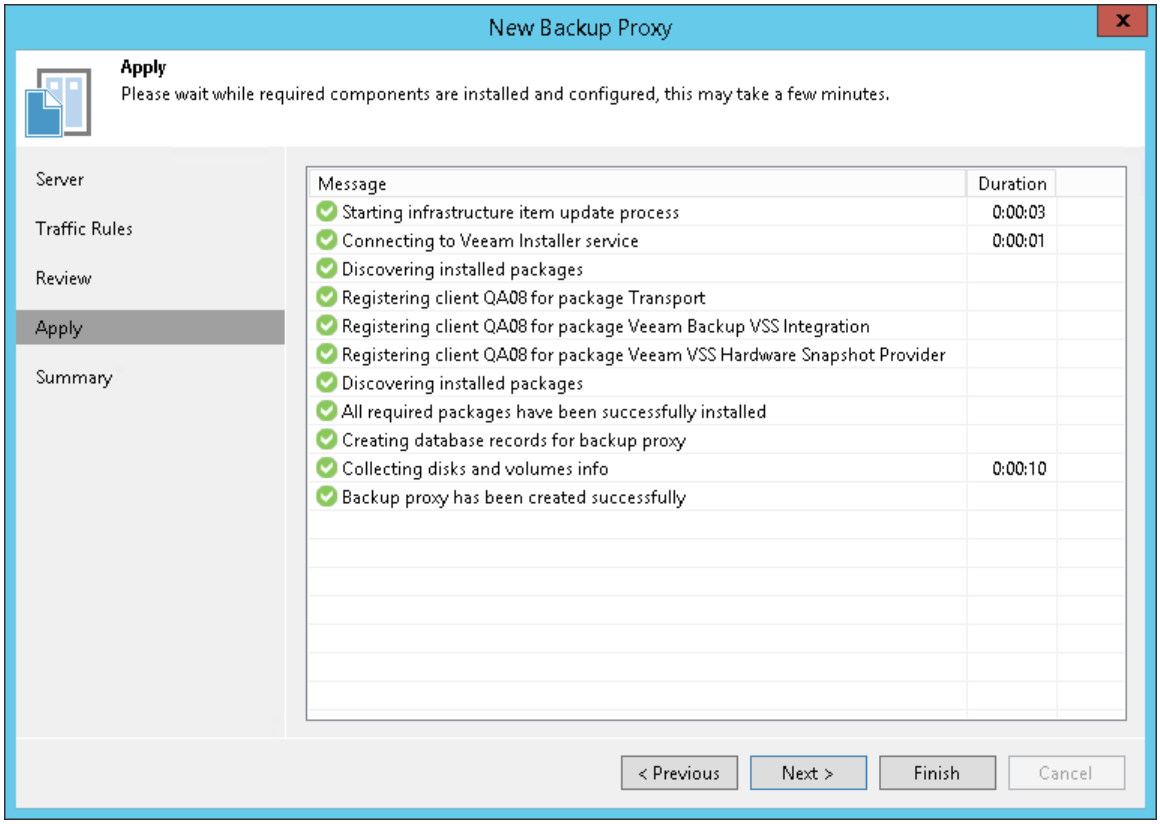
Step 4. Review Components to Install

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and click **Apply** to start installation of missing components.



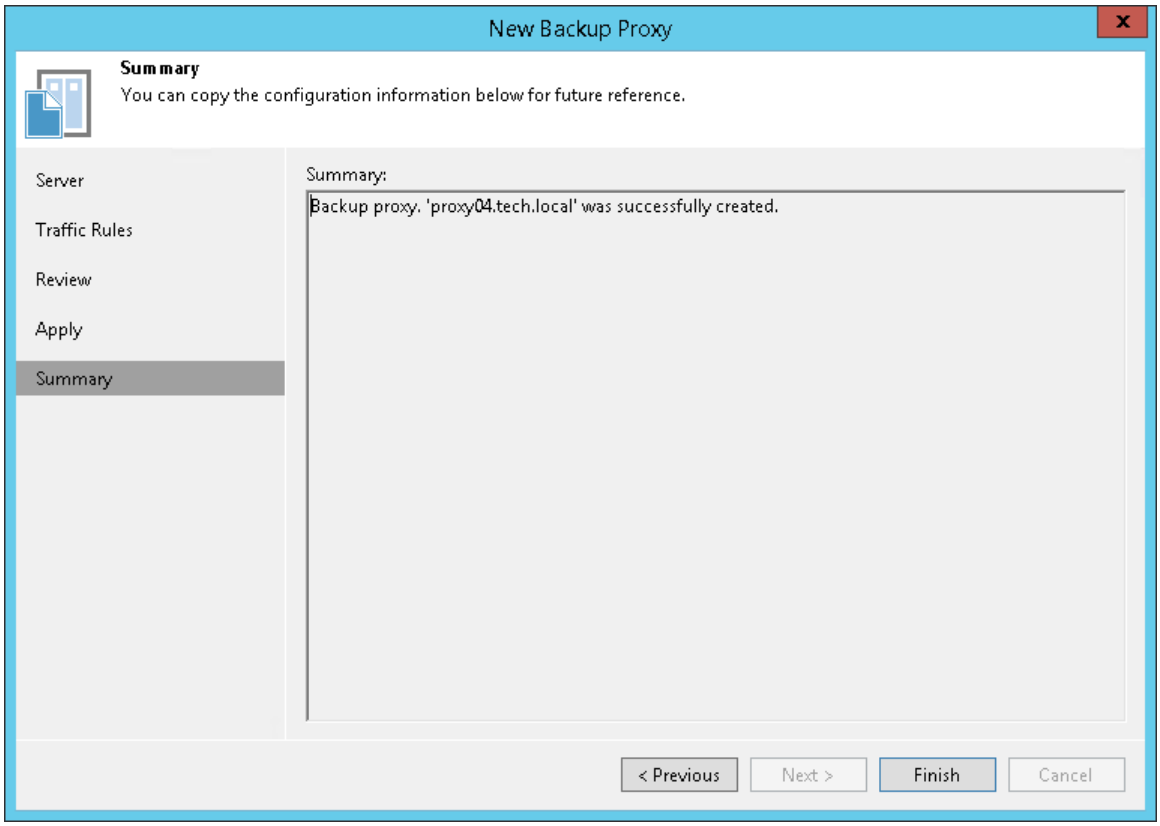
Step 5. Apply Backup Proxy Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of the backup proxy role assignment to the server.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup proxy and click **Finish** to exit the wizard.



Creating File Share Backup Jobs

To protect files and folders on the file share, you must configure a file share backup job. The backup job defines how, where and when to back up data from the file share. One job can be used to process one or more file shares. Jobs can be started manually or scheduled to run automatically at a specific time.

Before you create a file share backup job, check prerequisites.

Before You Begin

Before you create a file share backup job, check the following prerequisites:

- Backup infrastructure components that will take part in the file share backup process must be added to the backup infrastructure and properly configured. These include source file shares to back up, backup proxy, and all repositories, including cache, backup and archive repositories.
- The target backup repository must have enough free space to store created backup files. If you want to receive notifications on the repository running low on free space, configure global notification settings as described in Specifying Other Notification Settings.
- Make sure that repositories intended to store file share backups are not configured to store files in the WORM status. Otherwise, the backup jobs will fail when Veeam Backup & Replication cannot update the backup metadata files.
- If you plan to map a file share backup job to a backup that already exists in the backup repository, you must perform the rescan operation for this backup repository. Otherwise, Veeam Backup & Replication will not be able to recognize backup files in the backup repository.

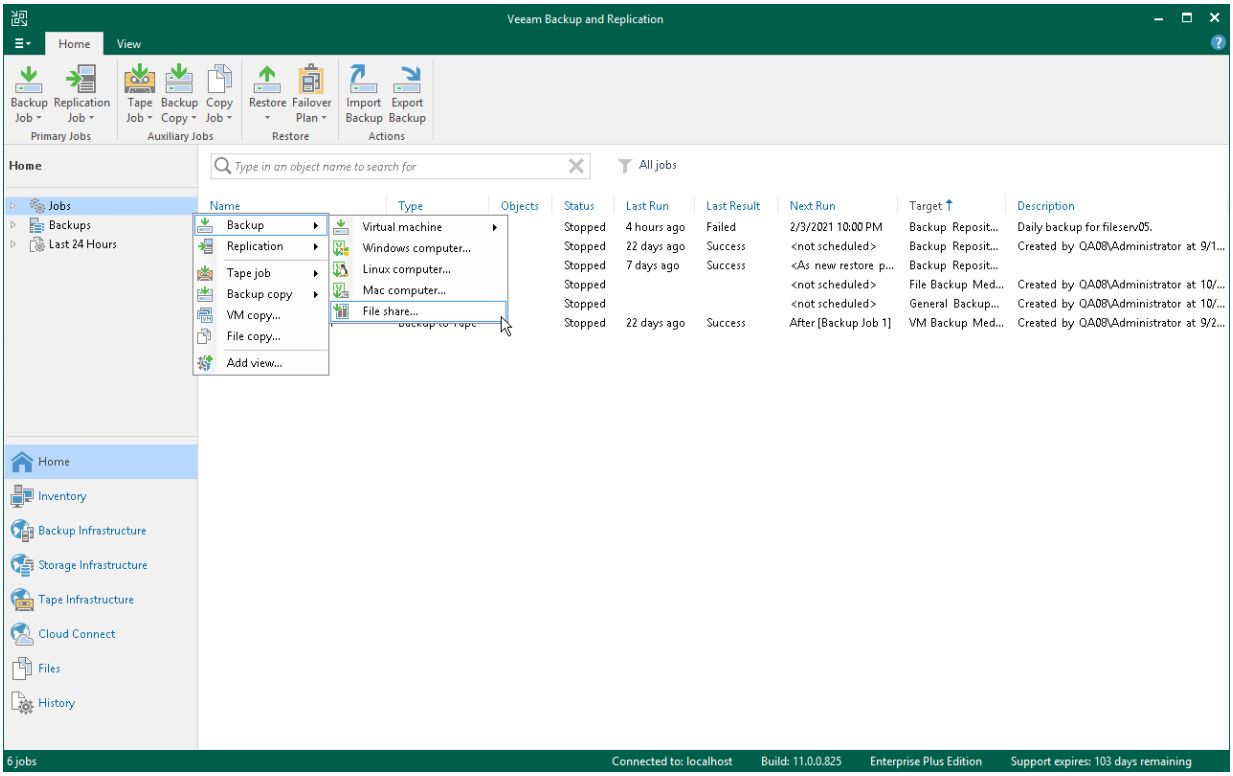
  For more information on how to rescan backup repositories, see Rescanning Backup Repositories.
- If you plan to use pre-job and/or post-job scripts, you must create scripts before you configure the file share backup job.
- For more information, see Storage Repositories in the **NAS Backup Support** section.

Antivirus software may significantly slow down file share backup jobs. To improve performance, we recommend you exclude the **c:\Program Files (x86)\Veeam\Backup Transport\x64\VeeamAgent.exe** process from the antivirus scan on machines running the file backup proxy and backup repository roles. Keep in mind that it can weaken the security of these machines.

Step 1. Launch New File Backup Job Wizard

To launch the **New File Share Backup Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Job** > **File Share**.
- Open the **Inventory** view. In the inventory pane, click **File Shares**. In the working area, click **Create Job**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **Backup** > **File Share**.
- You can quickly add the file share to an already existing job. Open the **Inventory** view. Under the **File Shares** node in the inventory pane, select type of the file share you want to back up. In the working area, right-click the file share you want to back up and select **Add to backup job** > *name of the job*.

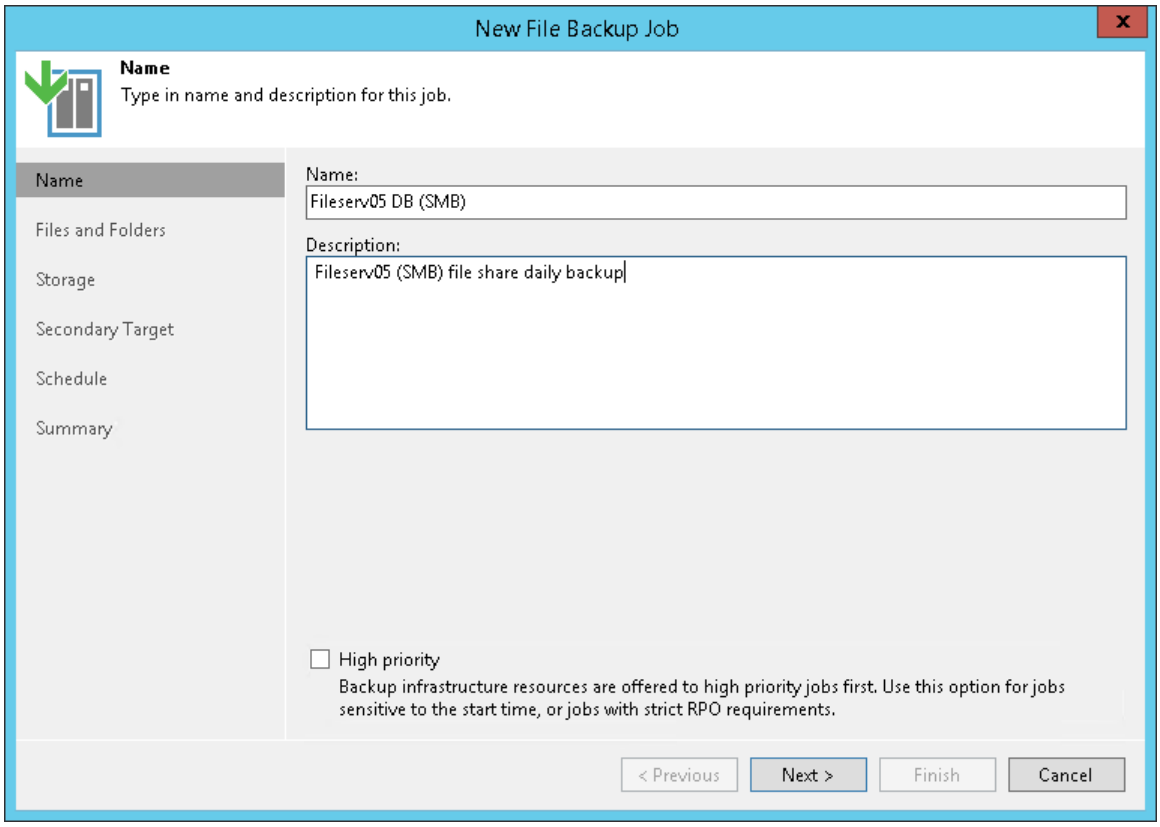


Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the file share backup job.

1. In the **Name** field, enter a name for the file share 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 Files and Folders to Back Up

At the **Files and Folders** step of the wizard, select files and folders that you want to back up.

1. Click **Add**.
2. From the **Server** list, select a file share on which the necessary files or folders reside.

   > **NOTE**
   >
   > If you plan to protect file shares residing on the enterprise storage system, you can choose between two different approaches. For more information, see NAS Backup Integration with Storage Systems.

3. In the **Folders** tree, select folders you want to backup.

   > **NOTE**
   >
   > If you select a NAS filer from the **Server** list, but the **Folders** tree is empty, make sure that the storage system rescan was performed and finished. For more information about the storage system rescan process, see Storage Discovery Process.

   To select multiple folders, hold [Ctrl] and click necessary folders. Although different folders of the same share form separate records in the table, they will be processed by one job task.

   Mind the following:

   - Hard links in file shares are protected with content included.
   - Symbolic links in file shares are protected as links, without the content they refer to.

4. In the **Files and Folders** screen, use the **Up** and **Down** buttons on the right to move sources up or down. The sources of the job are processed in the order in which the list displays them.
5. If you add a folder to the job, all the folder contents will be processed.

   If necessary, you can choose only specific files from the added folder.

Including/Excluding Files and Folders

To filter files and folders that you want to back up from the selected folder, you can specify the file name and extension masks.

To specify filters:

1. Select a file share folder in the list and click **Advanced**.
2. In the **File Filters** window, use **Include masks** and **Exclude masks** fields to filter the folder contents:
   - To include content, you can create name masks or use exact file names. For example:
     - To backup only PDF files, enter `*.PDF` in the **Include masks** field and click **Add**.
     - To backup all files with name **sales_report.xlsx**, enter `sales_report.xlsx` in the **Include masks** field and click **Add**.
   - To exclude content, you can use name masks, exact file names, or specify the path to the folder or the whole file share to exclude. For example:
     - To exclude PDF files from processing, enter `*.PDF` in the **Exclude masks** field and click **Add**.

- To exclude all files with name **pricelist.xlsx** from processing, enter `pricelist.xlsx` in the **Exclude masks** field and click **Add**.
- To exclude folder **2016** from processing, you can enter the full path to it in the **Exclude masks** field and click **Add**. For example, for an NFS file share this path looks like: `QA04:/NFS04/Documents/2016`, for an SMB file share — `\\fileserv05\Documents\2016`).

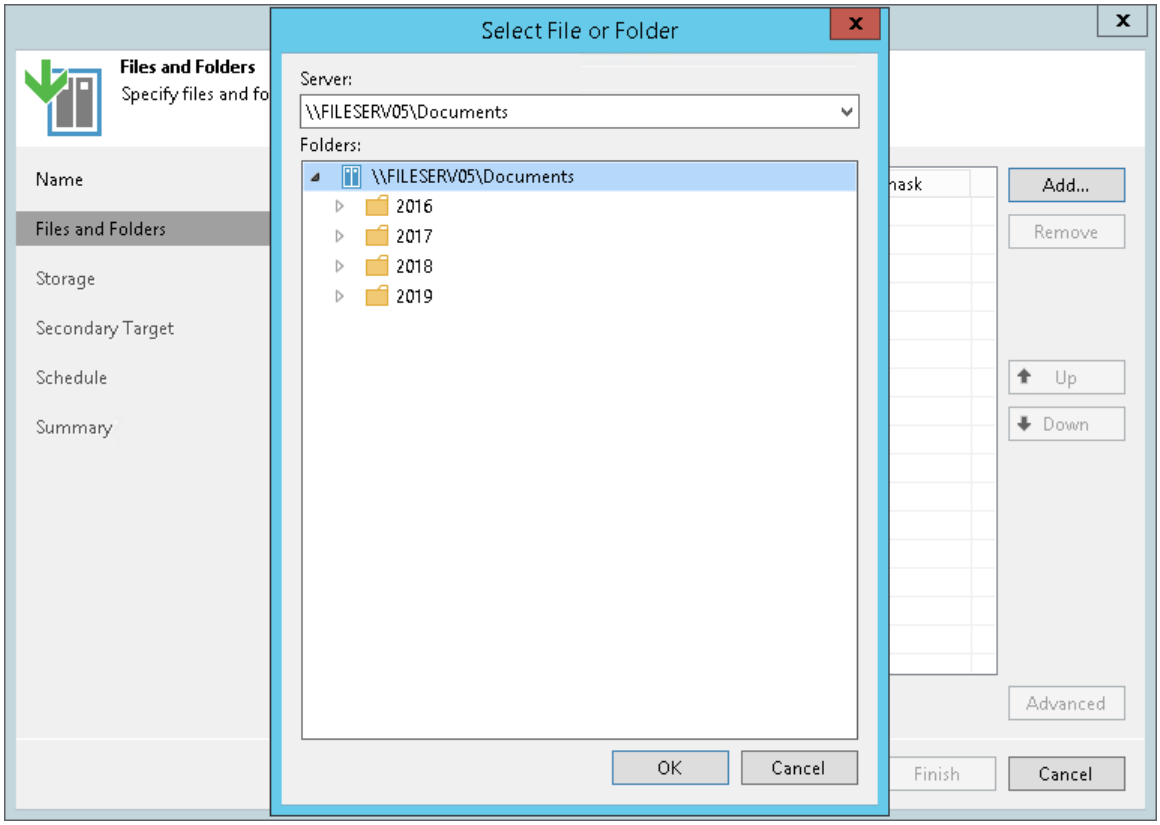
  Alternatively, you can specify a relative path to the folder or file to exclude. In case of the NAS filer, this is the only option to specify a path to exclude. For example, to exclude folders **call_records** (where call records are located) from all file shares residing on the NAS filer, specify `/call_records` (this mask will exclude the **call_records** folder from processing for all NFS file shares on this NAS filer) and `\call_records` (this mask will exclude the **call_records** folder from processing for all SMB file shares on this NAS filer).

  You can exclude a whole file share from processing. For example, you add the `\\Server\` SMB server to the file share backup job, but you want to exclude the `\\Server\Sharing` file share from processing. To exclude this file share from processing, enter the `\\Server\Sharing` path in the **Exclude masks** field and click **Add**. The job will back up all file shares that Veeam Backup & Replication detects on this server, but will skip the excluded file share.

NOTE
Mind the following:

- Include and exclude masks are case sensitive.
- You cannot exclude a whole file share from processing if the storage system, where file shares reside, is added to Veeam Backup & Replication as a NAS filer.

  As a workaround, you can exclude the volume(s) in the Access Options of the NAS filer itself.
- You cannot use mask with * to specify folders to exclude from processing. For example, mask `QA04:/NFS04/Documents/201*` will not work.
- You cannot mix different exclusion options, for example, you cannot use a mask to exclude files with certain extensions from the specific folder. For example, `QA04:/NFS04/Documents/2016/*`.xlsx will not work.



Step 4. Define Target Backup Storage Settings

At the **Storage** step of the wizard, define the backup and archive repositories where the file share backup job must store backup files, and settings for moving files and folders to these repositories. To learn what storage types you can assign the role of the backup or archive repository to, see NAS Backup Support.

> **NOTE**
>
> Mind that if you use the option of limiting the number of file versions to keep that is configured in File Version Settings, Veeam Backup & Replication first applies those file-version retention settings and only after that applies time-based retention settings specified at this step.

1. From the **Backup repository** drop-down list, select a repository where backup files must be stored. When you select a backup repository, Veeam Backup & Replication automatically checks the amount of free space left. Make sure that you have enough free space to store backups.
2. You can map the job to a specific backup stored in the backup repository. Backup job mapping allows you to move backup files to a new backup repository and to point the job to existing backups on this new backup repository. You can also use backup job mapping if the configuration database got corrupted and you need to reconfigure backup job settings.

   To map the job to a backup, click the **Map backup** link. In the opened **Select Backup** window, select a 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.
3. Use the **Keep all file versions for the last** field to specify how long copies of all recent file versions in the selected file share must be kept in the backup repository. You can restore the entire file share to any restore point within the period specified in this setting.

   If, for example, **Keep all file versions for the last** is set to 30 days, the backup repository will store all file versions that appeared at the file share during the last 30 days. At the scheduled time on the 31st day, the file share backup job first backs up new file versions and saves them to the backup repository. Right after that, file versions older than 30 days (created on the 1st day) are either deleted from the backup repository or moved to the archive repository if the **Keep previous file version for** check box is selected and the archive retention is configured.
4. If you need to keep versions of some files for a longer time after they are moved from the backup repository, you can configure archiving options to move file versions to a cheaper archive storage, for example, an object storage.

   To use the archive repository, select the **Keep previous file versions for** check box. In the fields on the right, specify the number of months or years during which backup files must be retained.

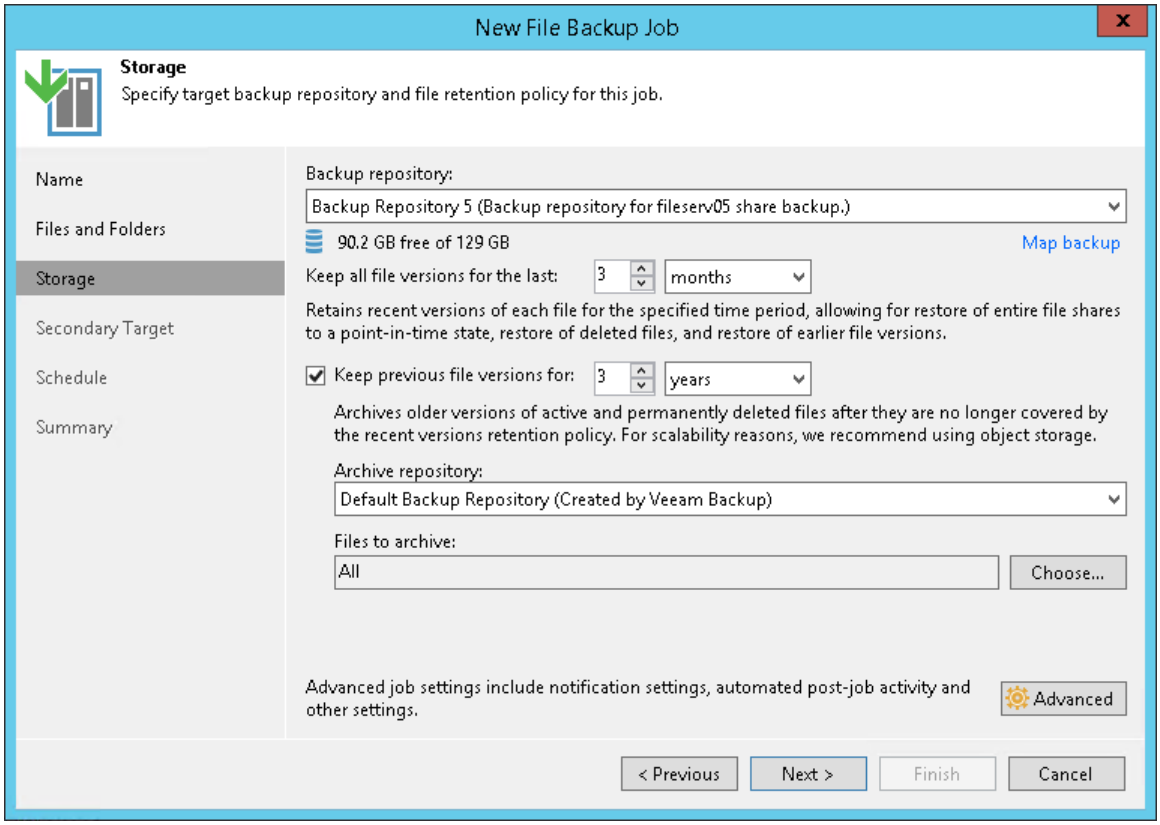
   This setting denotes a time period starting from the creation of the backup files in the backup repository, not from the moment when the file versions are moved from the backup repository to the archive repository.
5. From the **Archive repository** drop-down list, select the storage to be used as a repository to store archived files and folders.

   By default, all files deleted from the backup repository will be moved to the archive repository. If you do not need all the files in the archive, you can choose what files to keep.
6. To specify what files must be archived or excluded from the archive, do the following:

   a. Click **Choose** to open the **File Archive Settings** window.

   b. Under **Files to archive**, specify what files must be archived:

- **All files** – select this option to archive all files moved from the backup repository to the archive repository.
- **All files except the following extensions** – select this option to exclude files with certain extensions from the selection to be archived. Specify extensions for files to exclude from the selection. Files with the specified extensions will not be copied from the backup repository to the archive repository.
- **Files with the following extensions only** – select this option to archive files with certain extensions only. Specify extensions for files to archive. Files with these extensions will be copied from the backup repository to the archive repository.



Step 5. Specify Advanced Backup Settings

At the **Storage** step of the wizard, specify advanced settings for the file share backup job:

- File version settings
- ACL handling settings
- Storage settings
- Maintenance settings
- Script settings
- Notification settings

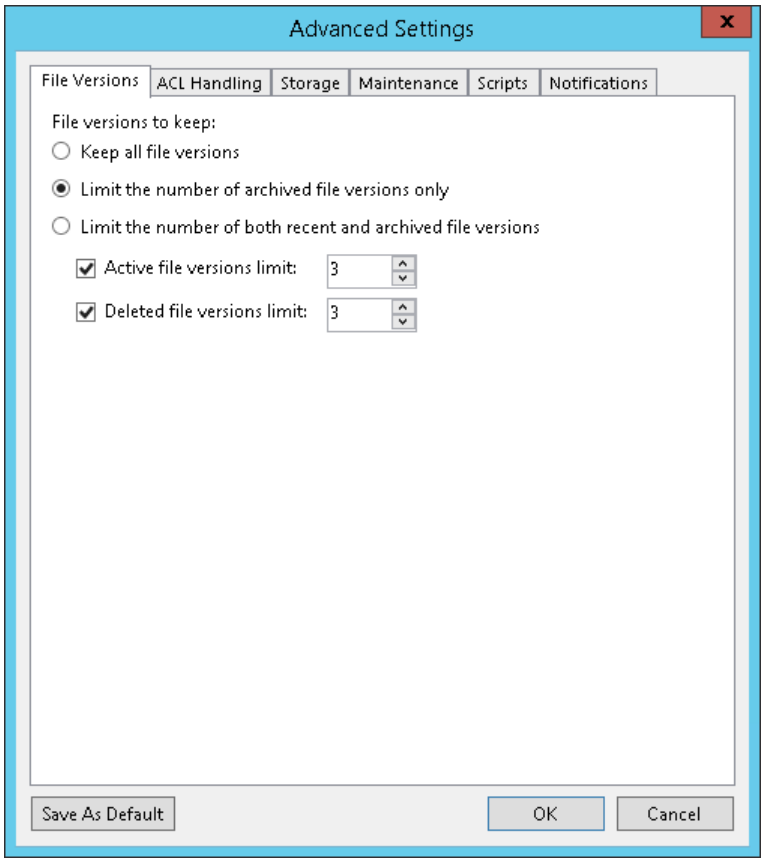
> **TIP**
>
> After you specify necessary settings for the backup job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup job, Veeam Backup & Replication will automatically apply the default settings to the new job.

File Version Settings

To configure how many file versions should be kept for protected files:

1. At the **Storage** step of the wizard, click **Advanced**.
2. On the **File Versions** tab, specify to what file versions the settings should apply:
   - Select **Keep all file versions** to keep all file versions for the time period specified in the main window at the **Storage** step.
   - Select **Limit the number of archived file versions only** to limit archived file versions to the numbers specified below.
   - Select **Limit the number of both recent and archived file versions** to limit recent and archived file versions to the numbers specified below.
3. After you choose what file versions to keep, specify how many file versions to keep:
   - Select **Active file versions limit** to keep the specified number of versions for files currently existing in the source file share. Specify how many versions of active files copied from the backup repository to store.

- Select **Deleted file versions limit** to keep the specified number of versions for files deleted from the source file share. Specify how many versions of deleted files copied from the backup repository to store.



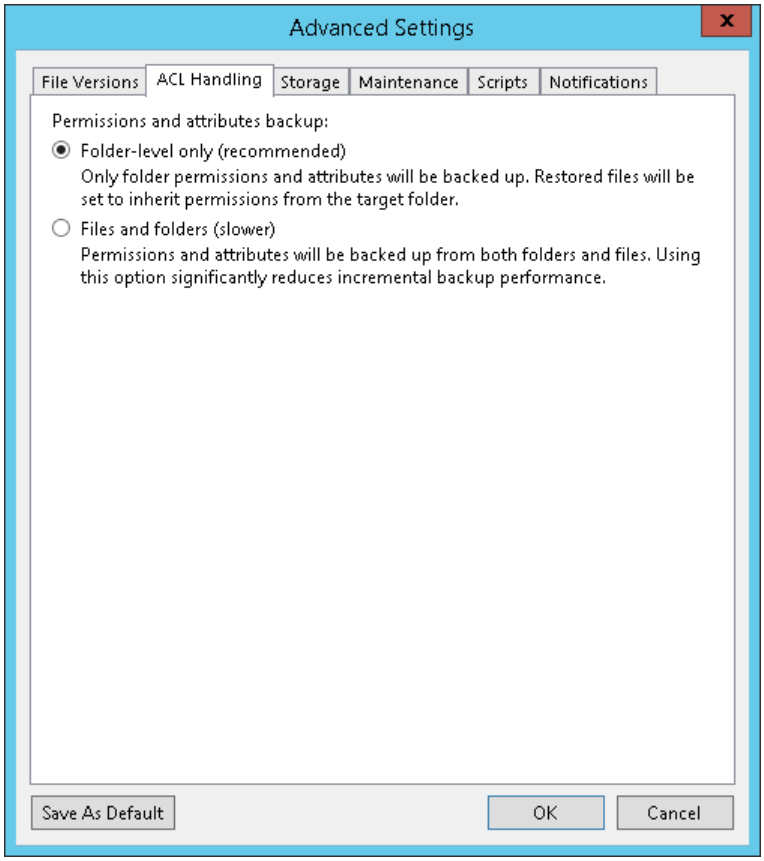
ACL Handling Settings

To specify how the backup job will process permissions and attributes:

1. At the **Storage** step of the wizard, click **Advanced**.
2. On the **ACL Handling** tab, specify how the backup job will process permissions and attributes:
   - Select **Folder-level only (recommended)** to back up permissions and attributes from folders only. The restored files will inherit permissions from the target folder.
   - Select **Files and folders (slower)** to back up permissions and attributes from both folders and individual files. This option can significantly reduce the backup performance.

NOTE

Mind that Veeam Backup & Replication does not collect ACL handling settings of the source file share root folder, so you cannot restore them. Before restoring an entire file share, you will have to specify required ACL handling settings for the root folder of the target file share.



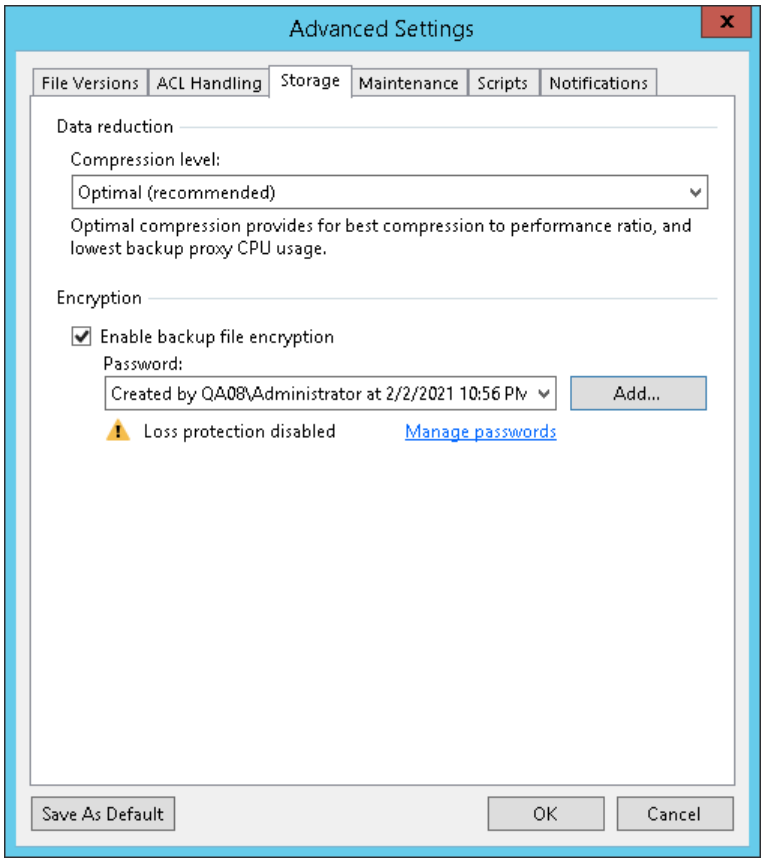
Storage Settings

To specify advanced storage settings for the file share backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. On the **Storage** tab, specify data reduction and encryption settings:
   - From the **Compression level** list, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.
   - 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 and does not have the Veeam Universal License or a legacy socket-based Enterprise or Enterprise Plus license installed, 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 back up all the files of the file share to a new backup file irrespective of whether they changed or not. The created backup files and subsequent backup files will be encrypted with the specified password.
>
> If you enable encryption for an existing job, Veeam Backup & Replication does not encrypt the previous backup chain created with this job.



Maintenance Settings

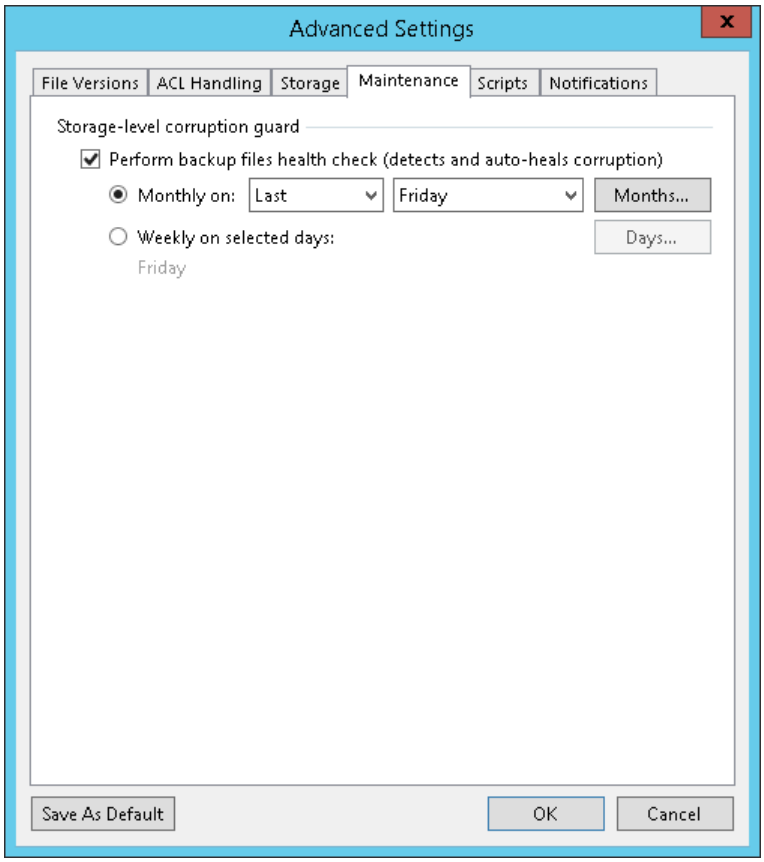
You can instruct Veeam Backup & Replication to periodically perform a health check for the latest restore point in the backup chain. The health check helps make sure that the restore point is consistent, and you will be able to restore data from this restore point.

During the health check, Veeam Backup & Replication performs a CRC check for metadata and a hash check for data blocks in the file share backup files to verify their integrity. For more information, see Health Check for File Share Backup Files.

To configure the health-check settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. On the **Maintenance** tab, select **Perform backup files health check** to enable the health check option. It allows ensuring that all data and metadata is backed up correctly.

3. Specify the time schedule for the health check.



Script Settings

You can configure custom scripts to run before or after the file share backup job. For example, you can configure scripts to take a VSS snapshot before running the job and to delete it after completing the job.

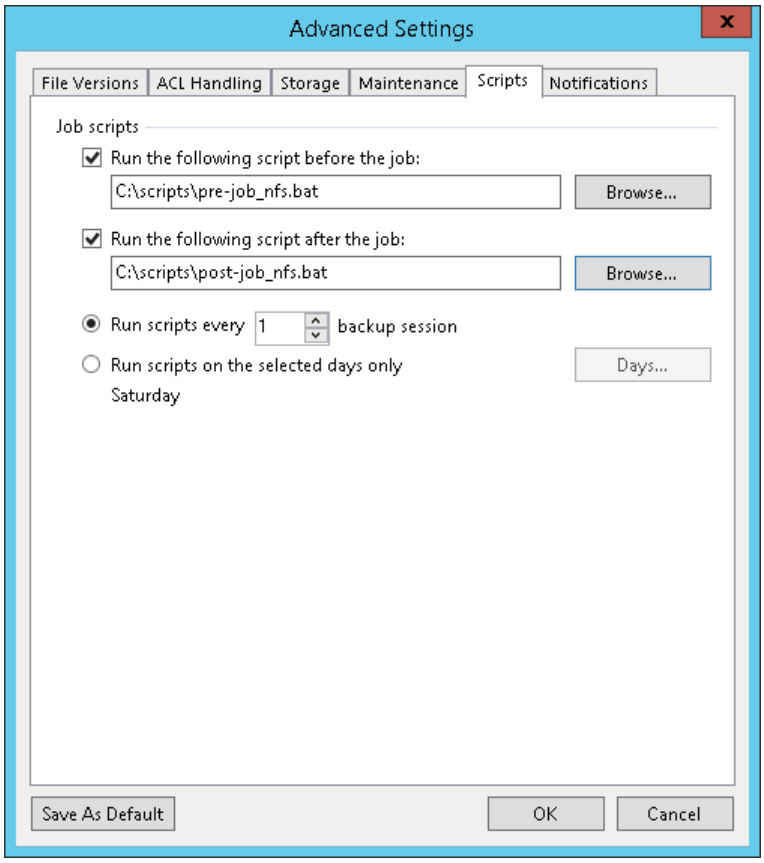
To specify script settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. If you want to execute custom scripts, 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.



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**. 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 file shares in the job) and *%Issues%* (number of files shares in the job that have finished with the *Warning* or *Failed* status).

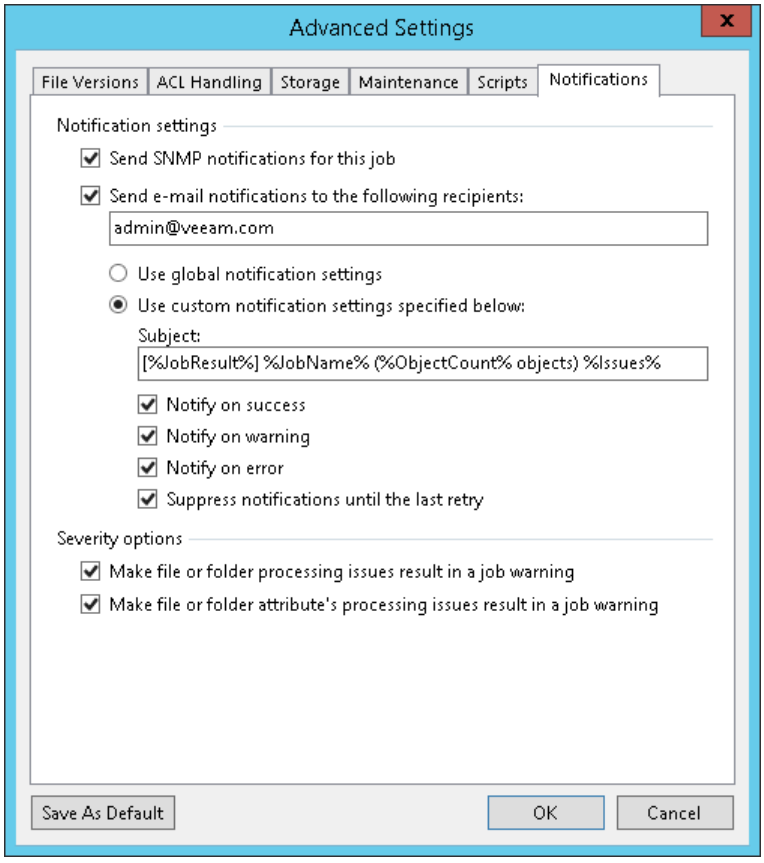
  ii. Select the **Notify on success**, **Notify on warning**, and/or **Notify on error** check boxes to receive email notification if the job completes successfully, completes with a warning or fails.

  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.

6. Specify severity options if required:

   - Select the **Make file or folder processing issues result in a job warning** to receive a warning at the end of the job processing session if any issues with file or folder processing occur.

   - Select the **Make file or folder attribute's processing issues result in a job warning** to receive a warning at the end of the job processing session if any issues with processing of file or folder attributes occur.



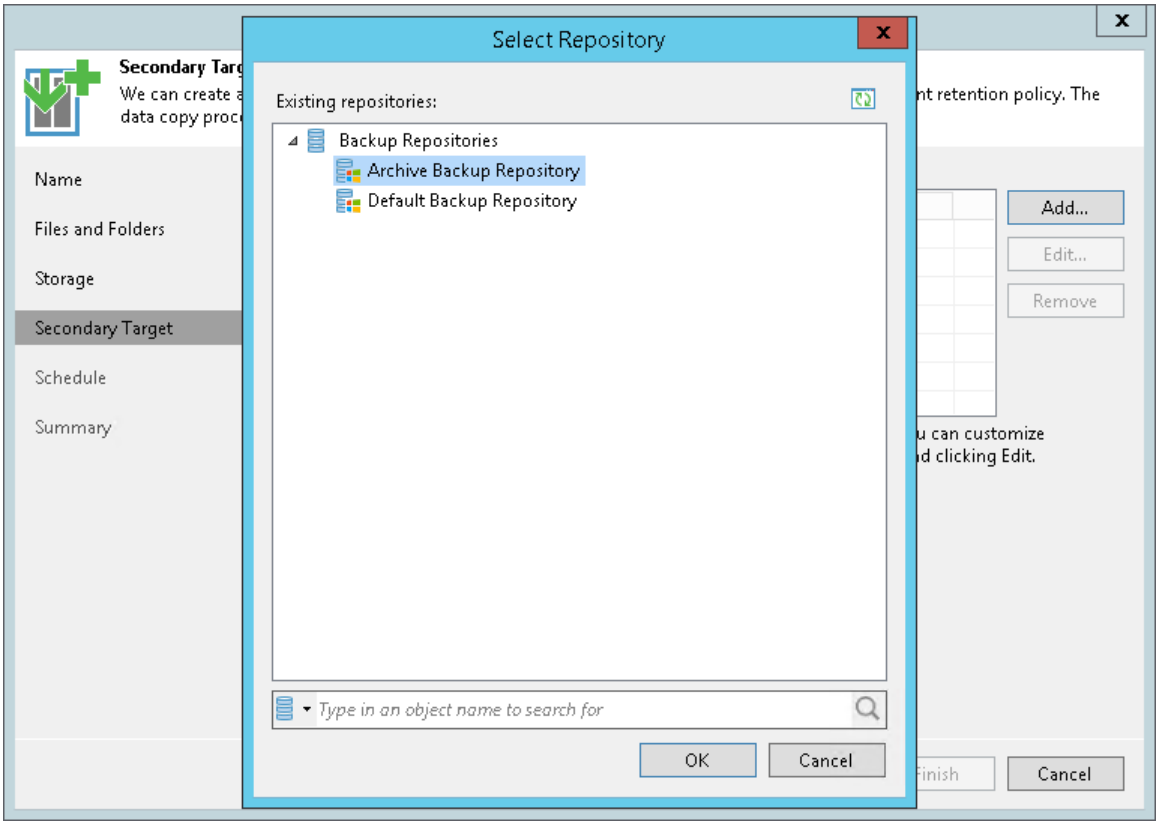
Step 6. Specify Secondary Repository

At the **Secondary Target** step of the wizard, you can specify a secondary repository that will be used to store additional copies of backup files from the backup storage for redundancy. If you add a secondary repository, Veeam Backup & Replication will create a separate job for backup copy to it. The data copy process will start automatically after each primary job runs.

> **NOTE**
>
> To learn what storage types you can assign the role of the secondary repository to, see NAS Backup Support.

To add a secondary repository:

1. Click **Add**.
2. From the list of existing repositories, select a repository that will keep additional copy of the backup files. You can add several secondary repositories for copying files of the primary backup job. To quickly find the repository, use the search field at the bottom of the wizard.



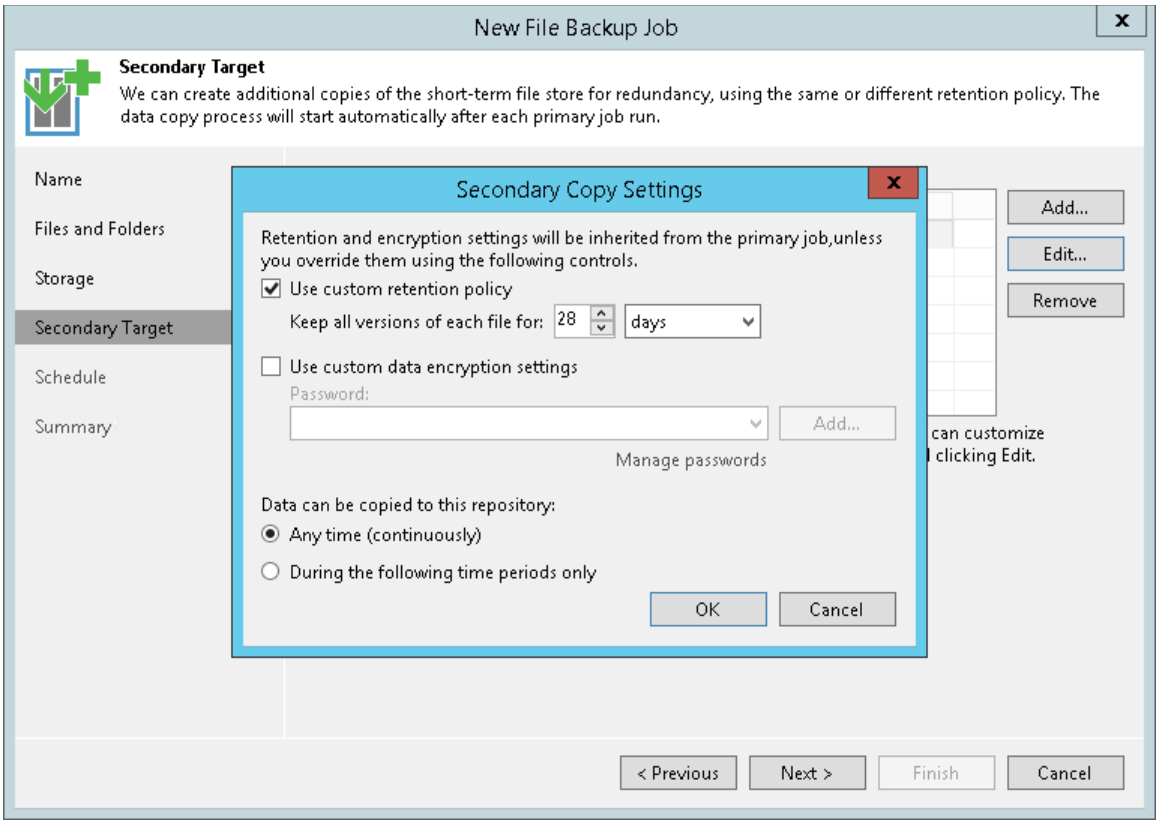
3. By default, retention and encryption settings for the secondary target repository are inherited from the primary job. To customize them, select the necessary repository in the **Secondary repositories** list and click **Edit**.
   - To enable custom retention settings:
     i. Select **Use custom retention policy**.
     ii. Specify how long all versions of each file will be kept in the secondary repository.
   - To specify encryption settings different from those of the primary repository:

i. Select **Use custom data encryption settings**.

ii. 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 and does not have the Veeam Universal License or a legacy socket-based Enterprise or Enterprise Plus license installed, 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.

- Configure time intervals at which the data can be copied to the secondary repository.
  - If you select the **Any time (continuously)** option, Veeam Backup & Replication will copy backup files to the secondary repository as soon as the primary file share backup job completes.
  - If you want to specify time periods when copying backup to the secondary repository is permitted, select the **During the following time periods only** option and configure allowed and prohibited hours.



Step 7. 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 file share 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 file shares 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.

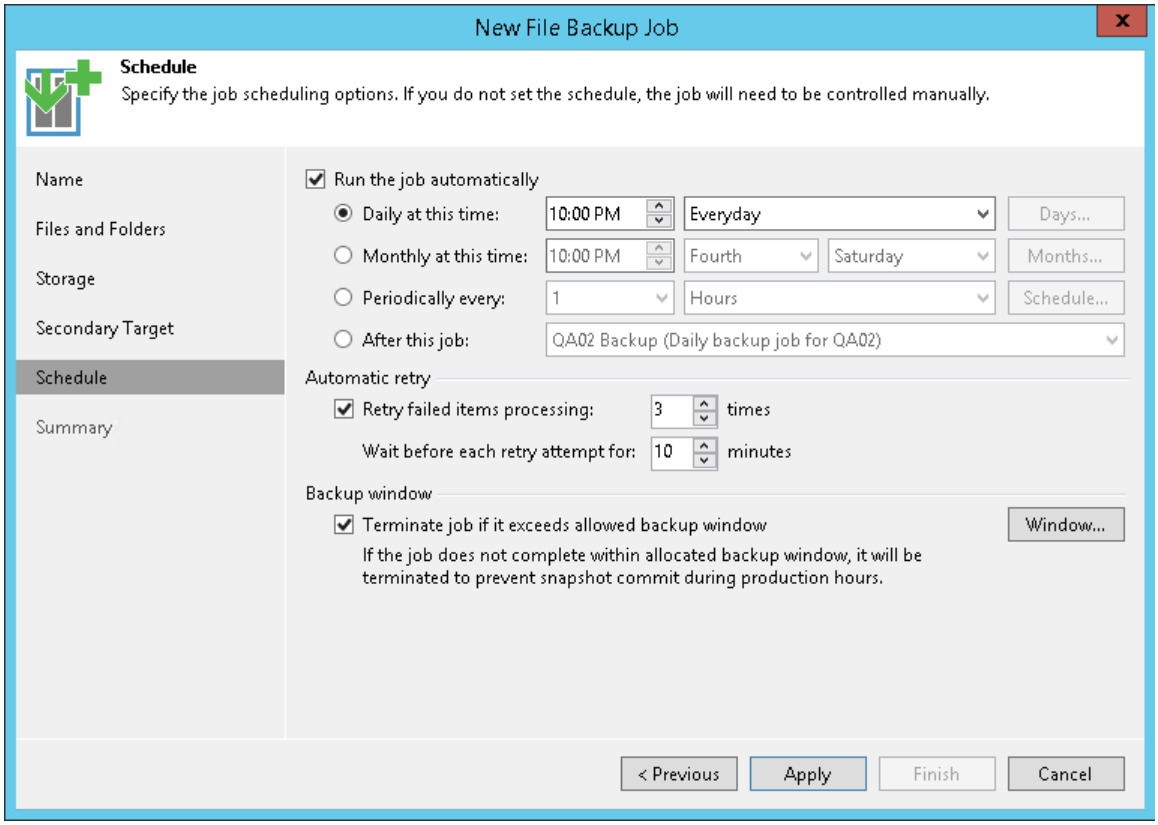
4. In the **Backup window** section, define the time interval within which the backup job must complete. The backup window prevents the job from overlapping with production hours and ensures that the job does not provide unwanted overhead on the production environment. To set up a backup window for the 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

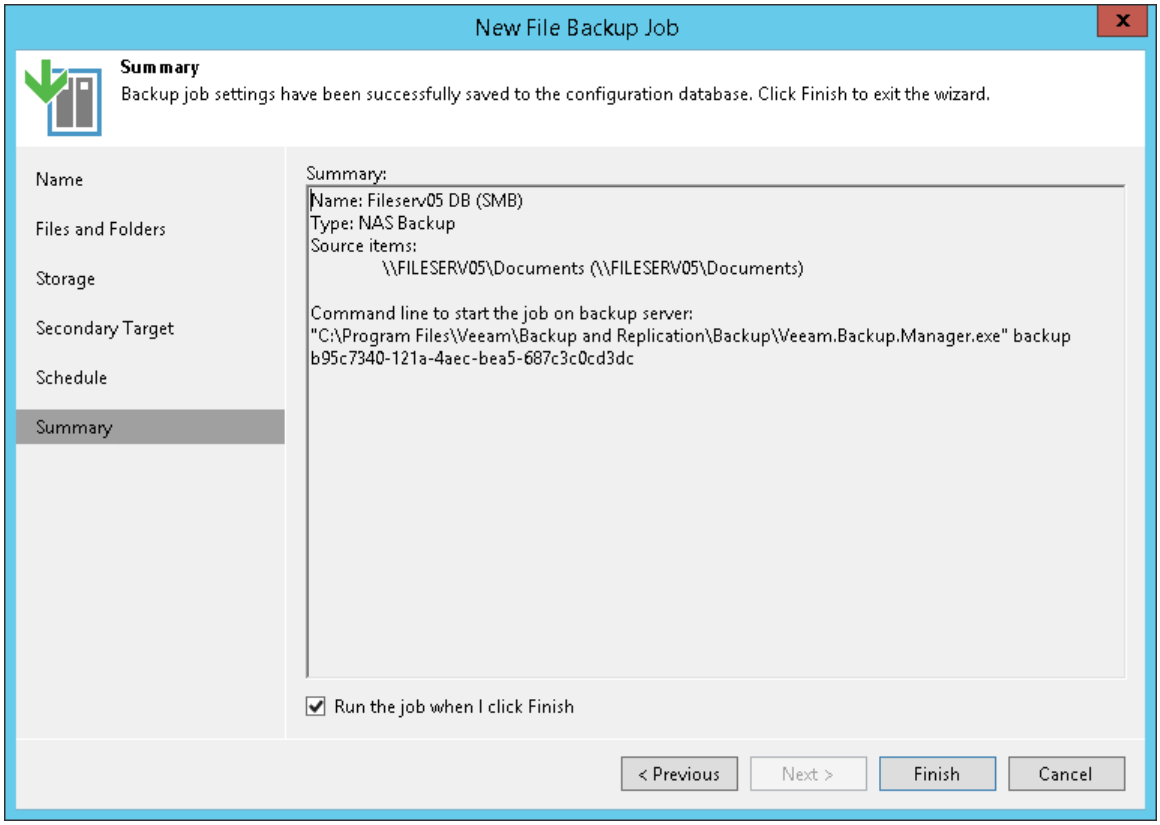
The **After this job** function will automatically start a job if the first job in the chain is started automatically by schedule. If you start the first job manually, Veeam Backup & Replication will display a notification. You will be able to choose whether Veeam Backup & Replication must start the chained job as well.



Step 8. 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.

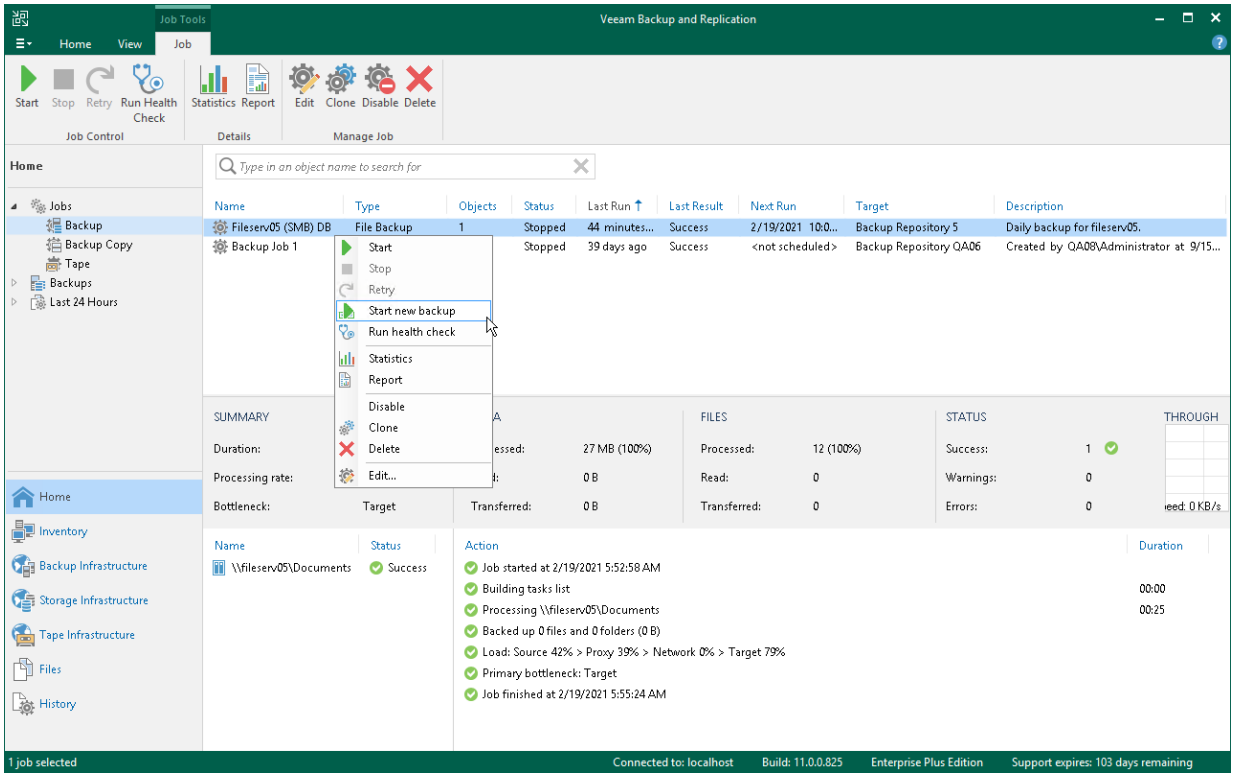


Starting New Backup Chain

You can start a new backup chain for the file share backup job. Veeam Backup & Replication then creates a new active full backup that starts the new chain for the entire file share. All existing backup files are moved to the **Disk (Imported)** node under the **Backups** node in the Veeam Backup & Replication Console. Data files are stored to the same folder in the backup repository. The data files for the new backup chain are stored to a new separate folder in the backup repository.

To start a new backup chain:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select a job of the **File Backup** type, press and hold the [Ctrl] key, right-click the job and select **Start new backup**.



Performing Health Check and Repair for File Share Backup Files

In this section you will learn how to perform:

- Health check for file share backup files
- Repair of file share backup files

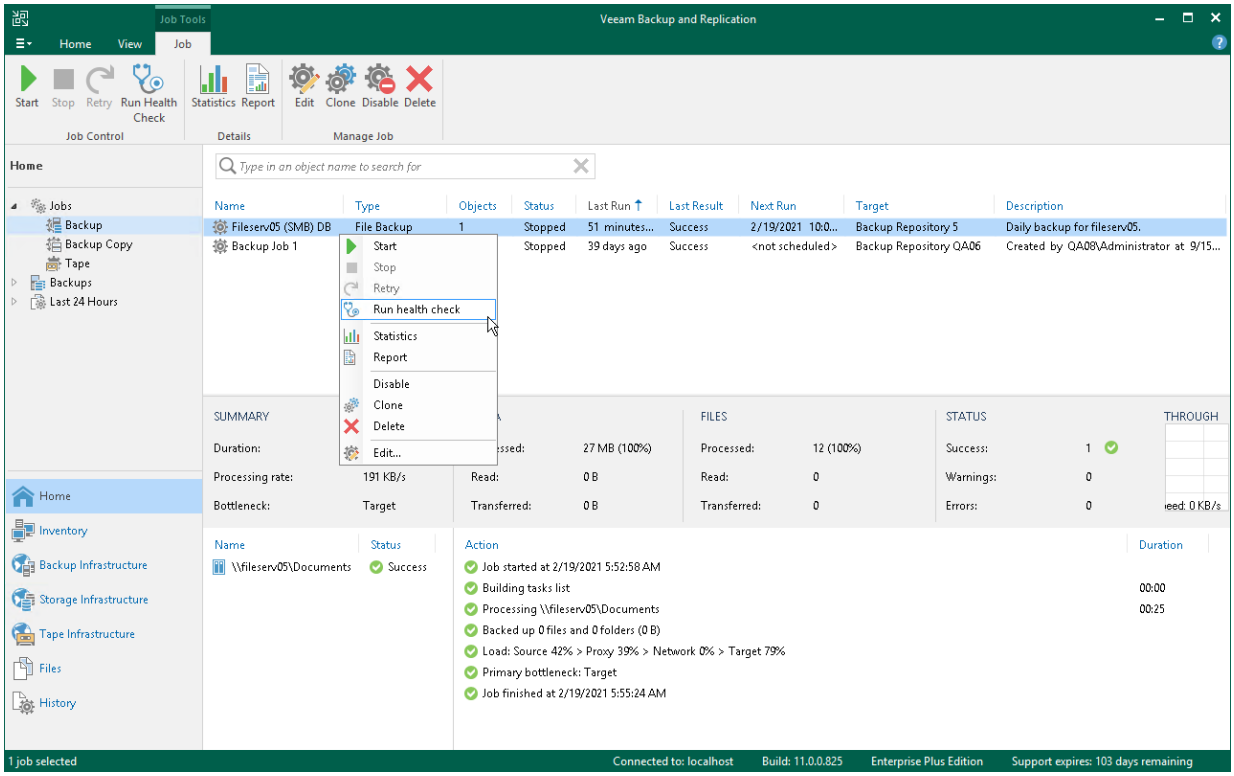
Health Check for File Share Backup Files

You can manually 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 data blocks in backup files 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.

To run the health check:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select a job of the **File Backup** type and click **Run Health Check** on the ribbon or right-click the job and select **Run health check**.

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. To learn how to configure periodic health check, see Maintenance Settings.

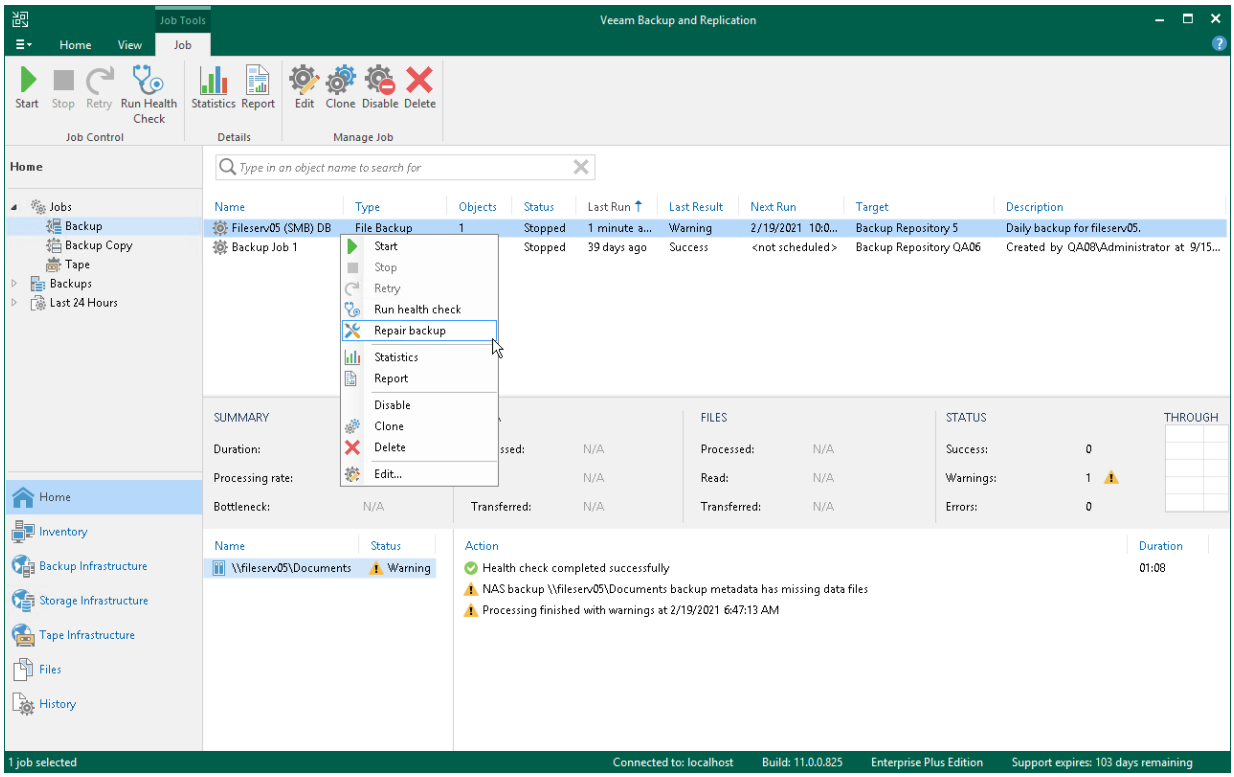


Repair of File Share Backup Files

If during the health check Veeam Backup & Replication detects some inconsistency in the file share backup files, you can run the backup repair procedure to fix the issues.

To run the backup repair:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select a job of the **File Backup** type, right-click the job and select **Repair backup**.



Viewing File Share Backup Properties

You can view summary information about the file share backup. The summary information provides the following data:

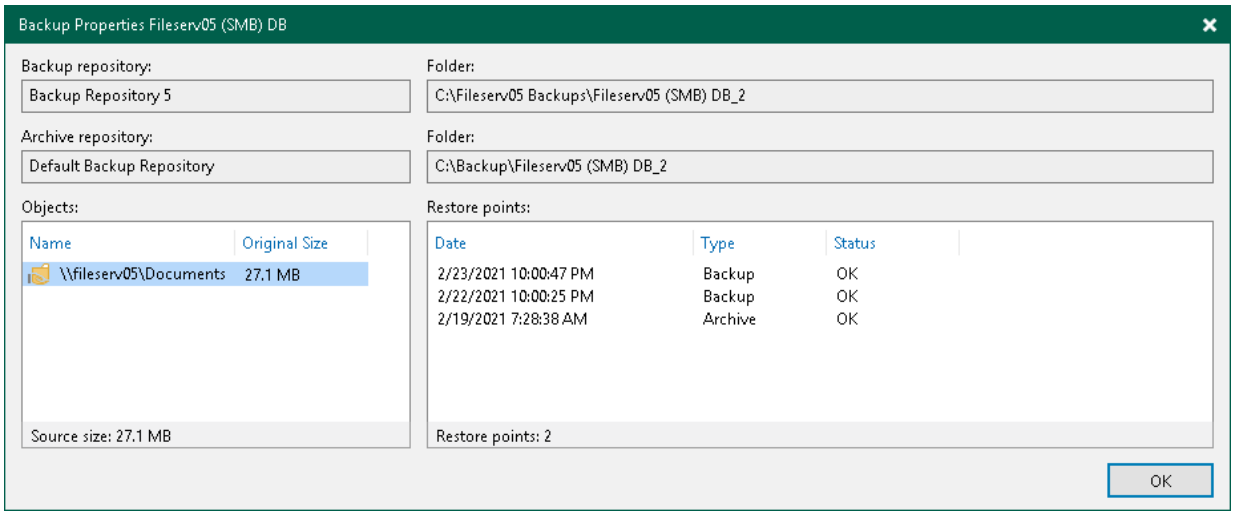
- Name and path to the backup repository that stores file share backup files.
- Name and path to the archive repository that stores archived backup data.
- Path to the file share backup source and its original size.
- Available restore points: date of their creation, their type (**Backup** or **Archive**) and status.

  For the **Backup** type, the table shows all restore points stored in the backup repository. You can restore the file share to the state as of any of these points. To learn how to restore the file share, see File Share Data Recovery.

  For the **Archive** type, the table shows only a single record. The time stamp of this record denotes the date and time of the restore point, which was created in the backup repository and files of which were the latest to be moved to the archive repository according to retention settings. To learn how to restore files from the archive, see Restoring Backup Files from Archive Repository.

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.



Converting Backups from Non-Root to Root Shared Folders

Starting from version 11, Veeam Backup & Replication allows you to add a server root folder as a source for file share backup jobs. In this case, all changes to separate shared folders residing on this server will be reflected in the file share backup job where the root shared folder of this server is added. You can even add shared root folders using different protocols to one file share backup job and thus protect all file shares that are or will be added on the server.

If you previously had several separate non-root shared folders residing on the same server and want to switch to using a single root shared folder to cover the same shares, you do not have to run full backups to update data of protected shares. Instead, you can convert existing backups and update existing file share backup jobs to protect single root shared folders comprising all other non-root shared folders residing on the same server. Perform the conversion with extreme caution.

To convert backups from non-root to root shared folders, do the following:

1. Disable file share backup jobs protecting file shares, for which you want to convert backups. To do that, right-click the required job in the **Jobs** node of the inventory pane in the **Home** view and select **Disable**. Alternatively, you can click **Disable** on the ribbon.
2. Make sure that your backup infrastructure has a root share (for example, NFS or SMB) added for the whole server or storage system where existing non-root shares reside. These shares must reside on the same server or storage system. The correspondence of the shares must be full except for the host name.
3. Run the `Convert-VBRNASBackupRootFormat` PowerShell cmdlet to convert backups created by one file share backup job for separate non-root shared folders residing on the same server into the backup created for the server single root folder with all the non-root shared folders of the same type under it. For more information, see the description of the `Convert-VBRNASBackupRootFormat` cmdlet in the Veeam PowerShell Reference.

   As a result, the backup will be moved from **Backups** > **Disk** node to **Backups** > **Disk (Orphaned)** node in the inventory pane of the **Home** view.

   At this step, you can check if the cmdlet has correctly converted the backup. To do that, check if backup object names in the **Disk (Orphaned)** node have changed and now show the path to the server root folder. If object names have not changed and show the paths to multiple separate non-root shared folders as before, continuing the conversion process can lead to the unwanted result. For example, when you enable the backup job for the converted backup, it will back up all shared folders under root folder not with an incremental run, but with a full run instead, which may lead to extra costs.
4. Use the Edit File Backup Job wizard to edit the file backup job that protects the file shares:
   a. At the Files and Folders step of the wizard, remove the existing non-root shared folders from the job and add the server root folder instead.
   b. At the Storage step of the wizard, map the job to the backup that was converted at step 2.
5. Enable file share backup jobs protecting file shares, for which you converted backups. To do that, right-click the required job from the **Jobs** node of the inventory pane in the **Home** view and clear selection **Disable**. Alternatively, you can click **Disable** on the ribbon.

Converting Backups from SMB or NFS Shares to NAS Filer Shares

You can use enterprise storage systems integrated with Veeam Backup & Replication both to host simple SMB or NFS shares and to act as NAS filer shares.

To use all the advantages of NAS filer shares, for example the native file change tracking technology, you can convert backups created for existing SMB or NFS shares into the format of NAS filer shares. After that you can continue to protect the file shares as NAS filer shares by running existing file share backup jobs and by using existing backup files. Perform the conversion with extreme caution.

To convert SMB or NFS shares into NAS filer shares, do the following:

1. Disable file share backup jobs protecting SMB or NFS shares, for which you want to convert backups. To do that, right-click the required job in the **Jobs** node of the inventory pane in the **Home** view and select **Disable**. Alternatively, you can click **Disable** on the ribbon.
2. Make sure that you have created NAS filer, which corresponds to existing SMB or NFS shares, added to the Veeam Backup & Replication inventory. The NAS filer and these shares must reside on the same storage system. The correspondence of the shares must be full except for the host name.
3. Run the `Convert-VBRNASBackupSANFormat` PowerShell cmdlet to convert the format of the file share backup to provide support of NAS filer shares.

   ```
   $nasbackup = Get-VBRNASBackup -name "File Backup Job 1"
   $netapp = Get-NetAppHost -name "pdc-ontap-1"
   $netapp_filer = Get-VBRNASServer -SANEntity $netapp
   Convert-VBRNASBackupSANFormat -Backup $nasbackup -Server $netapp_filer
   ```

   For more information, see the description of the `Convert-VBRNASBackupSANFormat` cmdlet in the Veeam PowerShell Reference.

   As a result, the backup will be moved from **Backups** > **Disk** node to **Backups** > **Disk (Orphaned)** node in the inventory pane of the **Home** view.

   At this step, you can check if the cmdlet has correctly converted the backup. To do that, check if backup object names in the **Disk (Orphaned)** node have changed and now show the path to the NAS filer share. If object names have not changed and show the path to the SMB or NFS share as before, continuing the conversion process can lead to the unwanted result. For example, when you enable the backup job for the converted backup, it will back up the NAS filer share not with an incremental run, but with a full run instead, which may lead to extra costs.
4. Use the Edit File Backup Job wizard to edit the file backup job that protects the file shares:
   a. At the Files and Folders step of the wizard, remove the existing SMB and NFS shares from the job and add the corresponding NAS filer shares.
   b. At the Storage step of the wizard, map the job to the backup that was converted at step 2.
5. Enable file share backup jobs protecting file shares, for which you converted backups. To do that, right-click the required job from the **Jobs** node of the inventory pane in the **Home** view and clear selection of **Disable**. Alternatively, you can click **Disable** on the ribbon.

Starting from version 11a (build 11.0.1.1261), Veeam Backup & Replication supports conversion of backups created with backup copy jobs. To continue the old backup chain created with the backup copy job, do the following:

1. Disable file share backup jobs protecting SMB or NFS shares, for which you want to convert backups. To do that, right-click the required job in the **Jobs** node of the inventory pane in the **Home** view and select **Disable**. Alternatively, you can click **Disable** on the ribbon.
2. Make sure that you have created NAS filer, which corresponds to existing SMB or NFS shares, added to the Veeam Backup & Replication inventory. The NAS filer and these shares must reside on the same storage system. The correspondence of the shares must be full except for the host name.
3. Run the `Convert-VBRNASBackupSANFormat` PowerShell cmdlet to convert the format of the file share backup to provide support of NAS filer shares.

   ```
   $backup_copy = Get-VBRNASBackup -name "File Backup Job 1 (Copy)"
   $netapp = Get-NetAppHost -name "pdc-ontap-1"
   $netapp_filer = Get-VBRNASServer -SANEntity $netapp
   Convert-VBRNASBackupSANFormat -Backup $backup_copy -Server $netapp_filer
   ```

   For more information, see the description of the `Convert-VBRNASBackupSANFormat` cmdlet in the Veeam PowerShell Reference.
4. Use the Edit File Backup Job wizard to remove the secondary target storage for the file share backup job that protects the file shares:
   a. At the Secondary Target step of the wizard, remove the required repository selected as a secondary target for the file share backup job.
   b. Go through all the wizard steps without running the job. Click **Finish**.
5. Use the Edit File Backup Job wizard to add the secondary target storage back for the file share backup job that protects the file shares:
   a. At the Secondary Target step of the wizard, add the required repository as a secondary target for the file share backup job.
   b. Go through all the wizard steps. Click **Finish**.
6. Enable file share backup jobs protecting file shares, for which you converted backups. To do that, right-click the required job from the **Jobs** node of the inventory pane in the **Home** view and clear selection of **Disable**. Alternatively, you can click **Disable** on the ribbon.

   File share backup copy will automatically map to the file share backup copy job. After that, the backup copy job will back up new points of the main file share backup job if they were created.

File Share Data Recovery

You can restore data previously backed up with file share backup jobs. You can restore the following data:

- SMB file share files and folders
- NFS file share files and folders
- Files and folders of a managed Microsoft Windows server
- Files and folders of a managed Linux server

Veeam Backup & Replication offers several recovery options for different recovery scenarios:

- Instant file share recovery allows you to publish a point-in-time file share state as a read-only SMB file share to enable users to instantly access all protected files.
- Restore of the entire file share allows you to recover all files and folders of the file share to one of the restore points.
- Rollback to a point in time allows you to restore only changed files to one of the restore points.
- Restore of files and folders allows you to select files and folders to restore to one of the restore points.
- Restore of files from an archive repository allows you to select archived files to restore to one of the restore points.

Performing Instant File Share Recovery

You can use the instant file share recovery feature to publish a point-in-time file share state as a read-only SMB file share to enable users to instantly access all protected files.

Before you perform instant file share recovery, check prerequisites.

Before You Begin

Before you perform instant file share recovery, consider the following:

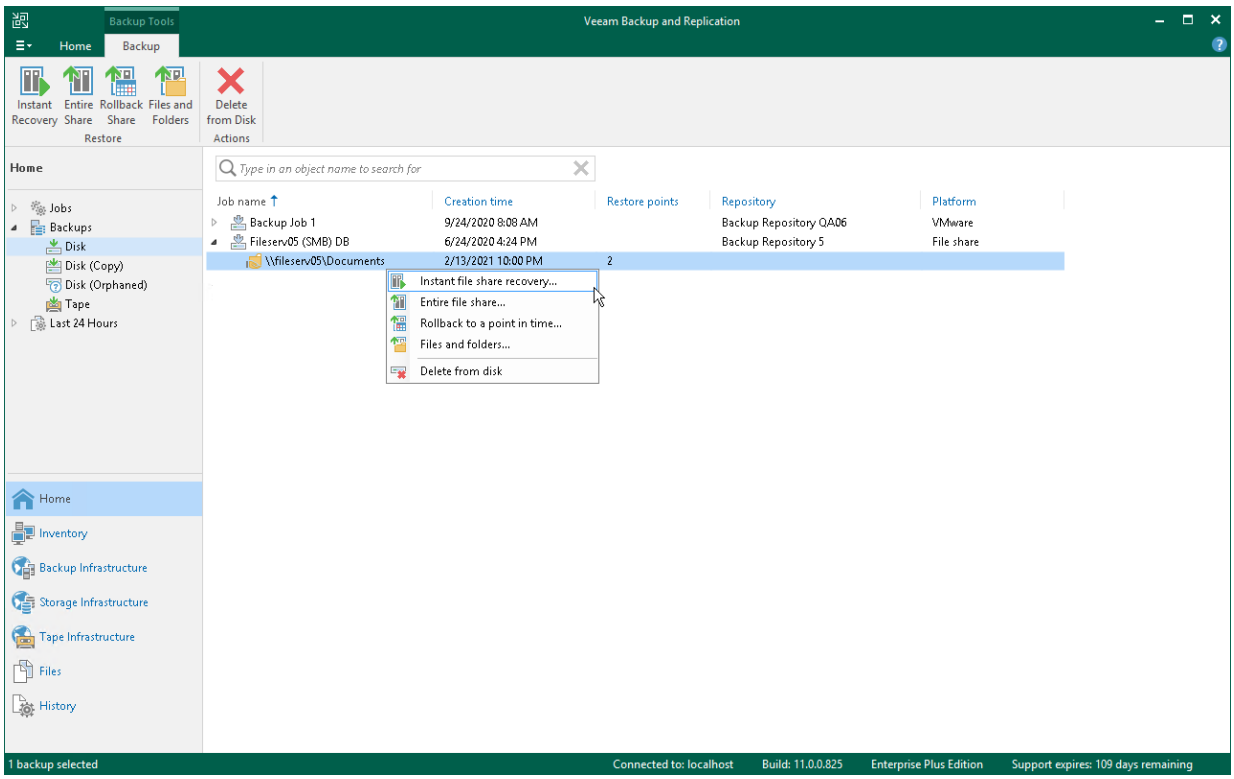
- Instant file share recovery is supported for SMB file shares only.
- File shares recovered with instant recovery are available in the read-only mode.

Step 1. Launch Instant File Share Recovery Wizard

To launch the **Instant File Share Recovery** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **File Share** > **Instant file share recovery**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share that you want to restore and click **Backup** > **Instant Recovery** on the ribbon.
  - Right-click the file share that you want to restore and select **Instant file share recovery**.

  You can perform the instant file share recovery by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane.



Step 2. Select File Share to Restore

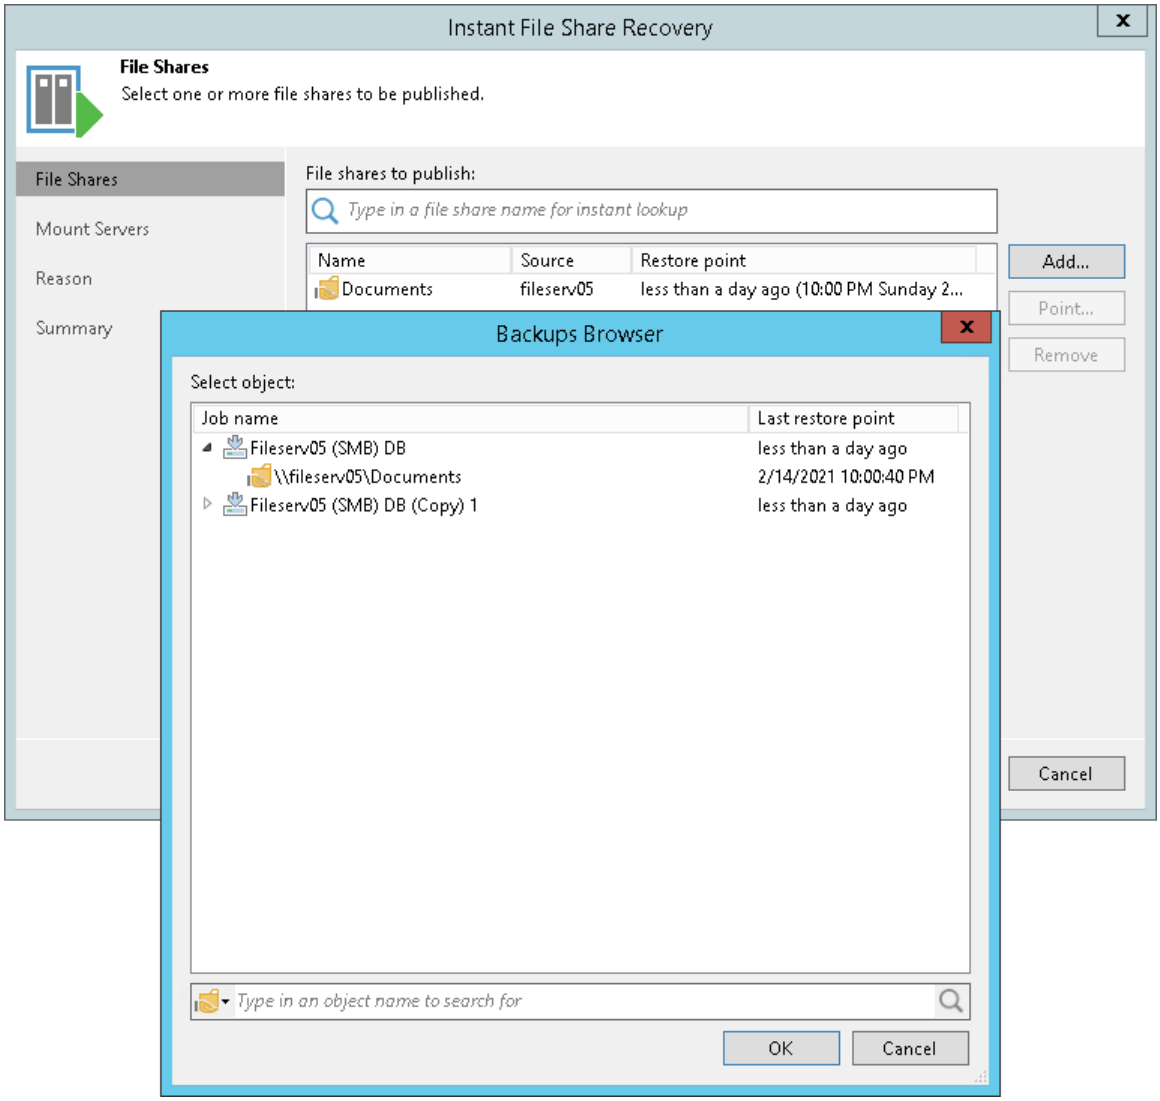
At the **File Shares** step of the wizard, select the file share for which you want to perform instant recovery:

1. Click **Add**.
2. In the **Backups Browser** window, expand the necessary backup job to select the required file share to restore.

To quickly find a file share, you can use the search field at the bottom of the window.

1. Enter a file share name or a part of it in the search field.
2. Press [Enter] to start the search.

Alternatively, you can use the **File shares to publish** search field to quickly search the required file share and add it to the list of file shares to publish.



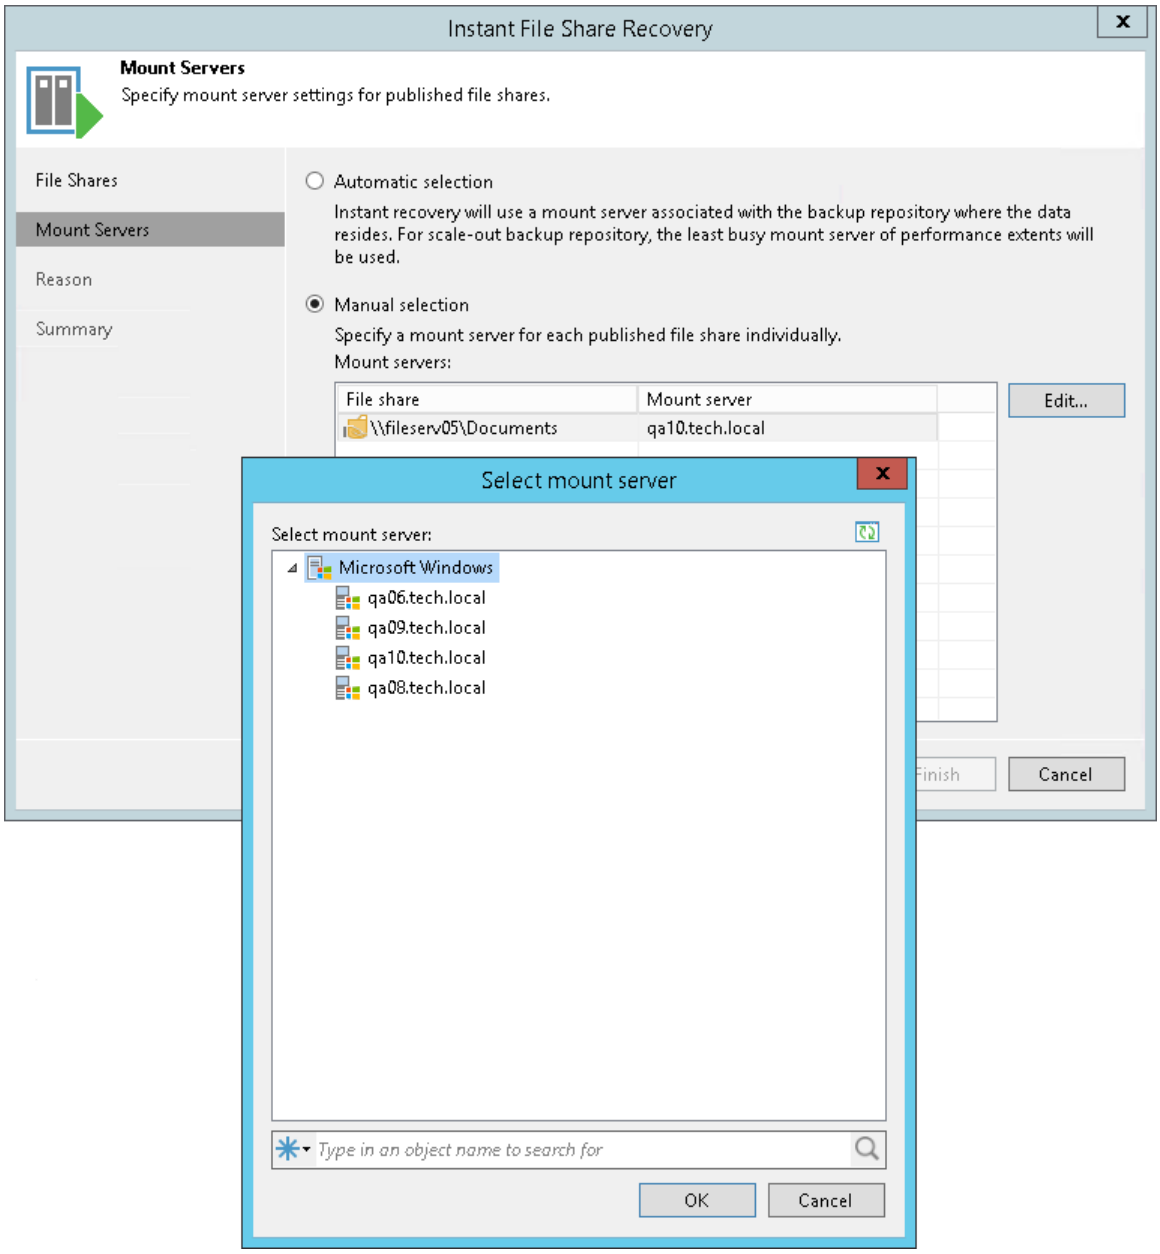
Step 3. Specify Mount Server Settings

At the **Mount Servers** step of the wizard, specify mount server settings for published file shares.

- If you select the **Automatic selection** option, Veeam Backup & Replication will automatically choose the mount server where to restore file shares. The job will use the mount server from the repository where the backup files reside. For scale-out backup repositories, Veeam Backup & Replication will use the least occupied mount server.
- If you select the **Manual selection** option, you can specify which mount server to use to individually publish each file share:
  a. In the **Mount servers** list, select a file share for which you want to assign a mount server.
  b. Click **Edit**. Alternatively you can double-click the required file share in the list.
  c. In the **Select mount server** window, select a mount server to use to publish the chosen file share.

     To quickly find a mount server, you can use the search field at the bottom of the window.

d. Click **OK** to confirm selection.

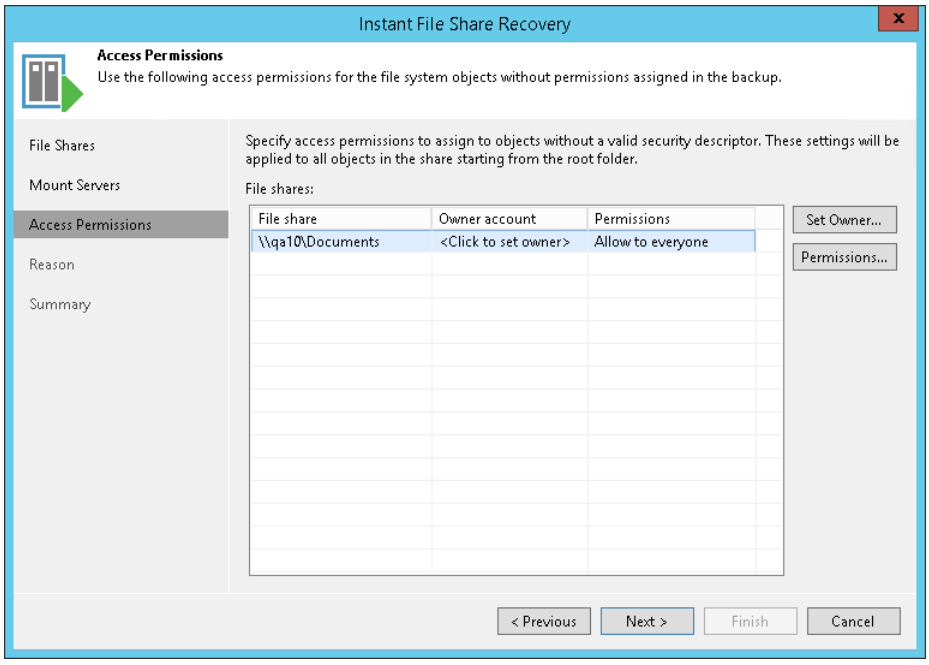


Step 4. Specify Access Permissions

After you specify file shares and mount servers, Veeam Backup & Replication validates them. If Veeam Backup & Replication detects missing security descriptors on the file shares, it adds the **Access Permissions** step to the wizard. At this step you can specify the owner account and permissions for the file share.

1. From the **File shares** list, select a file share for which you want to specify an owner account and permissions.
2. Click **Set Owner** and specify the owner account for the file share.
3. Click **Permissions** and configure access permissions for the file share. The following options are available:
   - Deny to everyone
   - Allow to everyone
   - Allow to the following accounts or groups only

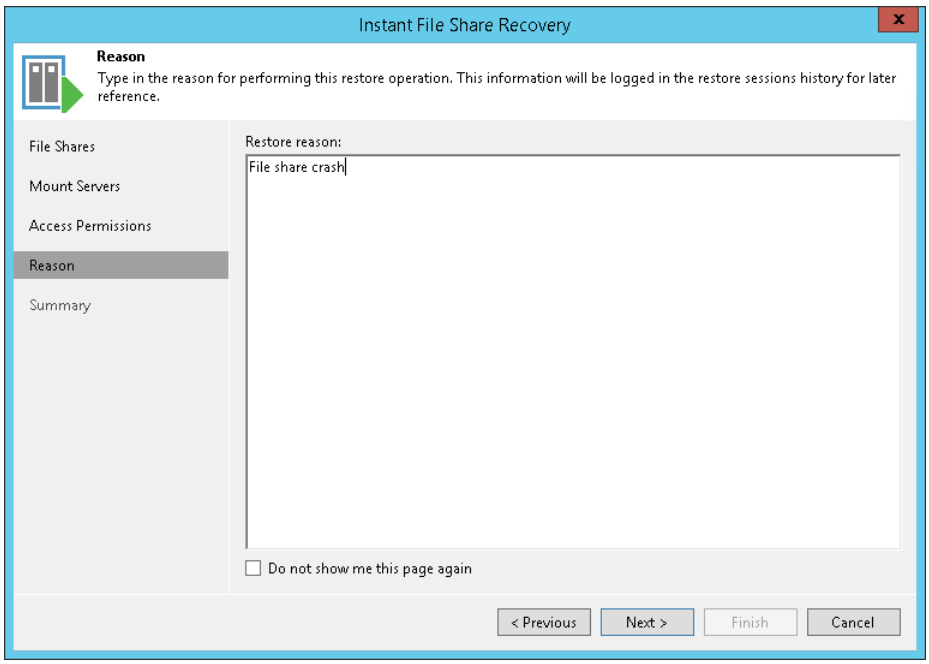
     Use **Add** and **Remove** buttons to configure accounts and groups to which you want to grant permissions for accessing the file share.



Step 5. Specify Reason for Recovery

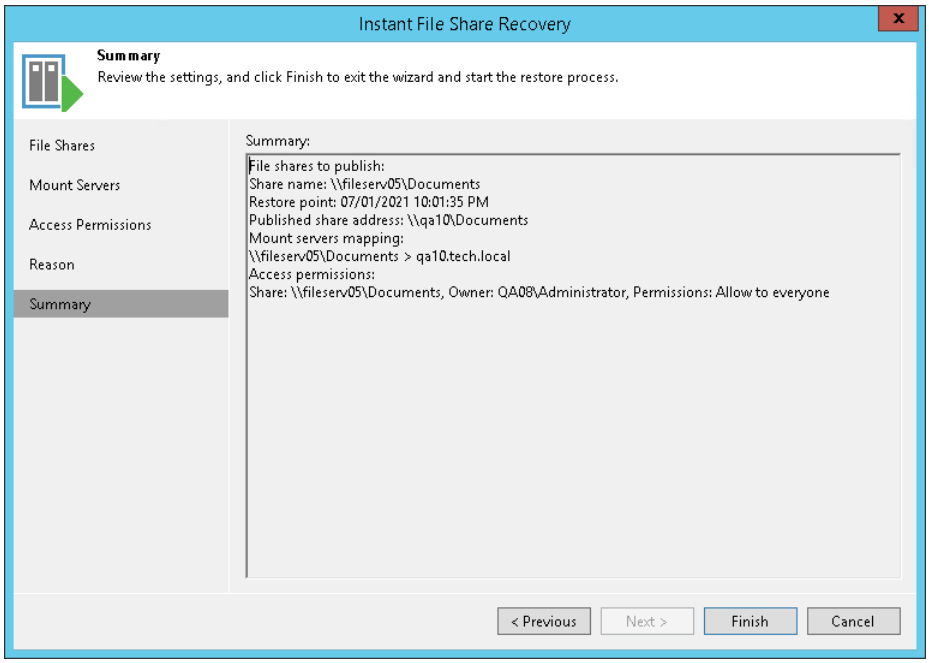
At the **Reason** step of the wizard, specify the reason for performing instant file share recovery. You can leave the field blank.

If you do not want to specify the reason for instant file share recovery in the future, select the **Do not show me this page again** check box.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review the instant file share recovery settings and click **Finish**. Veeam Backup & Replication will publish the file share to the specified mount servers.



Restoring Entire File Share

You can restore the entire file share from the backup to a specific restore point. That can be helpful, for example, if your file share device gets out of order and you need to restore the entire file share to the original or other location.

Before you restore an entire file share, check prerequisites.

Before You Begin

Before you restore the entire file share, consider the following:

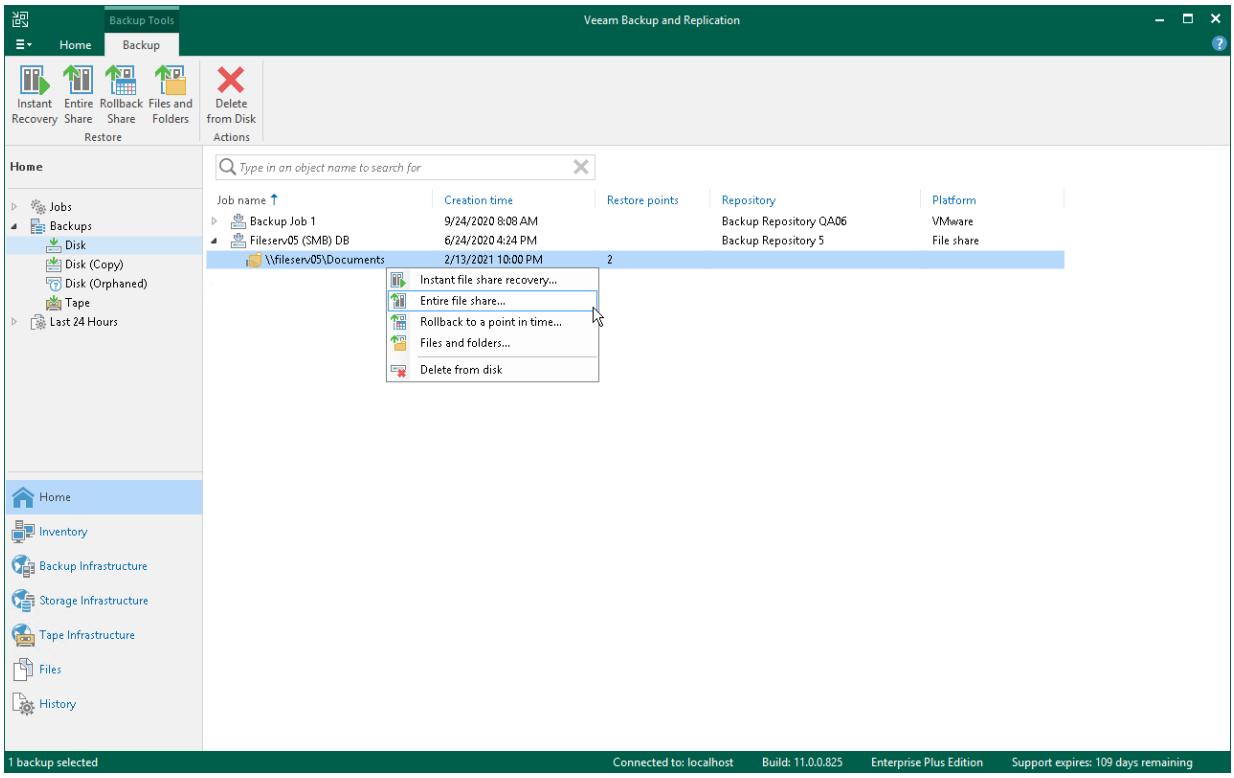
- You can restore files and folders of the file share from a backup that has at least one successfully created restore point.
- The file share on which you plan to save restored files and folders must be added to the backup infrastructure.

Step 1. Launch File Restore Wizard

To launch the **File Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **File Share** > **Restore entire share**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share that you want to restore and click **Backup** > **Entire share** on the ribbon.
  - Right-click the file share that you want to restore and select **Entire file share**.

  You can restore the file share to the state as of a specific restore point by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane.



Step 2. Select File Share to Restore

At the **File Shares** step of the wizard, select the file share that you want to restore:

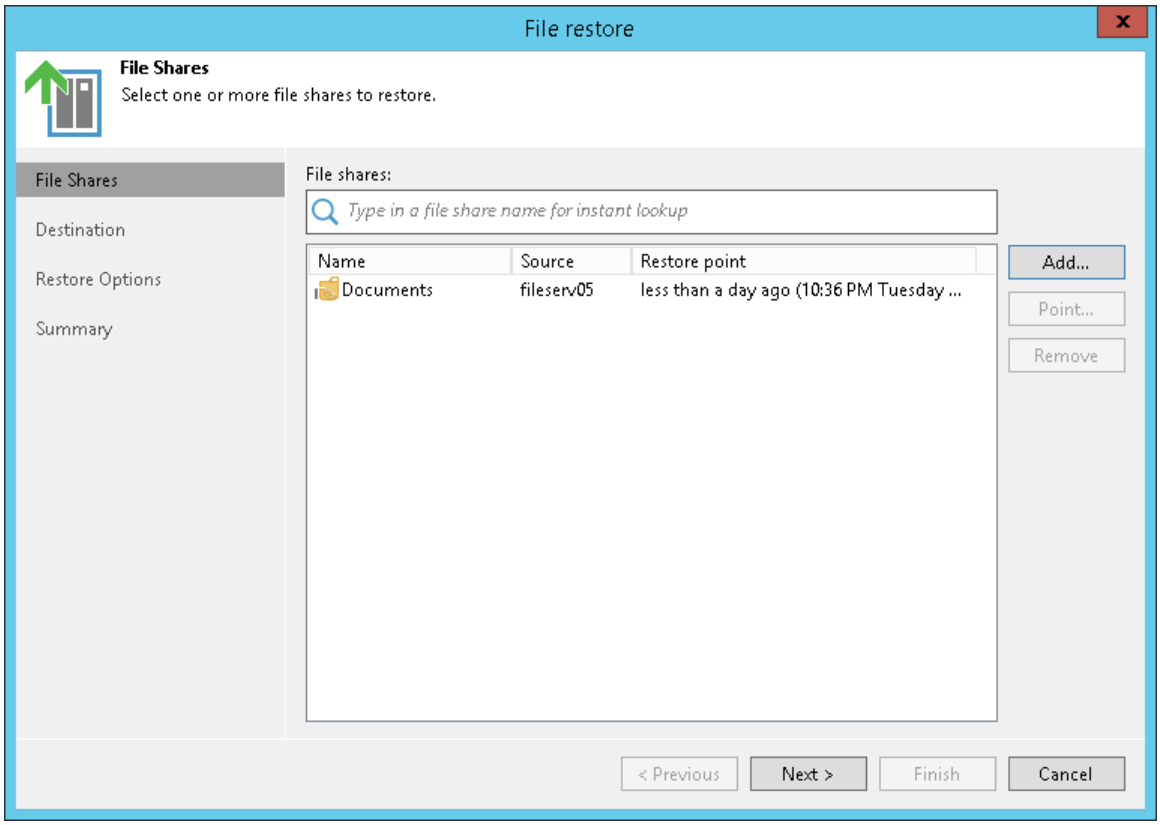
1. Click **Add**.
2. In the **Backups Browser** window, select the file share backup job and a file share in it that you want to restore. Click **OK**.
3. In the **File shares** table, select the file share to select a restore point to restore to. Click **Point**.
4. In the **Select Restore Point** window, select the restore point to which you want to restore the file share. To select the required restore point, do one of the following:
   - Use the **Restore point** slider.
   - Click the date link under the **Restore point** slider. In the calendar in the right pane of the **Restore points** window, select the date when the required restore point was created. The list of restore points on the left pane displays restore points created on the selected date. Select the point to which you want to restore the file share.

   In the **Files in backup** tree, you can see what folders and files are covered by the selected restore point and the date when each of them was modified.

   Click **OK**.

To quickly find a file share, you can use the search field at the top of the window. Enter a file share name or a part of it in the search field and press [Enter].

To exclude the file share from the restore process, select the file share in the table and click **Remove**.



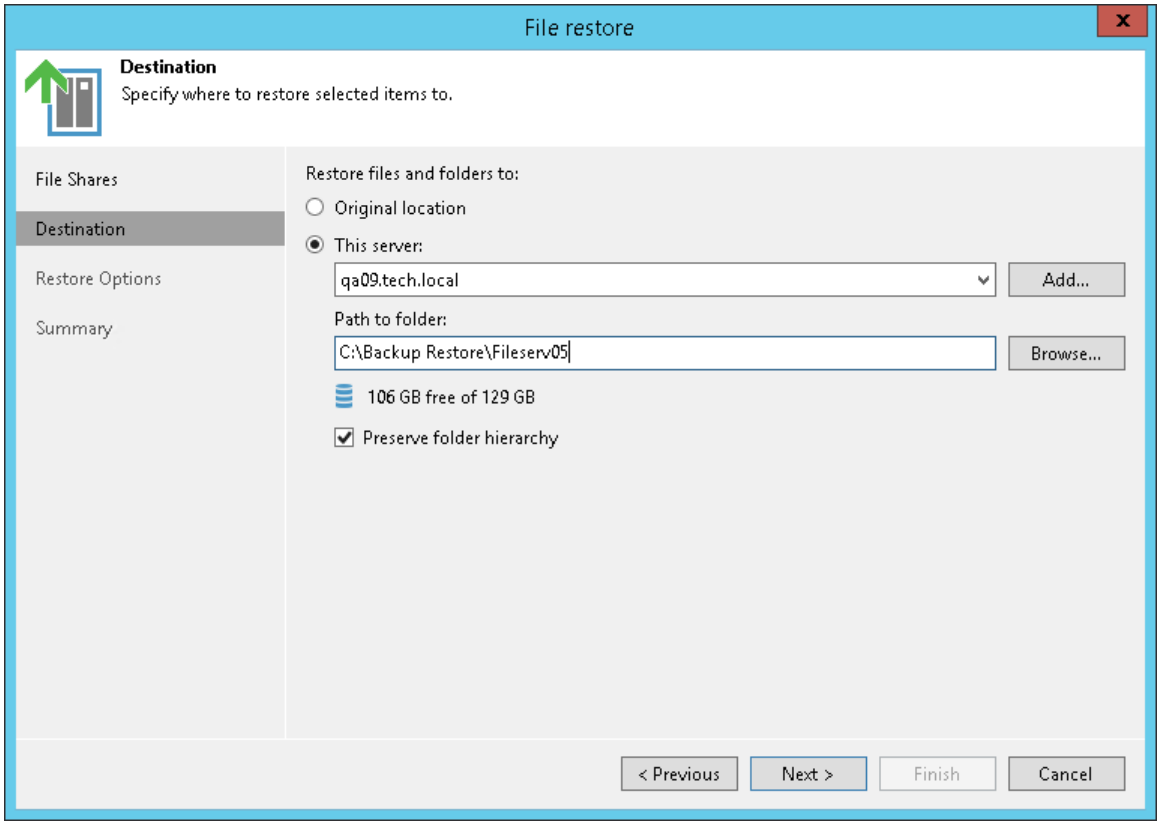
Step 3. Specify Destination for Data Restore

At the **Destination** step of the wizard, specify the location to which you want to restore the file share.

- Select **Original location** to restore data to the location where it resided originally. This type of restore is only possible if the original device is connected to Veeam Backup & Replication and powered on.
- Select **This server** to restore data to another location:
  a. In the **This server** field, select a file share to which files must be restored. You can select any file share added to the backup inventory. If the required file share is missing in the drop-down list, click **Add** and add a new file share to Veeam Backup & Replication. For more information on how to add a new file share, see File Share.
  b. In the **Path to folder** field, specify a path to the folder on the selected file share where files must be restored.

     To select a specific folder on the file share to restore files to, click **Browse**, in the **Select Folder** window, select the target location for the file share.

     If you want to restore the file share to a new folder, click **New Folder** at the bottom of the window, enter the folder name and click **OK** to confirm the new folder creation.
  c. Select **Preserve folder hierarchy** to keep the folder hierarchy of the original file share in the new location.



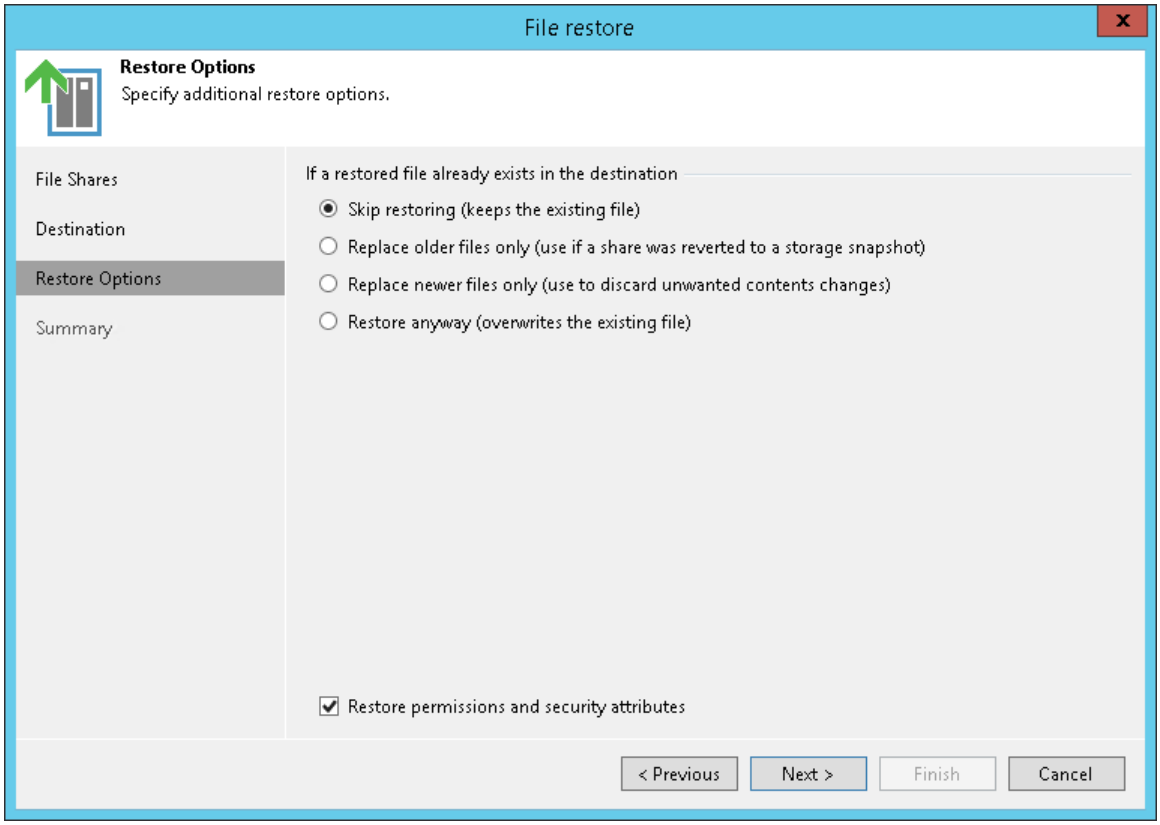
Step 4. Specify Restore Options

At the **Restore Options** step of the wizard, specify overwrite options in case the file with the same name already exists in the target folder:

- **Skip restoring (keeps the existing file)**. Select this option if you do not want to overwrite the existing file with the restored one.
- **Replace older files only (use if a share was reverted to a storage snapshot)**. Select this option if you want to overwrite the existing file only if it is older than the restored file.
- **Replace newer files only (use to discard unwanted contents changes)**. Select this option if you want to overwrite the existing file only if the restored file is older than the source share file.
- **Restore anyway (overwrites the existing file)**. Select this option if you want to overwrite the existing file with the restored file in all cases.

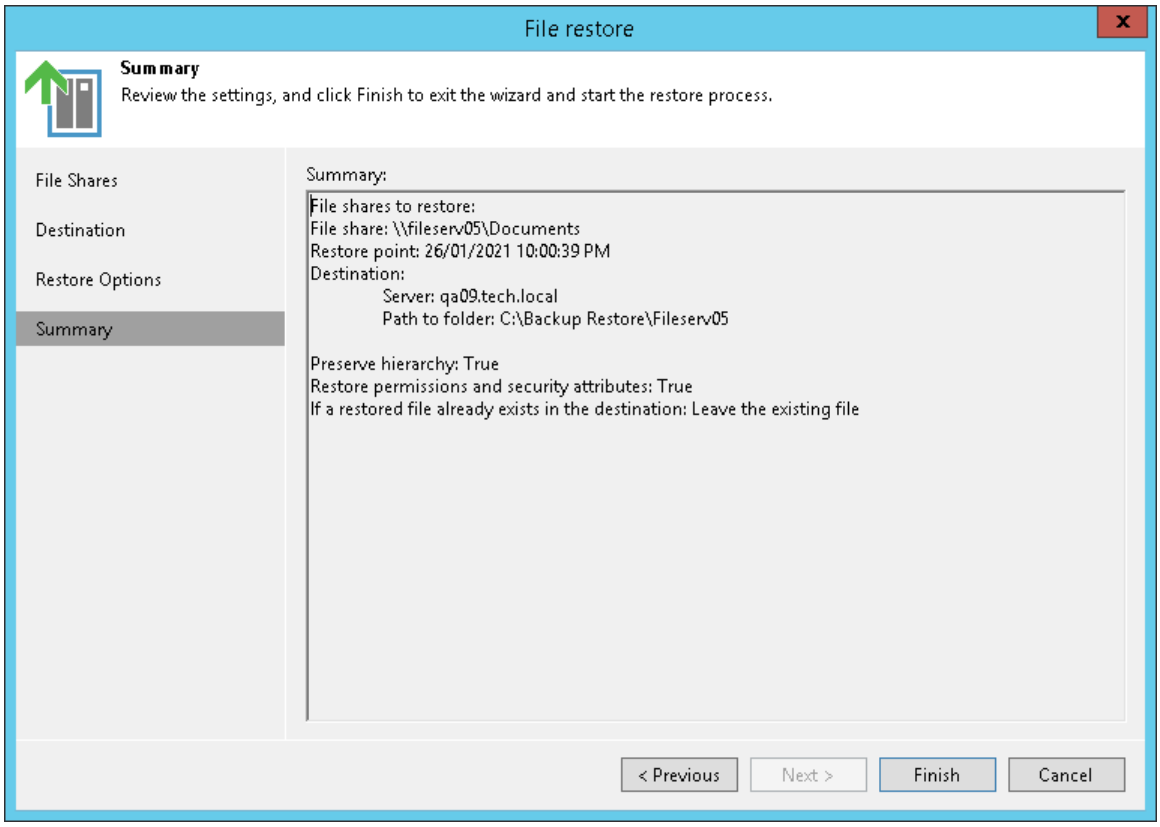
Select the **Restore permissions and security attributes** check box if you want the restored files to keep their original ownership and security permissions. If you do not select this check box, Veeam Backup & Replication will change security settings. The user account under which the Veeam Backup Service runs will be set as the owner of the restored objects. Access permissions will be inherited from the target folder to which the objects are restored.

> **NOTE**
>
> Mind that Veeam Backup & Replication does not collect ACL handling settings of the source file share root folder, so you cannot restore them. Before restoring an entire file share, you will have to specify required ACL handling settings for the root folder of the target file share.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review the file share restore settings and click **Finish**. Veeam Backup & Replication will restore the file share to the specified location.



Rolling Back to a Point in Time

You can roll back changes made to files and folders on the file share to a specific restore point. This option can be useful, for example, when the original file share was attacked by ransomware. In this case you can roll back all the files that were changed by the ransomware to the state before the attack.

Before you roll back the file share to a point in time, check prerequisites.

Before You Begin

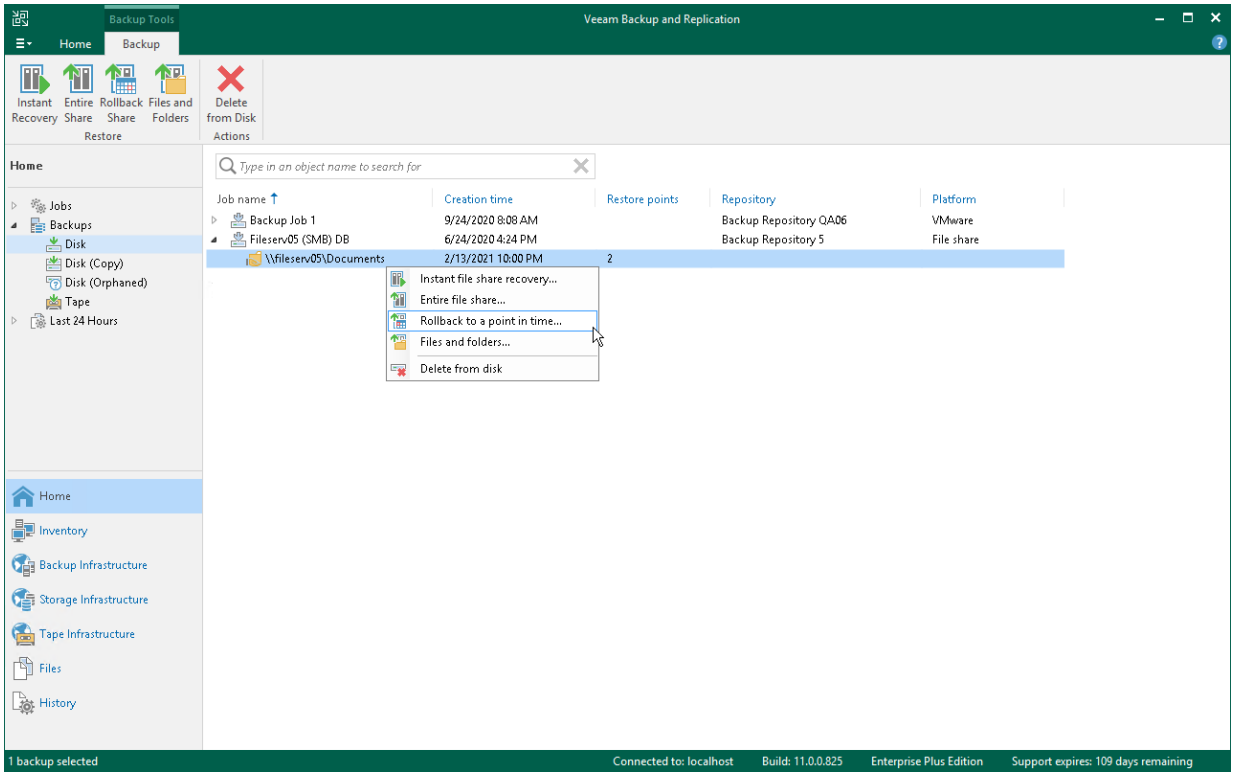
Before you restore files and folders to a point in time, mind that you can restore files and folders of the file share from a backup that has at least one successfully created restore point.

Step 1. Launch File Restore Wizard

To launch the **File Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **File Share** > **Rollback to a point in time**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share backup whose files you want to restore and click **File Share** > **Rollback to a point in time** on the ribbon.
  - Right-click the file share backup whose files you want to restore and select **Restore** > **Rollback to a point in time**.

  You can roll back the file share to a point in time by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane.



Step 2. Select Object to Restore

At the **File Shares** step of the wizard, select the file share, whose files you want to roll back:

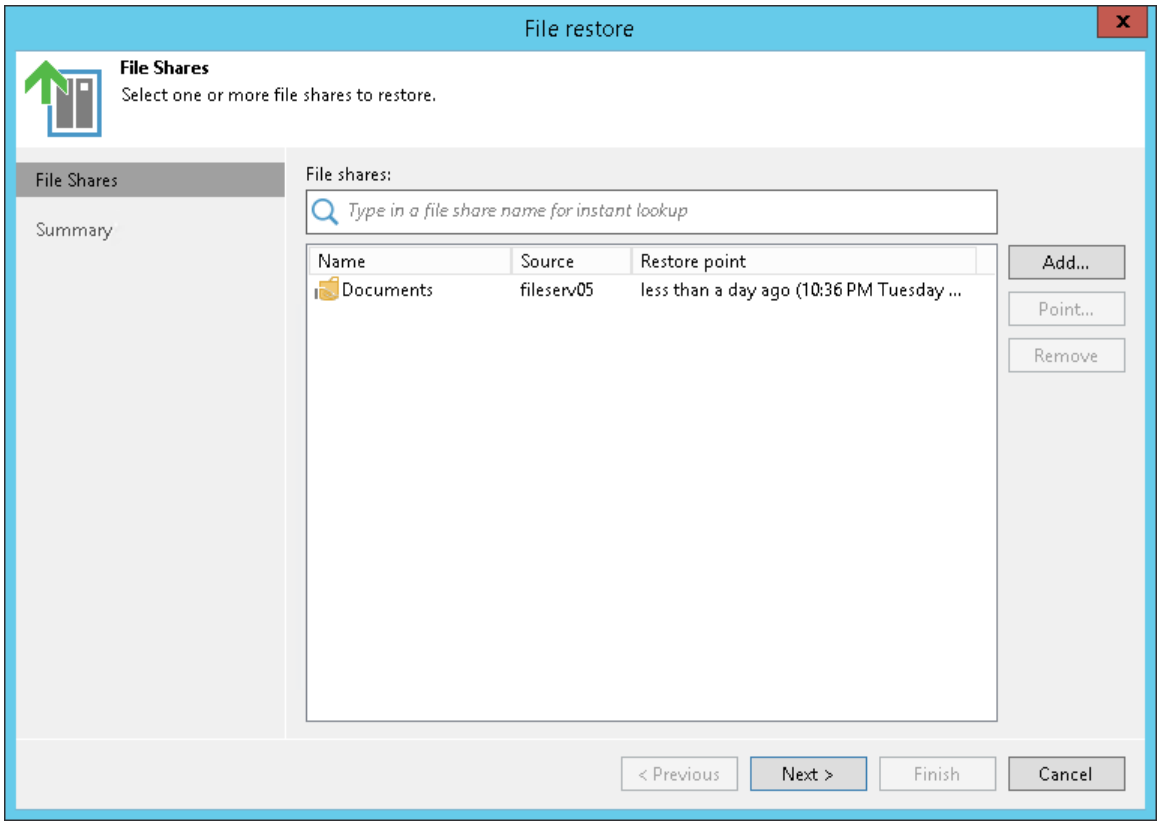
1. Click **Add**.
2. In the **Backups Browser** window, select the file share backup job and a file share in it, whose files you want to roll back. Click **OK**.
3. In the **File shares** table, select the file share to select a restore point to rollback to. Click **Point**.
4. In the **Select Restore Point** window, select the restore point to which you want to roll back the files. To select the required restore point, do one of the following:
   - Use the **Restore point** slider.
   - Click the date link under the **Restore point** slider. In the calendar in the right pane of the **Restore points** window, select the date when the required restore point was created. The list of restore points on the left pane displays restore points created on the selected date. Select the point to which you want to roll back the files to.

   In the **Files in backup** tree, you can see what folders and files are covered by the selected restore point and the date when each of them was modified.

   Click **OK**.

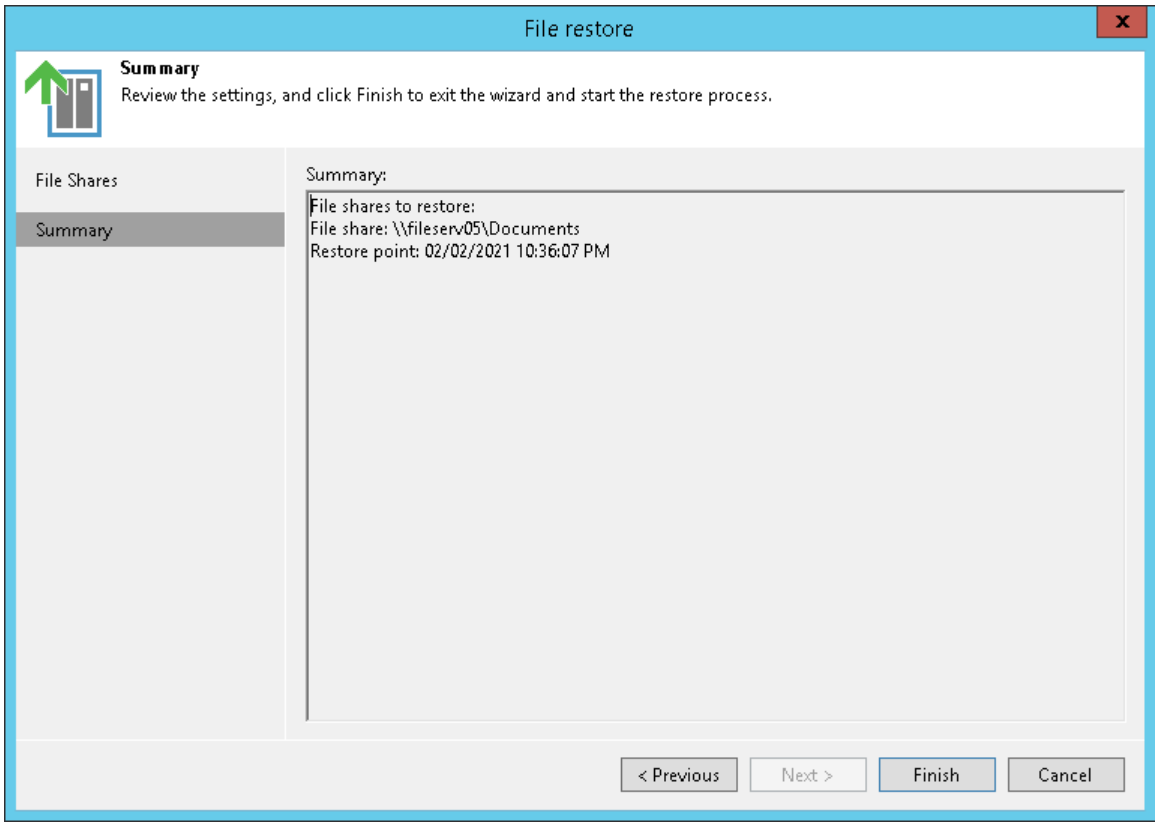
To quickly find a file share, you can use the search field at the top of the window. Enter a file share name or a part of it in the search field and press [Enter].

To exclude the file share from the rollback process, select the file share in the table and click **Remove**.



Step 3. Finish Working with Wizard

At the **Summary** step of the wizard, review the file share restore settings and click **Finish**. Veeam Backup & Replication will restore the files to the specified point in time.



Restoring Specific Files and Folders

You can restore specific files and folders to the original or a new location. This option can be useful, for example, if you need to get an older version of some files and folders from the backup.

When you restore specific files, you can extract file versions not only from the backup repository, but also from the archive repository. For more information, see Restoring Backup Files from Archive Repository.

> **NOTE**
>
> Mind that from the archive repository you can restore files only. Restore of whole folders from the long-term repository is not supported.

Besides, you can restore multiple versions of the same file.

Before you restore specific files and folders, check prerequisites.

Before You Begin

Before you restore files and folders from the backup, mind the following:

- You can restore files and folders from a backup that has at least one created restore point, even if it is incomplete.
- The file share on which you plan to save restored files and folders must be added to the backup infrastructure.

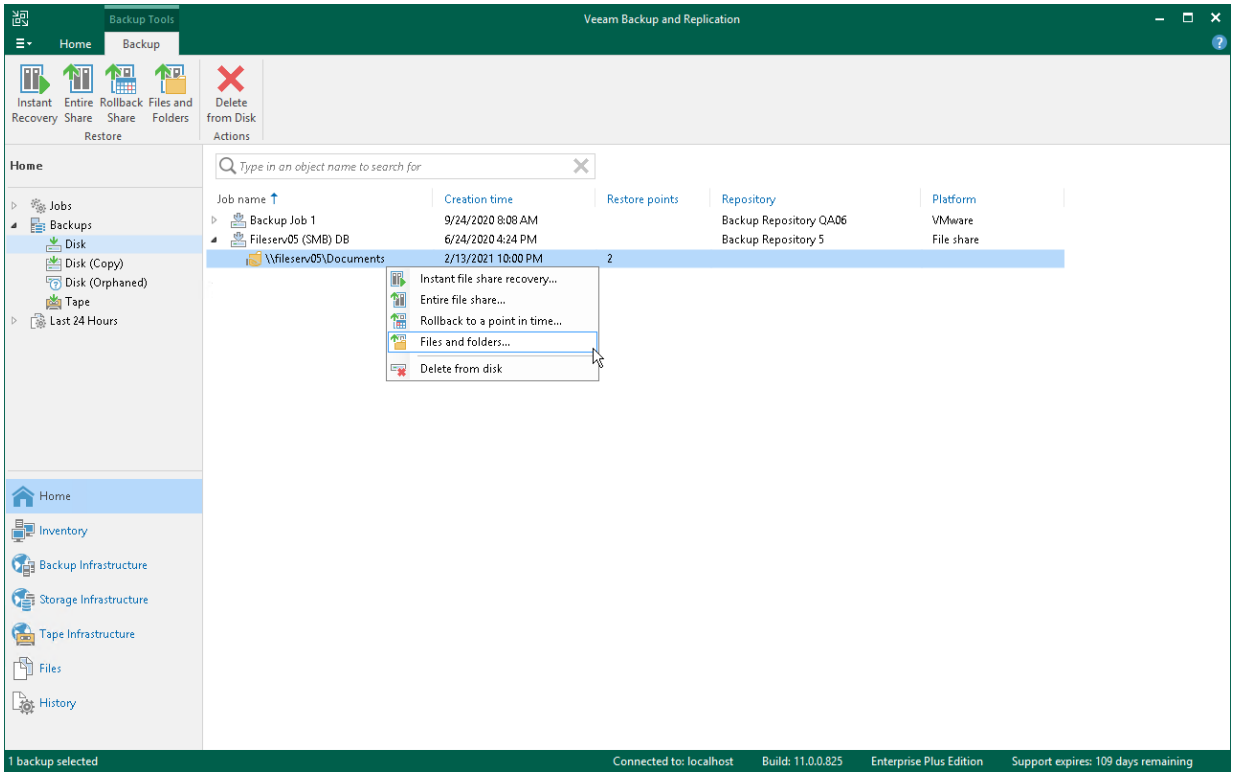
Step 1. Launch File Restore Wizard

To launch the **File Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **File Share** > **Restore individual files and folders**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share backup whose files you want to restore and click **File Share** > **Restore files and folders** on the ribbon.
  - Right-click the file share backup whose files you want to restore and select **Restore** > **Files and folders**.

  In this case, you will pass directly to the Backup Browser.

  You can restore files and folders from a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane.



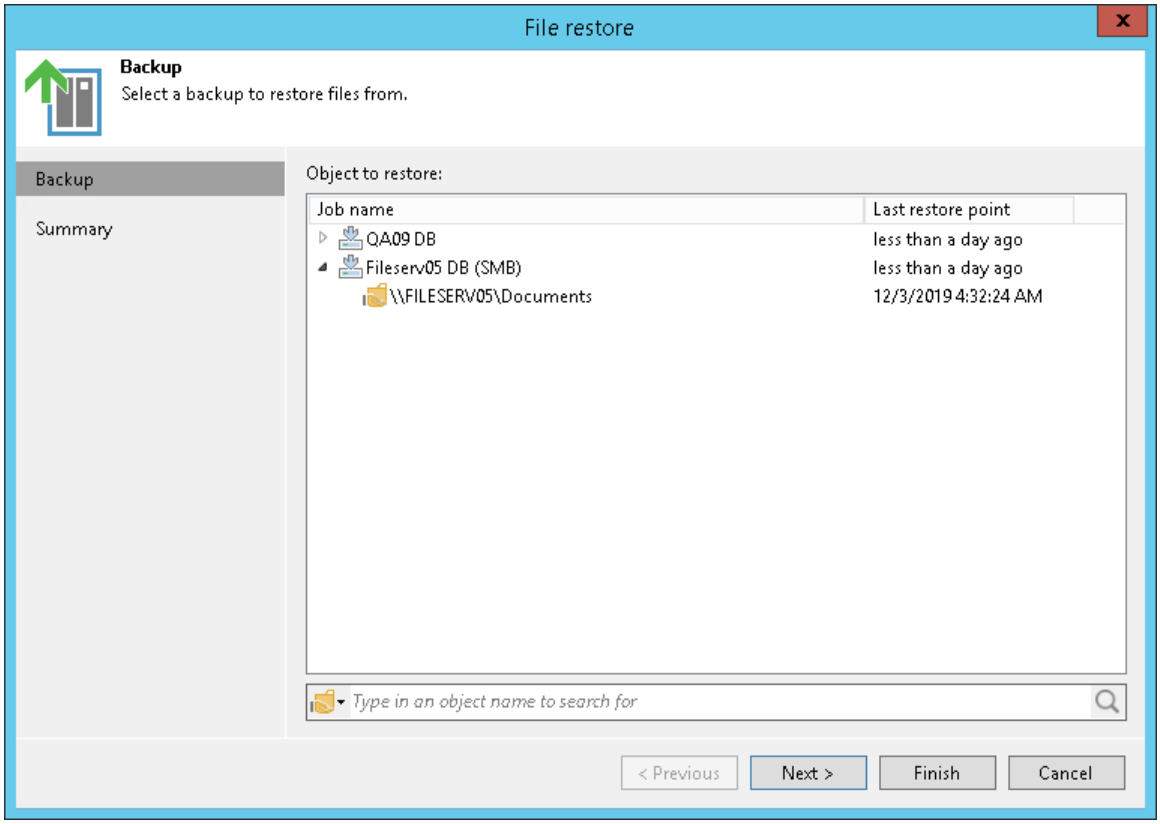
Step 2. Select Object to Restore

At the **Backup** step of the wizard, select the file share backup whose files you want to restore:

1. In the **Object to restore** list, expand the necessary backup job.
2. Select the file share.

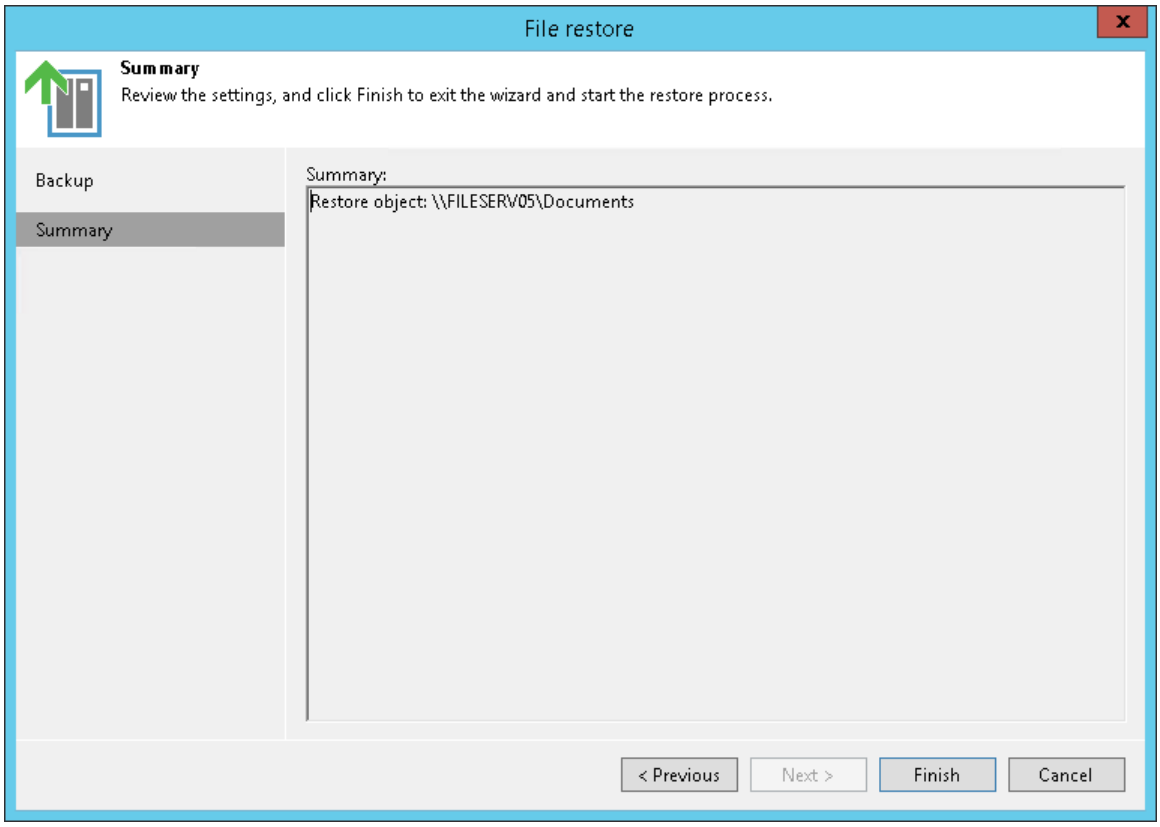
To quickly find a file share, you can use the search field at the bottom of the window.

1. Enter a file share name or a part of it in the search field.
2. Click the **Start search** button on the right or press [Enter].



Step 3. Verify Restore Object Settings

At the **Summary** step of the wizard, review selected restore object and click **Finish** to switch to the Backup Browser and select files and folders to restore.



Step 4. Select Files and Folders to Restore

In the **Backup Browser**, select files and folders to restore. Backup browser has three representations of restore points. You can select one of the following options in the **Restore Point** group on the ribbon to display file versions:

- **Latest** – the **Backup Browser** shows the latest versions of files and folders on the file share.
- **All Time** – the **Backup Browser** shows all files and folders ever backed up by the backup job. This option retrieves file versions stored both in the backup and archive repositories. This representation additionally shows how many file versions of each file are stored in the backup and the date when the latest file version was created.

  After you select this option:

  - If you restore a whole folder, you will be prompted to the Select Restore Point to Use step to select a restore point to restore files and folders from.

    > **NOTE**
    >
    > Mind that during file-level restore you cannot restore whole folders from the archive restore points. If the files are already moved to the archive repository, you have to restore them one by one.

  - If you restore a single file, you will be prompted to the Select File Version to Restore step to select a file version to restore the file from.
- **Selected** – the **Backup Browser** shows versions of files and folders backed up as of the certain restore point. Select the restore point in the list on the right of the **Selected** option on the ribbon.

You can use the search field at the top of the working area to search for specific files and folders.

> **NOTE**
>
> In order to keep the operation of the Backup Server stable, the number of retrieved search result records is limited to 1000. Therefore, if you work with backup folders that store large volumes of data, it is recommended to narrow the search criteria to fit into the limitation.

You can restore files and folders to their original location or a new location.

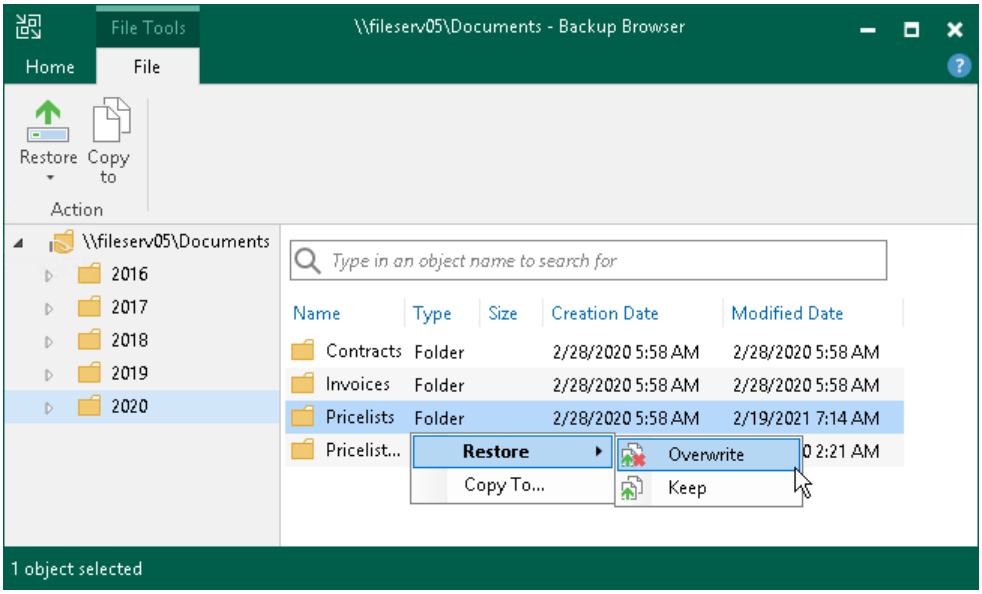
Restoring Files to Original Location

To restore files and folders to the original location, in the Veeam Backup browser right-click a file or folder and select one of the following commands:

- To overwrite the original file on the file share with the file restored from the backup, select **Restore** > **Overwrite**.

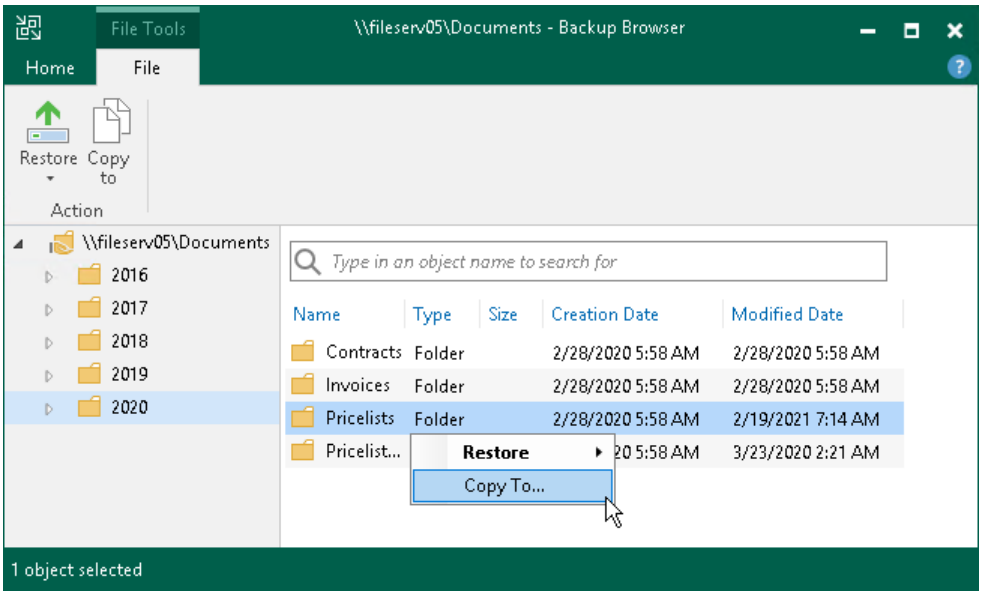
- To save the file restored from the backup next to the original file, select **Restore** > **Keep**.

  Veeam Backup & Replication will add the *_RESTORED_YYYYMMDD_HHMMSS* suffix to the original file name and store the restored file in the same folder where the original file resides.



Saving Files to New Location

To save files and folders to a new location, right-click the necessary file or folder and select **Copy To**. If you restore files and folders to a new location, you will be prompted to the Specify Destination for File Restore step to specify a new destination.



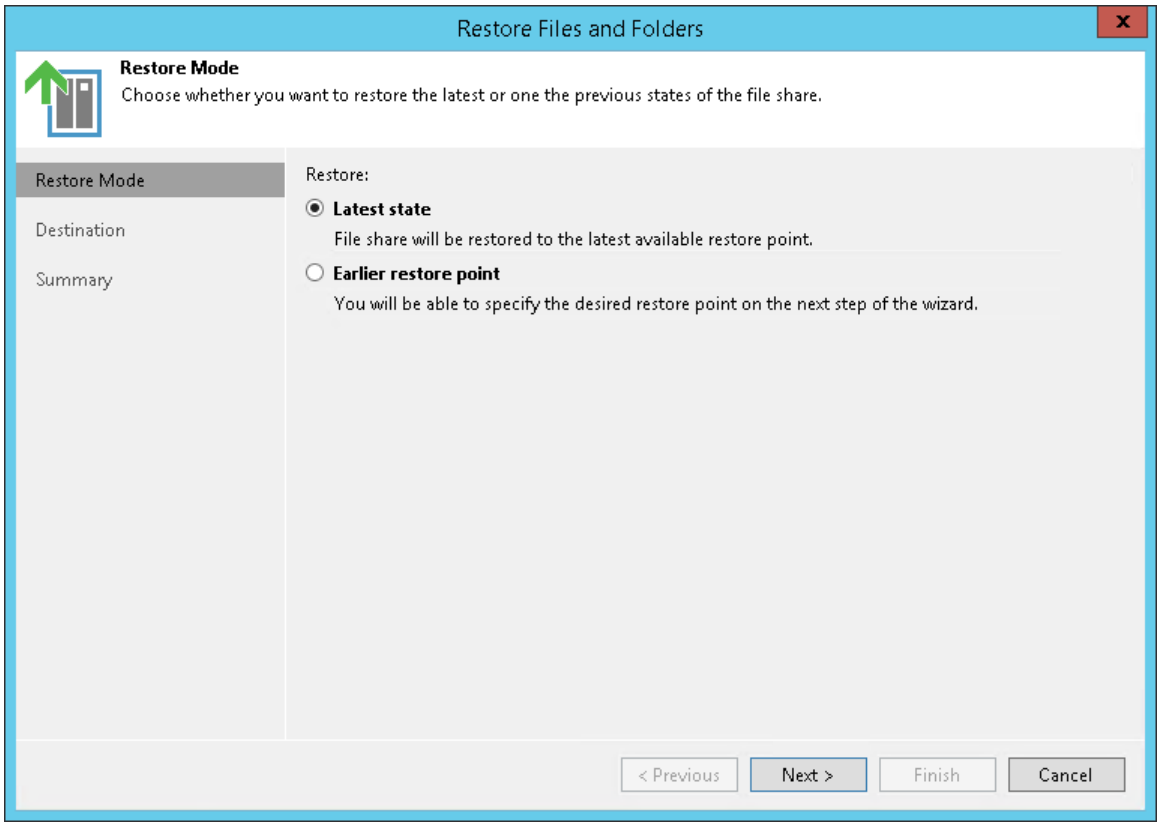
Step 5. Select Restore Point

The **Restore Mode** step is required if you use the **All Time** option at the Select Files and Folders to Restore step and the selected folders have more than one restore point.

Specify the restore point to which you want to restore files and folders:

- To restore the folder to the latest available restore point, select **Latest state**.
- To select a specific restore point, select **Earlier restore point**, click **Next** and specify the restore point. To select the required restore point, do one of the following:
  - Use the **Restore point** slider.
  - Click the date link under the **Restore point** slider. In the calendar in the left pane of the **Restore points** window, select the date when the required restore point was created. The list of restore points on the left pane displays restore points created on the selected date. Select the point to which you want to restore the file share.

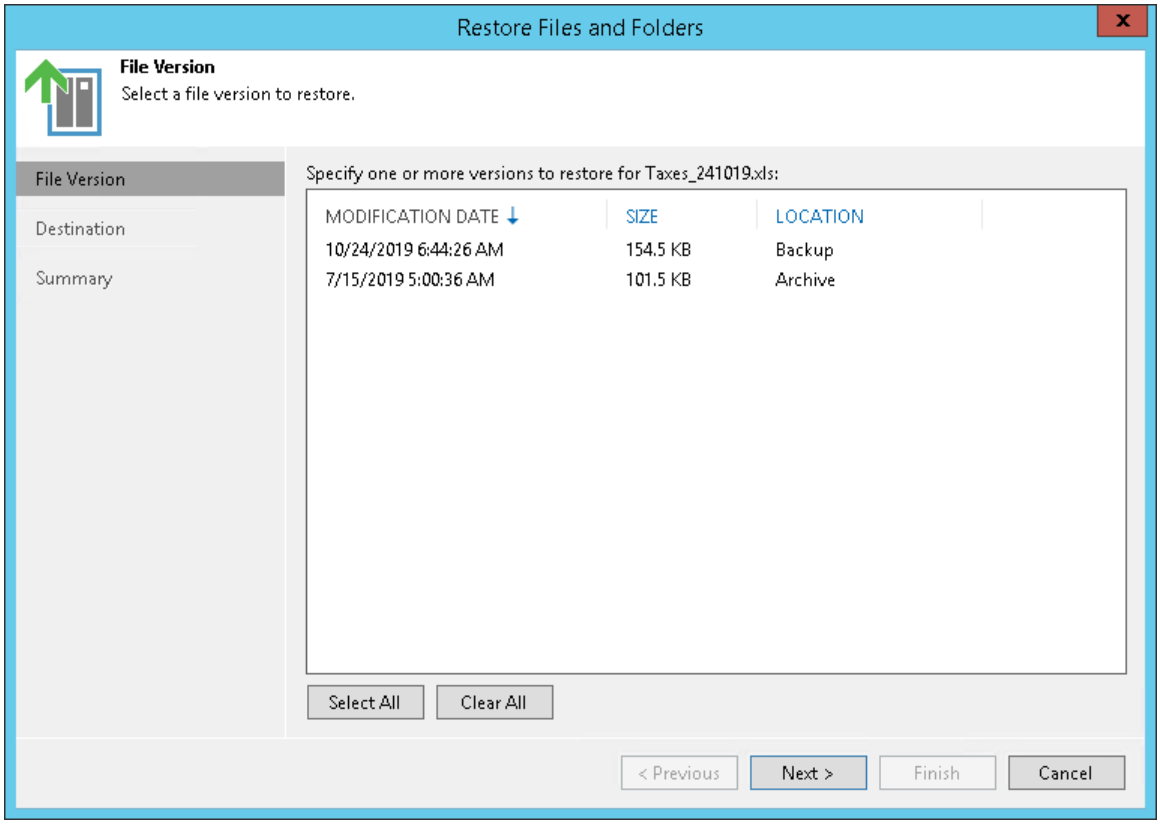
  In the **Files in backup** tree, you can see what folders and files are covered by the selected restore point and the date when files and folders were modified.



Step 6. Select File Version to Restore

The **File Version** step is required if you use the **All Time** option at the Select Files and Folders to Restore step and the selected files have more than one file version.

Select one or more versions to restore. You can restore files both from the backup repository and archive repository. To select several file versions, hold [Ctrl] and select multiple records in the table. Restore of multiple file versions can be helpful, for example, when you need to search for a specific version of the file, but you do not know for sure which one contains required changes.



Step 7. Specify Destination for File Restore

The **Destination** step is required if you choose the **Copy To** option at the Select Files and Folders to Restore step. Specify the destination where the restored files must be stored:

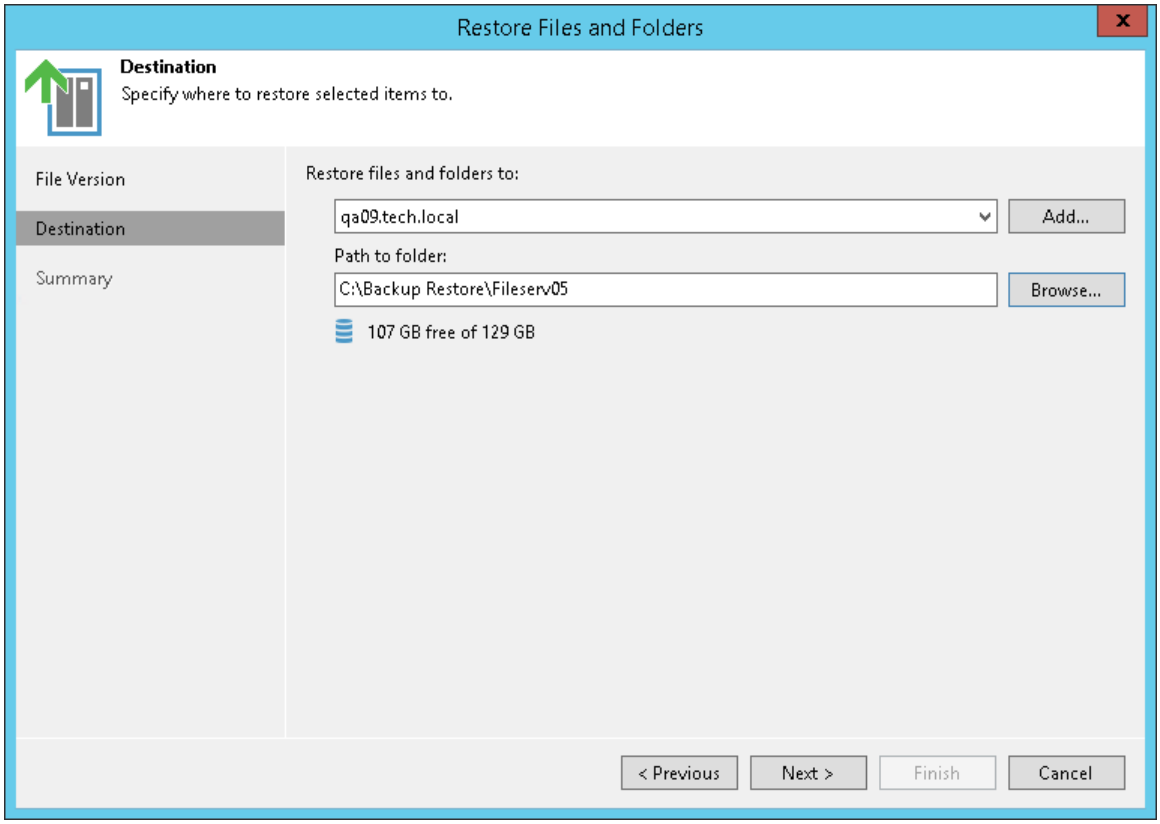
1. In the **Restore files and folders to** field, select a file share to which the files must be restored. All file shares added to the inventory of Veeam Backup & Replication are available. If the required file share is missing in the drop-down list, click **Add** and add a new file share to Veeam Backup & Replication.

   For more information on how to add a new file share, see Adding File Share.

2. In the **Path to folder** field, specify a path to the folder on the selected file share where files must be restored.

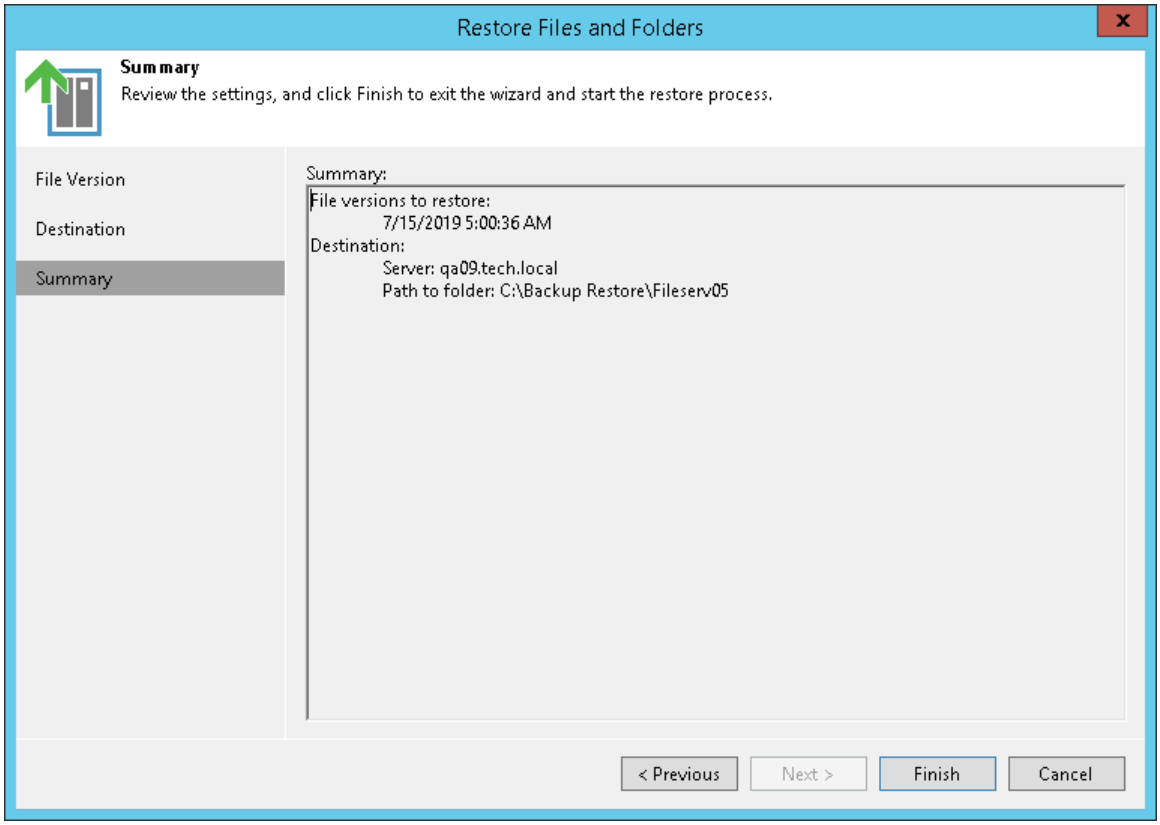
   To create a dedicated folder for restored files, click **Browse**. In the **Select Folder** window, select the target location for the file share.

   If you want to restore the file share to a new folder, click **New Folder** at the bottom of the window. Confirm the new folder creation.



Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review the file restore settings and click **Finish**.



Restoring Backup Files from Archive Repository

You can restore any file from the archive repository to the state of any file version stored in the archive. Depending on the circumstances, such a restore can require different actions.

> **NOTE**
>
> Mind that from the archive repository you can restore files only. Restore of whole folders from the long-term repository is not supported.

Regular Restore from Archive

To perform a regular restore from the archive repository when you have all required backup data stored both in the backup repository and archive repository, follow the instructions given in Restoring Specific Files and Folders. Mind that to restore data from the archive repository, you must select the **All Time** option for selecting files and folders to restore.

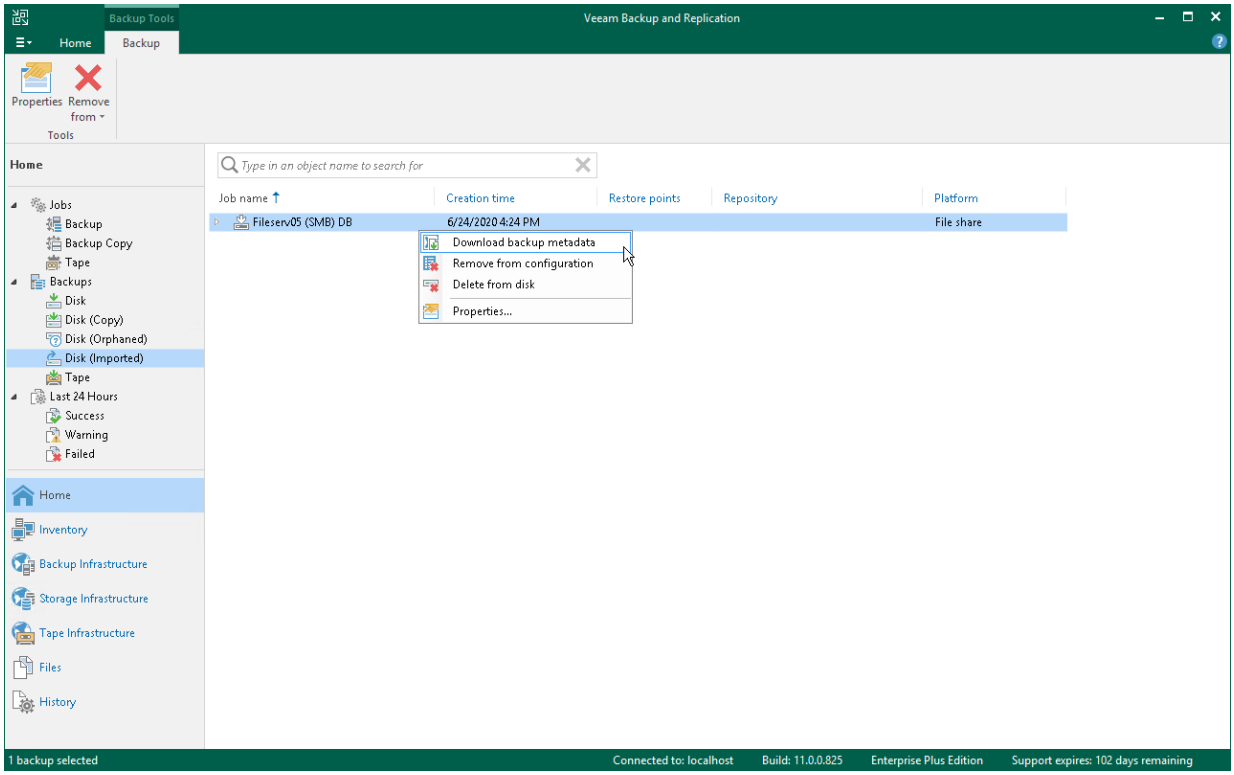
Emergency Restore from Archive

There can be different situations where backups are not available in the backup repository: for example, backup files are broken or deleted from the backup repository, backup configuration files are removed from the configuration, archive repository is added into backup infrastructure of another backup server, or backup archive is moved from one archive repository to another one. In these cases, you can restore specific files and folders from the backup archive.

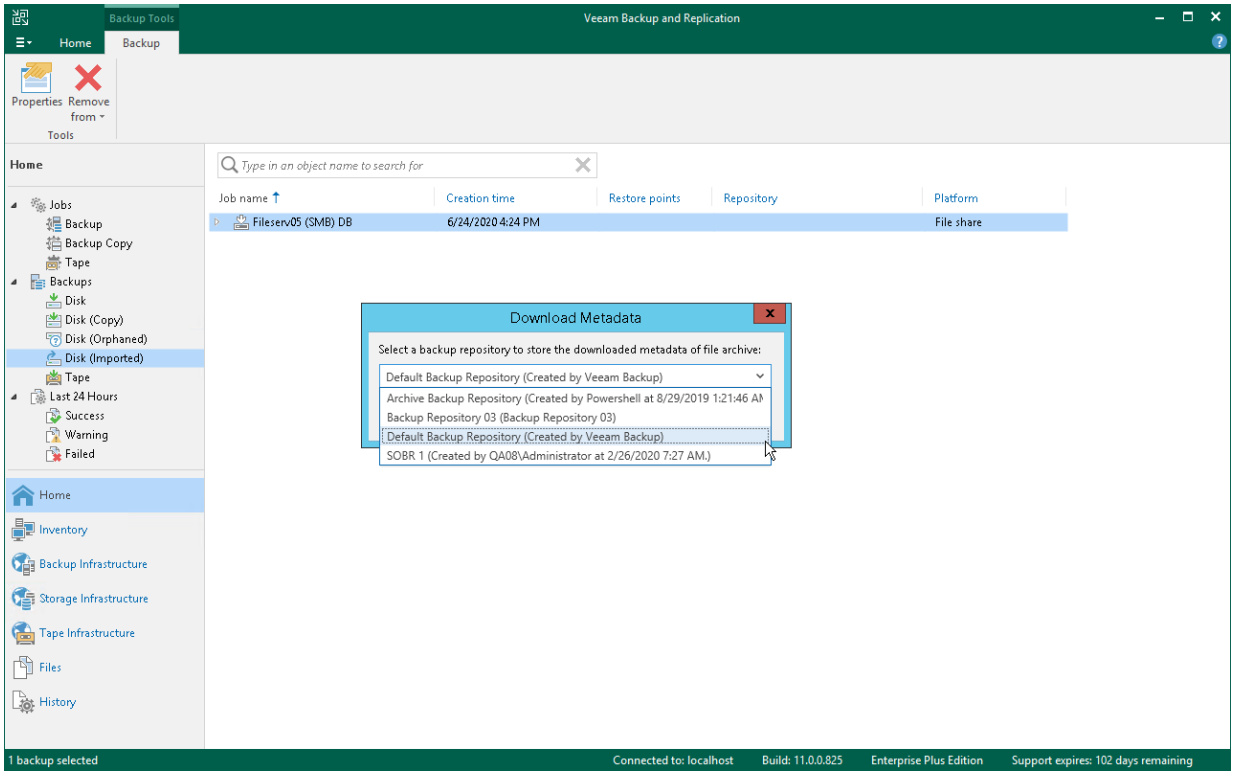
To restore files from the archive repository, for example, from an object storage:

1. If necessary, add the storage that keeps the required archive to the backup infrastructure as described in Adding Backup Repositories.
2. Rescan the added archive repository as described in Rescanning Backup Repositories.
3. Download metadata for the archive backup:
    a. Locate the required file backup archive under **Backups - Disk (Imported)** node in the **Home** view.

b. Right-click the file backup and select **Download backup metadata**.



c. From the drop-down list, select a backup repository to store the downloaded metadata of files archive and click **OK**.



4. Restore files from the archive backup as described in Restoring Specific Files and Folders.

Integration with Veeam Backup for AWS

Veeam Backup & Replication allows you to create and manage data protection and restore tasks for Amazon Elastic Compute Cloud (Amazon EC2) environments. For this, Veeam Backup & Replication uses a AWS Plug-in for Veeam Backup & Replication component. This component extends Veeam Backup & Replication functionality and provides access to Veeam Backup for AWS from the Veeam Backup & Replication console. For more information, see the Integration with Veeam Backup for AWS Guide.

Integration with Veeam Backup for Microsoft Azure

Veeam Backup & Replication allows you to create and manage data protection and restore tasks for Microsoft Azure environments. For this, Veeam Backup & Replication uses a Microsoft Azure Plug-in for Veeam Backup & Replication component. This component extends Veeam Backup & Replication functionality and provides access to Veeam Backup for Microsoft Azure from the Veeam Backup & Replication console. For more information, see the Integration with Veeam Backup for Microsoft Azure Guide.

Integration with Google Cloud Platform

> **NOTE**
> This functionality is available in Veeam Backup & Replication starting from version 11a (build 11.0.1.1261).

Veeam Backup & Replication allows you to create and manage data protection and restore tasks for Google Cloud Platform (GCP) environments. For this, Veeam Backup & Replication uses a Google Cloud Platform Plug-in for Veeam Backup & Replication component. This component extends Veeam Backup & Replication functionality and provides access to Veeam Backup for Google Cloud Platform from the Veeam Backup & Replication console. For more information, see the Integration with Veeam Backup for Google Cloud Platform Guide.

Integration with Veeam Backup for Nutanix AHV

Veeam Backup & Replication allows you to manage data protection and restore tasks for Nutanix AHV environments. For this, Veeam Backup & Replication uses an additional component: Veeam Backup for Nutanix AHV.

Veeam Backup for Nutanix AHV extends Veeam Backup & Replication functionality and allows you to back up and restore Nutanix AHV VMs. For more information, see the Veeam Backup for Nutanix AHV User Guide.

Integration with Veeam Backup for Red Hat Virtualization

Veeam Backup & Replication allows you to manage data protection and restore tasks for Red Hat Virtualization (RHV) hosts. For this, Veeam Backup & Replication uses an additional component: Veeam Backup for RHV.

Veeam Backup for RHV extends Veeam Backup & Replication functionality and allows you to back up and restore RHV VMs. For more information, see the Veeam Backup for RHV User Guide.

Integration with Kasten K10

> **NOTE**
> This functionality is available in Veeam Backup & Replication starting from version cumulative patch P20211211 for v11a.

Veeam Backup & Replication allows you to manage data protection and restore tasks for backups exported with K10 policies. For this, Veeam Backup & Replication uses the Veeam Backup Repositories for Kasten K10 solution.

Veeam Backup Repositories for Kasten K10 extends Veeam Backup & Replication functionality and allows you to export backups created by K10 policies to backup repositories. For more information, see the Integration with Veeam Backup Repositories for Kasten K10 User Guide.

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).

For more information about Veeam Agents, see the Veeam Agent Management Guide.

Veeam Cloud Connect

If you want to store your data in the cloud, you can connect to the service provider and write VM backups to cloud repositories or create VM replicas on cloud hosts.

For more information about Veeam Cloud Connect, see the Veeam Cloud Connect Guide.

Advanced VMware vSphere Features

Veeam Backup & Replication lets you leverage the following VMware vSphere features and functionality during data protection and disaster recovery operations:

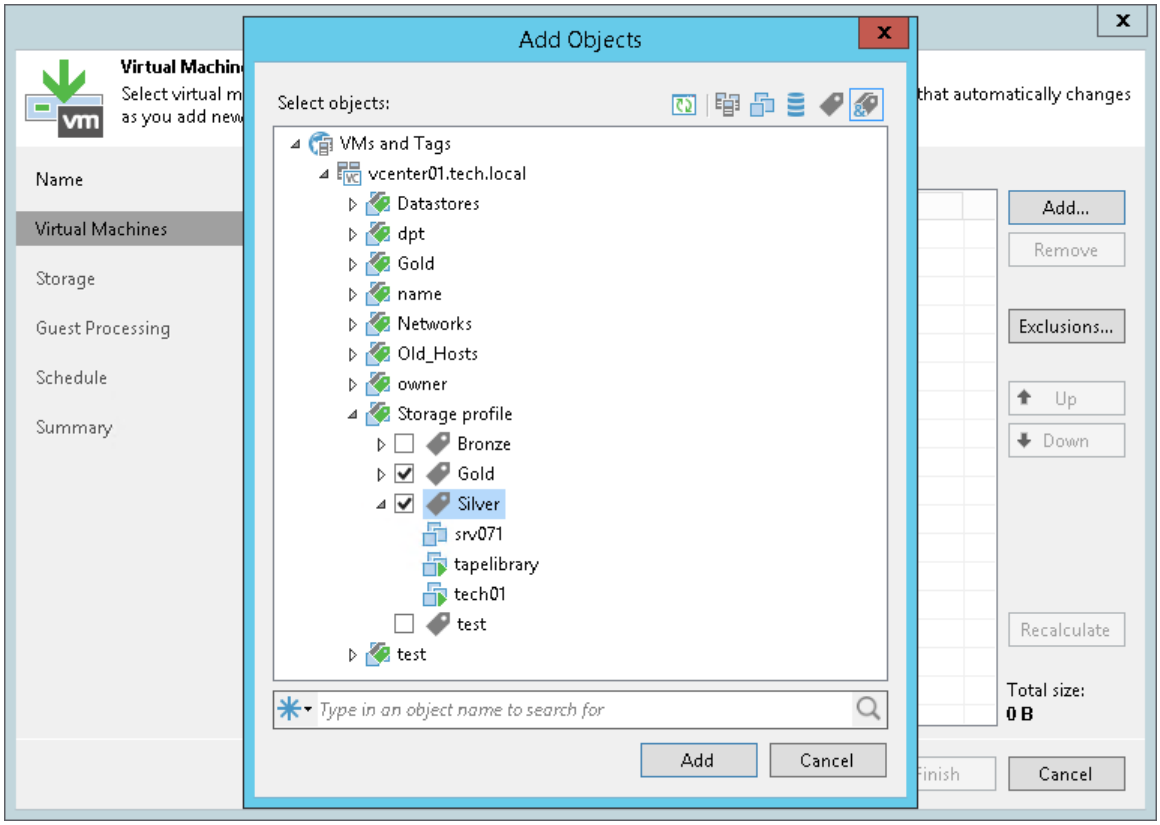
- VM tags
- Encrypted VMs
- Storage policies

VM Tags

If you use vCenter Server tags to categorize objects in the virtual infrastructure, you can filter objects that you add to data protection and disaster recovery jobs and tasks by these tags. Use of tags facilitates object management. You can quickly configure jobs and tasks for VMs that belong to a specific category, for example, a certain department or SLA level.

To add objects by tags, switch to the **VMs and Tags** or **Tag combination** view in the **Add Objects** window. Veeam Backup & Replication will display objects categorized by tags.

In the **VMs and Tags** view, you can select only one tag and those objects that have this tag will be processed by a job. In the **Tags combination** view, you can select multiple tags and those object that have all the selected tags will be processed.



Requirements and Recommendations for VM Tags

When you work with tags in Veeam Backup & Replication, mind the following requirements and recommendations:

- The *VirtualCenter.FQDN* parameter in the **Advanced Settings** of vCenter Server must contain the real fully qualified domain name of the vCenter Server host.
- A certificate installed on vCenter Server must contain the real fully qualified domain name of the vCenter Server host.
- The fully qualified domain name of the vCenter Server host must be accessible and resolved to its IP (and vice versa) from machines on which Veeam Backup & Replication services are installed (at least the Veeam Backup Service and Veeam Broker Service).

- A user account used for specific data protection and disaster recovery operations must have sufficient permissions on the vCenter Server. For more information, see Full VM Restore, Replica Failback and Cumulative Permissions sections in the Required Permissions Reference.
- If VM tags are not displayed in the Veeam Backup & Replication console for some reason, try restarting VMware vSphere services that are responsible for the tags functionality. In VMware vSphere earlier than 6.5, you must restart the vCenter Inventory Service. Starting from VMware vSphere 6.5, vCenter Inventory Services functionality is replaced by the vCenter Content Library and other services that are part of vCenter Server 6.5.

  When you upgrade to VMware vSphere 6.5, data from vCenter Inventory Service is migrated to the new database support services in vCenter Server 6.5. The vCenter Inventory Service may remain in the list of services, however, it is no longer used.
- If you are using the vSphere Fault Tolerance feature, assign VM tags in one of the following ways in order to provide proper VM backup processing by Veeam Backup & Replication:
  - Assign the required VM tag to the primary VM. In case of the failover to the secondary VM, the VM tag is copied to the secondary VM and this VM is used as a source for continuing Veeam Backup & Replication jobs.
  - Assign the required VM tag to the VM container that comprises the primary VM in the virtual infrastructure inventory. The primary VM inherits this tag. In case of the failover to the secondary VM, the VM tag is copied to the secondary VM and this VM is used as a source for continuing Veeam Backup & Replication jobs.

  Thus, there is no need to assign the VM tag to the secondary VM. Moreover, if you manually assign the same tag both to the primary VM and the secondary VM and use this tag to add objects to a VM backup job, this will result in the backup failure.

Limitations for VM Tags

The VM tags support functionality has the following limitation:

Veeam Backup & Replication ignores the cardinality setting for VM tag categories. For example, you create a tag category *Priority* and set cardinality to **One tag per object**. In the tag category, you create two tags: *Normal* and *High*. You assign the Normal tag to a VM folder and the *High* tag to a VM in this folder. If you now configure a job that will process objects with the *Normal* tag, the VM with the *High* tag will also be added to this job (since Veeam Backup & Replication regards that VMs and templates in the VM container inherit the tag assigned to the container).

To overcome this situation, you can add to the list of exclusions the tag assigned to objects that you do not want to process.

Encrypted VMs

Veeam Backup & Replication provides support for VMware vSphere encrypted VMs.

- Backup of encrypted VMs
- Restore of encrypted VMs
- Replication of encrypted VMs
- Failback of encrypted VM replicas

Backup of Encrypted VMs

To back up VMware encrypted VMs, the backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption policy, and assign it to VMs in advance.
- The backup proxy must be working in the **Virtual appliance** or **Network** transport mode:
  - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.
  - The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

Restore of Encrypted VMs

Veeam Backup & Replication supports the following restore options:

- Restore encrypted VM as encrypted or unencrypted.
- Restore unencrypted VM as encrypted.

The backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption storage policy, and assign it to VMs in advance.
- The backup proxy must be working in the **Virtual appliance** or **Network** transport mode:
  - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.
  - The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

If you restore a VM as an encrypted one to the specified location, ensure that the target datastore is under the **VM Encryption Policy** node.

If a VM has several disks, you can optionally restore some disks as encrypted and some disks as unencrypted. Keep in mind, that even if one disk is restored as encrypted, the VM configuration file must also be placed on a datastore under the **VM Encryption Policy** node.

Replication of Encrypted VMs

To replicate VMware encrypted VMs, the backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption policy, and assign it to VMs in advance. Ensure that you use a common KMS or the KMS clusters at both sites use common encryption keys.

  > **NOTE**
  >
  > If you do not set up KMS, the replication job will not fail, but replicated VMs will not be encrypted in this case.

- Source and target backup proxies must be working in the **Virtual appliance** or **Network** transport mode:
  - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.
  - The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

To replicate a VM as an encrypted one, place disks and the configuration file of the VM replica on datastores compatible with the VM encryption policy:

- At the **Destination** step of the wizard, click **Choose** near the **Datastore** field.
- In the **Select Datastore** window, select a datastore under the **VM Encryption Policy** node.

> **NOTE**
>
> Multi-OS file-level restore for encrypted VM replicas is not supported.

Failback of Encrypted VM Replicas

To fail back VMware encrypted VMs replicas, the backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption policy and assign it to VMs in advance. Ensure that you use a common KMS or the KMS clusters at both sites use common encryption keys.
- Source and target backup proxies must be working in the **Virtual appliance** or **Network** transport mode:
  - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.
  - The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

If you fail back an encrypted VM replica to the specified location, ensure that the target datastore is under the **VM Encryption Policy** node.

Storage Profiles

During backup, Veeam Backup & Replication preserves information about the storage policy associated with the VM, and stores this information to the backup file or replica metadata. When you restore the VM to its original location, Veeam Backup & Replication also restores information about the VM storage policy. The restored VM gets automatically associated with the original storage policy.

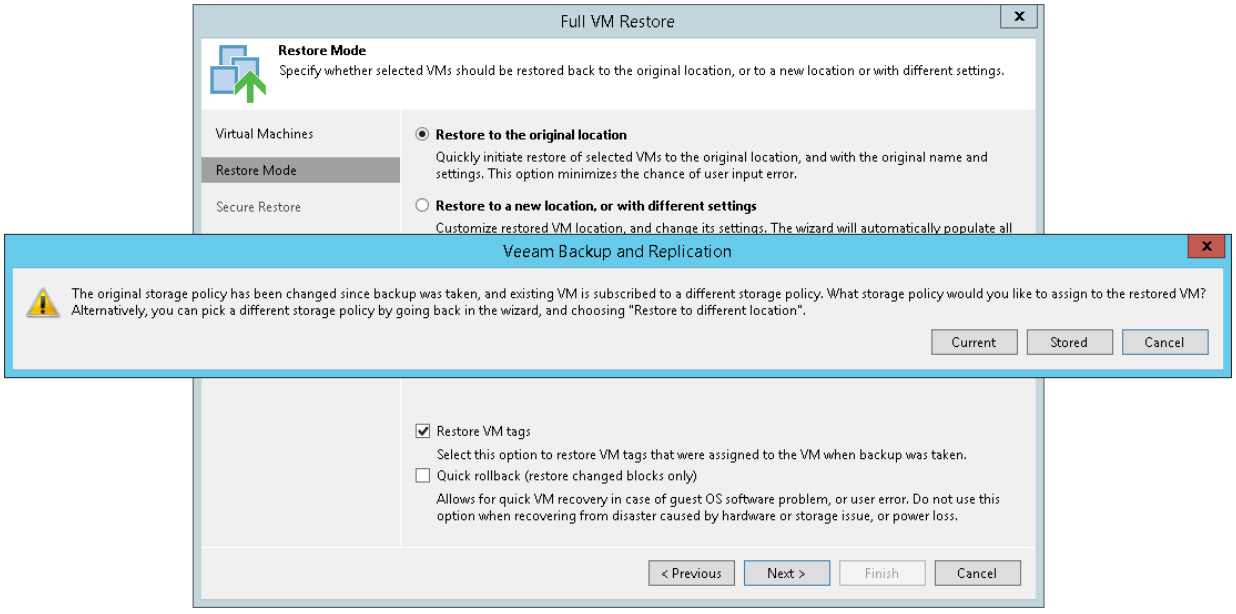
Veeam Backup & Replication restores the storage policy when you perform the following operations:

- Entire VM restore
- VM failback

Veeam Backup & Replication restores the storage policy only if you restore the VM to the original location. If you restore the VM to a new location, Veeam Backup & Replication does not preserve the storage policy for the VM.

In some cases, the original storage policy may be changed or missing by the time when you restore the VM. For example, the storage policy may be deleted, the original VM in the production environment may be associated with another storage policy and so on. In such situation, Veeam Backup & Replication displays a warning and lets you choose one of the following scenarios:

- Associate the VM with the current storage policy — the restored VM will be associated with the profile with which the original VM in the production environment is currently associated.
- Associate the VM with the default storage policy — the restored VM will be associated with the profile that is set as default for the target datastore.
- Associate the VM with the profile stored in the backup file — the restored VM will be associated with the profile that was assigned to the original VM at the moment of backup, and whose information is stored in the backup file.



Veeam Backup & Replication Utilities

You can use the following Veeam Backup & Replication utilities to perform advanced administration tasks in your backup infrastructure:

- Extract.exe Utility
- Veeam.Backup.DBConfig.exe Utility
- Veeam Backup Validator

Extract Utility

Veeam Backup & Replication comes with an extract utility that can be used to recover machines from backup files. The extract utility does not require any interaction with Veeam Backup & Replication and can be used as an independent tool on Linux and Microsoft Windows machines.

The extract utility can be helpful, for example, if it is written to the tape next to machine backup files. In this case, you get a possibility to recover machines from backups at any moment of time even if backups are removed from Veeam Backup & Replication or Veeam Backup & Replication is not installed.

> **IMPORTANT**
>
> 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.

The extract utility can be used in two interfaces:

- Graphic user interface (GUI)
- Command-line interface working in the interactive and regular mode

The extract utility is located in the installation folder of Veeam Backup & Replication, by default: `%PROGRAMFILES%\Veeam\Backup and Replication\Backup`. The folder contains three files for the extract utility:

- `Veeam.Backup.Extractor.exe` — utility working in GUI (can be used on Microsoft Windows machines only)
- `extract.exe` — utility working in the command-line interface, a version for Microsoft Windows
- `extract` — utility working in the command-line interface, a version for Linux

Using Extract Utility in GUI

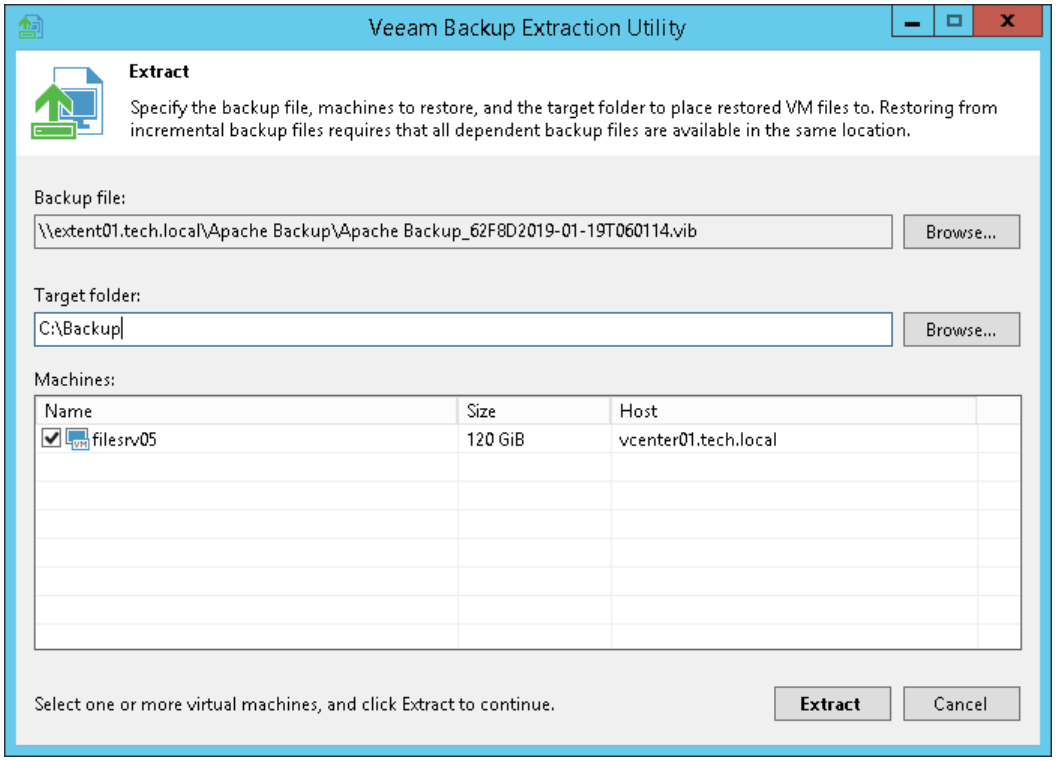
To restore machine data in the extract utility GUI:

1. Run the `Veeam.Backup.Extractor.exe` file from the installation folder of Veeam Backup & Replication.
2. In the **Backup file** field, specify a path to the backup file from which you want to restore machine data.
3. If the backup file is encrypted, the extract utility will require you to provide a password to unlock the backup file. Enter the password that was used for backup file encryption.
4. In the **Target folder** field, specify a path to the destination folder where machine data must be restored.
5. From the **Machines** list, select machines whose data you want to restore.
6. Click **Extract**. Machine data will be restored to the specified folder.

IMPORTANT

If you restore machine data in the extract utility GUI, consider the following:

- The extract utility can be started on Microsoft Windows machines only.
- If you plan to start the extract utility on the machine other than the backup server, make sure that you copy the `Veeam.Backup.Extractor.exe` file together with the `extract.exe` file from the product installation folder and store these files to the same folder on the destination machine. In the opposite case, the extract utility will fail to start.



Using Extract Utility in Interactive Mode

To start the extract utility in the interactive mode, run the `extract.exe` file from the product installation folder(in case of a Linux machine, run the `extract` file).

You will have to sequentially enter the following arguments:

1. A path to the backup file from which the machine must be restored. After you enter the path, the extract utility will display a list of all machines included in the backup and their description.
2. A name of the machine that you want to restore. If there is more than one machine with the specified name in the backup, you will be asked to specify the host on which the backed-up machine resides. If you want to restore all machines from the backup, press [Enter] on the keyboard.
3. If the backup was encrypted, password that was used to encrypt the backup file.
4. An output directory to which machines must be restored. If you want to restore machines to the current directory, press [Enter] on the keyboard.
5. The operation confirmation. Press [Y] on the keyboard to restore a machine to the directory you specified. If you want to abort the operation, press [Enter] on the keyboard.

Using Extract Utility from Command Line

To run the extract utility from the command line, do one of the following:

- In the command line, change the current directory to the directory where the extract utility is located.
- Add the directory where the extract utility is located to the PATH variable.

Then you can perform the following actions:

- Run the extract utility in the interactive mode
- Display help information for the utility usage
- Display the list of all VMs in the backup file
- Getting encryption status of a backup file
- Restore all or selected VMs from the backup

Running Extract Utility in Interactive Mode

This command runs the extract utility in the interactive mode.

Syntax

```
extract.exe [-password backupkey] [pathtobackup]
```

Parameters

Parameter	Description	Required/Optional
password	Password for the encrypted backup file.	Required for encrypted backup files
pathtobackup	Path to the backup file from which machines must be restored.	Optional

Displaying Help Information for Utility Usage

This command prints all variants of the extract utility usage along with required and optional parameters.

Syntax

```
extract.exe -help
```

Displaying List of Machines in Backup

This command displays the list of all machines in the backup file from which you want to perform restore.

Syntax

```
extract.exe -dir [-vm vmname] [-host hostname] [-password backupkey] pathtoback
up
```

Parameters

Parameter	Description	Required/Optional
vm	Name of the machine that you want to restore. Use this parameter to filter machines in the backup.	Optional
host	Name of the host on which the initial machine resides. Specify this parameter to filter machines that have the same name but reside on different hosts. Note: This parameter must be specified if the `vm` parameter is used.	Optional
password	Password for the encrypted backup file.	Required for encrypted backup files
pathtobackup	Path to the backup file from which the machine must be restored.	Required

Getting Encryption Status of Backup File

This command gets the encryption status of the backup file: encrypted or not encrypted.

Syntax

```
extract.exe –getEncryptionStatus pathtobackup
```

Parameters

Parameter	Description	Required/Optional
pathtobackup	Path to the backup file from which the machine must be restored.	Required

Restoring VMs from Backup

This command restores data for all machines or for the selected machine from the backup file.

Syntax

```
extract.exe -restore [-vm vmname] [-host hostname] [-password backupkey] pathto
backup [outputdir] [-log]
```

Parameters

Parameter	Description	Required/Optional
vm	Name of the machine that you want to restore. Use this parameter to filter machines in the backup. If you want to restore all machines from the backup file, do not specify this parameter.	Optional
host	Name of the host on which the initial machine resides. Specify this parameter to filter machines that have the same name but reside on different hosts. Note: This parameter must be specified if the `vm` parameter is used.	Optional
pathtobackup	Path to the backup file from which the machine must be restored.	Required
password	Password for the encrypted backup file.	Required for encrypted backup files
outputdir	Path to the directory to which machine data must be restored. If this parameter is not specified, the machine will be restored to the current directory.	Optional

Parameter	Description	Required/Optional
log	SwitchParameter. Enables log creation. The log file will be created in the current directory.	Optional

Veeam Configuration Database Connection Utility

Veeam Backup & Replication comes with the `Veeam.Backup.DBConfig.exe` utility that allows you to manage connection settings for Veeam Backup & Replication and/or Veeam Backup Enterprise Manager configuration database. Using this utility, you can:

- Connect to a different database on the same or another Microsoft SQL Server instance. If you specify a database that does not exist yet, it will be created on the selected server.
- Change authentication method for database connection. Possible methods are Microsoft Windows authentication and Microsoft SQL Server authentication.

> **NOTE**
>
> The `Veeam.Backup.DBConfig.exe` utility supports only connection to configuration databases of the current version.

Using Veeam.Backup.DBConfig.exe Utility

You can launch the `Veeam.Backup.DBConfig.exe` utility from the **Start** menu by clicking **Configuration Database Connection Settings**.

Alternatively, you can use the `Veeam.Backup.DBConfig.exe` file located in the installation folder. By default, the path to the folder is the following: `%PROGRAMFILES%\Common Files\Veeam\Backup and Replication\DBConfig`

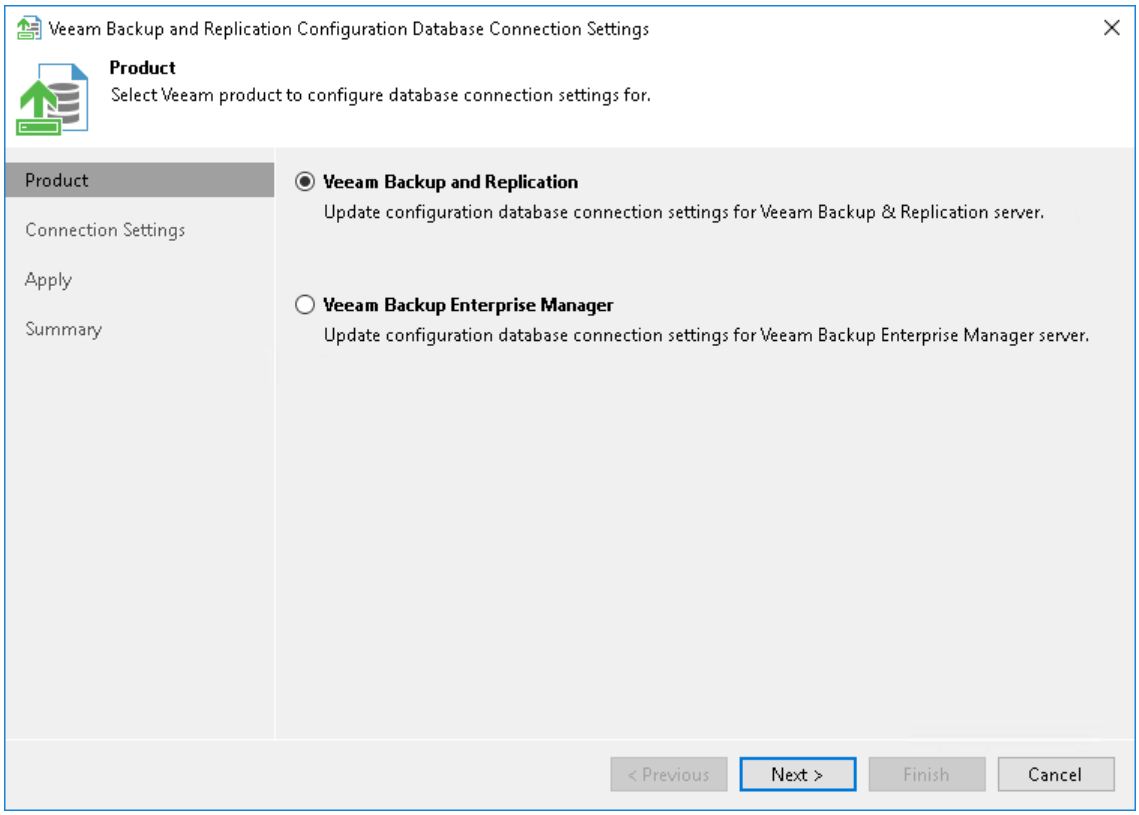
To run the utility, you must have administrative rights on the local machine, as long as the utility makes changes to the registry. If prompted at the launch, choose **Run as administrator**.

To manage connection settings for Veeam Backup & Replication or Veeam Backup Enterprise Manager configuration database, use the launched **Veeam Backup & Replication Configuration Database Connection Settings** wizard.

Step 1. Select Product

At the **Product** step of the wizard, select the database whose settings you want to configure.

The utility detects what server is installed on the local machine (backup server, Veeam Backup Enterprise Manager server or both) and displays available products for your choice. If Veeam Backup Enterprise Manager is not installed on the local machine, you will only have an opportunity to change Veeam Backup & Replication database settings (and vice versa). In this case, the **Product** step of the wizard will be skipped.



Step 2. Specify Connection Settings

At the **Connection Settings** step of the wizard, provide the connection settings for the selected database.

1. Specify the Microsoft SQL Server\instance and database name to which you want the Veeam Backup & Replication installation to connect. Both local and remote Microsoft SQL Server instances are supported. Microsoft SQL Server instances available on the network are shown in the **Server name** list. If necessary, click **Refresh** to get the latest information.

   If a database with the specified name does not exist on the selected Microsoft SQL Server instance, it will be created anew.

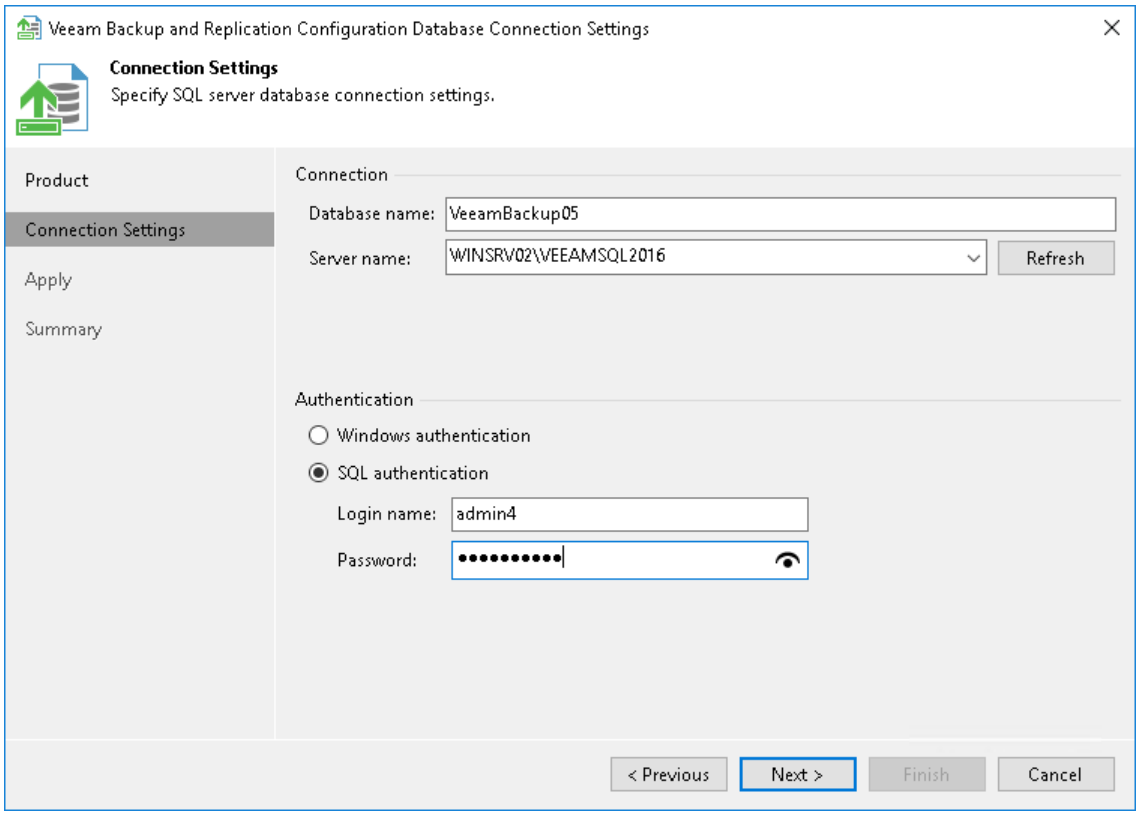
   > **NOTE**
   >
   > If you change the database to which Veeam Backup & Replication must be connected, make sure that the database to which Veeam Backup & Replication is currently connected is available. If not, you must stop the Veeam Backup Service on the machine where Veeam Backup & Replication is installed.

2. Select the authentication method that will be used for database connection:
   - If you plan to use the Microsoft Windows authentication, consider that the current service account will be used (that is, the account under which the Veeam Backup Service is running).
   - If you plan to use the Microsoft SQL Server authentication, provide a login name and password. To view the entered password, click and hold the eye icon on the right of the **Password** field.

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.



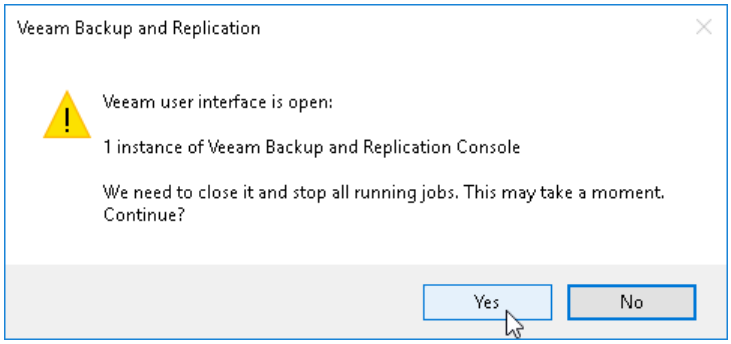
Step 3. Apply Connection Settings

Before proceeding, the utility validates the specified settings to make sure that the user account has enough privileges to access the database.

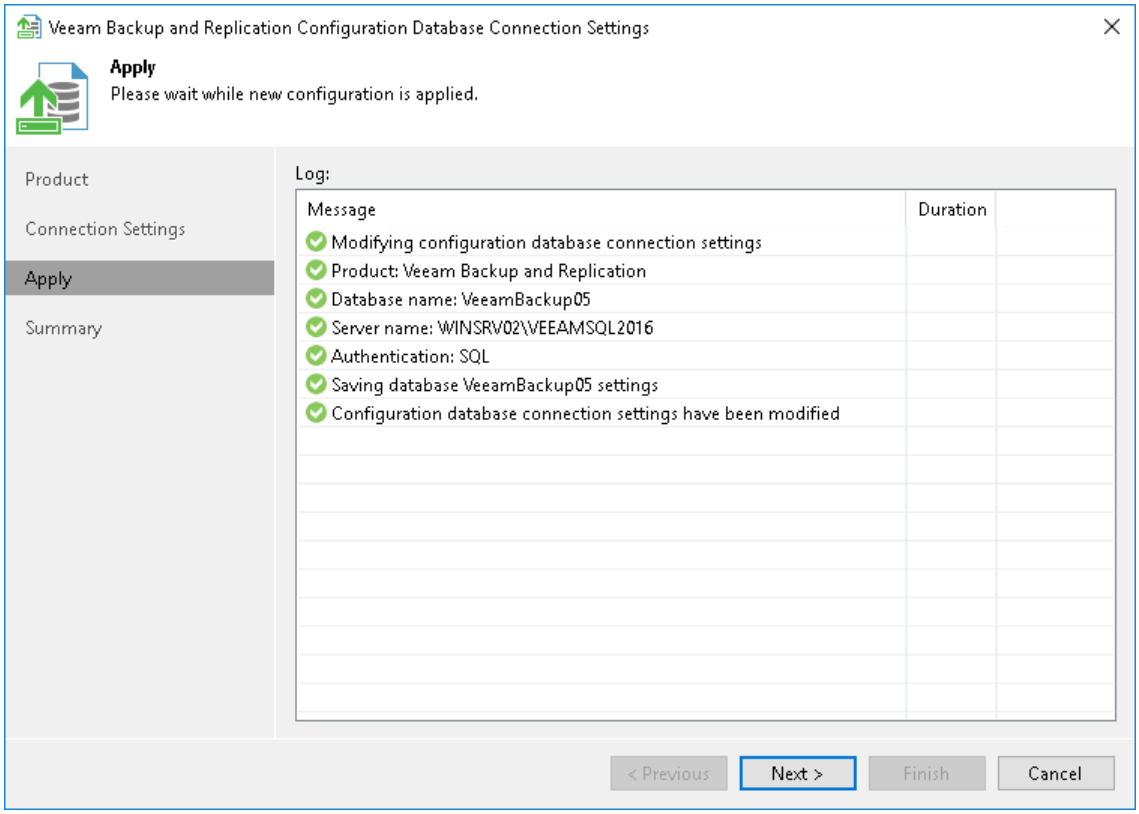
- If you have selected the Microsoft Windows authentication method, the utility will check the privileges of the current user account (that is, the account under which the utility is running) to connect to specified Microsoft SQL Server.
- If you have selected the Microsoft SQL Server authentication method, the utility will check the privileges of the account you have specified.

To ensure that these accounts (as well as the account under which the Veeam Backup Service is running) have sufficient privileges for database access, you can contact your database administrator. Refer to the list of required permissions for Veeam Backup & Replication for detailed information.

For the new settings to be applied, the utility needs to stop Veeam Backup & Replication services and jobs that are currently running. Before proceeding to the **Apply** step, you must confirm the operation. After you confirm the operation by clicking **Yes**, Veeam Backup & Replication will force services and jobs to stop, and will apply database connection settings. For example, if you are configuring Veeam Backup & Replication database settings, the following prompt will be displayed.



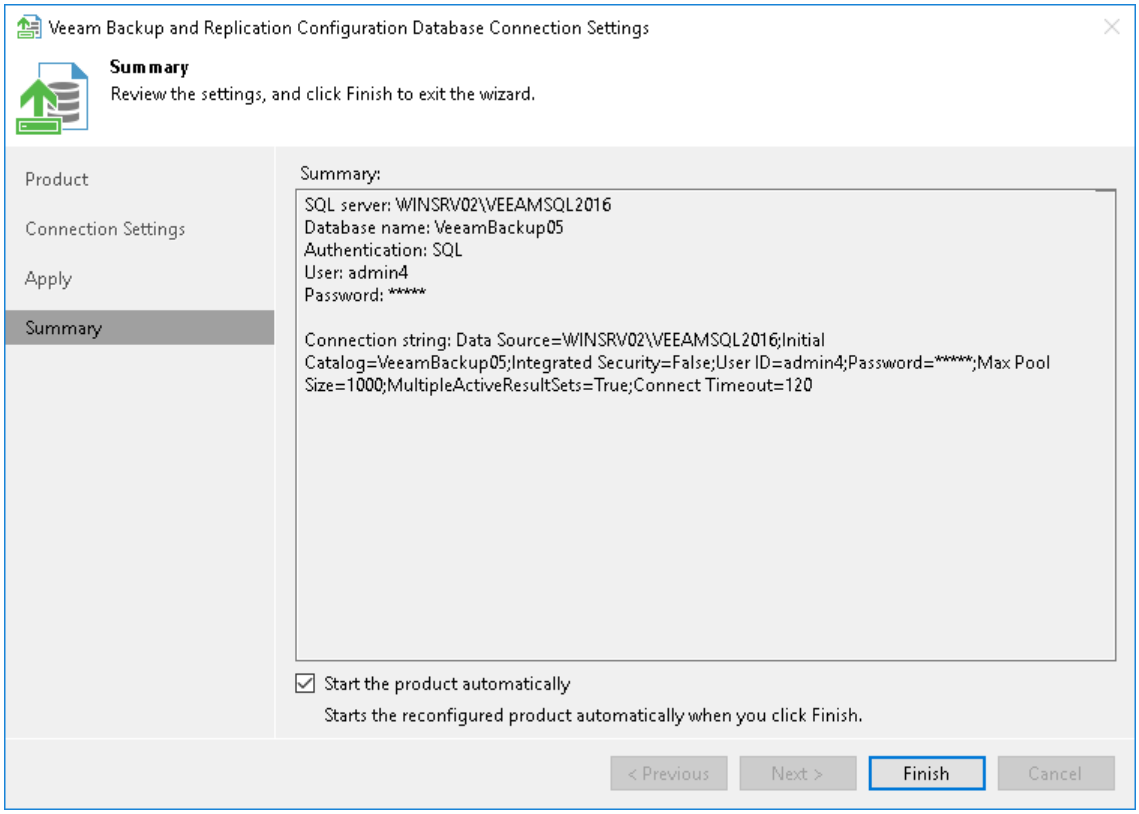
Wait for the operation to complete and click **Next** to proceed to the **Summary** step of the wizard. Previously stopped services will be started again at this moment.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, view the information about the changes in database connection settings. If you were configuring Veeam Backup & Replication database settings and you want the Veeam backup management console to be opened automatically after you finish working with the wizard, select the **Start the product automatically** check box.

> **NOTE**
>
> The **Start the product automatically** option is not available for Veeam Backup Enterprise Manager.



Veeam Backup Validator

Veeam Backup Validator is an utility that verifies the integrity of a backup file without extracting VM data. Veeam Backup Validator is a command-prompt CRC check utility that tests a backup at the file level. You may need this utility to check whether backup files were damaged: for example, after hardware failures occurred in a backup storage side or if backup files were transferred over network.

For integrity validation, Veeam Backup Validator uses the checksum algorithm. When Veeam Backup & Replication creates a backup of a VM, it calculates a checksum for every data block in the backup file and attaches these checksums to data blocks. Veeam Backup Validator recalculates checksums for data blocks and compares them against the initial checksum values. If the results match, the backup file is viable.

Using Veeam Backup Validator

Veeam Backup Validator is located on the backup server in the installation folder of Veeam Backup & Replication — by default, `%ProgramFiles%\Veeam\Backup and Replication\Backup\Veeam.Backup.Validator.exe`. If the default path was changed, you can find the actual path in the following registry value: `[HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication] CorePath`.

To run the utility, open the command prompt on the backup server and change the current folder to the folder where Veeam Backup Validator is located.

> **IMPORTANT**
>
> Mind the following:
>
> - To run Veeam Backup Validator, you must use an account with administrative rights on the local machine.
> - You cannot use Veeam Backup Validator to verify integrity of NAS backups.
> - You cannot use Veeam Backup Validator to verify integrity of backups created in Veeam Cloud Connect repositories. For more information on Veeam Cloud Connect repositories, see the Cloud Repository section in the Veeam Cloud Connect Guide.
> - [For backups stored in scale-out backup repositories] Veeam Backup Validator can validate only 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.

Syntax

Veeam Backup Validator provides parameter sets that allow you to:

- Display Veeam Backup Validator help information.

```
Veeam.Backup.Validator /?
```

- Validate integrity of the content of all VMs or selected VMs in the specified backup.

```
Veeam.Backup.Validator.exe /backup:backupname|backupid [/vmname:vmname] [/
point:pointid] [/date:pointdate] [/time:pointtime] [/silence] [/skip] [/re
port:reportpath [/format:xml|html]]
```

- Validate integrity of the VM content in the specified backup file.

```
Veeam.Backup.Validator.exe /file:backupfile{1..*} [/username:username /pas
sword:password] [/vmname:vmname]
[/silence] [/skip] [/report:reportpath [/format:xml|html]]
```

Parameters

Parameter	Description	Required/Optional	Parameter Type	Notes
/backup: backupname| backupID	Specifies a name or an ID* of a backup or backup copy job that you want to validate.	Required	String	For a backup copy job in the immediate mode, you must specify the name of its child job — task that copies backup job added as a source to the backup copy job. For example, if the "My Copy" backup copy job copies "Daily Backup" backup job, then you must set the parameter value to "My Copy\Daily Backup". If you want to validate the whole backup copy job, you must run the utility for each child job.
/file:backupfile{1..*}	Specifies backup files (VBM, VBK, VIB) to be validated.	Required	String	Mind the following: • If the file is located on a network share, make sure you specify a full path, for example: `\\172.16.16.198\TestShare\Empty VM encryptedD2017-09-22T172639.vbk` • Mapped network drives are not supported. For example, you cannot specify `z:\\172.16.16.198\TestShare`.
/username: username /password: password	To access files on network shares. Specifies account credentials that the utility uses to access network shares.	Required for network share	String	If you want to validate files located on different shares, make sure this account has access rights to all these shares.

Parameter	Description	Required/Optional	Parameter Type	Notes
/vmname:vmname	Specifies a name of the VM in the backup file to be validated.	Optional	String	If not specified, Veeam Backup Validator will check all VMs in the backup file.
/point:pointID	Specifies an ID* of the restore point to be validated.	Optional	String	If not specified, Veeam Backup Validator will verify the latest restore point, that is, all backup files the restore point consists of.
/date:pointdate	Specifies the date when the validated restore point was created.	Optional	Date	Make sure to specify the date in the same format as used on the Veeam Backup server. For example: • For the `mm/dd/yyyy` format, specify 08.30.2012. • For the `dd/mm/yyyy` format, specify 30.08.2012.
/time:pointtime	Specifies the approximate time when the validated restore point was created.	Optional	Time	—
/silence	Defines whether to run validation in the silence mode.	Optional	Boolean	—

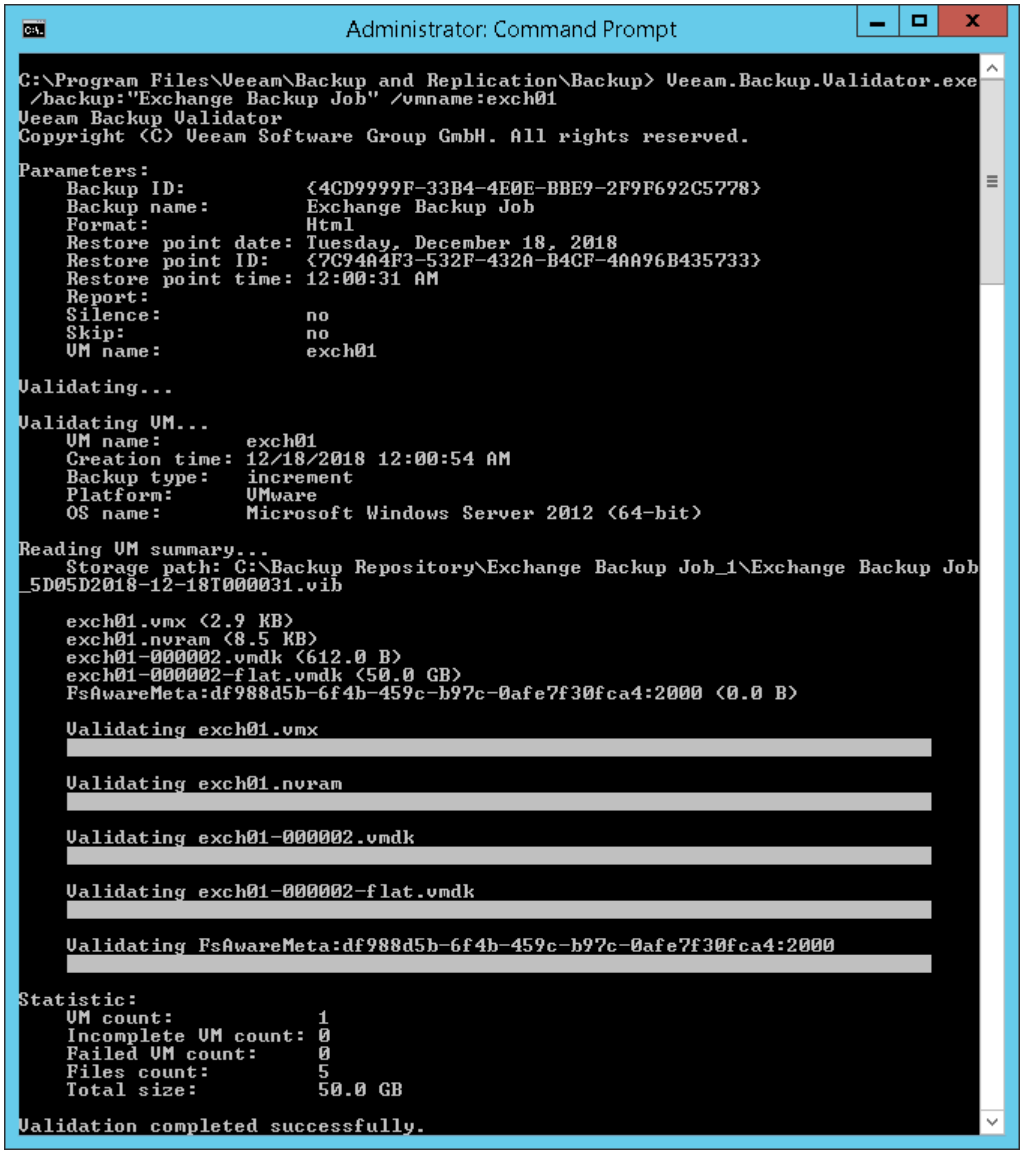
Parameter	Description	Required/Optional	Parameter Type	Notes
/skip	Defines whether to skip from processing VMs listed in the `vmname` parameter.	Optional	Boolean	In the `vmname` parameter, list all VMs that you want to skip.
/report: reportpath [/format:xml|html]	Specifies a path where to store a report on validation results. The utility will generate a report on validation results and store it at the specified path.	Optional	String	Supported report formats are HTML and XML.

* You can get IDs of backup jobs and restore points from the Veeam Backup & Replication database using scripts or Management Studio.

Example 1

This command validates the exch01 VM in the Exchange Backup Job file.

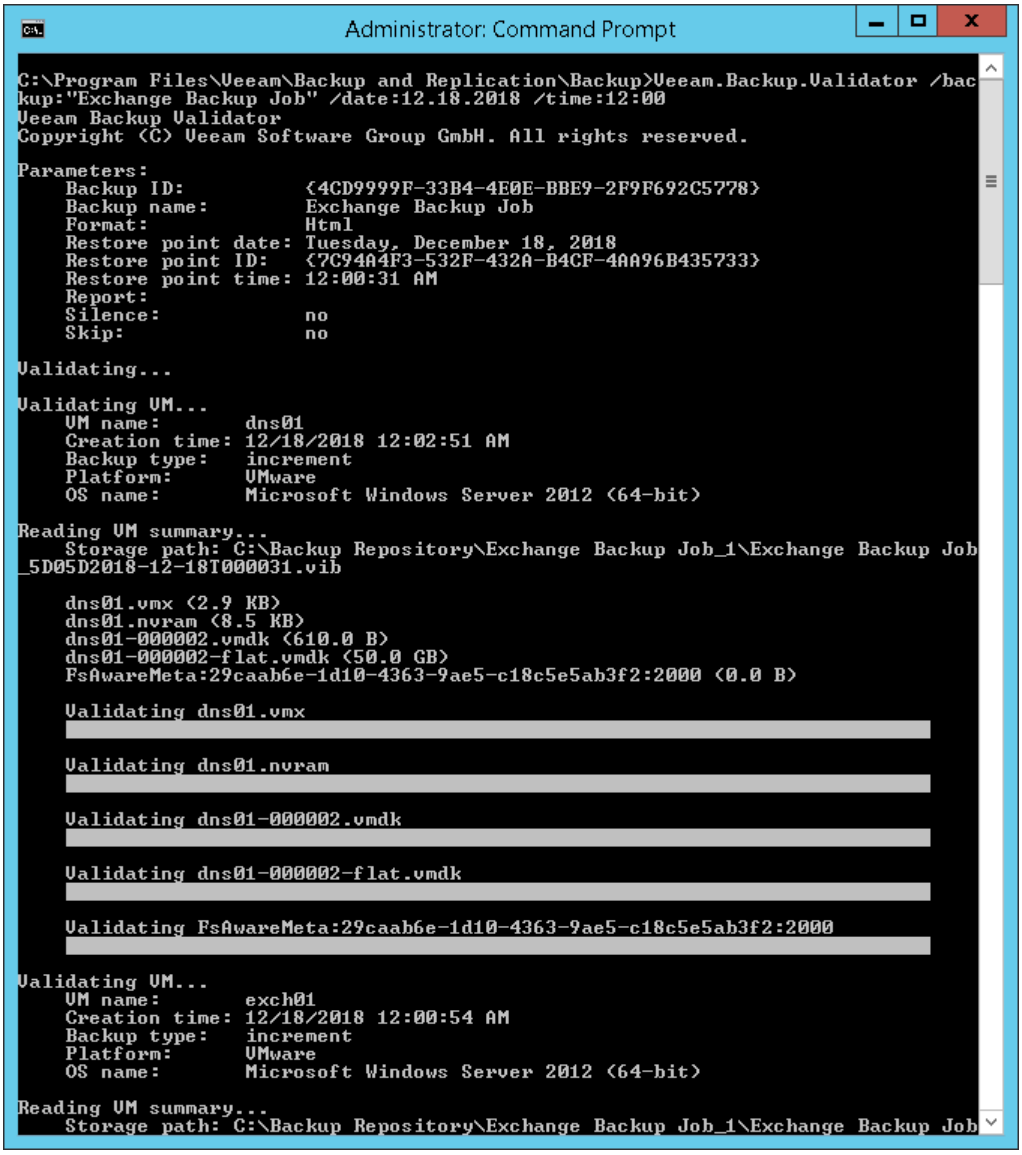
```
Veeam.Backup.Validator /backup:"Exchange Backup Job" /vmname:exch01
```



Example 2

This command validates all VMs in the Exchange Backup Job file created on December 18, 2018 around 12:00 AM.

```
Veeam.Backup.Validator /backup:"Exchange Backup Job"/date:12.18.2018 /time:12:00
```


Administrator: Command Prompt
C:\Program Files\Veeam\Backup and Replication\Backup>Veeam.Backup.Validator /backup:"Exchange Backup Job" /date:12.18.2018 /time:12:00
Veeam Backup Validator
Copyright (C) Veeam Software Group GmbH. All rights reserved.
Parameters:
Backup ID: {4CD9999F-33B4-4E0E-BBE9-2F9F692C5778}
Backup name: Exchange Backup Job
Format: Html
Restore point date: Tuesday, December 18, 2018
Restore point ID: {7C94A4F3-532F-432A-B4CF-4AA96B435733}
Restore point time: 12:00:31 AM
Report:
Silence: no
Skip: no
Validating...
Validating VM...
VM name: dns01
Creation time: 12/18/2018 12:02:51 AM
Backup type: increment
Platform: VMware
OS name: Microsoft Windows Server 2012 (64-bit)
Reading VM summary...
Storage path: C:\Backup Repository\Exchange Backup Job_1\Exchange Backup Job_5D05D2018-12-18T000031.vib
dns01.vmx (2.9 KB)
dns01.nvram (8.5 KB)
dns01-000002.vmdk (610.0 B)
dns01-000002-flat.vmdk (50.0 GB)
FsAwareMeta:29caab6e-1d10-4363-9ae5-c18c5e5ab3f2:2000 (0.0 B)
Validating dns01.vmx
Validating dns01.nvram
Validating dns01-000002.vmdk
Validating dns01-000002-flat.vmdk
Validating FsAwareMeta:29caab6e-1d10-4363-9ae5-c18c5e5ab3f2:2000
Validating VM...
VM name: exch01
Creation time: 12/18/2018 12:00:54 AM
Backup type: increment
Platform: VMware
OS name: Microsoft Windows Server 2012 (64-bit)
Reading VM summary...
Storage path: C:\Backup Repository\Exchange Backup Job_1\Exchange Backup Job

Veeam Backup & Replication Events

This section provides information about events written by Veeam Backup & Replication to Microsoft Windows Event Log on the backup server.

- License
- Repository
- Proxy Server
- WAN Accelerator
- Veeam Backup & Replication Jobs
- Restore
- SureBackup
- Tape

About this Document

This document provides information about events written by Veeam Backup & Replication 11.x to Windows Event Log on the Veeam backup server.

Intended Audience

This document is intended for IT managers, virtual infrastructure administrators, and other personnel responsible for the product deployment and operation. This document assumes that you have a good understanding of Veeam Backup & Replication.

General

The following table lists Veeam Backup & Replication events related to the general settings, credentials, passwords and so on.

Event ID	Message Summary	Message Details	Severity
25300	Credentials record added	Credentials <*Credentials name*> have been added.	Info, Warning, Error
25400	Credentials record updated	Credentials <*Credentials name*> have been updated.	Info, Warning, Error
25500	Credentials record deleted	Credentials <*Credentials name*> have been deleted.	Info, Warning, Error
31000	General options modified	General options have been modified.	Info, Warning, Error
31100	Network traffic rules updated	Network traffic rules have been modified.	Info, Warning, Error
31500	Configuration backup job updated	Configuration backup job has been modified.	Info, Warning, Error
31800	Encryption password deleted	Encryption password <*Password hint*> has been deleted.	Info, Warning, Error
31600	Password added to the password manager	Encryption password <*Password id*> has been added.	Info, Warning, Error
31700	Password stored in the password manager changed	Encryption password <*Password id*> has been changed.	Info, Warning, Error

License

The following table lists Veeam Backup & Replication events related to the product license usage.

Event ID	Message Summary	Message Details	Severity
24010	License installed	*<License type>* License key for Veeam Backup & Replication *<Edition>* has been installed.	Info
24020	License expiring	*<License type>* License key for Veeam Backup & Replication *<Edition>* is about to expire in *<Number of days>* Days.	Warning
24022	License evaluation expiring	*<License type>* Evaluation license key for Veeam Backup & Replication *<Edition>* is about to expire in *<Number of days>* Days.	Warning
24030	License expired	*<License type>* License key for Veeam Backup & Replication *<Edition>* has expired.	Error
24040	License support expiring	Support contract for Veeam Backup & Replication is about to expire in *<Number of days>* Days.	Warning
24050	License support expired	Support contract for Veeam Backup & Replication has expired. Contact Veeam sales representative to renew your support contract.	Error
24060	License exceeded	License exceeded.	Error
24070	License grace period entered	License grace period entered.	Info
24080	Part of merged license deleted	*<License type>* license that was a part of the merged Veeam Backup & Replication license has been removed.	Info

Servers and Hosts

The following table lists Veeam Backup & Replication events related to virtualization servers and hosts.

Event ID	Message Summary	Message Details	Severity
25600	Host added	New hypervisor host has been added.	Info, Warning, Error
25700	Host deleted	Hypervisor host has been deleted.	Info, Warning, Error
25800	Host updated	Hypervisor host has been updated.	Info, Warning, Error
28300	Host added	Host <*Host name*> (<*Host type*>) has been created.	Info, Warning, Error
28400	Host updated	Host <*Host name*> (<*Host type*>) has been modified.	Info, Warning, Error
28500	Host deleted	Host <*Host name*> (<*Host type*>) has been deleted.	Info, Warning, Error

Repositories

The following table lists Veeam Backup & Replication events related to repositories: backup repositories, scale-out repositories and so on.

Event ID	Message Summary	Message Details	Severity
21220	Repository server established connection	Connection to backup repository "*<Repository name>*" has been restored.	Info
21224	Repository server lost connection	Backup server has lost connection to backup repository "*<Repository name>*".	Error
28000	Repository created	Backup repository *<Repository name>* (*<Repository type>*) has been created.	Info
28100	Repository modified	Backup repository *<Repository name>* (*<Repository type>*) has been modified.	Info
28200	Repository deleted	Backup repository *<Repository name>* (*<Repository type>*) has been deleted.	Info
29700	Object storage repository created	Object storage repository "*<Repository name>*" has been created.	Info
29800	Object storage repository modified	Object storage repository "*<Repository name>*" has been modified.	Info
29900	Object storage repository deleted	Object storage repository "*<Repository name>*" has been deleted.	Info
30000	Scale-out backup repository created	Scale-out backup repository "*<Repository name>*" has been created.	Info
30100	Scale-out backup repository modified	Scale-out backup repository "*<Repository name>*" has been modified.	Info
30200	Scale-out backup repository deleted	Scale-out backup repository "*<Repository name>*" has been deleted.	Info
32000	External repository created	External repository *<Repository name>* (*<Repository type*) has been created.	Info

Event ID	Message Summary	Message Details	Severity
32100	External repository modified	External repository *<Repository name>* (*<Repository type>*) has been modified.	Info
32200	External repository deleted	External repository *<Repository name>* (*<Repository type>*) has been deleted.	Info
40002	Management session finished	The storage management session for repository *<Repository name>* has finished: *<State>*. Repository name: *<Repository name>* Job type: *<Offload/Archiving/Download>* Session result: *<Successful/Warning/Failed>*	Info, Warning, Error

Proxy Server

The following table lists Veeam Backup & Replication events related to backup proxy servers.

Event ID	Message Summary	Message Details	Severity
21210	Proxy server established connection	Connection to backup proxy "<*Proxy name*>" has been restored.	Info
21214	Proxy server lost connection	Backup server has lost connection to backup proxy "<*Proxy name*>".	Error
27700	Backup proxy created	Backup proxy <*Proxy name*> (<*Proxy type*>) has been created.	Info
27800	Backup proxy modified	Backup proxy <*Proxy name*> (<*Proxy type*>) has been modified.	Info
27900	Backup proxy deleted	Backup proxy <*Proxy name*> (<*Proxy type*>) has been deleted.	Info

WAN Accelerator

The following table lists Veeam Backup & Replication events related to WAN accelerators.

Event ID	Message Summary	Message Details	Severity
21230	WAN accelerator established connection	Connection to WAN accelerator "*<Wan name>*" has been restored.	Info
21234	WAN accelerator lost connection	Backup server has lost connection to WAN accelerator "*<Wan name>*".	Error
27100	WAN accelerator created	WAN accelerator *<Wan name>* has been created.	Info
27200	WAN accelerator modified	WAN Accelerator *<Wan name>* has been modified.	Info
27300	WAN accelerator deleted	WAN Accelerator *<Wan name>* has been deleted.	Info
33000	WAN accelerator cache populated	Cache of WAN accelerator *<Wan name>* has been populated.	Info
33001	WAN accelerator cache cleared	Cache of WAN accelerator *<Wan name>* has been cleared.	Info

Jobs

The following table lists Veeam Backup & Replication events related to the Veeam Backup & Replication jobs (including backup, replication, backup copy, file copy, VM copy and Quick Migration).

Event ID	Message Summary	Message Details	Severity	Source
110	Backup job started	*<Job type>* job "*<Job name>*" has been started.	Info	Backup
150	Backup task finished	VM "*<Task name>*" task has finished with "*<State name>*" state.	Info	Backup, Replication, SureBackup
190	Backup job finished	The *<Job type>* job "*<Job name>*" has finished with *<State name>* state.	Info, Warning, Error	Backup
410	Backup copy job started	*<Job type>* job "*<Job name>*" has been started.	Info	Backup copy
450	Backup copy task finished	VM *<Task name>* task has finished with "*<State name>*" state.	Info, Warning, Error	Backup copy
490	Backup copy job finished	The *<Job type>* job "*<Job name>*" has finished with *<State name>* state.	Info, Warning, Error	Backup copy
510	File copy job started	*<Job type>* job "*<Job name>*" has been started.	Info	File copy
512	VM copy job started	*<Job type>* job "*<Job name>*" has been started.	Info	VM copy
590	File copy job finished	The *<Job type>* job "*<Job name>*" has finished with *<State name>* state.	Info, Warning, Error	File copy

Event ID	Message Summary	Message Details	Severity	Source
592	VM copy job finished	The *<Job type>* job "*<Job name>*" has finished with *<State name>* state.	Info, Warning, Error	VM copy
610	Quick migration job has been finished	QuickMigration job finished with "*<State name>*" state.	Info, Warning, Error	Quick migration
710	Veeam Agent backup job started	Agent Backup job "*<Job name>*" has been started.	Info	Veeam Agent backup
790	Veeam Agent backup job finished	Agent Backup job "*<Job name>*" finished with "*<State name>*". Job details: "*<Details information>*".	Info, Warning, Error	Veeam Agent backup
810	Rescan job started (Veeam Agent)	Rescan job "*<Job name>*" has been started.	Info	Rescan job (Veeam Agent)
890	Rescan job finished (Veeam Agent)	Rescan job "*<Job name>*" has been finished with "*<State name>*". Job details: "*<Details information>*".	Info, Warning, Error	Rescan job (Veeam Agent)
10010	Restore point created	VM "*<Vm name>*" restore point has been created.	Info	
10050	Restore point removed	Restore point for VM "*<Vm name>*" has been removed according to the configured retention policy.	Info	
10090	Restore point updated	VM "*<Vm name>*" restore point has been modified.	Info	

Event ID	Message Summary	Message Details	Severity	Source
23010	Job created	Job [*<Job name>*] has been created.	Info	
23050	Job modified	Job *<Job name>* has been modified.	Info, Warning, Error	
23090	Job deleted	The *<Job type>* job "*<Job name>*" has been deleted.	Info	Backup, Replication, Backup copy
23110	Job objects added	"*<object count>*" objects have been created for "*<Job name>*".	Info	Backup, Replication, Backup copy
25900	Failover plan created	New failover plan has been created.	Info	Replication
26000	Failover plan updated	Failover plan has been updated.	Info	Replication
26100	Failover plan deleted	Failover plan has been deleted.	Info	Replication
26600	Failover plan started	Failover plan was started.	Info	Replication
26700	Failover plan stopped	Failover plan was stopped.	Info	Replication
32120	Job objects deleted	"*<object count>*" objects have been deleted for "*<Job name>*".	Info	Backup, Replication, Backup copy
23130	Job objects modified	"*<object count>*" objects have been modified for "*<Job name>*".	Info	Backup, Replication, Backup copy

Event ID	Message Summary	Message Details	Severity	Source
23410	Second destination job assigned	"<*Job name*>" job has been assigned as secondary destination for "<*Job count*>" jobs	Info	Backup copy
23420	Second destination job removed	"<*Job name*>" job has been removed as secondary destination for "<*Job count*>" jobs.	Info	Backup copy
40002	Offload/Download/Archiving job finished	The storage management session for repository <*Repository name*> has finished: <*State name*>.	Info, Warning, Error	Offload/Download/Archiving job

Nas Backup

The following table lists Veeam Backup & Replication events related to NAS backup functionality.

Event ID	Message Summary	Message Details	Severity
151	Task finished	Source "<*Task name*>" task has finished with "<*State name*>" state.	Info
28900	File share created	File share <*Server name*> (<*Server type*>) has been created.	Info
28910	File share modified	File share <*Server name*> (<*Server type*>) has been modified.	Info
28920	File share deleted	File share <*Server name*> (<*Server type*>) has been deleted.	Info
28930	File server created	File server <*Server name*> has been created.	Info
28940	File server modified	File server <*Server name*> has been modified.	Info
28950	File server deleted	File server <*Server name*> has been deleted.	Info

Restore

The following table lists Veeam Backup & Replication events related to data recovery.

Event ID	Message Summary	Message Details	Severity
210	Restore session started	Restore session has been initiated by "*<User name>*".	Info
250	Restore VI task finished	Restore for "*<Object name>*" has finished with *<State name>* state	Info, Warning, Error
251	Restore HV task finished	Restore for "*<Object name>*" has finished with *<State name>* state	Info, Warning, Error
290	Restore session finished	The restore session has finished with *<State name>* state.	Info, Warning, Error
Event ID	Message Summary	Message Details	Severity
210	Restore session started	Restore session has been initiated by "*<User name>*".	Info
250	Restore VI task finished	Restore for "*<Object name>*" has finished with *<State name>* state	Info, Warning, Error
290	Restore session finished	The restore session has finished with *<State name>* state.	Info, Warning, Error

SureBackup

The following table lists Veeam Backup & Replication events related to SureBackup functionality.

Event ID	Message Summary	Message Details	Severity
310	SureBackup job started	*<Job type>* job "*<Job name>*" has been started.	Info
350	SureBackup VI verification task finished	VM *<Vm name>* verification task has finished with *<State name>* state.	Info, Warning, Error
351	SureBackup VI validation task finished	VM *<Vm name>* validation task has finished with *<State name>* state.	Info, Warning, Error
390	SureBackup job finished	The *<Job type>* job "*<Job name>*" has finished with *<State name>* state.	Info, Warning, Error
23210	SureBackup job created	The SureBackup job "*<Job Name>*" has been created.	Info
23220	SureBackup job modified	The SureBackup job "*<Job Name>*" has been modified.	Info
23230	SureBackup job deleted	The SureBackup job "*<Job Name>*" has been deleted.	Info
23310	SureBackup objects added	"*<object count>*" objects have been created for "*<Job Name>*".	Info
23320	SureBackup objects deleted	"*<object count>*" objects have been deleted from "*<Job Name>*".	Info
23330	SureBackup objects modified	"*<object count>*" objects have been modified for "*<Job Name>*".	Info
30300	Application group added	Application group [*<Group name>*] has been added.	Info
30400	Application group modified	Application group *<Group name>* has been modified.	Info
30500	Application group deleted	Application group *<Group name>* has been deleted.	Info

Event ID	Message Summary	Message Details	Severity
30600	Virtual lab created	Virtual lab <*Lab name*> has been created.	Info
30700	Virtual lab modified	Virtual lab <*Lab name*> has been modified.	Info
30800	Virtual lab deleted	Virtual lab <*Lab name*> has been deleted.	Info

Tape

The following table lists Veeam Backup & Replication events related to tape archiving.

Event ID	Message Summary	Message Details	Severity
114	Tape backup job started	*<Job type>* job "*<Job name>*" has been started.	Info
194	Tape backup job finished	The *<Job type>* job "*<Job name>*" has finished with *<State name>* state.	Info, Warning, Error
10014	Tape restore point created	VM "*<Task name>*" restore point has been offloaded to the tape *<Barcode>*.	Info
20114	Tape updated	Tape *<Barcode>* has been renamed.	Info
23440	Tape backup job created	<*Job type*> "<*Job name*>" has been created.	Info
23450	Tape backup job modified	<*Job type*> "<*Job name*>" has been modified.	Info
23490	Tape backup job deleted	<*Job type*> "<*Job name*>" has been deleted.	Info
28600	Tape server created	Tape server <*Server name*> has been created.	Info
28700	Tape server modified	Tape server <*Server name*> has been modified.	Info
28800	Tape server deleted	Tape server <*Server name*> has been deleted.	Info

Veeam Agent for Linux

Version 5.0

User Guide

January, 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 **8**

ABOUT THIS DOCUMENT **9**

OVERVIEW **10**

Solution Architecture 11

Standalone and Managed Operation Modes 12

Data Backup 13

Backup Types 14

How Backup Works 20

Backup Job 25

Backup Repository 29

Backup Chain 31

Backup of Database Systems 36

Data Restore 43

Volume-Level Restore 44

File-Level Restore 45

Veeam Recovery Media 46

Veeam Recovery Media Versions 47

Drivers in Veeam Recovery Media 49

Integration with Veeam Backup & Replication 50

Backup to Veeam Cloud Connect Repository 52

MANAGING VEEAM AGENT IN VEEAM BACKUP & REPLICATION **53**

PLANNING AND PREPARATION **54**

System Requirements 55

Used Ports 62

INSTALLATION AND CONFIGURATION **64**

Before You Begin 65

Installing Veeam Agent for Linux 66

Installing Veeam Agent for Linux in Offline Mode 70

Upgrading Veeam Agent for Linux 75

Granting Permissions to Users 79

Performing Initial Setup 80

Step 1. Accept License Agreements 81

Step 2. Create Custom Veeam Recovery Media 82

Step 3. Install Product License 85

Configuring Advanced Settings 87

Uninstalling Veeam Agent for Linux 89

GETTING STARTED **90**
GETTING TO KNOW USER INTERFACE **91**
Veeam Agent for Linux Control Panel 92
Command Line Interface 94
Viewing Help 96
LICENSING **97**
Product Editions 98
License Agreement 99
Installing License 100
Viewing License Information 102
Removing License 103
License Expiration 104
Managing License with Command Line Interface 105
Accepting License Agreements 106
Installing License 107
Viewing License Information 108
Removing License 109
PERFORMING BACKUP **110**
Creating Custom Veeam Recovery Media 111
Creating Backup Jobs 113
Creating Backup Job with Backup Job Wizard 114
Creating Backup Job with Command Line Interface 163
Starting and Stopping Backup Jobs 206
Starting Backup Job from Control Panel 207
Starting Backup Job from Command Line Interface 208
Creating Active Full Backups 209
Stopping Backup Job 210
Managing Backup Jobs 212
Viewing List of Backup Jobs 213
Viewing Backup Job Settings 214
Editing Backup Job Settings 216
Deleting Backup Job 222
Managing Backup Repositories 224
Creating Backup Repository 225
Viewing List of Backup Repositories 229
Editing Backup Repository 230
Rescanning Veeam Backup Repository 232
Deleting Backup Repository 233
Managing Veeam Backup & Replication Servers 234

Connecting to Veeam Backup Server 235
Viewing List of Veeam Backup Servers 237
Viewing Backup Server Details 238
Editing Connection to Veeam Backup Server 239
Updating List of Veeam Backup Repositories 241
Deleting Connection to Veeam Backup Server 242
Managing Service Providers 243
Connecting to Service Provider 244
Viewing List of Service Providers 245
Editing Connection to Service Provider 246
Updating List of Cloud Repositories 249
Deleting Connection to Service Provider 250
Managing Backups 251
Viewing Backups 252
Viewing Backup Details 253
Viewing Restore Points in Backup 255
Importing Backups 256
Deleting Backups 259
PERFORMING RESTORE 260
Restoring from Veeam Recovery Media 261
Restoring Volumes 262
Restoring Files and Folders 302
Restoring Volumes with Command Line Interface 333
Before You Begin 334
Restoring from Backup 335
Restoring from Restore Point 340
Restoring Files and Folders with Recovery Wizard 346
Before You Begin 347
Step 1. Launch File Level Restore Wizard 348
Step 2. Select Backup and Restore Point 349
Step 3. Save Restored Files 353
Step 4. Stop Backup Mount Session 354
Restoring Files and Folders with Command Line Interface 355
Before You Begin 356
Restoring from Backup 357
Restoring from Restore Point 365
Exporting Backup to Virtual Disk 373
Exporting Backups 374
Exporting Restore Points 376

Restoring Data from Encrypted Backups 378
REPORTING 382
Viewing Job Session Progress 383
Viewing Real-Time Job Session Statistics 384
Viewing Job Session Result 386
Viewing Session Status 388
Viewing Session Logs 390
MANAGING CONFIGURATION DATABASE 391
Exporting Configuration Database 392
Importing Configuration Database 393
EXPORTING PRODUCT LOGS 394
Exporting Logs with Control Panel 395
Exporting Logs with Command Line Interface 397
GETTING SUPPORT 398
USING WITH VEEAM BACKUP & REPLICATION 399
Setting Up User Permissions on Backup Repositories 401
Managing License 404
Managing Instance Consumption by Veeam Agents 405
Assigning License to Veeam Agent 406
Viewing Licensed Agents and Revoking License 407
Performing Data Protection Tasks 409
Backing Up to Backup Repositories 410
Backing Up to Cloud Repositories 411
Performing Backup Copy for Veeam Agent Backups 413
Archiving Veeam Agent Backups to Tape 415
Performing Restore Tasks 416
Restoring Veeam Agent Backup to VMware vSphere VM 417
Restoring Veeam Agent Backup to Microsoft Hyper-V VM 419
Restoring Files and Folders 420
Restoring Application Items 421
Exporting Disks 422
Exporting Restore Point to Full Backup File 429
Restoring to Microsoft Azure 430
Restoring to Amazon EC2 431
Restoring to Google Cloud Platform 432
Performing Administration Tasks 433
Importing Veeam Agent Backups 434
Enabling and Disabling Veeam Agent Backup Jobs 436
Deleting Veeam Agent Backup Jobs 437

Viewing Veeam Agent Backup Job Statistics 438
Removing Veeam Agent Backups 439
Viewing Veeam Agent Backup Properties 441
Configuring Global Settings 442
Assigning Roles to Users 443
APPENDIX A. REQUIREMENTS FOR VEEAM-NOSNAP 444

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 Veeam Agent for Linux.

Intended Audience

The user guide is intended for anyone who wants to use Veeam Agent for Linux 5.0 or later to protect their machine.

Overview

Veeam Agent for Linux is a data protection and disaster recovery solution for physical endpoints and virtual machines running Linux-based operating systems.

Veeam Agent can be used by IT administrators who run Linux infrastructure to protect different types of computers and devices: servers, desktops and laptops. The solution runs inside the guest OS and does not need access to virtualization infrastructure components. Thus, Veeam Agent can be used to protect Linux server instances deployed in the public cloud, for example, in Microsoft Azure environment.

> **NOTE**
>
> Veeam Agent can operate in either standalone or managed mode. Depending on the mode, Veeam Agent provides different features and limitations. To learn more, see Standalone and Managed Operation Modes.

Veeam Agent offers a variety of features to protect your data. You can create an entire system image backup, back up specific machine volumes or individual directories and files. Backups can be stored on a local hard drive, on an external hard drive, in a network shared folder or on a Veeam backup repository.

In case of a disaster, you can perform the following restore operations:

- Start the OS from the Veeam Recovery Media and use standard Linux command line tools to diagnose and fix problems.
- Perform bare metal restore.
- Restore necessary data from backups to its original location or a new location.

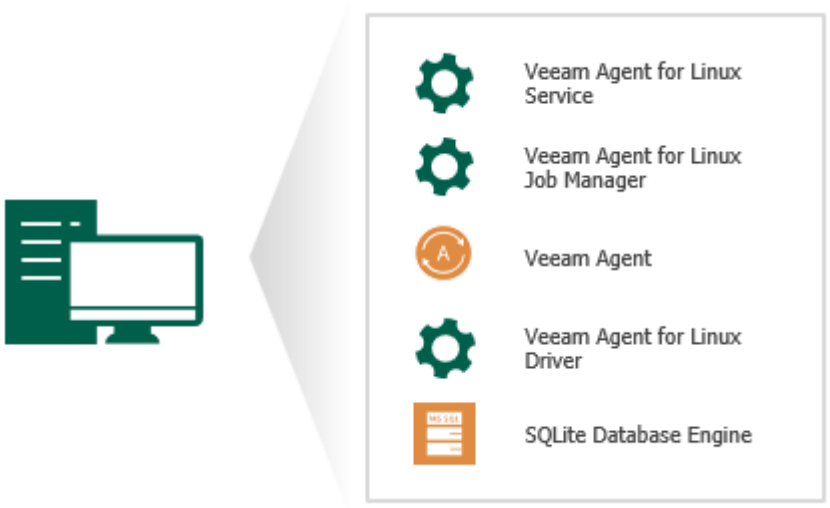
Veeam Agent integrates with Veeam Backup & Replication. Backup administrators who work with Veeam Backup & Replication can perform advanced tasks with Veeam Agent backups: restore files and disks from backups, manage backup jobs configured in Veeam Agent or backups created with these jobs.

Solution Architecture

Veeam Agent for Linux is set up on a Linux-based physical endpoint or virtual machine whose data you want to protect.

When you install the product, Veeam Agent deploys the following components:

- *Veeam Agent for Linux Service (veeamservice)* is a service responsible for managing all tasks and resources in Veeam Agent. The *veeamservice* component is registered as a daemon in the Linux OS upon the product installation. The service is started automatically when you start the OS, and runs in the background.
- *Veeam Agent for Linux Job Manager (veeamjobman)* is a process started by *Veeam Agent for Linux Service* for every backup job session.
- *Veeam Agent* that communicates with the *Veeam Agent for Linux Service* and *Veeam Agent for Linux Job Manager*. *Veeam Agent* is started by *Veeam Agent for Linux Manager* to perform data transfer operations of any kind: copy data from the backed-up volume to the backup location during backup, from the backup location to the target volume during restore, perform data compression, and so on.
- *Veeam Agent for Linux Driver (veeamsnap)* is a Veeam driver (Linux kernel module) used to create volume snapshots in the Linux OS and keep track of changed data blocks.
- To store its configuration data, Veeam Agent uses the SQLite database engine. SQLite requires only few files to install and takes little resources to run on a Linux OS.



Standalone and Managed Operation Modes

Veeam Agent for Linux can operate in two modes: *standalone mode* and *managed mode*.

- **Standalone mode** — in this mode, Veeam Agent for Linux operates as a standalone product. To use Veeam Agent operating in the standalone mode, you must manually install the product directly on the machine whose data you want to protect.

  For Veeam Agent for Linux operating in the standalone mode, data protection, disaster recovery, and administration tasks are performed in Veeam Agent by a user. You can also use Veeam Agent operating in the standalone mode with Veeam Backup & Replication. In this scenario, you can use a Veeam backup repository as a target location for Veeam Agent backups and use the Veeam Backup & Replication console to perform a number of tasks with Veeam Agent backup jobs and backups.

  The current User Guide covers subjects related to Veeam Agent for Linux operating in the standalone mode only.

- **Managed mode** — in this mode, Veeam Agent for Linux operates under control from Veeam Backup & Replication. To use Veeam Agent operating in the managed mode, you must deploy the product in one of the following ways:

  - remotely from Veeam Backup & Replication
  - manually using external tools

    To learn more about managed Veeam Agent deployment, see the Protected Computers Discovery and Veeam Agent Deployment section in the Veeam Agent Management Guide

  For Veeam Agent for Linux operating in the managed mode, data protection, disaster recovery, and administration tasks are performed by a backup administrator in the Veeam Backup & Replication console.

  Features and limitations of Veeam Agent for Linux operating in the managed mode are different from those in the standalone mode. To learn about Veeam Agent operating in the managed mode, see the Veeam Agent Management Guide.

Data Backup

It is recommended that you regularly back up data stored on your machine. Backup creates a safety copy of your data. If any kind of disaster strikes, you can restore your data from the backup and be sure that you will not lose the necessary information.

You can set up Veeam Agent to perform automatic scheduled backups (triggered at specific time of the day), or you can choose to back up data manually when needed. You can back up the entire computer image, specific computer volumes or individual directories and files.

You can set up Veeam Agent to create multiple backups — with individual backup scope, upon individual schedule or in different locations. This functionality is available if Veeam Agent operates in the Server edition. To learn more about editions, see Product Editions.

Backups created with Veeam Agent can be saved to the following locations:

- Removable storage device
- Local computer drive
- NFS or SMB (CIFS) network shared folder
- Backup repository managed by a Veeam backup server
- Veeam Cloud Connect repository

Backup Types

Veeam Agent for Linux lets you create the following backup types:

- Volume-level backup
- File-level backup

Volume-Level Backup

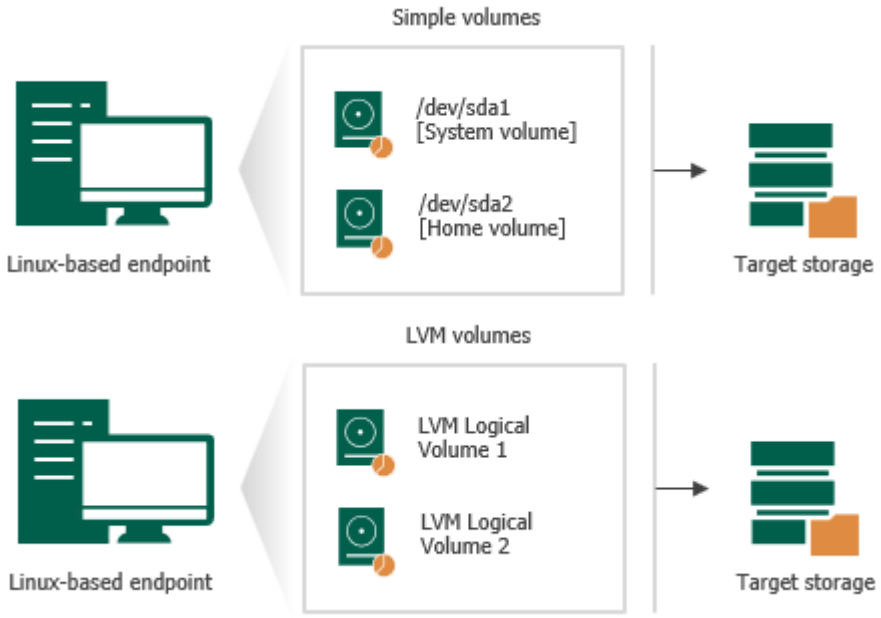
You can set up Veeam Agent for Linux to create volume-level backup. The volume-level backup captures the whole image of a data volume on your computer. You can use the volume-level backup to restore a computer volume, specific files and folders on the volume or perform bare metal recovery.

Veeam Agent for Linux supports backup of the following types of computer volumes:

- Simple volumes
- LVM logical volumes
- Btrfs subvolumes

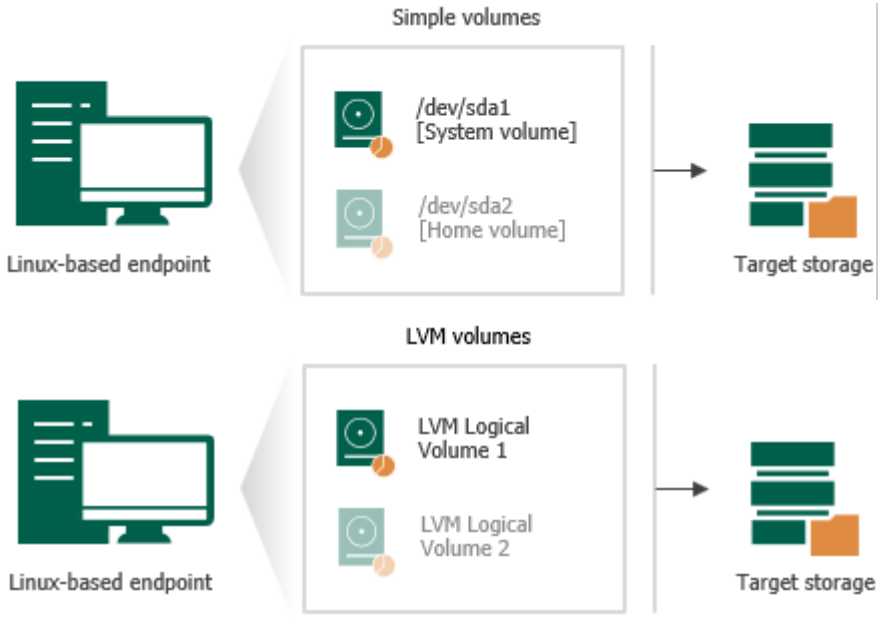
You can back up all computer volumes or specific computer volumes.

- When you back up the entire computer image, Veeam Agent captures the content of all volumes on your computer. The resulting backup file contains all volume data and Linux OS system data: system partition, partition table and bootloader.



- When you back up a specific computer volume, Veeam Agent captures only the data that resides on this specific volume: files, folder, application data and so on.

  If you choose to back up the system volume (volume to which the root file system is mounted), Veeam Agent automatically includes the bootloader into the backup scope.

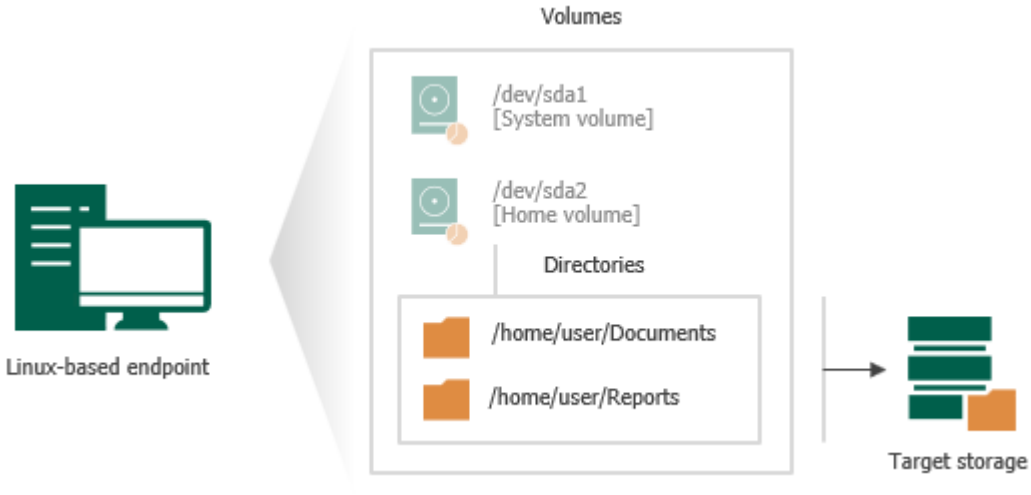


File-Level Backup

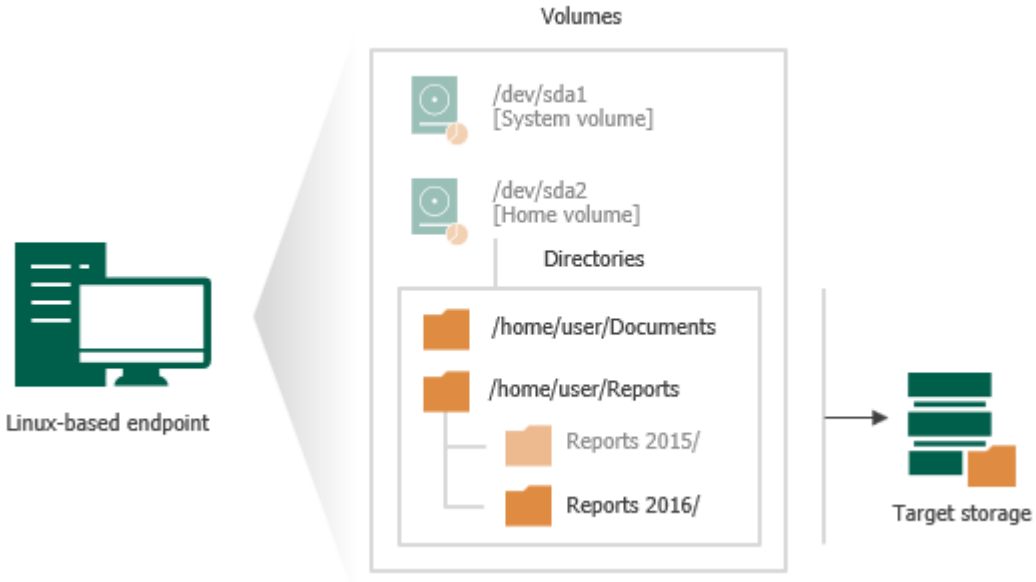
You can set up Veeam Agent for Linux to create file-level backup. The file-level backup captures only data of individual directories and files on the computer. You can use the file-level backup to restore files and directories that you have added to the backup scope.

With Veeam Agent for Linux, you can specify which files and directories to back up:

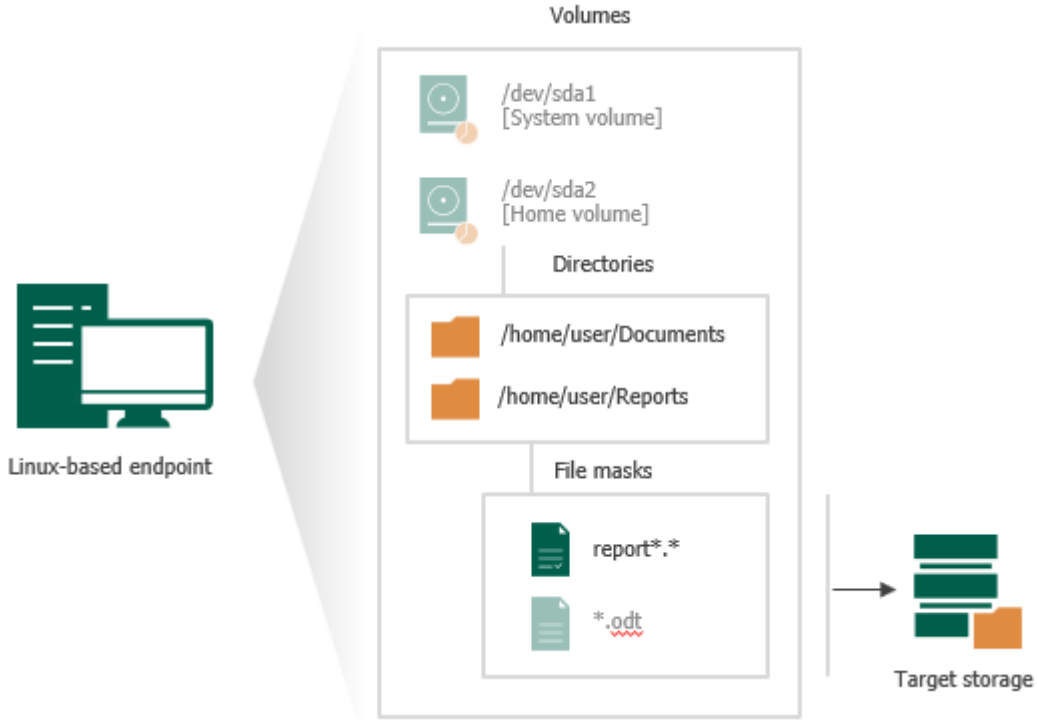
- You can include individual directories in the backup. When you include a directory in the backup, its subdirectories are automatically included in the backup too. When you recover from such backup, you will be able to restore directories that you have selected to back up, all subdirectories of these directories and files in these directories.



- You can exclude from the backup some subdirectories of the directories that are included in the backup. When you recover from such backup, you will be able to restore directories that you have selected to back up, specific subdirectories of these directories and files in these directories.



- You can include or exclude files of a specific type in/from the backup. You can specify file names explicitly or use UNIX wildcard characters to define include and exclude file name masks. When you recover from such backup, you will be able to restore directories that you have selected to back up with files whose names match the specified include masks.



Snapshot-Less File-Level Backup

You can set up Veeam Agent for Linux to create file-level backup in the snapshot-less mode. This allows you to back up data that resides in any file system mounted to the root file system of the Veeam Agent computer. For example, you can use the snapshot-less mode to back up data that resides in a file system that is not supported for snapshot-based backup with Veeam Agent, such as UFS, ZFS, GFS, GFS2 or OCFS2. You can also use it to back up data that resides in an NFS or CIFS network shared folder.

To create backups in the snapshot-less mode, you must enable this mode in the properties of the file-level backup job. To learn more, see Creating Backup Jobs.

In the snapshot-less mode, Veeam Agent does not create a snapshot of the backed-up volume. Instead, when the backup process starts, Veeam Agent reads files and directories that you selected to back up, and copies backed-up data to the target location.

> **IMPORTANT**
>
> During backup in the snapshot-less mode, Veeam Agent does not track whether files and directories have changed in their original location since the time when the backup process started. To make sure that data in the backup is in the consistent state, you must not perform write operations in the file system that contains the backed-up data until the backup process completes.

How Backup Works

Veeam Agent for Linux performs performs backup differently depending on the backup type:

- Volume-level backup
- File-level backup

How Volume-Level Backup Works

During volume-level backup, Veeam Agent performs the following operations for every backup job session:

1. When a new job session starts, Veeam Agent creates a backup file in the target location.
2. In the backup file, Veeam Agent creates a disk for each backed-up disk. In disks, Veeam Agent creates blank partitions that have the same size and location as partitions in backed-up disks.
3. Veeam Agent creates a snapshot of the volume whose data you want to back up. The snapshot is created on the volume that has enough free disk space to contain the snapshot data. To create a snapshot, Veeam Agent uses the *Veeam Agent for Linux Driver*.

   The snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup. If a data block is about to change on disk during backup, Veeam Agent will copy this block to the snapshot. After the data block is overwritten on the source location, its original copy will remain intact in the snapshot.

   > **NOTE**
   >
   > Consider the following:
   >
   > - If you instruct Veeam Agent to back up a database system, Veeam Agent prepares databases for backup before creating a snapshot of the volume. To learn more, see Backup of Database Systems.
   > - During backup of data that resides in the Btrfs file system, Veeam Agent does not use its driver to create a snapshot. Instead, Veeam Agent leverages Btrfs capabilities to create a Btrfs snapshot.

4. [For incremental backup] Veeam Agent uses the *Veeam Agent for Linux Driver* to detect what blocks have changed on the volume since the previous job session. The driver keeps this information as a changed block tracking map in the RAM of your computer.

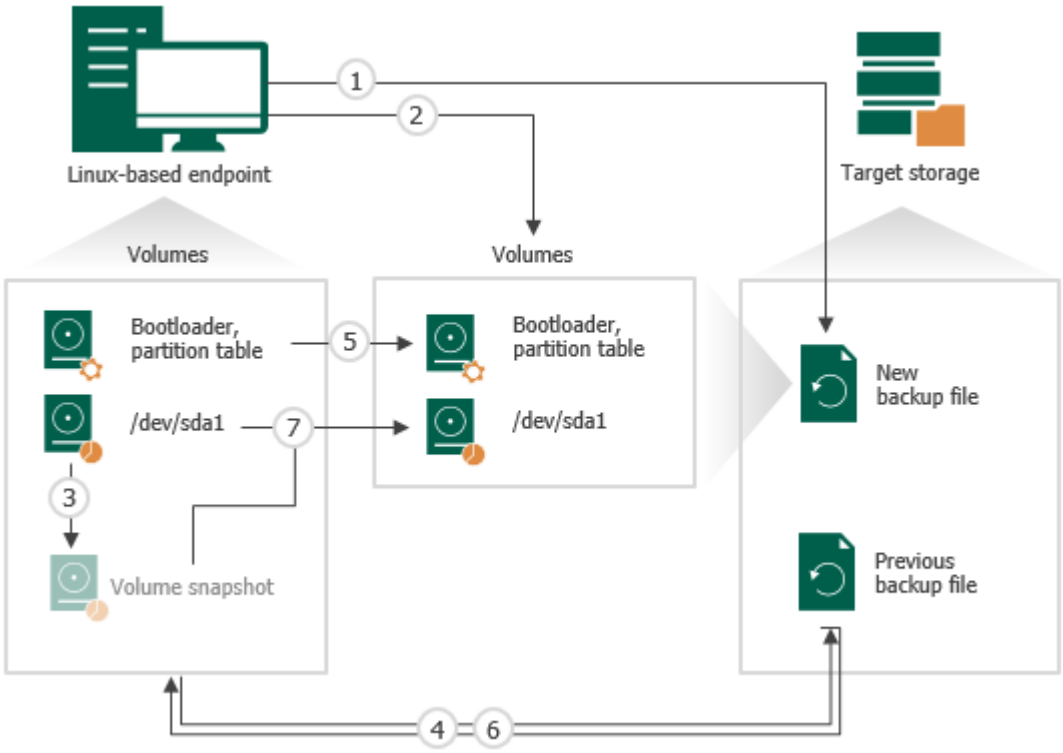
   Mind that every time the driver is unloaded or the Veeam Agent computer is rebooted, the changed block tracking map is reset as well. In such case, to detect what data blocks have changed since the previous job session, Veeam Agent rescans the entire data added to the backup scope and creates a new changed block tracking map. In this case, backup requires greater time.

   To learn about full and incremental backup, see Backup Chain.

5. Veeam Agent copies the partition table and bootloader located on the hard disk to the backup file in the target location.
6. [For incremental backup] Veeam Agent and calculates checksums for each data block and compares them with checksums from the backup file created during the previous job session. If checksums do not match, Veeam Agent will copy the data block to the target location during the next backup process step.

7. Veeam Agent copies data from the following sources:
   - Data that did not change on disk during backup is transferred from the source volume.
   - Data that changed on disk during backup is transferred from the snapshot.

   After all the data is transferred, Veeam Agent removes the snapshot.



How File-Level Backup Works

During file-level backup, Veeam Agent performs the following operations for every backup job session:

1. When a new job session starts, Veeam Agent creates a backup file in the target location.
2. In the backup file, Veeam Agent creates a disk. The disk contains a volume with the ext4 file system.
3. Veeam Agent creates a snapshot of the volume which data you want to back up. The snapshot is created on the volume that has enough free disk space to contain the snapshot data. To create a snapshot, Veeam Agent uses the *Veeam Agent for Linux Driver*.

   The snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup. If a data block is about to change on disk during backup, Veeam Agent will copy this block to the snapshot. After the data block is overwritten on the source location, its original copy will remain intact in the snapshot.

   > **TIP**
   >
   > Mind the following:
   >
   > - You can also set up Veeam Agent to create a file-level backup in the snapshot-less mode. This mode allows you to back up data that resides in any file system mounted to the root file system of the Veeam Agent computer. However, Veeam Agent does not track whether source files have changed since the backup process start. To learn more, see Snapshot-Less File-Level Backup.
   > - Compared to the volume-level backup, the file-level backup, Veeam Agent does not provide changed block tracking mechanism and does not split source files into data blocks. As a result, if you plan to back up a significant amount of data, the file-level backup will require greater time, and created backup files will have greater size.
   >
   >   For example, you have a 1 GB file, and since the previous backup session only one data block of this file has changed. In case of the file-level backup, Veeam Agent will send the whole 1 GB file to the target again.

4. [For incremental backup] To detect files that changed on the Veeam Agent computer since the previous backup session, Veeam Agent reads file metadata and compares last modification time of files in the original location and files in the backup created during the previous job session. If the file has modification time later than the previous job session start time, Veeam Agent considers the file as changed.

   To learn about full and incremental backup, see Backup Chain.

5. [For incremental backup] Veeam Agent calculates checksums for each data block and compares them with checksums from the backup file created during the previous job session. If checksums do not match, Veeam Agent will copy the data block to the target location during the next backup process step.

6. Veeam Agent copies data that you selected for backup to the target location. As part of this process, Veeam Agent performs the following operations:

   a. Enumerates all files in the source location.

   b. For each enumerated file, creates a target file in the volume inside the backup file.

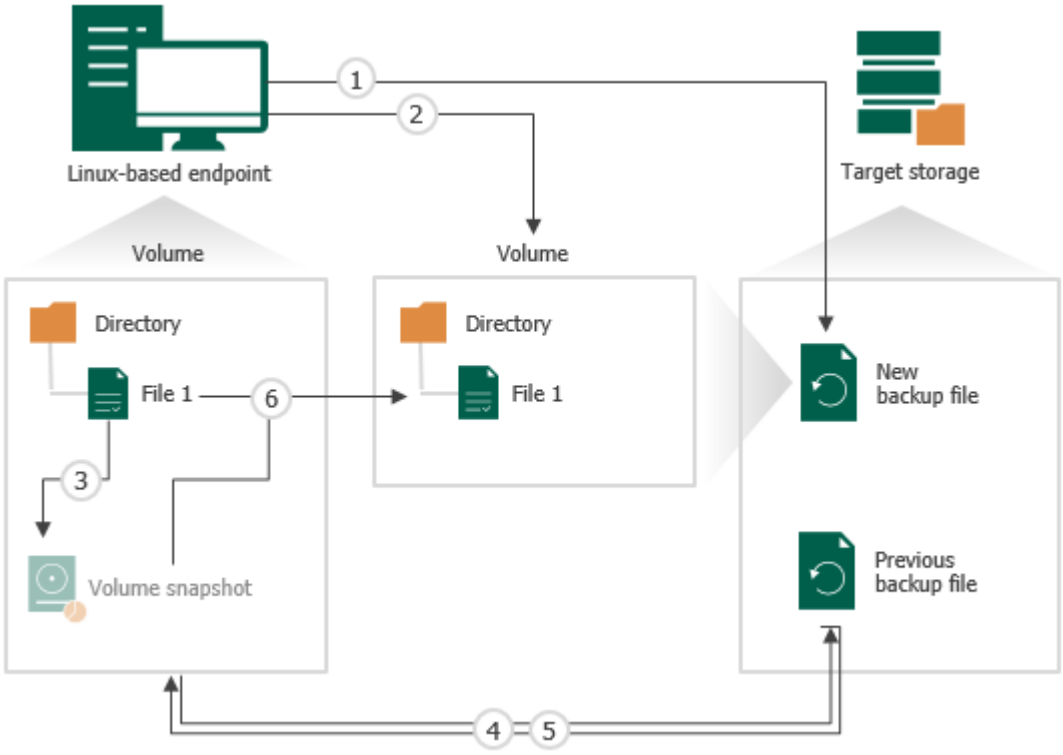
   c. Opens the source and the target files.

   d. Copies file data to the target location from the following sources:

      - Data blocks that did not change on disk during backup are transferred from the source volume.
      - Data blocks that changed on disk during backup are transferred from the snapshot.

   e. Closes the source and target files.

   After all backed-up files and directories are transferred, Veeam Agent removes the snapshot.



Backup Job

To back up your data, you must configure a backup job. The backup job settings define what data you want to back up, what the target location and retention policy for created backups are and how to back up your data. If necessary, you can re-configure the backup job and change its settings at any time.

> **NOTE**
>
> You cannot change the backup job type from volume-level to file-level, and vice versa.

In Veeam Agent for Linux, you can configure several backup jobs with different settings. For example, you can configure one backup job to create volume-level backup and another backup job to create file-level backup. You can configure backup jobs targeted at different backup locations to keep several copies of your backed-up data. You can also configure several backup jobs with individual schedule to fine-tune automatic backup creation process.

> **NOTE**
>
> You can create more than one backup job only if Veeam Agent operates in the Workstation or Server edition. To learn more, see Product Editions.

Veeam Agent launches the backup job according to the schedule you define. You can schedule the job to start at specific time daily or on specific week days. You can also start a backup job manually to perform backup on demand when needed.

Backup job scheduling settings are configured globally for all accounts of the Linux OS. For every backup job, Veeam Agent creates a record in its database and in the `crontab` configuration file of the root account. As a result, Veeam Agent can start a backup job automatically regardless of the currently running user session.

You can define schedule for a job in Veeam Agent or edit the `crontab` file directly to fine-tune the schedule. To learn more, refer to the Cron job scheduler documentation.

Backup Job Scripts

You can instruct Veeam Agent for Linux to run custom scripts within the backup job session:

- Pre-job and post-job scripts — Veeam Agent runs these scripts before the backup job starts and after the backup job completes. You can use pre-job and post-job scripts, for example, to configure email notifications about jobs performed by Veeam Agent.
- Pre-freeze and post-thaw scripts (in the Server edition only) — Veeam Agent runs these scripts before and after creating a snapshot. For example, the pre-freeze script may quiesce the file system and application data to bring the Linux OS to a consistent state before Veeam Agent creates a snapshot. After the snapshot is created, the post-thaw script may bring the file system and applications to their initial state.

Script settings are enabled at the job level. If Veeam Agent operates in the Server edition and you want to configure multiple backup jobs, you can specify individual scripts for each job.

Scripts must be created beforehand. You must specify paths to them in the job settings. Veeam Agent supports scripts in the SH file format.

Pre-Job and Post-Job Scripts

You can instruct Veeam Agent for Linux to run custom pre-job and post-job scripts. Veeam Agent executes the pre-job script directly before the backup job starts. After the backup job completes, Veeam Agent executes the post-job script.

Veeam Agent starts the backup job regardless of the pre-job script result. If the pre-job script fails to execute, Veeam Agent will always start the backup job. Then, after the backup job completes, Veeam Agent will execute the post-job script.

The script is considered to be executed successfully if "0" is returned.

The default time period for script execution is 10 minutes. After this period expires, Veeam Agent stops executing the script and displays a warning message in the job session. If the script fails to execute before the timeout expires, Veeam Agent does not display warning messages in the job session.

Pre-Freeze and Post-Thaw Scripts

You can instruct Veeam Agent for Linux to run custom pre-freeze and post-thaw scripts. Veeam Agent executes the pre-freeze script before creating a snapshot. After the snapshot is created, Veeam Agent executes the post-thaw script.

The script is considered to be executed successfully if "0" is returned.

The default time period for script execution is 10 minutes. After this period expires, Veeam Agent stops executing the script.

By default, if the pre-freeze or post-thaw script fails to execute, Veeam Agent does not start the backup job. However, you can instruct Veeam Agent to ignore errors that occur during the script execution process. To allow Veeam Agent to start backup jobs regardless of the script execution result, in the `/etc/veeam/veeam.ini` configuration file, uncomment the `ignoreFreezeThawFailures` parameter and set its value to `true`.

If Veeam Agent is set up to ignore script errors, and the pre-freeze or post-thaw script fails to execute, Veeam Agent will start the backup job. After the job successfully completes, Veeam Agent will display the *Warning* status for the job session.

NOTE

You can specify pre-freeze and post-thaw scripts only if Veeam Agent for Linux operates in the Server edition. If these scripts were enabled for the job while Veeam Agent operated in the Server edition, and then Veeam Agent has switched to another edition (for example, to the Free edition after the license has expired), the backup job will fail. You will need to delete the existing job and create a new backup job without pre-freeze and post-thaw scripts enabled.

File System Indexing

You can instruct Veeam Agent for Linux to create an index of files and directories located on the Veeam Agent computer during backup. File indexing allows you to search for specific files inside Veeam Agent backups and perform 1-click restore in Veeam Backup Enterprise Manager.

File indexing is enabled at the job level. You can specify granular indexing settings for each job.

IMPORTANT

Indexing mechanism does not recognize file exclusion masks. If you specify masks to exclude certain files in a file-level backup job, Veeam Agent for Linux will nevertheless index all files located in the directories that have been selected for backup.

For example, you have included the `/home` directory into the backup and specified the `*.pdf` exclusion mask. The *Index everything* option is enabled for the backup job. In this case, when you browse the resulting backup in Veeam Backup Enterprise Manager, PDF files will be displayed in the `/home` directory as if they were backed up.

Requirements for File System Indexing

- Veeam Agent for Linux must have either Workstation or Server license installed.
- The following utilities must be installed on the computer: `mlocate`, `gzip` and `tar` (standard utilities for majority of Linux distributions). These utilities are provided along with the product in the product installation media.

NOTE

Consider the following:

- 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 Restore section in the Veeam Backup Enterprise Manager User Guide.
- If SELinux is enabled in the Linux OS, file system indexing may fail.

Backup Repository

A backup job configured in Veeam Agent for Linux creates backup files in a backup repository. A backup repository is a directory on the storage where you want to keep backup files. You can use the following types of disk-based storage to create a backup repository:

- Local (internal) storage of the protected machine (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share.
- Veeam Backup & Replication 11 or later backup repository.
- Veeam Backup & Replication 11 or later cloud repository.

> **IMPORTANT**
>
> A backup repository must be created on a separate volume from a volume whose data you plan to back up.

Veeam Agent for Linux works with backup storage differently depending on the way you configure and start backup jobs — with the Veeam Agent control panel or command line interface.

Backup Location and Control Panel

If you use the Veeam Agent control panel to perform backup tasks, you do not have to deal with backup repositories. When you specify a target location for backup in the Backup Job wizard, Veeam Agent configures the backup repository automatically. Veeam Agent saves path to the specified backup location, assigns to this location a unique name and ID and saves this information in the database. The information is used by Veeam Agent and is not displayed in the control panel.

If you target a backup job at the network shared folder, every time the backup job starts, Veeam Agent will automatically mount the shared folder to the `/tmp/veeam` directory in the computer file system and create a backup file in this directory. After the backup job completes, Veeam Agent will automatically unmount the network shared folder.

You can target several backup jobs to individual backup locations or use the same target location for several backup jobs. This may be useful if you want to back up different types of data to separate locations or to keep all backed-up data at one place.

Backup Repository and Command Line Interface

If you work with Veeam Agent for Linux using the command line interface, you must deal with backup repositories depending on the target location selected for the backup job.

If you target a backup job at a local directory or network shared folder, you must create a repository before you configure a backup job:

- In case of a local directory, you specify a name for the repository and a local directory in which Veeam Agent will create backup files. To learn more, see Creating Repository in Local Directory.
- In case of a network shared folder, you specify a name for the repository, a path to the network shared folder in which Veeam Agent will create backup files, a type of the network shared folder and additional mounting options.

  Every time the backup job starts, Veeam Agent will automatically mount the shared folder to the `/tmp/veeam` directory in the computer file system and create a backup file in this directory. After the backup job completes, Veeam Agent will automatically unmount the network shared folder. To learn more, see Creating Repository in NFS Share and Creating Repository in SMB Share.

  If the directory to which the shared folder should be mounted resides on the backed-up volume, the backup job may fail.

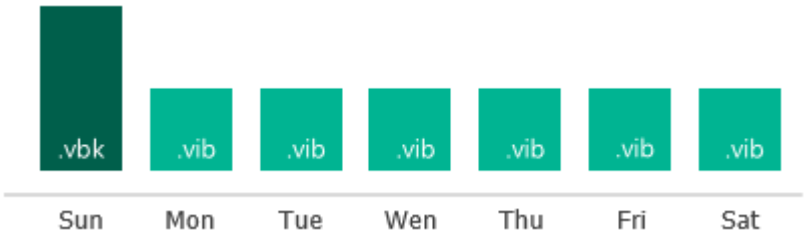
If you target a backup job at a Veeam backup repository or cloud repository, you do not need to create repositories. Before configuring the backup job, you must connect to the Veeam backup server or Veeam Cloud Connect service provider. To learn more, see Connecting to Veeam Backup Server and Connecting to Service Provider.

You can configure several backup repositories and target different backup jobs at these repositories. This may be useful if you want to back up different types of data to separate locations or to keep several copies of your backed-up data.

Backup Chain

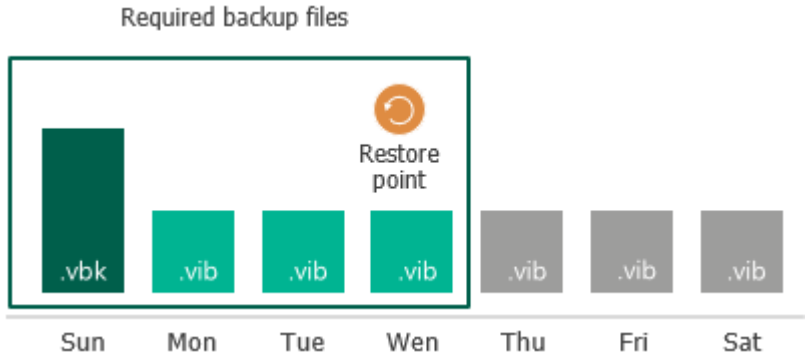
Every backup job session produces a new backup file in the target location. Backup files make up a backup chain. The backup chain can contain files of two types: full backups and incremental backups.

- During the first backup job session, Veeam Agent for Linux performs full backup. Veeam Agent for Linux copies all data that you have chosen to back up (entire volumes and directories) and stores the resulting full backup file (VBK) in the target location. The full backup takes significant time to complete and produces a large backup file: you have to copy the whole amount of data.
- During subsequent backup job sessions, Veeam Agent for Linux performs incremental backups. It copies only new or changed data relatively to the last backup job session and saves this data as an incremental backup file (VIB) in the target location. Incremental backups typically take less time than full backup: you have to copy only changes, not the whole amount of data.



After several backup cycles, you have a chain of backup files in the target location: the first full backup file and subsequent incremental backup files. Every backup file contains a restore point for backed up data. A restore point is a "snapshot" of your data at a specific point in time. You can use restore points to roll back your data to the necessary state.

To recover data to a specific restore point, you need a chain of backup files: a full backup file plus a set of incremental backup files following this full backup file. If some file from the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, it is recommended that you do not delete separate backup files manually. To learn more, see Deleting Backups.



Types of Backup Files

Veeam Agent for Linux produces backup files of the following types:

- VBK — full backup file.
- VIB — incremental backup file.
- VBM — backup metadata file. The backup metadata file is updated with every backup job session. It contains information about the computer on which the backup was created, every restore point in the backup chain, how restore points are linked to each other and so on. The backup metadata file is required for performing file-level and volume-level restore operations.

Backup Retention Policy

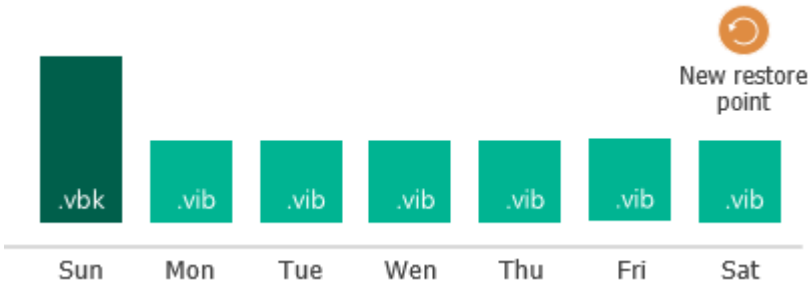
Restore points in the backup chain are not kept forever. They are removed according to the retention policy. The retention policy helps maintain the life cycle of restore points and make sure that backup files do not consume the whole disk space.

Veeam Agent for Linux retains the number of latest restore points defined by the user. During every backup job session, Veeam Agent for Linux checks if there is any obsolete restore point in the backup chain. If some restore point is obsolete, it is removed from the chain.

Removing Backups by Retention

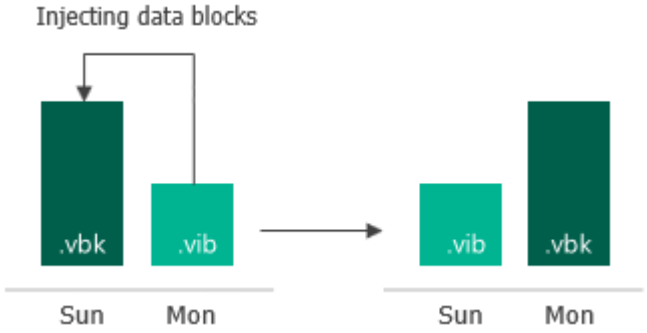
When removing obsolete restore points, Veeam Agent for Linux does not simply delete backup files from disk. It transforms the backup chain so that the backup chain always contains a full backup file on which subsequent incremental backup files are dependent. To maintain the consistency of the backup chain, Veeam Agent for Linux uses the following rotation scheme:

1. During every backup job session Veeam Agent for Linux adds a backup file to the backup chain and checks if there is an obsolete restore point.

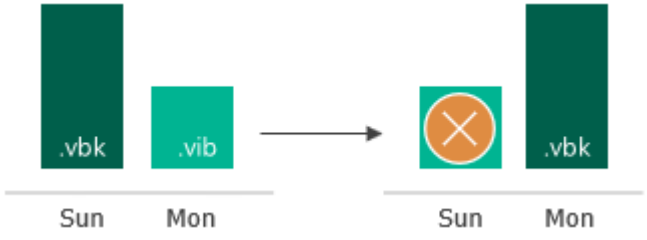


2. If an obsolete restore point exists, Veeam Agent for Linux transforms the backup chain. As part of this process, it performs the following operations:

   a. Veeam Agent for Linux re-builds the full backup file to include in it data of the incremental backup file that follows the full backup file. To do this, Veeam Agent for Linux injects into the full backup file data blocks from the earliest incremental backup file in the chain. This way, a full backup 'moves' forward in the backup chain.



   b. The earliest incremental backup file is removed from the chain as redundant: its data has already been injected into the full backup file, and the full backup file includes data of this incremental backup file.

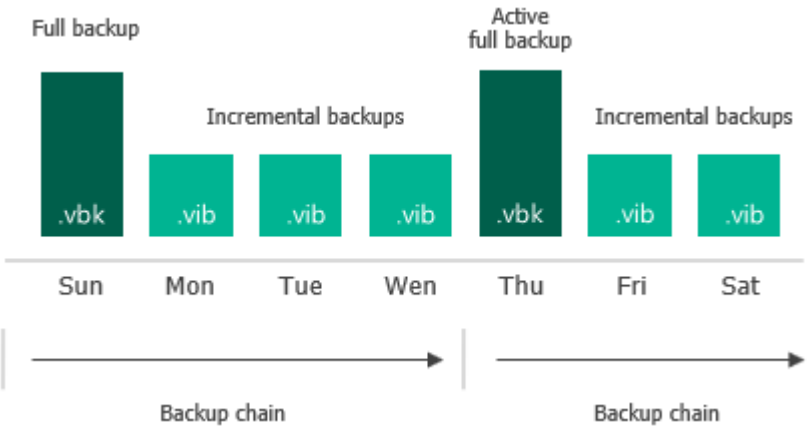


Active Full Backup

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 Agent for Linux lets you create active full backups.

When Veeam Agent for Linux performs active full backup, it produces a full backup file and adds this file to the backup chain.

The active full backup resets the backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file and its subsequent incremental backup files remain on the disk. After the last incremental backup file created prior to the active full backup becomes outdated, Veeam Agent for Linux automatically deletes the previous backup chain. To learn more, see Retention Policy for Active Full Backups.



You can create active full backups manually or schedule a backup job to create active full backups periodically. To do this, you can use the Veeam Agent for Linux control panel or command line interface.

- To learn how to configure active full backup schedule and create active full backups with the Veeam Agent for Linux control panel, see Active Full Backup Settings and Starting Backup Job from Control Panel.
- To learn how to configure active full backup schedule and create active full backups with the Veeam Agent for Linux command line interface, see Configuring Active Full Backup Schedule and Creating Active Full Backups.

Active Full Backup Schedule

You can schedule a backup job to create active full backups periodically.

Active full backup schedule depends on the regular backup schedule. For every backup job with enabled backup schedule, Veeam Agent saves active full schedule settings to the `crontab` configuration file of the root account.

- In case active full backup is scheduled on a week day, Veeam Agent modifies the `crontab` record associated with the regular schedule.

  For example, the regular backup schedule is set to Monday and Tuesday at 15:00. Active full backup schedule is set to Friday. In this case, the `crontab` record for the regular backup schedule will contain information that the job must start on Monday, Tuesday and Friday at 15:00.

- In case active full backup is scheduled on a day of the month, Veeam Agent creates a separate `crontab` record that instructs the backup job to run on this day at the same time as it must run upon the regular schedule.

Keep in mind that if the job is not scheduled to run automatically, Veeam Agent will not create a crontab record. Veeam Agent will save schedule settings for active full backup in its database only. As a result, active full backup will not run automatically.

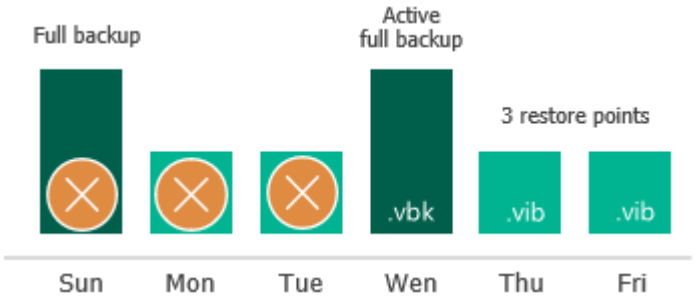
For information about how to configure job schedule, see Configuring Backup Schedule and Configuring Active Full Backup Schedule.

Retention Policy for Active Full Backups

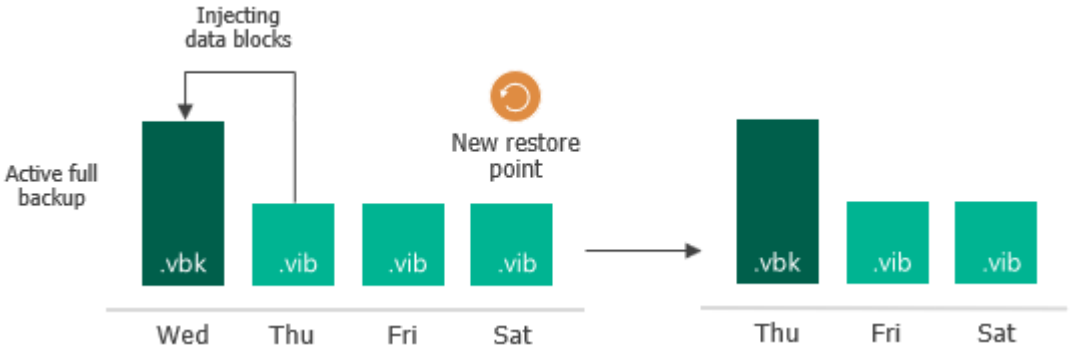
To be able to restore data from a Veeam Agent backup, you need to have a full backup file and a chain of subsequent incremental backup files on the disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you create an active full backup, in some days there will be more restore points on the disk than specified by retention policy settings. Veeam Agent for Linux 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. A full backup file is created on Sunday, incremental backup files are created on Monday and Tuesday, and an active full backup is created on Wednesday. Although the backup chain now contains 4 restore points, Veeam Agent for Linux will not delete the previous backup chain. Veeam Agent for Linux will wait for the next 2 incremental backup files to be created, and only then will delete the whole previous chain, which will happen on Friday. As a result, although the retention policy is set to 3 restore points, the actual number of backup files on the disk will be greater for some time.



Veeam Agent for Linux treats the active full backup in the same way as a regular full backup. If some restore point becomes obsolete, Veeam Agent for Linux will re-build the full backup file to include in it data of the incremental backup file that follows the full backup file. After that, Veeam Agent for Linux will remove the earliest incremental backup file from the chain as redundant.



Backup of Database Systems

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run one of the following database systems:

- Oracle database system
- MySQL database system
- PostgreSQL database system

> **IMPORTANT**
>
> Consider the following:
>
> - You can specify settings for database system processing only if Veeam Agent for Linux operates in the Server edition.
> - Veeam Agent supports processing of database systems for the volume-level backup only.
> - Veeam Agent does not support processing of multiple database systems on one Veeam Agent computer.
> - Veeam Agent does not support 32-bit database systems installed on a 64-bit Linux OS.

To process database systems with Veeam Agent for Linux, enable application-aware processing for the backup job:

- If you work with Veeam Agent using the Veeam Agent control panel, configure application-specific settings at the **Advanced** step of the Backup Job wizard. To learn more, see Specify Advanced Backup Settings.
- If you work with Veeam Agent using the command line interface, create the backup job, then specify application-specific settings for this job. To learn more, see Creating Volume-Level Backup Job and Configuring Database Processing Settings.

Oracle Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the Oracle database system.

> **NOTE**
>
> You can use Veeam Explorer for Oracle to restore Oracle databases from a Veeam Agent for Linux backup. For information about item-level recovery of Oracle systems, see Veeam Explorers User Guide.

How Oracle Processing Works

To ensure that the backed-up data is in the consistent state, Veeam Agent for Linux performs the Oracle database system processing. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent obtains information about Oracle databases that run on the Veeam Agent machine.
2. Veeam Agent connects to the Oracle database and operates depending on the database state:
   - Shutdown state
   - ARCHIVELOG mode
   - NOARCHIVELOG mode

After Veeam Agent for Linux finishes database system processing, Veeam Agent proceeds to the next step of the backup process. To learn more, see How Backup Works.

Processing of Database in Shutdown State

If the database is shut down, Veeam Agent skips it and connects to the next database. The skipped database will be included in the backup. To restore such database, you must restore the entire volume that contains the database. You cannot restore the database as an independent item using Veeam Explorer for Oracle. To learn more, see Veeam Explorers User Guide.

Veeam Agent displays a warning message about the database that is shut down in the job session logs. The backup job does not fail.

Processing of Database in ARCHIVELOG Mode

If the Oracle database is running in the ARCHIVELOG mode, the Oracle database system keeps archived logs that allow to recover all committed transactions of the database. To learn more, see Oracle documentation.

If the database operates in the ARCHIVELOG mode, Veeam Agent performs the following operations:

1. Veeam Agent switches the database to the backup mode. Veeam Agent changes the database state using the Oracle functionality.
2. Veeam Agent creates a snapshot of the volume.
3. Veeam Agent returns the database to the initial state.

Processing of Database in NOARCHIVELOG Mode

If the Oracle database is running in the NOARCHIVELOG mode, the Oracle database system does not create archived logs. Logs that are created before the database is switched to NOARCHIVELOG remain untouched. In this mode, you can only restore the database to the point of full database backup. You cannot recover transactions subsequent to that full database backup.

If the database operates in the NOARCHIVELOG mode, Veeam Agent performs the following operations:

1. Veeam Agent shuts down the database. Veeam Agent changes the database states using the Oracle functionality.
2. Veeam Agent creates a snapshot of the volume.
3. Veeam Agent returns the database to the initial state.

Authentication Methods

Veeam Agent for Linux can connect to the Oracle database system and perform Oracle archived logs backup and/or deletion using one of the following account types:

- *System account* — Veeam Agent uses the account of the Veeam Agent machine OS. To connect to the Oracle database system, the account must be a member of the group that owns configuration files for the Oracle database (for example, the oinstall group).
- *Oracle account* — Veeam Agent uses the Oracle account. To connect to the Oracle database system, the account must have SYSDBA rights.

Archived Log Processing

> **NOTE**
>
> Veeam Agent for Linux operating in the standalone mode supports the backup of archived logs, but does not support a separate schedule and retention policy for the backup of archived logs. This functionality is available only for Veeam Agent for Linux operating in the managed mode within the Veeam Agent management scenario. For more information, see Oracle Archived Log Settings section in the Veeam Agent Management Guide.

In the ARCHIVELOG mode, the Oracle database system stores database archived logs to a certain location on the machine that runs the database system, as specified by the database administrator. Veeam Agent allows you to set up the following ways of archived logs processing:

- *Delete logs older than the specified time (in hours)*. After the backup job completes, Veeam Agent deletes archived logs that are older than the specified time from the Veeam Agent machine. This helps make sure that logs do not overflow the storage space on the processed machine.
- *Delete oldest logs larger than the specified size (in GB)*. After the backup job completes, Veeam Agent checks whether the total size of archived logs exceeds the specified size. After that, Veeam Agent deletes oldest archived logs that exceed the specified size from the processed machine. This helps make sure that logs do not overflow the storage space on the Veeam Agent machine.
- *Do not delete archived logs*. Log files remain untouched on the Veeam Agent machine.

Veeam Agent processes archive logs via Oracle Call Interface (OCI).

Requirements and Limitations for Oracle Processing

- Oracle Database versions 11g – 21c are supported for all operating systems supported by Veeam Agent for Linux. To learn more, see System Requirements.
- Automatic Storage Management (ASM) is not supported.
- Oracle Real Application Clusters (RAC) are not supported.
- Oracle Grid Infrastructure is not supported.
- Oracle Database Express Edition is not supported.
- SAP on Oracle is not supported.
- Oracle Database architectures with Data Guard are not supported.

MySQL Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the MySQL database system.

How MySQL Processing Works

To ensure that the backed-up data is in the consistent state, Veeam Agent for Linux performs the MySQL database system processing. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent connects to the MySQL database system and obtains the list of tables.
2. Veeam Agent locks base tables that are based on the MyISAM storage engine. Veeam Agent changes the table state using the MySQL functionality. Tables based on the InnoDB storage engine do not require locking.

   Keep in mind that Veeam Agent supports processing of tables based on the MyISAM and InnoDB storage engines only. Veeam Agent does not support tables that use other storage engines.
3. Creates a snapshot of the volume.
4. Unlocks tables locked at the step 2.

After Veeam Agent unlocks tables, Veeam Agent proceeds to the next step of the backup process. To learn more, see How Backup Works.

Authentication Methods

> **IMPORTANT**
>
> To process the MySQL database system, the MySQL account must have the following privileges:
>
> - SELECT for all tables. If the MySQL account does not have the SELECT privilege for the table, Veeam Agent cannot access table metadata. As a result, Veeam Agent does not process the table. To learn more, see MySQL documentation.
> - LOCK TABLES. If the MySQL account does not have the LOCK TABLES privilege, Veeam Agent does not process tables based on the MyISAM storage engine.
>
> To obtain information about privileges that are assigned to the account, use MySQL functionality, for example, the SHOW GRANTS statement. To learn more, see MySQL documentation.

Veeam Agent for Linux can connect to the MySQL database system using one of the following methods:

- *Password* — Veeam Agent uses the MySQL account credentials that you specify in the backup job settings.
- *Password file* — Veeam Agent uses the MySQL account credentials that are stored in the `.my.cnf` password file. To learn more about password file configuration, see Preparing Password File for MySQL Processing.

Requirements and Limitations for MySQL Processing

- Veeam Agent for Linux supports processing of MySQL database system version 5.6 – 8.0.
- Configurations with multiple MySQL installations and/or instances on the same machine are not supported.
- MySQL Cluster versions are not supported.

PostgreSQL Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the PostgreSQL database system.

How PostgreSQL Processing Works

To ensure that the backed-up data is in the consistent state, Veeam Agent performs the PostreSQL database processing. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent forces a write-ahead log checkpoint. This command fastens the database system restore. To learn more, see PostgreSQL documentation.
2. Creates a snapshot of the volume and proceeds to the next step of the backup process. To learn more, see How Backup Works.

Authentication Methods

Veeam Agent for Linux can connect to the PostgreSQL database system using one of the following methods:

- *Database user with password* — Veeam Agent uses the PostgreSQL account credentials that you specify in the backup job settings.
- *Database user with password file* — Veeam Agent the PostgreSQL database system to use account credentials that are stored in the `.pgpass` password file. To learn more about password file configuration, see Preparing Password File for PostgreSQL Processing.
- *System user without password* — Veeam Agent uses the peer authentication. In the peer authentication method, Veeam Agent for Linux uses the OS account as the PostgreSQL database user name.

Requirements and Limitations for PostgreSQL Processing

- Veeam Agent for Linux supports processing of the PostgreSQL database system version 9.5 – 14.0.
- Configurations with multiple PostgreSQL installations and/or instances on the same server are not supported.

Data Restore

Veeam Agent for Linux offers two data restore scenarios:

- You can perform volume-level restore to recover the entire system image of your computer or specific computer volumes. To learn more, see Volume-Level Restore.
- You can perform file-level restore to recover individual files and directories. To learn more, see File-Level Restore.

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.

Keep in mind that you cannot browse the volume in the backup and select individual files and directories for restore. For granular file-level restore, you can use the File-Level Restore option.

A volume can be restored to its original location or new location. If you restore the volume to its original location, Veeam Agent for Linux overwrites data on the original volume. If you restore the volume to a new location, and the target disk contains any data, Veeam Agent for Linux overwrites data in the target location with data retrieved from the backup.

Limitations for Volume-Level Restore

Volume restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Linux swap space is hosted.
- You cannot restore a volume to the volume where the backup file used for restore is located.

To overcome the first two limitations, you can create a Veeam Recovery Media and use the **Volume Restore** wizard for volume-level restore. To learn more, see Veeam Recovery Media.

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.
    - 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.
    - 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 files and directories in the mounted backup and copy them to their initial location or to a new location.

Veeam Recovery Media

Veeam Agent for Linux lets you use the Veeam Recovery Media — a recovery image of the Linux OS that provides an alternative way to boot your computer.

The recovery image includes a custom Linux OS with the limited functionality. It comprises Linux kernel and a set of GNU/Linux utilities necessary to boot the computer and perform basic administration tasks. If the OS installed on the computer fails to start for some reason, you can boot the recovery image OS. After booting, you can do the following:

- You can restore data from a backup to your computer. For this scenario, you must have a backup created with Veeam Agent for Linux.
- You can use Linux OS tools to diagnose problems and fix errors on your computer.

The recovery image can be helpful if one of the following errors occur:

- The OS on the computer fails to start.
- You want to perform bare metal restore from the backup on the computer without the OS and other software installed.
- You want to restore the system volume of the computer and so on.

Veeam Recovery Media is distributed as an ISO image. The ISO image file can be downloaded from this Veeam webpage. You can burn the ISO image file to different kinds of media:

- Removable storage devices such as USB drives or SD cards
- CD/DVD/BD

> **NOTE**
>
> Consider the following:
>
> - You can also download the Veeam Recovery Media ISO image from the Veeam software repository. In addition to the regular recovery image, the repository offers a legacy recovery image intended for Veeam Agent computers that run an earlier version of Linux kernel with custom drivers. To learn more, see Veeam Recovery Media Versions.
> - For information about how to burn the ISO image to a removable storage device, as well as workaround for accessing the Veeam recovery UI, see this Veeam KB article.

When you boot from the Veeam Recovery Media, you can use the recovery environment to fix the OS system errors on your computer or restore data from the backup. Veeam Agent for Linux offers a set of tools for the computer system image and data recovery:

- Restore volumes — the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
- Restore files — the File Level Restore wizard to restore files and folders to the original location or to a new location.
- Exit to shell — Linux shell prompt with standard utilities to diagnose problems and fix errors.

Veeam Recovery Media Versions

Veeam Agent for Linux offers the following versions of Veeam Recovery Media:

- Regular recovery image suitable for the majority of Linux machines.
- Legacy recovery image that allows you to perform bare metal recovery of Linux machines that run an earlier version of Linux kernel with custom drivers.

Regular Recovery Image

The regular recovery image is intended for Veeam Agent computers that run Linux kernel version 3.16 and later. This version of Veeam Recovery Media is based on the Debian live install image and allows you to boot your computer using Linux kernel of either stable Debian release or testing release. In Veeam Agent for Linux 5.0, this Veeam Recovery Media offers Linux kernel versions 4.19.0 and 5.9.0.

You can download the regular recovery image from the following sources:

- Veeam website

  Recovery image ISO files downloaded from the Veeam website have the following names:

  - `veeam-recovery-media-5.0.2.4567_i386.iso` — for Veeam Agent computers based on the x86 architecture
  - `veeam-recovery-media-5.0.2.4567_x86_64.iso` — for Veeam Agent computers based on the x64 architecture

- Veeam software repository

  Recovery image ISO files downloaded from the Veeam software repository have the following names:

  - `veeam-recovery-i386-5.0.0.iso` — for Veeam Agent computers based on the x86 architecture
  - `veeam-recovery-amd64-5.0.0.iso` — for Veeam Agent computers based on the x64 architecture

The size of the regular recovery image file depends on the Veeam Agent computer architecture: 460 MB for x86 computers and 490 MB for x64 computers.

In the scenario where you instruct Veeam Agent to download and patch the recovery image during initial product setup on a computer that runs Linux kernel version 3.16 and later, Veeam Agent downloads the regular recovery image from the Veeam software repository.

Consider the following:

- The regular recovery image supports EFI.
- The regular recovery image supports Secure Boot.
- The regular recovery image contains a larger set of software packages than the legacy recovery image. You can also install additional software from a software repository using the `apt` package manager. Keep in mind that to install software from a repository, the computer booted from the Veeam Recovery Media must have internet connection.

  By default, the list of repositories contains the following package sources: DebianStable, DebianBackports, StableUpdates, DebianSecurity. You can also add custom repositories if necessary. For information about how to do this, see Debian documentation.

Legacy Recovery Image

The legacy recovery image is intended for Veeam Agent computers that run Linux kernel version earlier than 3.16 with custom drivers. It allows you to create a custom Veeam Recovery Media for such computers. In the scenario where you instruct Veeam Agent to download and patch the recovery image during initial product setup on such a computer, Veeam Agent downloads this version of the recovery image from the Veeam software repository. You can also download it manually from the repository.

The name and size of the legacy recovery image ISO file is `veeam-recovery-media-5.0.0.iso`. The size of the ISO file depends on the Veeam Agent computer architecture: 136 MB for x86 computers and 139 MB for x64 computers.

The legacy recovery image does not support EFI by default. You can enable EFI support for a custom Veeam Recovery Media based on this version of recovery image. To learn more, see Create Custom Veeam Recovery Media and Creating Custom Veeam Recovery Media.

Drivers in Veeam Recovery Media

The generic Veeam Recovery Media available for download from the Veeam website or Veeam software repository contains the following data:

1. Set of files required to start the recovery image OS from the recovery media.
2. Set of Veeam tools for the computer system image and data recovery.
3. Set of Linux command line tools to diagnose problems and fix errors on your computer. For the regular recovery image, in addition to the standard set of tools, you can install custom software from a software repository.
4. Drivers required to run hardware and devices on your computer in a regular way. The regular recovery image contains drivers included in the Linux kernel versions 4.19.0 and 5.9.0. The legacy recovery image contains drivers included in the Linux kernel version 5.1.15.

   When you boot your computer from the Veeam Recovery Media, drivers from the Veeam Recovery Media are automatically loaded on the recovery image OS.

   If your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media, you can create a custom recovery image. Veeam Agent will copy the Linux kernel running on your computer with its currently loaded modules and include them into the custom Veeam Recovery Media. To learn more, see Creating Custom Veeam Recovery Media.

Integration with Veeam Backup & Replication

If you plan to use Veeam Agent for Linux 5.0 with Veeam Backup & Replication, you must install Veeam Backup & Replication 11a on the Veeam backup server. Veeam Agent for Linux 4.0 supports integration with Veeam Backup & Replication 11 or later as well.

You can store backup files created with Veeam Agent for Linux on backup repositories managed by Veeam Backup & Replication. To do this, you must select a Veeam Backup & Replication backup repository as a target location in the properties of the backup job. To store Veeam Agent backups, you can use a simple backup repository or a scale-out backup repository.

> **NOTE**
>
> Consider the following:
>
> - The current guide covers subjects related to Veeam Agent for Linux operating in the standalone mode.
> - You can also use Veeam Backup & Replication to manage Veeam Agent for Linux on computers in your infrastructure. As part of the Veeam Agent management scenario, you can remotely deploy Veeam Agent to your computers, as well as configure and manage Veeam Agent backup jobs in Veeam Backup & Replication. To learn more, see Veeam Agent Management Guide.
> - If you create a backup job with the Veeam Agent command line interface, you need to specify a Veeam backup repository in the backup job settings. Veeam backup repository appears in the list of backup repositories after you connect to a Veeam backup server. To learn more, see Managing Veeam Backup & Replication Servers.

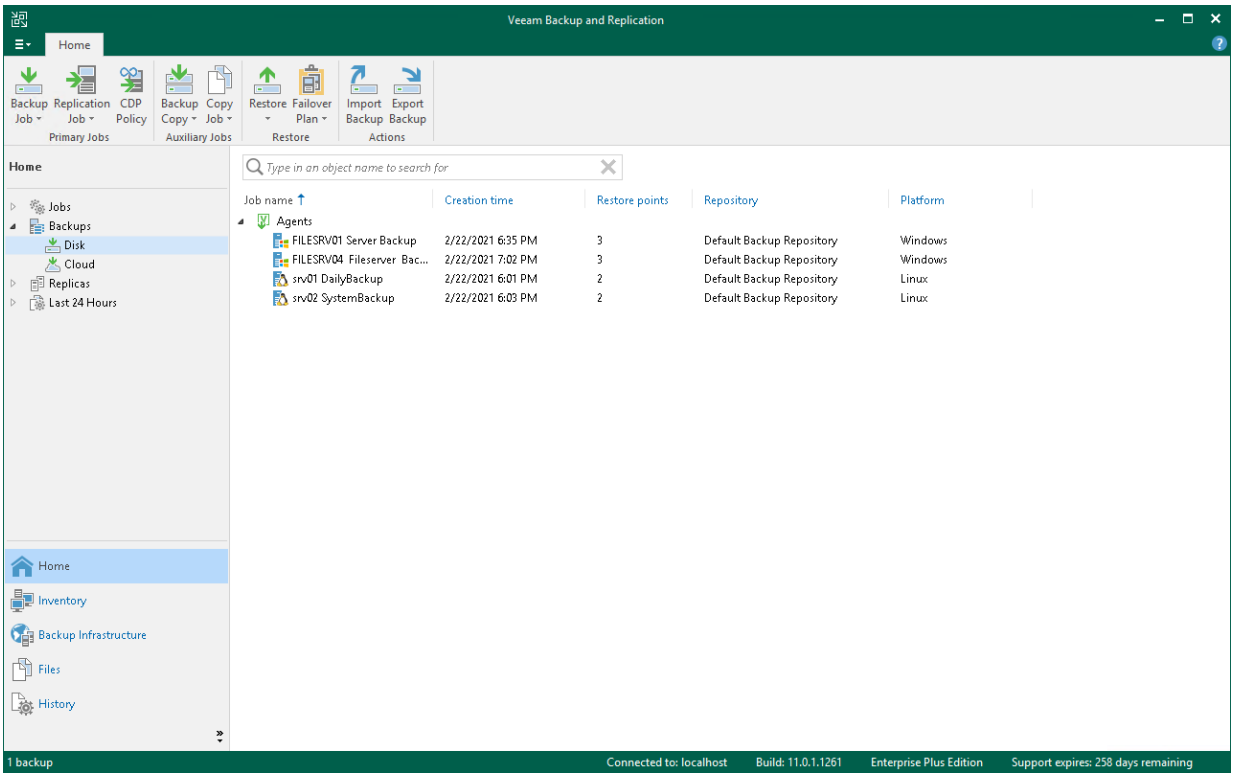
Veeam Agent for Linux works with the Veeam Backup & Replication backup repository as with any other backup repository. Backup files are stored to a separate folder; you can perform standard restore operations using these files.

Information about Veeam Agent backups stored on the Veeam Backup & Replication backup repositories, backup jobs and sessions becomes available in the Veeam Backup & Replication console:

- The Veeam Agent for Linux backup job is displayed in the list of jobs in Veeam Backup & Replication.
- Backup files created with Veeam Agent for Linux are displayed in the list of backups, under the **Agents** node.
- Performed job sessions are available in the **History** view of Veeam Backup & Replication.

Backup administrators working with Veeam Backup & Replication can perform a set of operations with Veeam Agent backups:

- Perform data protection operations: copy Veeam Agent backups to secondary backup repositories and archive these backups to tape.
- Perform restore operations: restore individual files and directories, application items from Veeam Agent backups; restore computer disks and convert them to the VMDK, VHD or VHDX format; restore to Microsoft Azure and Amazon EC2.
- Perform administrative tasks: disable and delete Veeam Agent backup jobs, remove Veeam Agent backups and so on.



Backup to Veeam Cloud Connect Repository

If you want to store your data in the cloud, you can connect to a Veeam Cloud Connect service provider (SP) and create Veeam Agent backups in a cloud repository. To do this, you must provide credentials of the tenant (or subtenant) account that you obtained from the SP, and select a cloud repository as a target for backup files in the properties of the backup job. To learn more, see Veeam Cloud Connect Repository Settings.

NOTE

Consider the following:

- You can create Veeam Agent backups in a cloud repository if the SP backup server runs Veeam Backup & Replication 11 or later.
- Backup to a cloud repository is available if Veeam Agent for Linux operates in the Workstation or Server edition.

Managing Veeam Agent in Veeam Backup & Replication

Veeam Backup & Replication lets you automate management of Veeam Agent on multiple computes in your infrastructure. You can deploy Veeam Agent for Linux, configure Veeam Agent backup jobs and perform other data protection and administration tasks on remote computers. To use the Veeam Agent management functionality in Veeam Backup & Replication, you must install Veeam Backup & Replication on the Veeam backup server.

To learn more, see Veeam Agent Management Guide.

Planning and Preparation

Before you install Veeam Agent for Linux, make sure that the target computer meets the system requirements and all required ports are open.

System Requirements

The protected Linux computer must meet requirements listed in the table below.

NOTE

The following system requirements apply to the following Veeam Agent for Linux configuration:

- Veeam Agent version is 5.0.2.

  To learn system requirements for Veeam Agent previous versions, see Release Notes for Veeam Agent for Linux 5.0 and Release Notes for Veeam Agent for Linux 5.0.1.
- Veeam Agent is operating in the standalone mode.

  To learn about system requirements for Veeam Agent managed by Veeam Backup & Replication, see the System Requirements section in the Veeam Agent Management Guide.
- Veeam Agent is installed with `veeam` and `veeamsnap` packages.

  To learn about system requirements for Veeam Agent installed using `veeam-nosnap` package, see Appendix A. Requirements for veeam-nosnap.

Specification	Requirement
Hardware	CPU: x86-64 processor (i386 or later). Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. Disk Space: 100 MB free disk space for product installation. Network: 10 Mbps or faster network connection to a backup target. System firmware: BIOS or UEFI. Disk layout: MBR or GPT. For virtual machines: Only full virtualization type is supported. Containers and paravirtualized instances are not supported. Oracle VM virtual machines are supported with limitations.

Specification	Requirement
OS	**Important!** Check considerations and limitations that apply to the list of supported OSes. Linux kernel version 2.6.32 or later is supported. Both 64-bit and 32-bit versions (if applicable) of the following distributions are supported: • Debian 9.0 – 11.1 • Ubuntu 14.04, 16.04, 18.04, 20.04, 20.10, 21.04, 21.10 • RHEL 6.0 – 8.5 • CentOS 7.0 – 8.5 • 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) • SLES 11 SP4, 12 SP2 – 15 SP3 • SLES for SAP 11 SP4, 12 SP2 – 15 SP3 • Fedora 33, 34, 35 • openSUSE Leap 15.2 – 15.3 • openSUSE Tumbleweed

Specification	Requirement
File System	**Important!** Check considerations and limitations that apply to the list of supported file systems. Veeam Agent for Linux supports consistent snapshot-based data backup for the following file systems: • Btrfs (for OSes that run Linux kernel 3.16 or later) • Ext 2/3/4 • F2FS • FAT16 • FAT32 • HFS • HFS+ • JFS • NILFS2 • NTFS • ReiserFS • XFS The supported file system (except for Btrfs) can reside on a simple volume or LVM2 volume; volumes protected with encryption software such as dm-crypt are supported. Btrfs is supported only if it resides directly on a physical device with no additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) below or above it. Data that resides on other file systems and volumes (including NFS and SMB shares) can be backed up using the snapshot-less mode. For details, see Snapshot-Less File-Level Backup.

Specification	Requirement
Software	**Important!** Check considerations and limitations that apply to the list of supported components. Protected computer must have the following components installed: • dkms • gcc • make • perl • linux-headers (for Debian-based systems) • kernel-headers (for RedHat-based systems) • kernel-devel (for RedHat-based systems) • kernel-uek-devel (for Oracle Linux with UEK) • libudev • libacl • libattr • lvm2 • libfuse • libncurses5 • dmidecode • libmysqlclient • libpq5 • python3 • efibootmgr (for UEFI-based systems) • isolinux (for Debian-based systems) • syslinux (for RedHat-based systems) • btrfs-progs (for backup of Btrfs file system) • mksquashfs (for custom Veeam Recovery Media) • unsquashfs (for custom Veeam Recovery Media) • wget (for custom Veeam Recovery Media) • xorriso (for custom Veeam Recovery Media with EFI support) For file system indexing, the following utilities are required: `mlocate`, `gzip` and `tar`. These utilities are provided along with the product in the product installation media.

Considerations and Limitations

OS

- Linux kernel version 2.6.32 or later is supported as long as you use kernels supplied by your distribution. Consider the following limitations:
  - Fedora and openSUSE Tumbleweed are supported up to kernel 5.15.7.
  - Linux kernel 2.6.32-754.6.3 in RHEL and Oracle Linux (RHCK) is not supported.
  - If you work with the Veeam Agent for Linux 5.0 build earlier than 5.0.1.4493 and the protected computer is running Linux kernel version 5.8 or later, only snapshot-less file-level backup is supported. For details, see Snapshot-Less File-Level Backup.
- Only GA versions of the supported distributions that have been released before the current version of Veeam Agent for Linux are supported.
- Use the `dkms` packages with the following distributions instead of the pre-built binary `veeamsnap` kernel module packages:
  - Debian 9.0 – 11.1
  - Ubuntu 14.04, 16.04, 18.04, 20.04, 21.04, 21.10
  - 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)
  - Fedora 33, 34, 35
  - openSUSE Tumbleweed
- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.
- For cloud-based installations that use customized kernels (such as Linux distributions deployed from AWS Marketplace or Azure Marketplace), the `veeamsnap` kernel module has experimental support status. For details about experimental support, see this Veeam KB article.
- Pre-built binary `veeamsnap` kernel module packages require kernel 2.6.32-131.0.15 or later for RHEL 6 (excluding 2.6.32-279.el6.i686) and 3.10.0-123 or later for CentOS / RHEL 7.0 – 7.9.
- RHEL, CentOS and Oracle Linux (RHCK) are supported up to certain kernel versions. For details, see this Veeam KB article.
- Ubuntu 21.10 minimal install is not supported.

File System

- Veeam Agent for Linux does not back up the following objects:
  - LVM snapshots
  - volumes that reside on USB devices and SD cards
- Total size of all file systems included in a file-level backup must not exceed 218 TB. Size of a file included in a file-level backup must not exceed 16 TB.

- Veeam Agent supports backup of extended attributes with the following limitations:
  - Veeam Agent backs up extended attributes only with the following public namespaces: system, security, trusted, and user.
  - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

    For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the ea_inodes feature. Backups created using the ea_inodes feature cannot be mounted on kernel versions up to 4.12.
- Backup of file and directory attributes (for example, a — append only, c — compressed, and so on) is not supported.
- Each volume included in a backup must have a unique UUID.
- The `veeamsnap` module provides RAM-based changed block tracking (CBT) mechanism. Every time the module is unloaded or Veeam Agent for Linux computer is rebooted, CBT data is reset. As a result, Veeam Agent reads the entire data added to the backup scope to detect what blocks have changed since the last job session, and incremental backup requires greater time.
- You cannot back up an entire system image or specific volumes of machines used as cluster nodes. Only snapshot-less file-level backup of cluster nodes is supported. That includes backup of machines that use shared disks, clustered file systems or clustered LVM.
- Certain limitations for EMC PowerPath configuration apply. To learn more, see this Veeam KB article.
- BFQ I/O scheduler is not supported.
- Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.

Software

> **IMPORTANT**
>
> Linux user account used to work with Veeam Agent for Linux must have the `/bin/bash` shell set as the default shell.

- The following packages are not required for CentOS, RHEL and SLES distributions if a pre-built binary `veeamsnap` package is to be installed.
  - dkms
  - gcc
  - make
  - perl
  - kernel-headers (for RedHat-based systems)
  - kernel-devel (for RedHat-based systems)

  For details, see Installing Veeam Agent for Linux.
- Version of the following packages varies according to the Linux kernel version that you use:
  - linux-headers (for Debian-based systems)
  - kernel-headers (for RedHat-based systems)

    - kernel-devel (for RedHat-based systems)
- For openSUSE and SLES distributions, either of the following packages is required: `libncurses5` or `libncurses6`.
- The `dmidecode` package is required for Veeam Agent management — a valid BIOS UUID must be obtainable either from `dmidecode | grep -i uuid` or from `/sys/class/dmi/id/product_uuid`. Each Veeam Agent that consumes a license installed in Veeam Backup & Replication must have a unique BIOS UUID. If a valid UUID cannot be obtained, Veeam will generate it automatically.
- The `libmysqlclient` package is required to process MySQL database system located on the Veeam Agent server. For details, see Backup of MySQL Database. Package version varies according to the MySQL database system version that you use.
- The `libpq5` package is required to process PostgreSQL database system located on the Veeam Agent server. For details, see Backup of PostgreSQL Database.
- The `python3` package or another RPM package providing a `/usr/bin/python3` binary is required for CentOS, RHEL 7.0 and later distributions if a pre-built binary `kmod-veeamsnap` package is to be installed.
- The `btrfs-progs` package version 3.16 or later is required.
- The `lvm2` package version up to 2.03.12 is required. Later versions are not supported.

Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see File System.

Backup Target

Backup can be performed to the following types of storage:

- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.
- Veeam Backup & Replication 11 or later backup repository (including deduplication appliances).
- Veeam Cloud Connect 11 or later cloud repository.

Network

Consider the following:

- Veeam Agent for Linux should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 addresses.

Used Ports

Make sure that you open ports listed below to enable proper work of Veeam Agent for Linux.

From	To	Protocol	Port	Notes
Veeam Agent computer	Veeam backup server	TCP	10006	Default port used for communication with the Veeam backup server. Data between the Veeam Agent for Linux computer and backup repositories is transferred directly, bypassing Veeam backup servers.
	Shared folder SMB (CIFS) share	TCP UDP	137 to 139, 445	Ports used as a data transmission channel from the Veeam Agent for Linux computer to the target SMB (CIFS) share. Ports 137 to 139 are used by backup infrastructure components to communicate using NetBIOS.
	Shared folder NFS share	TCP UDP	111, 2049	Standard NFS ports used as a data transmission channel from the Veeam Agent for Linux computer to the target NFS share.
	Veeam Agent computer	TCP	2500 to 3300	Default range of ports used for communication between Veeam Agent for Linux components during data transmission. For every TCP connection that a backup job uses, one port from this range is assigned. Ports must be open for incoming and outgoing traffic. Established connections must be allowed.
Communication with Veeam Backup & Replication Repositories				
Veeam Agent computer	Linux server performing the role of a backup repository	TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a backup job uses, one port from this range is assigned.
	Microsoft Windows server performing the role of a	TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic RPC port range. For more information, see this Microsoft article.

From	To	Protocol	Port	Notes
	backup repository	TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a backup job uses, one port from this range is assigned.
Communication with Veeam Cloud Connect Repositories				
Veeam Agent computer	Cloud gateway	TCP	6180	Port on the cloud gateway used to transport Veeam Agent data to the Veeam Cloud Connect repository.
	Certificate Revocation Lists	TCP	80 or 443 (most popular)	Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider. Generally, information about CRL locations can be found on the CA website.

IMPORTANT

The list of ports required for computers booted from the Veeam Recovery Media is the same as the list of ports required for Veeam Agent computers.

Installation and Configuration

You can install Veeam Agent for Linux on any Linux-based endpoint whose data you plan to protect — virtual machine or physical device (server, desktop or laptop).

Before You Begin

Before you start the installation process, check the following prerequisites:

1. The computer on which you plan to install Veeam Agent for Linux must satisfy system requirements specified in this document. To learn more, see System Requirements.
2. To install Veeam Agent for Linux software packages, you must use the `root` account or any user account that has super user (root) privileges on the computer where you plan to install the product.
3. If you have used the Beta version of Veeam Agent for Linux, you must remove Veeam Agent for Linux software packages prior to installing the release version of the product. To learn more, see Uninstalling Veeam Agent for Linux.

Installing Veeam Agent for Linux

> **NOTE**
>
> To make UEFI systems with Secure Boot work with the pre-built `veeamsnap` kernel module, Veeam Agent requires the Veeam public key enrolled to the MOK list. The key is available in the `veeamsnap-ueficert-5.0.2.4567-1.noarch` package residing on the Veeam repository. Veeam Agent requests the key enrollment during the package installation. After that, you must reboot computer to enroll the key into the UEFI database.
>
> After the package installation, you can check that the key enrollment is planned for the next reboot with the following command: `mokutil -N`. If the command output shows that the key enrollment is not planned, you can do the following:
>
> 1. Request the enrollment of the public key manually with the following command: `mokutil --import veeamsnap-ueficert.crt`.
> 2. Reboot the Veeam Agent computer to enroll the key into the UEFI database.
> 3. Check that the key is successfully enrolled with the following command: `mokutil -l`

To install Veeam Agent for Linux:

1. Download the Veeam software repository installation package (`veeam-release`) from the this Veeam webpage, and save the downloaded package on the computer where you plan to install the product.

   > **TIP**
   >
   > If the computer where you want to install Veeam Agent for Linux is not connected to the internet, you can download and install Veeam Agent for Linux packages manually. To learn more, see Installing Veeam Agent for Linux in Offline Mode.

2. Navigate to the directory where you have saved the `veeam-release` package and install the package with the following commands:

   For CentOS / RHEL / Oracle Linux / Fedora

   ```
   rpm -ivh ./veeam-release* && yum check-update
   ```

   For openSUSE / SLES

   ```
   zypper in ./veeam-release* && zypper refresh
   ```

   For Debian / Ubuntu

   ```
   dpkg -i ./veeam-release* && apt-get update
   ```

3. Install Veeam Agent for Linux packages from the Veeam software repository. To install Veeam Agent for Linux, you can use a package manager of your choice that works with software packages in your Linux distribution. For example, use the following commands:

For CentOS / RHEL / Fedora

```
yum install veeam
```

NOTE

[For CentOS / RHEL] If the `dkms` package is already installed in the OS, you can install Veeam Agent with one of the following commands:

- `yum install veeam`

  With this command, the `veeamsnap` kernel module will be installed from the source RPM package using `dkms`.
- `yum install kmod-veeamsnap veeam`

  With this command, the non-DKMS version of the `veeamsnap` kernel module will be installed from the pre-built `kmod-veeamsnap` binary package.

For Oracle Linux

```
yum install veeamsnap
yum install veeam
```

For openSUSE Tumbleweed

```
zypper in veeam
```

For openSUSE Leap 15.3

```
zypper in veeamsnap-kmp-default
zypper in veeam
```

For SLES with Default kernel

```
zypper in veeamsnap-kmp-default
zypper in veeam
```

For SLES with Trace kernel

```
zypper in veeamsnap-kmp-trace
zypper in veeam
```

For SLES with Xen kernel

```
zypper in veeamsnap-kmp-xen
zypper in veeam
```

For SLES with PAE kernel

```
zypper in veeamsnap-kmp-pae
zypper in veeam
```

For SLES with Preempt kernel

```
zypper in veeamsnap-kmp-preempt
zypper in veeam
```

For Debian / Ubuntu

```
apt-get install veeam
```

Managing Package Dependencies

The following dependency packages may require special handling in case you see installation errors:

- The `dkms` package is not present in default repositories for some Linux distributions. You should obtain it from third-party repositories:
  - EPEL repository (for CentOS / RHEL / Oracle Linux / Fedora)
  - Packman repository (for openSUSE). To learn more, see Installing dkms in openSUSE.

    For SLES, the `dkms` package is not available in the Packman repository. You must use the package intended for openSUSE. To learn more, see this Veeam KB article.
- Extended kernels, such as `kernel-pae`, `kernel-uek` and other, require appropriate `kernel-devel` packages to be installed, for example, `kernel-pae-devel`, `kernel-uek-devel`, and so on.

  Version of the `kernel-devel` package must match your current kernel version. To check your current kernel version, run the `uname -r` command.

  [For RHEL and derivatives] If the `yum` package manager installs packages that do not match your current kernel version, you should either update your system or fetch older versions of the required packages from the CentOS Vault repository.

Installing dkms in openSUSE

In openSUSE systems, while installing the `dkms` package, you may see an error similar to the following:

```
Problem: nothing provides kernel-devel needed by dkms-2.2.0.3-14.1.noarch
Solution 1: do not install dkms-2.2.0.3-14.1.noarch
Solution 2: break dkms-2.2.0.3-14.1.noarch by ignoring some of its dependencies
```

To install the `dkms` package, do the following:

1. Make sure that you have an appropriate `kernel-devel` package installed and its version matches your kernel version. For example:

```
root@localhost:~> rpm -qa | grep kernel-default
kernel-default-devel-3.0.101-91.1
kernel-default-3.0.101-91.1
```

2. Install the `dkms` package ignoring dependencies:

```
zypper -n install --force dkms
```

3. Make sure that you have allowed unsupported modules. To learn more, see this SUSE webpage.

Installing Veeam Agent for Linux in Offline Mode

If a computer where you want to install Veeam Agent for Linux has no connection to the internet, for example, for security reasons, you can install Veeam Agent in the offline mode. In this scenario, you do not need to download and install the Veeam software repository installation package (`veeam-release`). Instead, you need to download all Veeam Agent packages from the Veeam software repository and install them on the target computer.

To install Veeam Agent:

1. On a computer that is connected to the internet, download Veeam Agent packages intended for your Linux distribution from the Veeam software repository.

   The repository has the following structure:

   - For Debian packages: *Package format* > *Package*

     For example, Veeam Agent packages in the Debian format reside in the /deb/veeam/ , /deb/veeamsnap/ and /deb/veeam-nosnap/ folders of the Veeam software repository.

   - For RPM packages: *Package format* > *Distribution* > *Version* > *Architecture*

     For example, Veeam Agent packages for 64-bit RHEL 7 reside in the /rpm/el/7/x86_64/ folder of the Veeam software repository, and packages for 64-bit SLES 15 SP2 reside in the /rpm/sles/SLE_15_SP2/x86_64/ folder.

2. Save Veeam Agent packages to a directory that can be accessed from the computer where you want to install the product, for example, a directory on a local drive or USB drive, or a network shared folder.
3. On the computer where you want to install Veeam Agent, navigate to the directory where you have saved the packages and install Veeam Agent:

   - Installing Veeam Agent for Linux in CentOS / RHEL
   - Installing Veeam Agent for Linux in Oracle Linux
   - Installing Veeam Agent for Linux in Fedora
   - Installing Veeam Agent for Linux in SLES
   - Installing Veeam Agent for Linux in openSUSE
   - Installing Veeam Agent for Linux in Debian / Ubuntu

> **TIP**
>
> You can also set up a local mirror of the Veeam software repository in your internal network and add this repository to the list of software sources on a computer where you want to install the product. These operations may differ depending on the Linux distribution and package manager that you use. To learn more, refer to the documentation of your Linux distribution.
>
> After you add a local repository to the list of software sources on a computer, you will be able to install and upgrade Veeam Agent in a regular way. To learn more, see Installing Veeam Agent for Linux and Upgrading Veeam Agent for Linux.

Installing Veeam Agent for Linux in CentOS / RHEL

To install Veeam Agent for Linux, use the following commands:

For 32-bit RHEL 6

```
rpm -i <...>/kmod-veeamsnap-5.0.2.4567-2.6.32_131.0.15.el6.i386.rpm
rpm -i <...>/veeam-5.0.2.4567-1.el6.i386.rpm
```

For 64-bit RHEL 6

```
rpm -i <...>/kmod-veeamsnap-5.0.2.4567-2.6.32_131.0.15.el6.x86_64.rpm
rpm -i <...>/veeam-5.0.2.4567-1.el6.x86_64.rpm
```

For CentOS 7 / RHEL 7

```
rpm -i <...>/kmod-veeamsnap-5.0.2.4567-1.el7.x86_64.rpm
rpm -i <...>/veeam-5.0.2.4567-1.el7.x86_64.rpm
```

For CentOS 8 / RHEL 8

```
rpm -i <...>/kmod-veeamsnap-5.0.2.4567-4.18.0_80.el8.x86_64.rpm
rpm -i <...>/veeam-5.0.2.4567-1.el8.x86_64.rpm
```

where:

<. . .> — path to a directory where you have saved Veeam Agent packages.

> **NOTE**
>
> The pre-built `veeamsnap` binaries require kernel 2.6.32-131.0.15 or later for RHEL 6 (excluding 2.6.32-279.el6.i686) and kernel 3.10.0-123 or later for CentOS / RHEL 7.0 - 7.7 to operate.

Installing Veeam Agent for Linux in Oracle Linux

To install Veeam Agent for Linux, use the following commands:

For Oracle Linux 6

```
rpm -i <...>/veeamsnap-5.0.2.4567-1.noarch.rpm
rpm -i <...>/veeam-5.0.2.4567-1.el6.x86_64
```

For Oracle Linux 7

```
rpm -i <...>/veeamsnap-5.0.2.4567-1.noarch.rpm
rpm -i <...>/veeam-5.0.2.4567-1.el7.x86_64.rpm
```

For Oracle Linux 8

```
rpm -i <...>/veeamsnap-5.0.2.4567-1.noarch.rpm
rpm -i <...>/veeam-5.0.2.4567-1.el8.x86_64.rpm
```

where:

<...> — path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in Fedora

To install Veeam Agent for Linux, use the following commands:

```
rpm -i <...>/veeamsnap-5.0.2.4567-1.noarch.rpm
rpm -i <...>/veeam-5.0.2.4567-1.fc24.x86_64.rpm
```

where:

<...> — path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in SLES

NOTE

If you use a Trace, Xen, PAE, or preempt kernel, you need to choose the corresponding `veeamsnap` package for installation instead of the default one, for example, `veeamsnap-kmp-trace-5.0.2.4567_3.0.101_63-2.1.i586.rpm` instead of `veeamsnap-kmp-default-5.0.2.4567_3.0.101_63-2.1.i586.rpm`.

To install Veeam Agent for Linux, use the following commands:

For 32-bit SLES 11 SP4

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k3.0.101_63-2.i586.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles11.i386.rpm
```

For 64-bit SLES 11 SP4

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k3.0.101_63-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles11.x86_64.rpm
```

For 64-bit SLES 12 SP2

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k4.4.21_69-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

For 64-bit SLES 12 SP3

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k4.4.73_5-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

For 64-bit SLES 12 SP4

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k4.12.14_94.41-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

For 64-bit SLES 12 SP5

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k4.12.14_120-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

For 64-bit SLES 15 SP0

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k4.12.14_23-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

For 64-bit SLES 15 SP1

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k4.12.14_195-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

For 64-bit SLES 15 SP2

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k5.3.18_22-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

For 64-bit SLES 15 SP3

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k5.3.18_57-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.sles12.x86_64.rpm
```

where:

<. . .> — path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in openSUSE

To install Veeam Agent for Linux, use the following commands:

For 32-bit openSUSE Tumbleweed

```
zypper in <...>/veeamsnap-5.0.2.4567-1.sles.noarch.rpm
zypper in <...>/veeam-5.0.2.4567-1.suse.i386.rpm
```

For 64-bit openSUSE Tumbleweed

```
zypper in <...>/veeamsnap-5.0.2.4567-1.sles.noarch.rpm
zypper in <...>/veeam-5.0.2.4567-1.suse.x86_64.rpm
```

For 64-bit openSUSE Leap 15.2

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k5.3.18_lp152.19-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.suse.x86_64.rpm
```

For 64-bit openSUSE Leap 15.3

```
zypper in <...>/veeamsnap-kmp-default-5.0.2.4567_k5.3.18_57-2.x86_64.rpm
zypper in <...>/veeam-5.0.2.4567-1.suse.x86_64.rpm
```

where:

<. . .> — path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in Debian / Ubuntu

To install Veeam Agent for Linux, use the following commands:

For 32-bit Debian / Ubuntu 14.04, 16.04, 18.04, 20.04, 20.10, 21.04, 21.10

```
dpkg -i <...>/veeamsnap_5.0.2.4567_all.deb
dpkg -i <...>/veeam_5.0.2.4567_i386.deb
```

For 64-bit Debian / Ubuntu 14.04, 16.04, 18.04, 20.04, 20.10, 21.04, 21.10

```
dpkg -i <...>/veeamsnap_5.0.2.4567_all.deb
dpkg -i <...>/veeam_5.0.2.4567_amd64.deb
```

where:

<. . .> — path to a directory where you have saved Veeam Agent packages.

Upgrading Veeam Agent for Linux

For Veeam Agent for Linux, upgrade to newer versions is supported. You can start the upgrade process when the new version becomes available.

During the upgrade process, configuration and backup files that were created with the previous version of Veeam Agent are not impacted in any way.

> **IMPORTANT**
>
> Before starting the upgrade process, make sure that there are no jobs running on the Veeam Agent computer.

The upgrade procedure differs depending on the Linux distribution that you use:

- Upgrading Veeam Agent for Linux in CentOS / RHEL
- Upgrading Veeam Agent for Linux in Fedora / Oracle Linux
- Upgrading Veeam Agent for Linux in openSUSE
- Upgrading Veeam Agent for Linux in SLES 11 SP4
- Upgrading Veeam Agent for Linux in SLES 12 SP2 – SLES 15 SP3
- Upgrading Veeam Agent for Linux in Debian / Ubuntu

Upgrading Veeam Agent for Linux in CentOS / RHEL

To upgrade Veeam Agent for Linux, use the following commands:

```
yum update veeam
rpm -e --nodeps dkms veeamsnap
yum install kmod-veeamsnap*
```

or

```
yum install kmod-veeamsnap*
yum update veeam
rpm -e dkms veeamsnap
```

* With these commands, a pre-built `veeamsnap` binary package will be installed in your system. To stay on the DKMS version of the `veeamsnap` kernel module, use the following command for upgrade:

```
yum update veeamsnap && yum update veeam
```

Upgrading Veeam Agent for Linux in Fedora / Oracle Linux

To upgrade Veeam Agent for Linux, use the following command:

```
yum update veeam
```

Upgrading Veeam Agent for Linux in openSUSE

To upgrade Veeam Agent for Linux, use the following commands:

For openSUSE Tumbleweed

```
zypper up veeam
```

For openSUSE Leap 15.2

```
zypper in veeamsnap-kmp-default-5.0.2.4567_k5.3.18_lp152.19-2.x86_64
zypper in --force veeamsnap-kmp-default-5.0.2.4567_k5.3.18_lp152.19-2.x86_64
```

For openSUSE Leap 15.3

```
zypper in veeamsnap-kmp-default-5.0.2.4567_k5.3.18_57-2.x86_64
zypper in --force veeamsnap-kmp-default-5.0.2.4567_k5.3.18_57-2.x86_64
```

With these commands, a pre-built `veeamsnap` binary package will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the `veeamsnap` kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper update veeam
```

Upgrading Veeam Agent for Linux in SLES 11 SP4

To upgrade Veeam Agent for Linux, use the following commands:

For Default kernel

```
zypper in veeamsnap-kmp-default
zypper in --force veeamsnap-kmp-default
```

For Xen kernel

```
zypper in veeamsnap-kmp-xen
zypper in --force veeamsnap-kmp-xen
```

For Trace kernel

```
zypper in veeamsnap-kmp-trace
zypper in --force veeamsnap-kmp-trace
```

For PAE kernel

```
zypper in veeamsnap-kmp-pae
zypper in --force veeamsnap-kmp-pae
```

With these commands, a pre-built `veeamsnap` binary package will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the `veeamsnap` module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper update veeam
```

Upgrading Veeam Agent for Linux in SLES 12 SP2 – SLES 15 SP3

To upgrade Veeam Agent for Linux:

1. Download the Veeam software repository installation package (`veeam-release`) from the Veeam Download page, and save the downloaded package on the computer where you plan to install the product.
2. Install the Veeam software repository installation package with the following command:

```
zypper in ./veeam-release* && zypper refresh
```

3. Upgrade Veeam Agent for Linux with the following commands:

   For Default kernel

```
zypper in veeamsnap-kmp-default
zypper in --force veeamsnap-kmp-default
```

   Preempt kernel

```
zypper in veeamsnap-kmp-preempt
zypper in --force veeamsnap-kmp-preempt
```

   With this command, a pre-built `veeamsnap` binary package will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the `veeamsnap` module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper update veeam
```

Upgrading Veeam Agent for Linux in Debian / Ubuntu

To upgrade Veeam Agent for Linux, use the following commands:

```
apt-get update
apt-get install veeam
```

Granting Permissions to Users

When you install Veeam Agent for Linux, the product program files are placed to the folders on the system volume. For full access to Veeam Agent files, super user (root) privileges are required. Rights to execute product files and run commands are also granted to users that belong to the `veeam` group.

The `veeam` group is automatically created by Veeam Agent at the process of the product installation. To let regular users work with Veeam Agent without the need to gain root privileges, you can add the necessary users to this group. Users in the `veeam` group will be able to execute Veeam Agent commands and perform backup and restore tasks under regular user account.

To add a user to the `veeam` group, in most of Linux distributions you can use the following command:

```
usermod -a -G veeam <username>
```

where:

`<username>` — name of the account to which you want to grant access to Veeam Agent.

For example:

```
root@srv01:~# usermod -a -G veeam user
```

IMPORTANT

Consider the following:

- To add a user to the `veeam` group, you must have super user (root) privileges in the Linux OS.
- After the user is added to the `veeam` group, the user must re-login to the Linux OS.
- Add only trusted users to the `veeam` group. Veeam Agent for Linux daemon runs and executes commands and scripts with the super user privileges. Thus, users who belong to this group can potentially escalate their privileges through the creative use of pre-freeze/post-thaw or pre-job/post-job scripts.

To check whether the user who is currently logged in to the Linux OS is added to the `veeam` group, you can use the following command:

```
groups
```

For example:

```
user@srv01:~$ groups
user adm cdrom sudo dip plugdev lpadmin sambashare veeam
```

Performing Initial Setup

After you install Veeam Agent for Linux, you can use the Veeam Agent for Linux control panel to perform initial product setup. When you launch the control panel for the first time, Veeam Agent displays the initial setup wizard. The wizard offers you to accept license agreements, install a license and create a custom Veeam Recovery Media that will include drivers of your Veeam Agent computer.

To perform initial setup, launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command. Then use the initial setup wizard to complete the following steps:

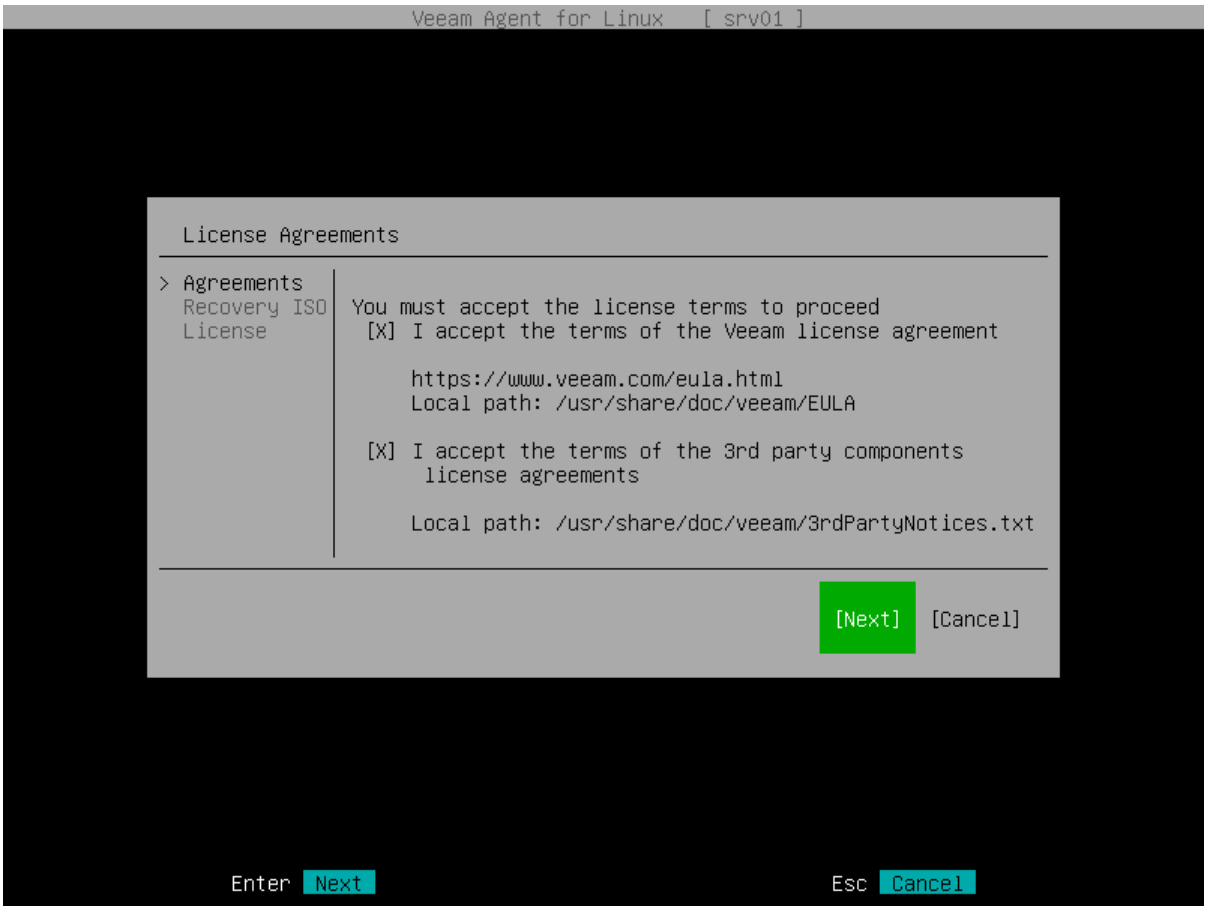
1. Accept Veeam and third-party license agreements.
2. Create a custom Veeam Recovery Media.
3. Install a license.

Step 1. Accept License Agreements

At the **Agreements** step of the initial setup wizard, accept the terms of the product license agreement and license agreements for third-party components of the product. You must accept the license agreements to start using the product. Until you accept the license agreements, you will not be able to perform backup and data recovery tasks with the Veeam Agent for Linux control panel and command line interface.

To accept the license agreements:

1. Make sure that the **I accept the terms of the Veeam license agreement** option is selected and press **Space**.
2. Select the **I accept the terms of the 3rd party components license agreements** option with the **Down** or **Tab** key and press **Space**.
3. Press **Enter**.



Step 2. Create Custom Veeam Recovery Media

At the **Recovery ISO** step of the initial setup wizard, specify settings for the custom Veeam Recovery Media.

In addition to the generic Veeam Recovery Media that is available for download at the Veeam website, you can create a custom Veeam Recovery Media. This option may be helpful if your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media. When you create the custom Veeam Recovery Media, Veeam Agent for Linux copies the Linux kernel running on your computer with its currently loaded modules and includes them into the custom recovery media.

Before you create a custom Veeam Recovery Media, check the following prerequisites:

- The Linux system must have the `genisoimage` package installed. For openSUSE and SLES 15 SP0 – SP3 distributions, the Linux system must have the `mkisofs` package installed instead.
- The Linux system must have the `mksquashfs` and `unsquashfs` utilities installed.
- For custom Veeam Recovery Media with EFI support, the Linux system must have the following packages installed:
  - `xorriso`
  - `isolinux` (or `syslinux`, if the software package repository of your Linux distribution lacks the `isolinux` package)
- For the scenario where you create a custom Veeam Recovery Media using the **Download and patch ISO** option, the Linux system must have the `wget` utilitiy installed.

> **TIP**
>
> If you do not want to create a custom Veeam Recovery Media at the process of initial product setup, switch to the **Next** button with the **Tab** key and press **Enter**. You will proceed immediately to the License step of the initial setup wizard.
>
> You can create the custom Veeam Recovery Media later, at any time you need, using the Veeam Agent for Linux command line interface. To learn more, see Creating Custom Veeam Recovery Media.

To specify settings for the custom Veeam Recovery Media:

1. Make sure that the **Patch Veeam Recovery Media ISO** option is selected and press **Space**.
2. If you want the Veeam Recovery Media to be able to boot on EFI-based systems, select the **EFI system** option with the **Tab** key and press **Space**.

   If you do not enable this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

3. Press **Tab** and select how you want to create a custom Veeam Recovery Media depending on the location of the generic recovery media ISO file:

   - If you did not download the generic Veeam Recovery Media from the Veeam website, make sure that the **Download and patch ISO** option is selected and press **Tab**. When you start the custom recovery media creation process, Veeam Agent for Linux will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the `/tmp/veeam` directory on the Veeam Agent computer and use this image to create the custom Veeam Recovery Media.

     Veeam Agent for Linux automatically downloads one of the following Veeam Recovery Media ISO files depending on the Veeam Agent computer architecture:

     - `veeam-recovery-i386-5.0.0.iso` — for x86 computers that run Linux kernel version 3.16 and later. The size of the downloaded ISO file is 460 MB.
     - `veeam-recovery-amd64-5.0.0.iso` — for x64 computers that run Linux kernel version 3.16 and later. The size of the downloaded ISO file is 490 MB.
     - `veeam-recovery-media-5.0.0.iso` — for x86 and x64 computers that run Linux kernel version earlier than 3.16. The size of the ISO image file depends on the Veeam Agent computer architecture: 136 MB for x86 computers and 139 MB for x64 computers.

   - If you have already downloaded a recovery media ISO file to a local directory on the Veeam Agent computer or to a network shared folder, select the **Patch local ISO** option with the **Down** key and press **Tab**. If you select this option, you will be able to specify a path to the ISO file of the generic Veeam Recovery Media.

     The name of the generic Veeam Recovery Media ISO file depends on the recovery image version, Veeam Agent computer architecture and the source from which you downloaded the ISO file: from the product download page or Veeam software repository. To learn more, see Veeam Recovery Media Versions.

4. If you have selected the **Patch local ISO** option, in the **Path to ISO** field, specify a path to the ISO file of the generic Veeam Recovery Media:

   a. Select the **Browse** option with the **Tab** key and press **Enter**.

   b. In the **Path to ISO** window, select the necessary directory and press **Enter**.

   c. Repeat the step 'b' until a path to the directory in which the recovery media ISO file resides appears in the **Current directory** field.

   d. In the directory where the recovery media ISO file resides, select the ISO file and press **Enter**.

5. In the **Save patched ISO to** field, specify a path to the resulting ISO file of the custom Veeam Recovery Media:

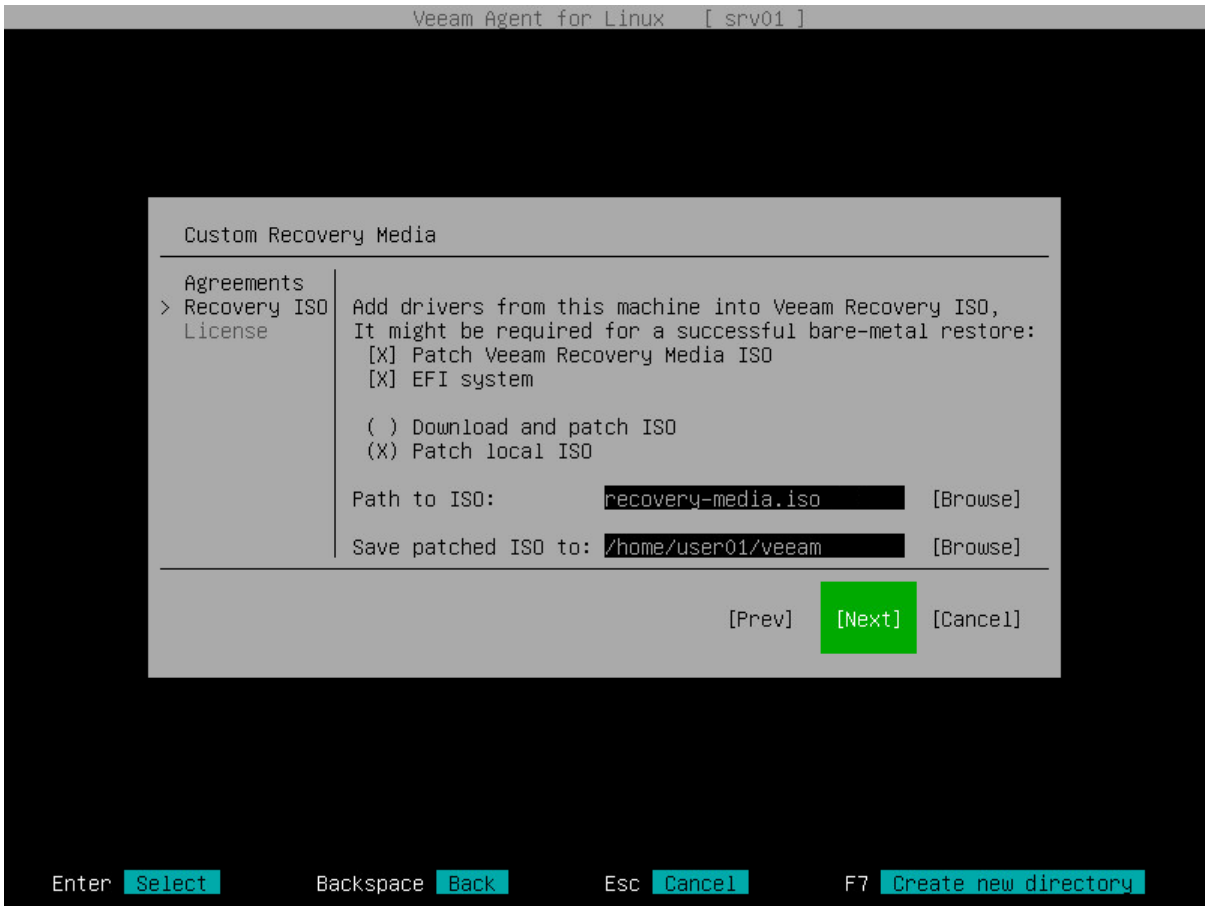
   a. Select the **Browse** option with the **Tab** key and press **Enter**.

   b. In the **Save patched ISO to** window, select the necessary directory and press **Enter**.

   c. Repeat the step 'b' until a path to the directory where you want to save the resulting custom recovery media ISO file appears in the **Current directory** field.

   d. Select the **OK** button with the **Tab** key and press **Enter**.

6. To start the custom recovery media creation process, select the **Next** button with the **Tab** key and press **Enter**.



Step 3. Install Product License

At the **License** step of the initial setup wizard, install the license. You can choose to install the license immediately or postpone this operation.

- If you choose to install the license, you can immediately browse for the license key on your computer and complete the license installation process.
- If you choose to postpone the license installation process, you will be able to install a license later at any time you need.

Until you install a license, Veeam Agent for Linux will operate in the Free edition. To learn more, see Product Editions.

> **NOTE**
>
> If you choose not to install a license and use Veeam Agent in the Free edition, Veeam Agent will display a notification offering to install a license every time you open the control panel. The notification will appear in the control panel until Veeam Agent completes the first backup job session.

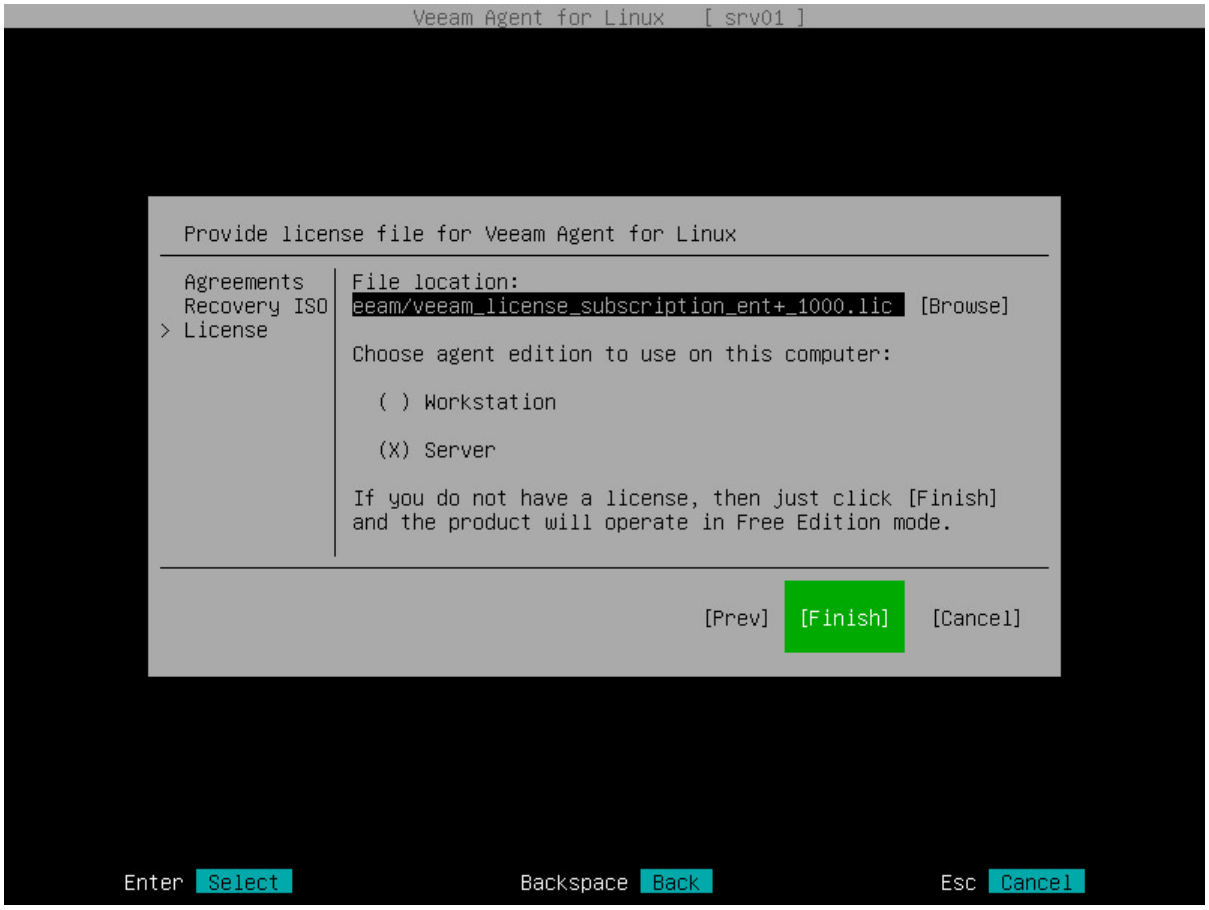
To install a license:

1. In the **File location** field, specify a path to the license key:
   a. Select the **Browse** option with the **Tab** key and press **Enter**.
   b. In the **Choose license file location** window, select the necessary directory and press **Enter**.
   c. Repeat the step 'a' until a path to the directory in which the license key resides appears in the **Current directory** field.
   d. In the directory where the license key resides, select the license key and press **Enter**.
2. In the **Choose agent edition to use on this computer** section, select the product edition in which Veeam Agent will operate and press **Enter** to install the license and finish working with the initial setup wizard.

TIP

Consider the following:

- If you do not want to install a license, to finish working with the initial setup wizard, switch to the **Finish** button with the **Tab** key and press **Enter**.
- You can view information about the installed license (expiration date, status of the license, current edition of the product and so on) in the Veeam Agent control panel or using the Veeam Agent command line interface. To learn more, see Viewing License.



Configuring Advanced Settings

Veeam Agent for Linux allows you to configure the following settings:

- HTTP proxy settings for Veeam Cloud Connect repository
- Connection settings for Veeam backup server

HTTP Proxy Settings for Veeam Cloud Connect Repository

If you want to use Veeam Agent for Linux to back up your data to a Veeam Cloud Connect repository, it might be required that you specify HTTP proxy settings for Veeam Agent.

Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider. In case it is not possible to establish a direct connection to CRLs, you must configure an HTTP proxy and specify settings to connect to the proxy in Veeam Agent.

To specify settings for an HTTP proxy, uncomment and edit the following lines in the `[cloudconnect]` section of the `/etc/veeam/veeam.ini` configuration file:

```
[cloudconnect]
...
# httpproxylogin= <username>
...
# httpproxypasswd= <password>
...
# httpproxyurl= <URL>
```

where:

- `<username>` — name of the account used to connect to the HTTP proxy.
- `<password>` — password of the account used to connect to the HTTP proxy.
- `<URL>` — URL of a proxy used for CRL checks.

> **NOTE**
>
> If the proxy does not require authentication, you do not need to specify the account name and password. Keep in mind that only the basic authentication method is supported for connection to a proxy.

For example:

```
[cloudconnect]
...
# HTTP proxy login
httpproxylogin= user01
# HTTP proxy password
httpproxypasswd= P@ssw0rd
# HTTP proxy URL for CRL checks
httpproxyurl= http://proxy.company.lan:3128
```

Connection Settings for Veeam Backup Server

If you want to connect Veeam Agent computer to Veeam backup server as a member of the protection group for pre-installed Veeam Agents, you must apply connection settings from the configuration file. The configuration file is one of the Veeam Agent for Linux setup files that you must obtain from your System Administrator. To learn more about protection group for pre-installed Veeam Agents, see the Protection Group Types section in the Veeam Agent Management Guide.

To connect Veeam Agent for Linux to Veeam backup server:

1. Get the configuration file from your System Administrator and upload this file on the Veeam Agent computer.
2. Navigate to the directory where you have saved the configuration file and run the following command:

```
veeamconfig mode setvbrsettings --cfg <file_name>.xml
```

   where `<file_name>` is a configuration file name.

Alternatively, you can specify the full path to the configuration file with the `--cfg` option.

For example:

```
user@srv01:~# veeamconfig mode setvbrsettings --cfg /home/Linux\ Servers\ Distr
ibs/Linux/LinuxServers.xml
```

Mind that the connection between Veeam backup server and Veeam Agent computer added as a member of the protection group for pre-installed Veeam Agents is not persistent. Veeam Agent synchronizes with Veeam Backup & Replication every 6 hours. To synchronize Veeam Agent immediately, run the following command from the Veeam Agent computer:

```
veeamconfig mode syncnow
```

Uninstalling Veeam Agent for Linux

To uninstall Veeam Agent for Linux, you need to remove the `veeam` and `veeamsnap` packages. For this operation, you can use any package manager that works with software packages in your Linux distribution.

The following examples show commands to uninstall Veeam Agent:

For CentOS / RHEL / Oracle Linux / Fedora

```
yum remove veeam veeamsnap
```

For openSUSE / SLES

```
zypper rm veeam veeamsnap
```

For Debian / Ubuntu

```
apt-get remove veeam veeamsnap
```

Getting Started

To protect your computer from a disaster of any kind, you must perform the following operations in Veeam Agent for Linux:

1. **Define what data you want to back up and configure the backup job**.

   Before you configure the backup job, you should decide on the following backup details:

   - Backup destination: where you want to store your backed-up data.
   - Backup scope: entire computer image, individual computer volumes or specific computer folders and files.
   - Backup schedule: how often you want to back up your data.

   After that, you can configure one or several backup jobs. The backup job captures the data that you have added to the backup scope and creates a chain of restore points in the target location. If your data gets lost or corrupted, you can restore it from the necessary restore point.

   In Veeam Agent, you can configure the backup job in one of the following ways:

   - With the Backup Job wizard
   - With the command line interface

2. **Monitor backup task performance**.

   You can use the Veeam Agent Control Panel to check how backup tasks are being performed, what errors have occurred during backup job sessions and so on. You can also use Veeam Agent command line interface to get information on backup and restore sessions status and view session logs. To learn more, see Reporting.

3. In case of a disaster, you can **restore the entire computer image or specific data** on the computer. With Veeam Agent, you can perform data recovery operations in several ways:

   - You can boot from the Veeam Recovery Media and perform volume-level restore or file-level restore.
   - You can perform volume-level restore with Veeam Agent command line interface.
   - You can perform file-level restore with the Veeam Agent File Level Restore wizard.
   - You can export backup to a VHD virtual disk and attach this disk to a virtual machine to recover your computer in virtual environment.

   To learn more, see Performing Restore.

Getting to Know User Interface

With Veeam Agent for Linux, you can perform backup, restore and configuration tasks in the following ways:

- With Veeam Agent control panel

  Veeam Agent control panel is a GUI-like user interface based on the `ncurses` programming library. With Veeam Agent control panel, you can perform all basic data protection tasks. You can configure a backup job, start and stop backup jobs, monitor backup job session performance and recover files and folders. When you perform restore tasks after booting from the Veeam Recovery Media, you can also perform volume-level restore with the Veeam Recovery Media wizard.

- With the command line interface

  With Veeam Agent command line interface, in addition to operations that can be performed with the Veeam Agent control panel, you can perform a set of advanced tasks. For example, you can:

  - Configure advanced settings for backup jobs: specify compression level and data block size.
  - Perform operations with backup repositories.
  - Perform volume-level restore without the need to boot from the Veeam Recovery Media.
  - Export backups to VHD virtual disks.
  - Monitor performance and status of any backup, restore and other data transfer session that was started in Veeam Agent.
  - View detailed information on every backup that was created with Veeam Agent.
  - Export/import Veeam Agent configuration database to/from a configuration file.

Veeam Agent for Linux Control Panel

Veeam Agent for Linux control panel is a GUI-like user interface that lets users perform main backup and restore tasks in an easy way. With Veeam Agent for Linux control panel, you do not need to work with Linux shell and remember numerous commands. However, some advanced Veeam Agent for Linux operations are not supported by the control panel and can be performed with the command line interface only.

> **IMPORTANT**
>
> You cannot use Veeam Agent for Linux control panel on terminals that do not support colors (for example, VT100).

To launch the Veeam Agent for Linux control panel, you can use the following commands:

```
veeamconfig ui
```

or

```
veeam
```

> **NOTE**
>
> Veeam Agent for Linux control panel is based on the `ncurses` programming library. To use the Veeam Agent for Linux control panel, you must have the `ncurses` library installed in your Linux OS. To learn more, see System Requirements.

When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent for Linux offers you to perform initial product setup. To learn more, see Performing Initial Setup.

After you perform initial product setup, before you configure the first backup job, you can use the Veeam Agent for Linux control panel to perform the following operations:

- Configure a new backup job.
- Restore files and folders from existing backup.
- Manage license and product logs.

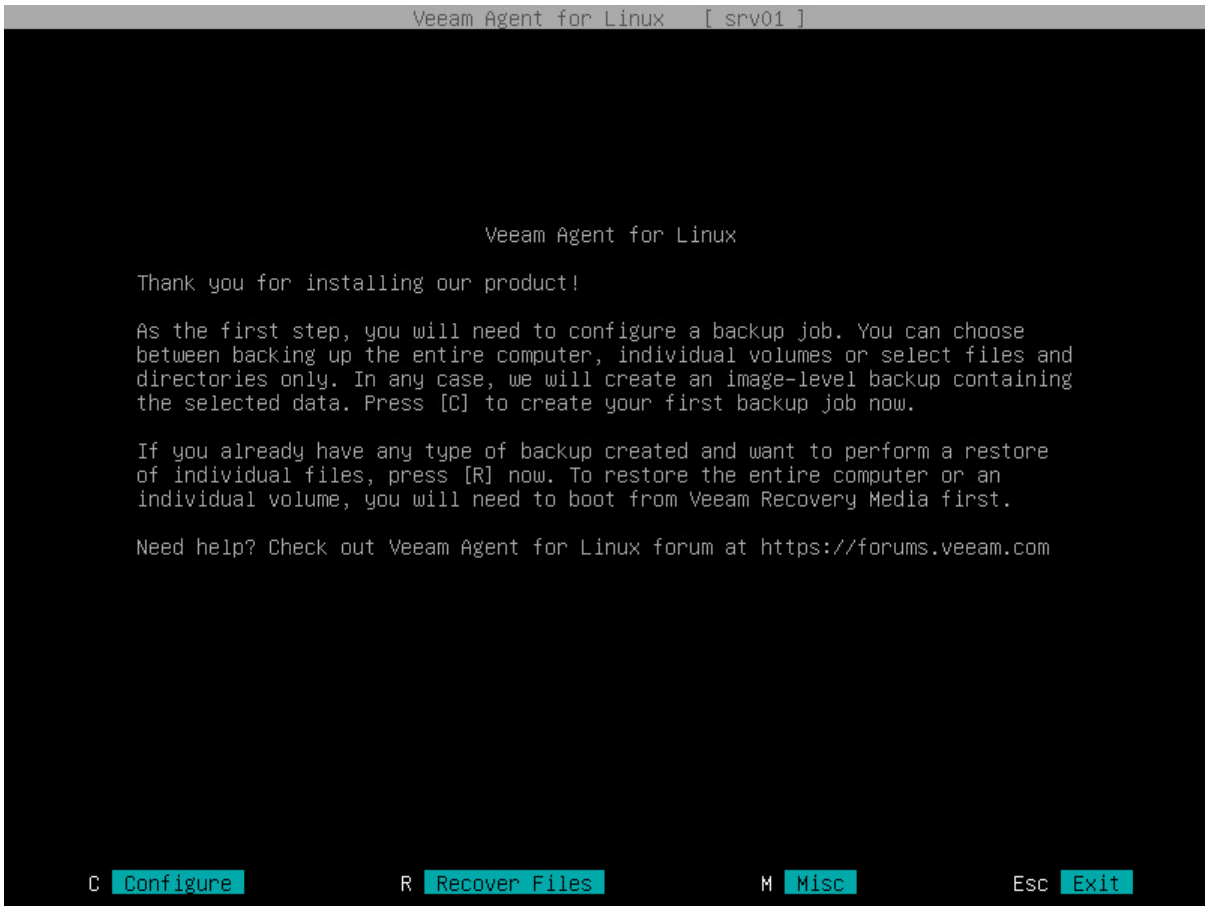
After you configure one or several backup jobs, you can also use the control panel to start a backup job and work with backup job sessions.

Navigating Veeam Agent for Linux Control Panel

In the Veeam Agent for Linux control panel, the use of a mouse is not supported. To start an operation, you need to use a specific key on your keyboard. For example, you can press the '**c**' key to start the backup job configuration, press the '**s**' key to start a backup job or press the '**r**' key to start the file-level restore process. Short help information on the currently available operations and keys is displayed at the bottom of the control panel.

To navigate the control panel, backup job configuration and file-level restore wizards, you can use the following keys:

- **Tab** — to switch between controls and buttons in the Backup Job wizard.
- **Up** and **Down** — to switch between items in a scrollable list.
- **Space** — to select the necessary item in a list. The selected item's mark may vary in different steps of the wizard.
- **Enter** — to proceed to the next step of a wizard or to view details of the backup job session selected in the list of sessions.
- **Backspace** — to return to the previous step of a wizard (you cannot use this button to change wizard steps when a text field is selected).
- **Esc** — to exit the currently used wizard or close the Veeam Agent for Linux control panel.



Command Line Interface

Veeam Agent command line interface is a powerful tool that lets users perform advanced operations that are not supported by the Veeam Agent control panel.

To work with Veeam Agent using command line interface, you can use a terminal console (TTY) or a terminal emulator of your choice. All tasks in Veeam Agent are performed with the `veeamconfig` command-line utility. To perform tasks with Veeam Agent, you should construct the necessary command and type it in the Linux shell prompt.

You can view short help information on every Veeam Agent command at any time you need. To learn more, see Viewing Help.

You should construct a command in the following format:

```
veeamconfig <command_1> <command_2> --<parameter_1> --<parameter_2> --<paramete
r_n>
```

where:

- `<command_1>` — command that defines a type of an object with which you want to perform a task. Currently, the following commands are available in Veeam Agent:
  - aap
  - agreement
  - backup
  - cloud
  - config
  - grablogs
  - help
  - job
  - license
  - mode
  - patchiso
  - point
  - repository
  - schedule
  - session
  - ui
  - version
  - vbrserver

- `<command_2>` — command that defines a task that you want to perform with an object of the specified type. For example, you can perform the following commands with backup repositories:
  - create
  - delete
  - edit
  - help
  - list
  - rescan
- `<parameter_1>, <parameter_2>, <parameter_n>` — parameters for the command that you want to execute. Commands may require one or several mandatory or optional parameters. Some commands, for example, `veeamconfig ui` and `veeamconfig [<command>] help` do not require parameters.

The following example shows the command that displays a list of backup repositories configured in Veeam Agent and the output of this command:

```
user@srv01:~$ veeamconfig repository list
Name          ID                                      Location            Typ
e   Backup server
Repository_1  {818e3a0f-8155-4a51-9430-248a203a43d1}  /home/backups       loca
l
Repository_2  {2155a2e7-a1e9-4347-9d8b-cf8f3a6f3fcb}  172.17.53.47/veeam  cif
s
```

Viewing Help

You can view short help information on the specific Veeam Agent command. To view help, use the following command:

```
veeamconfig <command> help
```

where:

<command> — name of the command for which you want to view help information.

For example:

```
user@srv01:~$ veeamconfig help
```

or

```
user@srv01:~$ veeamconfig job help
```

or

```
user@srv01:~$ veeamconfig job create help
```

You can also view the manual page for the veeamconfig utility. Use the following command:

```
man veeamconfig
```

Licensing

To work with Veeam Agent, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product.

If you want to use a commercial version of Veeam Agent, you must obtain a license and install it on the protected computer. If you do not install a license, the product will operate in the Free edition.

You can use the Veeam Agent control panel or Veeam Agent command line interface to install a license, monitor status of the installed license or remove the license if necessary.

Product Editions

Veeam Agent for Linux offers three product editions that define product functionality and operation modes:

- *Server* — a commercial edition that provides access to all product functions and is intended for performing data protection tasks on servers that run Linux OS. Veeam Agent for Linux can operate in the server edition if a commercial license that supports this edition is installed on the protected computer.
- *Workstation* — a commercial edition that offers limited capabilities that are sufficient for performing data protection tasks on desktop computers and laptops that run Linux OS. Veeam Agent for Linux can operate in the workstation edition if a commercial license that supports this edition is installed on the protected computer.
- *Free* — a free edition that offers the same capabilities as the Workstation edition but does not come with a commercial support program. In contrast to the workstation and server editions, the Free edition does not require a license.

For more information about product editions, pricing and features available for them, see this Veeam webpage.

> **TIP**
>
> To check in which edition Veeam Agent for Linux currently operates, you can use the Veeam Agent for Linux control panel or command line interface. To learn more, see Viewing License Information.

When you install a license on the protected computer, you can select in which edition Veeam Agent for Linux will operate: server edition or workstation edition (if both editions are supported by the license). If you use Veeam Agent for Linux with Veeam Backup & Replication, you must manage product licenses and editions from the Veeam Backup & Replication console. To learn more, see Managing License with Veeam Backup & Replication.

After the license expires, Veeam Agent for Linux automatically switches to the Free edition. To learn more, see License Expiration.

Limitations for Free and Workstation Editions

Compared to the Server edition of Veeam Agent for Linux, Free and Workstation editions have the following limitations:

1. [For Free edition] The number of backup jobs that you can configure in Veeam Agent for Linux is limited to one.
2. [For Free edition] You cannot use a Veeam Cloud Connect repository as a target location for backup files.
3. [For Workstation edition] The number of backup jobs that you can configure in Veeam Agent for Linux is limited to one backup job targeted at a local drive, network shared folder or Veeam backup repository plus unlimited number of backup jobs targeted at a Veeam Cloud Connect repository.
4. You cannot specify pre-freeze and post-thaw scripts in the backup job settings.
5. You cannot specify database system processing settings.

License Agreement

After you install Veeam Agent for Linux, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product. Until you accept the license agreements, you will not be able to perform backup and data recovery tasks with the Veeam Agent for Linux control panel and command line interface.

License agreements are located in the `/usr/share/doc/veeam` directory of the machine where you installed the product.

The process of accepting license agreements differs depending on the way you work with Veeam Agent for Linux — using the control panel or command line interface.

- When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent prompts you to accept the license agreements at the **Agreements** step of the initial setup wizard. To learn more, see Accept License Agreements.
- When you run a Veeam Agent for Linux command, for example, `veeamconfig repository create`, Veeam Agent prompts you to accept license agreements. To accept the license agreement, type `y` or `yes` in the command prompt and press **Enter**.

  Alternatively, you can accept license agreements using the dedicated commands. To learn more, see Accepting License Agreements.

Installing License

When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent for Linux offers you to install a license at the License step of the initial setup wizard. You can choose to install the license immediately or postpone this operation.

If you choose to postpone the license installation process, you can install a license later at any time you need. Until you install a license, Veeam Agent for Linux will operate in the Free edition. To learn more, see Product Editions.

> **NOTE**
>
> If you choose not to install a license and use Veeam Agent for Linux in the Free edition, Veeam Agent for Linux will display a notification offering to install a license every time you open the control panel. The notification will appear in the control panel until Veeam Agent for Linux completes the first backup job session.

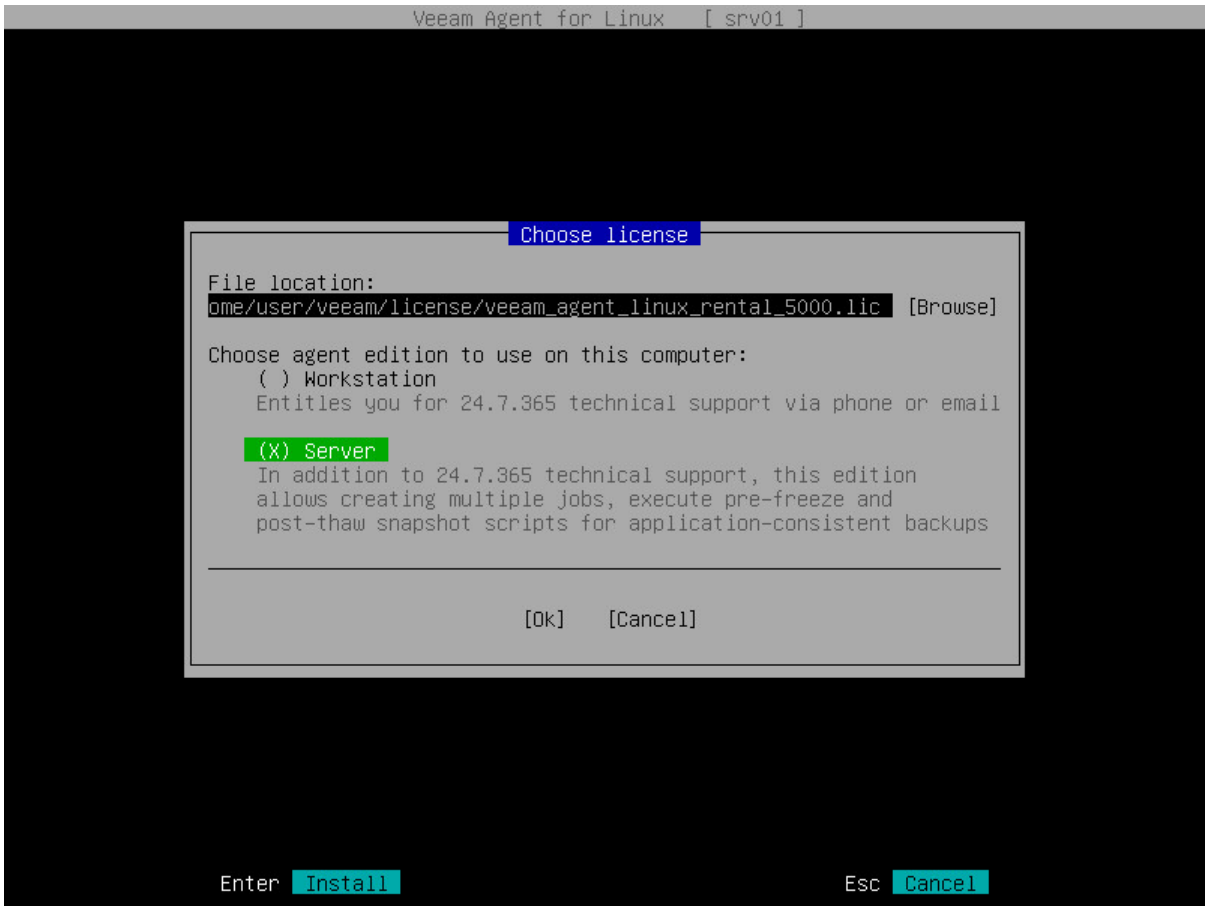
To install a license:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the '**m**' key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press **Enter**.
4. In the **Manage license** window, make sure that the **Install** button is selected and press **Enter**.
5. In the **Choose license** window, in the **File location** field, specify a path to the license key:
   a. Select the **Browse** option with the **Tab** key and press **Space** or **Enter**.
   b. In the **Choose license file location** window, select the necessary directory and press **Enter**.
   c. Repeat the step 'b' until a path to the directory in which the license key resides appears in the **Current directory** field.

> **TIP**
>
> If you chose to install the license immediately from the Veeam Agent for Linux welcome screen notification, you will pass to the **Choose license** step right from the notification window.

6. In the **Choose agent edition to use on this computer** section, select the product edition in which Veeam Agent for Linux will operate and press **Enter**. To learn more about editions, see Product Editions.
7. Veeam Agent for Linux will install the license and display a window notifying that the license is successfully installed. Press **Enter** to finish the license installation process.

> **TIP**
>
> After you install a license, you can view information about the license (expiration date, status of the license, current edition of the product and so on) in the **Manage license** window. You can also check information about the license using the Veeam Agent for Linux command line interface. To learn more, see Viewing License.

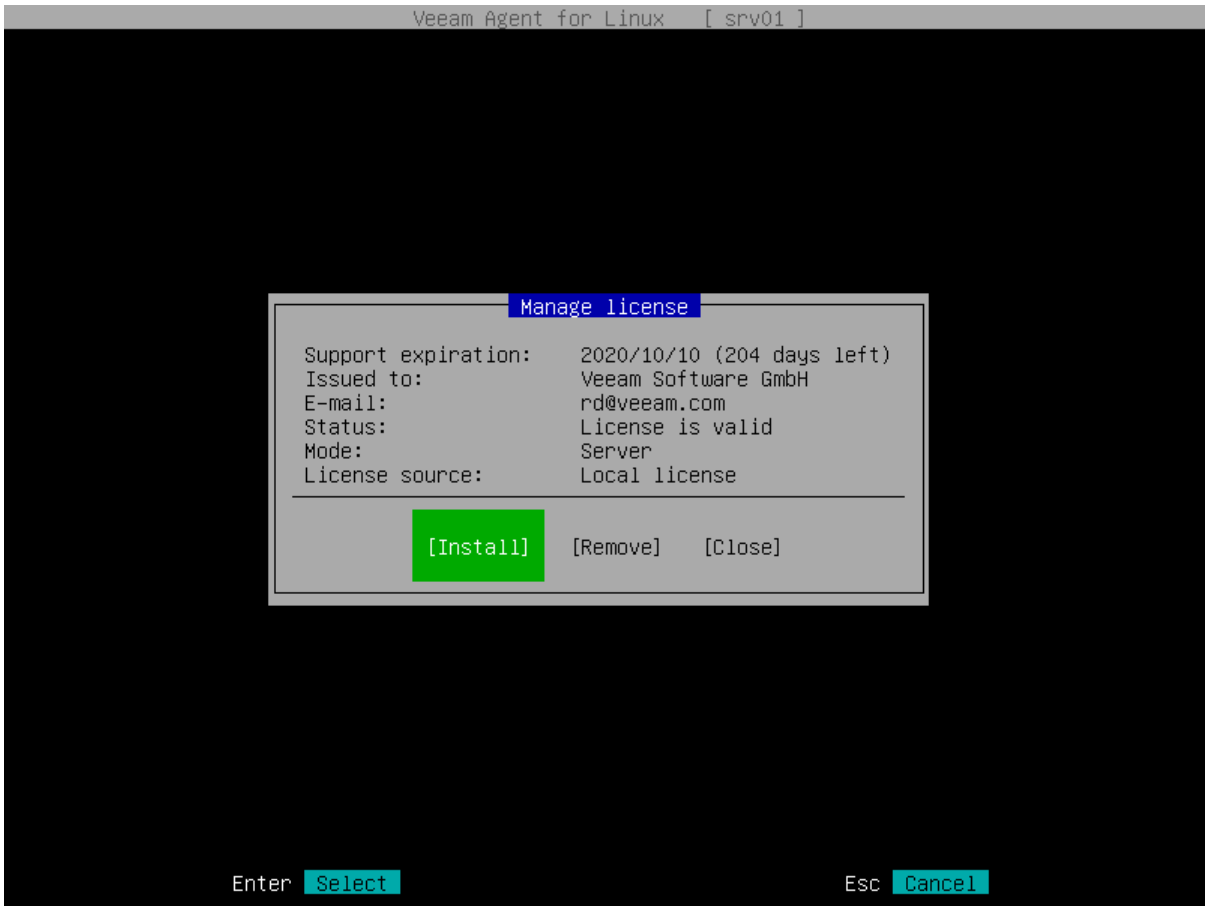

Veeam Agent for Linux [srv01]
Choose license
File location:
ome/user/veeam/license/veeam_agent_linux_rental_5000.lic [Browse]
Choose agent edition to use on this computer:
( ) Workstation
Entitles you for 24.7.365 technical support via phone or email
(X) Server
In addition to 24.7.365 technical support, this edition
allows creating multiple jobs, execute pre-freeze and
post-thaw snapshot scripts for application-consistent backups
[Ok] [Cancel]
Enter Install
Esc Cancel

Viewing License Information

To view information about the installed license, do the following:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the '**m**' key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press **Enter**.

Veeam Agent for Linux will display information about the license.



Removing License

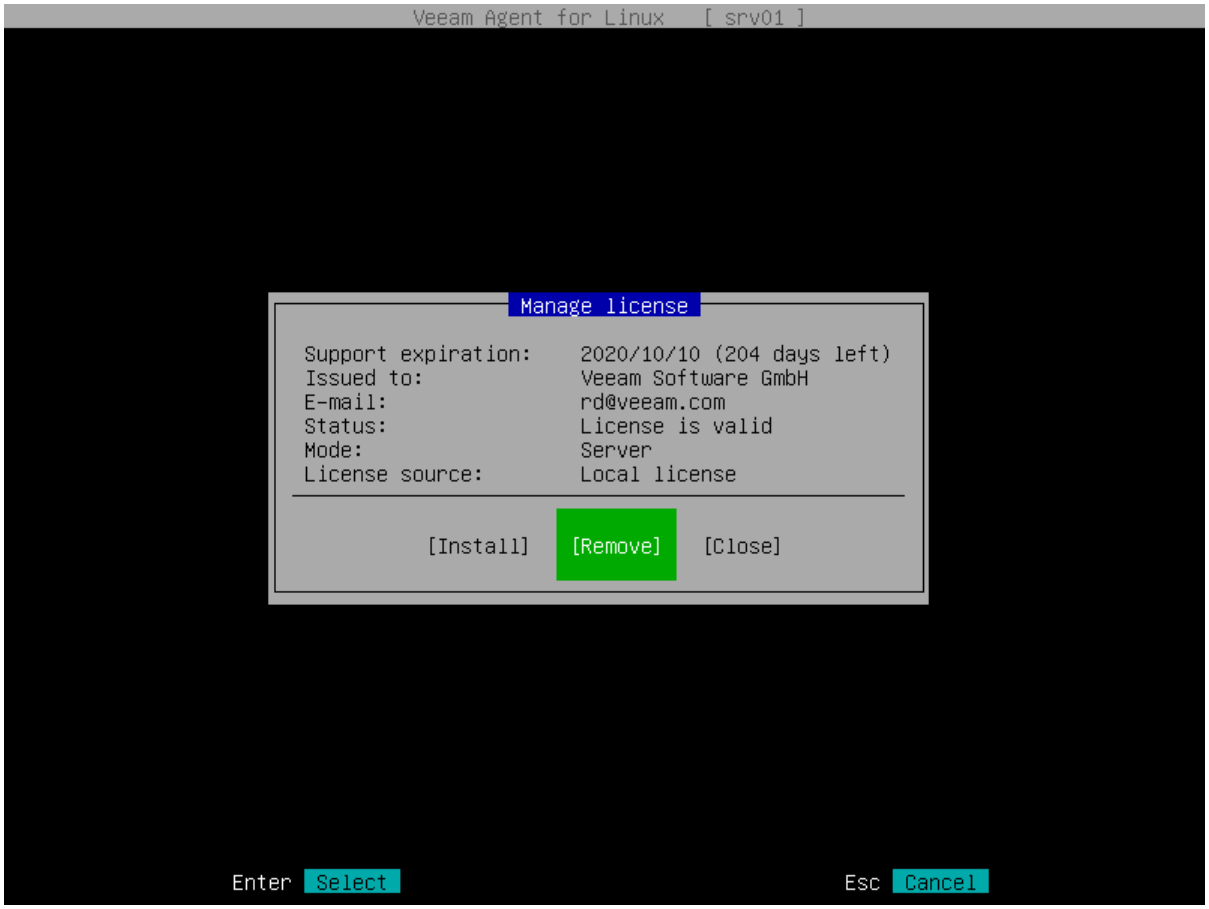
You can remove the license if necessary. To remove a license:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the '**m**' key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press **Enter**.
4. In the **Manage license** window, press **Tab** to select the **Remove** button, then press **Enter**.
5. Veeam Agent for Linux will remove the license and display a window notifying that the license is successfully removed. Press **Enter** to finish the license removal process.

NOTE

After you remove the license, Veeam Agent for Linux will continue to operate in the Free edition. Consider the following:

- If Veeam Agent for Linux operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will fail.
- If pre-freeze and/or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will fail.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will fail.



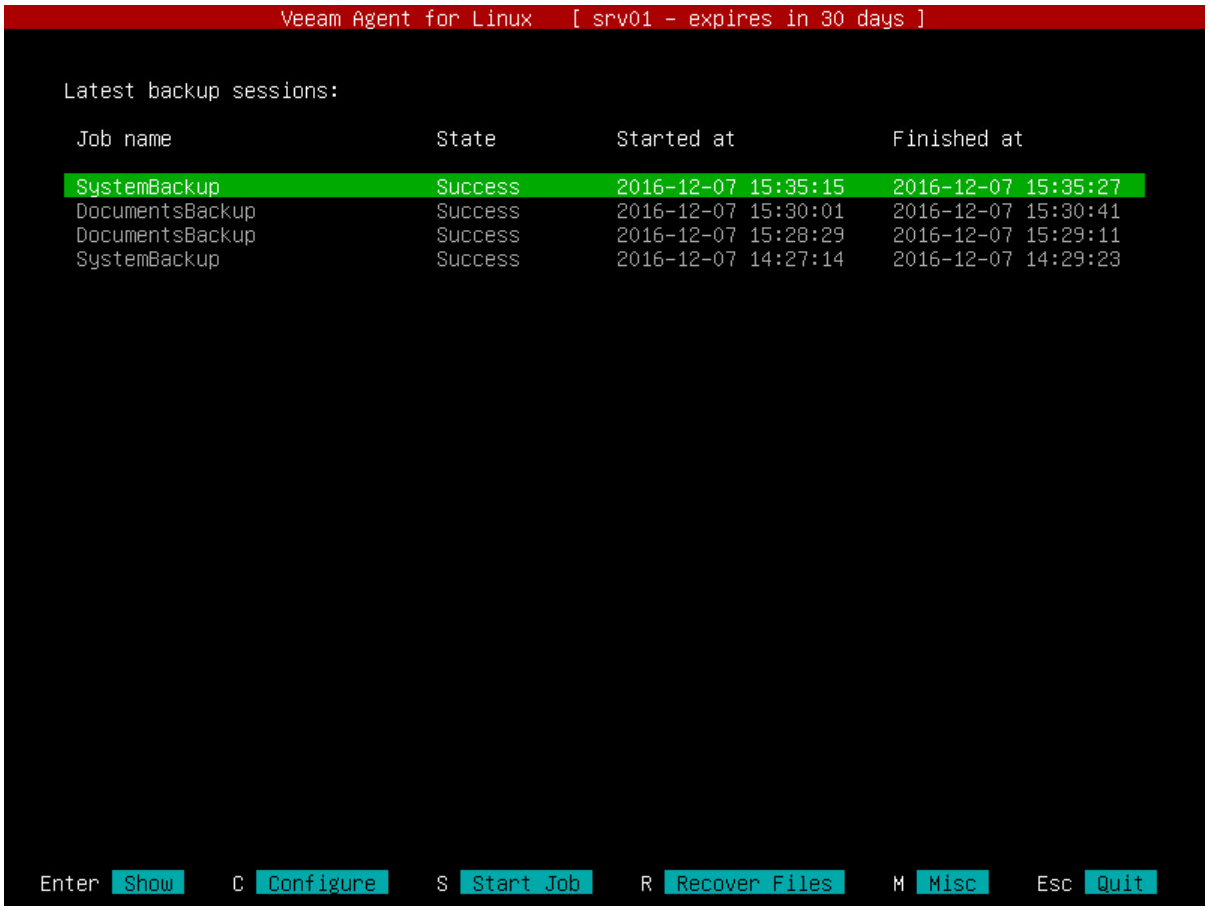
License Expiration

30 days before the license expiration date, Veeam Agent for Linux will display a warning at the top of the control panel. After the license expires, Veeam Agent for Linux will switch to the Free edition.

Consider the following:

- If Veeam Agent for Linux operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will be failing.
- If pre-freeze and/or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will be failing.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will be failing.

You can switch to the Free edition manually at any time if necessary. To learn more, see Removing License.



Managing License with Command Line Interface

You can use the Veeam Agent for Linux command line interface to perform the following operations with the license:

- Accept license agreements for the product itself and its third-party components.
- Install a license on the protected computer.
- View information about the license.
- Remove the license.

Accepting License Agreements

To work with Veeam Agent for Linux, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product. Until you accept license agreements, you can use the `veeamconfig` utility to run the following commands only:

- `veeamconfig agreement show`
- `veeamconfig help` (or `veeamconfig -h` or `veeamconfig --help`)
- `veeamconfig mode info`
- `veeamconfig mode reset`
- `veeamconfig version` (or `veeamconfig -v` or `veeamconfig --version`)
- `veeamconfig ui`

To accept license agreements, use the following command:

```
veeamconfig agreement accepteula && veeamconfig agreement acceptthirdpartylicen
ses
```

TIP

To check whether license agreements are accepted, use the following command: `veeamconfig agreement show`.

Installing License

To install a license, use the following command:

```
veeamconfig license install --path <path> --workstation
```

or

```
veeamconfig license install --path <path> --server
```

where:

- `<path>` — path to the license key file in the local file system of your computer.
- `workstation` or `server` — edition in which Veeam Agent will operate. To learn more about editions, see Product Editions.

Veeam Agent for Linux will install the license and display information about the license. You can also view this information later at any time. To learn more, see Viewing License Information.

For example:

```
user@srv01:~$ veeamconfig license install --path /home/user/veeam/license/veeam
.lic --server
License was installed successfully.
License information:
 License source: Local license
 Mode: Server
 Support expiration: 2019/09/20 (649 days left)
 Status: License is valid.
 Issued to: TechCompany
 E-mail: administrators@tech.com
```

TIP

You can also install a license using the Veeam Agent control panel. To learn more, see Installing License.

Viewing License Information

You can view information about the installed license. Use the following command:

```
veeamconfig license show
```

Veeam Agent for Linux will display information about the license. For example:

```
user@srv01:~$ veeamconfig license show
License information:
 License source: Local license
 Mode: Server
 Support expiration: 2019/09/20 (649 days left)
 Status: License is valid.
 Issued to: TechCompany
 E-mail: administrators@tech.com
```

Removing License

You can remove a license with the following command:

```
veeamconfig license remove
```

After you remove the license, Veeam Agent for Linux will continue to operate in the Free edition. Consider the following:

- If Veeam Agent operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will be failing.
- If pre-freeze and/or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will be failing.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will be failing.

Performing Backup

You can back up your data to protect the entire computer image, individual volumes or folders and files on your computer. To back up your data, you must configure a backup job. Depending on the product edition, Veeam Agent lets you configure one or several backup jobs targeted at the same or different backup repositories.

You can configure a backup job that will automatically back up your data by the defined schedule. You can also start a backup job manually at any time.

Creating Custom Veeam Recovery Media

In addition to the generic Veeam Recovery Media that is available for download at the Veeam website, you can create a custom Veeam Recovery Media. This option may be helpful if your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media. When you create a custom Veeam Recovery Media, Veeam Agent copies the Linux kernel running on your computer with its currently loaded modules and includes them into the custom recovery image.

Keep in mind that the custom recovery image comprises an unsigned Linux kernel. As a result, you cannot use it for UEFI systems with enabled Secure Boot.

You can create a custom Veeam Recovery Media in one of the following ways:

- With the Veeam Agent control panel. You can perform this operation at the process of initial product setup at the Recovery ISO step of the initial setup wizard.
- With the Veeam Agent command line interface. You can perform this operation at any time you need.

  If you create a custom Veeam Recovery Media using the command line interface, you can also specify a directory that contains additional drivers that you want to include in the recovery media. Veeam Agent will copy the content of the specified directory and add it to the root directory of the custom Veeam Recovery Media ISO file.

Before you create custom Veeam Recovery Media, check the following prerequisites:

- The Linux system must have the `genisoimage` package installed. For openSUSE and SLES 15 SP0 - SP3 distributions, the Linux system must have the `mkisofs` package installed.
- The Linux system must have the `mksquashfs` and `unsquashfs` utilities installed.
- For custom Veeam Recovery Media with EFI support, the Linux system must have the following packages installed:
  - `xorriso`
  - `isolinux` (or `syslinux`, if the software package repository of your Linux distribution lacks the `isolinux` package)

To create custom Veeam Recovery Media, use the following command:

```
veeamconfig patchiso --input <input_path> --output <output_path> --copy <additi
onal_path>
```

or

```
veeamconfig patchiso --efi --input <input_path> --output <output_path> --copy <
additional_path>
```

where:

- `<input_path>` — path to the ISO file of the generic Veeam Recovery Media.
- `<output_path>` — path to the resulting ISO file of the custom Veeam Recovery Media.
- `<additional_path>` — path to a directory with additional drivers that you want to include in the Veeam Recovery Media.

- `--efi` — option that defines whether custom Veeam Recovery Media should be able to boot on EFI-based systems. Without this option, custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

For example:

```
$ veeamconfig patchiso --input /mnt/veeam/iso/veeam-recovery-amd64-5.0.0.iso --
output /mnt/veeam/iso/veeam-recovery-media-srv01.iso --efi
```

Creating Backup Jobs

You can choose one of the following backup modes:

- Backup of an entire computer image
- Backup of specific computer volumes, for example, a system volume or secondary volume
- Backup of individual files and folders

[For Server Edition] You can configure one or several backup jobs to back up your data. Configuring several backup jobs may be useful in the following situations:

- You can configure separate backup jobs for volume-level backup and file-level backup.
- You can configure backup jobs targeted at different backup repositories to keep several copies of your backed-up data at different locations.
- You can configure several backup jobs and define individual schedule for every job to back up necessary data at the desired time.

With Veeam Agent, you can configure the backup job in one of the following ways:

- With the Backup Job wizard
- With the command line interface

Creating Backup Job with Backup Job Wizard

You can configure volume-level and file-level backup jobs with the Backup Job wizard.

Before You Begin

Before you configure the backup job, check the following prerequisites:

- The target location where you plan to store backup files must have enough free space.
- When you configure the backup job with the Backup Job wizard, Veeam Agent creates the job with default advanced settings: compression level and data block size. To specify these parameters explicitly, you should create a backup job with the command line interface.
- [For Veeam Backup & Replication repository targets] You can store created backups in a backup repository only if the backup server runs Veeam Backup & Replication 11 or later.
- [For Veeam Backup & Replication repository targets] If you plan to use a Veeam Backup & Replication repository as a target for backups, you must pre-configure user access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.

Backup has the following limitations:

- You cannot save the backup of entire computer on the local computer disk. Use an external hard drive or USB drive, network shared folder or backup repository as a target location.
- Veeam Agent does not back up data to which symbolic links are targeted. It only backs up the path information that the symbolic links contain. After restore, identical symbolic links are created in the restore destination.
- Keep in mind that Veeam Agent stops running the backup job after 21 days (504 hours).

Navigating Backup Job Wizard

The Backup Job wizard window comprises the following areas:

- The navigation pane, located on the left of the window, displays the list of wizard steps and currently selected step of the wizard
- The working area displays controls relating to a specific step of the wizard.
- The buttons area, located at the bottom of the window, displays buttons that you can use to switch between steps of the wizard (**Previous** and **Next**) and close the wizard (**Cancel** and **Finish**).

In the Backup Job wizard, the use of a mouse is not supported. To navigate the Backup Job wizard and associated dialog windows, you can use the following keys:

- **Tab** — to switch between displayed controls in the working area and buttons in the buttons area. The currently selected control or button is highlighted with a green color.
- **Up** and **Down** — to switch between items in a scrollable list.
- **Space** — to select the necessary item in a list. The selected item's mark may vary in different steps of the wizard.
- **Enter** — to proceed to the next step of the wizard or to open a directory.
- **Backspace** — to return to the previous step of a wizard.
- **Escape** — to cancel the backup job configuration and exit the wizard.

> **TIP**
>
> You can switch between steps of the Backup Job wizard in two ways. The easier and more comfortable way is to use the **Enter** key to proceed to the next step and **Backspace** key to return to the previous step of the wizard. You can also use the **Tab** key to select the **Next** or **Previous** button in the buttons area and then press **Enter** to switch to the next or previous step of the wizard respectively.

Step 1. Launch Backup Job Wizard

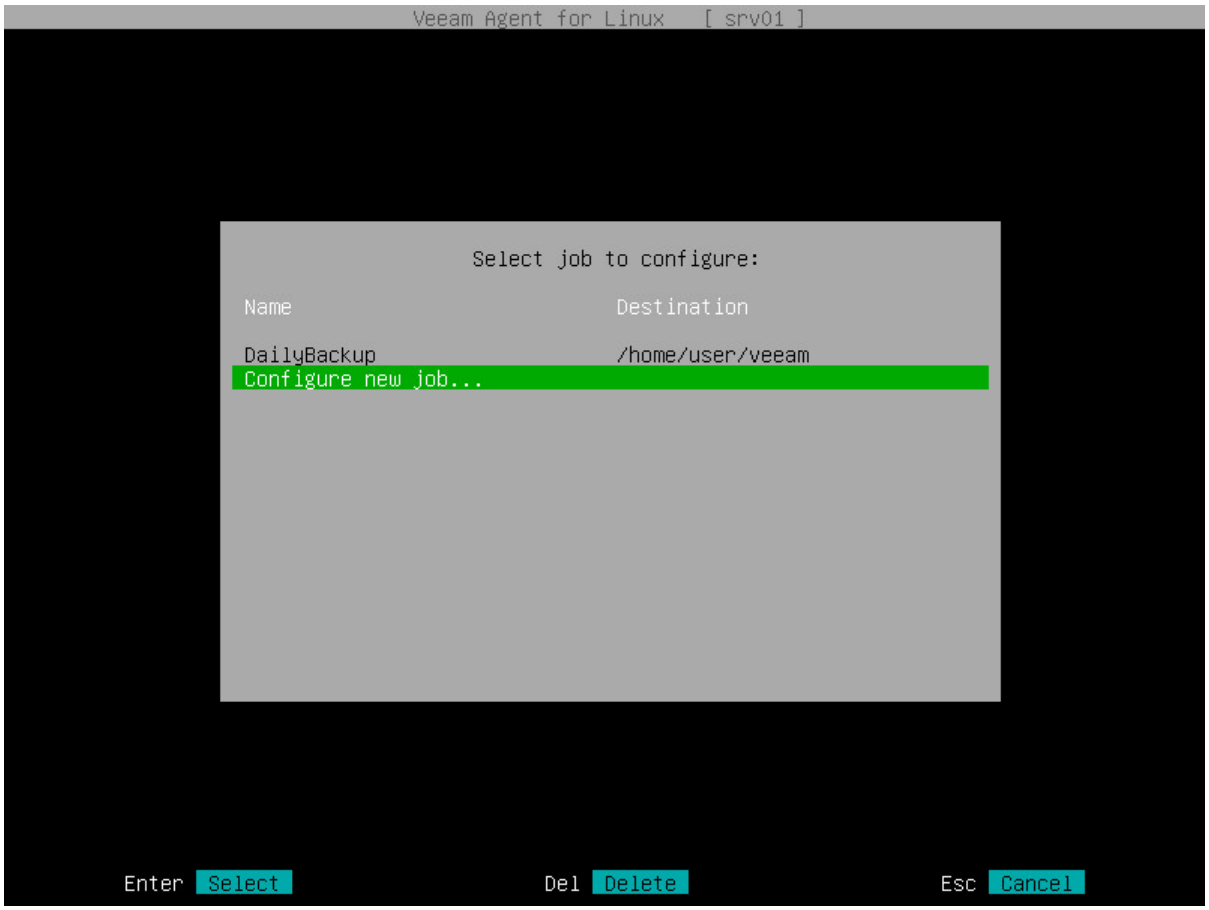
To launch the **Backup Job** wizard, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command:
2. If you have not configured any backup jobs yet, Veeam Agent will display a welcome screen. Press the '**c**' key to proceed to the Backup Job wizard and configure the backup job.
3. If you have already configured and performed a backup job, Veeam Agent will display the list of backup job sessions. When you press '**c**' to launch the Backup Job wizard, Veeam Agent will display a list of configured backup jobs. To configure a new backup job, select the **Configure new job** option and press **Enter**.

> **NOTE**
>
> The **Configure new job** option is not available if Veeam Agent for Linux operates in the Free edition and you have already configured one backup job.

To edit settings of a backup job that you have already configured, select the job in the list and press **Enter**. To learn more, see Editing Backup Job Settings.

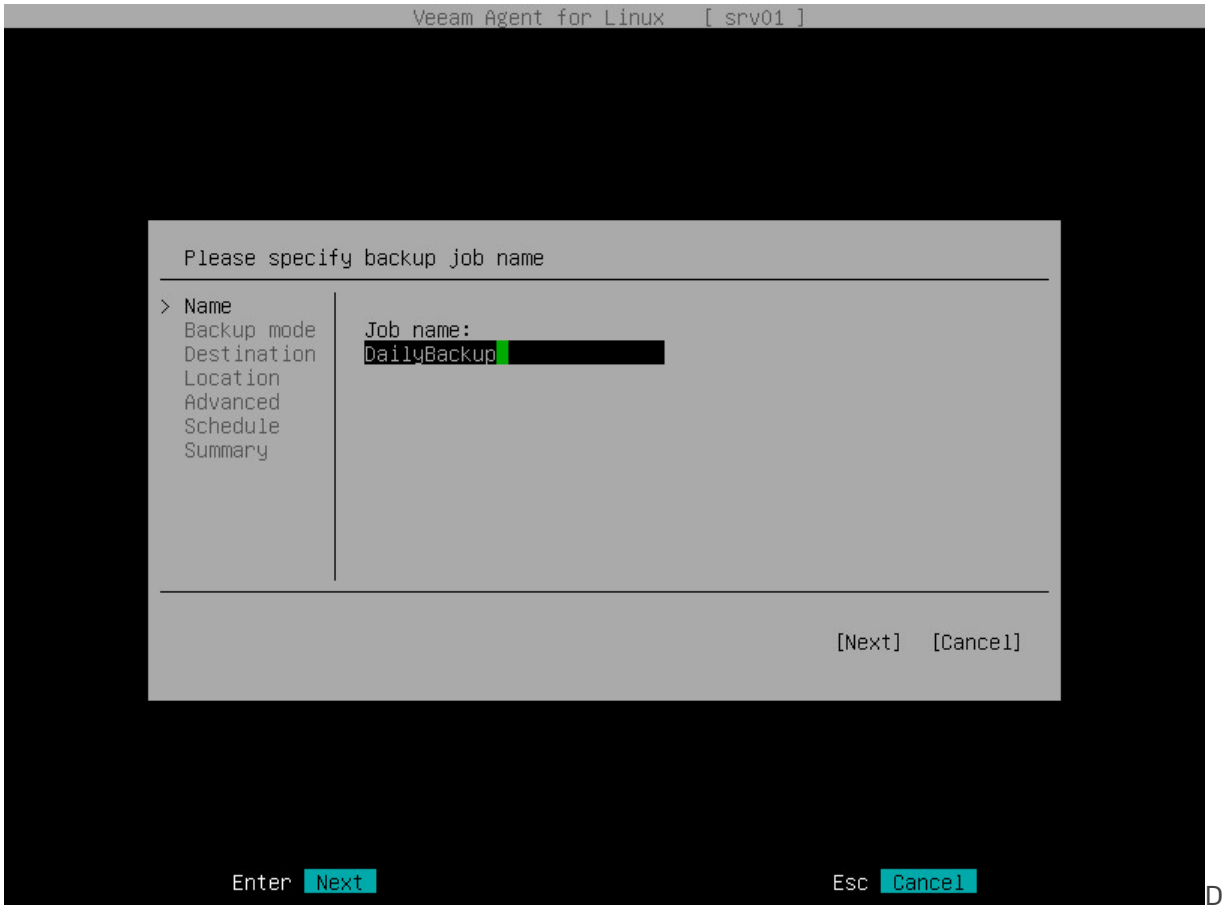
If you have decided not to create a backup job, press **Escape** to close the list of backup jobs and return to the welcome screen. After that, you can press **Escape** once again to return to the command line interface.



Step 2. Specify Backup Job Name

At the **Name** step of the wizard, in the **Job name** field, type the name for the backup job and press **Enter**.

> **TIP**
>
> To proceed to the next step of the wizard, you can also select the **Next** button with the **Tab** key and then press **Enter**.



Step 3. Select Backup Mode

At the **Backup mode** step of the wizard, select the mode in which you want to create a backup:

1. Select the necessary backup mode. You can select one of the following options:
   - **Entire machine** — select this option if you want to create a backup of the entire computer image. When you restore data from such backup, you will be able to recover the entire computer image as well as data on specific computer volumes: files, folders, application data and so on. With this option selected, you will pass to the Destination step of the wizard.
   - **Volume level backup** — select this option if you want to create a backup of specific computer volumes, for example, the system volume. When you restore data from such backup, you will be able to recover data on these volumes only: files, folders, application data and so on. With this option selected, you will pass to the Volumes step of the wizard.
   - **File level backup** — select this option if you want to create a backup of individual directories on your computer. With this option selected, you will pass to the Files step of the wizard.

2. [For file-level backup] If you want to perform backup in the snapshot-less mode, select **Disable snapshot**. With this option selected, Veeam Agent will not create a snapshot of the backed-up volumes during backup. This allows Veeam Agent to back up data residing in file systems that are not supported for snapshot-based backup with Veeam Agent. To learn more, see Snapshot-Less File-Level Backup.

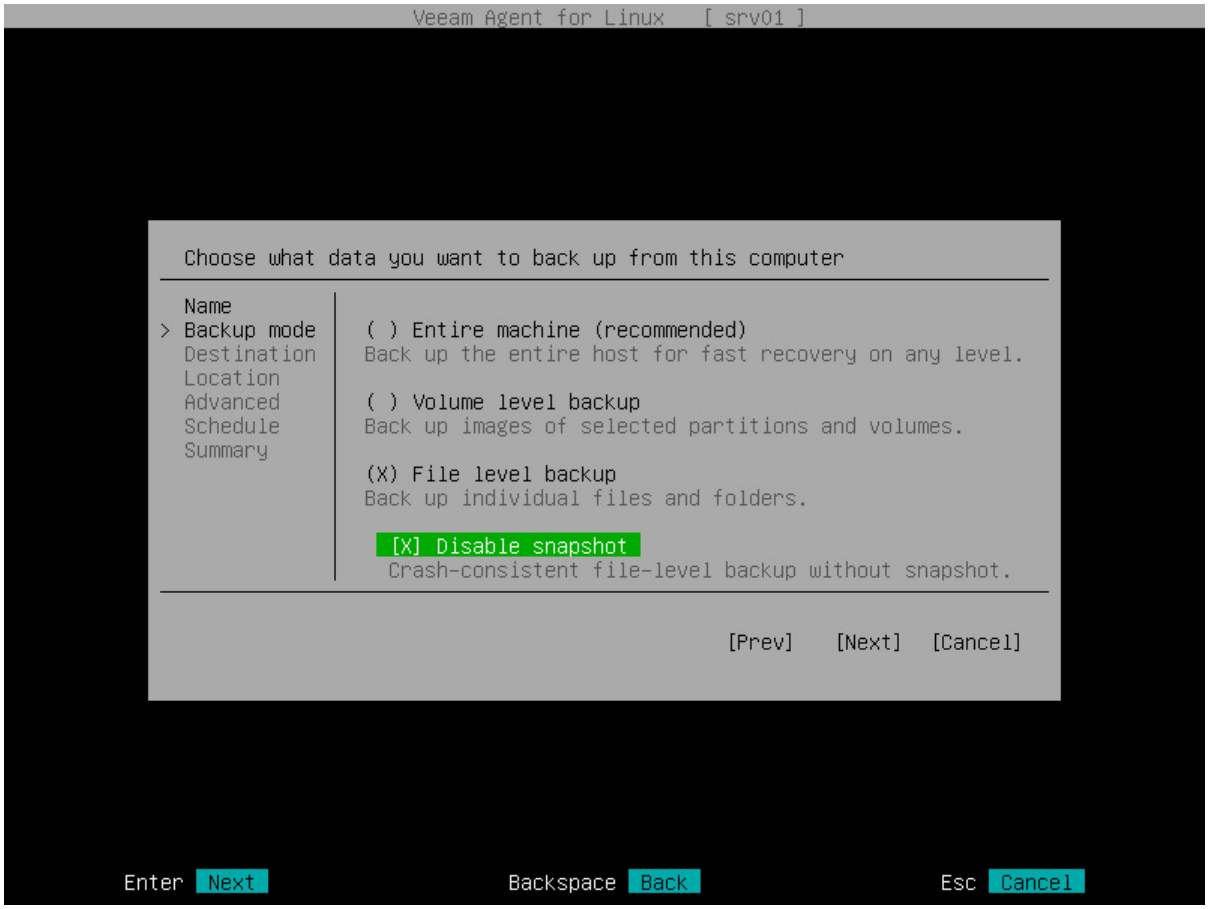
IMPORTANT

Consider the following:

- [For entire machine backup] Certain limitations for EMC PowerPath configuration apply. To learn more, see this Veeam KB article.
- [For volume-level backup] Volume-level backup job relies on a device name in the `/dev` directory. Device names in the `/dev` directory (e.g. `/dev/md-127`, `/dev/dm-1`) must stay persistent for backed-up volumes. Otherwise, the job will back up the wrong volume.
- [For file-level backup] If the backed-up file system has a complex folder structure with many hierarchy levels, during incremental backup, the inbound network traffic on the Veeam Agent computer may exceed by far the outbound traffic. Significant amount of data can be transferred to the Veeam Agent computer from the target backup location even if few files are changed since the previous job session.

TIP

File-level backup is typically slower than volume-level backup. If you plan to back up all folders with files on a specific volume, it is recommended that you configure volume-level backup instead of file-level backup.



Step 4. Specify Backup Scope Settings

Specify backup scope for the backup job:

- Select volumes to back up — if you have selected the **Volume level backup** option at the Backup Mode step of the wizard.
- Select folders to back up — if you have selected the **File level backup** option at the Backup Mode step of the wizard.

Selecting Volumes to Back Up

The **Volumes** step of the wizard is available if you have chosen to create a volume-level backup.

At this step of the wizard, you must specify the backup scope — define what volumes you want to include in the backup. Veeam Agent lets you include the following types of objects in the volume-level backup:

- Block devices (entire disks and individual volumes)
- Mount points
- LVM logical volumes and volume groups
- Btrfs storage pools and subvolumes

Selecting Devices

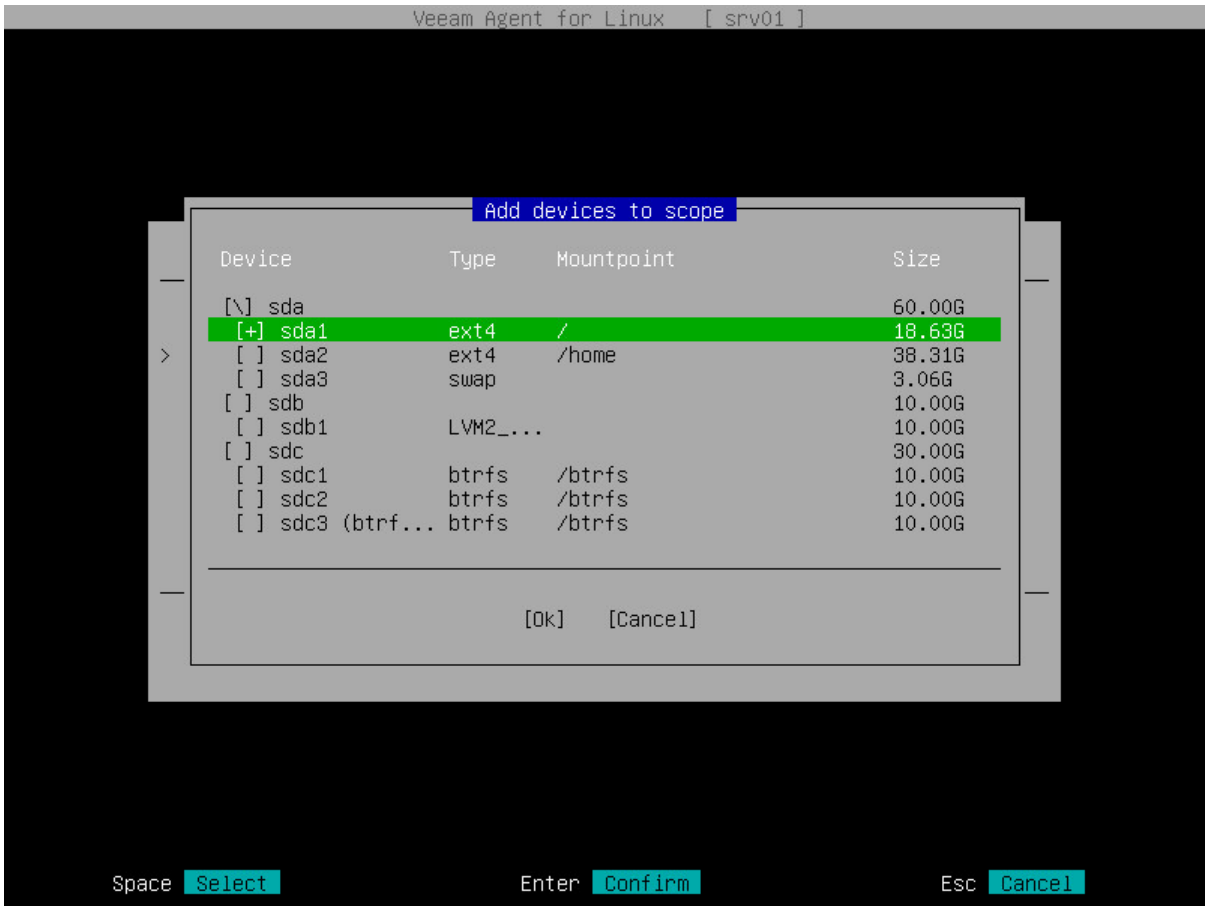
To add a block device to the backup scope, do the following:

1. At the **Volume** step of the wizard, make sure that the **Device** option is selected and press **Enter**.
2. In the **Add devices to scope** window, select individual volumes or entire computer disks that you want to include in the backup and press **Enter**.
   - To include individual volumes of your computer in the backup, select block devices that represent volumes that you want to back up, for example: *sda1* and/or *sda6*.
   - To include all volumes on a computer disk in the backup, select block devices that represent disks whose volumes you want to back up, for example: *sda* and/or *sdb*. All volumes on the selected disk will be automatically selected, too.

   To navigate the list of volumes and select the necessary items, use **Up**, **Down** and **Space** keys. To learn more, see Navigating Backup Job Wizard.

If you have created several system partitions, for example, a separate partition for the `/boot` directory, you should remember to include all of these partitions in the backup. Otherwise, Veeam Agent does not guarantee that the OS will boot properly when you attempt to recover from such backup.

> **NOTE**
>
> If you include a block device in the backup, and this block device is a physical volume assigned to an LVM volume group, Veeam Agent will include the whole LVM volume group in the backup.

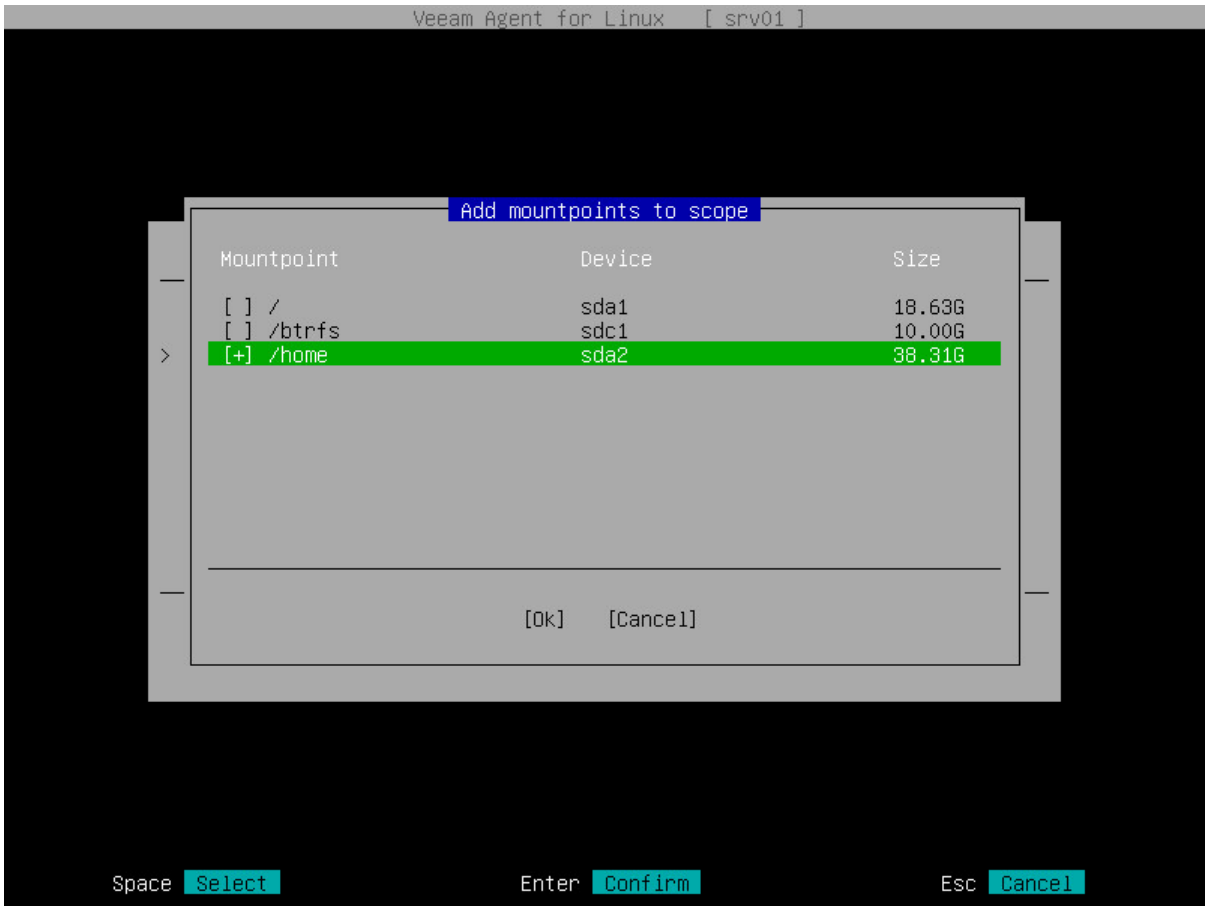


Selecting Mount Points

To add a mount point to the backup scope, do the following:

1. At the **Volume** step of the wizard, select the **Mountpoint** option and press **Enter**.
2. In the **Add mountpoints to scope** window, select mount points that you want to include in the backup and press **Enter**

   To navigate the list of mount points and select the necessary mount points, use **Up**, **Down** and **Space** keys. To learn more, see Navigating Backup Job Wizard.



Selecting LVM Volumes

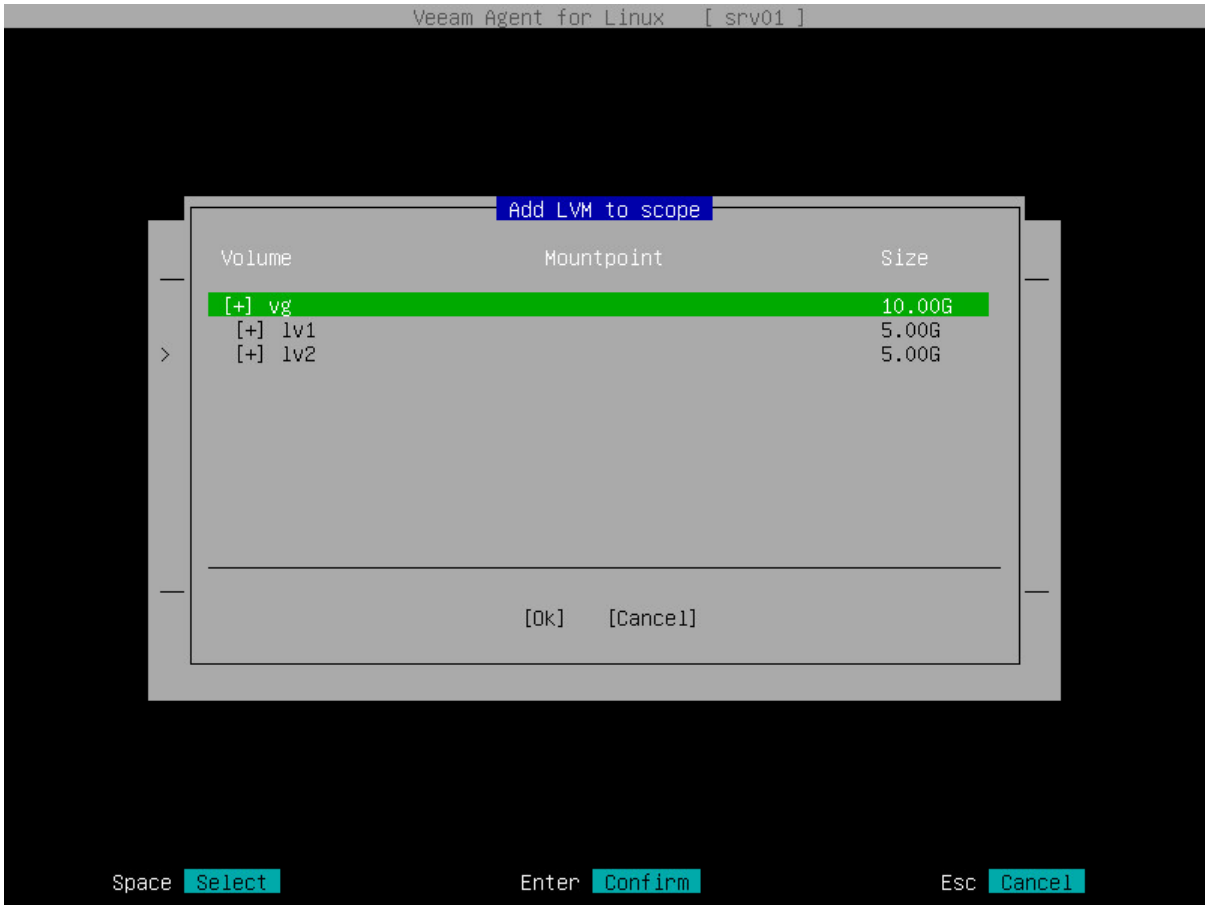
To add an LVM logical volume or volume group to the backup scope, do the following:

1. At the **Volume** step of the wizard, select the **LVM** option and press **Enter**.
2. In the **Add LVM to scope** window, select LVM logical volumes or volume groups that you want to include in the backup and press **Enter**.

   To navigate the list of LVM volumes and select the necessary items, use **Up**, **Down** and **Space** keys. To learn more, see Navigating Backup Job Wizard.

   If you include an LVM volume group in the backup, all LVM logical volumes in the selected volume group will be automatically selected, too.

> **NOTE**
>
> Veeam Agent does not back up LVM snapshots.



Selecting Btrfs Volumes

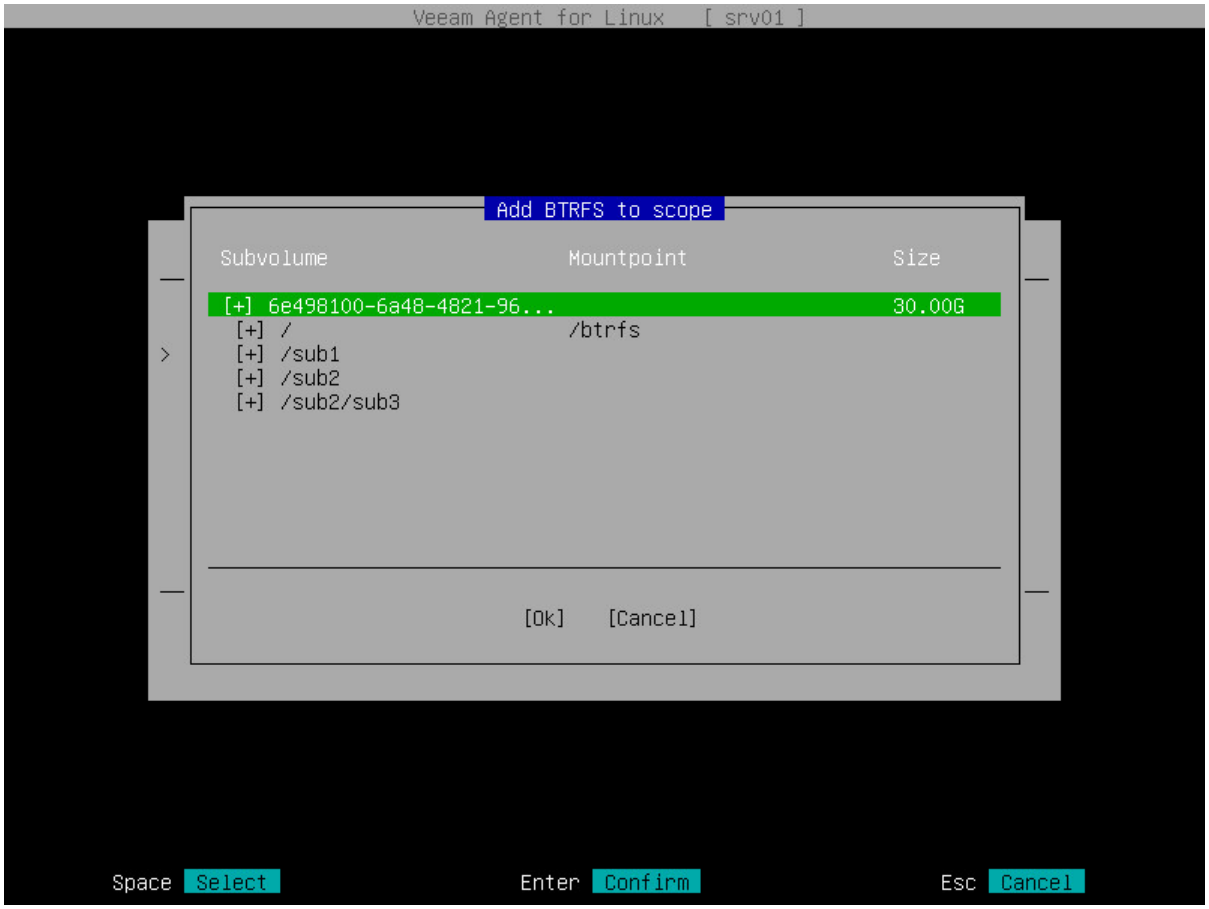
To add a Btrfs storage pool or subvolume to the backup scope, do the following:

1. At the **Volume** step of the wizard, select the **BTRFS** option and press **Enter**.
2. In the **Add BTRFS to scope** window, select Btrfs storage pools or subvolumes that you want to include in the backup and press **Enter**.

   To navigate the list of Btrfs pools and subvolumes and select the necessary items, use **Up**, **Down** and **Space** keys. To learn more, see Navigating Backup Job Wizard.

   Veeam Agent identifies Btrfs storage pools by UUIDs. If you include a Btrfs pool in the backup, all Btrfs subvolumes in the selected pool will be automatically selected, too.

> **NOTE**
>
> You cannot add read-only Btrfs snapshots to the backup scope.



Selecting Files and Directories to Back Up

The **Files** step of the wizard is available if you have chosen to create a file-level backup.

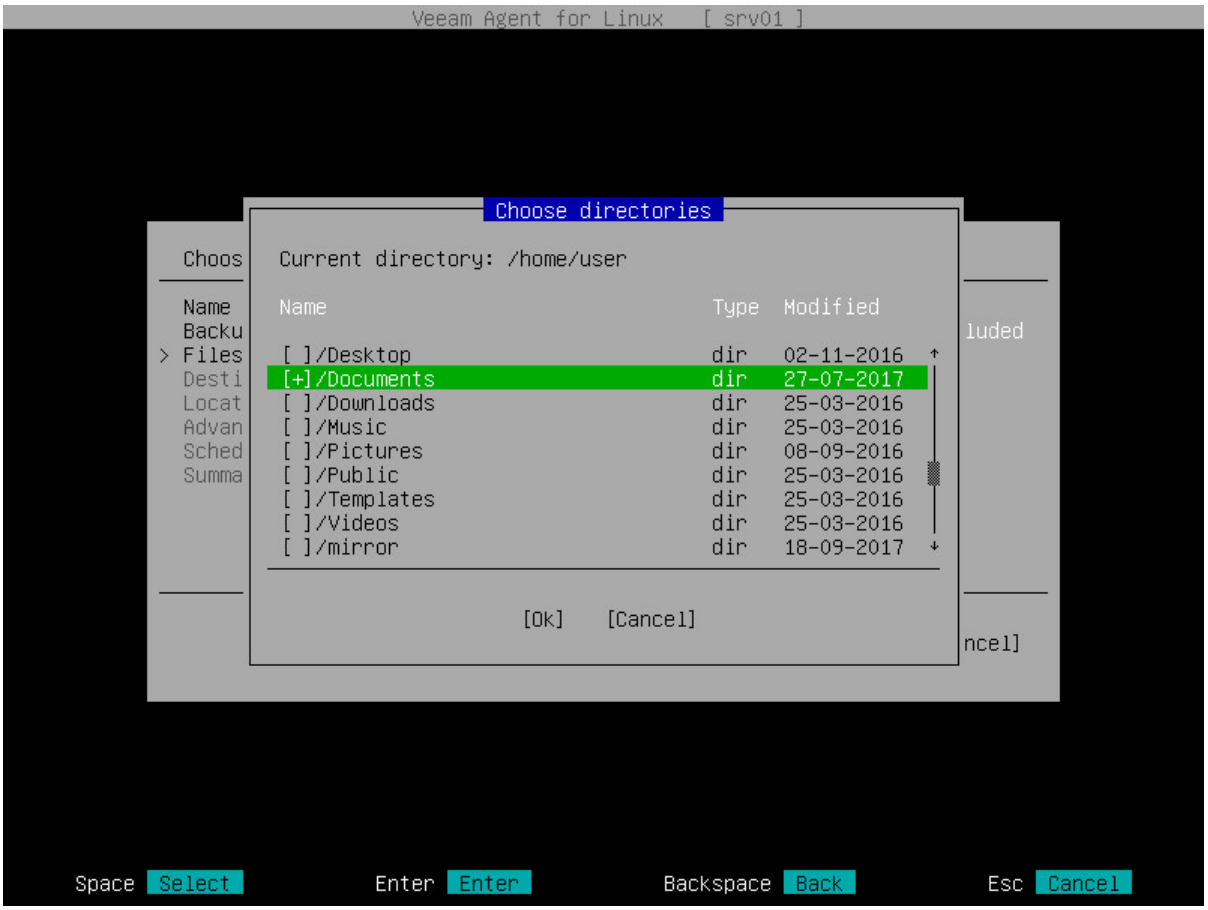
At this step of the wizard, you must specify the backup scope — define what directories with files you want to include in the backup.

In the file-level backup mode, you must include in the backup at least one directory. If you do not want to back up some subdirectories of the specified directory, you can exclude these directories from the backup.

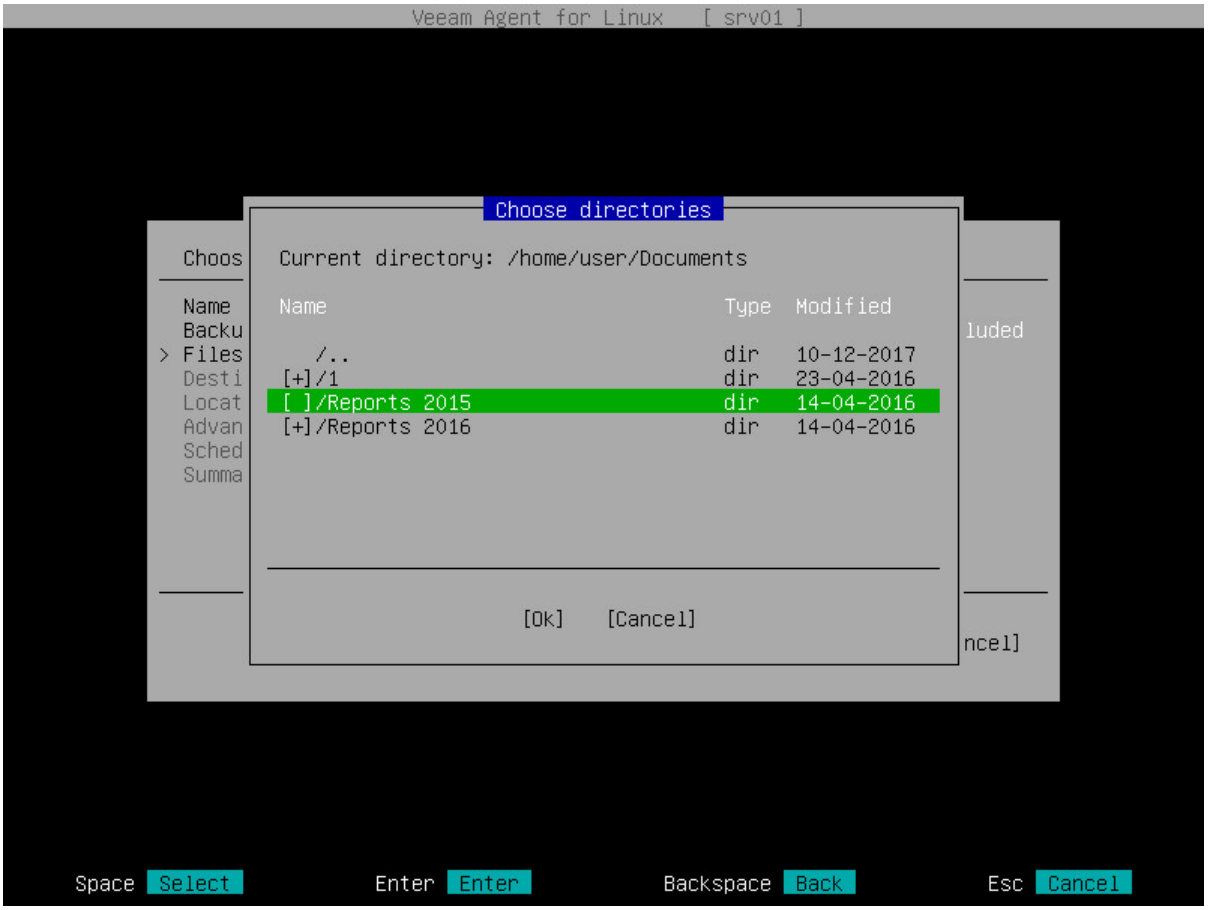
You can also include or exclude files of a specific type in/from the backup. You can specify file names explicitly or use UNIX wildcard characters to define file name masks. Veeam Agent will apply the specified file name masks to files in directories that are included in the backup.

To specify the backup scope:

1. At the **Files** step of the wizard, make sure that the **Add directories** option is selected and press **Enter**.
2. In the **Choose directories** window, select one or several directories that you want to include in the file-level backup.
   - To navigate the list of directories, use the **Up** and **Down** keys.
   - To browse for subdirectories, navigate to the necessary directory and press **Enter**.
   - To include a directory in the backup, navigate to the necessary directory and press **Space**. The included directory will be marked with the '+' character. All subdirectories of the selected directory will be included in the backup too.



3. Specify directories that you want to exclude from the file-level backup. To exclude a directory:
   a. Browse for subdirectories of a directory that you have included in the backup.
   b. Navigate to the directory that you want to exclude from the backup and press **Space**. The excluded directory will not be marked with the '+' character.



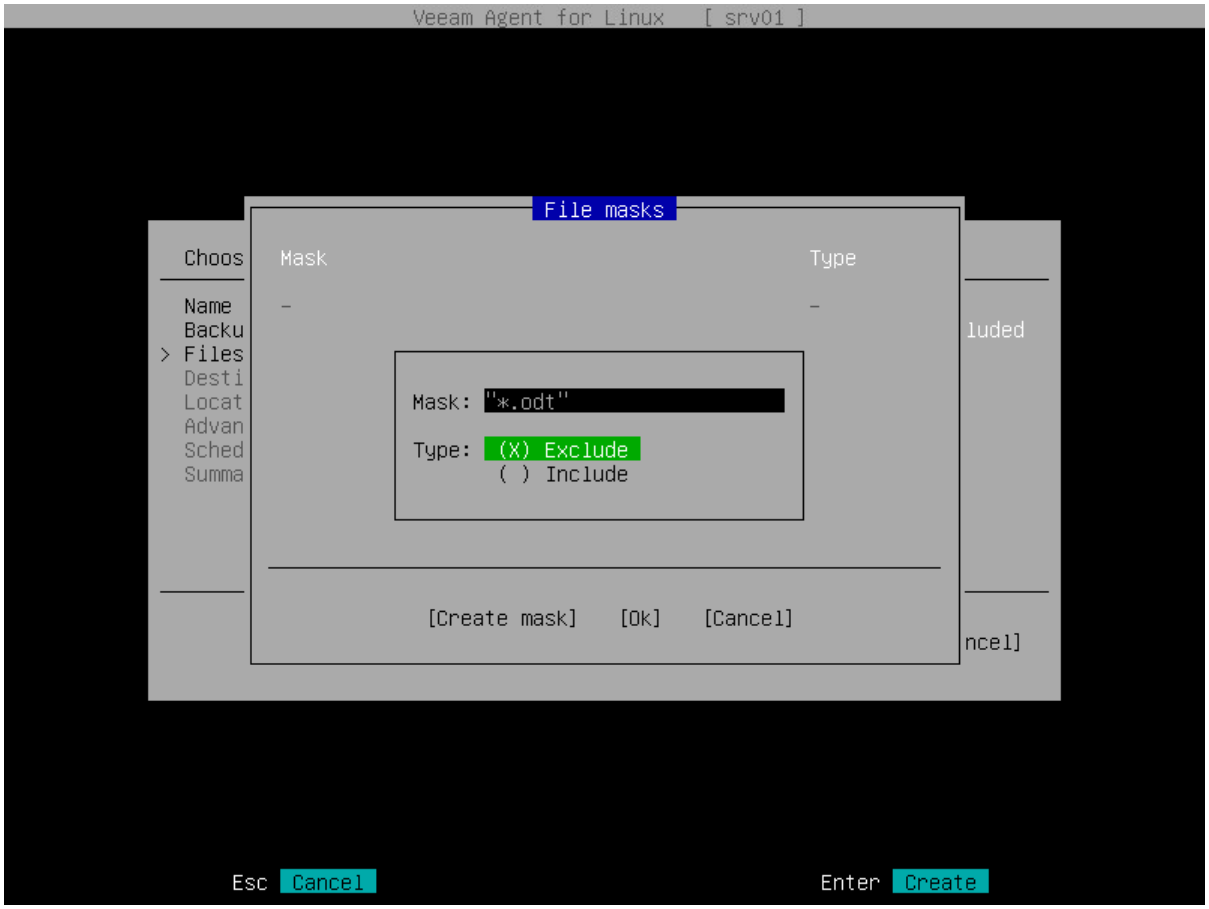
4. Switch to the **OK** button and press **Enter**. Veeam Agent will display a list of paths to the selected directories and the number of excluded subdirectories for each directory in the list.
5. Specify file name masks for files that you want to include or exclude in/from the backup:
   a. Select the **File Masks** option with the **Tab** key and press **Enter**.
   b. In the **File masks** window, make sure that the **Create Mask** button is selected and press **Enter**.
   c. In the **Mask** field, enter the file name mask, for example, `report.pdf`, `*filename*` or `*.odt`.

      Keep in mind that you must specify all names with masks in double quotation marks ("").
   d. In the **Type** field, select one of the following options:
      - **Exclude** — if you do not want to back up files whose names match the specified mask. Veeam Agent will back up all files in the directories selected for backup except for such files.
      - **Include** — if you want to back up files whose names match the specified mask. Veeam Agent will create a backup only for such files in the directories selected for backup.

You can use a combination of include and exclude masks. Keep in mind that exclude masks have a higher priority than include masks. For example, you can specify masks in the following way:

- Include mask: `report*.*`
- Exclude mask: `*.odt`

Veeam Agent will include in the backup all files whose name begins with `report` except for the files of the ODT format.



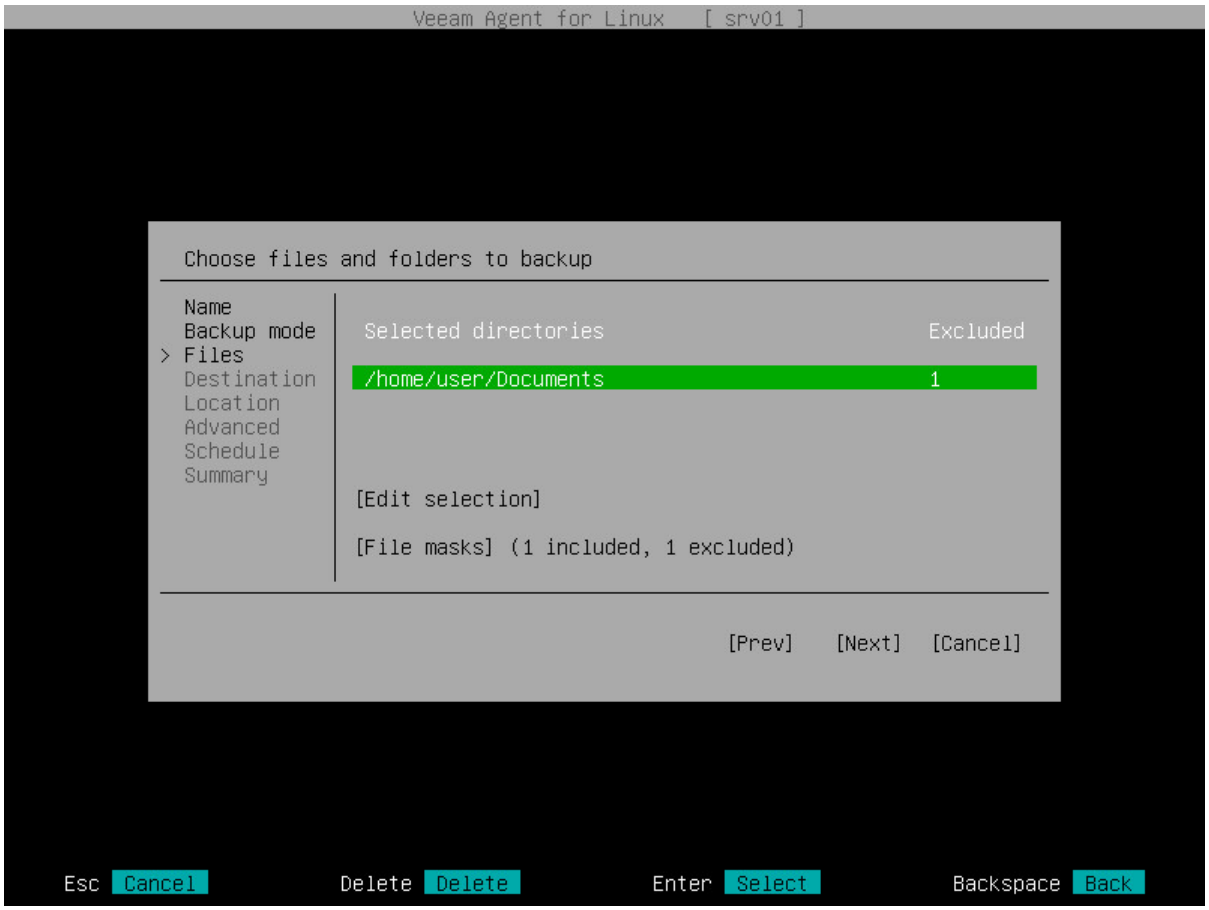
e. Press **Enter**. Veeam Agent will display in the **File masks** window the specified file mask and its type: *Include* or *Exclude*.

f. Repeat steps 'b'–'e' for each mask that you want to specify.

g. After you specify all file masks, switch to the **OK** button and press **Enter**.

TIP

To remove a file name mask, in the **File masks** window, select the necessary mask and press **Delete**.



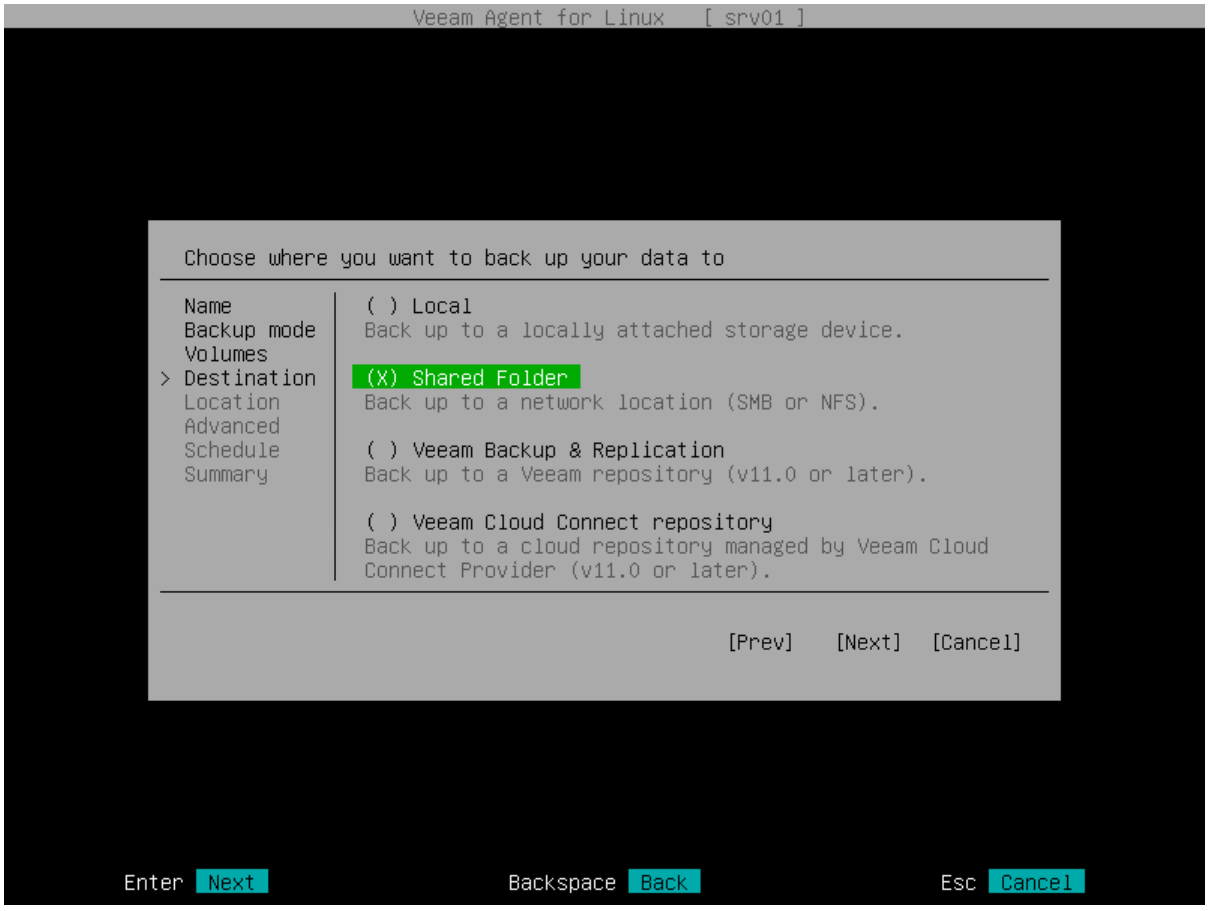
Step 5. Select Backup Destination

At the **Destination** step of the wizard, select a target location for the created backup.

You can select one of the following options:

- **Local** — select this option if you want to save the backup on a removable storage device attached to the computer or on a local computer drive. With this option selected, you will pass to the Location step of the wizard.
- **Shared Folder** — select this option if you want to save the backup in a network shared folder. With this option selected, you will pass to the Network step of the wizard.
- **Veeam Backup & Replication** — select this option if you want to save the backup on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the Veeam step of the wizard.
- **Veeam Cloud Connect repository** — select this option if you want to create the backup on a cloud repository exposed to you by the Veeam Cloud Connect service provider. With this option selected, you will pass to the Service Provider step of the wizard.

It is recommended that you store backups in the external location like USB storage device or network shared folder. You can also keep your backup files on the separate non-system local drive.



Step 6. Specify Backup Storage Settings

Specify backup storage settings for the backup job:

- Local storage settings — if you have selected the **Local storage** option at the Destination step of the wizard.
- Shared folder settings — if you have selected the **Shared folder** option at the Destination step of the wizard.
- Veeam backup repository settings — if you have selected the **Veeam backup repository** option at the Destination step of the wizard.
- Veeam Cloud Connect repository settings — if you have selected the **Veeam Cloud Connect repository** option at the Destination step of the wizard.

> **NOTE**
>
> The **Veeam Cloud Connect repository** option is available if Veeam Agent operates in the Workstation or Server edition.

Local Storage Settings

The **Location** step of the wizard is available if you have selected the **Local** option at the Destination step of the wizard. Specify location for the backup file and retention policy for the backup job:

1. To specify location for the backup file, browse to the directory where backup files must be saved:
   a. Select the **Browse** option with the **Tab** key and press **Space** or **Enter**.
   b. In the **Choose backup location** window, select the necessary directory and press **Enter**.
   c. Repeat the step 'b' until a path to the directory in which you want to save backup files appears in the **Current directory** field.
   d. To create a new directory, switch to the **Create Dir** button, press **Enter**, then type a name for the new directory and press **Enter**.
   e. Switch to the **OK** button and press **Enter**. Veeam Agent will display the path to the specified directory in the **Location** field.

   Alternatively, you can type a path to the directory in which you want to save backup files in the **Location** field.

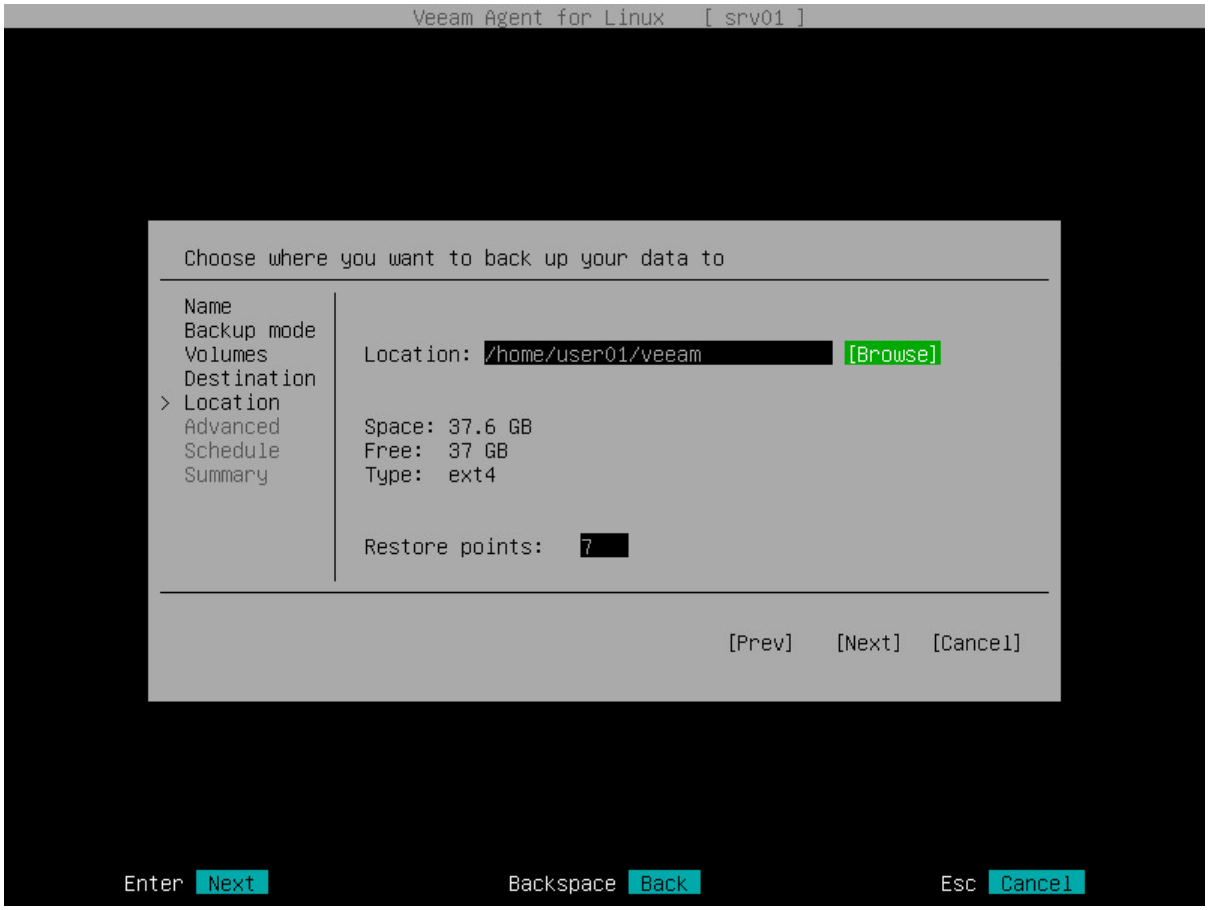
   After you specify location for the backup, Veeam Agent will display the following information on the volume where the directory selected for backup storage resides:

   - **Space** — total size of the volume on which the selected directory resides.
   - **Free** — free space on the volume where the selected directory resides.
   - **Type** — file system type of the volume on which the selected directory resides.
2. In the **Restore points** field, specify the number of backup files that you want to keep in the target location. By default, Veeam Agent keeps 7 latest backup files. When the number of restore points is exceeded, Veeam Agent for Linux will remove the earliest restore point from the backup chain.

IMPORTANT

Consider the following:

- The backup location must reside on a separate volume from a volume whose data you plan to back up.
- USB storage devices formatted as FAT32 do not allow storing files larger than 4 GB in size. For this reason, it is recommended that you do not use such USB storage devices as a backup target.



Shared Folder Settings

The **Network** step of the wizard is available if you have selected the **Shared Folder** option at the Destination step of the wizard.

To save backup files in a remote network location, Veeam Agent mounts to the local file system of your computer the network shared folder that you specify as a location for the backup. When you specify the network shared folder settings, Veeam Agent saves information about the network shared folder and its mount point in the database.

You do not need to mount the network shared folder in advance before every backup job run. Veeam Agent will do it automatically when the backup job is started manually or upon schedule.

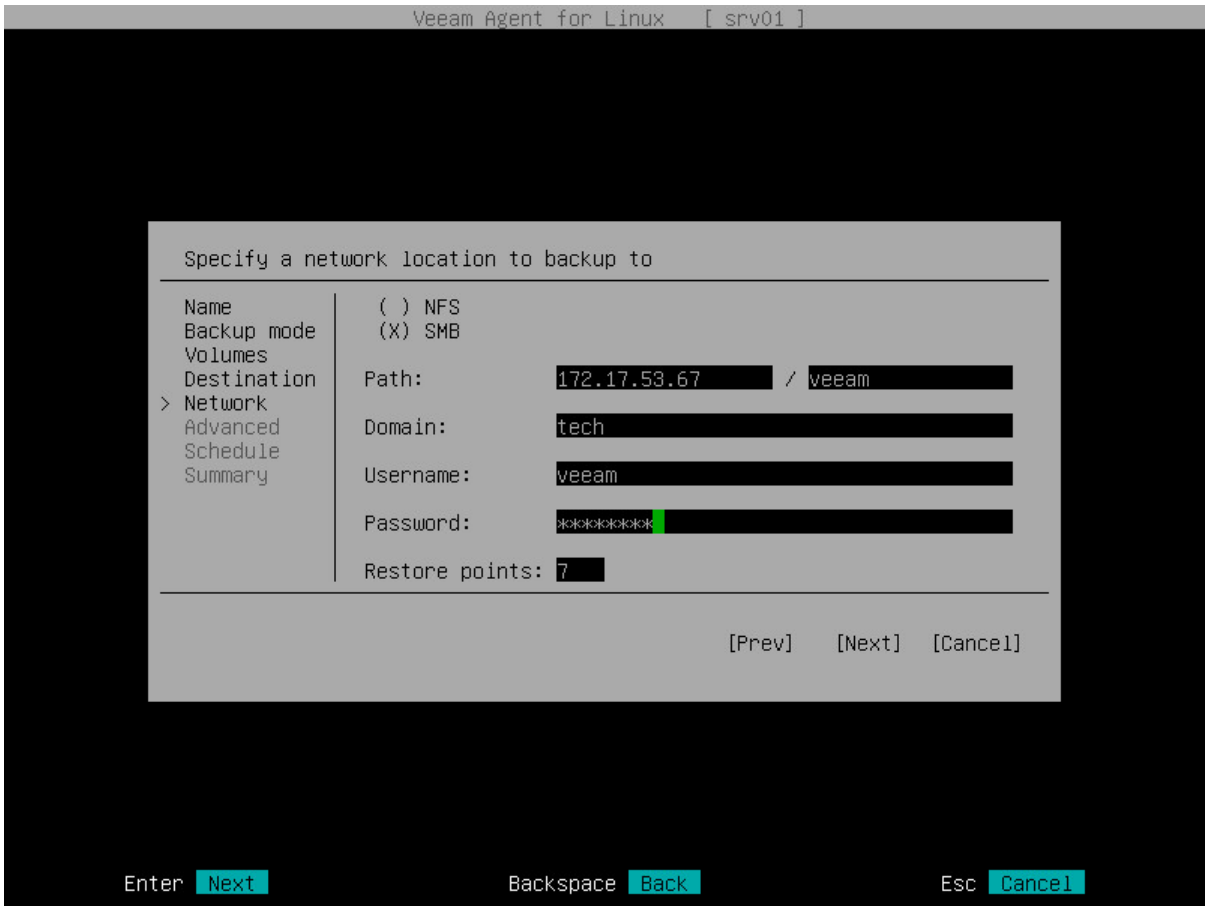
After the backup job completes, Veeam Agent will automatically unmount the network shared folder.

Specify shared folder settings:

1. Select the type of a network shared folder:
   - **NFS** — to connect to a network shared folder using the NFS protocol.
   - **SMB** — to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Path** field, specify the network shared folder name in the *SERVER/DIRECTORY* format: type an IP address or domain name of the server and the name of the network shared folder in which you want to store backup files.

   Every time the backup job starts, Veeam Agent will automatically mount the specified network shared folder to the `/tmp/veeam` directory in the computer file system. After the backup job completes, Veeam Agent will unmount the network shared folder.
3. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
4. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
5. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.

6. In the **Restore points** field, specify the number of backup files that you want to keep in the target location. By default, Veeam Agent keeps 7 latest backup files. When the number of restore points is exceeded, Veeam Agent will remove the earliest restore point from the backup chain.



Veeam Backup Repository Settings

If you have selected to store backup files on a Veeam Backup & Replication repository, specify settings to connect to the backup repository:

1. Specify backup server settings.
2. Select the Veeam backup repository.

Specifying Backup Server Settings

The **Veeam** step of the wizard is available if you have chosen to store backup files on a Veeam Backup & Replication repository.

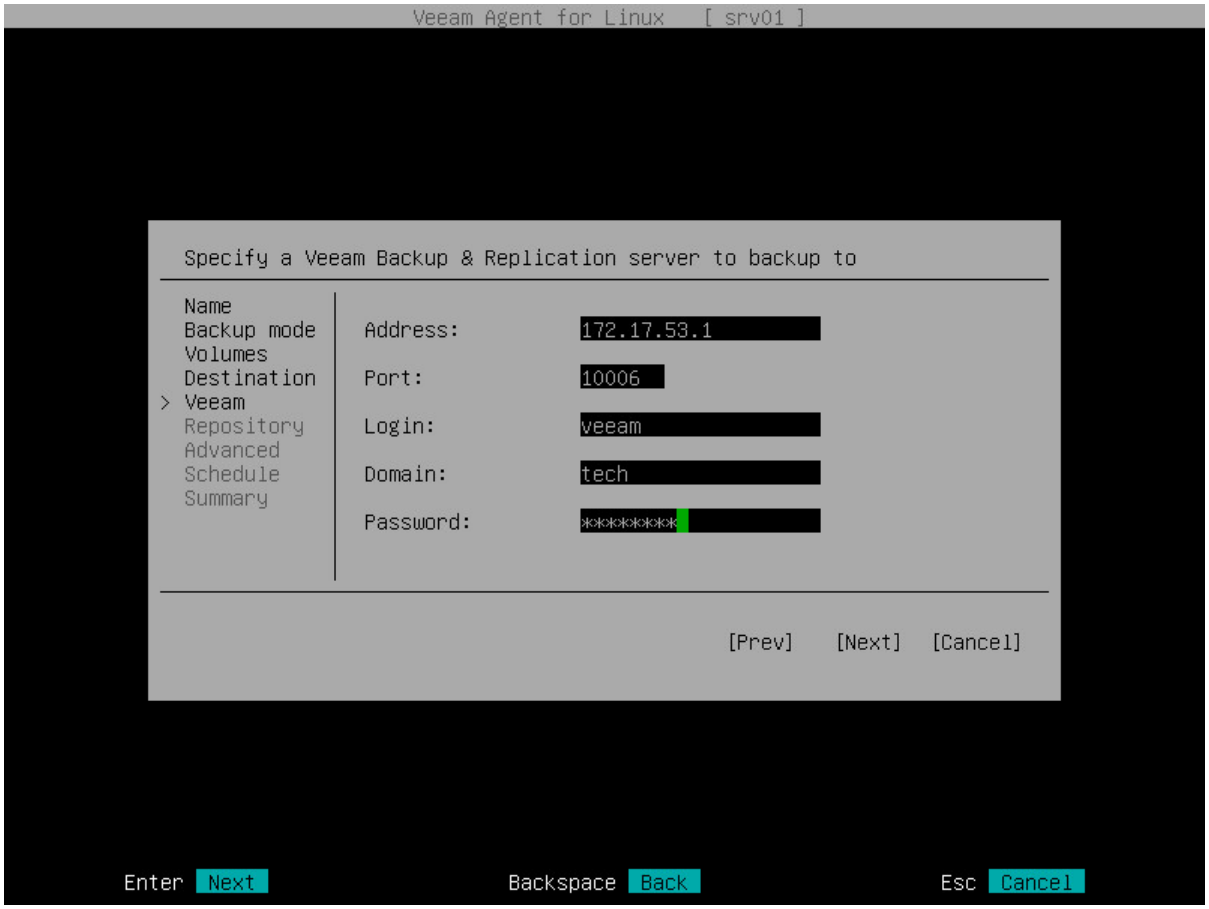
Specify settings for the Veeam backup server that manages the target backup repository:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
4. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.
5. In the **Password** field, type a password of the account that has access to the Veeam backup repository.

   Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see Setting Up User Permissions on Backup Repositories.

IMPORTANT

If you specify a DNS name of the Veeam backup server, make sure that the Veeam backup server name is resolved into IPv4 address on the machine where Veeam Agent is installed. The Veeam Backup Service in Veeam Backup & Replication listens on IPv4 addresses only. If the Veeam backup server name is resolved into IPv6 address, Veeam Agent will fail to connect to the Veeam backup server.



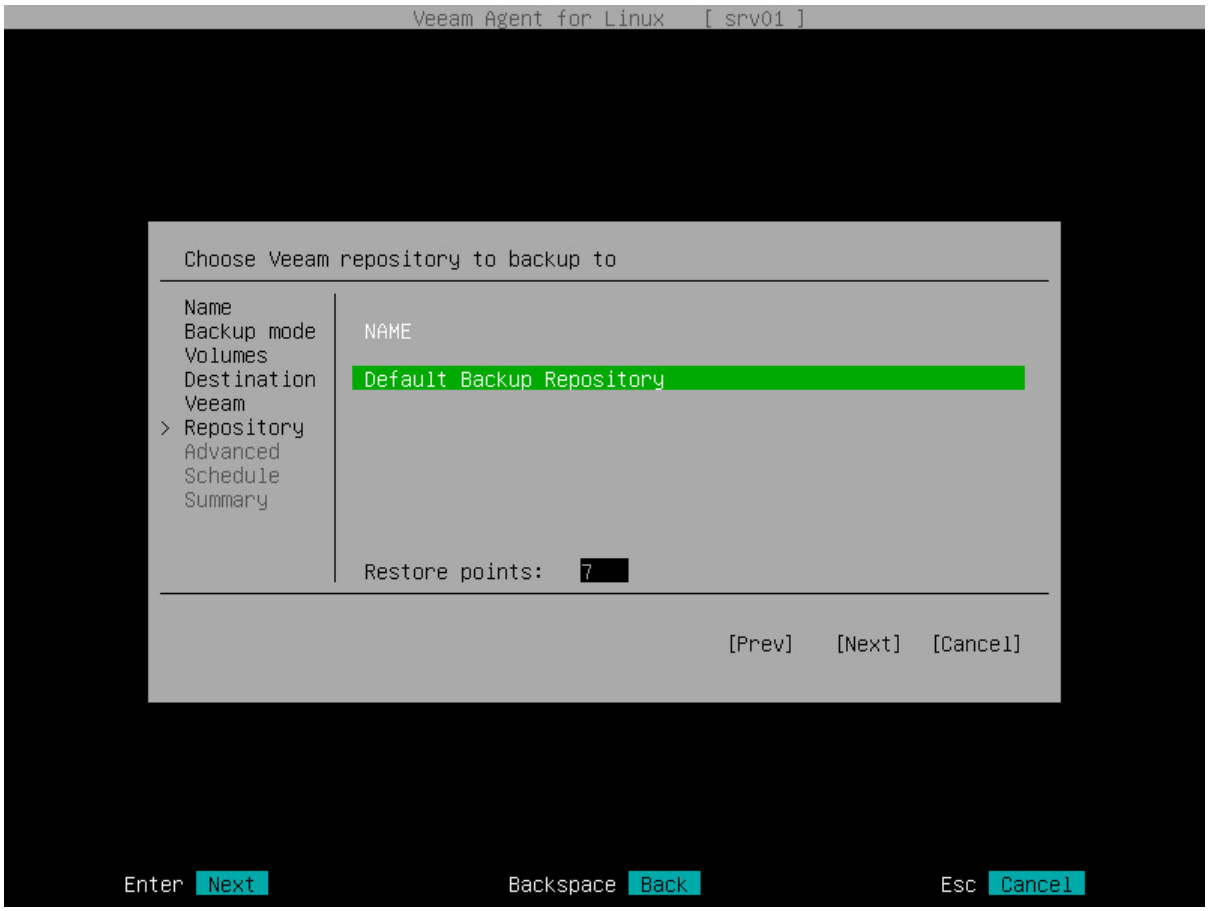
Selecting Backup Repository

The **Repository** step of the wizard is available if you have chosen to save backup files on a Veeam Backup & Replication repository.

Specify settings for the target backup repository:

1. From the list of available backup repositories, select a backup repository where you want to store backups. The list of backup repositories displays only those repositories on which you have permissions to store data. To learn more, see Setting Up User Permissions on Backup Repositories.
2. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

   To learn more, see Backup Retention Policy.



Veeam Cloud Connect Repository Settings

If you have selected to store backup files on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. Specify service provider settings.
2. Verify the TLS certificate and specify user account settings.
3. Select the cloud repository.

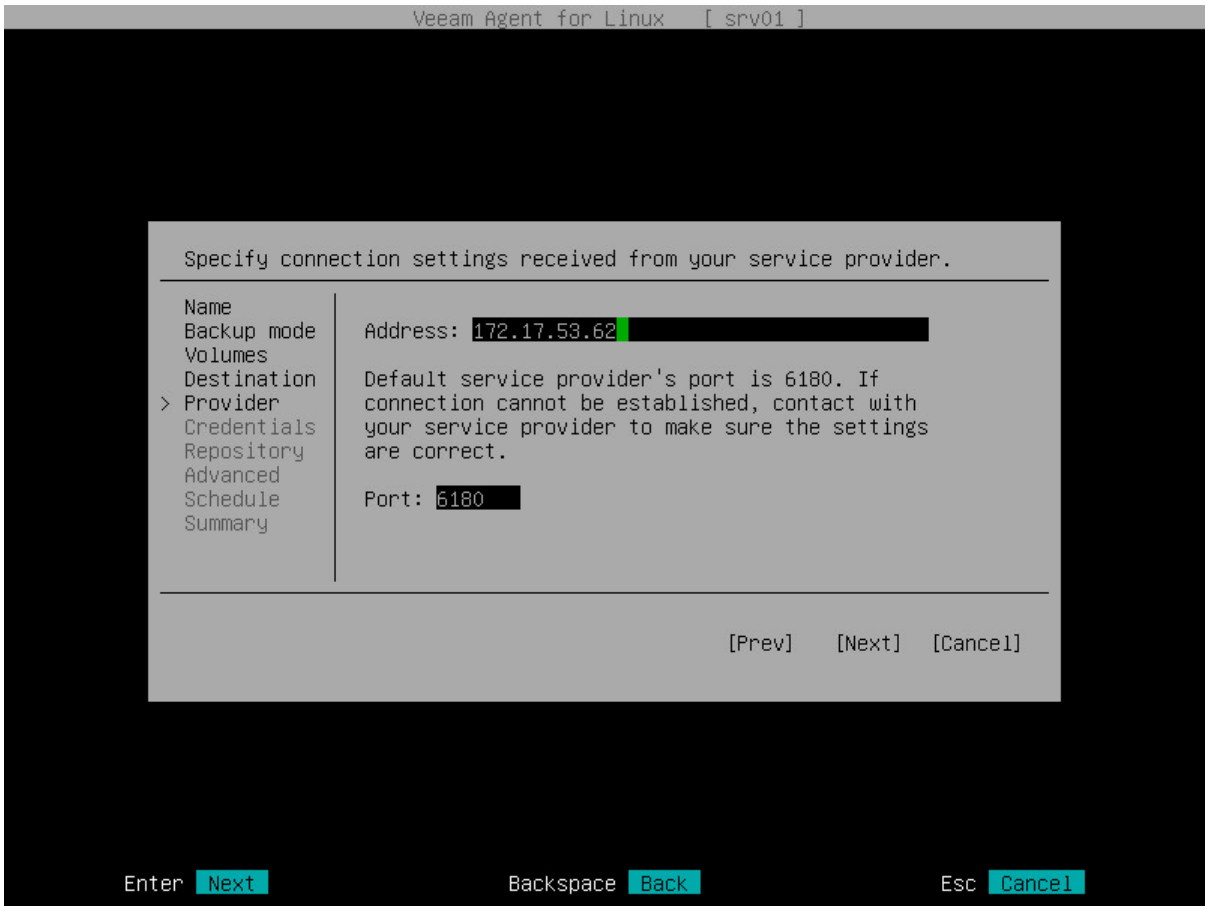
> **NOTE**
>
> The **Veeam Cloud Connect repository** option is available if Veeam Agent operates in the Workstation or Server edition.

Specifying Service Provider Settings

The **Provider** step of the wizard is available if you have chosen to save backup files on a Veeam Cloud Connect repository.

Specify settings for the cloud gateway that the Veeam Cloud Connect service provider (SP) or your backup administrator has provided to you:

1. In the **Address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.



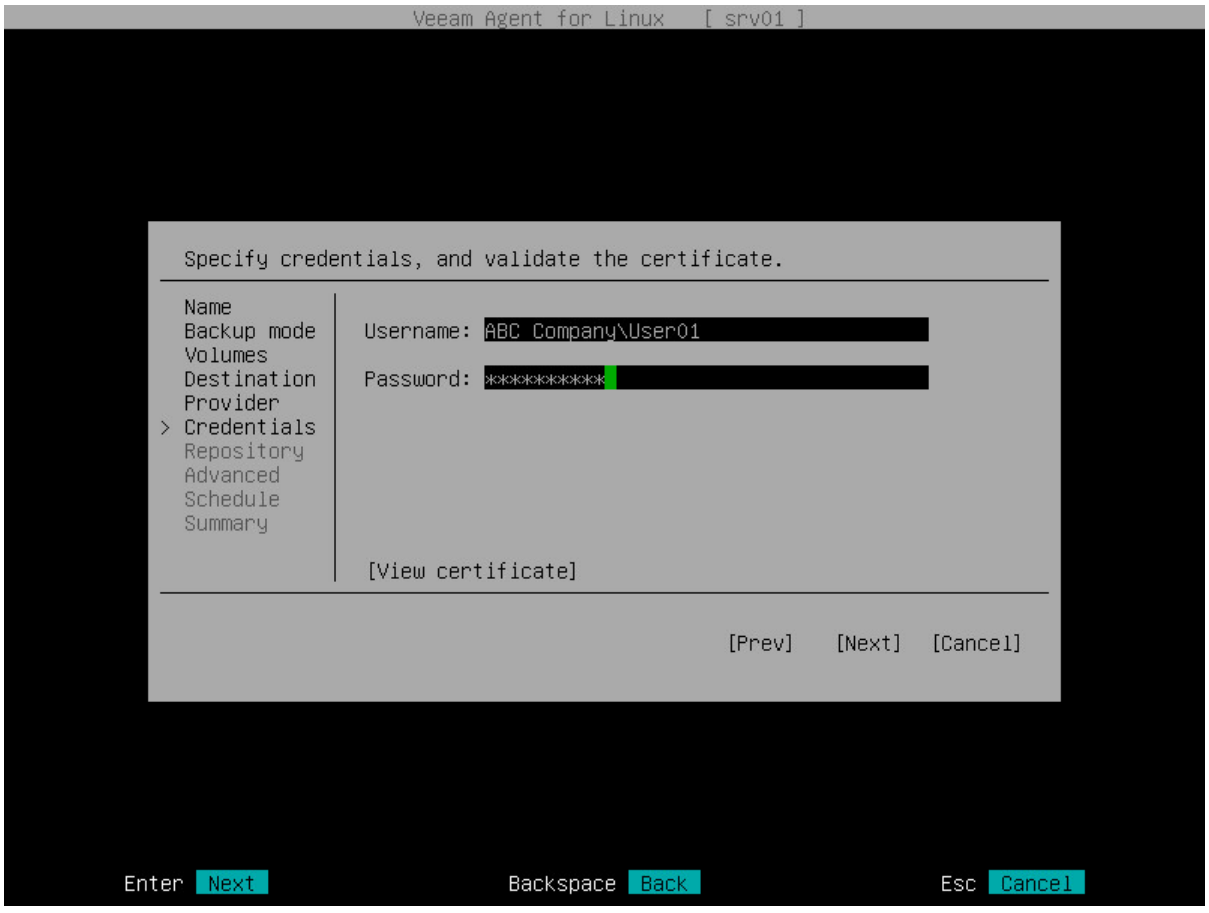
Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to save backup files on a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the cloud repository.

1. In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate:
   - [Optional] To verify the TLS certificate with a thumbprint, do the following:
     i. Select the **Verify thumbprint** button with the **Tab** key and press **Enter**.
     ii. Copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Thumbprint verification** field.
     iii. Switch to the **Verify** button and press **Enter**. Veeam Agent will check if the thumbprint you entered matches the thumbprint of the obtained TLS certificate.

     TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.
   - To accept the TLS certificate, select the **Accept** button with the **Tab** key and press **Enter**.
2. In the **Username** field, enter the name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.
3. In the **Password** field, provide a password for the tenant or subtenant account.



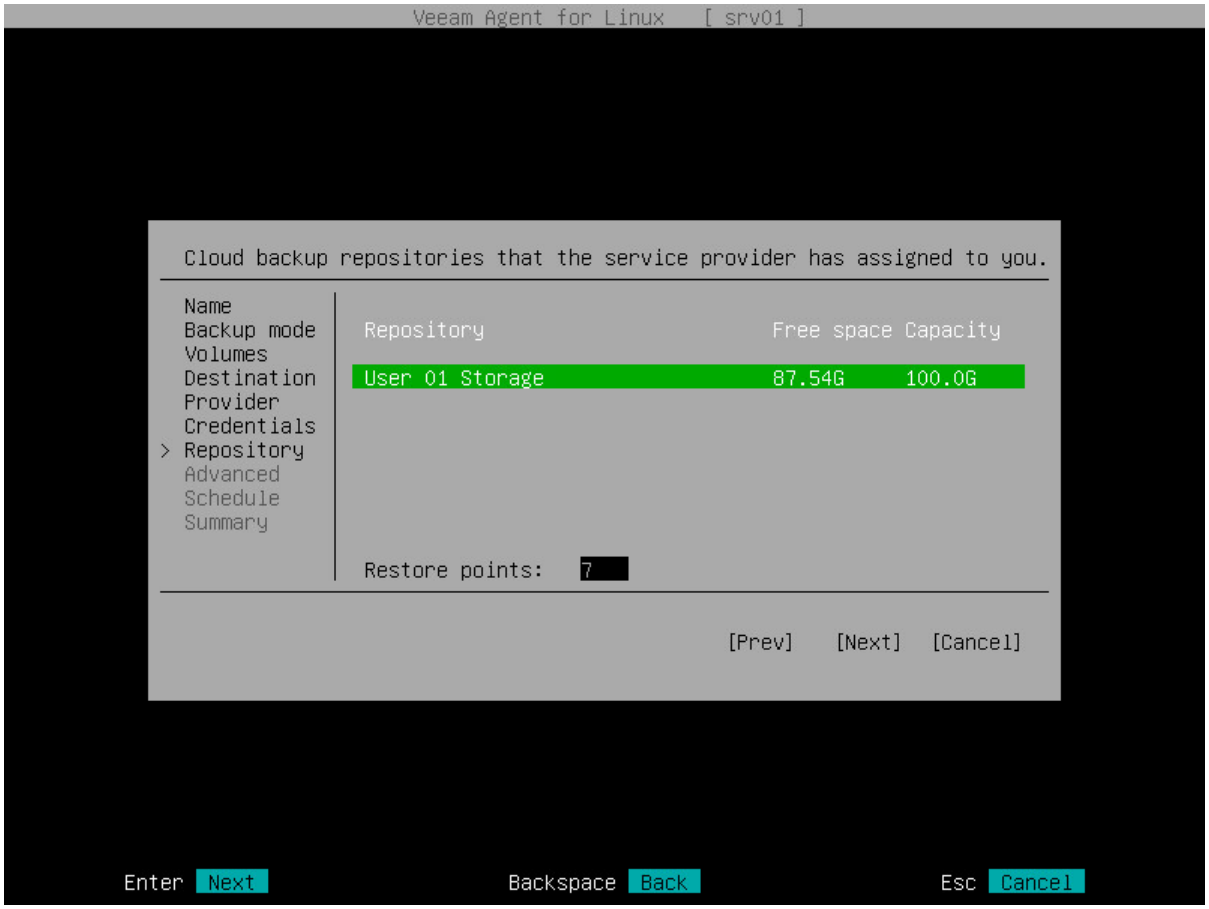
Selecting Cloud Repository

The **Repository** step of the wizard is available if you have chosen to save backup files on a cloud repository and specified settings to connect to the SP.

Specify settings for the cloud repository:

1. From the **Repository** list, select a cloud repository where you want to store created backups. The **Repository** list displays only those cloud repositories that can be accessed by the tenant or subtenant account that you use to connect to the service provider.
2. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

   To learn more, see Backup Retention Policy.



Step 7. Specify Advanced Backup Settings

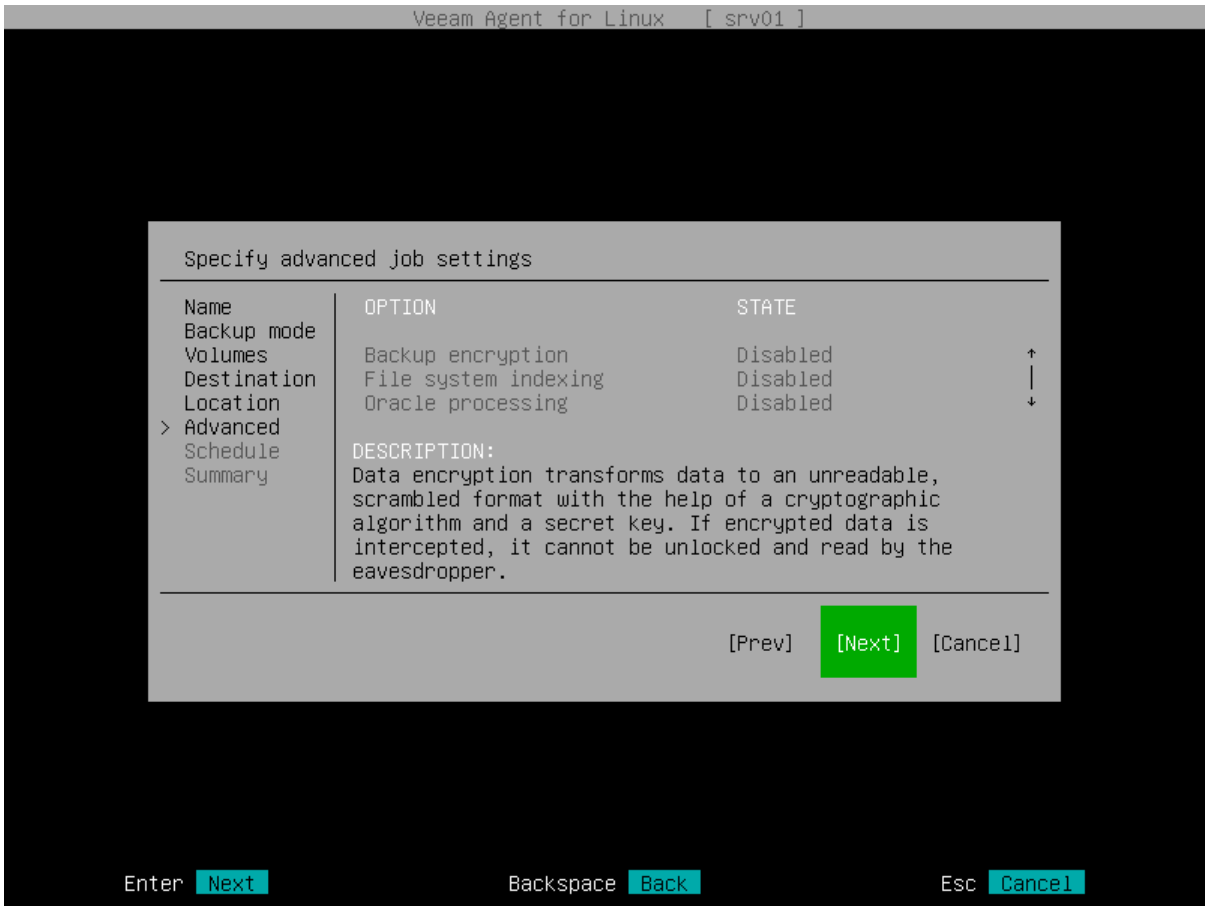
At the **Advanced** step of the wizard, specify advanced settings for the backup job:

- Data encryption settings
- File indexing settings
- Oracle database system processing settings
- MySQL database system processing settings
- PostgreSQL database system processing settings
- Active full backup settings
- Backup maintenance settings
- Script settings

NOTE

Consider the following:

- You cannot specify encryption settings for the backup job if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the Data Encryption section in the Veeam Backup & Replication User Guide.
- You can specify file indexing settings only if Veeam Agent operates in the Workstation or Server edition.
- You can specify settings for Oracle, MySQL or PostgreSQL database system processing only if Veeam Agent operates in the Server edition. The settings are available for a volume-level backup job only.
- You can specify backup maintenance settings only if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the Destination step of the wizard.
- You cannot specify data compression settings when you configure a backup job with the Backup Job wizard. If you want to specify these settings, consider creating the backup job with the Veeam Agent command line interface. To learn more, see Advanced Backup Job Settings.



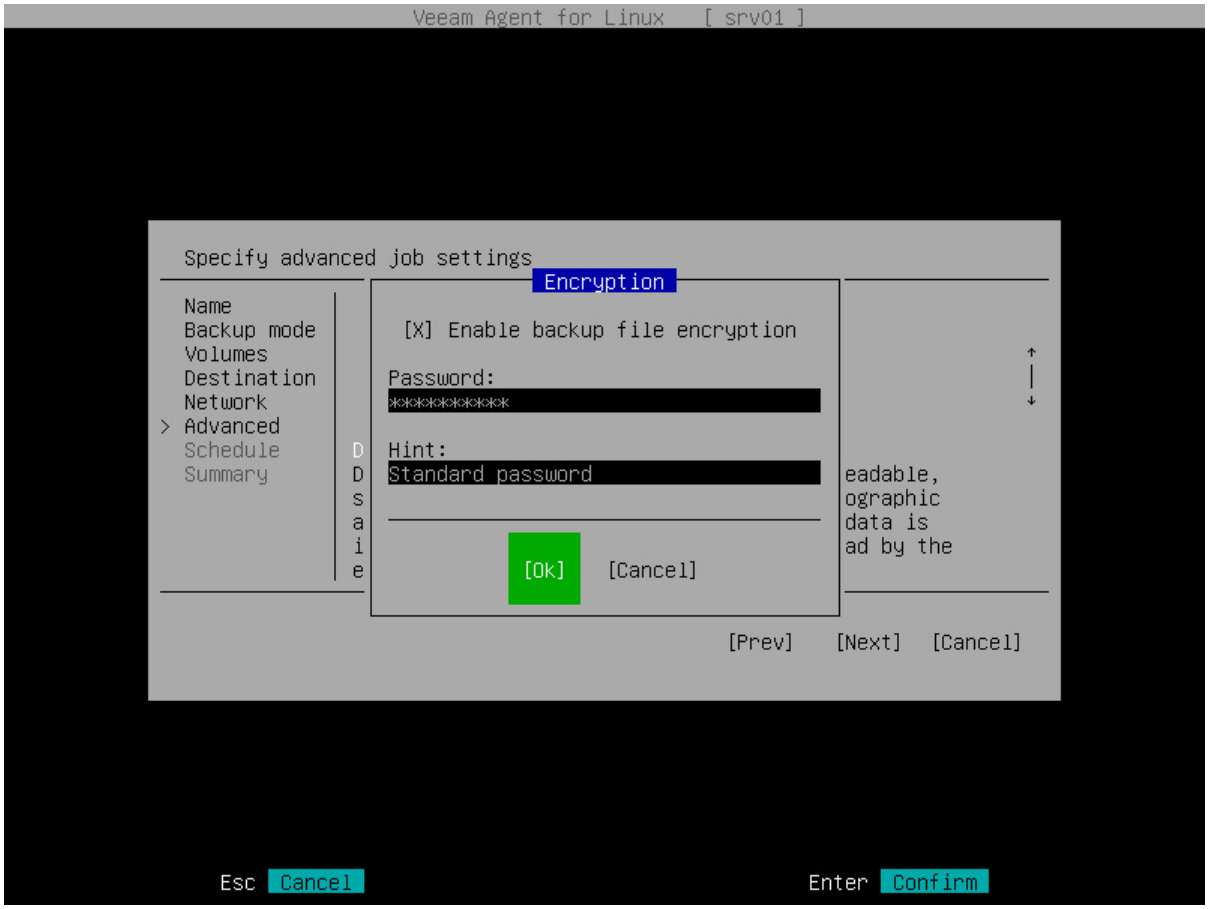
Data Encryption Settings

If you want to encrypt the content of backup files, at the **Advanced** step of the wizard, specify data encryption settings for the backup job:

1. Select the **Backup encryption** option with the **Tab** key and press **Enter**.

> **NOTE**
>
> The **Backup encryption** option is unavailable if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the Data Encryption section in the Veeam Backup & Replication User Guide.

2. In the **Encryption** window, make sure that the **Enable backup file encryption** option is selected and press **Space**.
3. In the **Password** field, type a password that you want to use for encryption.
4. In the **Hint** field, type a hint for the password. In case you lose the password, the specified hint will help you to remember the lost password.
5. Switch to the **Ok** button and press **Enter**.



File Indexing Settings

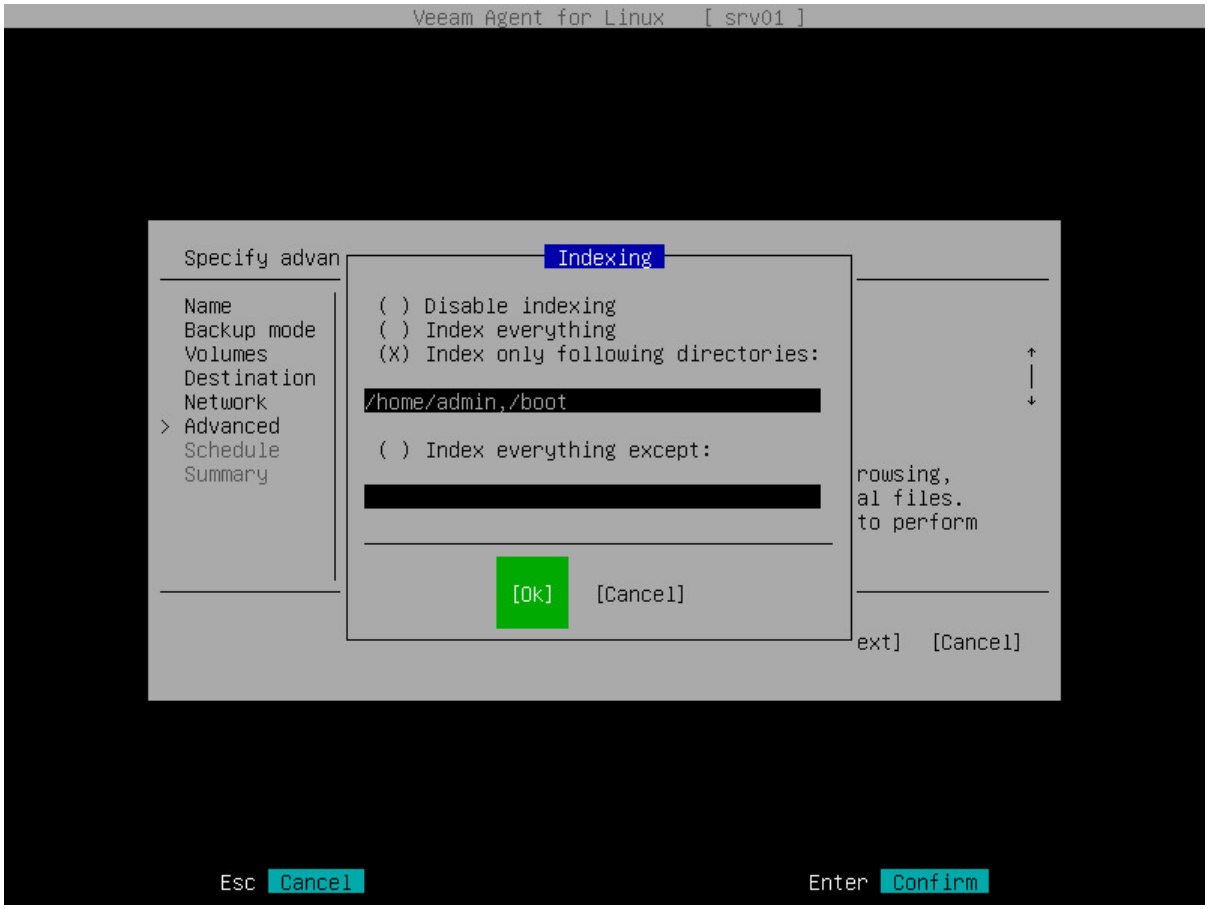
At the **Advanced** step of the wizard, specify file indexing settings for the backup job:

1. Select the **File system indexing** option with the **Tab** and **Down** keys and press **Enter**.

> **NOTE**
>
> The **File system indexing** option is available if Veeam Agent for Linux operates in the Workstation or Server edition.

2. In the **Indexing** window, specify the indexing scope:
   - Select **Index everything** if you want to index all files within the backup scope that you have specified at the Backup mode step of the wizard. Veeam Agent for Linux will index all files that reside:
     - On your computer OS (for entire machine backup)
     - On the volumes that you have selected for backup (for volume-level backup)
     - In the directories that you have selected for backup (for file-level backup)
   - [For entire machine and volume-level backups] Select **Index everything except** if you want to index all files within the specified backup scope except those files that reside in specific directories. Enter paths to directories whose files you do not want to index. To separate several paths, use the ',' (comma) character.
   - [For entire machine and volume-level backups] Select **Index only following directories** to define directories that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character.

3. Switch to the **Ok** button and press **Enter**.



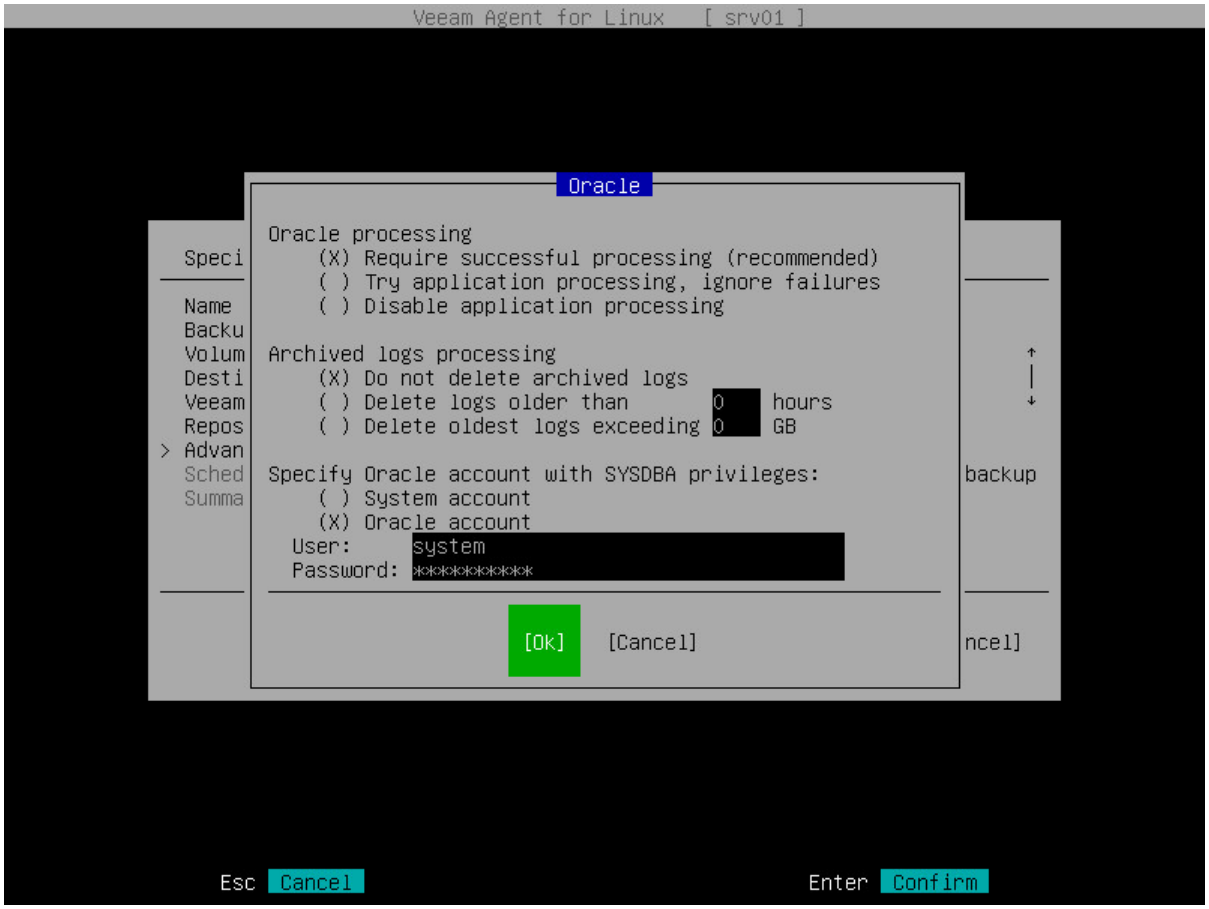
Oracle Database Processing Settings

At the **Advanced** step of the wizard, specify processing settings for the Oracle database system:

1. Select the **Oracle processing** option with the **Tab** and **Down** keys and press **Enter**.
2. In the **Oracle processing** section, select one of the following options:
   - **Require successful processing**. With this option selected, Veeam Agent will stop the backup process if an error occurs while processing the Oracle database system.
   - **Try application processing, ignore failures**. With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the Oracle database system.
3. In the **Archived logs processing** section, specify how Veeam Agent will process archived logs on the Oracle database:
   - Select **Do not delete archived logs** if you want Veeam Agent to keep archived logs. When the backup job completes, Veeam Agent will not delete archived logs.

     It is recommended that you select this option when you do not have databases running in the ARCHIVELOG mode. If the database is running in the ARCHIVELOG mode, archived logs may grow large and consume all disk space. In this case, the database administrator must take care of archived logs themselves.
   - Select **Delete logs older than <N> hours** or **Delete oldest logs exceeding <N> GB** if you want Veeam Agent to delete archived logs that are older than <N> hours or larger than <N> GB. Veeam Agent will wait for the backup job to complete successfully and then trigger archived logs truncation via Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next successful backup job session.

4. In the **Specify Oracle account with SYSDBA privileges** section, specify which account type Veeam Agent will use to connect to the database system.
   - Select **System account** if you want Veeam Agent to use an account of the Veeam Agent machine OS. The account must be a member of the group that owns configuration files for the Oracle database (for example, the oinstall group).
   - Select **Oracle account** if you want Veeam Agent to use an Oracle account. The account must have SYSDBA rights.



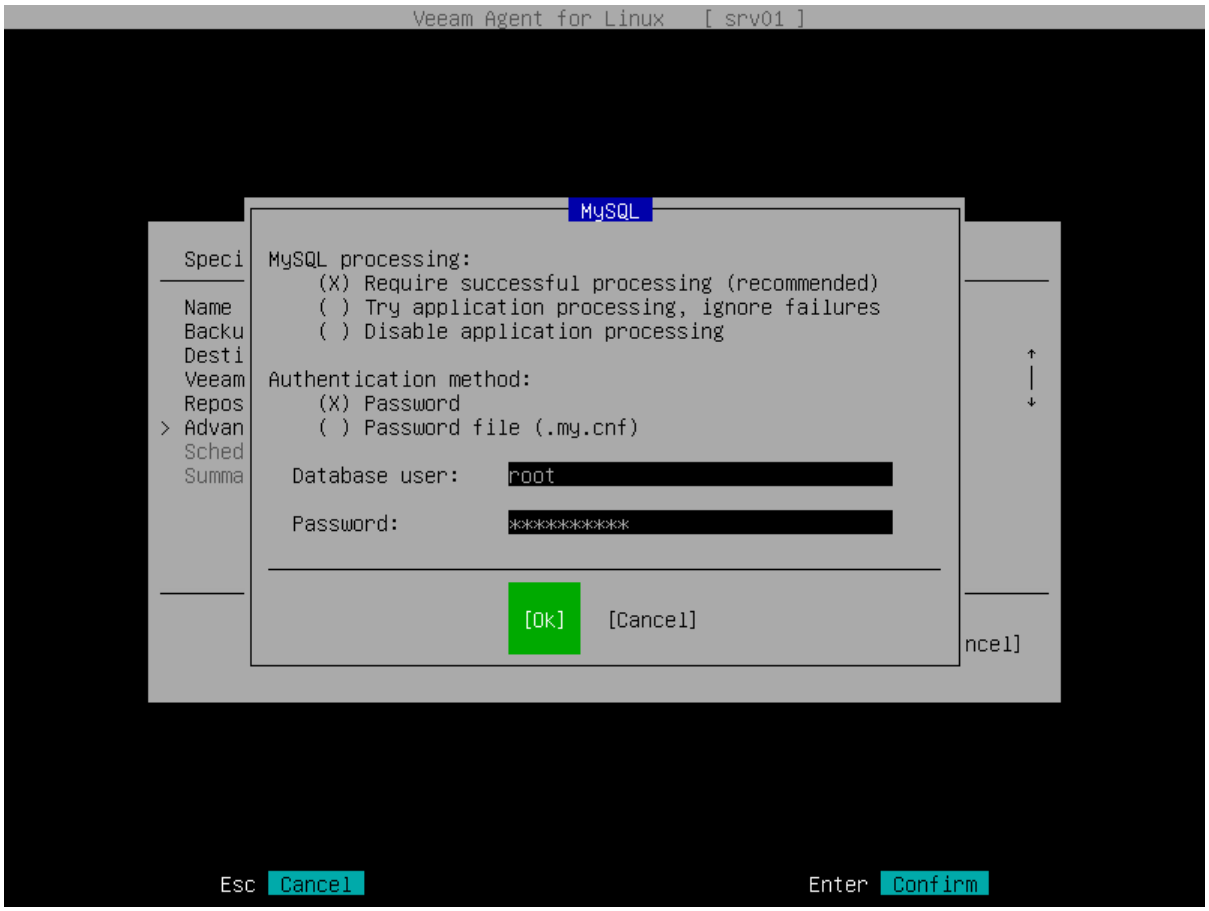
MySQL Database Processing Settings

> **IMPORTANT**
>
> To process the MySQL database system, the MySQL account must have the following privileges:
>
> - SELECT for all tables. If the MySQL account does not have the SELECT privilege for the table, Veeam Agent cannot access table metadata. As a result, Veeam Agent does not process the table. To learn more, see MySQL documentation.
> - LOCK TABLES. If the MySQL account does not have the LOCK TABLES privilege, Veeam Agent does not process tables based on the MyISAM storage engine.
> - RELOAD. This privilege is required to allow the MySQL account to perform FLUSH operations.
>
> To obtain information about privileges that are assigned to the account, use MySQL functionality, for example, the SHOW GRANTS statement. To learn more, see MySQL documentation.

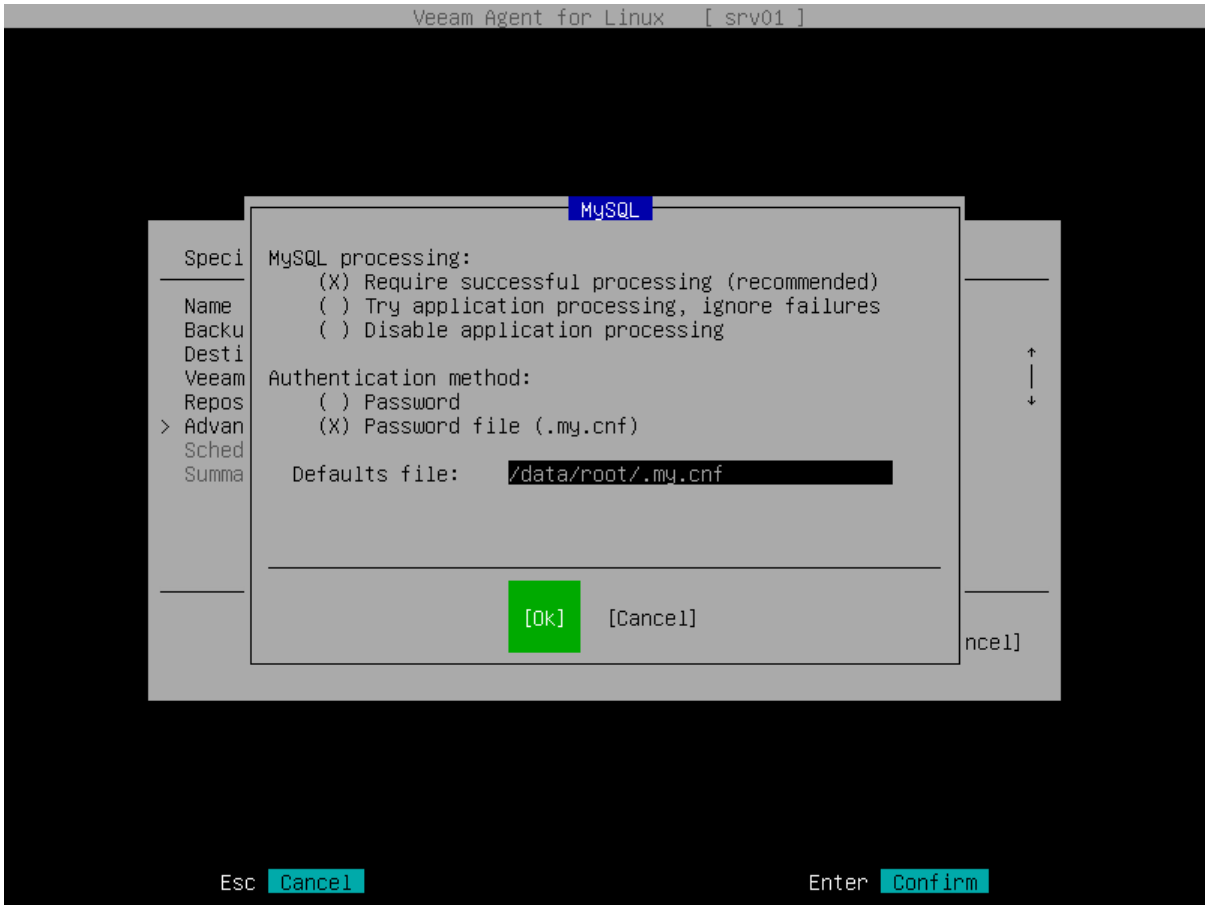
At the **Advanced** step of the wizard, specify processing settings for the MySQL database system:

1. Select the **MySQL processing** option with the **Tab** and **Down** keys and press **Enter**.
2. In the **MySQL processing** section, select one of the following options:
    - **Require successful processing**. With this option selected, Veeam Agent will stop the backup process if an error occurs when processing the MySQL database system.
    - **Try application processing, ignore failures**. With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the MySQL database system.

3. In the **Authentication method** section, specify how Veeam Agent will connect to the MySQL database:
   - Select **Password** if you want Veeam Agent to connect with the MySQL account name and password. With this option selected, you must specify account name and password in the backup job settings.



- Select **Password file** if you want Veeam Agent to connect with the MySQL account name and password that are stored in the `.my.cnf` password file. With this option selected, you must specify a path to the password file, but do not need to specify account credentials in the backup job settings. To learn more about password file configuration, see Preparing Password File for MySQL Processing.



Preparing Password File for MySQL Processing

You can use MySQL account credentials that are stored in the password file to connect Veeam Agent for Linux to the MySQL database system.

> **NOTE**
>
> Consider the following:
>
> - If you specify a custom path to the password file, specify a full path. Specifying relative paths is not supported.
> - The password file can also contain user-specific connection settings that Veeam Agent will apply to connect to the MySQL database system. For example, if you want to connect to the MySQL database system using the custom socket, specify the socket path in the password file. To learn more, see MySQL documentation.

If you want to use a password file for authentication, create a file. By default, Veeam Agent expects the password file to have the `.my.cnf` name and to be in the home directory of the `root` user. If the password file has a custom name or is stored in another directory, you can specify a custom path.

The password file must have the following contents:

```
[client]
user=<username>
password=<password>
```

where:

- <username> — name of the account that Veeam Agent will use to connect to the MySQL database system.
- <password> — password of the account that Veeam Agent will use to connect to the MySQL database system.

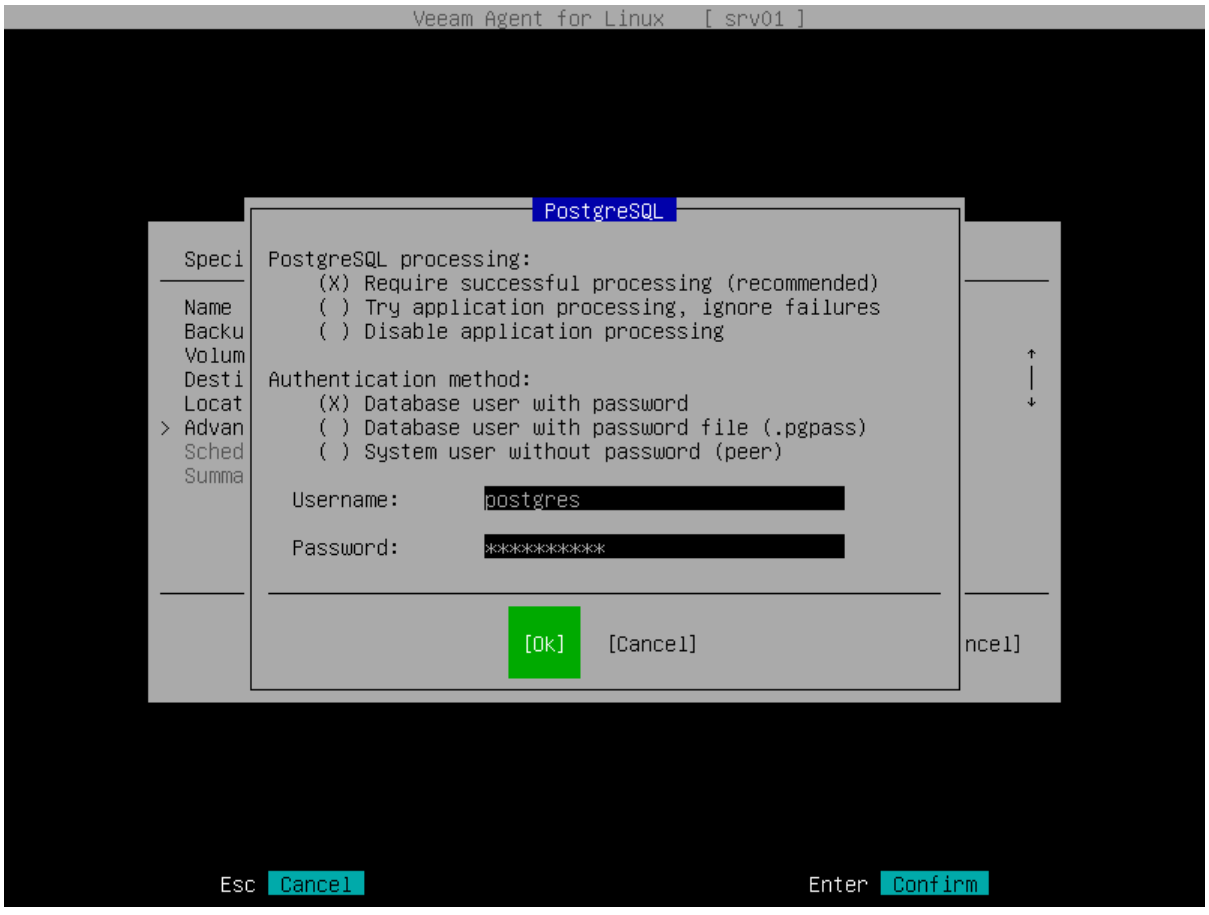
For example:

```
[client]
user=root
password=P@ssw0rd
```

PostgreSQL Database Processing Settings

At the **Advanced** step of the wizard, specify processing settings for the PostgreSQL database system:

1. Select the **PostgreSQL processing** option with the **Tab** and **Down** keys and press **Enter**.
2. In the **PostgreSQL processing** section, select one of the following options:
   - **Require successful processing**. With this option selected, Veeam Agent will stop the backup process if an error occurs when processing the PostgreSQL database system.
   - **Try application processing, ignore failures**. With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the PostgreSQL database system.
3. In the **Authentication method** section, specify how Veeam Agent for Linux will connect to the PostgreSQL database:
   - Select **Database user with password** if you want Veeam Agent to connect with the PostgreSQL account name and password. With this method selected, you must specify account name and password in the backup job settings.
   - Select **Database user with password file** if you want Veeam Agent to connect with the PostgreSQL account password that is stored in the `.pgpass` password file. With this method selected, you must specify account name only in the backup job settings. To learn more about password file configuration, see Password File for PostgreSQL.
   - Select **System user without password** if you want Veeam Agent to connect using a peer authentication method. In the peer authentication method, Veeam Agent uses the OS account as the PostgreSQL database user name. With this option selected, you must specify OS account in the backup job settings. To learn more about peer authentication, see PostgreSQL documentation.



Preparing Password File for PostgreSQL Processing

You can use PostgreSQL account credentials that are stored in the password file to connect Veeam Agent to the PostgreSQL database system.

If you want to use a password file for authentication, create the `.pgpass` file in the home directory of the root user.

The password file must have the following contents:

```
<hostname>:<port>:<database>:<username>:<password>
```

where:

- `<hostname>` — name of the host where the PostgreSQL database system is located.
- `<port>` — number of the free port that Veeam Agent will use to connect to the PostgreSQL database system.
- `<database>` — name of the PostgreSQL database.
- `<username>` — name of the account that Veeam Agent will use to connect to the PostgreSQL database system.
- `<password>` — password of the account that Veeam Agent will use to connect to the PostgreSQL database system.

For example:

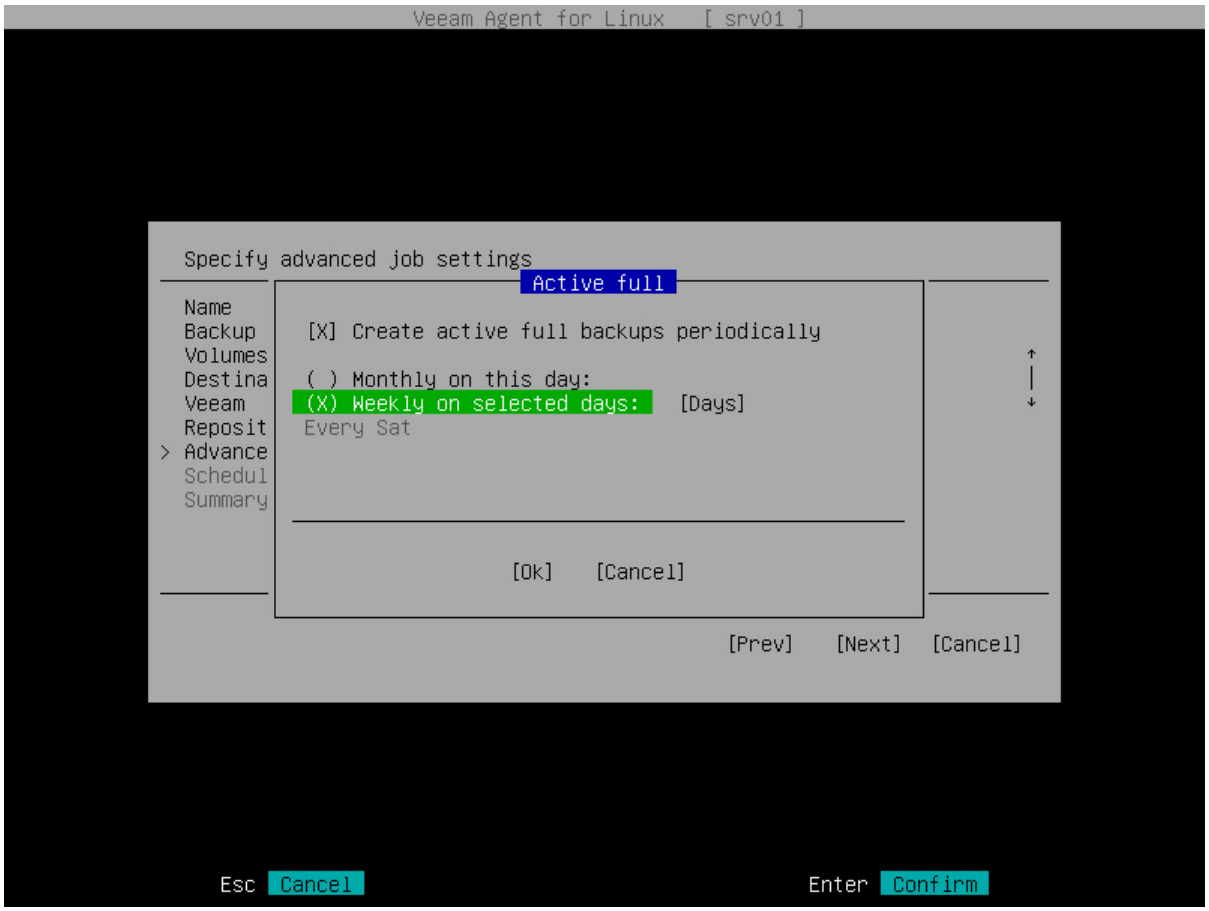
```
srv01:5432:mydb:postgres:P@ssw0rd
```

For more information about the password file, see PostgreSQL documentation.

Active Full Backup Settings

At the **Advanced** step of the wizard, specify active full backup settings for the backup job:

1. Select the **Active full** option with the **Tab** and **Down** keys and press **Enter**.
2. In the **Active full** window, make sure that the **Create active full backups periodically** option is selected and press **Space**.
3. Specify schedule for periodic active full backups:
   - If you want active full backups to run monthly, select the **Monthly on this day** option and specify the day of a month when Veeam Agent will perform active full backup.
   - If you want active full backups to run weekly, do the following:
     i. Select the **Weekly on selected days** option, then select **Days** with the **Tab** key and press **Enter**.
     ii. In the **Days** window, specify days on which Veeam Agent will perform active full backup. By default, Veeam Agent performs active full backup every Saturday. To select days, use **Up**, **Down**, **Right**, **Left** and **Space** keys.
     iii. Switch to the **Ok** button with the **Tab** key and press **Enter**.



Maintenance Settings

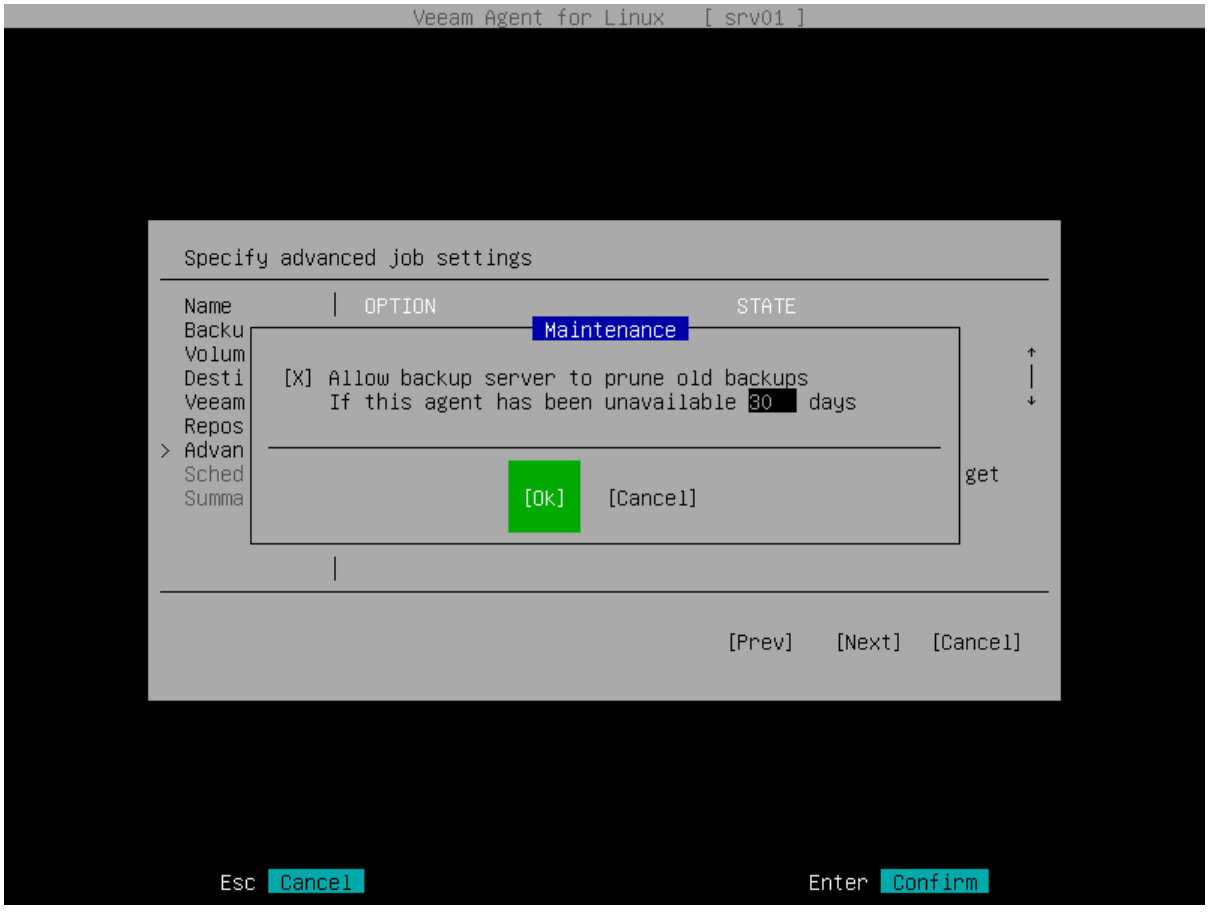
At the **Advanced** step of the wizard, specify backup maintenance settings for the backup job:

1. Select the **Maintenance** option with the **Tab** and **Down** keys and press **Enter**.

> **NOTE**
>
> The **Maintenance** option is available if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the Destination step of the wizard.

2. In the **Maintenance** window, make sure that the **Allow backup server to prune old backups** option is selected and press **Space**.
3. In the **If this agent has been unavailable <N> days** field, specify the number of days for which you want to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the period that you have specified. When this period is over, the backup will be removed from the target location.

   By default, the retention period for old backups is 30 days. Do not set this retention period to 1 day or a similar short interval. In the opposite case, the backup job may work not as expected and remove data that you still require.



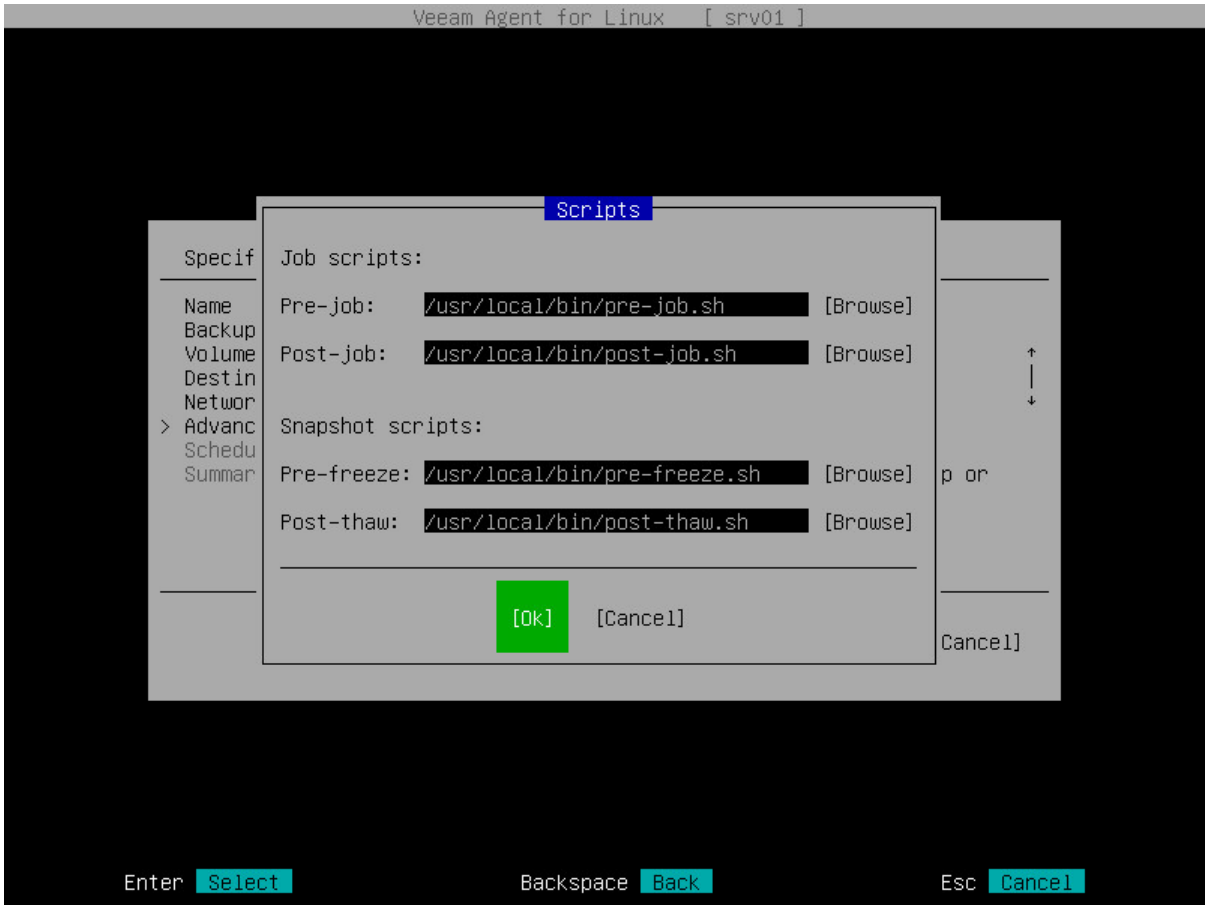
Script Settings

At the **Advanced** step of the wizard, specify script settings for the backup job:

1. Select the **Scripts** option with the **Tab** and **Down** keys and press **Enter**.
2. In the **Scripts** window, in the **Job scripts** section, specify custom scripts that you want to execute before and/or after the backup job:
   - In the **Pre-job** field, specify a path to the script that should be executed before the backup job starts.
   - In the **Post-job** field, specify a path to the script that should be executed after the backup job completes.
3. In the **Scripts** window, in the **Snapshot scripts** section, specify custom scripts that you want to execute before Veeam Agent creates a snapshot of the backed-up volume and/or after the snapshot is created:
   - In the **Pre-freeze** field, specify a path to the script that should be executed before Veeam Agent creates a volume snapshot.
   - In the **Post-thaw** field, specify a path to the script that should be executed after Veeam Agent creates a volume snapshot.
4. Switch to the **Ok** button and press **Enter**.

IMPORTANT

You can specify snapshot script settings only if Veeam Agent for Linux operates in the Server edition. To learn more about editions, see Product Editions.



Specifying Path to Script

You can specify a path to the executable file of the job or snapshot script in one of the following ways:

1. Type a path to the executable file.
2. Browse to the executable file:
   a. Select the **Browse** option with the **Tab** key and press **Enter**.
   b. In the **Choose script location** window, select the directory being a part of the path to the script and press **Enter**.
   c. Repeat the step 'b' until a path to the directory in which the executable file resides appears in the **Current directory** field.
   d. Select the necessary executable file and press **Enter**.

      Alternatively, you can switch to the **Ok** button and press **Enter**.

> **TIP**
>
> If you do not want to execute a script, you can leave the corresponding field blank and proceed to the next step of the wizard.

Step 8. Specify Backup Schedule

At the **Schedule** step of the wizard, specify the schedule according to which you want to perform backup.

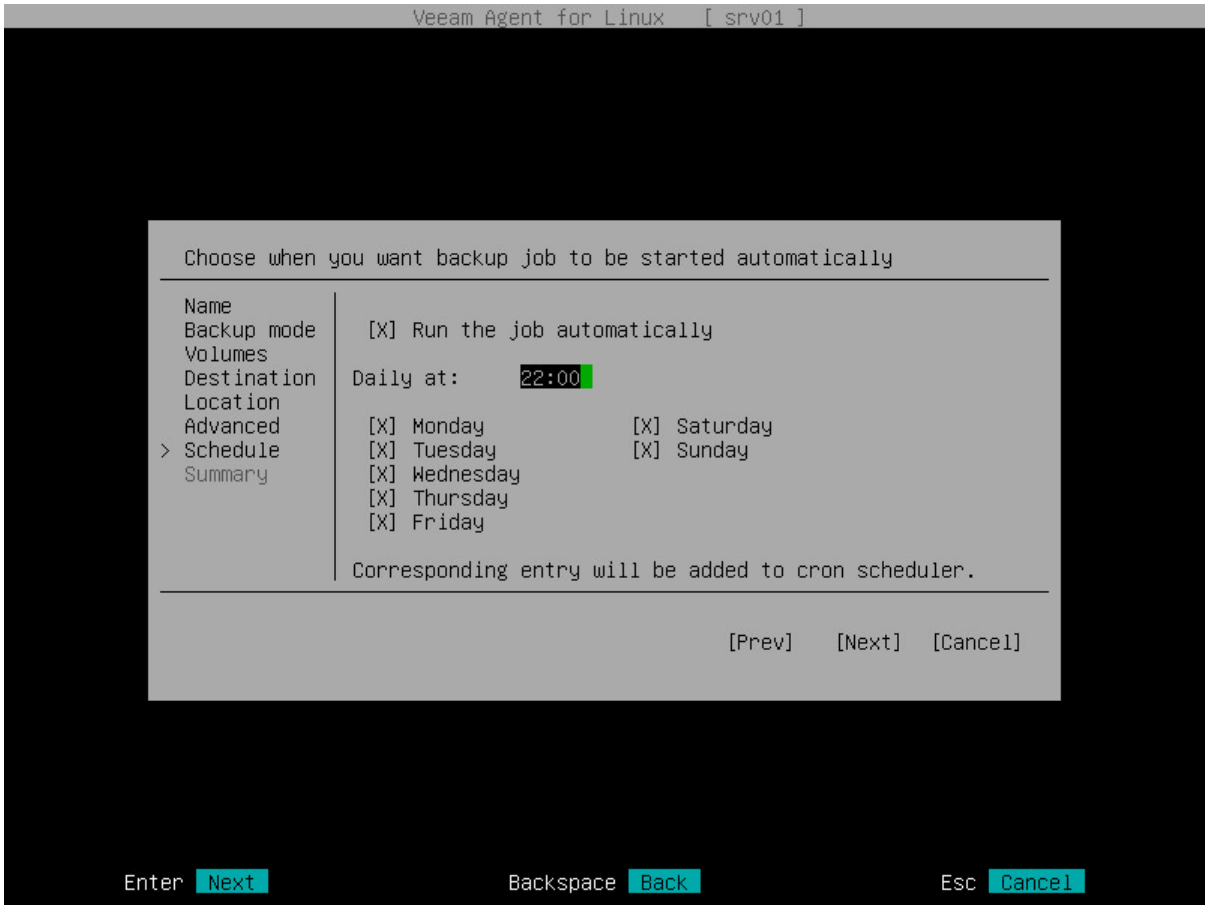
1. Make sure that the **Run the job automatically** check box is selected.

   If you want to configure the backup job without schedule, you can clear the **Run the job automatically** check box. In this case you will be able start the configured backup job manually at any time you need.

2. In the **Daily at** field, specify the time of day when the backup job must start.

3. In the list of days of the week, specify days on which the job must start. By default, Veeam Agent will start the backup job daily at the specified time. To exclude specific days from the schedule, clear check boxes next to the necessary days.

Veeam Agent for Linux will save the scheduling settings for the backup job in its database and in the `crontab` configuration file of the root account. Thus, Veeam Agent can start a backup job automatically regardless of the currently running user session. You can change schedule settings at any time in Veeam Agent or edit the `crontab` file manually to fine-tune the schedule. To learn more, refer to the Cron job scheduler documentation.

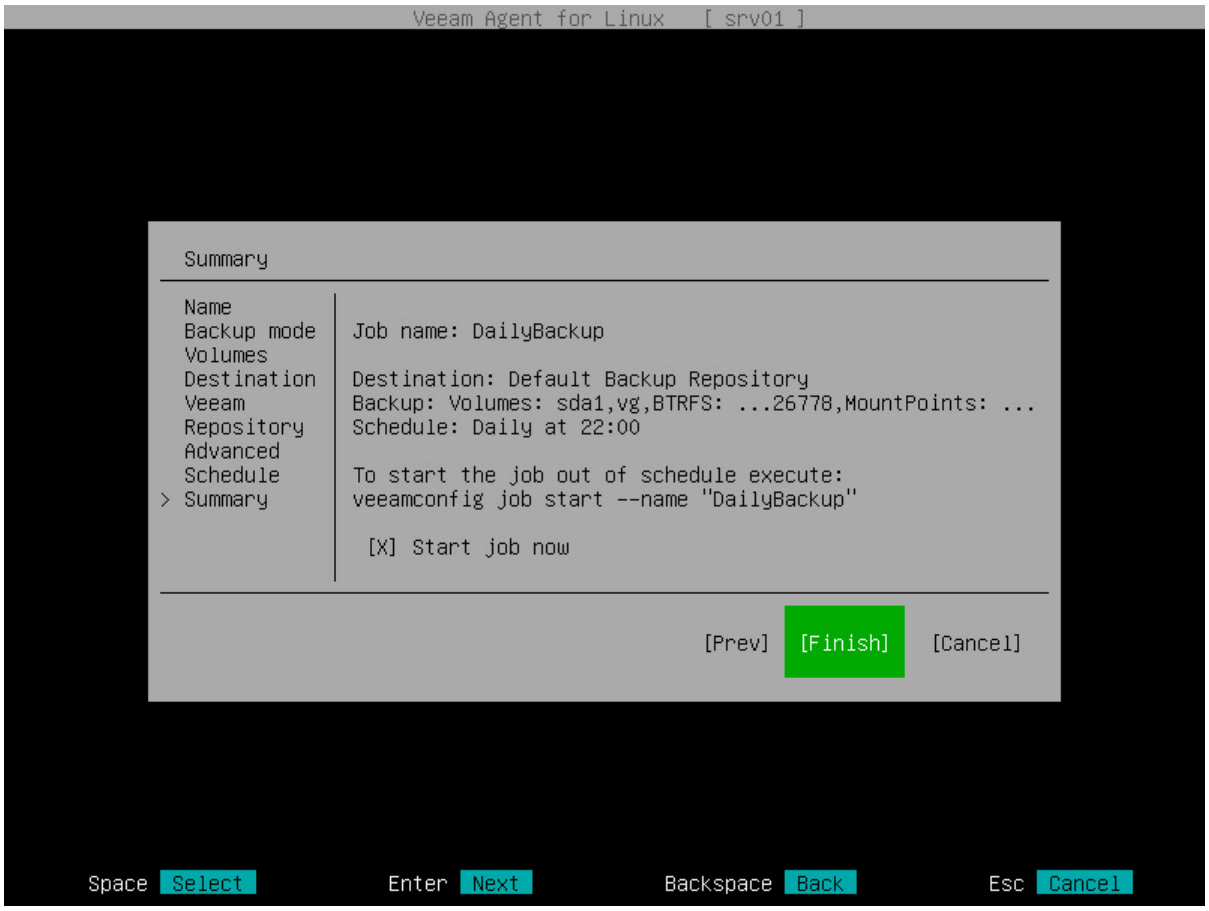


Step 9. Review Backup Job Settings

At the **Summary** step of the wizard, complete the backup job configuration process.

1. Review settings of the configured backup job.
2. To start the job after you close the wizard, make sure that the **Start job now** check box is selected.

   If you want to start the backup job later, you can clear the **Start job** now check box. You will be able to start the backup job manually at any time you need. To learn more, see Starting Backup Job.
3. Press **Enter to exit the wizard**.



What You Do Next

After you configure the backup job, you can start the backup job at any time you need. To learn more, see Starting Backup Job.

If some of your data gets lost or corrupted, you can do the following:

- Recover all computer volumes or specific volumes from the backup.
- Recover individual files and folders from the backup.

Creating Backup Job with Command Line Interface

You can configure the backup job with the command line interface. Using Veeam Agent for Linux commands, you can create volume-level and file-level backup jobs, specify advanced settings for the created backup job, define backup schedule and enable backup encryption.

Creating Volume-Level Backup Job

> **IMPORTANT**
>
> Volume-level backup job relies on a device name in the `/dev` directory. Device names in the `/dev` directory (for example, `/dev/md-127`, `/dev/dm-1`) must stay persistent for backed-up volumes. Otherwise, the job will back up the wrong volume.

You can create a volume-level backup of the entire computer image or specific volumes.

To back up the entire computer image, use the following command:

```
veeamconfig job create volumelevel --name <job_name> --reponame <repository_nam
e> --backupallsystem
<advanced_options> <schedule_options> <active_full_backup_options> <indexing_op
tions>
```

To back up specific volumes, use the following command:

```
veeamconfig job create volumelevel --name <job_name> --reponame <repository_nam
e> --objects <volume_to_backup> <advanced_options> <schedule_options> <active_f
ull_backup_options> <indexing_options>
```

where:

- `<job_name>` — name for the created backup job.
- `<repository_name>` — name of the backup repository that should be used as a target location for the backup job. The backup repository must be created in advance. To learn more, see Creating Backup Repository.

  If you want to create Veeam Agent backups in a Veeam backup repository of cloud repository, you need connect to the Veeam backup server or Veeam Cloud Connect service provider in advance, before configuring the backup job. To learn more, see Connecting to Veeam Backup Server and Connecting to Service Provider.
- `<volume_to_backup>` — object that should be included in backup:
  - For simple volumes — name of a block device that represents a volume or an entire disk that should be included in backup. You can specify entire disk to create backup of the entire computer image or individual computer volumes to create backup of specific volumes. If you want to back-up several disks or volumes, specify them one after another using the ',' (comma) character as a separator.

    > **NOTE**
    >
    > If you include a block device in the backup, and this block device is a physical volume assigned to an LVM volume group, Veeam Agent will include the whole LVM volume group in the backup.
  - For LVM volumes — name of an LVM logical volume that should be included in backup. If you want to back-up several LVM logical volumes, specify them one after another using the ',' (comma) character as a separator.
- `<advanced_options>` — advanced options for the backup job. To learn more, see Advanced Backup Job Settings.

- `<schedule_options>` — schedule options for the backup job. To learn more, see Schedule Settings.
- `<active_full_backup_options>` — active full backup schedule options for the backup job. To learn more, see Active Full Backup Schedule Settings.
- `<indexing_options>` — file system indexing options for the backup job. To learn more, see File System Indexing Settings.

For example:

```
$ veeamconfig job create --name SystemBackup --reponame Repository_01 --objects
/dev/sda1 --weekdays Mon,Sun --weekdays-full Thu
```

TIP

Consider the following:

- You can also specify backup schedule for the backup job after your create the backup job. For details, see Configuring Backup Schedule.
- You can also specify active full backup schedule for the backup job after your create the backup job. For details, see Configuring Active Full Backup Schedule.
- You can also specify database processing settings for the volume-level backup job. For details, see Configuring Database Processing Settings.

Advanced Backup Job Settings

You can specify the following advanced options for the backup job:

Option	Description and values
--compressionlevel	Data compression level. Possible values are: • *0* — No compression • *1* — Rle • *2* — Lz4 • *3* — ZlibLow • *4* — ZlibHigh The default value is *2*.
--blocksize	Data block size in kilobytes. Possible values are 256, 512, 1024, 4096 or 8192. The default value is *1024*.
--maxpoints	The number of restore points that you want to store in the backup location. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.
--prefreeze	Path to the script that should be executed before the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--postthaw	Path to the script that should be executed after the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--prejob	Path to the script that should be executed at the start of the backup job.
--postjob	Path to the script that should be executed after the backup job completes.
--setencryption	Defines that data encryption option is enabled for the job. When you use the `veeamconfig job create` command with the `--setencryption` option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.

Option	Description and values
--deleteold	The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999. If you do not specify the `--deleteold` option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually.

Schedule Settings

You can specify schedule options for the backup job to create backups daily or on specific weekdays at specific time:

Option	Description and values
--daily	Defines that the backup job must start daily at specific time.
--weekdays	Weekdays when the backup job must start separated by a comma (','). Possible values are: • *Mon* — Monday • *Tue* — Tuesday • *Wed* — Wednesday • *Thu* — Thursday • *Fri* — Friday • *Sat* — Saturday • *Sun* — Sunday
--at	Time of day when the backup job must start specified in the *HH:MM* format. For example: `20:00`.

After the backup job is created, Veeam Agent for Linux automatically enables backup schedule. To learn about how to configure backup schedule for an existing backup job, see Configuring Backup Schedule.

Active Full Backup Schedule Settings

You can specify schedule options for the backup job to create active full backups on specific weekdays or days of the month:

Option	Description and values
--weekdays-full	Weekdays when the backup job must create an active full backup separated by a comma (','). Possible values are: • *Mon* — Monday • *Tue* — Tuesday • *Wed* — Wednesday • *Thu* — Thursday • *Fri* — Friday • *Sat* — Saturday • *Sun* — Sunday
--thisday-full	Days of the month when the backup job must create an active full backup separated by a comma (','). Possible values are: 1-30 (or 1-31, depending on the month length).

After the backup job is created, Veeam Agent for Linux automatically enables active full backup schedule. To learn about how to configure active full backup schedule for an existing backup job, see Configuring Active Full Backup Schedule.

File System Indexing Settings

You can specify one the following file system indexing options for the backup job:

Option	Description and values
--indexall	Defines that Veeam Agent for Linux must index all files on the volumes included in backup.
--indexonly	Path to a directory that contains files that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character.
--indexexcept	Path to a directory that contains files that you do not want to index. You can specify one or more paths. To separate several paths, use the ',' (comma) character.

To learn more about file indexing, see File System Indexing.

Creating File-Level Backup Job

To create a file-level backup job, use the following command:

```
veeamconfig job create filelevel --name <job_name> --reponame <repository_name>
<objects> <advanced_options> <schedule_options> <active_full_backup_options> <i
ndexing_options> --nosnap
```

where:

- `<job_name>` — name for the created backup job.
- `<repository_name>` — name of the backup repository that should be used as a target location for the backup job. The backup repository must be created in advance. To learn more, see Creating Backup Repository

  If you want to create Veeam Agent backups in the Veeam backup repository, you should connect to the Veeam backup server in advance, before configuring the backup job. To learn more, see Connecting to Veeam Backup Server.
- `<objects>` — files and directories inclusion/exclusion options. To learn more, see File Inclusion Options.
- `<advanced_options>` — advanced options for the backup job. To learn more, see Advanced Backup Job Settings.
- `<schedule_options>` — schedule options for the backup job. To learn more, see Schedule Settings.
- `<active_full_backup_options>` — active full backup schedule options for the backup job. To learn more, see Active Full Backup Schedule Settings.
- `<indexing_options>` — file system indexing options for the backup job. To learn more, see File System Indexing Settings.
- `--nosnap` — option that instructs Veeam Agent for Linux to perform backup in the snapshot-less mode. With this option enabled, Veeam Agent for Linux will not create a snapshot of the backed-up volumes during backup. This allows Veeam Agent to back up data residing in file systems that are not supported for snapshot-based backup with Veeam Agent for Linux. Keep in mind that the snapshot-less file-level backup does not guarantee that data in the backup is consistent. To learn more, see Snapshot-Less File-Level Backup.

For example:

```
$ veeamconfig job create filelevel --name HomeFolderBackup --reponame NetworkRe
pository --includedirs /home/user --excludedirs /home/user/temp --excludemasks
"*.pdf"
```

TIP

Consider the following:

- You can also specify backup schedule for the backup job after your create the backup job. For details, see Configuring Backup Schedule.
- You can also specify active full backup schedule for the backup job after your create the backup job. For details, see Configuring Active Full Backup Schedule.

File Inclusion Options

When you create a file-level backup job, you must specify at least one directory that should be included in backup. If you do not want to back up some files and directories in the specified directory, you can exclude specific files and directories from backup.

To define the backup scope for the file-level backup job, you can use the following command-line options:

Option	Description and values
--includedirs	Full path to a directory that should be included in backup, for example: `/home/user`. You can specify one or several paths to directories in the computer file system. To separate several paths, use the ',' (comma) character, for example: `/home/user/Documents,/home/user/reports`. **Tip:** If you want to back up the root directory and specify the '/' (slash) character, Veeam Agent does not automatically include mount points in the backup scope. To include mount points, you need to specify paths to these mount points manually. For example, you have a file system mounted to the `/home/media` directory. If you add `/` as an object to the backup scope, Veeam Agent will not back up the mounted file system. To back up the root directory and the mounted file system, add the following objects to the backup scope: `/,/home/media`.
--excludedirs	Full path to a directory that should be excluded from backup. The directory specified with this option must be a subdirectory of the directory specified with the `--includedirs` option. To separate several paths, use the ',' (comma) character, for example, `/home/user/Documents,/home/user/reports`.

Option	Description and values
--includemasks	Mask for file name or path that should be included in backup. You can use the following UNIX wildcard characters for file name 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`. • '[]' — a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: `report_201[3456].pdf` or `report_201[3-6].pdf`. Keep in mind that you must specify all names with masks in double quotation marks (""). For example: `--includemasks "*.bak"`. If you want to use several file name masks, you must specify them in double quotation marks ("") and separated with a comma (,). For example: `--includemasks "*.bak,*.pdf"`. File inclusion option is applied to all directories that are specified with the `--includedirs` option. For example, if you include in backup the `/home/user/Documents` directory and files that match the `repor?.pdf` file name mask, Veeam Agent for Linux will back up the `/home/user/Documents/report.pdf` file and will not back up the `/home/user/reports/report.pdf` file.

Option	Description and values
--excludemasks	Mask for file name or path that should be excluded from backup. You can use the following UNIX wildcard characters for file name 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`. • '[]' — a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '–' character). For example: `report_201[3456].pdf` or `report_201[3-6].pdf`. Keep in mind that you must specify all names with masks in double quotation marks (""). For example: `--excludemasks "*.bak"`. If you want to use several file name masks, you must specify them in double quotation marks ("") and separated with a comma (,). For example: `--excludemasks *.bak,*.pdf"`. File exclusion option is applied to all directories that are specified with the `--includedirs` option and files that match file name masks specified with the `--includemasks` option. For example, you may want to specify the following backup scope for the backup job: • Include in backup the `/home/user/Documents` directory • Include files that match the `report.*` file name mask • Exclude files that match the `*.odt` file name mask. In this case, Veeam Agent for Linux will back up the `/home/user/Documents/report.pdf` file and will not back up `/home/user/Documents/report.odt` and `/home/user/reports/report.pdf` files.

Advanced Backup Job Settings

You can specify the following advanced options for the backup job:

Option	Description and values
--compressionlevel	Data compression level. Possible values are: • *0* — No compression • *1* — Rle • *2* — Lz4 • *3* — ZlibLow • *4* — ZlibHigh

Option	Description and values
--blocksize	Data block size in kilobytes. Possible values are 256, 512, 1024, 4096 or 8192. The default value is *1024*.
--maxpoints	The number of restore points that you want to store in the backup location. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.
--prefreeze	Path to the pre-freeze script that should be executed before the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--postthaw	Path to the post-thaw script that should be executed after the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--prejob	Path to the script that should be executed at the start of the backup job.
--postjob	Path to the script that should be executed after the backup job completes.
--setencryption	Defines that data encryption option is enabled for the job. When you use the `veeamconfig job create` command with the `--setencryption` option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.
--deleteold	The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999. If you do not specify the `--deleteold` option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually.

Schedule Settings

You can specify schedule options for the backup job to create backups daily or on specific weekdays at specific time:

Option	Description and values
--daily	Defines that the backup job must start daily at specific time.
--weekdays	Weekdays when the backup job must start separated by a comma (','). Possible values are: • *Mon* — Monday • *Tue* — Tuesday • *Wed* — Wednesday • *Thu* — Thursday • *Fri* — Friday • *Sat* — Saturday • *Sun* — Sunday
--at	Time of day when the backup job must start specified in the *HH:MM* format. For example: `20:00`.

After the backup job is created, Veeam Agent for Linux automatically enables backup schedule. To learn about how to configure backup schedule for an existing backup job, see Configuring Backup Schedule.

Active Full Backup Schedule Settings

You can specify schedule options for the backup job to create active full backups on specific weekdays or days of the month:

Option	Description and values
--weekdays-full	Weekdays when the backup job must create an active full backup separated by a comma (','). Possible values are: • *Mon* — Monday • *Tue* — Tuesday • *Wed* — Wednesday • *Thu* — Thursday • *Fri* — Friday • *Sat* — Saturday • *Sun* — Sunday
--thisday-full	Days of the month when the backup job must create an active full backup separated by a comma (','). Possible values are: 1-30 (or 1-31, depending on the month length).

After the backup job is created, Veeam Agent for Linux automatically enables active full backup schedule. To learn about how to configure active full backup schedule for an existing backup job, see Configuring Active Full Backup Schedule.

File System Indexing Settings

You can specify the following file system indexing option for the backup job:

Option	Description and values
--indexall	Defines that Veeam Agent for Linux must index all files in the directories included in backup.

To learn more about file indexing, see File System Indexing.

Configuring Backup Schedule

To run a backup job periodically without the user intervention, you can schedule it to start automatically. You can specify schedule settings individually for every job created in Veeam Agent for Linux. You can perform the following actions with the backup job schedule via the command-line interface:

- Specify schedule settings for the job.
- Enable schedule for the job.
- View the schedule defined for the job.
- Disable schedule for the job.

> **TIP**
>
> You can also specify backup schedule for the backup job when you create the job. For details, see Creating Volume-Level Backup Job and Creating File-Level Backup Job.

Specifying Backup Schedule

You can schedule a backup job to start at specific time daily or on specific week days.

Specifying Daily Schedule

To specify daily schedule settings for the backup job, use the following command:

```
veeamconfig schedule set --jobid <job_id> --daily --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --daily --at <time>
```

where:

- `<job_id>` — ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to configure the schedule.

> **TIP**
>
> To view IDs or names of all existent backup jobs, you can press the **Tab** key right after you type the `--jobid` or `--jobname` option.

- `<time>` — time of day when the backup job must start specified in the *HH:MM* format. For example: `20:00`.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6
c73f --daily --at 20:00
```

Specifying Schedule on Specific Days

To specify schedule settings for the backup job, use the following command:

```
veeamconfig schedule set --jobid <job_id> --weekdays <days> --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --weekdays <days> --at <time>
```

where:

- `<job_id>` — ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to configure the schedule.

> **TIP**
>
> To view IDs or names of all existent backup jobs, you can press the **Tab** key right after you type the `--jobid` or `--jobname` option.

- `<days>` — days when the backup job must start separated by a comma (','). For example: `Monday,Tuesday,Wednesday,Thursday,Friday`.
- `<time>` — time of day when the backup job must start specified in the *HH:MM* format. For example: `20:00`.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --weekdays Monday,
Tuesday,Wednesday,Thursday,Friday --at 20:00
```

Enabling Backup Schedule

To run a backup job automatically upon the defined schedule, you must enable the schedule for the job. Use the following command:

```
veeamconfig schedule enable --jobid <job_id>
```

or

```
veeamconfig schedule enable --jobname <job_name>
```

where:

- `<job_id>` — ID of the backup job for which you want to enable the schedule. You should look up the job ID in advance, for example, with the veeamconfig job list command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to enable the schedule.

For example:

```
user@srv01:~$ veeamconfig schedule enable --jobid 4849a3ae-1935-4969-98a3-d8acd
2f6c73f
```

You can disable the schedule for the job at any time. To learn more, see Disabling Backup Schedule.

Viewing Backup Schedule

To view the schedule defined for the backup job, use the following command:

```
veeamconfig schedule show --jobid <job_id>
```

or

```
veeamconfig schedule show --jobname <job_name>
```

where:

- `<job_id>` — ID of the backup job for which you want to view the schedule.
- `<job_name>` — name of the backup job for which you want to view the schedule.

Veeam Agent for Linux displays the following information about the backup job schedule:

Parameter	Description
Days	Days on which the backup job runs automatically.
At	Time of day when the backup job starts automatically.
Run automatically	State of backup schedule. Possible values: • Enabled • Disabled

For example:

```
user@srv01:~$ veeamconfig schedule show --jobid 4849a3ae-1935-4969-98a3-d8acd2f
6c73f
Days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday
At: 20:00
Run automatically: enabled
```

Disabling Backup Schedule

To disable the schedule for the backup job, use the following command:

```
veeamconfig schedule disable --jobid <job_id>
```

or

```
veeamconfig schedule disable --jobname <job_name>
```

where:

- `<job_id>` — ID of the backup job for which you want to disable the schedule.
- `<job_name>` — name of the backup job for which you want to disable the schedule.

For example:

```
user@srv01:~$ veeamconfig schedule disable --jobid 4849a3ae-1935-4969-98a3-d8ac
d2f6c73f
```

Configuring Active Full Backup Schedule

You can schedule a backup job to create active full backups periodically. You can specify active full schedule settings individually for every job created in Veeam Agent for Linux. You can perform the following actions with the active full backup schedule via the command-line interface:

- Specify active full backup schedule.
- Enable active full backup schedule.
- View active full backup schedule.
- Disable active full backup schedule.

> **TIP**
>
> You can also specify active full backup schedule for the backup job when when you create the job. For details, see Creating Volume-Level Backup Job and Creating File-Level Backup Job.

Specifying Active Full Backup Schedule

You can schedule a backup job to create active full backups on a specific day of the month or on specific week days.

Specifying Monthly Active Full Backup Schedule

To instruct Veeam Agent to perform active full backup on a specific day on the month, use the following command:

```
veeamconfig schedule activefull set --jobid <job_id> --thisday <day>
```

or

```
veeamconfig schedule activefull set --jobname <job_name> --thisday <day>
```

where:

- `<job_id>` — ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to configure the active full backup schedule.

> **TIP**
>
> To view IDs or names of all existent backup jobs, you can press the **Tab** key right after you type the `--jobid` or `--jobname` option.

- `<day>` — number of the day of the month when Veeam Agent must perform active full backup.

For example:

```
user@srv01:~$ veeamconfig schedule activefull set --jobname DailyBackup --thisd
ay 1
```

Specifying Weekly Active Full Backup Schedule

To instruct Veeam Agent to perform active full backup on specific week days, use the following command:

```
veeamconfig schedule activefull set --jobid <job_id> --weekdays <days>
```

or

```
veeamconfig schedule activefull set --jobname <job_name> --weekdays <days>
```

where:

- `<job_id>` — ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to configure the active full backup schedule.

> **TIP**
>
> To view IDs or names of all existent backup jobs, you can press the **Tab** key right after you type the `--jobid` or `--jobname` option.

- `<days>` — days when the backup job must create an active full backup separated by a comma (','). For example: `Monday,Friday`. The backup job will create an active full backup on the specified days at the time specified in the backup job schedule settings.

For example:

```
user@srv01:~$ veeamconfig schedule activefull set --jobname DailyBackup --weekd
ays Monday,Friday
```

Enabling Active Full Backup Schedule

After you specify active full backup schedule settings for the backup job, Veeam Agent automatically enables active full backup schedule for the job. You can also enable active full backup schedule manually, for example, if you previously disabled it for the backup job. To enable active full backup schedule, use the following command:

```
veeamconfig schedule activefull enable --jobid <job_id>
```

or

```
veeamconfig schedule activefull enable --jobname <job_name>
```

where:

- `<job_id>` — ID of the backup job for which you want to enable the active full backup schedule. You should look up the job ID in advance, for example, with the veeamconfig job list command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to enable the active full backup schedule.

For example:

```
user@srv01:~$ veeamconfig schedule activefull enable --jobname DailyBackup
```

You can disable the schedule for the job at any time. To learn more, see Disabling Backup Schedule.

Viewing Active Full Backup Schedule

To view the active full backup schedule defined for the backup job, use the following command:

```
veeamconfig schedule activefull show --jobid <job_id>
```

or

```
veeamconfig schedule activefull show --jobname <job_name>
```

where:

- `<job_id>` — ID of the backup job for which you want to view the active full backup schedule.
- `<job_name>` — name of the backup job for which you want to view the active full backup schedule.

Veeam Agent for Linux displays the following information about the active full backup schedule:

Parameter	Description
Every <value>	Days on which the backup job creates active full backups. For example: *Every Monday* or *Every 1 day of every month*.
Run automatically	State of the active full backup schedule. Possible values: • Enabled • Disabled

For example:

```
user@srv01:~$ veeamconfig schedule activefull show --jobname DailyBackup
Every Monday
Run automatically: enabled
```

Disabling Active Full Backup Schedule

To disable the active full backup schedule for the backup job, use the following command:

```
veeamconfig schedule activefull disable --jobid <job_id>
```

or

```
veeamconfig schedule activefull disable --jobname <job_name>
```

where:

- `<job_id>` — ID of the backup job for which you want to disable the active full backup schedule.
- `<job_name>` — name of the backup job for which you want to disable the active full backup schedule.

For example:

```
user@srv01:~$ veeamconfig schedule activefull disable --jobname DailyBackup
```

Configuring Database Processing Settings

You can enable database processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux. With database processing settings enabled, Veeam Agent will create transactionally consistent backups of Veeam Agent machines that run database systems.

You can perform the following actions with database processing settings via the command-line interface:

- Specify Oracle database processing settings
- Specify MySQL database processing settings
- Specify PostgreSQL database processing settings
- View database processing settings
- Disable database processing settings

Specifying Oracle Processing Settings

You can enable Oracle processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

To enable Oracle processing settings for the backup job, use the following command:

```
veeamconfig aap set oracle --jobid <job_id> <oracle_options>
```

or

```
veeamconfig aap set oracle --jobname <job_name> <oracle_options>
```

where:

- `<job_id>` — ID of the backup job for which you want to enable Oracle processing settings. You should look up the job ID in advance, before configuring Oracle processing settings, for example, with the `veeamconfig job list` command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to enable Oracle processing settings.
- `<oracle_options>` — Oracle processing settings for the backup job. To learn more, see Oracle Processing Settings.

> **TIP**
>
> To view IDs or names of all existent backup jobs, you can press the **Tab** key right after you type the --jobid or --jobname option.

Oracle Processing Settings

You can specify the following Oracle processing settings for the backup job:

Option	Description and values
--tryprocess	Defines that Veeam Agent must continue the backup process if errors occur when processing the Oracle database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the Oracle database system.

Option	Description and values
--prunelogs	The number of hours to keep Oracle archived logs or the size of archived logs to keep. • If you want Veeam Agent to delete archived logs that are older than <N> hours, specify the necessary value in the `<N>H` format, For example, `10H`. • If you want Veeam Agent to delete archived logs that are larger than <N> GB, specify the necessary value in the `<N>G` format. For example: `10G`. Veeam Agent will wait for the backup job to complete successfully and then trigger archived logs truncation via Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next successful backup job session.
--usroracleos	Name of the Veeam Agent machine OS account. To connect to the Oracle database system, the account must be a member of the group that owns configuration files for the Oracle database (for example, the oinstall group). You do not need this option if you want to use the Oracle account to connect to the database. Instead, specify the necessary account with the `--usroracledb` option.
--usroracledb	Name of the Oracle account. To connect to the Oracle database system, the account must have SYSDBA rights on the databases to be processed. You do not need this option if you want to use the OS account to connect to the database. Instead, specify the necessary account with the `--usroracleos` option.

For example:

```
user@srv01:~$ veeamconfig aap set oracle --jobid 29bc2e1a-e35c-4efb-8d37-b7177b
8ea75 --tryprocess --prunelogs 10G --usroracledb system
```

Specifying MySQL Processing Settings

You can enable MySQL processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

> **IMPORTANT**
>
> To process the MySQL database system, the MySQL account must have the following privileges:
>
> - SELECT for all tables. If you do not have the SELECT privilege for the table, Veeam Agent does not process the table.
> - LOCK TABLES. If you do not have the LOCK TABLES privilege, Veeam Agent does not process MyISAM tables.
> - RELOAD. This privilege is required to allow the MySQL account to perform FLUSH operations.
>
> To obtain information about privileges that are assigned to the account, use MySQL functionality, for example, the SHOW GRANTS statement. To learn more, see MySQL documentation.

To enable MySQL processing settings for the backup job, use the following command:

```
veeamconfig aap set mysql --jobid <job_id> <mysql_options>
```

or

```
veeamconfig aap set mysql --jobname <job_name> <mysql_options>
```

where:

- `<job_id>` — ID of the backup job for which you want to enable MySQL processing settings. You should look up the job ID in advance, before configuring MySQL processing settings, for example, with the `veeamconfig job list` command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to enable MySQL processing settings.
- `<mysql_options>` — MySQL processing settings for the backup job. To learn more, see MySQL Processing Settings.

> **TIP**
>
> To view IDs or names of all existent backup jobs, you can press the **Tab** key right after you type the --jobid or --jobname option.

MySQL Processing Settings

You can specify the following MySQL processing settings for the backup job:

Option	Description and values
--tryprocess	Defines that Veeam Agent must continue the backup process if errors occur when processing the MySQL database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the MySQL database system.
--usrmysqidb	Name of the MySQL account. Veeam Agent can connect to the MySQL database system in one of the following ways: • If you specify account name (`--usrmysqidb` option) only, Veeam Agent will prompt you to specify a password to access the MySQL database system. • If you specify account name and password (`--usrmysqidb` and `--password` options), Veeam Agent will access the MySQL database system. • If you do not specify account credentials (`--usrmysqidb` and `--password` options), Veeam Agent will use a password file to connect to the MySQL database system. To learn more about password file configuration, see Preparing Password File for MySQL Processing.
--password	Password of the MySQL account. If you do not specify the `--password` value, Veeam Agent will prompt you to specify a password to access the MySQL database. Keep in mind, if you specify the password using the `--password` option, password is stored in terminal in plain text.
--defaults-file	Path to a password file. You must specify a full path to a password file if you want Veeam Agent to use a password file located in specific directory. Specifying relative paths is not supported. With this method selected, you do not need to specify account credentials in the backup job settings. You do not need this option in the following cases: • Veeam Agent uses account name and password that are specified in the backup job settings to connect to the MySQL database. • Veeam Agent uses account credentials that are stored in the password file in `/root/.my.cnf`.

Examples

Authentication with password:

```
user@srv01:~$ veeamconfig aap set mysql --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8
ea75 --tryprocess --usrmysqidb root --password P@ssw0rd
```

Authentication with password file:

```
user@srv01:~$ veeamconfig aap set mysql --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8
ea75 --tryprocess
--defaults-file /data/root/.my.cnf --password P@ssw0rd
```

Preparing Password File for MySQL Processing

You can use MySQL account credentials that are stored in the password file to connect Veeam Agent for Linux to the MySQL database system.

> **NOTE**
>
> Consider the following:
>
> - If you specify a custom path to the password file, specify a full path. Specifying relative paths is not supported.
> - The password file can also contain user-specific connection settings that Veeam Agent will apply to connect to the MySQL database system. For example, if you want to connect to the MySQL database system using the custom socket, specify the socket path in the password file. To learn more, see MySQL documentation.

If you want to use a password file for authentication, create a file. By default, Veeam Agent expects the password file to have the `.my.cnf` name and to be in the home directory of the `root` user. If the password file has a custom name or is stored in another directory, you can specify a custom path.

The password file must have the following contents:

```
[client]
user=<username>
password=<password>
```

where:

- <username> — name of the account that Veeam Agent will use to connect to the MySQL database system.
- <password> — password of the account that Veeam Agent will use to connect to the MySQL database system.

For example:

```
[client]
user=root
password=P@ssw0rd
```

Specifying PostgreSQL Processing Settings

You can enable PostgreSQL processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

To enable PostgreSQL processing settings for the backup job, use the following command:

```
veeamconfig aap set postgres --jobid <job_id> <postgres_options>
```

or

```
veeamconfig aap set postgres --jobname <job_name> <postgres_options>
```

where:

- `<job_id>` — ID of the backup job for which you want to enable PostgreSQL processing settings. You should look up the job ID in advance, before configuring PostgreSQL processing settings, for example, with the `veeamconfig job list` command. To learn more, see Viewing List of Backup Jobs.
- `<job_name>` — name of the backup job for which you want to enable PostgreSQL processing settings.
- `<postgres_options>` — PostgreSQL processing settings for the backup job. To learn more, see PostgreSQL Processing Settings.

> **TIP**
>
> To view IDs or names of all existent backup jobs, you can press the **Tab** key right after you type the --jobid or --jobname option.

PostgreSQL Processing Settings

You can specify the following PostgreSQL processing settings for the backup job:

Option	Description and values
--tryprocess	Defines that Veeam Agent must continue the backup process if errors occur when processing the PostgreSQL database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the PostgreSQL database system.
--usrpgdb	Name of the PostgreSQL account. If you uses a password file to connect to the PostgreSQL database system, the `--usrpgdb` option allows to select the user from the password file. You do not need this option if you want to use a Veeam Agent machine OS account to connect to the PostgreSQL database system. Instead, specify the OS account with the `--usrpgos` option.
--password	Password of the PostgreSQL account. If you do not specify this option, Veeam Agent will prompt to enter the password. If you do not specify the password in prompt, Veeam Agent uses a password file to connect to the PostgreSQL database system. To learn more about password file configuration, see Preparing Password File for PostgreSQL Processing. Keep in mind, if you specify the password using the `--password` option, password is stored in terminal in plain text.
--usrpgos	Name of the OS account. Veeam Agent will use the name to connect to the PostgreSQL database system using the peer authentication method. In the peer authentication method, Veeam Agent uses the OS account as the PostgreSQL database user name. With this option selected, you must specify OS account only. To learn more about peer authentication, see PostgreSQL documentation. You do not need this option if you want to use a PostgreSQL account to connect to the database system. Instead, specify the PostgreSQL account with the `--usrpgdb` option.

For example:

```
user@srv01:~$ veeamconfig aap set postgres --jobid 29bc2e1a-e35c-4efb-8d37-b717
7b8ea75 --tryprocess --usrpgdb postgres --password P@ssw0rd
```

Preparing Password File for PostgreSQL Processing

You can use PostgreSQL account credentials that are stored in the password file to connect Veeam Agent to the PostgreSQL database system.

If you want to use a password file for authentication, create the `.pgpass` file in the home directory of the root user.

The password file must have the following contents:

```
<hostname>:<port>:<database>:<username>:<password>
```

where:

- `<hostname>` — name of the host where the PostgreSQL database system is located.
- `<port>` — number of the free port that Veeam Agent will use to connect to the PostgreSQL database system.
- `<database>` — name of the PostgreSQL database.
- `<username>` — name of the account that Veeam Agent will use to connect to the PostgreSQL database system.
- `<password>` — password of the account that Veeam Agent will use to connect to the PostgreSQL database system.

For example:

```
srv01:5432:mydb:postgres:P@ssw0rd
```

For more information about the password file, see PostgreSQL documentation.

Viewing Database Processing Settings

To view database processing settings defined for the backup job, use the following command:

```
veeamconfig aap show --jobid <job_id>
```

or

```
veeamconfig aap show --jobname <job_name>
```

where:

- `<job_id>` — ID of the backup job for which you want to view database processing settings.
- `<job_name>` — name of the backup job for which you want to view database processing settings.

Veeam Agent for Linux displays the following information about database processing settings:

- Oracle processing settings
- MySQL processing settings
- PostgreSQL processing settings

Oracle Processing Settings

Parameter	Description
Oracle processing	Oracle processing settings status. Possible values: • *Required* — Oracle processing settings are enabled for the job. If an error occurs when processing the Oracle database system, Veeam Agent will stop the backup process. • *Try* — Oracle processing settings are enabled for the job. If an error occurs when processing the Oracle database system, Veeam Agent will continue the backup process. • *Disabled* — Oracle processing settings are disabled for the job using command line interface.
Account used for processing	Account used to connect to the Oracle database. Possible values: • *System account (username: <username>)* — if Veeam Agent connects to the Oracle database system with the account of the Veeam Agent machine OS. • *Oracle account (username: <username>)* — if Veeam Agent connects to the Oracle database system with the Oracle account. where *<username>* is a name of the user account that Veeam Agent will use to connect to the Oracle database.
Delete logs over <N> Gb	Veeam Agent displays this information if Veeam Agent is set to delete archived logs that are larger than <N> GB.
Delete logs older <N> Hr	Veeam Agent displays this information if Veeam Agent is set to delete archived logs that are older than <N> hours.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
Oracle processing: required
  Account used for processing: Oracle account (username: sys)
  Delete logs over 10 Gb
```

MySQL Processing Settings

Parameter	Description
MySQL processing	MySQL processing settings status. Possible values: • *Required* — MySQL processing settings are enabled for the job. If an error occurs when processing a MySQL database, Veeam Agent will stop the backup process. • *Try* — MySQL processing settings are enabled for the job. If an error occurs when processing a MySQL database, Veeam Agent will continue the backup process. • *Disabled* — MySQL processing settings are disabled for the job using command line interface.
Account used for processing	Veeam Agent displays this information if Veeam Agent is set to connect to the MySQL database system with the account name and password.
Path to a password file	Veeam Agent displays this information if Veeam Agent is set to connect to the MySQL database system with the account credentials that are stored in the password file.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
MySQL processing: required
  Account used for processing: username: root
```

PostgreSQL Processing Settings

Parameter	Description
PostgreSQL processing	PostgreSQL processing settings status. Possible values: • *Required* — PostgreSQL processing settings are enabled for the job. If an error occurs when processing a PostgreSQL database, Veeam Agent will stop the backup process. • *Try* — PostgreSQL processing settings are enabled for the job. If an error occurs when processing a PostgreSQL database, Veeam Agent will continue the backup process. • *Disabled* — PostgreSQL processing settings are disabled for the job using command line interface.
Account used for processing	Account used to connect to the PostgreSQL database. Possible values: • *<username> (password)* — Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the PostgreSQL account. • *<username> (file)* — Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the password file. • *<username> (peer)* — Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the Veeam Agent machine OS account. where *<username>* is a name of the account that Veeam Agent will use to connect to the PostgreSQL database.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
PostgreSQL processing: required
  Account used for processing: postgres (password)
```

Disabling Database Processing Settings

To disable database processing settings defined for the backup job, use the following command:

```
veeamconfig aap disable <db_sys> --jobid <job_id>
```

or

```
veeamconfig aap disable <db_sys> --jobname <job_name>
```

where:

- `<db_sys>` — name of the database system that you want to disable. Possible values:
  - *oracle* — Oracle database processing to be disabled.
  - *mysql* — MySQL database processing to be disabled.
  - *postgres* — PostgreSQL database processing to be disabled.
- `<job_id>` — ID of the backup job for which you want to disable database processing settings.
- `<job_name>` — name of the backup job for which you want to disable database processing settings.

For example:

```
user@srv01:~$ veeamconfig aap disable oracle --jobid 29bc2e1a-e35c-4efb-8d37-b7
177b8ea759
Oracle processing was disabled.
```

Starting and Stopping Backup Jobs

You can start a backup job manually at any time you need, for example, if you want to create an additional restore point for Veeam Agent backup and do not want to change the job schedule. You can also stop the running backup job before the job session completes, if necessary.

You can start and stop backup jobs in one of the following ways:

- With the Veeam Agent control panel
- With the Veeam Agent command line interface

Starting Backup Job from Control Panel

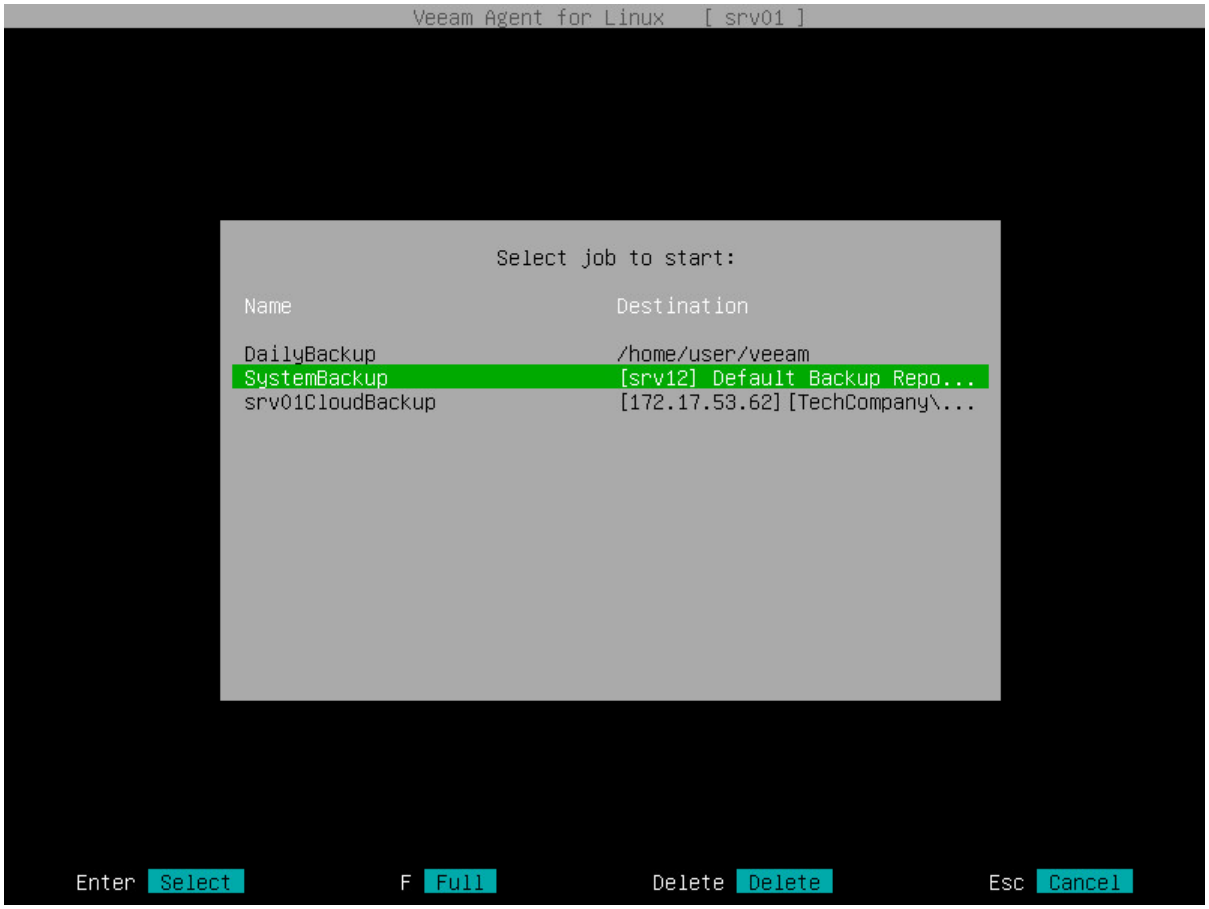
You can start a backup job with the Veeam Agent control panel.

> **NOTE**
>
> Veeam Agent can currently perform only one backup job at a time. You cannot start a backup job when another backup job is already running.

To start a backup job:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. Press the '**s**' key to open the **Select job to start** dialog window.
3. Select the necessary backup job in the list and start the job in on of the following ways:
   - To start an incremental backup job session, press **Enter**.
   - To create an active full backup, press '**f**'.
4. Veeam Agent will immediately start the backup job and display a notification window informing that the job has been started. Press **Enter** to close the window and proceed to the list of backup job sessions.

   You can monitor the backup job performance in the Veeam Agent control panel. To learn more, see Viewing Real-Time Job Session Statistics.



Starting Backup Job from Command Line Interface

You can start a backup job with the command line interface. When you start a backup job, Veeam Agent initiates a new backup job session and provides you with a Session ID. You can monitor the progress of the backup job session or view the session status.

> **NOTE**
>
> Veeam Agent can currently perform only one backup job at a time. You cannot start a backup job when another backup job is already running.

To start a backup job, use the following command:

```
veeamconfig job start --name <job_name>
```

or

```
veeamconfig job start --id <job_id>
```

where:

- `<job_name>` — name of the backup job that you want to start.
- `<job_id>` — ID of the backup job that you want to start.

> **TIP**
>
> Consider the following:
>
> - You can use the `veeamconfig job start` command with the `--nosnap` option to start a file-level backup job. In this case, Veeam Agent will not create a snapshot of the backed-up volume during the backup job session. Keep in mind that the snapshot-less file-level backup does not guarantee that data in the backup is consistent. To learn more, see Snapshot-Less File-Level Backup.
> - You can use the `veeamconfig job start` command with the `--activefull` option to create active full backups. To learn more, see Creating Active Full Backups.

For example:

```
$ veeamconfig job start --name SystemBackup
Backup job has been started.
Session ID: [{381532f7-426a-4e89-b9fc-43d98942c71a}].
Logs stored in: [/var/log/veeam/Backup/SystemBackup/Session_20161207_162608_{38
1532f7-426a-4e89-b9fc-43d98942c71a}].
```

You can check the backup job session status or view the backup job session log using the Veeam Agent command line interface.

You can also monitor the backup job performance in the Veeam Agent control panel. To learn more, see Viewing Real-Time Job Session Statistics.

Creating Active Full Backups

You can create an ad-hoc full backup — active full backup, and add it to the backup chain on the target storage. 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 on the target storage until it is removed from the backup chain according to the retention policy.

Before you create an active full backup, check the following prerequisites:

- The backup job must be configured.
- You cannot create an active full backup if a backup task of any type is currently running.

To perform active full backup, use the following command:

```
veeamconfig job start --name <job_name> --activefull
```

or

```
veeamconfig job start --id <job_id> --activefull
```

where:

- `<job_name>` — name of the backup job that you want to start to create an active full backup.
- `<job_id>` — ID of the backup job that you want to start to create an active full backup.

For example:

```
$ veeamconfig job start --name SystemBackup --activefull
Backup job has been started.
Session ID: [{ce864e24-8211-4df7-973a-741adce96fe7}].
Logs stored in: [/var/log/veeam/Backup/SystemBackup/Session_20180611_150046_{ce
864e24-8211-4df7-973a-741adce96fe7}].
```

You can view the progress for the active full backup session in the same way as for any other backup job session. In particular, you can check the backup job session status or view the backup job session log using the Veeam Agent command line interface.

You can also monitor the backup job performance in the Veeam Agent control panel. To learn more, see Viewing Real-Time Job Session Statistics.

Stopping Backup Job

You can stop the running backup job before the job session completes, for example, if the backup process is about to take long, and you do not want the job to produce workload on the production environment during business hours.

When you stop a backup job, the job session will finish immediately. Veeam Agent will not produce a new restore point during the session, and the session will finish with the *Failed* status.

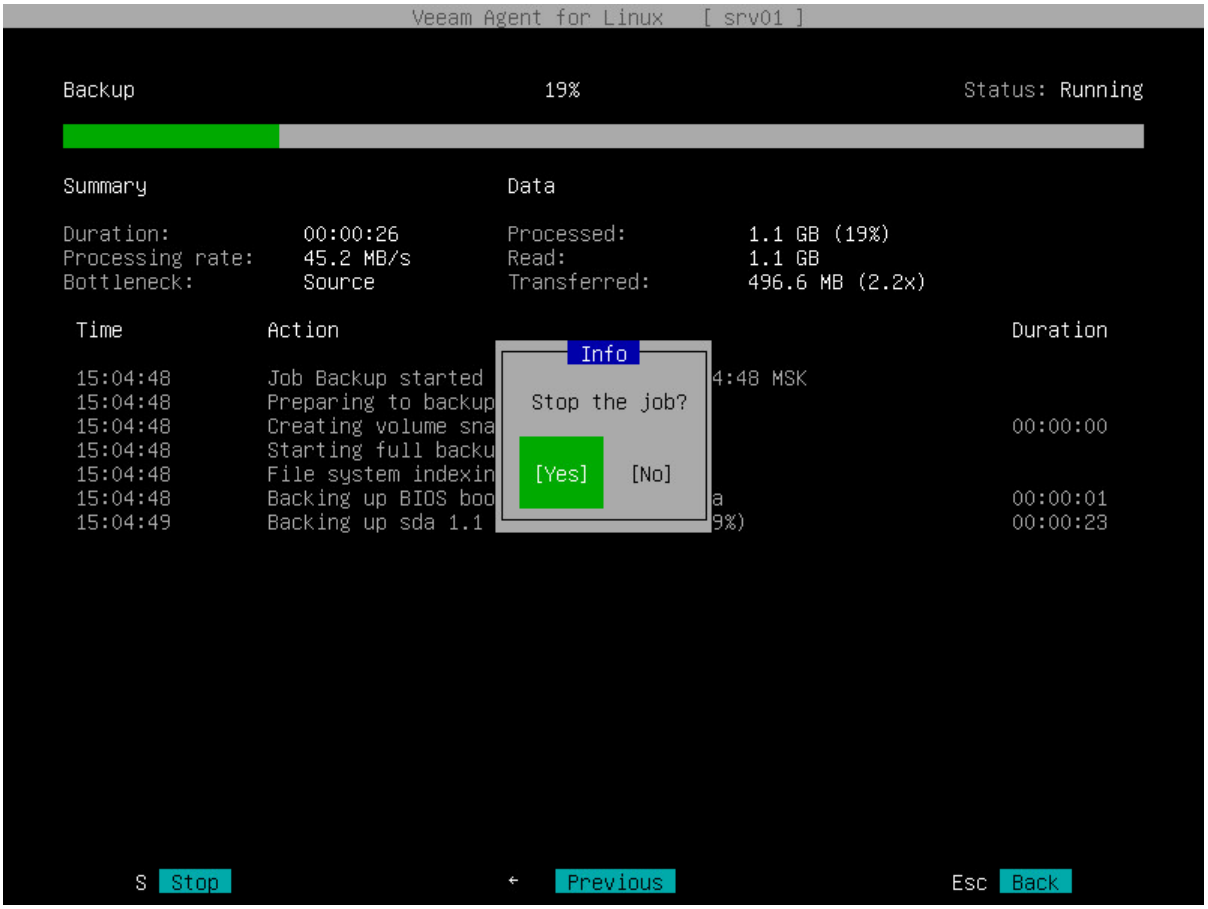
You can stop a job in one of the following ways:

- With the control panel
- With the command line interface

Stopping Job from Control Panel

To stop a backup job:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, select the currently running session with **Up** and **Down** keys and press **Enter**.
3. In the session statistics window, press '**s**'.
4. In the displayed window, make sure that the **Yes** button is selected and press **Enter**.



Stopping Job from Command Line Interface

To stop a backup job, use the following command:

```
veeamconfig session stop --id <session_id>
```

or

```
veeamconfig session stop --force --id <session_id>
```

where:

- `<session_id>` — ID of the currently running backup job session that you want to stop.
- `--force` — with this option enabled, Veeam Agent will immediately stop the backup session even if it is unable to stop the *veeamjobman* process for some reason.

For example:

```
$ veeamconfig session stop --id 381532f7-426a-4e89-b9fc-43d98942c71a
Session has stopped.
```

Managing Backup Jobs

You can perform the following actions with backup jobs configured in Veeam Agent for Linux:

- View the list of configured backup jobs.
- View information about the backup job settings.
- Edit the backup job settings.
- Delete a backup job.

Viewing List of Backup Jobs

To view a list of backup jobs configured in Veeam Agent for Linux, use the following command:

```
veeamconfig job list
```

In the list of backup jobs, Veeam Agent for Linux displays the following information:

Parameter	Description
Name	Name of the backup job.
ID	ID of the backup job.
Repository	Name of the backup repository that is specified as a backup storage for the backup job.

For example:

```
user@srv01:~$ veeamconfig job list
Name                ID                                      Repository
SystemBackup        {2495911e-58db-4452-b4d1-f53dcfbc600e}  Repository_1
DocumentsBackup     {bcf821e6-b35f-4d57-b1c3-d3a477605cb9}  Repository_1
HomePartitionBackup {2aaa8c71-2434-4f12-a168-3d8e225fa416}  Repository_2
```

Viewing Backup Job Settings

To view detailed information about the backup jobs settings, use the following command:

```
veeamconfig job info --name <job_name>
```

or

```
veeamconfig job info --id <job_id>
```

where:

- `<job_name>` — name of the backup job for which you want to view settings.
- `<job_id>` — ID of the backup job for which you want to view settings.

Veeam Agent for Linux displays the following information about the backup job:

Parameter	Description
ID	ID of the backup job.
Name	Name of the backup job.
Repository ID	ID of the backup repository that is specified as a backup storage for the backup job.
Repository name	Name of the backup repository that is specified as a backup storage for the backup job.
Creation time	Date and time of the backup job creation.
Compression	Data compression level. Possible values are: • *0* — No compression • *1* — Rle • *2* — Lz4 • *3* — ZlibLow • *4* — ZlibHigh
Max Points	Number of restore points to keep on disk. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.
Index	File system indexing options defined for the backup job.

Parameter	Description
Objects for backup	Backup scope specified for the backup job.

For example:

```
user@srv01:~$ veeamconfig job info --name SystemBackup
Backup job
  ID: {2495911e-58db-4452-b4d1-f53dcfbc600e}
  Name: SystemBackup
  Repository ID: {4557ef7a-9c44-4f28-b8d0-44d78e5ddd5d}
  Repository name: Repository_1
  Creation time: 2017-04-06 13:29:03
  Options:
     Compression: Lz4
     Max Points: 7
     Index all mounted filesystems on the volumes selected for backup
  Objects for backup:
  Include Disk: sda1
```

Editing Backup Job Settings

If you want to change settings of the backup job, you can edit it at any time. For example, you may want to edit the backup job to add a new directory to the backup scope or change the target location.

To edit a backup job, use the following command:

For volume-level backup jobs

```
veeamconfig job edit volumelevel <option> for --name <job_name>
```

or

```
veeamconfig job edit volumelevel <option> for --id <job_id>
```

For file-level backup jobs

```
veeamconfig job edit filelevel <option> for --name <job_name>
```

or

```
veeamconfig job edit filelevel <option> for --id <job_id>
```

where:

- `<option>` — option that you want to edit for the job. You can specify one or several options at a time. To learn more about available options, see Backup Job Options.
- `<job_name>` — name of the backup job that you want to edit.
- `<job_id>` — ID of the backup job that you want to edit.

For example:

```
user@srv01:~$ veeamconfig job edit volumelevel --name SystemVolumeBackup for --
name SystemVolume
```

Backup Job Options

You can use the following options to edit parameters for the backup job:

Option	Description and values
--compressionlevel	Data compression level. Possible values are: • *0* — No compression • *1* — Rle • *2* — Lz4 • *3* — ZlibLow • *4* — ZlibHigh
--blocksize	Data block size in kilobytes. Possible values are 256, 512, 1024, 4096 or 8192. The default value is *1024*.
--maxpoints	Number of restore points that you want to store in the backup location. By default, Veeam Agent keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent will remove the earliest restore point from the backup chain.
--prefreeze	Pre-freeze command that should be executed before the snapshot creation.
--postthaw	Post-thaw command that should be executed after the snapshot creation.
--objects	Object that should be included in backup: • For simple volumes — name of a block device that represents a volume or an entire disk that should be included in backup. You can specify entire disk to create backup of the entire computer image or individual computer volumes to create backup of specific volumes. If you want to back-up several disks or volumes, specify them one after another using a ',' (comma) character as a separator. • For LVM volumes — name of an LVM logical volume that should be included in backup. If you want to back-up several LVM logical volumes, specify them one after another using a ',' (comma) character as a separator. This option is available for volume-level backup jobs only.
--includedirs	Full path to a directory that should be included in backup, for example: `/home/user`. The option is available for file-level backup jobs only. You can specify one or several paths to directories in the computer file system. To separate several paths, use a ',' (comma) character, for example: `/home/user/Documents,/home/user/reports`.

Option	Description and values
--excludedirs	Full path to a directory that should be excluded from backup. The option is available for file-level backup jobs only. The directory specified with this option must be a subdirectory of the directory specified with the `--includedirs` option. To separate several paths, use a ',' (comma) character, for example, `/home/user/Documents,/home/user/reports`.
--includemasks	Mask for file name or path that should be included in backup. The option is available for file-level backup jobs only. You can use the following UNIX wildcard characters for file name 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`. • '`[]`' — a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '`-`' character). For example: `report_201[3456].pdf` or `report_201[3-6].pdf`. To separate several masks, use a ',' (comma) character, for example, `report.*,reports.*`. File inclusion option is applied to all directories that are specified with the `--includedirs` option. For example, if you include in backup the `/home/user/Documents` directory and files that match the `repor?.pdf` file name mask, Veeam Agent will back up the `/home/user/Documents/report.pdf` file and will not back up the `/home/user/reports/report.pdf` file.

Option	Description and values
--excludemasks	Mask for file name or path that should be excluded from backup. The option is available for file-level backup jobs only. You can use the following UNIX wildcard characters for file name 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`. • '[]' — a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '–' character). For example: `report_201[3456].pdf` or `report_201[3-6].pdf`. To separate several masks, use a ',' (comma) character, for example, `report.*,reports.*`. File exclusion option is applied to all directories that are specified with the `--includedirs` option and files that match file name masks specified with the `--includemasks` option. For example, you may want to specify the following backup scope for the backup job: • Include in backup the `/home/user/Documents` directory • Include files that match the `report.*` file name mask • Exclude files that match the `*.odt` file name mask. In this case, Veeam Agent will back up the `/home/user/Documents/report.pdf` file and will not back up `/home/user/Documents/report.odt` and `/home/user/reports/report.pdf` files. If you want to use several name masks, you must specify them in double quotation marks, for example: `veeamconfig job create filelevel --name BackupJob1 --reponame vault13 --includedirs /home --includemasks "*.bak,*.pdf"`.
--indexnothing	Defines that file system indexing options are disabled for the backup job.
--indexall	Defines that Veeam Agent must index all files on the volumes included in backup.
--indexonly	Path to a directory that contains files that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character. The option is available for volume-level backup jobs only.
--indexexcept	Path to a directory that contains files that you do not want to index. You can specify one or more paths. To separate several paths, use the ',' (comma) character. The option is available for volume-level backup jobs only.

Option	Description and values
--setencryption	Defines that data encryption option is enabled for the job. You can use this option to enable encryption for the existing backup job or change a password used for encryption for the backup job. When you use the `veeamconfig job edit` command with the `--setencryption` option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.
--resetencryption	Defines that data encryption option is disabled for the job. You can use this option to disable encryption for the existing backup job.
--deleteold	The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location.Possible values are: 1-999. If you do not specify the `--deleteold` option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually. If you specified the value earlier and want to disable this setting, specify the *false* value for this option: `--deleteold false`. After the next successful backup session, this setting will be disabled for the backup in the target location.
--nosnap	Defines whether Veeam Agent must perform backup in the snapshot-less mode. Possible values: • *true* — if you use this option, Veeam Agent will create a snapshot of the backed-up volumes during file-level backup. • *false* — if you use this option, Veeam Agent will not create a snapshot of the backed-up volumes during file-level backup. Keep in mind that the snapshot-less file level backup does not guarantee that data in the backup is consistent. To learn more, see Snapshot-Less File-Level Backup.

NOTE

Consider the following:

- If you change the target location for the backup job, during the next backup job session Veeam Agent for Linux will perform full data backup. All subsequent backup sessions will produce incremental backups — Veeam Agent for Linux will copy only changed data to the target location and add a new incremental backup file to the backup chain.
- If you change the backup scope for the backup job, during the next backup job session Veeam Agent for Linux will create a new incremental backup. The backup will contain all data blocks pertaining to new data added to the backup scope and changed data blocks pertaining to original data in the backup scope (data that was processed by the job at the time before you changed the backup scope).
- If you enable or disable encryption for the existing backup job that has already created one or more restore points, during the next job session, Veeam Agent for Linux will create active full backup.
- Full backup takes much more time than incremental backup. If you change the target location, you can copy an existing backup chain to the new location manually. In this case, the new backup job session will produce an incremental backup file and add it to the backup chain.

Deleting Backup Job

You can delete a backup job configured in Veeam Agent for Linux. When you delete a backup job, backup files created by this job remain intact on the backup repository.

You can delete backup jobs in one of the following ways:

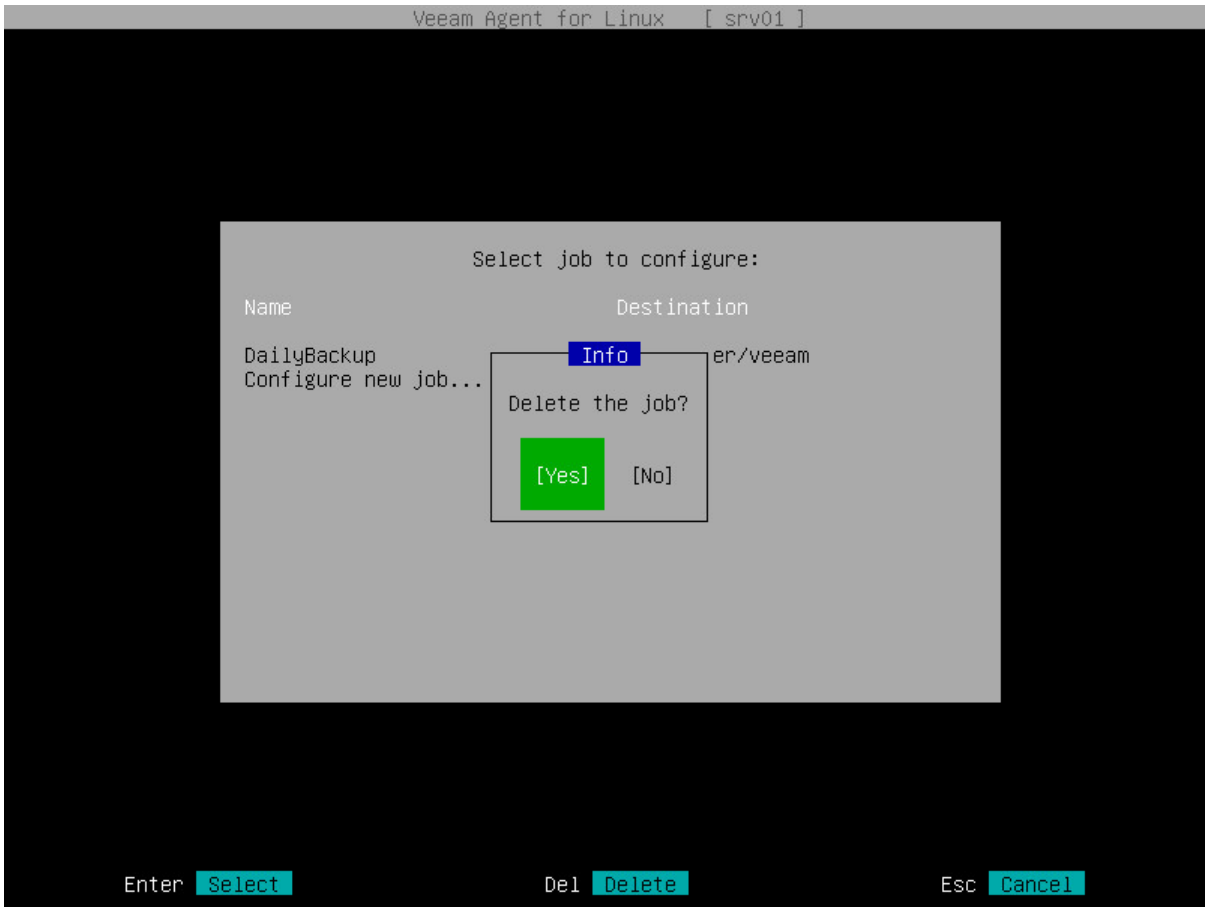
- With the Veeam Agent for Linux control panel
- With the Veeam Agent for Linux command line interface

Deleting Backup Job with Control Panel

You can delete a backup job with the Veeam Agent control panel.

To delete a backup job:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. Press the '**c**' key to open the **Select job to configure** dialog window or the '**s**' key to open the **Select job to start** dialog window.
3. Select the necessary backup job in the list and press **Delete**.
4. In the displayed notification window, make sure that the **Yes** button is selected and press **Enter**.



Deleting Backup Job with Command Line Interface

You can delete a backup job with the Veeam Agent command line interface. To delete a backup job, use the following command:

```
veeamconfig job delete --name <job_name>
```

or

```
veeamconfig job delete --id <job_id>
```

where:

- `<job_name>` — name of the backup job that you want to delete.
- `<job_id>` — ID of the backup job that you want to delete.

For example:

```
$ veeamconfig job delete --name SystemBackup
```

Managing Backup Repositories

A backup repository is a storage location where Veeam Agent for Linux keeps backup files. You can use the following types of disk-based storage as a target location for a backup job:

- Local (internal) storage of the protected machine (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share.
- Veeam Backup & Replication 11 or later backup repository.
- Veeam Backup & Replication 11 or later cloud repository.

Creating Backup Repository

Veeam Agent for Linux creates backup files in a backup repository. When you create a backup job with the Backup Job wizard, you must specify a target location for the backup. Veeam Agent will create a backup repository in the specified location and save information about this repository in the database.

> **IMPORTANT**
>
> A backup repository must be created on a separate volume from a volume whose data you plan to back up.

If you want to create backups in local directory or network shared folder, you must create a repository. To learn more, see the following sections:

- Creating a repository in a local directory
- Creating a repository in an NFS network shared folder
- Creating a repository in an SMB network shared folder

If you want to create Veeam Agent backups in a Veeam backup repository of cloud repository, you do not need to create repositories. Before configuring the backup job, you need to connect to the Veeam backup server or Veeam Cloud Connect service provider. To learn more, see the following sections:

- Connecting to Veeam Backup Server
- Connecting to Service Provider

Creating Repository in Local Directory

To create a repository in a local directory, use the following command:

```
veeamconfig repository create --name <repository_name> --location <path_to_repo
sitory>
```

where:

- `<repository_name>` — name of the repository.
- `<path_to_repository>` — path to the directory in which backup files will be stored.

For example:

```
$ veeamconfig repository create --name VeeammBackup --location /home/backups
```

Creating Repository in NFS Share

To create a repository in an NFS share, use the following command:

```
veeamconfig repository create --name <repository_name> --type nfs --location <p
ath_to_repository>
--options <mounting_options>
```

where:

- `<repository_name>` — name of the backup repository.
- `<path_to_repository>` — path to the network shared folder where backup files will be stored in the *SERVER:/DIRECTORY* format.
- `<mounting_options>` — additional options that Veeam Agent will use to mount the network shared folder to the Veeam Agent machine file system. You can use the standard Linux `mount` command content as mounting options. This parameter is optional.

For example:

```
$ veeamconfig repository create --name VeeamBackup --type nfs --location srv01:
/VeeamRepository --options vers=3,hard,retry=1
```

> **TIP**
>
> If you mount a network shared folder to a directory in the Veeam Agent machine file system in advance, you can create the backup repository in the same way as in a local directory. For details, see Creating Repository in Local Directory.

Creating Repository in SMB Share

To create a repository in an SMB share, use the following command:

```
veeamconfig repository create --name <repository_name> --type smb --location <p
ath_to_repository>
--username <user_name> --password --domain <domain> --options <mounting_options
>
```

where:

- `<repository_name>` — name for the backup repository.
- `<path_to_repository>` — path to the network shared folder where backup files will be stored in the *//SERVER/DIRECTORY* format.
- `<user_name>` — account name that Veeam Agent will use to access the SMB network shared folder.
- `<domain>` — domain in which the account that has access permissions on the shared folder is registered.
- `<mounting_options>` — options that Veeam Agent will use to mount the network shared folder to the Veeam Agent machine file system. You can use the standard Linux `mount` command content as mounting options. This parameter is optional.

You can specify account name and domain for the SMB network shared folder using the `--username` and `--domain` parameters. If a password is required to access the network shared folder, you must also specify the `--password` parameter. When you run the `veeamconfig repository create` command, Veeam Agent will prompt you to type a password of the specified account.

Alternatively, you can specify account name, password and domain for the network shared folder as values for the `--options` parameter. Mind that these values will override values of the `--username`, `--password` and `--domain` parameters.

Examples

Command with `--username`, `--password` and `--domain` parameters:

```
$ veeamconfig repository create --name VeeamBackup --type smb --location //srv0
2/VeeamRepository --username Administrator --password --domain srv02
```

Command with `--options` parameter:

```
$ veeamconfig repository create --name VeeamBackup --type smb --location //srv0
2/VeeamRepository --options username=Administrator,password=P@ssw0rd,domain=srv
02,port=666
```

TIP

If you mount a network shared folder to a directory in the Veeam Agent machine file system in advance, you can create the backup repository in the same way as in a local directory. For details, see Creating Repository in Local Directory.

Viewing List of Backup Repositories

To view backup repositories configured in Veeam Agent for Linux, use the following command:

```
veeamconfig repository list
```

Veeam Agent will display a list of backup repositories.

You can view the following information about backup repositories:

Parameter	Description
Name	Name of the backup repository.
ID	ID of the backup repository.
Location	Directory in the local file system specified as a target location for backup files.
Type	Type of the backup repository. Possible values: • Local • Backup server
Backup server	Backup server on which Veeam backup repository added to Veeam Agent is configured.

For example:

```
user@srv01:~$ veeamconfig repository list
Name         ID                                       Location        Type   Bac
kup server
BackupVol01  {818e3a0f-8155-4a51-9430-248a203a43d1}  /home/backups   loca
l
BackupVol02  {2155a2e7-a1e9-4347-9d8b-cf8f3a6f3fcb}  /home/backups2  loca
l
```

Editing Backup Repository

You can edit settings for a backup repository created with Veeam Agent for Linux.

You can edit the following parameters for the backup repository:

- Name of the backup repository
- Location of the backup repository

> **NOTE**
>
> Consider the following:
>
> - If you change location for the backup repository that is already used by a backup job and contains backup files, during the next backup job run, Veeam Agent will create a new backup chain in the new repository location.
> - You can temporarily change backup repository location if you want to create an ad hoc full backup in addition to the backup chain created by the backup job in the original repository location.

Changing Backup Repository Name

To change a name for the backup repository, use the following command:

```
veeamconfig repository edit --name <new_name> for --name <old_name>
```

or

```
veeamconfig repository edit --name <new_name> for --id <id>
```

where:

- `<old_name>` — current name of the backup repository.
- `<new_name>` — desired name for the backup repository.
- `<id>` — ID of the backup repository.

For example:

```
user@srv01:~$ veeamconfig repository edit --name LocalRepository for --name Rep
ository_1
```

Changing Backup Repository Location

To change location for the backup repository, use the following command:

```
veeamconfig repository edit --location <path> for --name <name>
```

or

```
veeamconfig repository edit --location <path> for --id <id>
```

where:

- `<path>` — desired path for the backup repository.
- `<name>` — current name of the backup repository.
- `<id>` — ID of the backup repository.

For example:

```
user@srv01:~$ veeamconfig repository edit --location /home/veeam for --id 34587
97-3ffe-45bc-870e-c5628643bbb3
```

Changing Backup Repository Name and Location

You can change a name and location for the backup repository at the same time, for example:

```
user@srv01:~$ veeamconfig repository edit --name LocalRepository --location /ho
me/veeam for --name Repository_1
```

Rescanning Veeam Backup Repository

If Veeam Agent for Linux fails to display backups stored in the Veeam Backup & Replication backup repository for some reason, you can rescan the Veeam backup repository. Veeam Agent will try to reconnect to the Veeam backup server and refresh the list of backups in the backup repository.

To rescan a Veeam backup repository, use the following command:

```
veeamconfig repository rescan --id <repository_id>
```

or

```
veeamconfig repository rescan --name <repository_name>
```

where:

- `<repository_id>` — ID of the backup repository that you want to rescan.
- `<repository_name>` — name of the backup repository that you want to rescan.

For example:

```
user@srv01:~$ veeamconfig repository rescan --name [vbr01]BackupVol01
```

You can also rescan all Veeam backup repositories managed by the backup server to which Veeam Agent is connected with the following command:

```
veeamconfig repository rescan --all
```

TIP

With commands mentioned above, you can also rescan local backup repositories. This may be useful, for example, after information about a backup stored in the local repository is deleted from the Veeam Agent configuration database, or after you copy a backup to the local repository.

Deleting Backup Repository

You can delete a backup repository configured with Veeam Agent for Linux. When you delete a backup repository, Veeam Agent removes record of the deleted repository from its database. Backup files created by a backup job targeted at the deleted backup repository remain intact on the backup storage.

To delete a backup repository, use the following command:

```
veeamconfig repository delete --id <repository_id>
```

or

```
veeamconfig repository delete --name <repository_name>
```

where:

- `<repository_id>` — ID of the backup repository that you want to delete.
- `<repository_name>` — name of the backup repository that you want to delete.

For example:

```
user@srv01:~$ veeamconfig repository delete --name Repository_1
```

NOTE

You cannot delete a backup repository that is specified as a backup storage location in the backup job settings.

Managing Veeam Backup & Replication Servers

You can store backup files created with Veeam Agent for Linux on backup repositories managed by Veeam Backup & Replication. To do this, you must connect to a Veeam backup server. After that, you can specify a Veeam backup repository as a target location for backup files in the properties of the backup job.

Connecting to Veeam Backup Server

To create Veeam Agent backups on a backup repository managed by Veeam Backup & Replication, you must connect to a Veeam backup server.

> **IMPORTANT**
>
> Currently, Veeam Agent for Linux can be connected to one Veeam Backup & Replication server only. If you want to create backups on the backup repository managed by another Veeam backup server, you need to delete currently used backup server and all jobs targeted at backup repositories managed by this backup server. To learn more, see Deleting Connection to Veeam Backup Server.
>
> If you add a connection to another backup server, backup jobs targeted at the original backup server will fail, and backups created on the Veeam backup repository will become unavailable in Veeam Agent. To continue using the original backup server, you need to delete the connection to the new backup server and re-create all backup jobs that use the original backup server.
>
> If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain on the backup repository.

To connect Veeam Agent for Linux to a Veeam backup server, use the following command:

```
veeamconfig vbrserver add --name <vbr_name> --address <vbr_address> --port <vbr
_port> --login <username> --domain <domain> --password <password>
```

where:

- `<vbr_name>` — name of the Veeam backup server that manages the backup repository.
- `<vbr_address>` — DNS name or IP address of the Veeam backup server.

  > **IMPORTANT**
  >
  > If you specify a DNS name of the Veeam backup server, make sure that the Veeam backup server name is resolved into IPv4 address on the machine where Veeam Agent is installed. The Veeam Backup Service in Veeam Backup & Replication listens on IPv4 addresses only. If the Veeam backup server name is resolved into IPv6 address, Veeam Agent will fail to connect to the Veeam backup server.

- `<vbr_port>` — port over which Veeam Agent must communicate with Veeam Backup & Replication. The default port used for communication with the Veeam backup server is 10006.
- `<username>` — a name of the account that has access to the Veeam backup repository.
- `<domain>` — a name of the domain in which the account that has access to the Veeam backup repository is registered.
- `<password>` — password of the account that has access to the Veeam backup repository.

  Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see Setting Up User Permissions on Backup Repositories.

For example:

```
user@srv01:~$ veeamconfig vbrserver add --name vbr01 --address 172.17.53.1 --po
rt 10006 --login veeam --domain tech --password P@ssw0rd
```

When Veeam Agent for Linux connects to a Veeam Backup & Replication server, Veeam Agent retrieves information about backup repositories managed by this Veeam backup server and displays them in the list of available backup repositories. You can then specify a Veeam backup repository as a target for a backup job.

> **TIP**
>
> To view the list of backup repositories, use the `veeamconfig repository list` command. To learn more, see Viewing List of Backup Repositories.

Viewing List of Veeam Backup Servers

To view a list of Veeam backup servers to which Veeam Agent for Linux is connected, use the following command:

```
veeamconfig vbrserver list
```

Veeam Agent will display the list of Veeam backup servers.

For the Veeam backup server in the list, Veeam Agent for Linux displays the following information:

Parameter	Description
Name	Name of the Veeam backup server.
ID	ID of the Veeam backup server in the Veeam Agent database.
Endpoint	IP address of the Veeam backup server and port over which Veeam Agent for Linux communicates with Veeam Backup & Replication.

For example:

```
user@srv01:~$ veeamconfig vbrserver list
Name         ID                                      Endpoint
vbr01        {0fc87c11-6a8d-48c1-8aeb-7f7655738796}  172.17.53.1:10006
```

Viewing Backup Server Details

You can view detailed information about the Veeam backup server to which Veeam Agent for Linux is connected. Use the following command:

```
veeamconfig vbrserver info --name <vbr_name>
```

or

```
veeamconfig vbrserver info --id <vbr_id>
```

where:

- `<vbr_name>` — name of the Veeam backup server.
- `<vbr_id>` — ID of the Veeam backup server in the Veeam Agent database.

Veeam Agent for Linux displays the following information about the Veeam backup server:

Parameter	Description
ID	ID of the Veeam backup server in the Veeam Agent database.
Name	Display name of the Veeam backup server.
Endpoint	IP address of the Veeam backup server and port over which Veeam Agent for Linux communicates with Veeam Backup & Replication.
Login	Name of the account that has access to the Veeam backup repository.
Domain	Name of the domain in which the account that has access to the Veeam backup repository is registered.

For example:

```
user@srv01:~$ veeamconfig vbrserver info --name vbr01
VBR server
  ID: {0fc87c11-6a8d-48c1-8aeb-7f7655738796}
  Name: vbr01
  Endpoint: 172.17.53.1:10006
  Login: veeam
  Domain: tech
```

Editing Connection to Veeam Backup Server

You can edit the following parameters for a connection to a Veeam backup server:

- Display name of the Veeam backup server
- IP address and port used to connect to the Veeam backup server
- Account to connect to the Veeam backup server

Changing Veeam Backup Server Name

To change a name for the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --name <new_vbr_name>
```

where:

`<new_vbr_name>` — desired name for the backup server.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --name vbr01
```

Changing IP Address and Port for Veeam Backup Server

To change the IP address and port used to connect to the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --address <vbr_address> --port <vbr_port>
```

where:

- `<vbr_address>` — DNS name or IP address of the Veeam backup server.

  > **IMPORTANT**
  >
  > If you specify a DNS name of the Veeam backup server, make sure that the Veeam backup server name is resolved into IPv4 address on the machine where Veeam Agent is installed. The Veeam Backup Service in Veeam Backup & Replication listens on IPv4 addresses only. If the Veeam backup server name is resolved into IPv6 address, Veeam Agent will fail to connect to the Veeam backup server.

- `<vbr_port>` — port over which Veeam Agent for Linux must communicate with Veeam Backup & Replication.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --address 172.17.53.1 --port 10006
```

Changing Account to Connect to Veeam Backup Server

> **NOTE**
>
> If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain on the backup repository.

To change an account whose credentials will be used to connect to the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --login <username> --domain <domain> --password
```

where:

- `<username>` — name of the account that has access to the Veeam backup repository.
- `<domain>` — name of the domain in which the account that has access to the Veeam backup repository is registered.

When you run the command, Veeam Agent will prompt you to enter the password of the specified account.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --login veeam --domain tech --password
Enter password:
```

Changing Several Backup Server Parameters

You can change several parameters for the connection to the Veeam backup server simultaneously. For example, the following command changes the name and connection settings for the Veeam backup server:

```
user@srv01:~$ veeamconfig vbrserver edit --name vbr02 --address 172.17.53.2 --p
ort 10006
```

Updating List of Veeam Backup Repositories

When you connect to a Veeam backup server, Veeam Agent for Linux retrieves information about backup repositories managed by this Veeam backup server and displays them in the list of available backup repositories. You can refresh information about available Veeam backup repositories manually at any time. This may be useful, for example, after a new backup repository was added on the Veeam backup server.

To update the list of backup repositories managed by the Veeam backup server, use the following command:

```
veeamconfig vbrserver resync
```

TIP

To view updated list of available Veeam backup repositories after resync, use the `veeamconfig repository list` command. To learn more, see Viewing List of Backup Repositories.

Deleting Connection to Veeam Backup Server

You can delete a connection to the Veeam backup server to which Veeam Agent is currently connected. When you delete a connection to a Veeam backup server, Veeam Agent removes record on the deleted backup server from its database. Veeam backup repositories managed by the deleted backup server are removed from the list of available backup repositories. Backup files created by backup jobs targeted these repositories remain intact on the backup storage.

You cannot delete a connection to a Veeam backup server if a backup repository managed by this backup server is used by a backup job. To remove such connection to a Veeam backup server, you first need to delete a reference to the Veeam backup repository in the job settings.

To delete a connection to the Veeam backup server, use the following command:

```
veeamconfig vbrserver delete --name <vbr_name>
```

or

```
veeamconfig vbrserver delete --id <vbr_id>
```

where:

- `<vbr_name>` — name of the Veeam backup server.
- `<vbr_id>` — ID of the Veeam backup server.

For example:

```
user@srv01:~$ veeamconfig vbrserver delete --name vbr01
```

Managing Service Providers

You can store backup files created with Veeam Agent for Linux on a cloud repository exposed to you by a Veeam Cloud Connect service provider. To do this, you must connect to a service provider. After that, you can specify a cloud repository as a target location for backup files in the properties of the backup job.

Connecting to Service Provider

To create Veeam Agent backups on a cloud repository, you must connect to a Veeam Cloud Connect service provider.

To connect Veeam Agent for Linux to a service provider, use the following command:

```
veeamconfig cloud add --name <sp_name> --address <sp_address> --port <sp_port>
--login <username> --password <password> --fingerprint <sp_thumbprint>
```

where:

- `<sp_name>` — name of the service provider to which you want to connect.
- `<sp_address>` — IP address or full DNS name of the cloud gateway that the SP or your backup administrator has provided to you.
- `<sp_port>` — port over which Veeam Agent must communicate with the cloud gateway. The default port used for communication with the cloud gateway is 6180.
- `<username>` — name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT/SUBTENANT* format.
- `<password>` — password of the tenant or subtenant account used to connect to the service provider.
- `<sp_thumbprint>` — thumbprint used to verify the TLS certificate that the SP has provided to you.

For example:

```
user@srv01:~$ veeamconfig cloud add --name SP --address 172.17.53.15 --port 618
0 --login TechCompany/User01 --password P@ssw0rd --fingerprint 92FA988A3D9E80EE
095DDAB75BF06B05DF6F205B
```

NOTE

When you enter the `veeamconfig cloud add` command, Veeam Agent will display information about the TLS certificate obtained from the SP. To accept the certificate, type `yes` in the command prompt and press **Enter**.

When Veeam Agent connects to the service provider, Veeam Agent retrieves information about cloud repositories available to the tenant or subtenant and displays them in the list of available backup repositories. You can then specify a cloud repository as a target for a backup job.

TIP

To view the list of available cloud repositories, use the `veeamconfig repository list` command. To learn more, see Viewing List of Backup Repositories.

Viewing List of Service Providers

To view a list of service providers to which Veeam Agent is connected, use the following command:

```
veeamconfig cloud list
```

Veeam Agent will display the list service providers.

For the service provider in the list, Veeam Agent for Linux displays the following information:

Parameter	Description
Name	Name of the service provider.
ID	ID of the service provider in the Veeam Agent database.
Address	IP address of the cloud gateway and port over which Veeam Agent communicates with the cloud gateway.
Gate servers	IP address of the cloud gateway and port over which Veeam Agent communicates with the cloud gateway.
Username	Name of the tenant or subtenant account used for connection to the service provider.

For example:

```
user@srv01:~$ veeamconfig cloud list
Name        ID                                      Address                Gate
servers  Username
SP          {0840f770-354d-426a-b5ce-1aa80f56cc08}  172.17.53.15:618
0                   TechCompany
```

Editing Connection to Service Provider

You can edit the following parameters for a connection to a Veeam Cloud Connect service provider:

- Display name of the Veeam Cloud Connect service provider
- IP address and port used to connect to the cloud gateway
- Account to connect to the service provider
- Thumbprint to connect to the service provider

Changing SP Name

To change a name for the SP, use the following command:

```
veeamconfig cloud edit --name <new_sp_name> for --name <old_sp_name>
```

or

```
veeamconfig cloud edit --name <new_sp_name> for --id <sp_id>
```

where:

- `<old_sp_name>` — current name of the SP.
- `<new_sp_name>` — desired name for the SP.
- `<sp_id>` — ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --name SP for --id 7d3022de-4f4d-4c70-85eb
-e8a946a555cd
```

Changing IP Address and Port for Cloud Gateway

To change the IP address and port of the cloud gateway provided by the SP, use the following command:

```
veeamconfig cloud edit --address <sp_address> --port <sp_port> for --name <sp_n
ame>
```

or

```
veeamconfig cloud edit --address <sp_address> --port <sp_port> for --id <sp_id>
```

where:

- <sp_address> — IP address or full DNS name of the cloud gateway that the SP or your backup administrator has provided to you.
- <sp_port> — port over which Veeam Agent must communicate with the cloud gateway. The default port used for communication with the cloud gateway is 6180.
- <sp_name> — name of the SP.
- <sp_id> — ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --address 172.17.53.67 --port 6180 for --n
ame SP
```

Changing Account to Connect to SP

To change an account whose credentials will be used to connect to the SP, use the following command:

```
veeamconfig cloud edit --login <username> --password <password> for --name <sp_
name>
```

or

```
veeamconfig cloud edit --login <username> --password <password> for --id <sp_id
>
```

where:

- <username> — name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT/SUBTENANT* format.
- <password> — password of the tenant or subtenant account used to connect to the service provider.
- <sp_name> — name of the SP.
- <sp_id> — ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --login ABC_Compan/User01 --password P@ssw
0rd for --name SP
```

Changing Thumbprint to Connect to SP

To change a thumbprint that will be used to connect to the SP, use the following command:

```
veeamconfig cloud edit --fingerprint <sp_thumbprint> for --name <sp_name>
```

or

```
veeamconfig cloud edit --fingerprint <sp_thumbprint> for --id <sp_id>
```

where:

- `<sp_thumbprint>` — thumbprint used to verify the TLS certificate and connect to the service provider.
- `<sp_name>` — name of the SP.
- `<sp_id>` — ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --fingerprint 92FA988A3D9E80EE095DDAB75BF0
6B05DF6F205B for --name SP
```

Updating List of Cloud Repositories

When you connect to the Veeam Cloud Connect service provider, Veeam Agent for Linux retrieves and saves to the database information about cloud repositories available to the tenant or subtenant whose account you use to connect to the SP. You can refresh information about available cloud repositories manually at any time. This may be useful, for example, after the SP changes backup resource settings for the tenant.

To update the list of cloud repositories, use the following command:

```
veeamconfig cloud resync
```

If the cloud repository currently used as a target location for Veeam Agent backups becomes unavailable, and Veeam Agent fails to reflect this change in its database for some reason, the `veeamconfig cloud resync` command may finish with errors. In this case, you can use the `--force` option to refresh information about available cloud repositories. For example:

```
veeamconfig cloud resync --force
```

With the `--force` option, Veeam Agent will retrieve the list of available cloud repositories from the service provider and save the new information about cloud repositories in the Veeam Agent database.

> **TIP**
>
> To view updated list of available cloud repositories after resync, use the `veeamconfig cloud list` command. To learn more, see Viewing List of Service Providers.

Deleting Connection to Service Provider

You can delete a connection to the service provider to which Veeam Agent for Linux is currently connected. When you delete a connection to a service provider, Veeam Agent removes the record on the deleted service provider from the database. Cloud repositories managed by the deleted service provider are removed from the list of available backup repositories. Backup files created by backup jobs targeted at these repositories remain intact on the cloud repository.

You cannot delete a connection to the service provider if a cloud repository managed by this service provider is used by a backup job. To remove such connection to a service provider, you first need to delete a reference to the cloud repository in the job settings.

To delete a connection to the service provider, use the following command:

```
veeamconfig cloud delete --name <sp_name>
```

or

```
veeamconfig cloud delete --id <sp_id>
```

where:

- `<sp_name>` — name of the service provider.
- `<sp_id>` — ID of the service provider.

For example:

```
user@srv01:~$ veeamconfig cloud delete --name SP
```

Managing Backups

You can perform the following operations with backups created by backup jobs configured in Veeam Agent for Linux:

- View backups
- View backup details
- View restore points in backup
- Export backup to a virtual disk
- Import backup to the Veeam Agent database
- Delete backup

Viewing Backups

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list --all
```

In the list of backups, Veeam Agent displays information about all Veeam Agent for Linux backups stored in all backup repositories. If Veeam Agent is connected to a Veeam backup server, all Veeam Agent for Linux backups that are kept on Veeam backup repositories managed by this server also appear in this list.

The `--all` parameter is optional. If you do not use this option, Veeam Agent displays information about backups created by the current Veeam Agent computer only.

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as *Imported* in the **Repository** column. For information about the import procedure, see Importing Backups.
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name                  Backup ID                                  Repositor
y    Created at
srv01 SystemBackup        {45f074d2-d2d9-423d-84e9-8f1798b08d4c}  Repository_
1  2016-11-11 17:37
srv01 DocumentsBackup     {ea64a7e5-038a-4c86-970a-6d59d4cf3968}  Repository_
1  2016-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b}  Repository_
2  2016-11-15 11:28
wrk01 SystemBackup        {951ac571-dd29-45ac-8624-79b8ccb45863}  Repository_
2  2016-11-13 15:26
wrk02 SystemBackup        {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167}  Repository_
3  2016-11-13 15:59
```

TIP

If you want to recover data from a backup that is stored in another location, for example, a backup created with another instance of Veeam Agent, you can import such backup into the Veeam Agent database on your computer. To learn more, see Importing Backups.

Viewing Backup Details

You can view detailed information about specific backup. To view backup details, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup_id> — ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Name	Name of the volume in the backup.
Device	Path to the block device file that represents the volume.
FS UUID	File system ID.
Offset	Position of the volume on the computer disk.
Size	Size of the volume in the backup.

For example:

```
user@srv01:~$ veeamconfig backup show --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
     Machine name: srv01 SystemBackup
        Name:       [sda1]
        Device:     [/dev/sda1]
        FS UUID:    [6945f2eb-e8bb-48fe-a276-5ba67b9030a5]
        Offset:     [1048576] bytes (2048 sectors)
        Size:       [9999220736] bytes (19529728 sectors)
```

For a file-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Backed up	Backup scope for the file-level backup job.

For example:

```
user@srv01:~$ veeamconfig backup show --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
     Machine name: srv01 DocsBackup
        File-level backup
        Backed up:
           /home/user/Documents
```

Viewing Restore Points in Backup

To view information about restore points in the backup, you can use one of the following commands:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where:

<backup_id> — ID of the backup for which you want to view information on restore points.

For example:

```
user@srv01:~$ veeamconfig backup info --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
```

or

```
user@srv01:~$ veeamconfig point list --backupid 4f75bb20-a6b6-4323-9287-1c6c8ce
ccb6b
```

Veeam Agent for Linux displays the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: • Full • Increment
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: • True • False

Importing Backups

You can import a backup created by Veeam Agent into the Veeam Agent database. For example, you may want to import a previously deleted backup or backup that was created in a network shared folder by Veeam Agent installed on another computer.

To import a backup:

1. Start the import process with the following command:

```
veeamconfig backup import --path <path>
```

where:

<path> — path to the VBM file of the backup that you want to import.

For example:

```
user@srv01:~$ veeamconfig backup import --path /home/share/BackupJob/Backu
pJob.vbm
Backup has been imported successfully.
Session ID: [{4031f058-766c-4f2c-a7ae-7257adb2929f}].
Logs stored in: [/var/log/veeam/Import/Session_{4031f058-766c-4f2c-a7ae-72
57adb2929f}].
```

2. You can monitor the import process and result by viewing the import session log with the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> — ID of the import session.

For example:

```
user@srv01:~$ veeamconfig session log --id 4031f058-766c-4f2c-a7ae-7257adb
2929f
2016-11-19 13:21:33 UTC {765af178-a9cc-4596-8bf2-03850c5da1ac} [info] Job
started at 2016-11-19 16:21:33
2016-11-19 13:21:33 UTC {6ae2922d-454b-4a8d-a11b-2b5c7a85029d} [info] Impo
rting backup
2016-11-19 13:21:33 UTC {783f40a7-ead7-4555-9c35-545d875990ee} [info] Back
up has been imported.
```

3. Imported backup will be displayed in the list of backups. To view the list of backups, use the following command:

```
veeamconfig backup list
```

For example:

```
user@srv01:~$ veeamconfig backup list
Job name            Backup ID                                    Repositor
y    Created at
srv01 SystemBackup  {45f074d2-d2d9-423d-84e9-8f1798b08d4c}  Repository_
1  2016-11-11 17:37
srv01 DocsBackup    {ea64a7e5-038a-4c86-970a-6d59d4cf3968}  Repository_
1  2016-11-11 18:30
srv01 HomeBackup    {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b}  Repository_
2  2016-11-15 11:28
BackupJob           {64957b1d-d219-456c-a9cd-9598292c10cd}  Importe
d      2016-11-19 19:12
```

Importing Encrypted Backups

You can import an encrypted backup created by Veeam Agent into the Veeam Agent database. This operation is required if you want to use the Veeam Agent command line interface to restore data from an encrypted backup created by Veeam Agent running on another computer.

To import an encrypted backup:

1. Start the import process with the following command:

   ```
   veeamconfig backup import --path <path>
   ```

   where:

   `<path>` — path to the VBM file of the backup that you want to import.

   For example:

   ```
   user@srv01:~$ veeamconfig backup import --path /home/share/srv15\ Backup/B
   ackup.vbm
   ```

2. Veeam Agent will prompt you to provide a password for the backup file. Type in the password and press the **Enter key** to import the backup.

   Veeam Agent displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.

   If you enter the correct password, Veeam Agent will decrypt the backup file and import it into the database.

   ```
   user@srv01:~$ veeamconfig backup import --path /home/share/srv15\ Backup/B
   ackup.vbm
   [Info] Backup srv15 Backup encrypted
   [Info] Press "Enter" to skip. Enter password to decrypt the backup:
   [Info] Hint: Standard password
   Password:
   Backup imported successfully
   ```

3. Imported backup will be displayed in the list of backups. To view the list of backups, use the following command:

   ```
   veeamconfig backup list
   ```

   For example:

   ```
   user@srv01:~$ veeamconfig backup list
   Job name            Backup ID                                    Repositor
   y    Created at
   srv15 Backup        {4b1f873c-857d-b984-4f22-6ce66bf62570}  Importe
   d      2018-06-12 20:20
   srv01 ServerBackup  {f212f641-54aa-40de-a0eb-8727be56760b}  Importe
   d      2018-06-12 20:04
   ```

Deleting Backups

Backup files created with Veeam Agent are removed automatically according to the retention policy settings. You can also remove backups from the target location and/or Veeam Agent configuration database manually if necessary.

Removing Backup from Configuration

To remove a backup from the Veeam Agent configuration database, use the following command:

```
veeamconfig backup delete --id <backup_id>
```

where `<backup_id>` is an ID of the backup that you want to delete.

The way Veeam Agent removes a backup from configuration depends on the backup location:

- If the backup resides in a local directory or network shared folder, Veeam Agent removes records about the deleted backup from the Veeam Agent database. Backup files themselves (VBK, VIB, VBM) remain in the backup repository.

  You can import information about the removed backup later to Veeam Agent and perform restore operations with the imported backup. To import information about the removed backup, use the `veeamconfig repository rescan --all` command.

- If the backup resides in a Veeam Backup & Replication repository, Veeam Agent removes records about the deleted backup from the Veeam Agent database and Veeam Backup & Replication database. Backup files themselves (VBK, VIB, VBM) remain in the backup repository.

  If you want to import information about the removed backup later to Veeam Agent and perform restore operations with this backup, you must contact backup administrator working with Veeam Backup & Replication. The administrator must rescan the backup repository that contained the backup in the Veeam Backup & Replication console. For details, see the Rescanning Backup Repositories section in the Veeam Backup & Replication User Guide.

  After rescan, the backup will be displayed in the list of backups on the Veeam Agent machine connected to the Veeam backup server.

Deleting Backup Files

To delete backup files from the target location and Veeam Agent database, use the following command:

```
veeamconfig backup delete --id <backup_id> --purge
```

where `<backup_id>` is an ID of the backup that you want to delete.

Veeam Agent for Linux will remove records about the deleted backup from the Veeam Agent database and, additionally, delete backup files themselves from the destination storage.

Performing Restore

If you experience a problem with your computer, your data gets lost or corrupted, you can use one of the following options to recover your data or bring the computer back to work:

- Restore from the Veeam Recovery Media
  - Restore volumes
  - Restore files and folders
- Restore volumes with the command line interface
- Restore files and folders:
  - Restore files and folders with the File Level Restore wizard
  - Restore files and folders with the command line interface
- Export data as VHD disks
- Restore data from encrypted backups

Restoring from Veeam Recovery Media

If the OS on your computer fails to start, you can use the Veeam Recovery Media to recover your computer. The Veeam Recovery Media will help you boot the computer in the limited mode. After booting, you can use a backup created with Veeam Agent for Linux to restore the whole system image of your computer, specific volumes on your computer or specific files and folders. You can also use standard Linux command line utilities to diagnose problems and fix errors.

> **IMPORTANT**
>
> If you plan to use the custom Veeam Recovery Media, Veeam Agent requires 3 GB RAM or more installed on the target computer or virtual machine. Memory consumption varies depending on size and number of modules included into the recovery media. To learn more, see Creating Custom Veeam Recovery Media.

Restoring Volumes

You can restore a specific computer volume or all volumes from the volume-level backup.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

Before You Begin

Before you boot from the recovery image and recover your data, check the following prerequisites and limitations:

- You must have a recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For volume-level restore, you can use a volume-level backup created with Veeam Agent for Linux. Make sure that the backup or system image is available on the computer drive (local or external), on a network shared folder or on the backup repository managed by a Veeam backup server.
- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a Veeam backup repository, you must have access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.
- You cannot restore a volume to the volume where the backup file that you use for restore is located.
- Veeam Recovery Media 3.0 and later versions do not support Oracle VM virtual machines. Use an earlier version of Veeam Recovery Media or contact Veeam Customer Support for workarounds.

Step 1. Boot from Veeam Recovery Media

To boot from the Veeam Recovery Media:

1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.

   [For removable storage device] Attach the removable storage device with the recovery image to your computer.

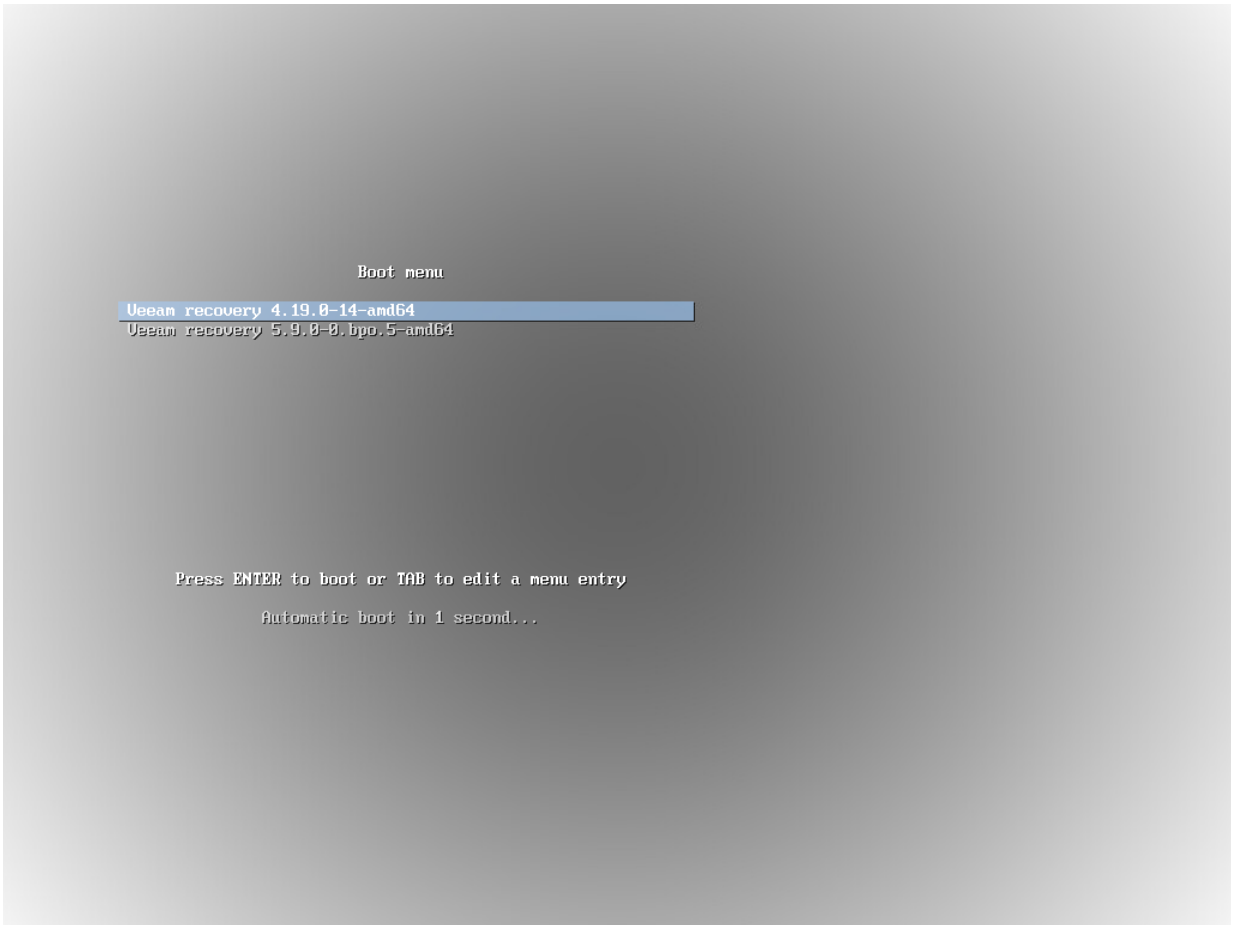
2. Start your computer.

3. [For regular recovery image] In the boot menu, select what Linux kernel version to use to boot your computer and specify boot options if necessary.

   You can select a Linux kernel version if you use generic Veeam Recovery Media downloaded from the Veeam website or Veeam software repository. If you created a custom Veeam Recovery Media, you will be prompted to boot using the Linux kernel of your Veeam Agent computer included in the recovery image.

   To specify boot options, press the **Tab** key and type the necessary options in the command prompt.

   > **NOTE**
   >
   > For the legacy recovery image, the boot menu is unavailable. After you start your computer, Veeam Agent will immediately start loading files from the Veeam Recovery Media.

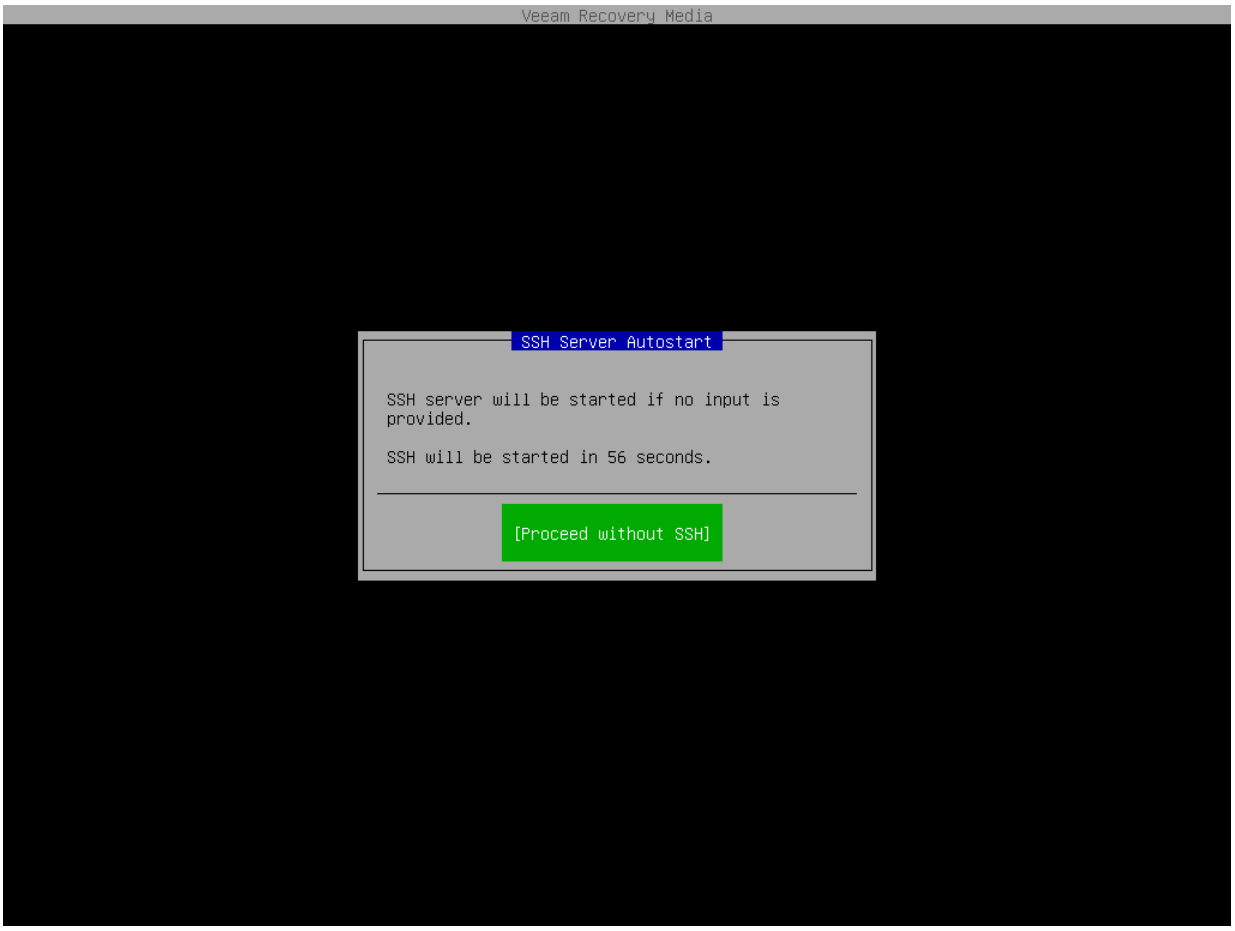


4. Wait for Veeam Agent to load files from the Veeam Recovery Media.

5. After the recovery image OS has loaded, choose whether you want to start the SSH server. The SSH server allows you to connect to the Veeam Recovery Media from a remote machine.

   The Veeam Recovery Media starts the SSH server automatically after a time-out. The default value for the time-out is 60 seconds.

   If you do not want to start the SSH server, make sure that the **Proceed without SSH** button is selected and press **Enter**. You will proceed immediately to the step 7.



6. After the SSH server has started, review settings to connect to the Veeam Recovery Media and press **Enter**.

   The Veeam Recovery Media displays the following connection settings:

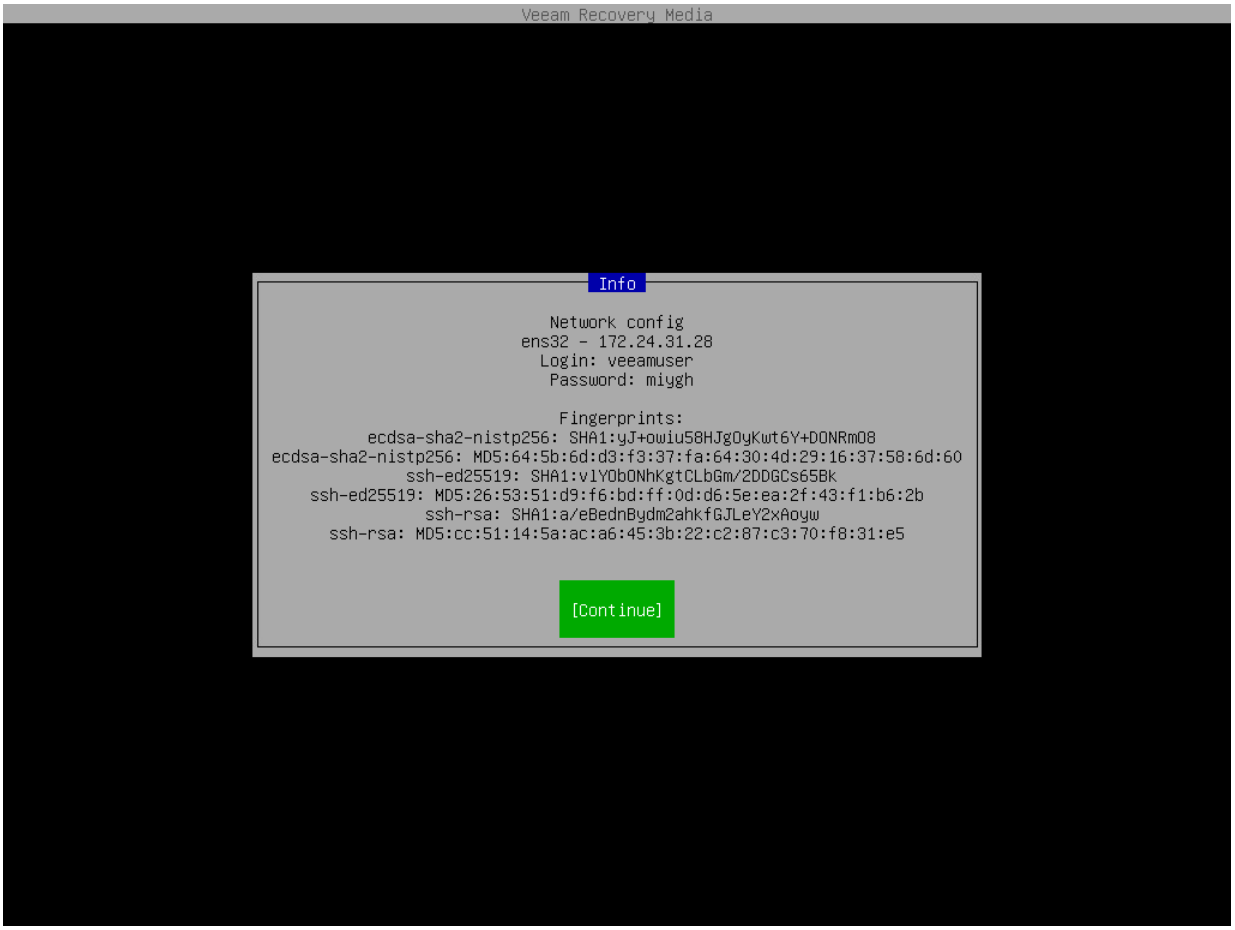
   - IP address of the computer booted from the Veeam Recovery Media
   - User name and password of the account used to connect to the Veeam Recovery Media
   - Fingerprints of the computer booted from the Veeam Recovery Media

NOTE

The user name of the account used to work with the Veeam Recovery Media depends on the recovery image version:

- *veeamuser* — for the regular recovery image
- *root* — for the legacy recovery image

If you want to use command-line utilities built in the regular recovery image, use the `sudo` command to provide the *veeamuser* account with privileges of the *root* account.

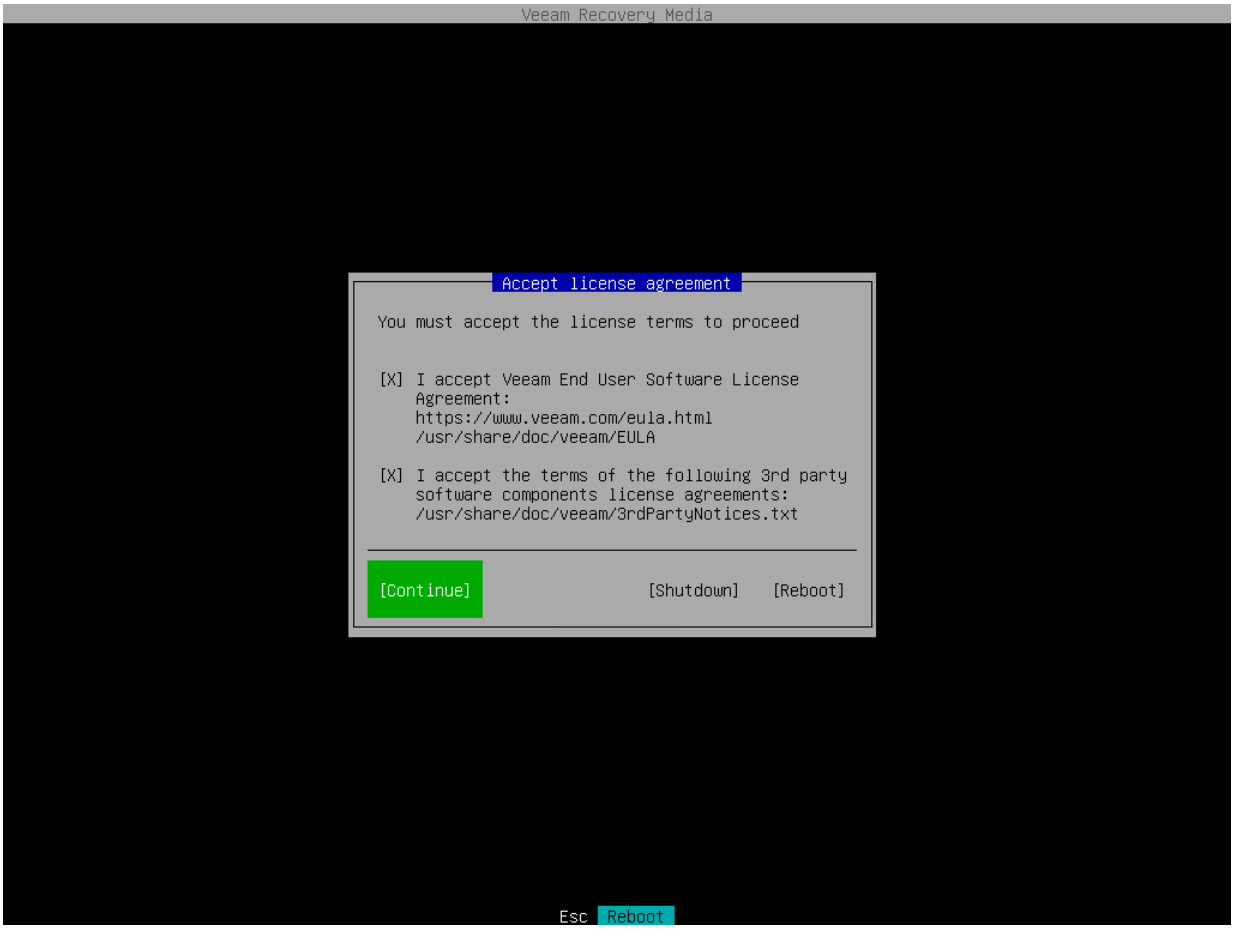


7. Accept the terms of the product license agreement and license agreements for third-party components of the product:

   a. Make sure that the **I accept Veeam End User Software License Agreement** option is selected and press **Space**.

   b. Select the **I accept the terms of the following 3rd party software components license agreements** option with the **Tab** key and press **Space**.

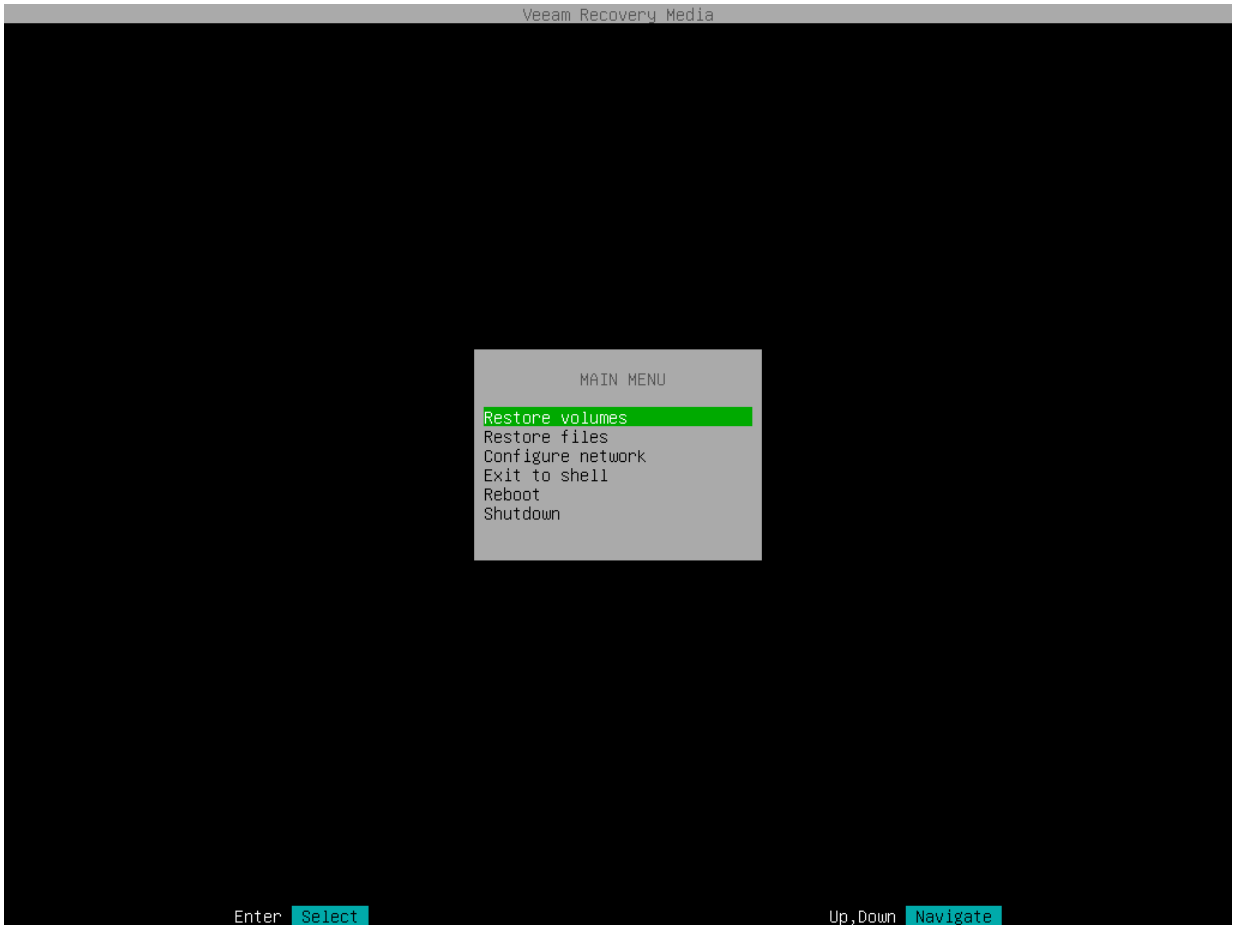
c. Switch to the **Continue** button with the **Tab** key and press **Enter**.



8. Make sure that network settings are specified correctly and configure the network adapter if necessary. To learn more, see Configure Network Settings.
9. Choose the necessary recovery option. Veeam Agent offers the following tools:
   - **Restore volumes** — the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
   - **Restore files** — the File Level Restore wizard to restore files and folders to the original location or to a new location.
   - **Exit to shell** — Linux shell prompt with standard utilities to diagnose problems and fix errors.

TIP

To stop working with the Veeam Recovery Media and shut down or restart your computer, in the Veeam Recovery Media main menu, select the **Reboot** or **Shutdown** option and press **Enter**.

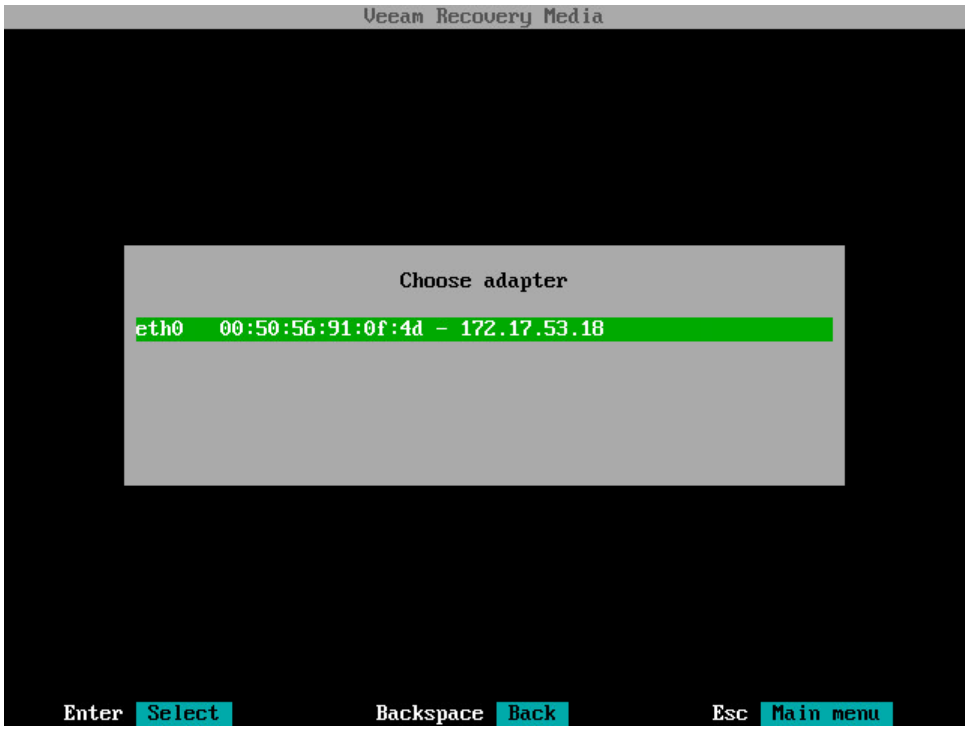


Step 2. Configure Network Settings

To open the **Network settings** dialog, in the Veeam Recovery Media main menu, select the **Configure network** option and press **Enter**.

Veeam Agent for Linux will display network adapters that are available on the system. If there is a DHCP server in your network, Veeam Agent will configure the network settings automatically and display the IP address assigned to the network adapter. You can then press **Esc** to return to the Veeam Recovery Media main menu and launch the Volume Restore wizard.

You can manually configure TPC/IP v4 settings for network adapters if necessary. To learn more, see Specifying Network Settings.



Specifying Network Settings

You can manually configure TPC/IP v4 settings for the network adapter. To configure network settings:

1. In the **Choose adapter** list, select the network adapter that you want to use to connect to the network shared folder or Veeam backup repository where the backup resides and press **Enter**.
2. In the **Configure adapter** dialog, select the **Manual** option and press **Enter**.
3. In the **Adapter settings** dialog, specify the following network settings:
   - IP address
   - Subnet mask
   - Default gateway
   - DNS server
4. Select the **Apply** button and press **Enter**.

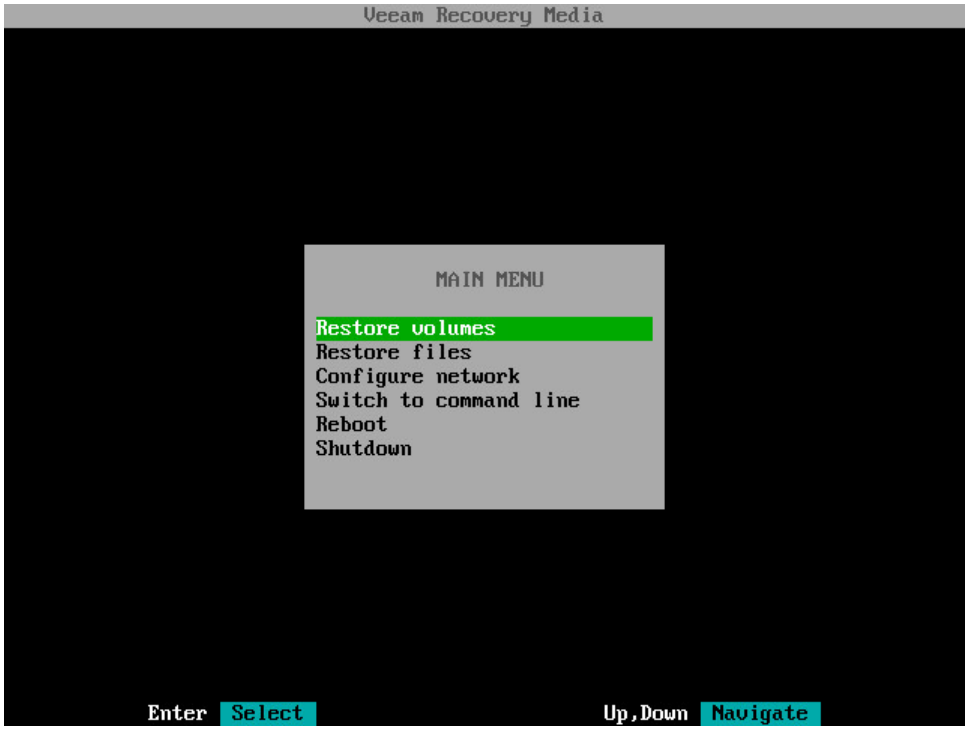
If there is a DHCP server in your network, you can return to automatic IP addressing:

1. In the **Choose adapter** list, select the necessary network adapter and press **Enter**.
2. In the **Configure adapter** dialog, select the **Auto** option and press **Enter**. Veeam Agent will automatically configure network settings for the adapter.



Step 3. Launch Volume Restore Wizard

To launch the volume restore wizard, in the Veeam Recovery Media main menu, select **Restore volumes** and press **Enter**.



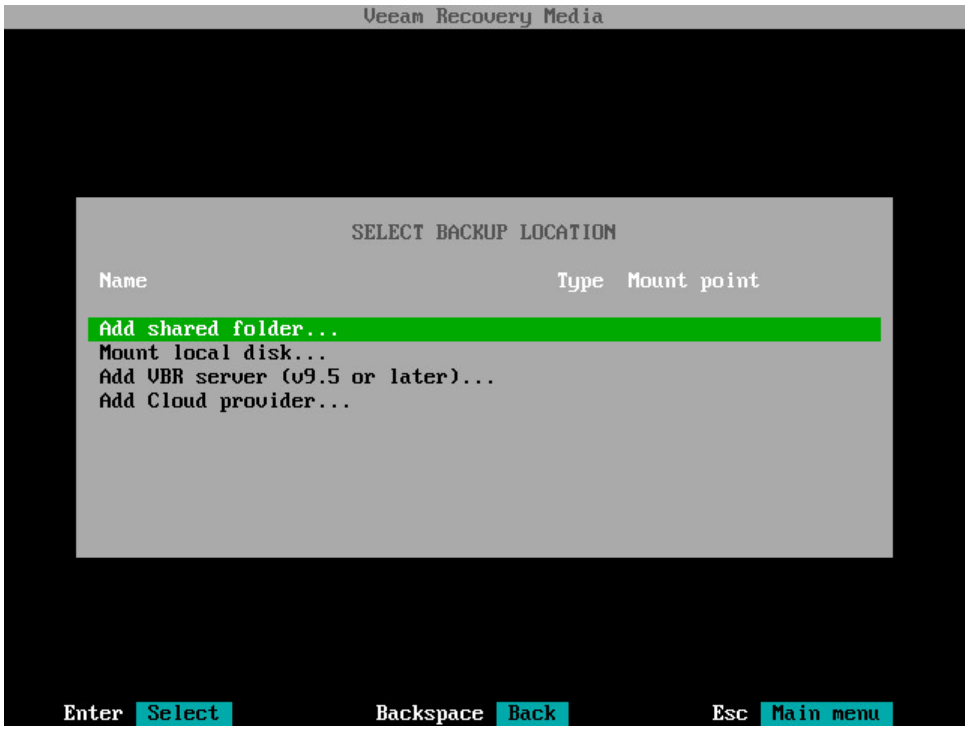
Step 4. Select Backup Location

At the **Select backup location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

To recover data from backup, you need to mount the backup storage on which the backup file resides to the recovery image OS file system. Veeam Agent for Linux automatically mounts external USB drives that are connected to the computer and displays them in the list of available backup locations. You can select the necessary device and press **Enter** to pass to the Browse for backup files step of the wizard.

If the backup file is located in a network shared folder, on a local drive or on a Veeam backup repository, select one of the following options:

- **Add shared folder** — select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the Mount shared folder step of the wizard.
- **Mount local disk** — select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. With this option selected, you will pass to the Select local disk step of the wizard.
- **Add VBR server** — select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the Specify backup server parameters step of the wizard.
- **Add Cloud provider** — select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the Specify Cloud provider parameters step of the wizard.



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.

Specifying Shared Folder Settings

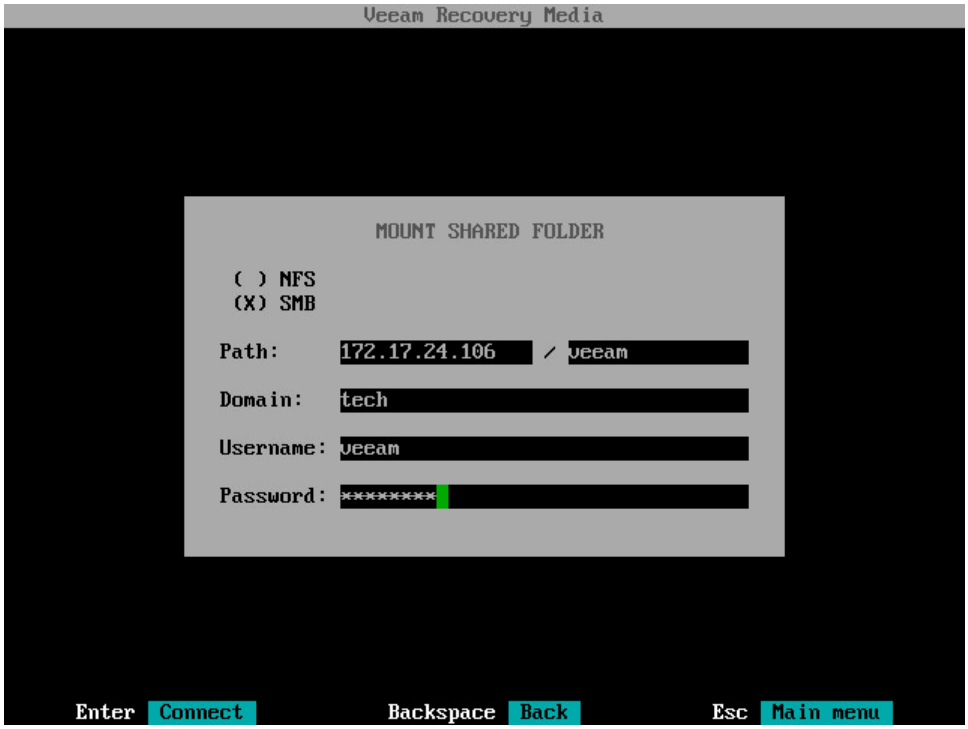
The **Mount shared folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. Select the type of a network shared folder:
   - **NFS** — to connect to a network shared folder using the NFS protocol.
   - **SMB** — to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Path** field, specify the network shared folder name in the *SERVER/DIRECTORY* format: type an IP address or domain name of the server and the name of the network shared folder in which the backup file resides.
3. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
4. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
5. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.
6. Press **Enter** to connect to the network shared folder. Veeam Agent will mount the specified network shared folder to the `/media` directory of the recovery image OS file system and display content of the network shared folder.

TIP

You can mount several network shared folders to work with backup files that are stored in different locations if needed. To do this, return to the Select Backup Location step of the wizard and select the **Add shared folder** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted network shared folders and browse for a backup file located on the necessary storage.

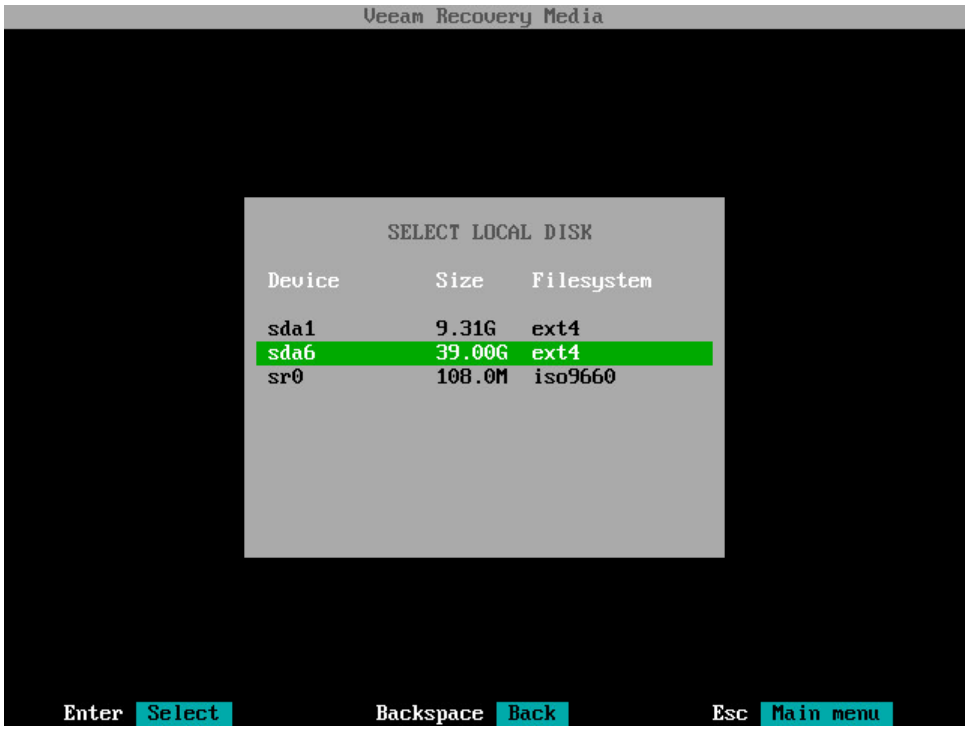


Selecting Local Drive

The **Select local disk** step of the wizard is available if you have selected to restore data from a backup file located on a computer drive.

In the list of devices, select the necessary disk or disk partition and press **Enter**. Veeam Agent will mount the selected device to the `/media` directory of the recovery image OS file system and display content of the directory.

> **TIP**
>
> You can mount several devices to work with backup files that are stored in different locations if needed. To do this, return to the Select Backup Location step of the wizard and select the **Mount local disk** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted devices and browse for a backup file located on the necessary storage.

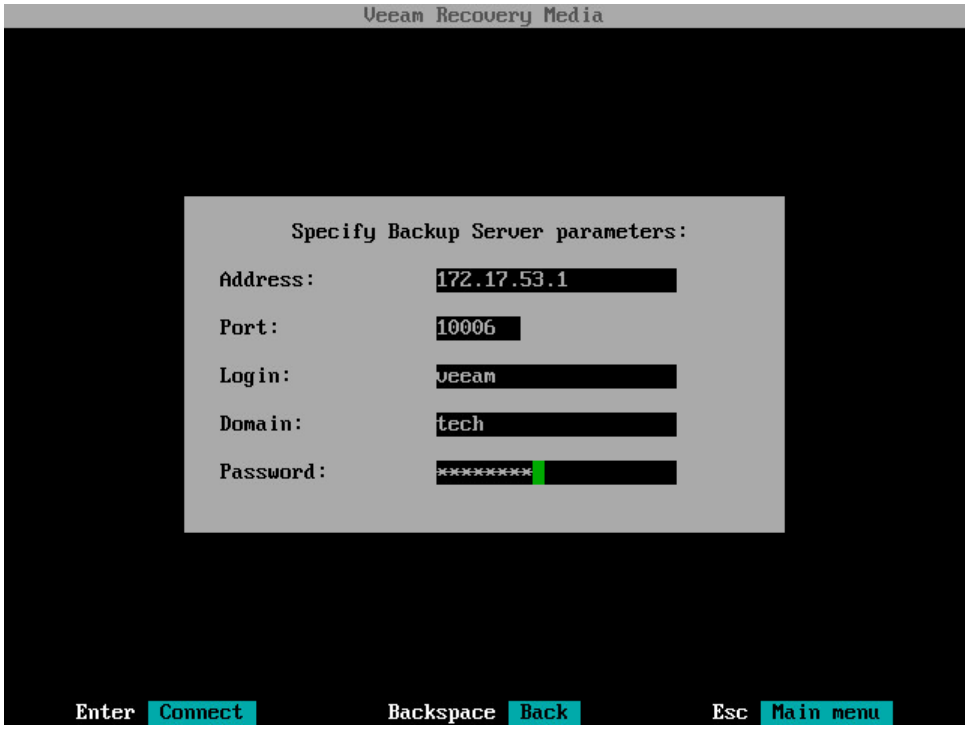


Specifying Backup Server Settings

The **Specify Backup Server parameters** step of the wizard is available if you have selected to restore data from a backup repository managed by the Veeam backup server.

Specify settings for the Veeam backup server that manages the backup repository where the backup file resides:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
4. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.
5. In the **Password** field, type a password of the account that has access to the Veeam backup repository.
6. Press **Enter**. Veeam Agent will connect to the Veeam backup server, and you will pass immediately to the Backup step of the wizard.



Specifying Service Provider Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. Specify service provider settings.
2. Verify the TLS certificate.
3. Specify user account settings.

Specifying Service Provider Settings

The **Specify Cloud provider parameters** step of the wizard is available if you have selected to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.
3. Press **Enter**. Veeam Agent will connect to the service provider and display the Certificate details window.

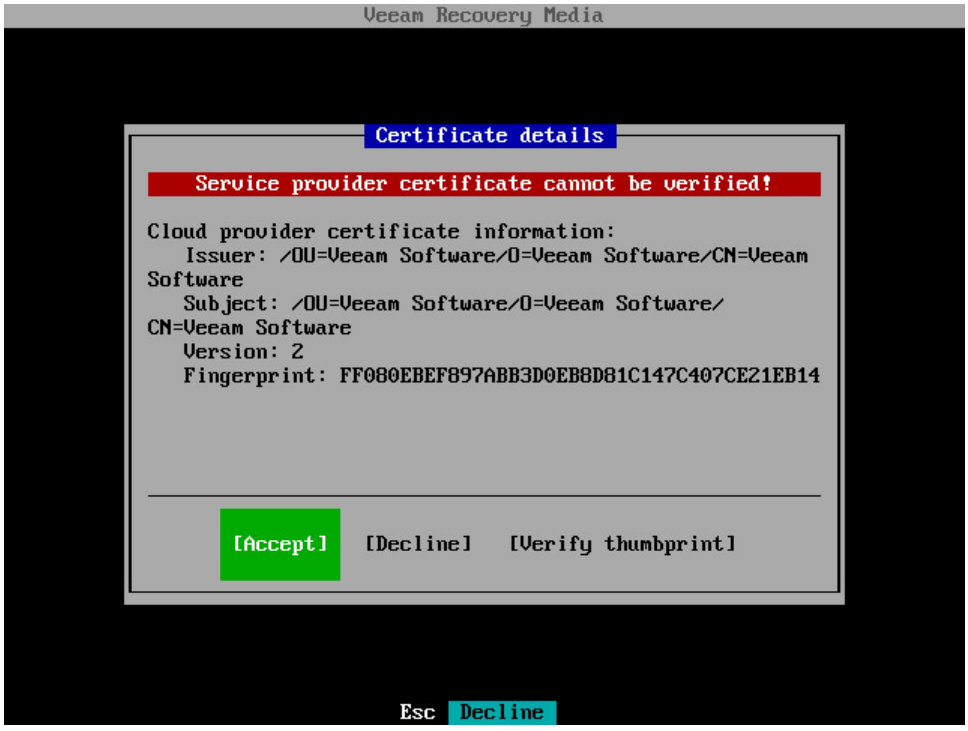


Verifying TLS Certificate

In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate.

- To accept the TLS certificate, select the **Accept** button with the **Tab** key and press **Enter**.
- [Optional] To verify the TLS certificate with a thumbprint, do the following:
  a. Select the **Verify thumbprint** button with the **Tab** key and press **Enter**.
  b. In the **Thumbprint verification** field, enter the thumbprint that you obtained from the SP.
  c. Switch to the **Verify** button and press **Enter**. Veeam Agent will check if the thumbprint that you entered matches the thumbprint of the obtained TLS certificate.

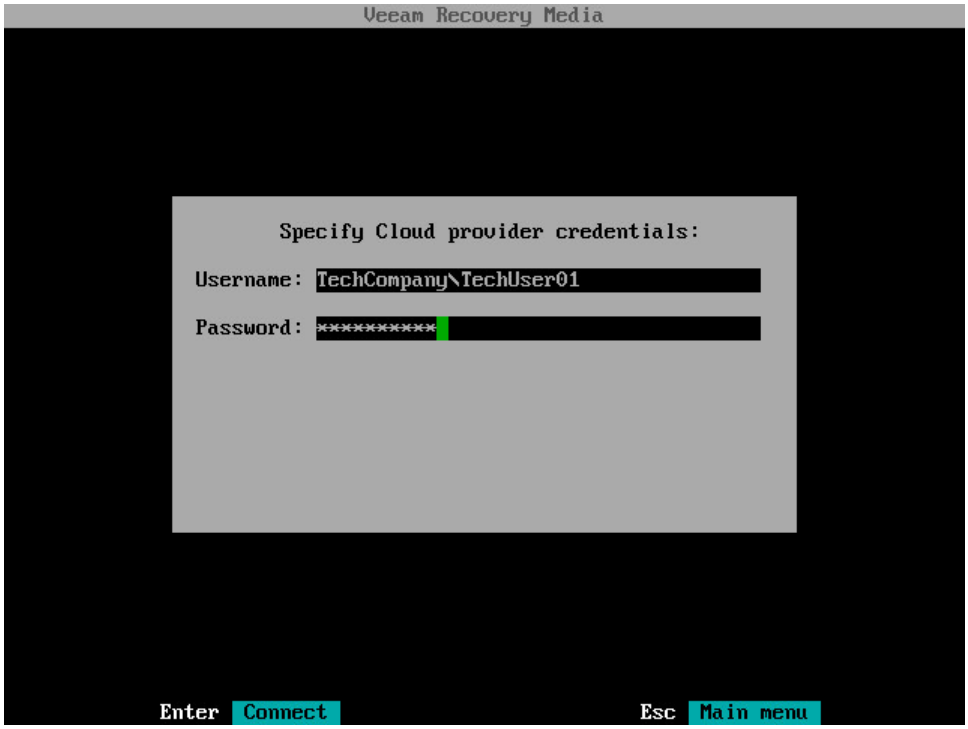
  TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.



Specifying User Account Settings

The **Specify Cloud provider credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

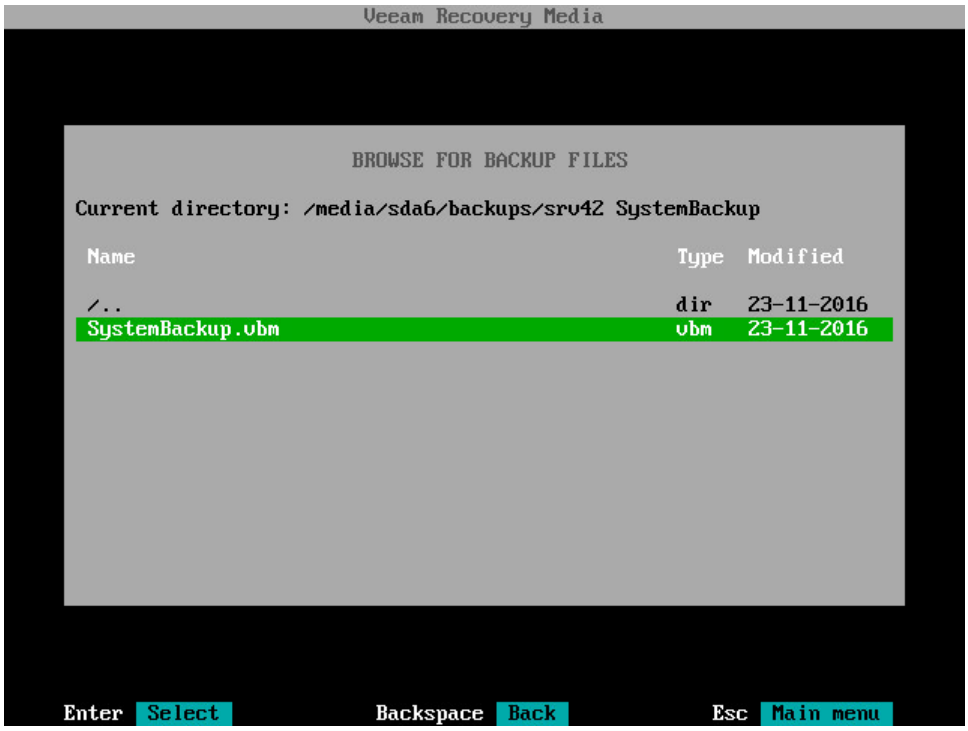
1. In the **Username** field, enter the name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.
2. In the **Password** field, provide a password for the tenant or subtenant account.
3. Press **Enter**. Veeam Agent will connect to the cloud repository, and you will pass immediately to the Backup step of the wizard.



Step 6. Browse for Backup File

At the **Browse for backup files** step of the wizard, select the backup file that you plan to use for volume-level restore:

1. In the file system tree, select a directory in which the backup file you plan to use for restore resides:
   - Use **Up** and **Down** arrow keys to select a directory.
   - Use the **Enter** key to open the necessary directory.
2. In the directory where the backup file resides, select the backup file and press **Enter**.



Step 7. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.

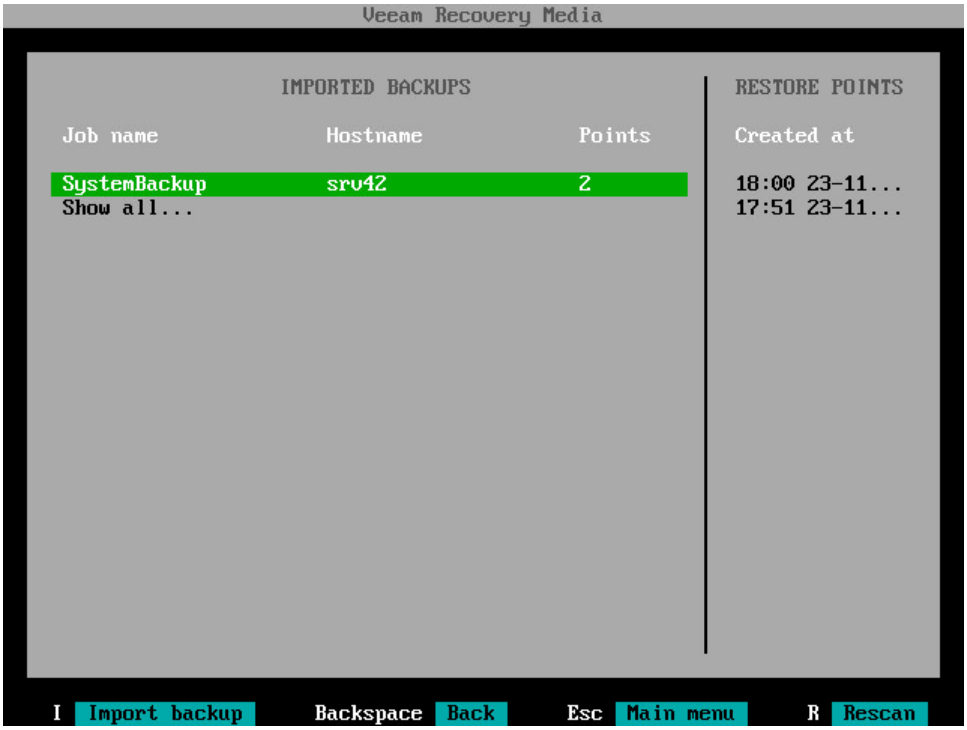
The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays information about backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported backups** pane, ensure that the backup from which you want to recover data is selected and press **Enter**.

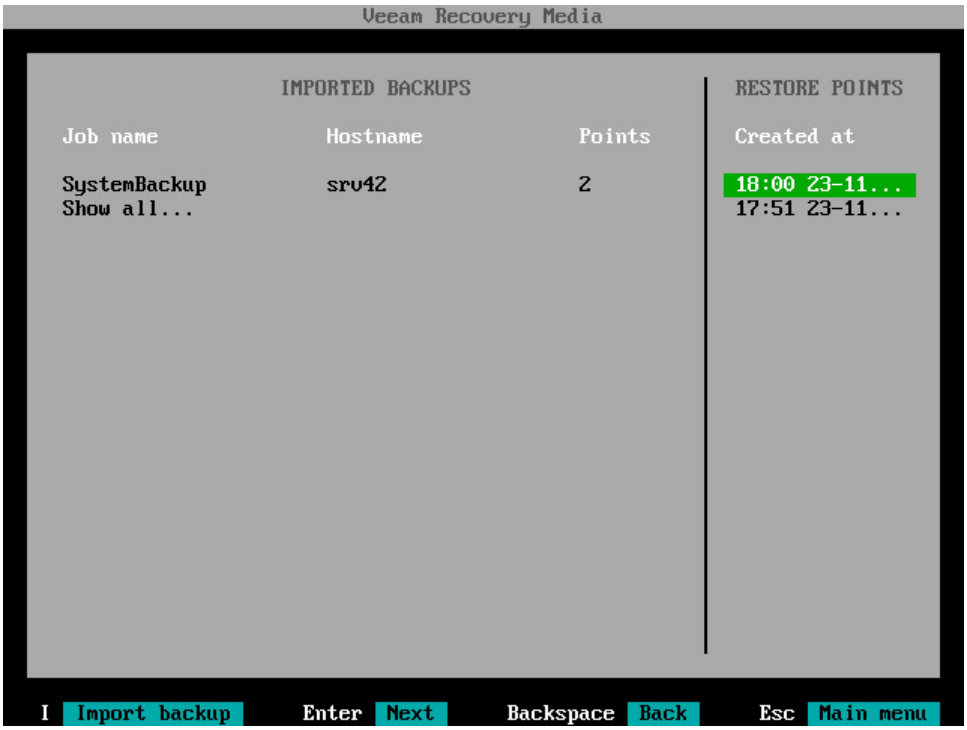
   If you want to select another backup, press the '**i**' key and browse for the necessary backup file. To learn more, see Locate Backup File.



2. In the **Restore points** pane, select with **Up** and **Down** keys the restore point from which you want to recover data and press **Enter**.

NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To lean more, see Restoring Data from Encrypted Backups.



Step 8. Map Restored Disks

At the **Disk Mapping** step of the wizard, select what volumes you want to restore and map volumes from the backup to volumes on your computer.

> **IMPORTANT**
>
> It is strongly recommended that you change disk mapping settings only if you have experience in working with Linux disks and partitions. If you make a mistake, your computer data may get corrupted.

You can map volumes in the backup (source volumes) and volumes on your computer (target volumes) in one of the following ways:

- Map a source volume to a target volume
- Map a target volume to a source volume

As well as individual volumes, you can also map entire disks:

- Map a source disk to a target disk
- Map a target disk to a source disk

If you choose to restore an entire disk, Veeam Agent will try to map all volumes that reside on this disk.

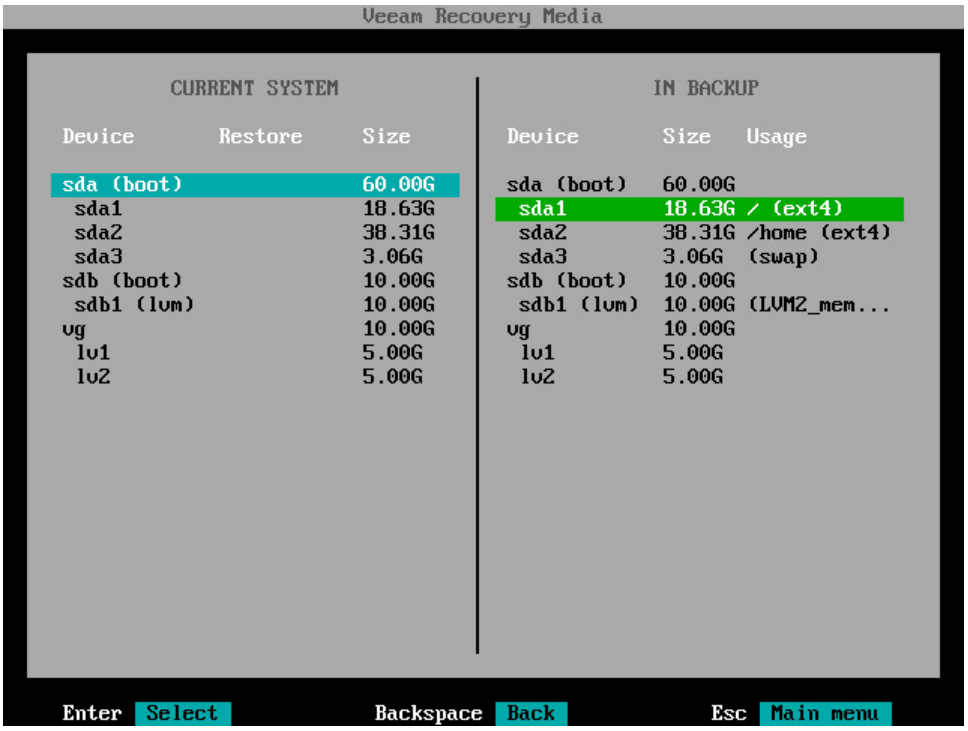
If you want to restore Btrfs subvolumes, you must map subvolumes in the backup to a Btrfs pool on the Veeam Agent computer. To learn more, see Mapping Btrfs Subvolumes.

Mapping Source Volume to Target Volume

The **In backup** pane of the **Veeam Recovery Media** wizard contains a list of disks and volumes in the backup. You can select volumes in the backup that you want to restore to your computer and specify mapping rules for these volumes.

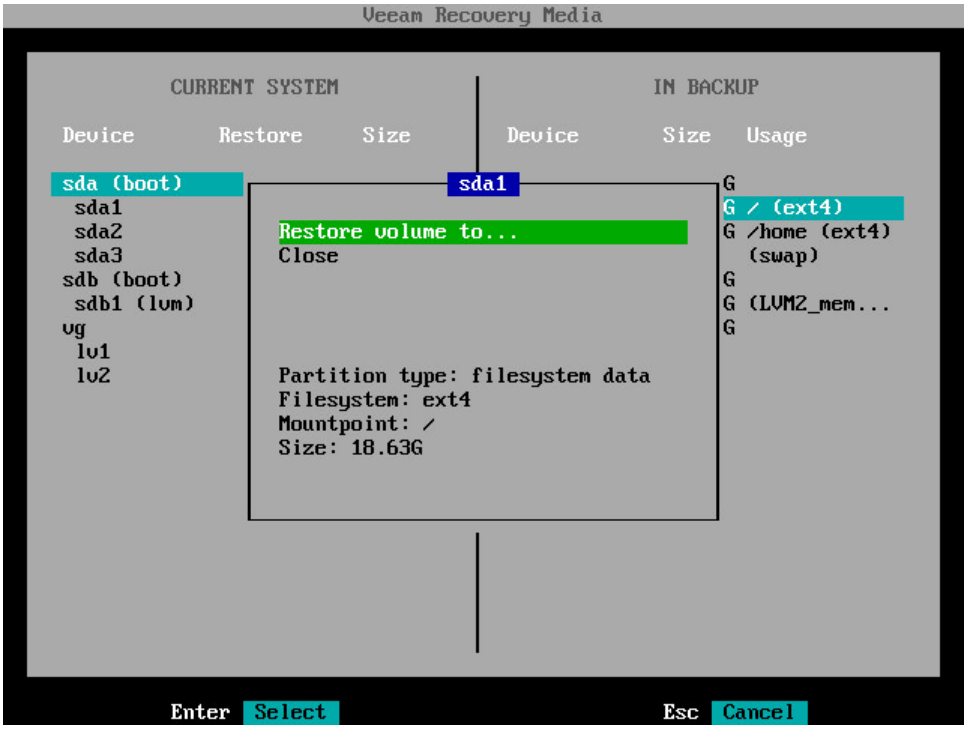
To map a source volume to a target volume:

1. In the **In backup** pane, select a volume in the backup whose data you want to recover and press **Enter**.

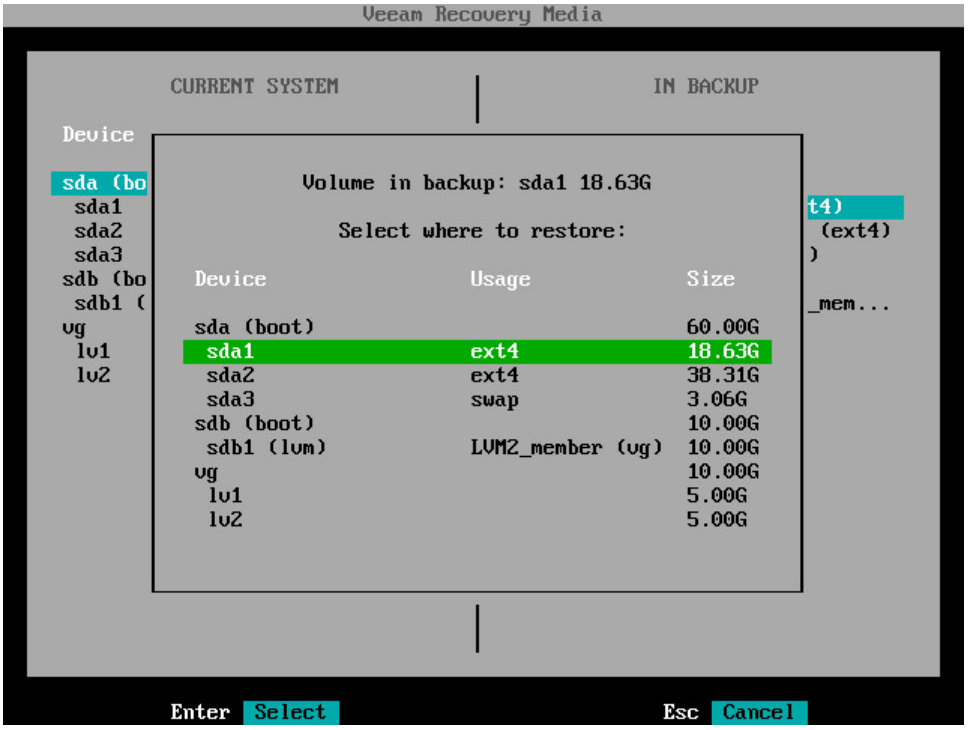


2. Veeam Agent for Linux will display a window with information on the selected volume (partition type, file system type, mount point and volume size) and a list of available operations:
   - **Restore volume to** — select this option if you want to restore the selected volume to your computer.
   - **Close** — select this option if you want to close the window and select another volume.

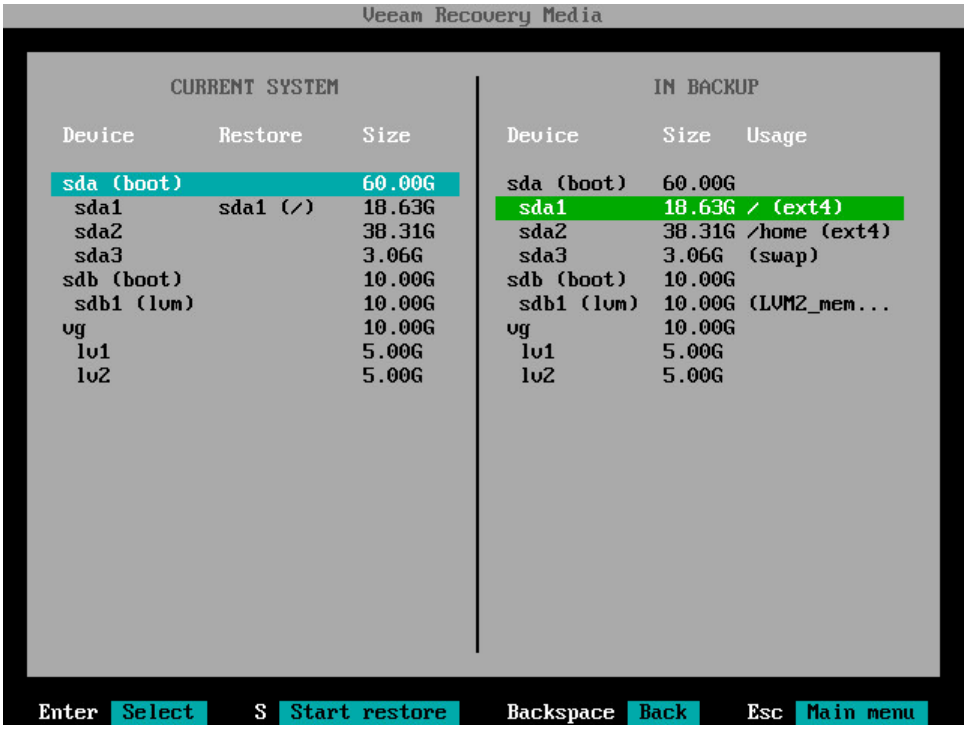
3. Select the **Restore volume to** option and press **Enter**.



4. Veeam Agent for Linux will display a list of volumes on your computer. Select the volume that you want to restore and press **Enter**.



5. In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volume from backup will be restored to the target volume.



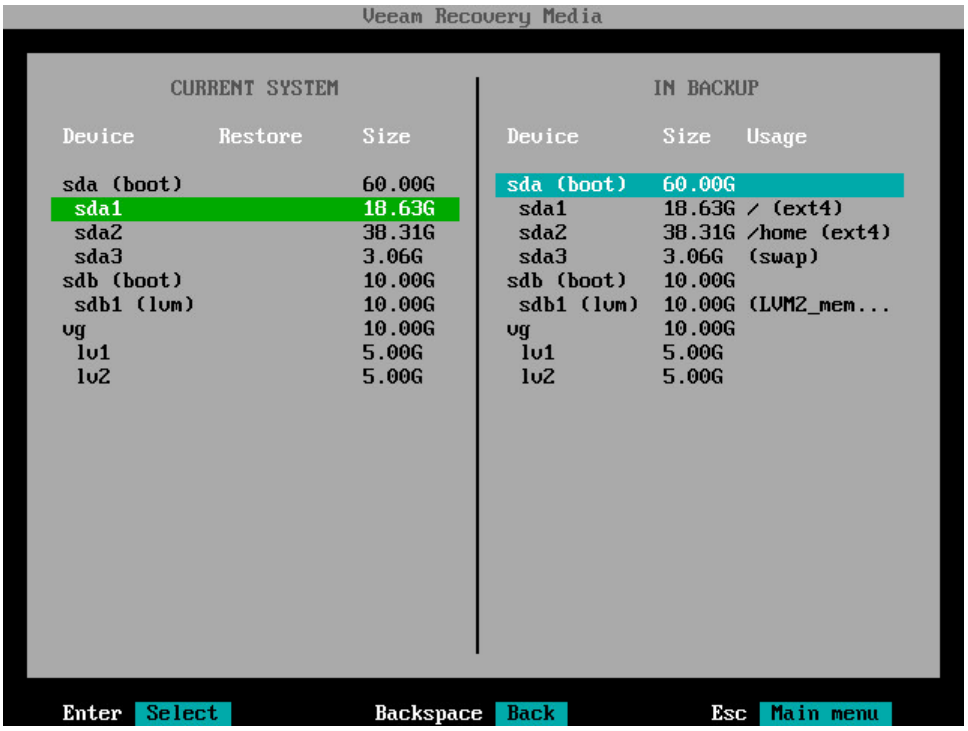
6. Repeat steps 1-5 for all volumes that you want to restore.
7. Press '**s**' to start the restore process.

Mapping Target Volume to Source Volume

The **Current system** pane of the **Veeam Recovery Media** wizard displays a partition table of your computer booted from the Veeam Recovery Media. In this pane, you can select volumes on your computer which you want to restore and specify mapping rules for these volumes. If necessary, you can edit the disk layout before restoring volumes.

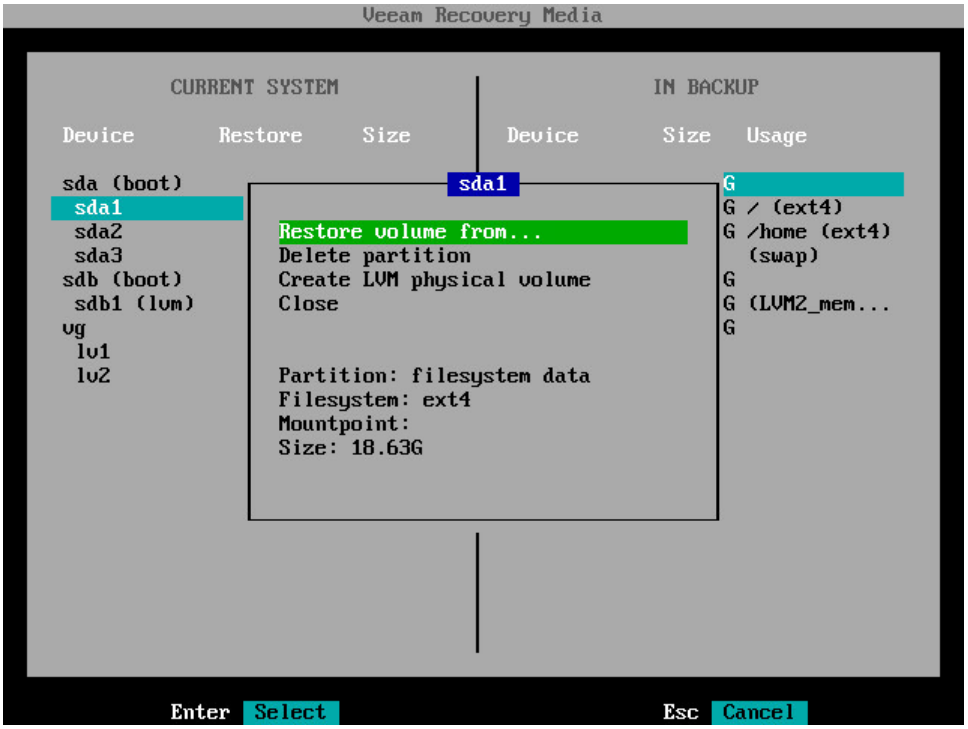
To map a target volume to a source volume:

1. In the **Current system** pane, select a volume on your computer whose data you want to recover and press **Enter**.

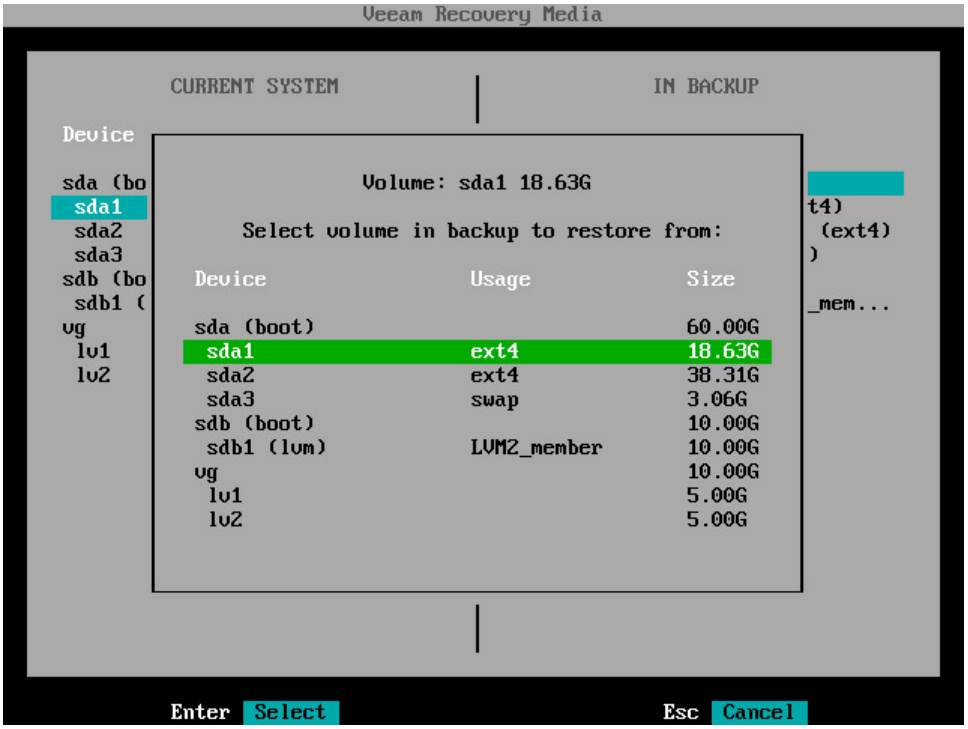


2. Veeam Agent for Linux will display a window with information on the selected volume (partition type, file system type, mount point and volume size) and a list of available operations:
   - **Restore volume from** — select this option if you want to recover the selected volume from the backup.
   - **Delete partition** [for simple volumes] or **Delete volume** [for LVM volumes] — select this option if you want to change the disk layout before restoring a volume. After you delete a partition or volume, you will be able to create a new partition or volume of the desired size and map a volume in the backup to the volume on your computer.
   - [For simple volumes] **Create LVM physical volume** — select this option if you want to create an LVM physical volume on the selected disk partition. In the created physical volume, you will be able to create a volume group and restore to this volume group LVM logical volumes from the backup.
   - **Close** — select this option if you want to close the window and select another volume.

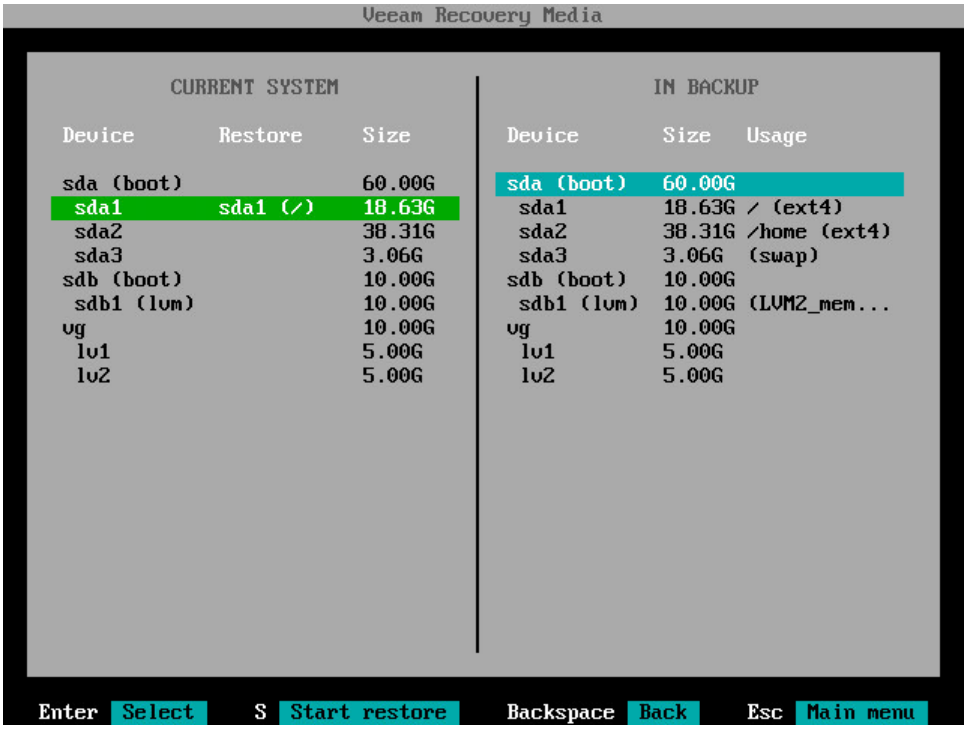
3. Select the **Restore volume from** option and press **Enter**.



4. Veeam Agent for Linux will display a window with a list of volumes in the backup. Select the volume that you want to restore and press **Enter**.



5. In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volume from backup will be restored to the target volume.



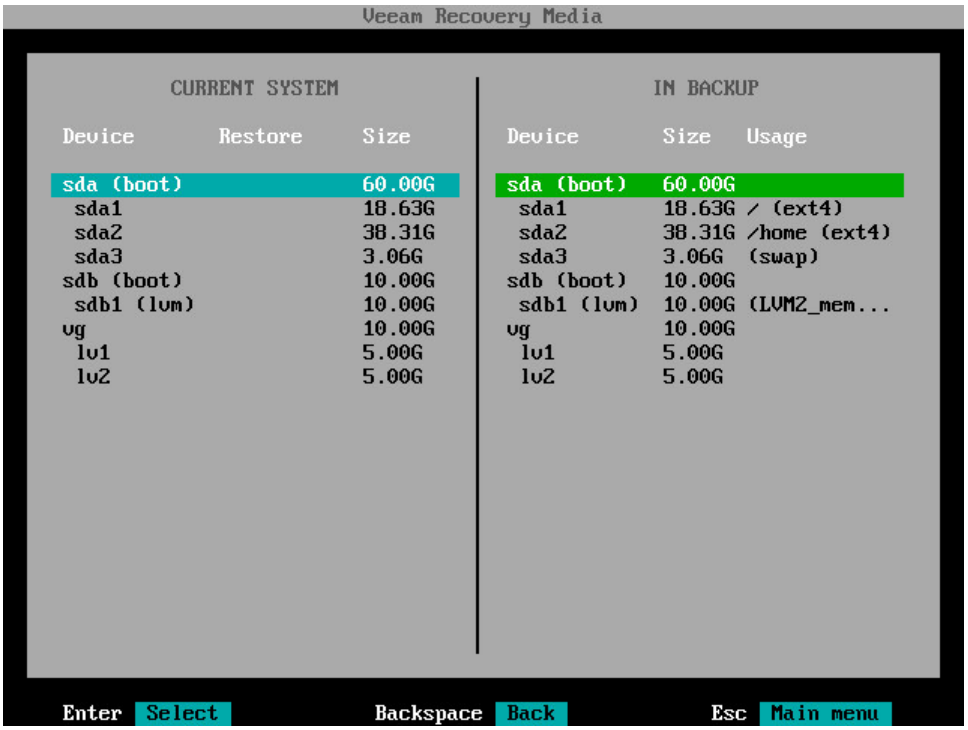
6. Repeat steps 1-5 for all volumes that you want to restore.
7. Press '**s**' to start the restore process.

Mapping Source Disk to Target Disk

The **In backup** pane of the **Veeam Recovery Media** wizard contains a list of disks and volumes in the backup. As well as individual volumes, you can select for restore entire computer disks.

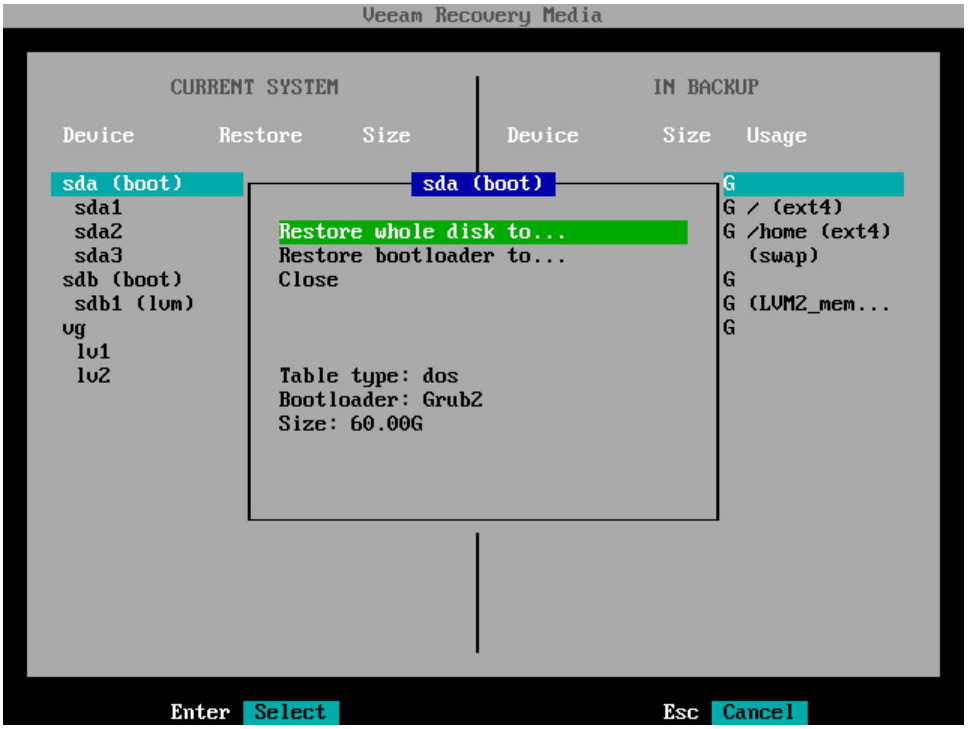
To map a source disk to a target disk:

1. In the **In backup** pane, select a disk in the backup volumes on which you want to recover and press **Enter**.

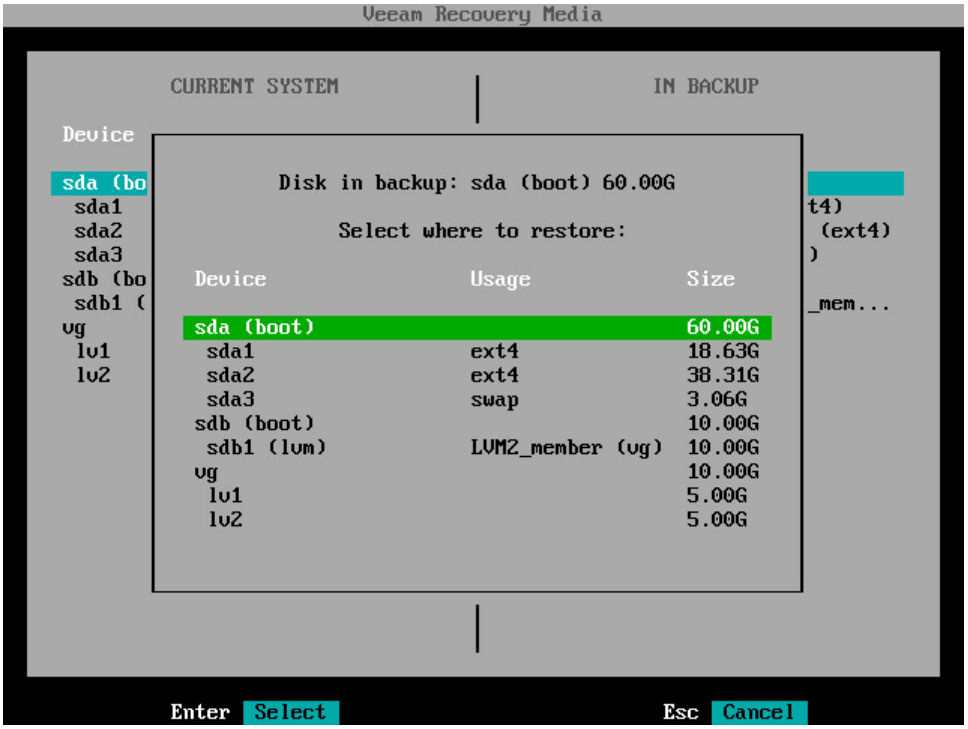


2. Veeam Agent for Linux will display a window with information on the selected disk (partition table type, bootloader type and disk size) and a list of available operations:
   - **Restore whole disk to** — select this option if you want to restore all volumes on the selected disk in the backup to your computer.
   - **Restore bootloader to** — select this option if you want to restore a bootloader from the disk in the backup to your computer.
   - **Close** — select this option if you want to close the window and select another disk or volume.

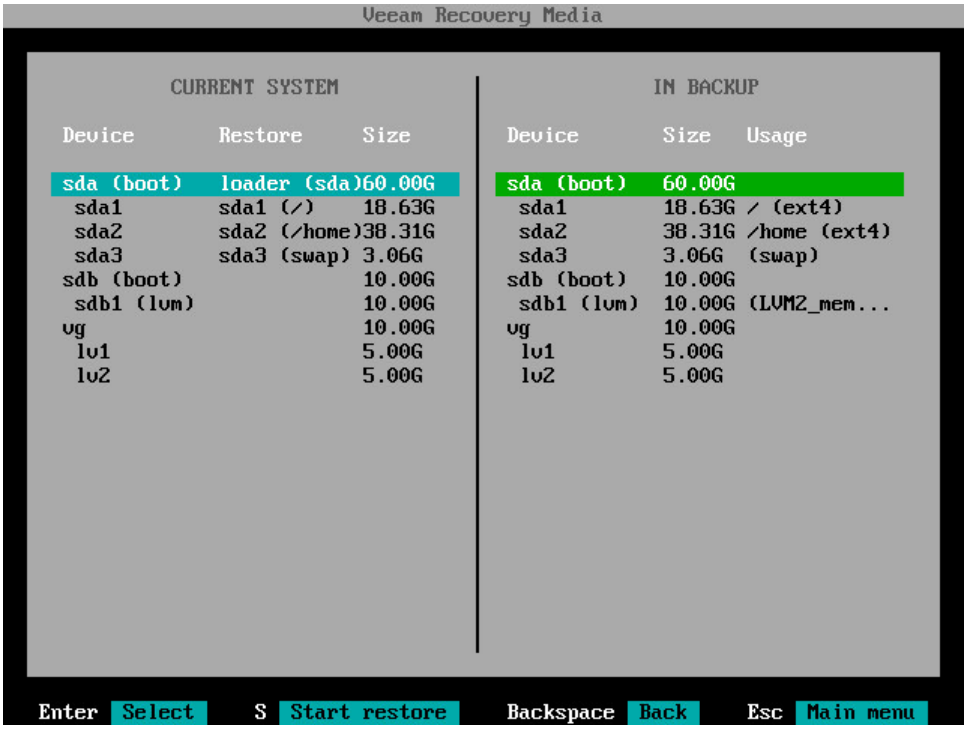
3. To restore volumes that reside on the selected disk, select the **Restore whole disk to** option and press **Enter**.



4. Veeam Agent for Linux will display a list of disks and volumes on your computer. Select the disk whose volumes you want to restore and press **Enter**.



5. In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volumes from the disk in the backup will be restored to the target disk.



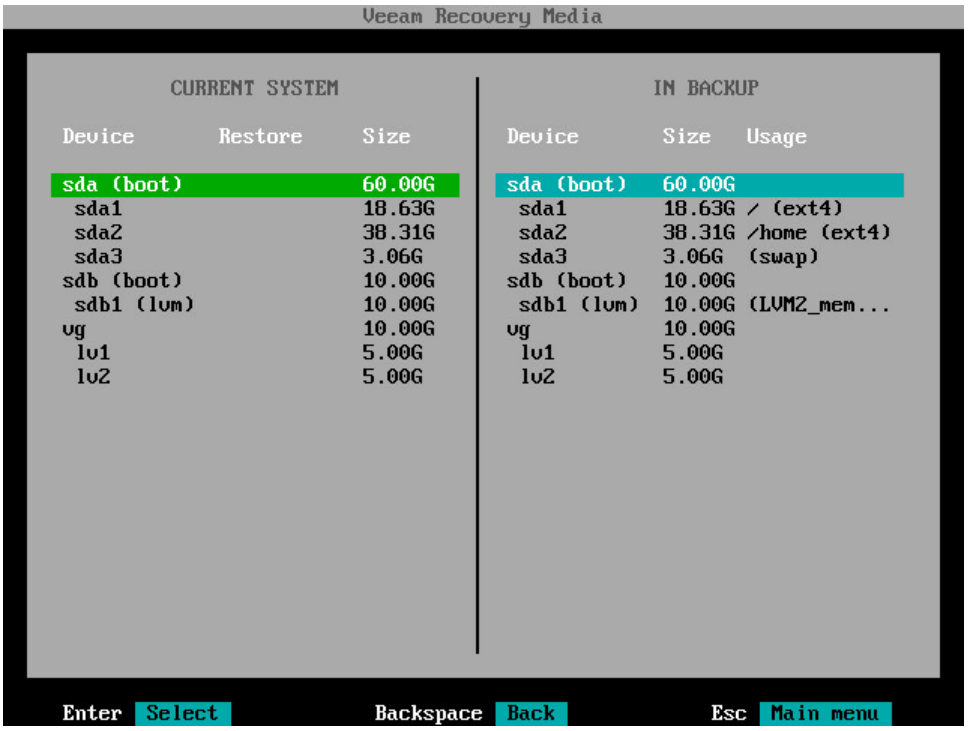
6. Repeat steps 1–5 for all computer disks whose volumes you want to restore.
7. Press '**s**' to start the restore process.

Mapping Target Disk to Source Disk

The **Current system** pane of the **Veeam Recovery Media** wizard displays a partition table of your computer booted from the Veeam Recovery Media. As well as individual volumes, you can select for restore entire computer disks. If necessary, you can edit the disk layout before restoring volumes.

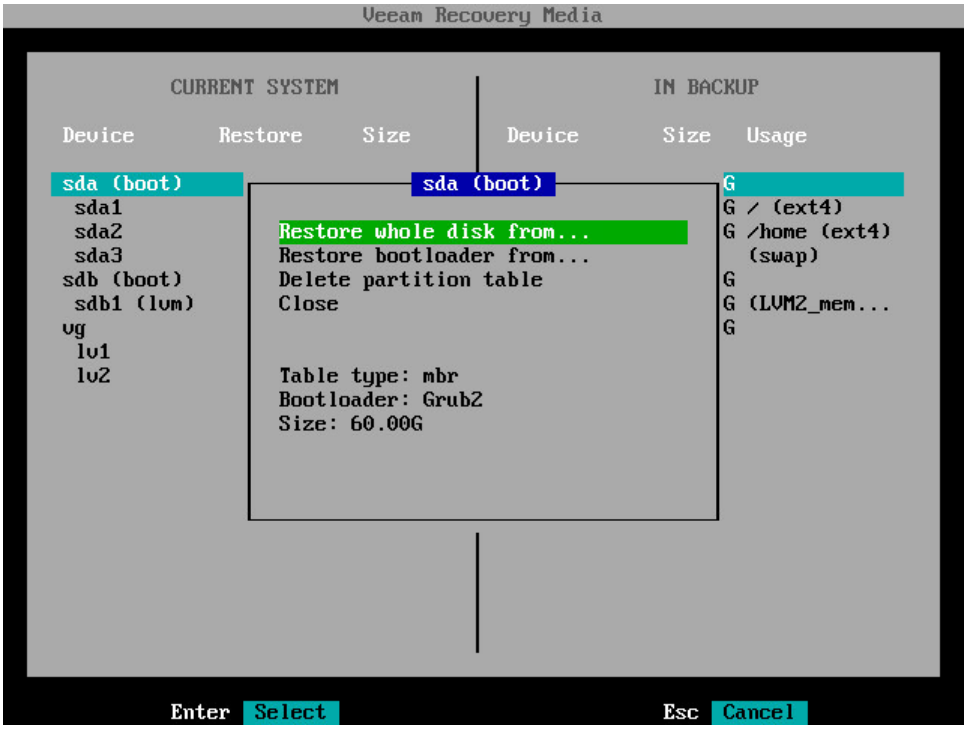
To map a target disk to a source disk:

1. In the **Current system** pane, select a disk on your computer to which you want to restore volumes and press **Enter**.

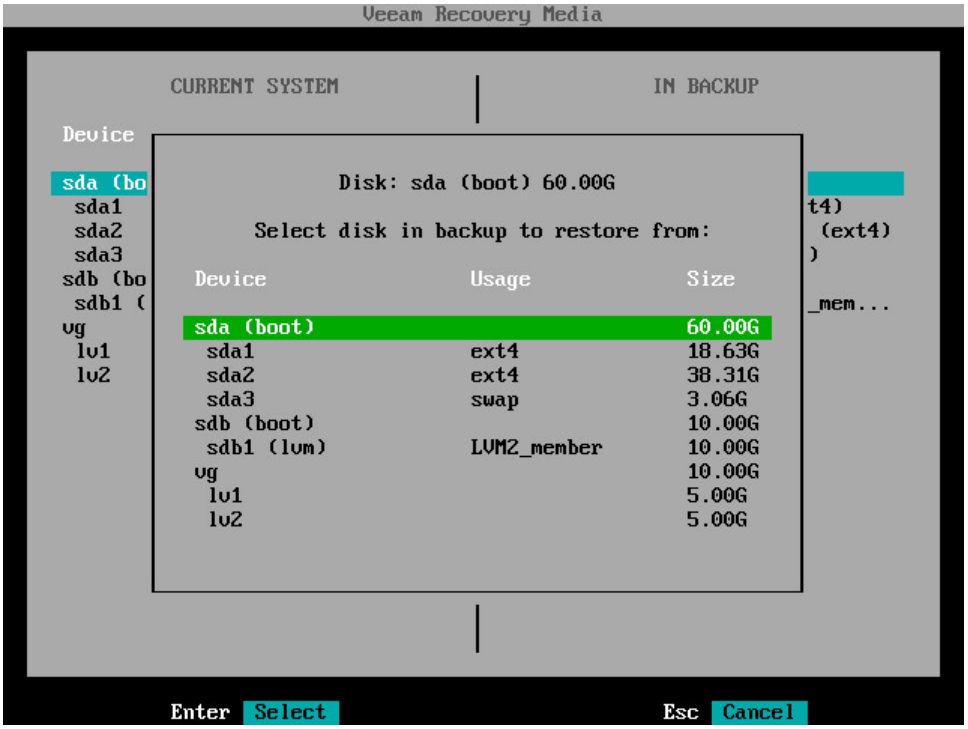


2. Veeam Agent for Linux will display a window with information on the selected disk (partition table type, bootloader type and disk size) and a list of available operations:
   - **Restore whole disk from** — select this option if you want to restore to the selected disk all volumes from a disk in the backup.
   - **Restore bootloader from** — select this option if you want to restore to the selected disk a bootloader from a disk in the backup.
   - **Delete partition table** — select this option if you want to change the disk layout before restoring volumes. After you delete a partition table, you will be able to create a new partition table, create disk partitions and volumes of the desired size, and map volumes in the backup to volumes on your computer.
   - **Close** — select this option if you want to close the window and select another disk or volume.

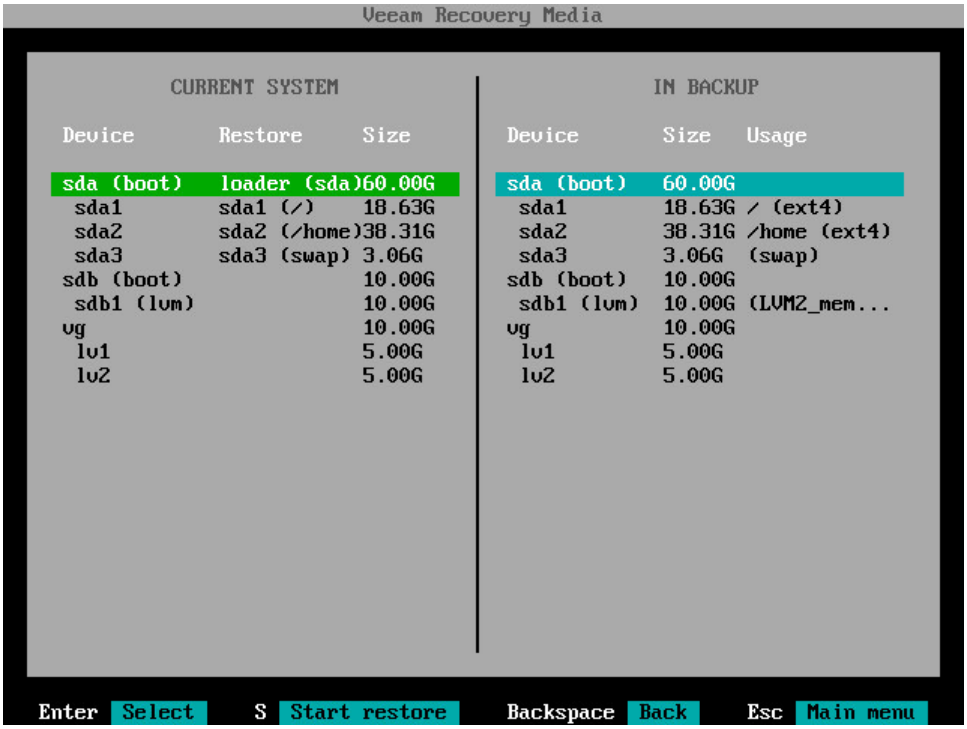
3. To restore volumes to the selected disk, select the **Restore whole disk from** option and press **Enter**.



4. Veeam Agent for Linux will display a list of disks and volumes in the backup. Select the disk whose volumes you want to restore and press **Enter**.



5. In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volumes from the disk in the backup will be restored to the target disk.



6. Repeat steps 1–5 for all disks whose volumes you want to restore.
7. Press '**s**' to start the restore process.

Mapping Btrfs Subvolumes

If the backup contains Btrfs file system data, in the **In backup** pane of the **Veeam Recovery Media** wizard, Veeam Agent displays the list of backed-up Btrfs subvolumes. Information about the original Btrfs pool that contained these subvolumes is not included in the backup.

You can restore from the backup all Btrfs subvolumes or selected subvolumes. To restore a subvolume, you must specify a target Btrfs pool — a Btrfs pool on the computer where you perform restore using the Veeam Recovery Media.

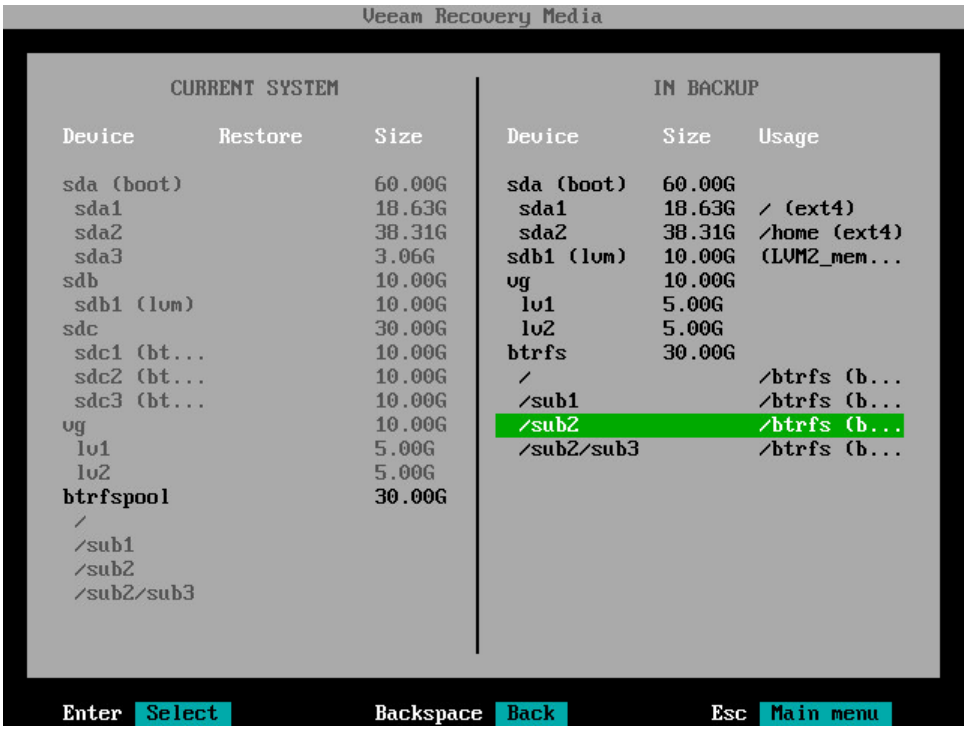
You can restore Btrfs subvolumes to the original Btrfs pool or new Btrfs pool. If the target Btrfs pool contains a subvolume with the same name as the name of the subvolume that you selected for restore, Veeam Agent will automatically map these subvolumes. During the restore process, Veeam Agent will overwrite data on the target subvolume with the data retrieved from the backup.

> **NOTE**
>
> Veeam Agent for Linux does not check whether the target Btrfs pool has enough disk space to restore the selected subvolumes. If the total size of the restored data is larger than the size of the target Btrfs pool, after the restore process completes, the restored data will be corrupted.

To map a source Btrfs subvolume to a target Btrfs pool:

1. In the **In backup** pane, select a subvolume in the backup whose data you want to restore and press **Enter**.

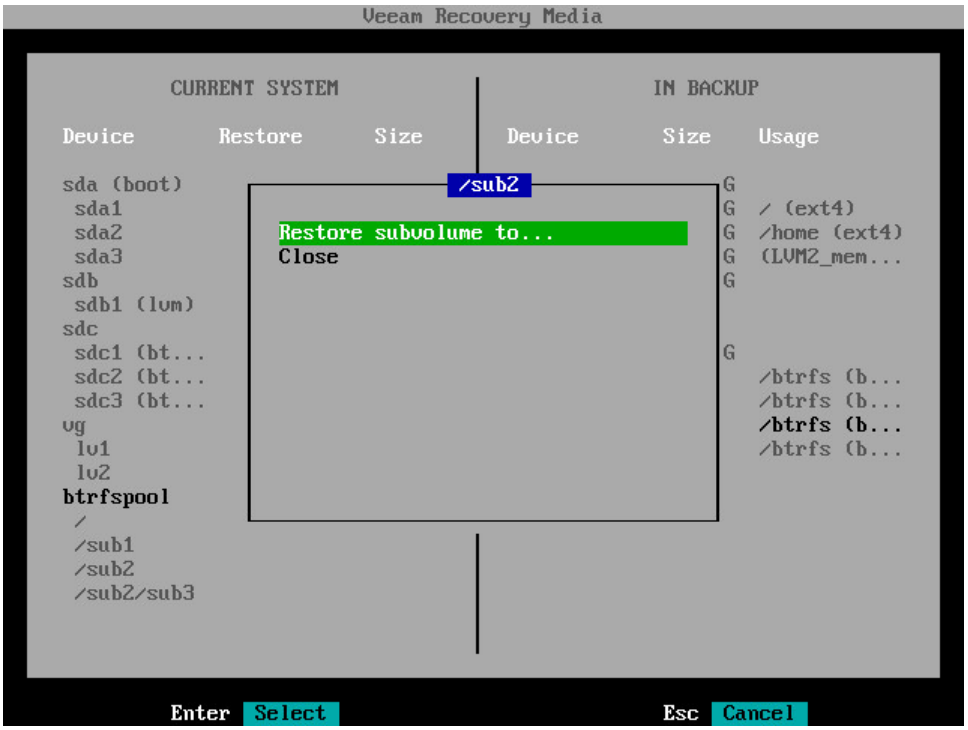
   You can also choose to restore all subvolumes from the backup at once. To do this, in the **In backup** pane, select **btrfs** and press **Enter**.



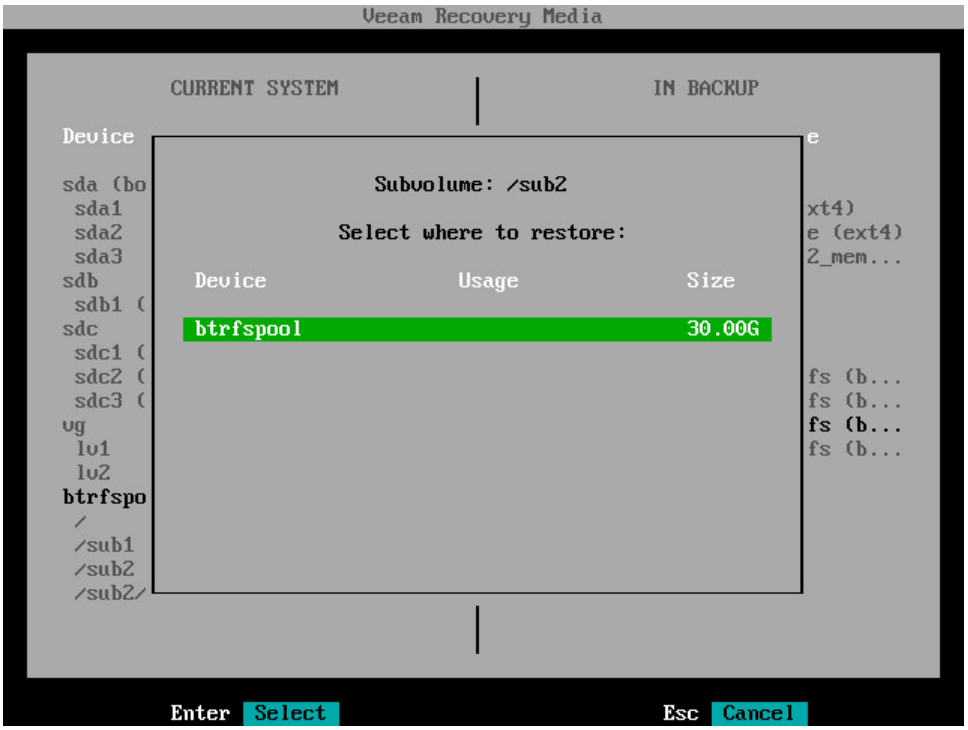
2. In the displayed window, select the necessary option for Btrfs restore and press **Enter**. The available options depend on what Btrfs subvolumes you selected for restore: all subvolumes or specific subvolume.

   - **Restore subvolume to** — this option is available if you chose to restore a specific Btrfs subvolume from the backup. Select this option to restore the selected subvolume to your computer.
   - **Restore btrfs to** — this option is available if you chose to restore all Btrfs subvolumes from the backup. Select this option to restore subvolumes to your computer.

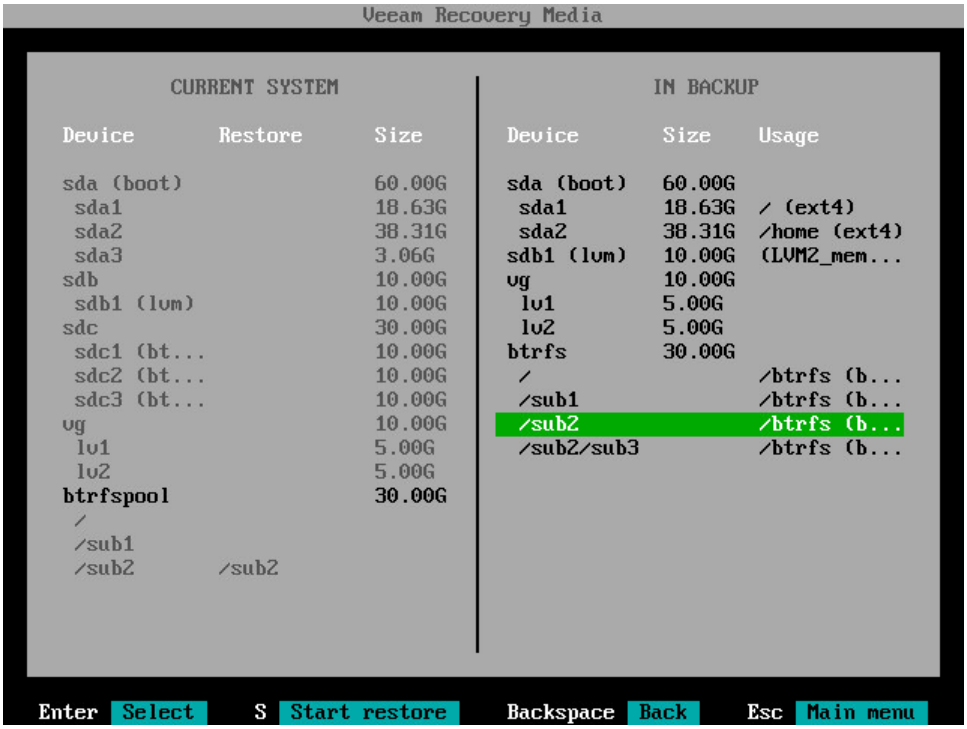
- **Close** — select this option if you want to close the window and select another subvolume.



3. Veeam Agent for Linux will display a list of Btrfs pools on your computer. Select the Btrfs pool where you want to restore data from the backup and press **Enter**.



4. In the **Current system** pane, in the **Restore** column, Veeam Agent for Linux will display which subvolume from backup will be restored to the target Btrfs pool.



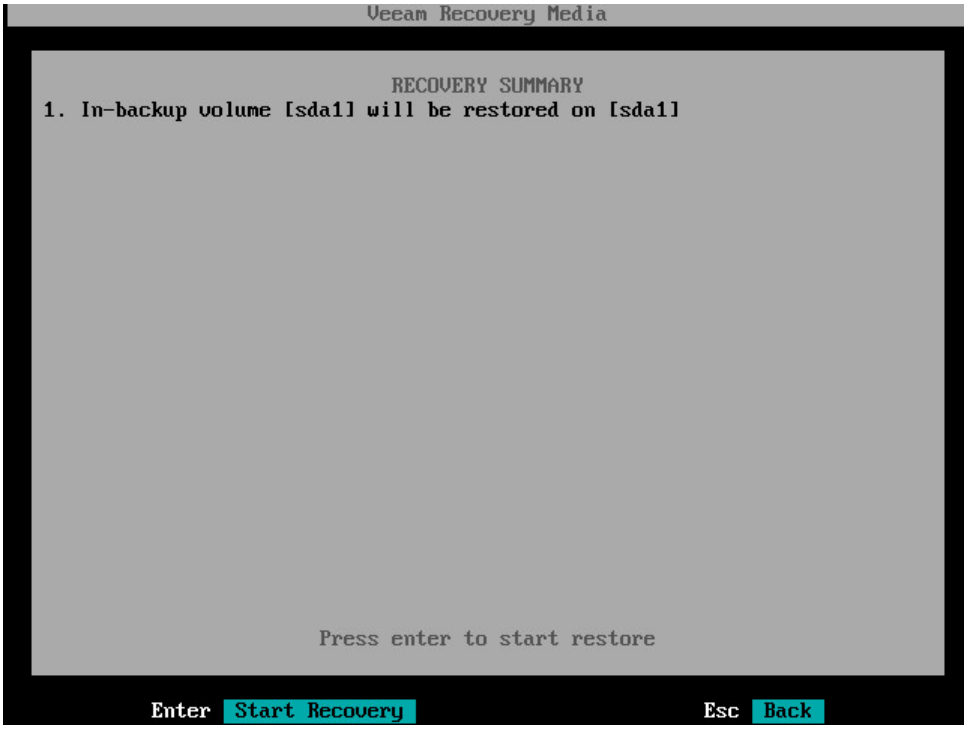
5. If you want to restore more than one subvolume, repeat steps 1–4 for all subvolumes that you want to restore.

6. Press '**s**' to start the restore process.

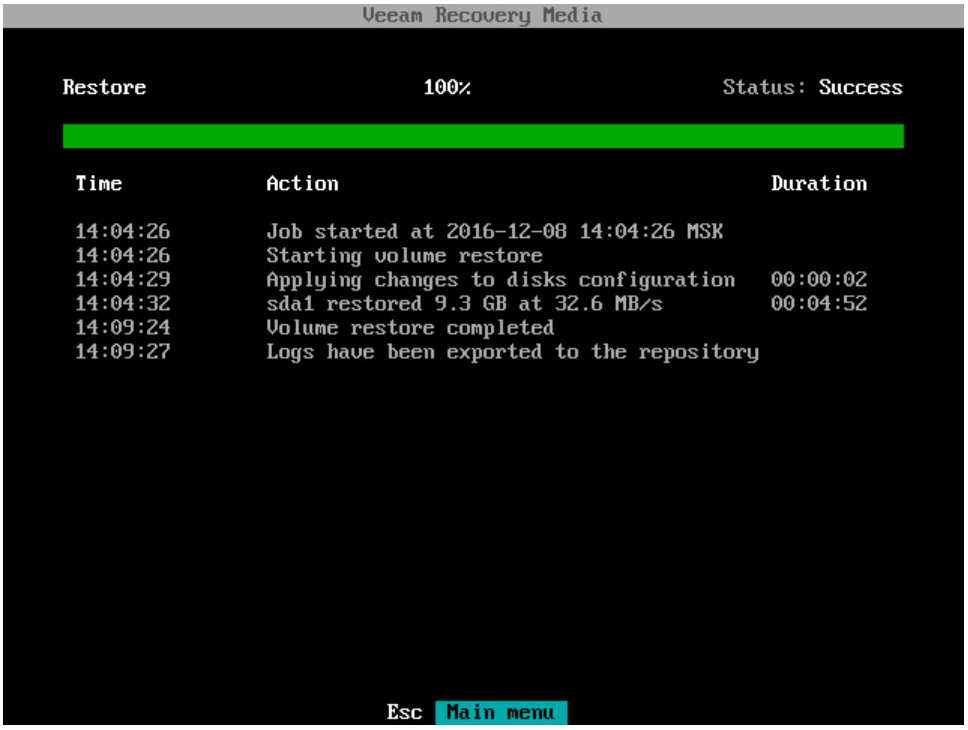
Step 9. Complete Restore Process

At the **Recovery summary** step of the wizard, complete the procedure of volume-level restore.

1. Review the specified recovery settings.



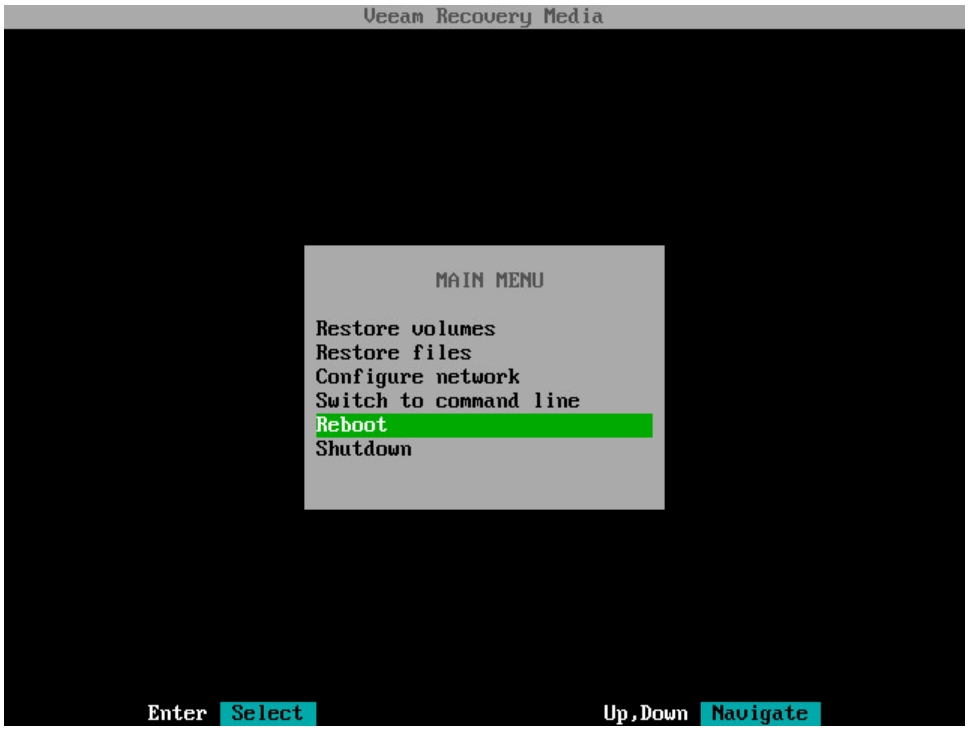
2. Press **Enter** to start the volume-level restore process. Veeam Agent for Linux will perform partition re-allocation operations if necessary, restore the necessary data from the backup and overwrite data on your computer with it.



Step 10. Finish Working with Veeam Recovery Media

When the restore operation completes, finish working with the Veeam Recovery Media and start your operating system.

1. Press **Esc** to return to the Veeam Recovery Media main menu.
2. Eject the media or removable storage device with the recovery image.
3. In the Veeam Recovery Media main menu, select the **Reboot** option and press **Enter**.
4. Wait for your Linux operating system to start.



Restoring Files and Folders

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore with the Veeam Recovery Media, Veeam Agent publishes the backup content directly into the recovery image OS file system and displays it in the file browser. You can restore files and folders to their initial location or copy files and folders to a new location.

Before You Begin

Before you boot from the recovery image and recover your data, check the following prerequisites:

- You must have a recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For data recovery, you can use a volume-level or file-level backup created with Veeam Agent for Linux. Make sure that the backup or system image is available on the computer drive (local or external), on a network shared folder or on the backup repository managed by a Veeam backup server.
- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a Veeam backup repository, you must have access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.
- Veeam Recovery Media 3.0 and later versions do not support Oracle VM virtual machines. Use an earlier version of Veeam Recovery Media or contact Veeam Customer Support for workarounds.

Step 1. Boot from Veeam Recovery Media

To boot from the Veeam Recovery Media:

1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.

   [For removable storage device] Attach the removable storage device with the recovery image to your computer.

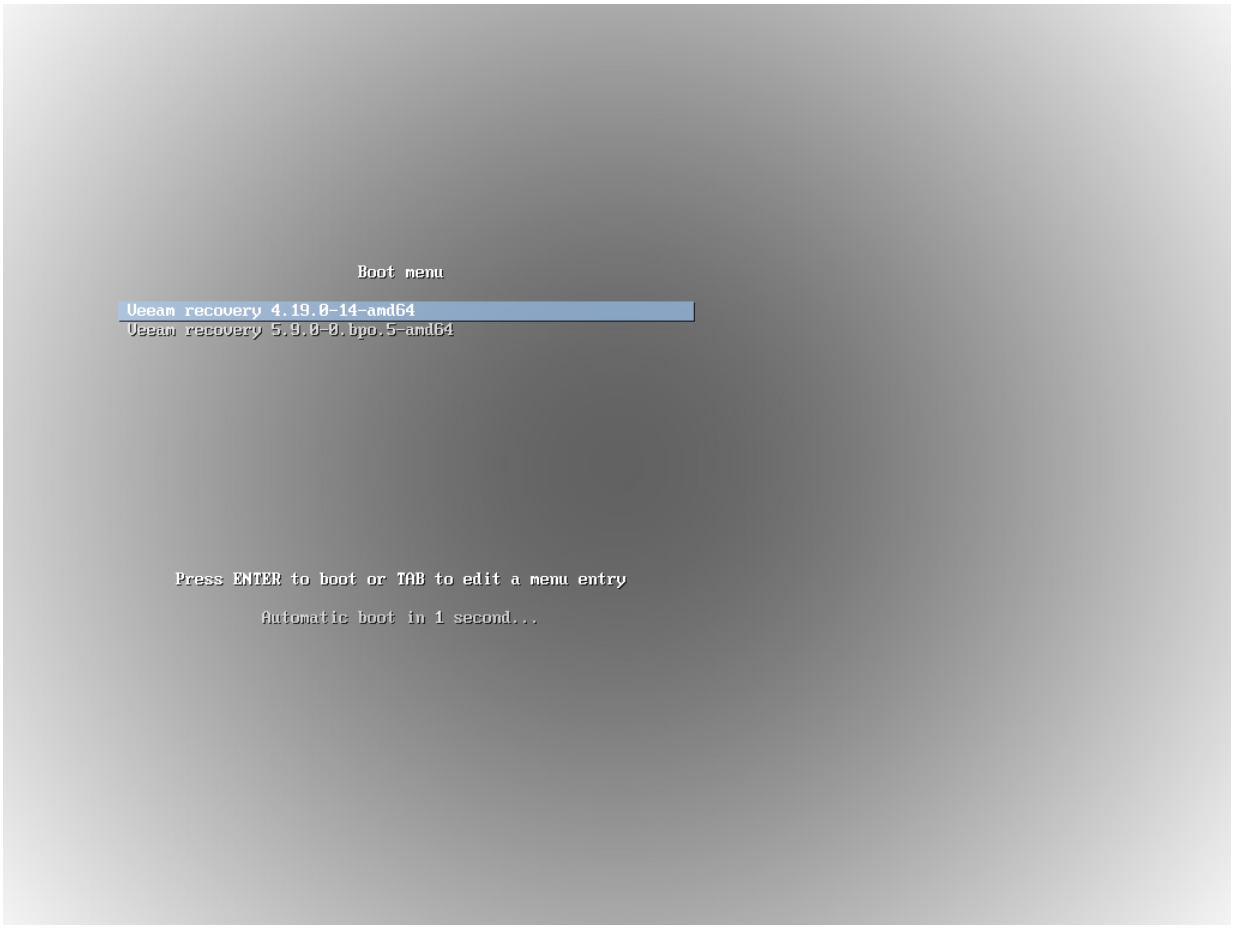
2. Start your computer.

3. [For regular recovery image] In the boot menu, select what Linux kernel version to use to boot your computer and specify boot options if necessary.

   You can select a Linux kernel version if you use generic Veeam Recovery Media downloaded from the Veeam website or Veeam software repository. If you created a custom Veeam Recovery Media, you will be prompted to boot using the Linux kernel of your Veeam Agent computer included in the recovery image.

   To specify boot options, press the **Tab** key and type the necessary options in the command prompt.

   > **NOTE**
   >
   > For the legacy recovery image, the boot menu is unavailable. After you start your computer, Veeam Agent will immediately start loading files from the Veeam Recovery Media.

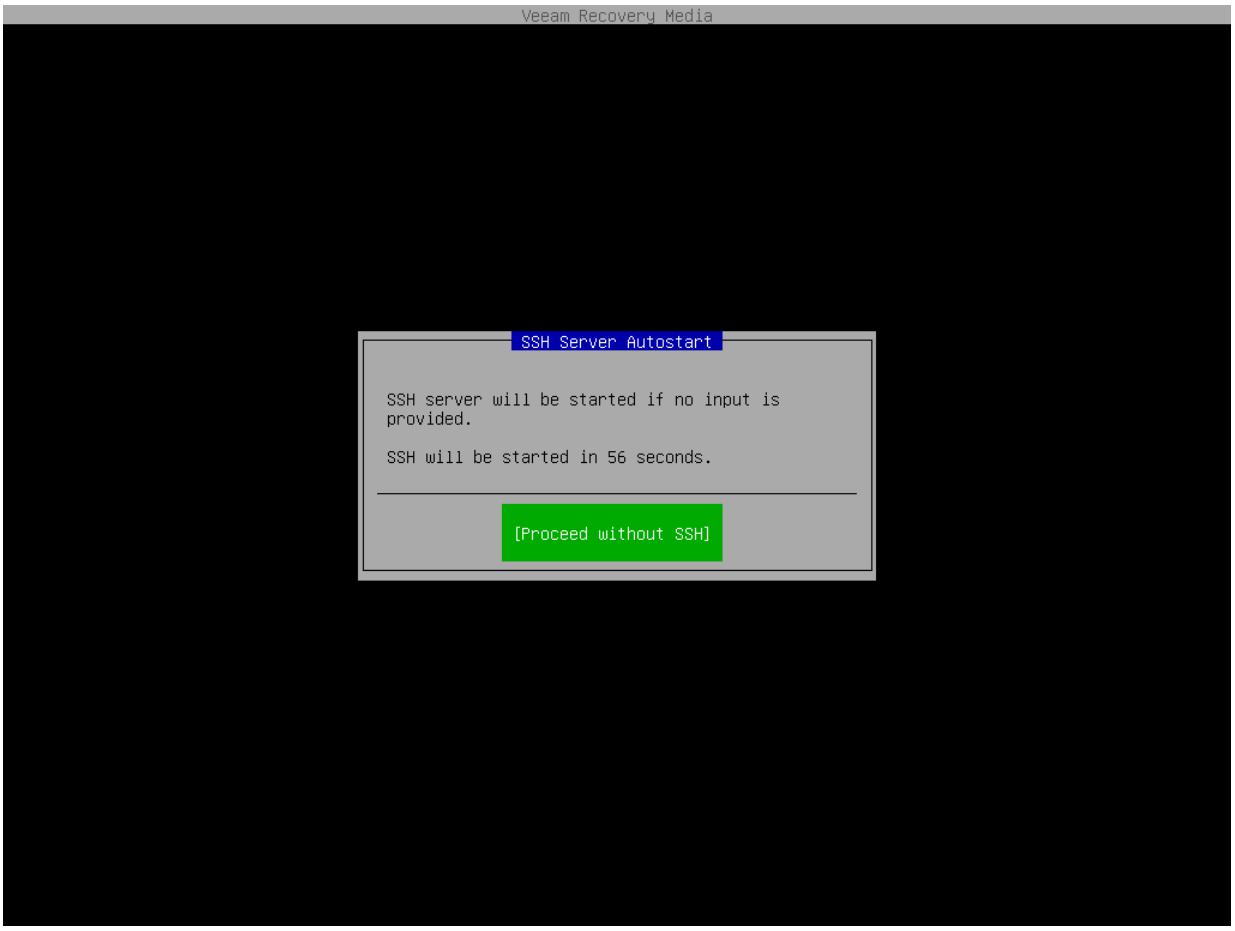


4. Wait for Veeam Agent to load files from the Veeam Recovery Media.

5. After the recovery image OS has loaded, choose whether you want to start the SSH server. The SSH server allows you to connect to the Veeam Recovery Media from a remote machine.

   The Veeam Recovery Media starts the SSH server automatically after a time-out. The default value for the time-out is 60 seconds.

   If you do not want to start the SSH server, make sure that the **Proceed without SSH** button is selected and press **Enter**. You will proceed immediately to the step 7.



6. After the SSH server has started, review settings to connect to the Veeam Recovery Media and press **Enter**.

   The Veeam Recovery Media displays the following connection settings:

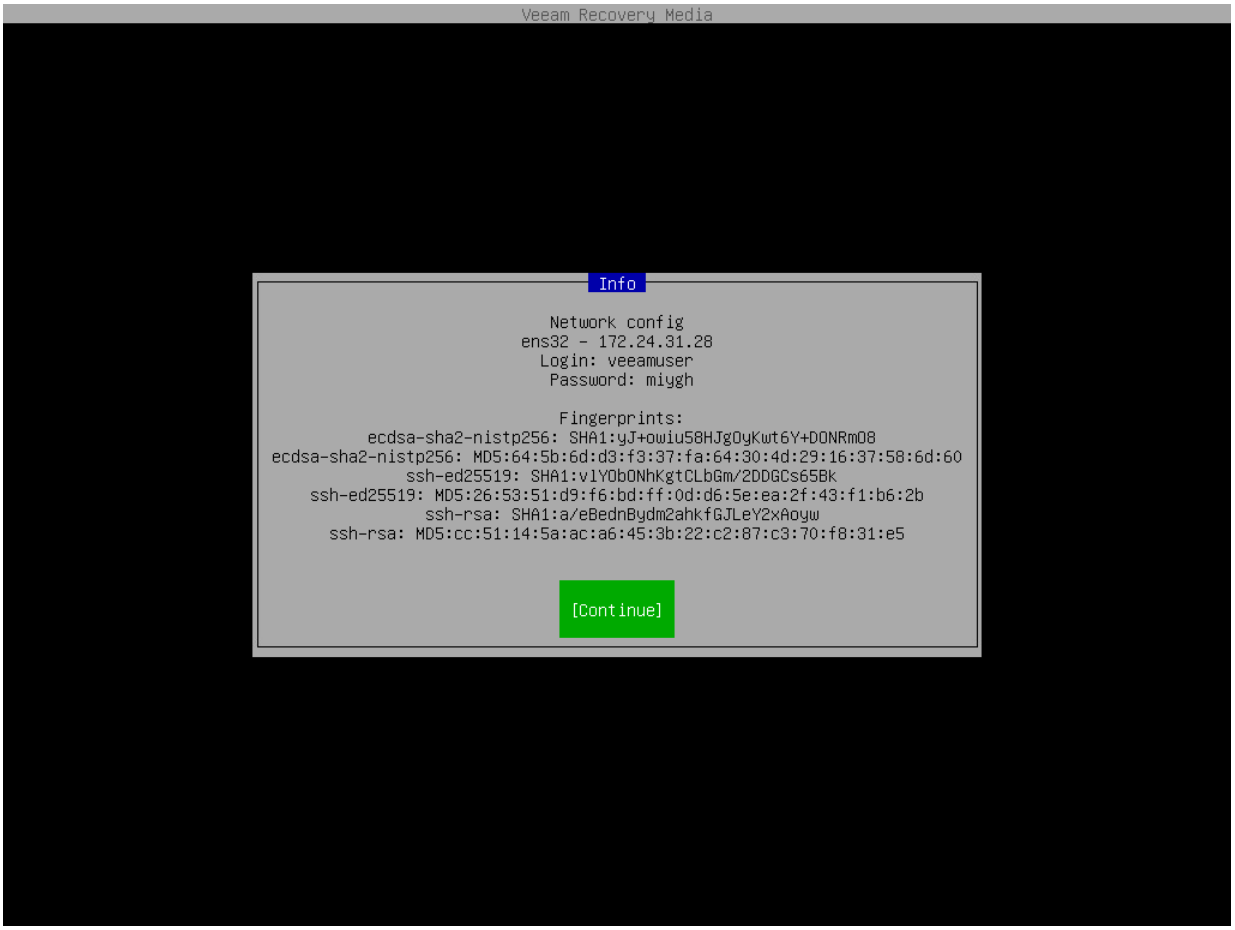
   - IP address of the computer booted from the Veeam Recovery Media
   - User name and password of the account used to connect to the Veeam Recovery Media
   - Fingerprints of the computer booted from the Veeam Recovery Media

NOTE

The user name of the account used to work with the Veeam Recovery Media depends on the recovery image version:

- *veeamuser* — for the regular recovery image
- *root* — for the legacy recovery image

If you want to use command-line utilities built in the regular recovery image, use the `sudo` command to provide the *veeamuser* account with privileges of the *root* account.

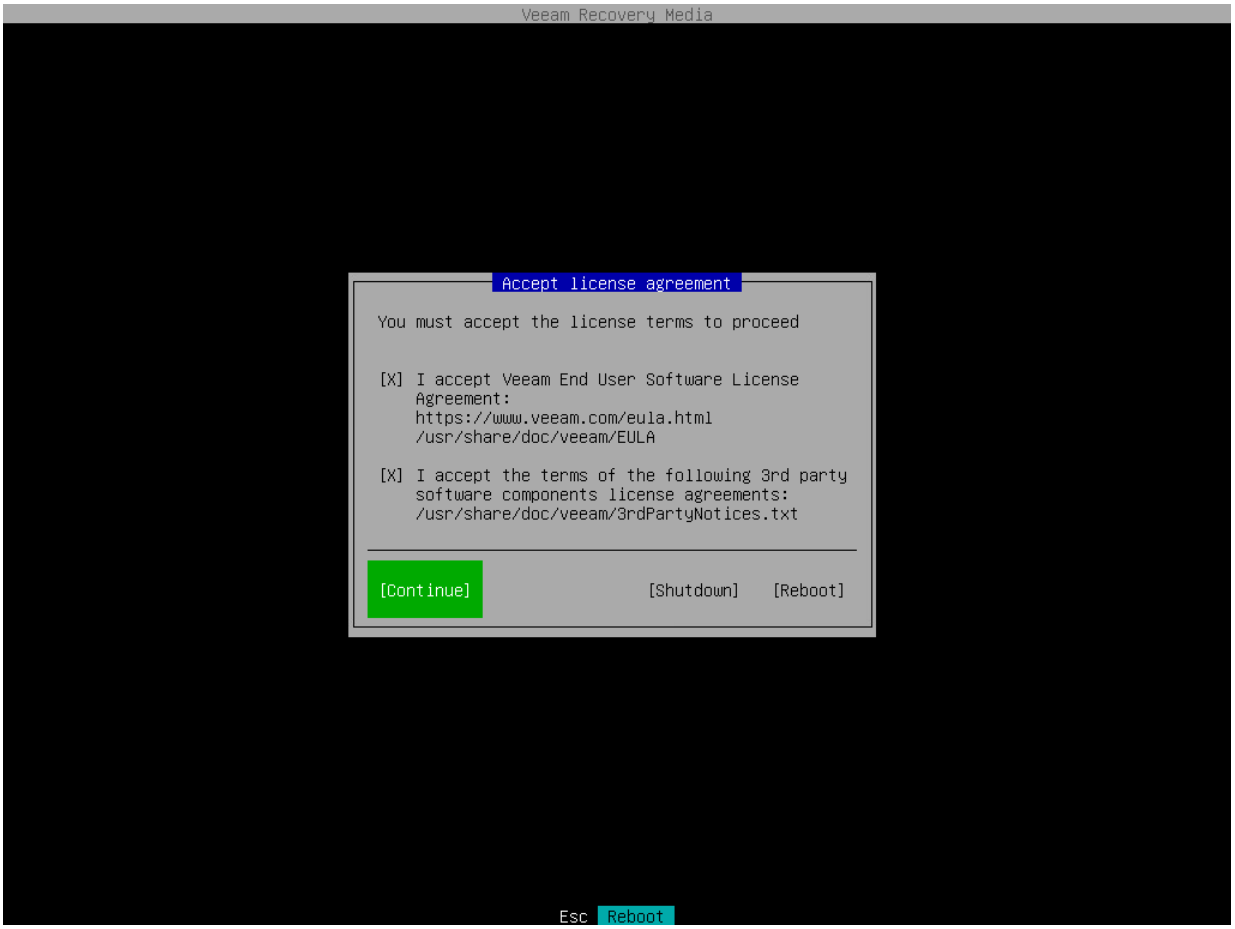


7. Accept the terms of the product license agreement and license agreements for third-party components of the product:

   a. Make sure that the **I accept Veeam End User Software License Agreement** option is selected and press **Space**.

   b. Select the **I accept the terms of the following 3rd party software components license agreements** option with the **Tab** key and press **Space**.

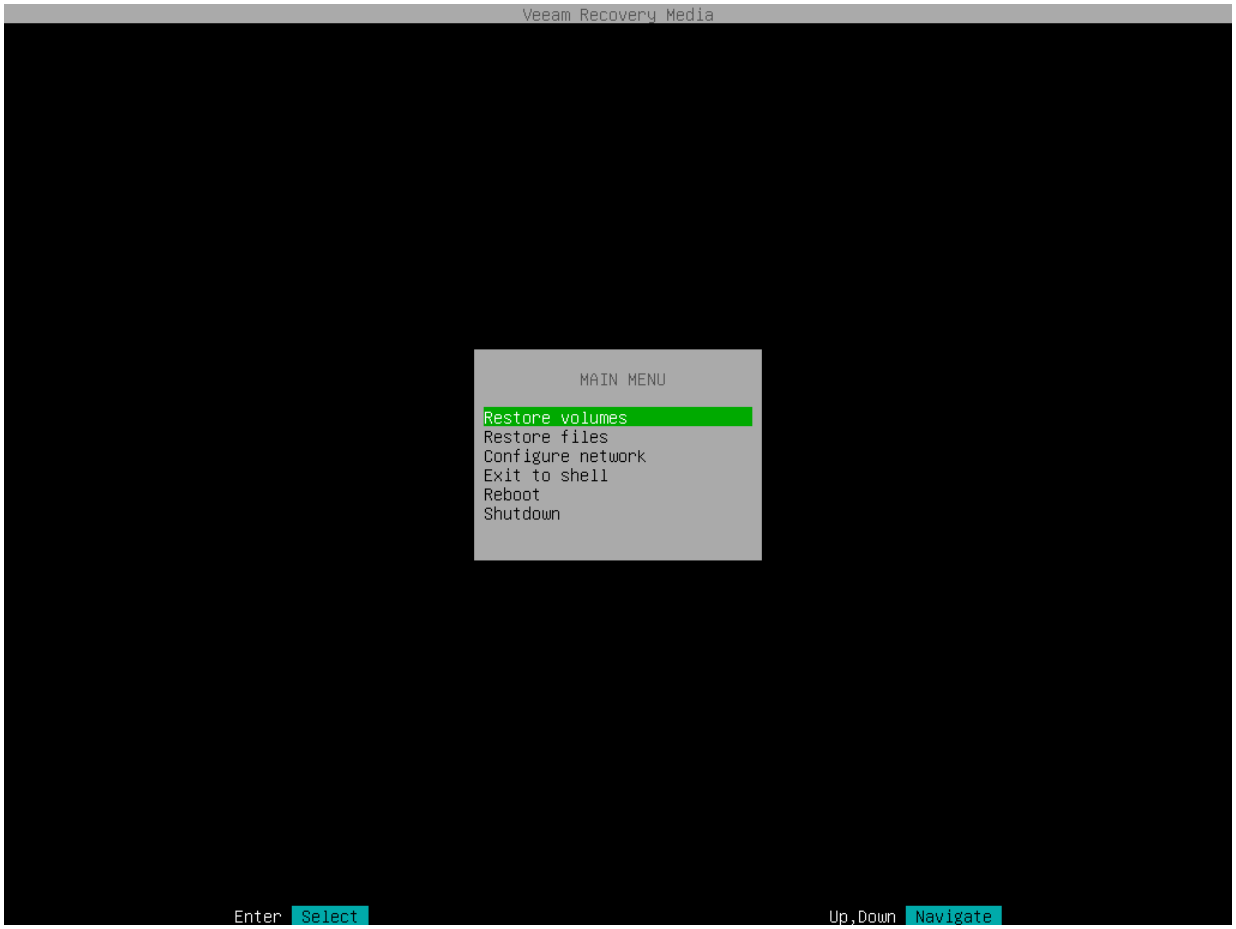
c. Switch to the **Continue** button with the **Tab** key and press **Enter**.



8. Make sure that network settings are specified correctly and configure the network adapter if necessary. To learn more, see Configure Network Settings.
9. Choose the necessary recovery option. Veeam Agent offers the following tools:
   - **Restore volumes** — the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
   - **Restore files** — the File Level Restore wizard to restore files and folders to the original location or to a new location.
   - **Exit to shell** — Linux shell prompt with standard utilities to diagnose problems and fix errors.

TIP

To stop working with the Veeam Recovery Media and shut down or restart your computer, in the Veeam Recovery Media main menu, select the **Reboot** or **Shutdown** option and press **Enter**.

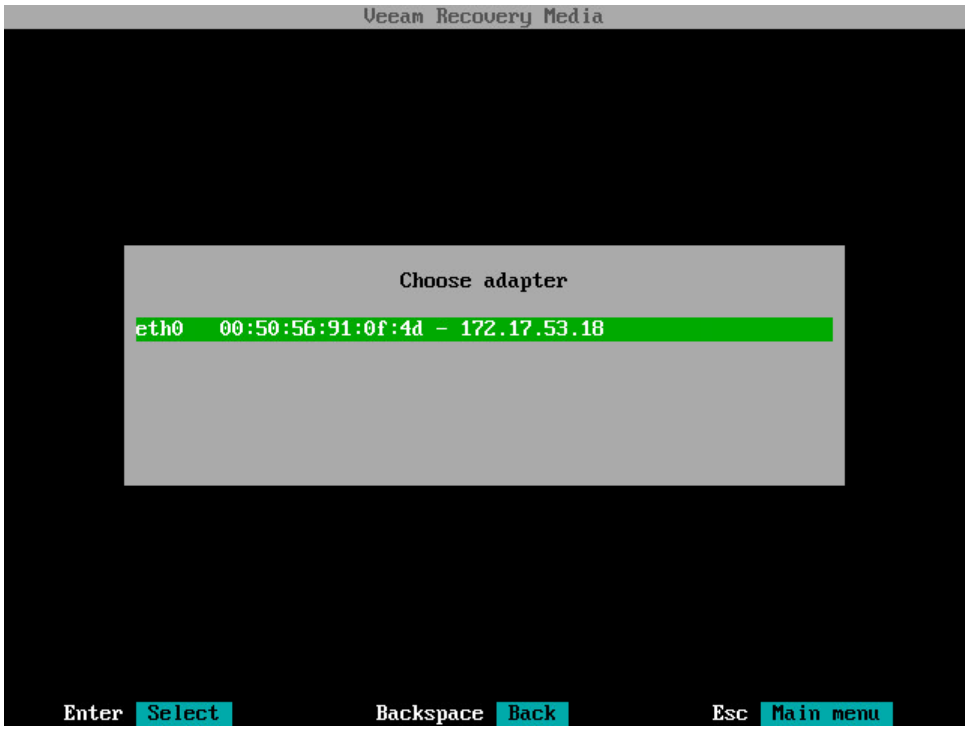


Step 2. Configure Network Settings

To open the **Network settings** dialog, in the Veeam Recovery Media main menu, select the **Configure network** option and press **Enter**.

Veeam Agent for Linux will display network adapters that are available on the system. If there is a DHCP server in your network, Veeam Agent will configure the network settings automatically and display the IP address assigned to the network adapter. You can then press **Esc** to return to the Veeam Recovery Media main menu and launch the Volume Restore wizard.

You can manually configure TPC/IP v4 settings for network adapters if necessary. To learn more, see Specifying Network Settings.



Specifying Network Settings

You can manually configure TPC/IP v4 settings for the network adapter. To configure network settings:

1. In the **Choose adapter** list, select the network adapter that you want to use to connect to the network shared folder or Veeam backup repository where the backup resides and press **Enter**.
2. In the **Configure adapter** dialog, select the **Manual** option and press **Enter**.
3. In the **Adapter settings** dialog, specify the following network settings:
   - IP address
   - Subnet mask
   - Default gateway
   - DNS server
4. Select the **Apply** button and press **Enter**.

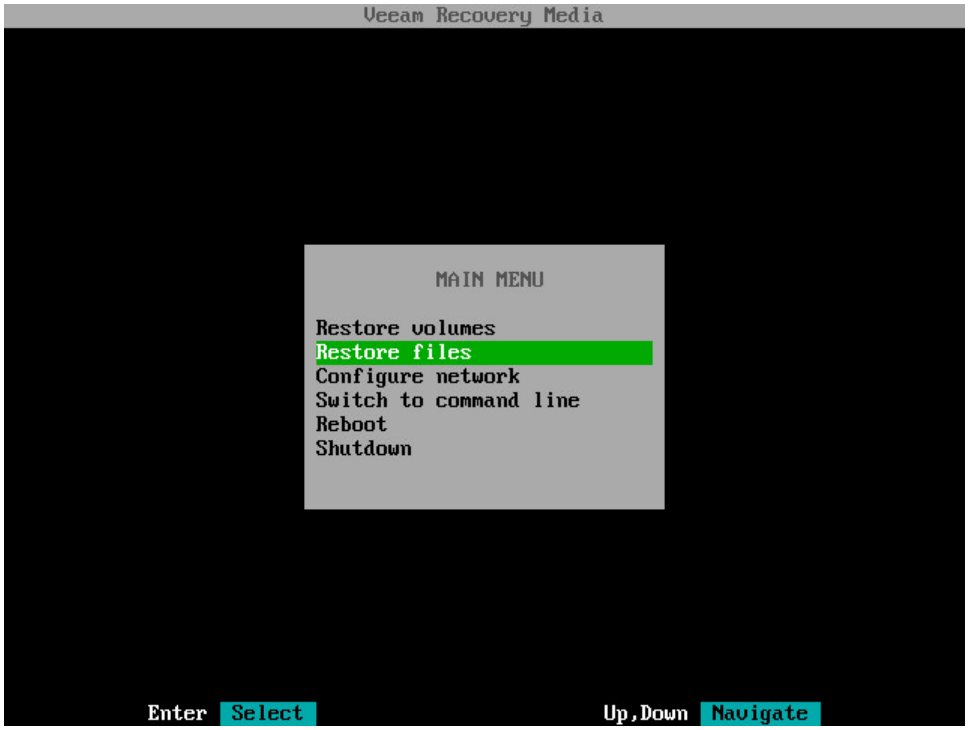
If there is a DHCP server in your network, you can return to automatic IP addressing:

1. In the **Choose adapter** list, select the necessary network adapter and press **Enter**.
2. In the **Configure adapter** dialog, select the **Auto** option and press **Enter**. Veeam Agent will automatically configure network settings for the adapter.



Step 3. Launch File Level Restore Wizard

To launch the file-level restore wizard, in the Veeam Recovery Media main menu, select **Restore files** and press **Enter**.



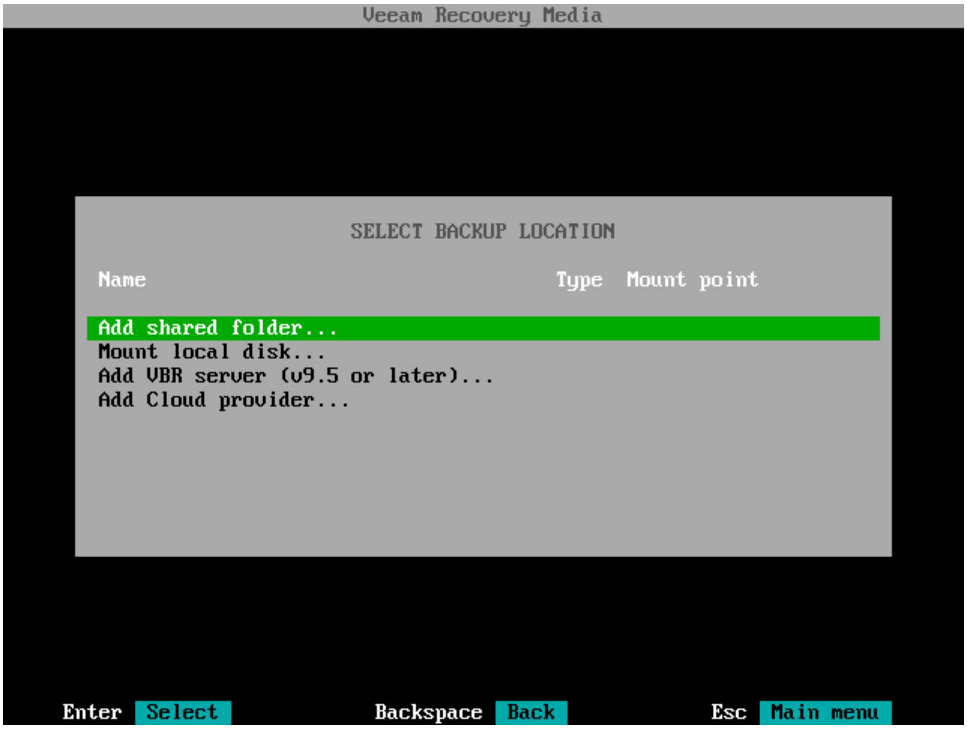
Step 4. Select Backup Location

At the **Select backup location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

To recover data from backup, you need to mount the backup storage on which the backup file resides to the recovery image OS file system. Veeam Agent automatically mounts external USB drives that are connected to the computer and displays them in the list of available backup locations. You can select the necessary device and press **Enter** to pass to the Browse for backup files step of the wizard.

If the backup file is located in a network shared folder or on a local drive, select one of the following options:

- **Add shared folder** — select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the Mount shared folder step of the wizard.
- **Mount local disk** — select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. With this option selected, you will pass to the Select local disk step of the wizard.
- **Add VBR server** — select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the Specify backup server parameters step of the wizard.
- **Add Cloud provider** — select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the Specify Cloud provider parameters step of the wizard.



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.

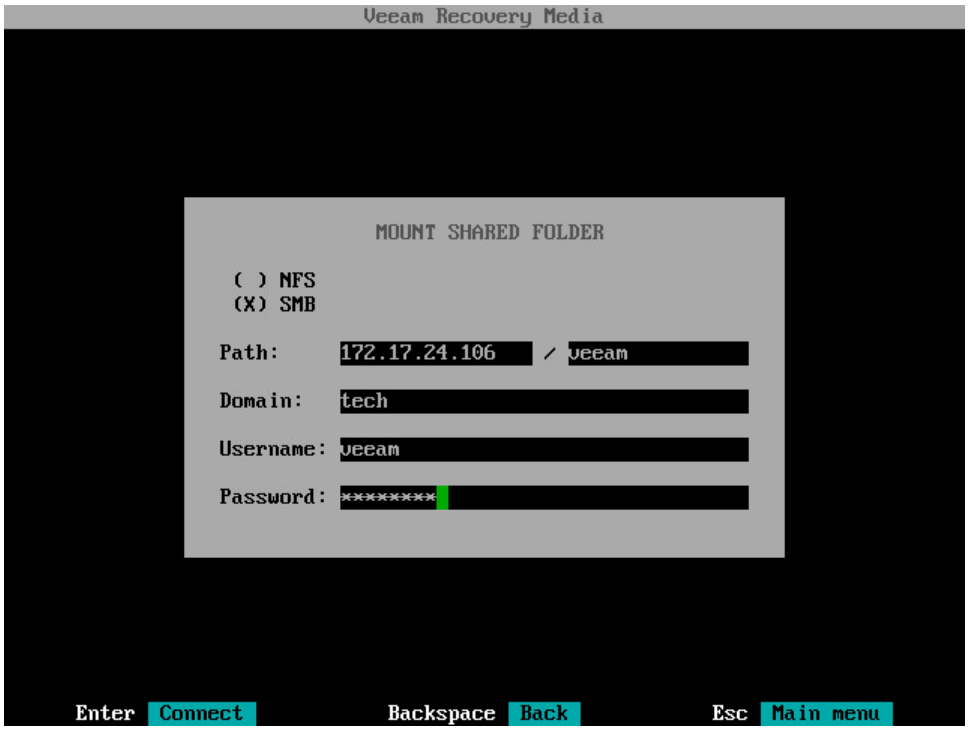
Specifying Shared Folder Settings

The **Mount shared folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. Select the type of a network shared folder:
   - **NFS** — to connect to a network shared folder using the NFS protocol.
   - **SMB** — to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Path** field, specify the network shared folder name in the *SERVER/DIRECTORY* format: type an IP address or domain name of the server and the name of the network shared folder in which the backup file resides.
3. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
4. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
5. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.

> **TIP**
>
> You can mount several network shared folders to work with backup files that are stored in different locations if needed. To do this, return to the Select Backup Location step of the wizard and select the **Add shared folder** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted network shared folders and browse for a backup file located on the necessary storage.

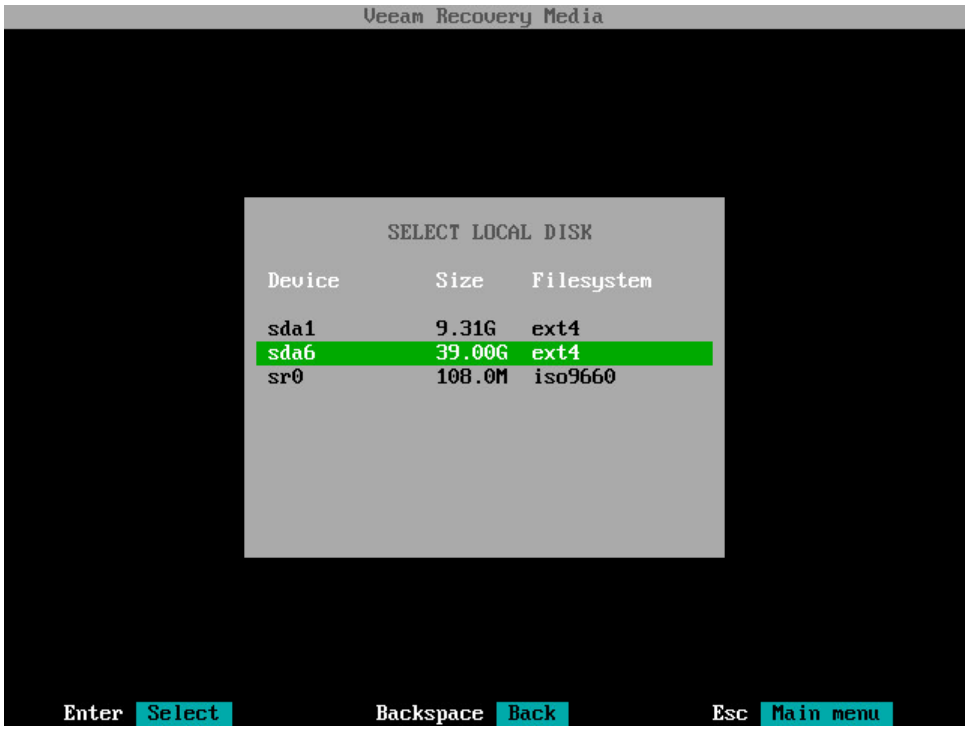


Selecting Local Drive

The **Select local disk** step of the wizard is available if you have selected to restore data from a backup file located on a computer drive.

In the list of devices, select the necessary disk or disk partition and press **Enter**. Veeam Agent will mount the selected device to the `/media` directory of the recovery image OS file system and display content of the directory.

> **TIP**
>
> You can mount several devices to work with backup files that are stored in different locations if needed. To do this, return to the Select Backup Location step of the wizard and select the **Mount local disk** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted devices and browse for a backup file located on the necessary storage.

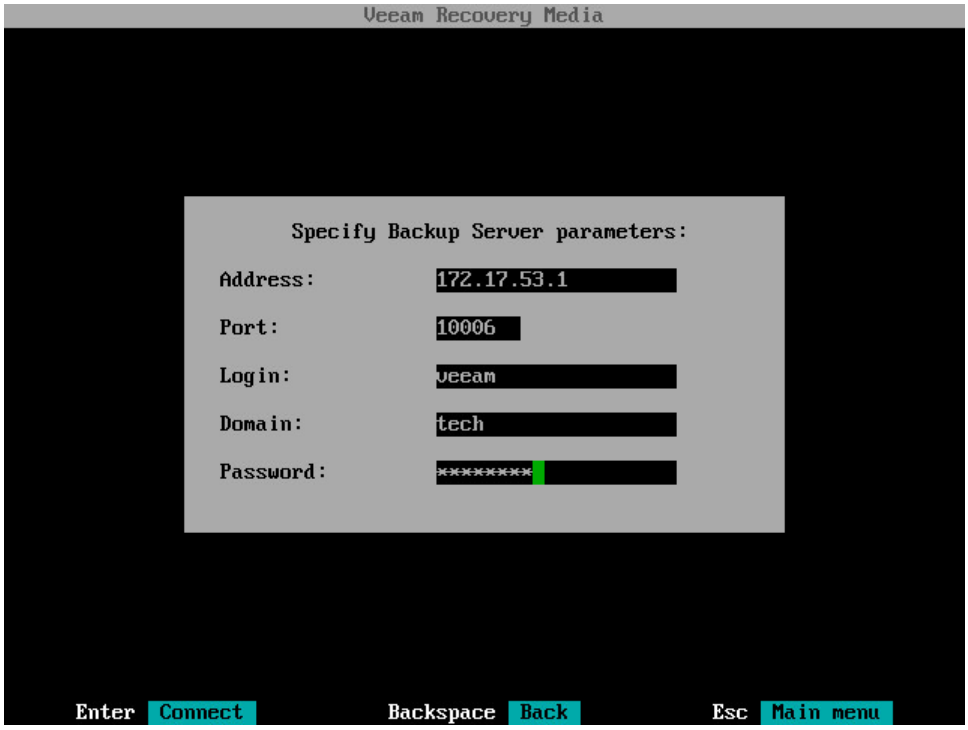


Specifying Backup Server Parameters

The **Specify backup server parameters** step of the wizard is available if you have selected to restore data from a backup repository managed by the Veeam backup server.

Specify settings for the Veeam backup server that manages the backup repository where the backup file resides:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
4. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.
5. In the **Password** field, type a password of the account that has access to the Veeam backup repository.
6. Press **Enter**. Veeam Agent will connect to the Veeam backup server, and you will pass immediately to the Backup step of the wizard.



Specifying Service Provider Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. Specify service provider settings.
2. Verify the TLS certificate.
3. Specify user account settings.

Specifying Service Provider Settings

The **Specify Cloud provider parameters** step of the wizard is available if you have selected to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.
3. Press **Enter**. Veeam Agent will connect to the service provider and display the Certificate details window.



Verifying TLS Certificate

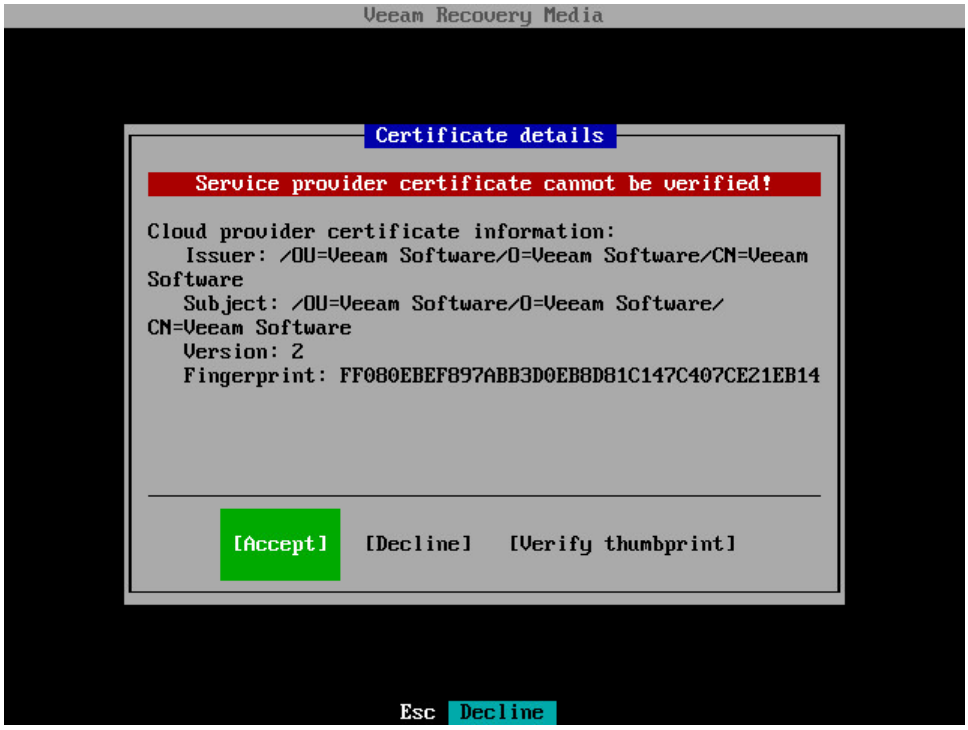
In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate.

- To accept the TLS certificate, select the **Accept** button with the **Tab** key and press **Enter**.
- [Optional] To verify the TLS certificate with a thumbprint, do the following:
    a. Select the **Verify thumbprint** button with the **Tab** key and press **Enter**.

b. In the **Thumbprint verification** field, enter the thumbprint that you obtained from the SP.

c. Switch to the **Verify** button and press **Enter**. Veeam Agent for Linux will check if the thumbprint that you entered matches the thumbprint of the obtained TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

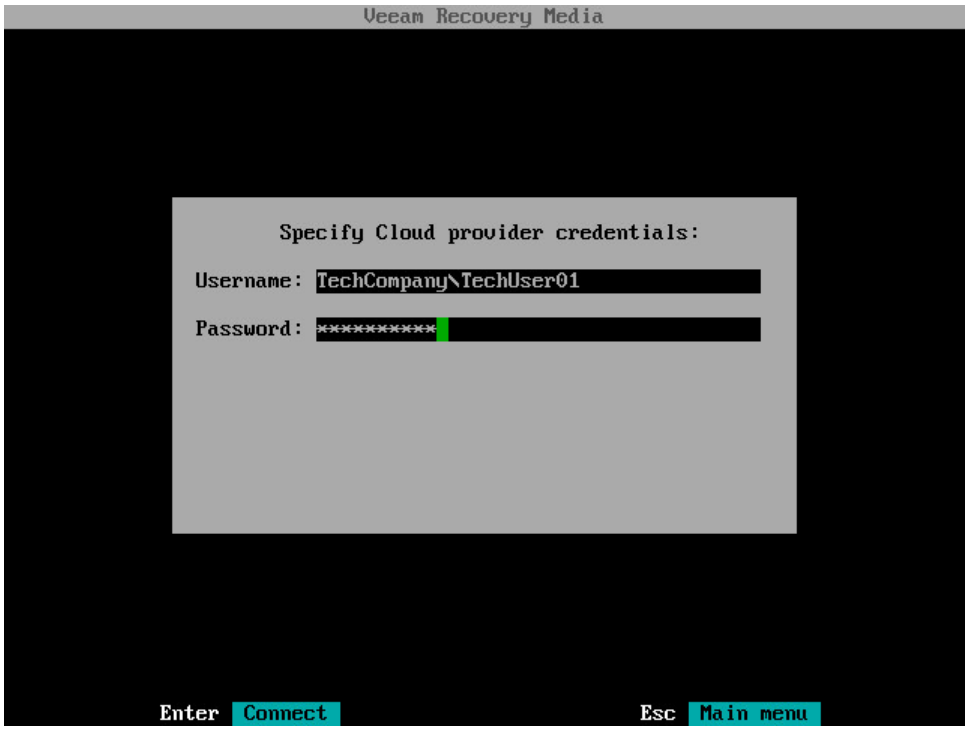


Specifying User Account Settings

The **Specify Cloud provider credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

1. In the **Username** field, type a name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.
2. In the **Password** field, provide a password for the tenant or subtenant account.

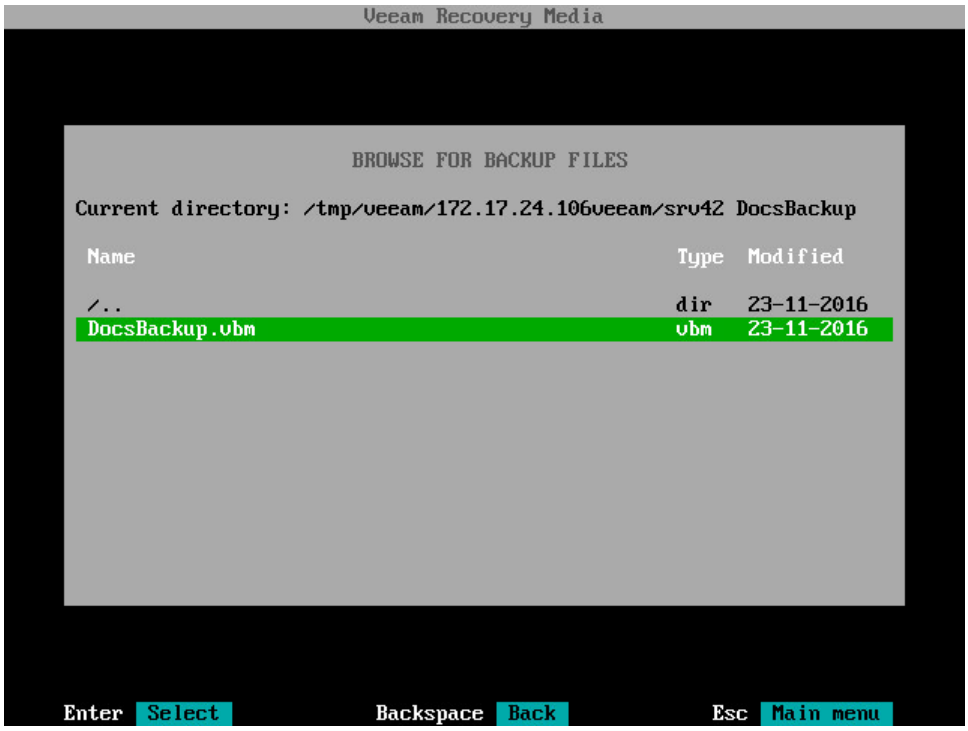
3. Press **Enter**. Veeam Agent for Linux will connect to the cloud repository, and you will pass immediately to the Backup step of the wizard.



Step 6. Browse for Backup File

At the **Browse for backup files** step of the wizard, select the backup file that you plan to use for volume-level restore:

1. In the file system tree, select a directory in which the backup file you plan to use for restore resides:
   - Use **Up** and **Down** arrow keys to select a directory.
   - Use the **Enter** key to open the necessary directory.
2. In the directory where the backup file resides, select the backup file and press **Enter**.



Step 7. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.

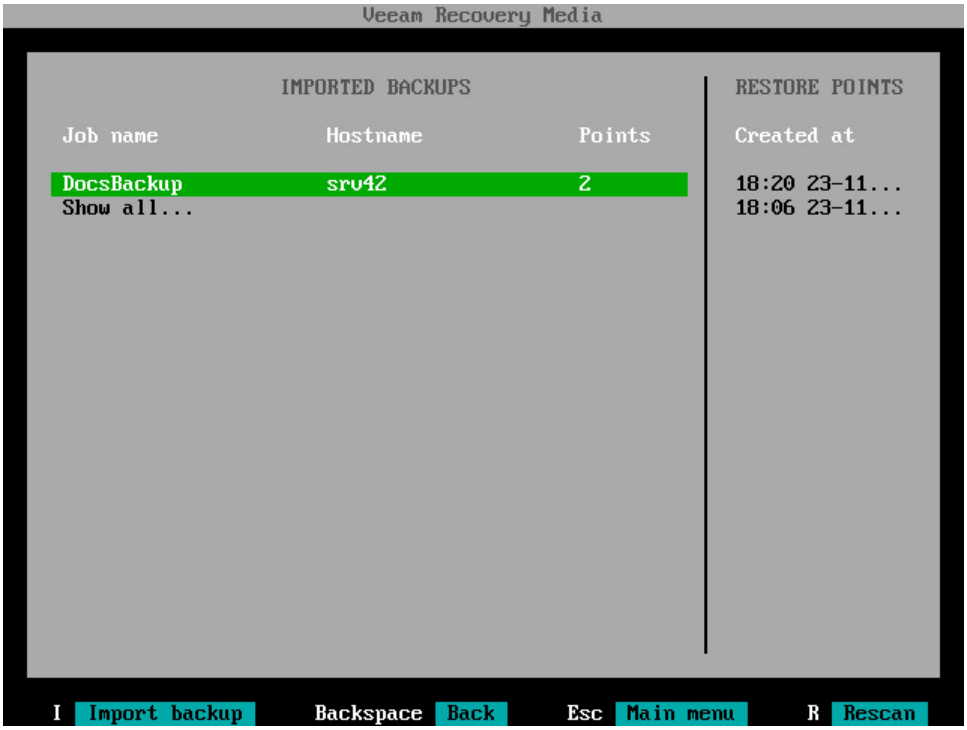
The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays information about backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported backups** pane, ensure that the backup from which you want to recover data is selected and press **Enter**.

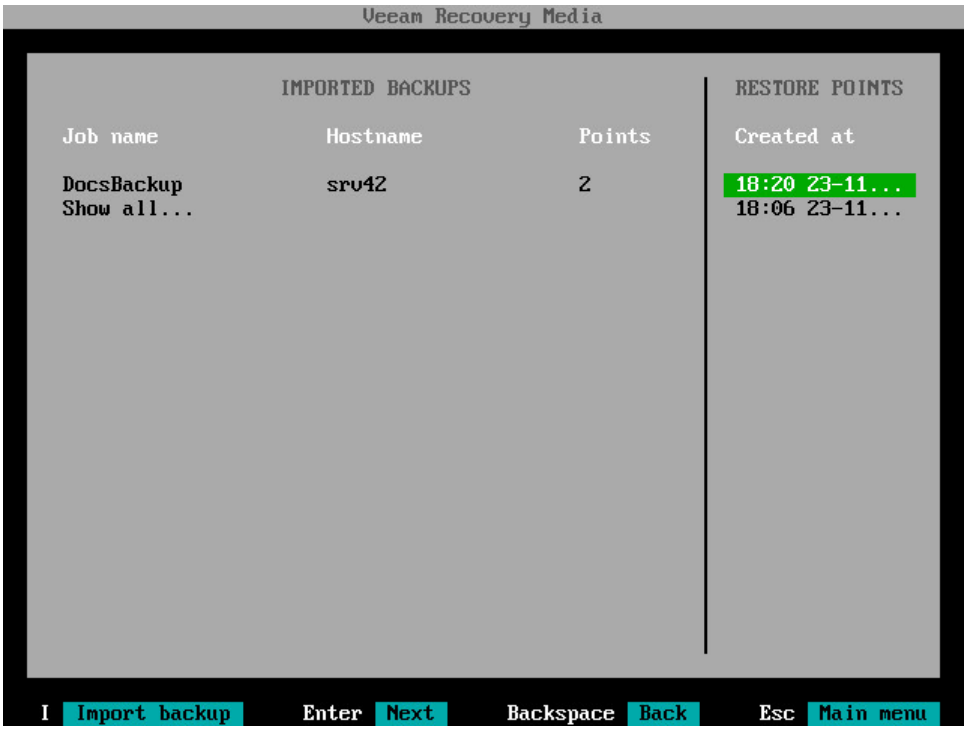
   If you want to select another backup, press the '**i**' key and browse for the necessary backup file. To learn more, see Locate Backup File.



2. In the **Restore points** pane, select with **Up** and **Down** keys the restore point from which you want to recover data and press **Enter**.

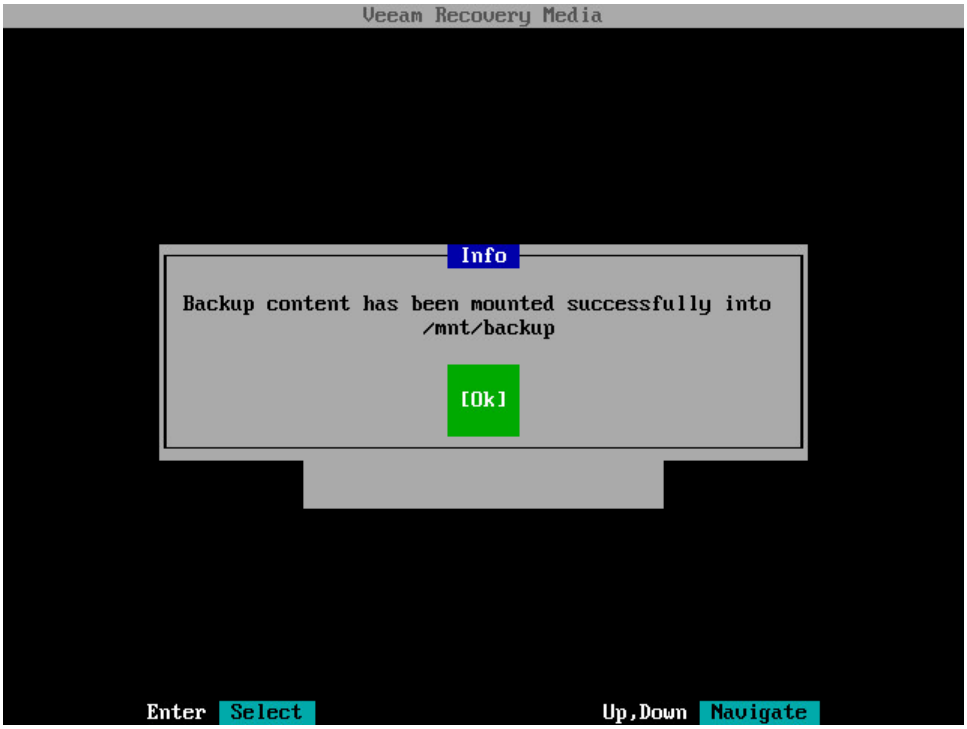
NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To lean more, see Restoring Data from Encrypted Backups.



3. Veeam Agent will mount the content of the backup file to the `/mnt/backup` directory in the recovery image OS file system and display a notification window with the corresponding message. Press **Enter** to proceed to the File Level Restore wizard menu, open the file manager and save restored files.

   When you perform file-level restore with the File Level Restore wizard, Veeam Agent always mounts the backup to the `/mnt/backup` directory. If you want to specify another directory for backup mount, you can perform file-level restore with the Veeam Agent command line interface. To learn more, see Restoring Files and Folders with Command Line Interface.



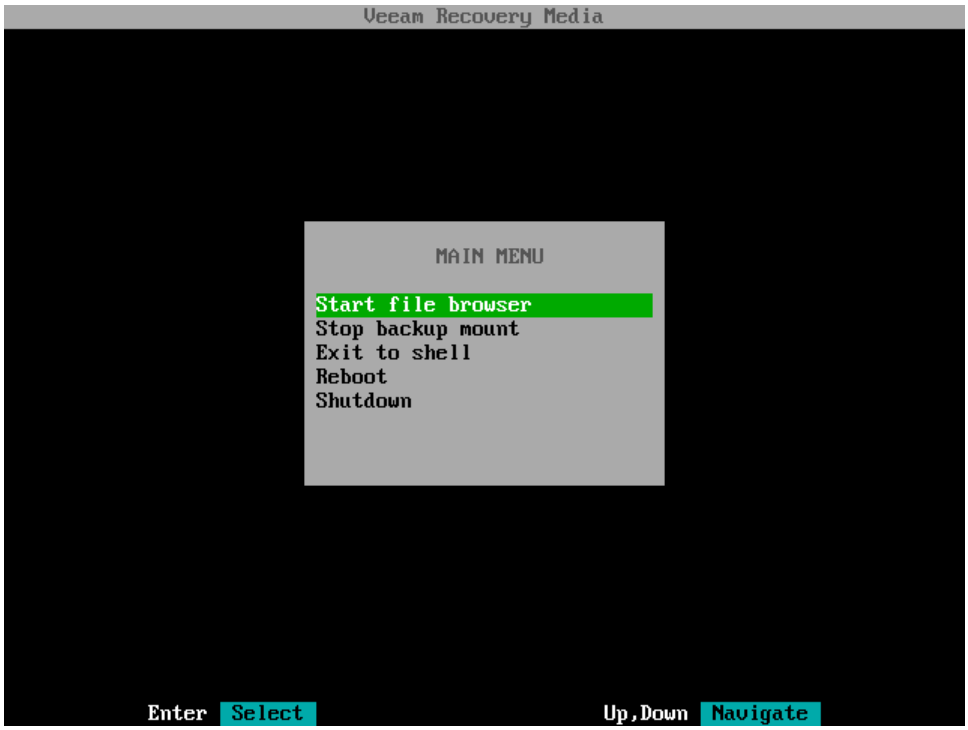
Step 8. Save Restored Files

When the backup file content is mounted to the recovery image OS file system, Veeam Agent opens the File Level Restore wizard menu displaying a list of available operations.

You can perform the following operations with file-level backup:

- **Start file commander** — select this option if you want to start the file manager and work with restored files and folders. To learn more, see Working with Midnight Commander.
- **Stop backup mount** — select this option if you want to stop the backup mount session and unmount the backup file content from the `/mnt/backup` directory of the recovery image OS file system. To learn more, see Stopping Backup Mount Session.
- **Exit to shell** — select this option if you want to open the Linux shell prompt and use common Linux command-line tools.

> **TIP**
>
> To stop working with the Veeam Recovery Media and shut down or restart your computer, in the File Level Restore wizard menu, select the **Reboot** or **Shutdown** option and press **Enter**.



Working with Midnight Commander

To work with restored files and folders, you can use Midnight Commander — a file manager that is included into the Veeam Recovery Media. With the Midnight Commander file manager, you can browse the mounted backup content and file system on your computer, and save restored files and folders to the original location or to a new location.

To launch the file manager, in the File Level Restore wizard menu, select **Start file browser** and press **Enter**.

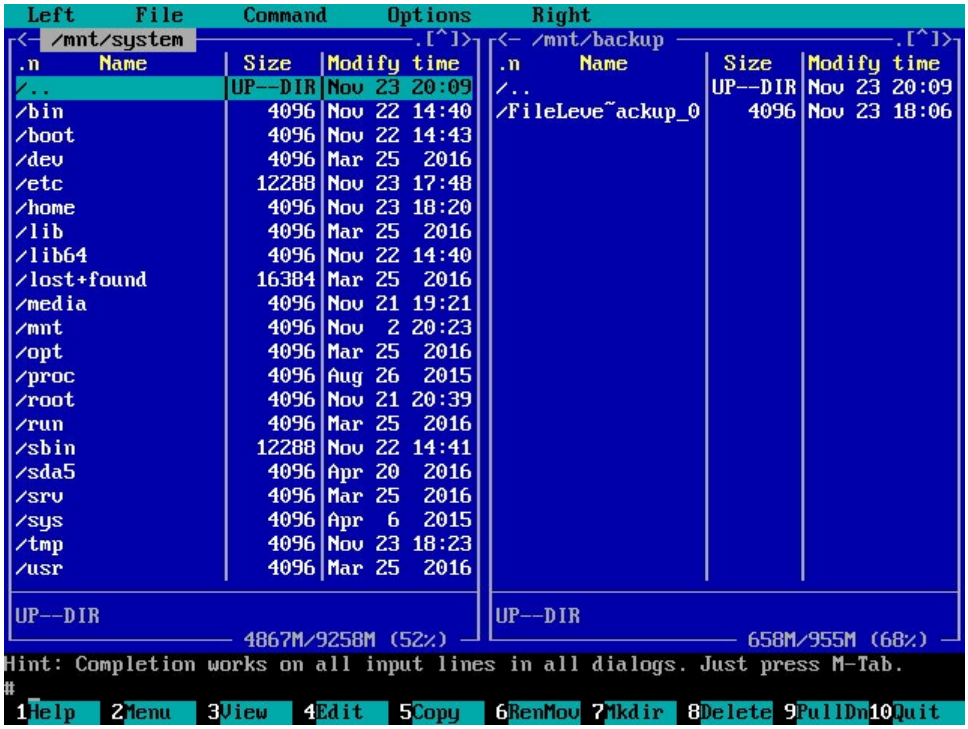
When you launch Midnight Commander, Veeam Agent displays in the file manager the directory with the backup content and your computer's file system:

- In the left pane, Veeam Agent displays a directory of your computer's file system mounted under the `/mnt/system` directory of the recovery image OS file system. By default, Veeam Agent mounts to the recovery image OS file system the following volumes of your computer:
  - If you use a volume-level backup for file-level restore, Veeam Agent detects the partition table in the backup, mounts to the `/mnt/system` directory block devices that represent volumes of your computer with the same names as volumes in the backup. For example, if your volume-level backup contains `/dev/sda1` and `/dev/sda6` volumes with `/` and `/home` mount points, Veeam Agent will mount to the `/mnt/system` directory both root (`/`) and `/home` partitions.
  - If you use a file-level backup for file-level restore, Veeam Agent mounts to the `/mnt/system` directory only the system volume of your computer, for example, `/dev/sda1`. If you want to save restored files and folders to a directory on another computer volume or to a network shared folder, you need to mount this volume or folder manually. To mount a target storage for restored files:
    i. In Midnight Commander, press **F10** to close the file manager.
    ii. In the **File Level Restore** wizard menu, select the **Exit to shell** option and press **Enter**.
    iii. Mount the target storage for the restored files and folders with the `mount` command.
- In the right pane, Veeam Agent displays a directory in which the backup content is mounted. Veeam Agent mounts the backup content under the `/mnt/backup` folder.

While the Midnight Commander file manager is open, you can perform the following operations with restored files and folders:

- Save files to initial location
- Save files to a new location

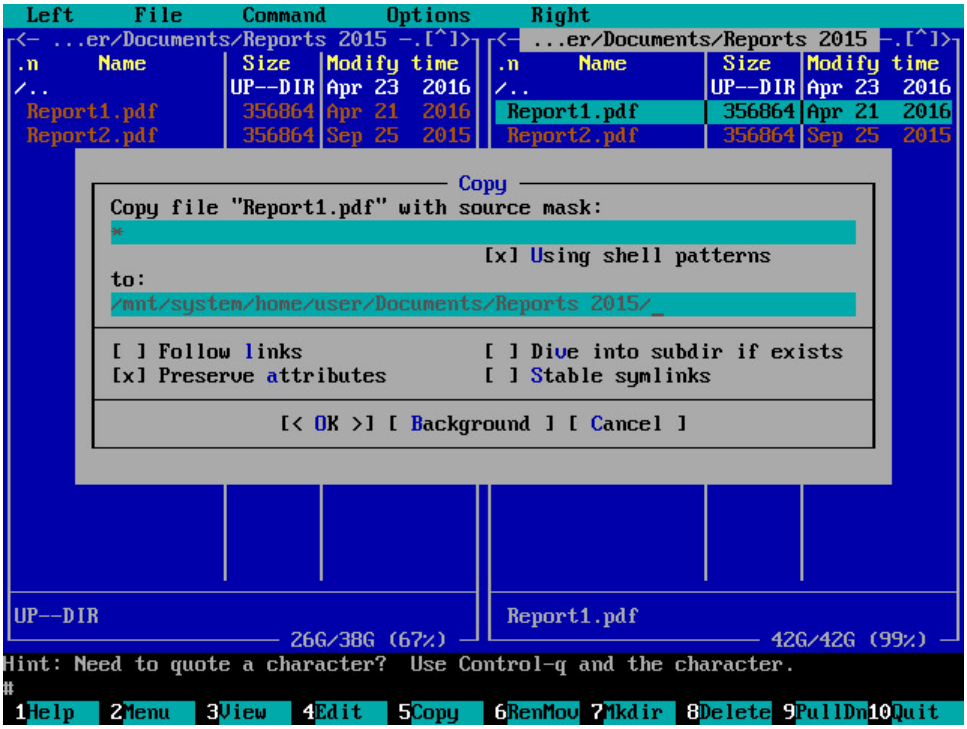
After you finish working with files and folders, finish working with the Veeam Recovery Media.



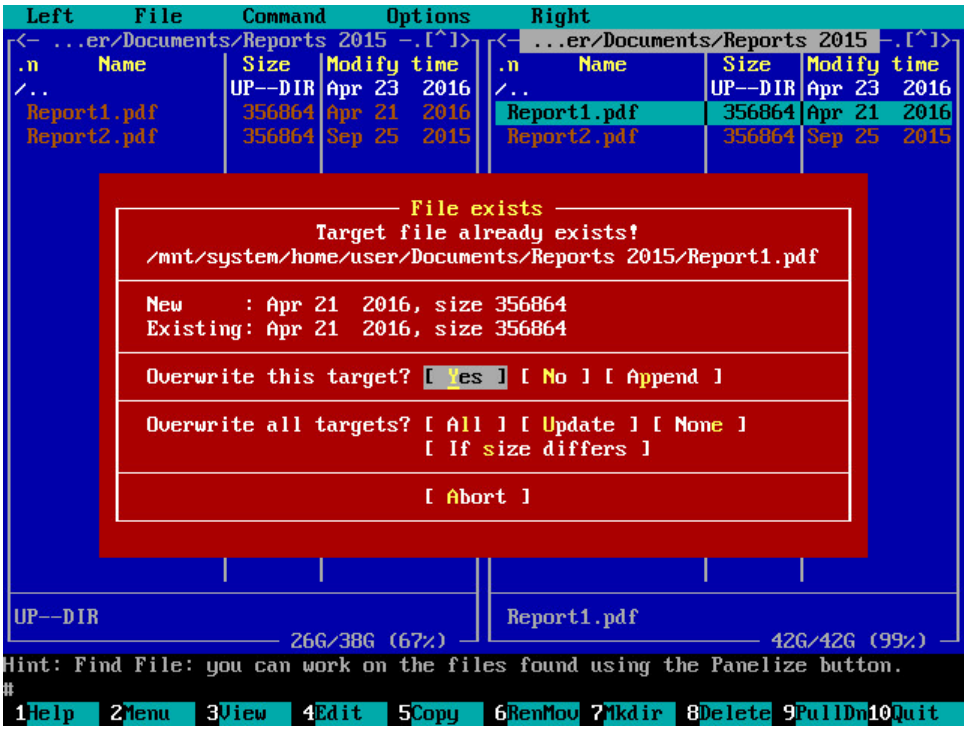
Saving Files to Initial Location

To save restored files or folders to their initial location on your computer, do the following:

1. In the left pane of the file manager window, open the directory in your computer's file system in which the backed-up file or folder that you want to restore originally resided.
2. In the right pane of the file manager window, open the directory that contains the file or folder in the backup that you want to restore to its original location.
3. Select the file or folder that you want to restore and press **F5**.
4. In the **Copy** dialog window, review the file or folder copy settings, select **Ok** and press **Enter**.



5. If the file or folder you want to restore exists in its original location, Midnight Commander will display a warning. In the warning window, select the necessary operation with the target file or folder and press **Enter**. Midnight Commander will save the file or folder in its original location.

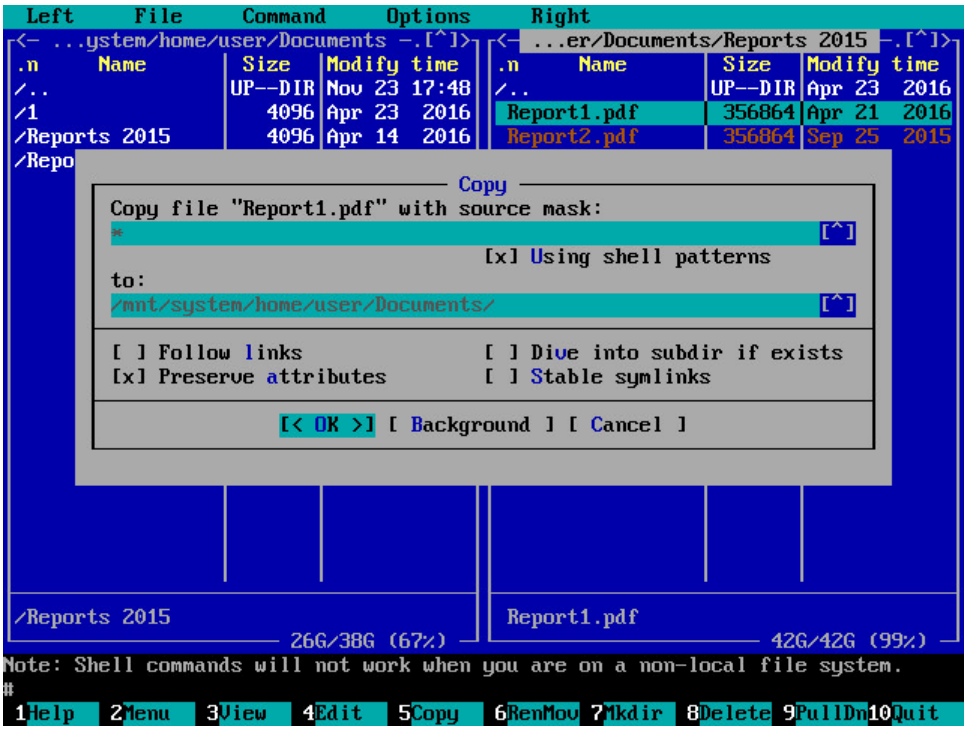


6. After you finish working with files and folders, press **F10** to close the file manager.

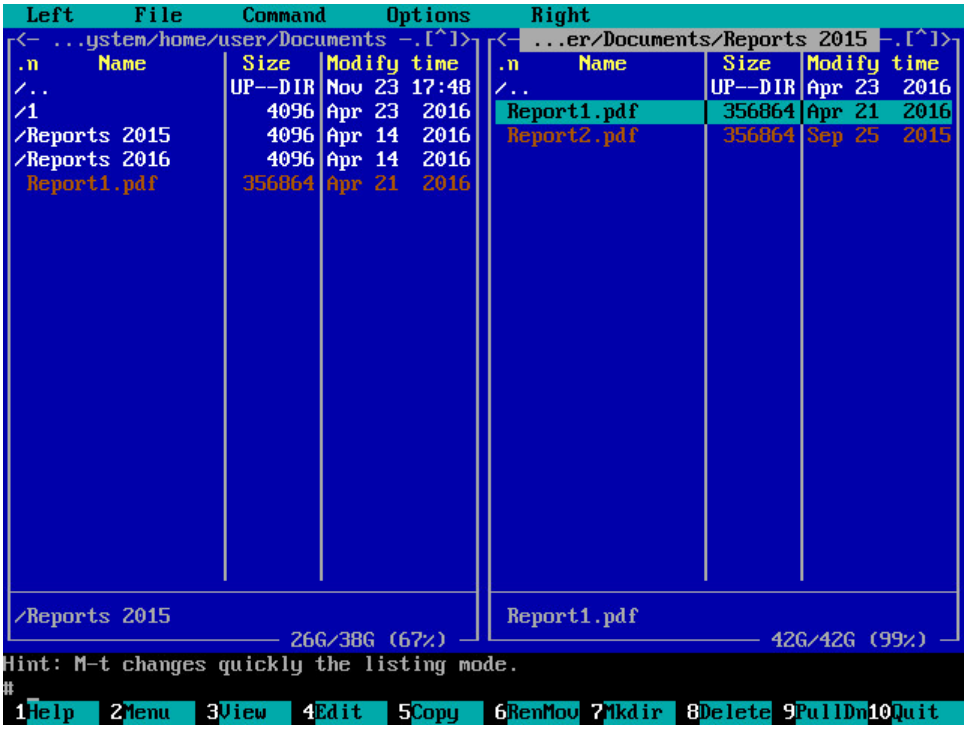
Saving Files to New Location

To save restored files or folders to a new location on your computer or to a network shared folder, do the following:

1. In the left pane of the file manager window, open the directory in your computer's file system in which you want to restore a file or folder.
2. In the right pane of the file manager window, open the directory that contains the file or folder in the backup that you want to restore.
3. Select the file or folder that you want to restore and press **F5**.
4. In the **Copy** dialog window, review the file or folder copy settings, select **Ok** and press **Enter**.



5. Midnight Commander will save the file or folder to the specified location.

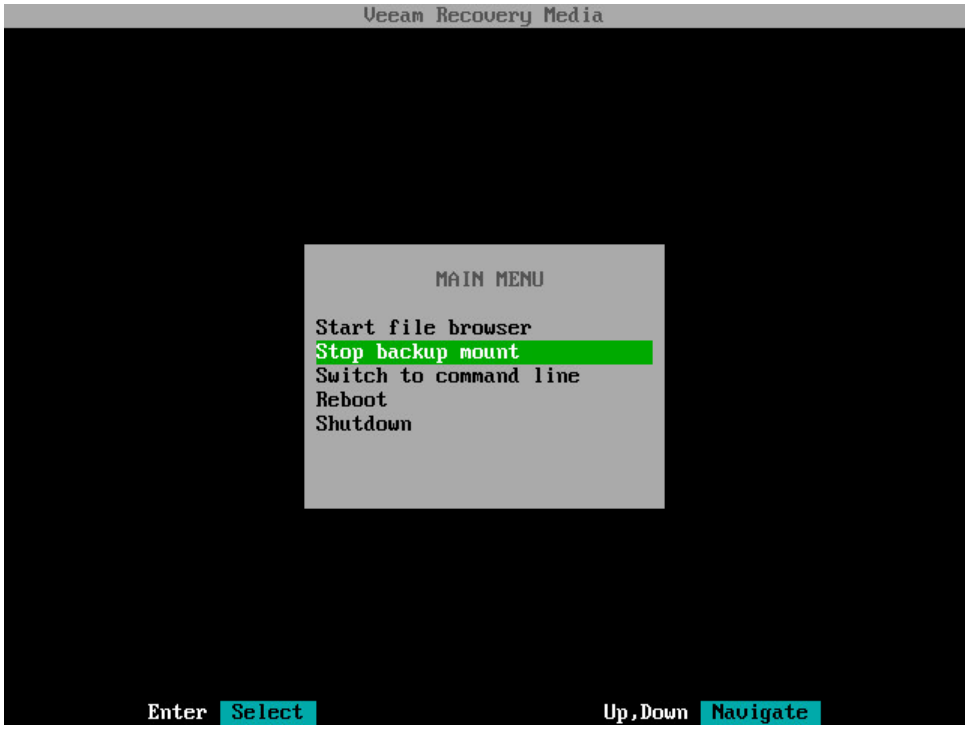


6. After you finish working with files and folders, press **F10** to close the file manager.

Stopping Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. To unmount a backup, you need to stop the backup mount session. This may be required, for example, if you want to stop working with files and folders in one backup and mount another backup for file-level restore.

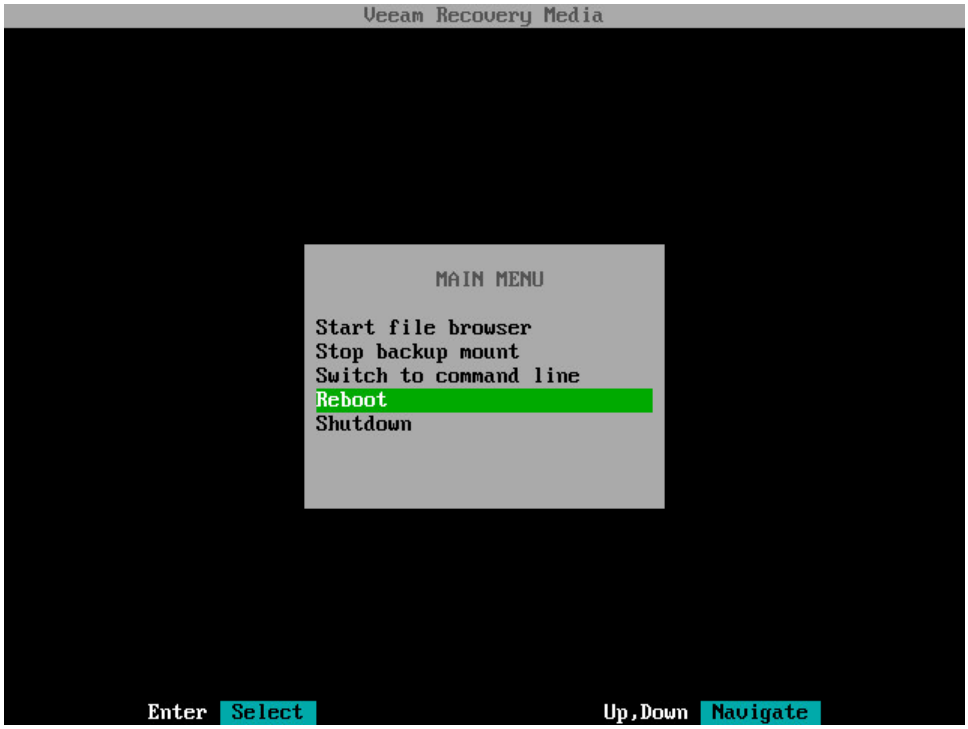
To stop the backup mount session with the Veeam Recovery Media, in the File Level Restore wizard menu, select the **Stop backup mount** option and press **Enter**. Veeam Agent will stop the backup mount session, unmount the backup from the `/mnt/backup` directory of the recovery image OS file system, exit the File Level Restore wizard and display the Veeam Recovery Media main menu.



Step 9. Finish Working with Veeam Recovery Media

When the restore operation completes, finish working with the Veeam Recovery Media and start your operating system.

1. Eject the media or removable storage device with the recovery image.
2. In the File Level Recovery wizard menu or Veeam Recovery Media main menu, select the **Reboot** option and press **Enter**.
3. Wait for your Linux operating system to start.



Restoring Volumes with Command Line Interface

You can restore a specific computer volume or all volumes from the volume-level backup.

> **NOTE**
>
> You cannot use the Veeam Agent for Linux command line interface to restore Btrfs subvolumes.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

You can use Veeam Agent commands to restore volumes from a backup or restore point:

- Restore from backup

  When you restore a volume from the backup, Veeam Agent will automatically select the latest restore point in the backup. The volume will be restored to the state in which the volume was at the time when the latest restore point was created.

- Restore from a restore point

  When you restore a volume from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

Before You Begin

Before you begin the volume-level restore process, check the following prerequisites:

- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders and on Veeam backup repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.

Volume-level restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Linux swap space is hosted.
- You cannot restore a volume to the volume where the backup file that you use for restore is located.

To overcome the first two limitations, you can boot from the recovery image and use the Veeam Recovery Media tools for volume-level restore. To learn more, see Restoring from Veeam Recovery Media.

Restoring from Backup

With Veeam Agent command line interface, you can restore volumes from the backup. When you restore a volume from the backup, Veeam Agent automatically selects the latest restore point in the backup and restores the volume to the state in which the volume was at the time when the latest restore point was created.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list --all
```

In the list of backups, Veeam Agent displays information about all Veeam Agent for Linux backups stored in all backup repositories. If Veeam Agent is connected to a Veeam backup server, all Veeam Agent for Linux backups that are kept on Veeam backup repositories managed by this server also appear in this list.

The `--all` parameter is optional. If you do not use this option, Veeam Agent displays information about backups created by the current Veeam Agent computer only.

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as *Imported* in the **Repository** column. For information about the import procedure, see Importing Backups.
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name                   Backup ID                                  Repositor
y    Created at
srv01 SystemBackup         {45f074d2-d2d9-423d-84e9-8f1798b08d4c}  Repository_
1  2016-11-11 17:37
srv01 DocumentsBackup      {ea64a7e5-038a-4c86-970a-6d59d4cf3968}  Repository_
1  2016-11-11 18:30
srv01 HomePartitionBackup  {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b}  Repository_
2  2016-11-15 11:28
wrk01 SystemBackup         {951ac571-dd29-45ac-8624-79b8ccb45863}  Repository_
2  2016-11-13 15:26
wrk02 SystemBackup         {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167}  Repository_
3  2016-11-13 15:59
```

TIP

If you want to recover data from a backup that is stored in another location, for example, a backup created with another instance of Veeam Agent, you can import such backup into the Veeam Agent database on your computer. To learn more, see Importing Backups.

Step 2. Explore Backup Content

To view detailed information about specific backup, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup_id> — ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Name	Name of the volume in the backup.
Device	Path to the block device that represents the volume.
FS UUID	File system ID.
Offset	Position of the volume on the computer disk.
Size	Size of the volume in the backup.

For example:

```
user@srv01:~$ veeamconfig backup show --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
     Machine name: srv01 HomePartitionBackup
        Name:       [sda6]
        Device:     [/dev/sda6]
        FS UUID:    [4967f2eb-e8bb-48fe-a694-5ba67b9030a5]
        Offset:     [11813257216] bytes (23072768 sectors)
        Size:       [41872785408] bytes (81782784 sectors)
```

Step 3. Start Restore Process

To start the process of volume-level restore from the backup, use the following command:

```
veeamconfig backup restore --id <backup_id> --targetdev <target_volume> --backu
pdev <volume_in_backup>
```

where:

- `<backup_id>` — ID of the backup.
- `<target_volume>` — path to a block device that represents a volume on your computer that you want to recover.
- `<volume_in_backup>` — path to a block device that represents a volume in the backup.

  This parameter is optional. If you do not specify this parameter, Veeam Agent will restore from the backup a volume that has the same name as a `<target_volume>`.

For example:

```
user@srv01:~$ veeamconfig backup restore --id 4f75bb20-a6b6-4323-9287-1c6c8cecc
b6b --targetdev /dev/sdb --backupdev /dev/sda6
Restoring backup.
Backup: 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Devices:
   Device in current system: [/dev/sdb]  In backup: [/dev/sda6];
You are sure? (y/n)
y
Volume restore from backup has been started.
Session ID: [{0b72ef45-4c88-4639-b940-ad3828b1cd4e}].
Logs stored in: [/var/log/veeam/Restore/Session_{0b72ef45-4c88-4639-b940-ad3828
b1cd4e}].
```

IMPORTANT

You can restore a backed-up volume only to a target volume that is not used by your Linux OS (that does not have file system mount points). For example, you can add a new disk to your computer and restore a volume in the backup to this disk. To restore a volume to its original location or to another volume used by Linux OS, you should boot from the Veeam Recovery Media and perform volume-level restore with the Volume Restore wizard.

Step 4. Monitor Restore Process

You can monitor the restore process by viewing the restore session log in the command line interface.

To view Veeam Agent for Linux session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

`<session_id>` — ID of the restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 0b72ef45-4c88-4639-b940-ad3828b1cd4e
2016-11-27 11:04:04 UTC {b141f32a-3e77-45a6-b55a-c100a1464d67} [info] Job start
ed at 2016-11-27 14:04:04
2016-11-27 11:04:04 UTC {9b60ac03-2de0-4fe2-a00e-bec556d98ee8} [info] Starting
volume restore
2016-11-27 11:04:07 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [processing] sdb
2016-11-27 11:04:15 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 512.0
kB at 58.6kB/s (0%)
...
2016-11-27 11:14:35 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB
at 10.6MB/s (97%)
2016-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB
at 10.6MB/s (100%)
2016-11-27 11:14:37 UTC {00add723-cbfa-4cc8-b299-d2349a051d6f} [warn] /dev/sdb
has a duplicate filesystem UUID
2016-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb resto
red 6.5GB at 10.6MB/s
2016-11-27 11:14:37 UTC {8b8742a2-1c80-4e14-bbf1-45a3612bc3a7} [info] Volume re
store completed
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see Viewing Session Status.

Restoring from Restore Point

With Veeam Agent command line interface, you can restore volumes from the specific restore point. When you restore a volume from the restore point, you can select the necessary restore point in the backup to recover data to a desired point in time.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list --all
```

In the list of backups, Veeam Agent displays information about all Veeam Agent for Linux backups stored in all backup repositories. If Veeam Agent is connected to a Veeam backup server, all Veeam Agent for Linux backups that are kept on Veeam backup repositories managed by this server also appear in this list.

The `--all` parameter is optional. If you do not use this option, Veeam Agent displays information about backups created by the current Veeam Agent computer only.

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as *Imported* in the **Repository** column. For information about the import procedure, see Importing Backups.
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name                   Backup ID                                  Repositor
y    Created at
srv01 SystemBackup         {45f074d2-d2d9-423d-84e9-8f1798b08d4c}  Repository_
1  2016-11-11 17:37
srv01 DocumentsBackup      {ea64a7e5-038a-4c86-970a-6d59d4cf3968}  Repository_
1  2016-11-11 18:30
srv01 HomePartitionBackup  {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b}  Repository_
2  2016-11-15 11:28
wrk01 SystemBackup         {951ac571-dd29-45ac-8624-79b8ccb45863}  Repository_
2  2016-11-13 15:26
wrk02 SystemBackup         {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167}  Repository_
3  2016-11-13 15:59
```

TIP

If you want to recover data from a backup that is stored in another location, for example, a backup created with another instance of Veeam Agent, you can import such backup into the Veeam Agent database on your computer. To learn more, see Importing Backups.

Step 2. Explore Restore Points

To view information about restore points in the backup, use the following command:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where

`<backup_id>` — ID of the backup for which you want to view information on restore points.

You can view the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: • Full • Increment.
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: • True • False

For example:

```
user@srv01:~$ veeamconfig backup info --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Job name                   OIB ID                                   Type       C
reated at       Is corrupt
srv01 HomePartitionBackup  {23cb927d-5e2d-42fe-a4a4-e5f254a6413e}  Full       2
016-11-15 11:28  false
srv01 HomePartitionBackup  {25e31075-4c30-4d67-86a6-293c0887f4eb}  Increment  2
016-11-15 11:58  false
srv01 HomePartitionBackup  {9375140d-720a-4d3e-a69b-ab9cf60d53fa}  Increment  2
016-11-27 13:15  false
```

or

```
user@srv01:~$ veeamconfig point list --backupid 4f75bb20-a6b6-4323-9287-1c6c8ce
ccb6b
Job name                   OIB ID                                   Type       C
reated at       Is corrupt
srv01 HomePartitionBackup  {23cb927d-5e2d-42fe-a4a4-e5f254a6413e}  Full       2
016-11-15 11:28  false
srv01 HomePartitionBackup  {25e31075-4c30-4d67-86a6-293c0887f4eb}  Increment  2
016-11-15 11:58  false
srv01 HomePartitionBackup  {9375140d-720a-4d3e-a69b-ab9cf60d53fa}  Increment  2
016-11-27 13:15  false
```

Step 3. Start Restore Process

To start the process of volume-level restore from the specific restore point, use the following command:

```
veeamconfig point restore --id <point_id> --targetdev <target_volume> --backupd
ev <volume_in_backup>
```

where:

- `<point_id>` — ID of the restore point.
- `<target_volume>` — path to a block device that represents a volume on your computer that you want to recover.
- `<volume_in_backup>` — path to a block device that represents a volume in the backup.

  This parameter is optional. If you do not specify this parameter, Veeam Agent will restore from the backup a volume that has the same name as a `<target_volume>`.

For example:

```
user@srv01:~$ veeamconfig point restore --id 9375140d-720a-4d3e-a69b-ab9cf60d53
fa --backupdev /dev/sda6 --targetdev /dev/sdb
Restoring point.
Restore point: 9375140d-720a-4d3e-a69b-ab9cf60d53fa
Devices:
   Device in current system: [/dev/sdb]  In backup: [/dev/sda6];
You are sure? (y/n)
y
Volume restore by point has been started.
Session ID: [{697d9348-9001-4845-8764-3cc4fb3f296b}].
Logs stored in: [/var/log/veeam/Restore/Session_{697d9348-9001-4845-8764-3cc4fb
3f296b}].
```

IMPORTANT

You can restore a backed-up volume only to a target volume that is not used by your Linux OS (that does not have file system mount points). For example, you can add a new disk to your computer and restore a volume in the backup to this disk. To restore a volume to its original location or to another volume used by Linux OS, you should boot from the Veeam Recovery Media and perform volume-level restore with the Volume Restore wizard.

Step 4. Monitor Restore Process

You can monitor the restore process by viewing the restore session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

`<session_id>` — ID of the restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 697d9348-9001-4845-8764-3cc4fb3f296b
2016-11-27 10:35:47 UTC {b9604775-d265-4537-b98e-848fd77c7375} [info] Job start
ed at 2016-11-27 13:35:47
2016-11-27 10:35:47 UTC {ed66a1f6-5216-4596-a7b5-be10dd10c32f} [info] Starting
volume restore
2016-11-27 10:35:50 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [processing] sdb
2016-11-27 10:35:59 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb 512.0
kB at 59.1kB/s (0%)
...
2016-11-27 10:46:27 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb 6.5GB
at 10.5MB/s (100%)
2016-11-27 10:46:28 UTC {dae118c8-eb7c-4e14-9832-f0bfd089b329} [warn] /dev/sdb
has a duplicate filesystem UUID
2016-11-27 10:46:28 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb resto
red 6.5GB at 10.5MB/s
2016-11-27 10:46:28 UTC {a21a89d9-d0ca-4f5c-8399-28ae599f2f1c} [info] Volume re
store completed
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see Viewing Session Status.

Restoring Files and Folders with Recovery Wizard

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent publishes the backup content directly into the computer file system. You can browse to files and folders in the backup, restore files and folders to their initial location, copy files and folders to a new location or simply target applications to restored files and work with them as usual.

Before You Begin

Before you begin the file-level restore process, check the following prerequisites:

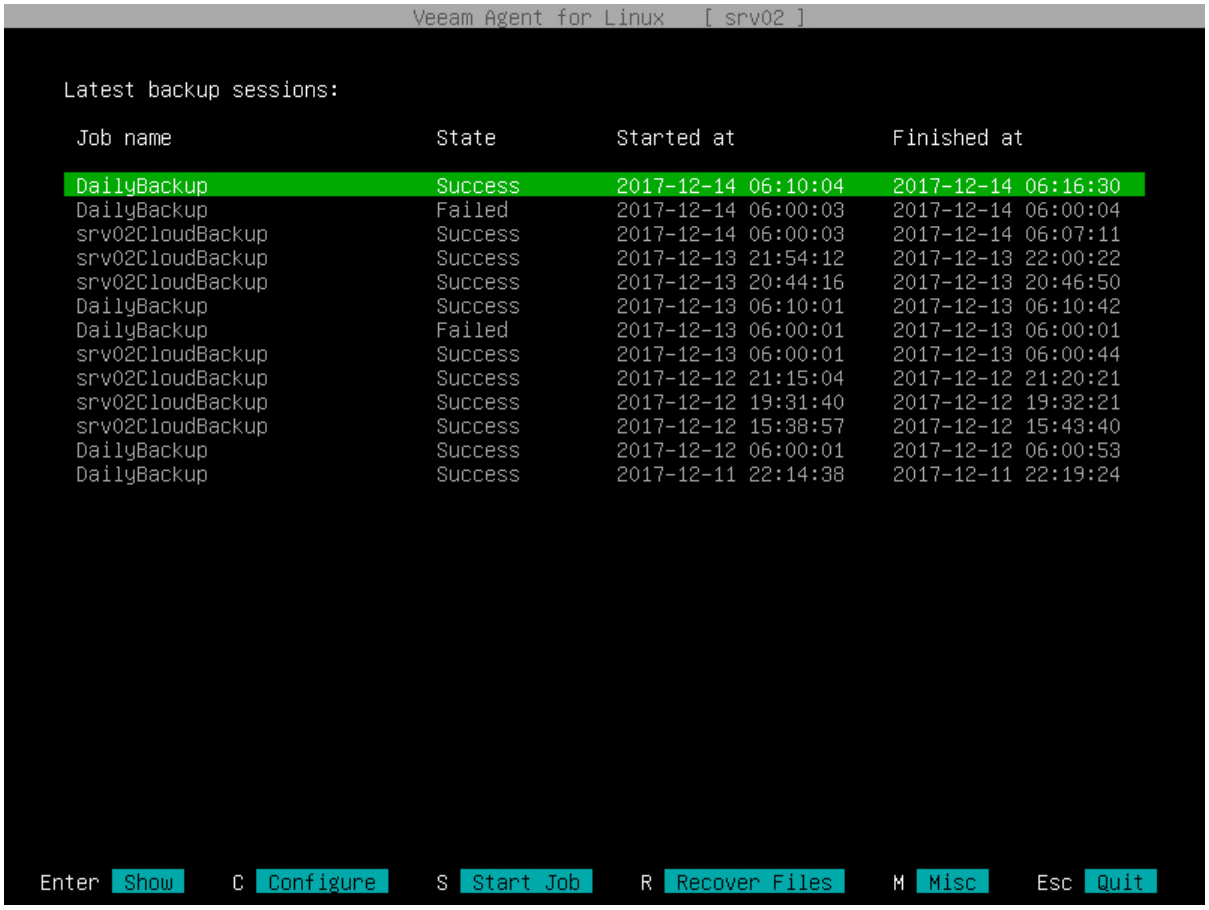
- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.
- [For backups of Btrfs file system] A machine on which you perform file-level restore must run the same or later Linux kernel version as the machine on which the backup was created.

  For example, you created a backup of a machine that runs Linux kernel version 4.14. If you perform file-level restore from this backup on another machine that runs Linux kernel 2.6, the file-level restore process will fail.

Step 1. Launch File Level Restore Wizard

To launch the **File Level Restore** wizard, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the '**r**' key to proceed to the File Level Restore wizard.



Step 2. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.

The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays available backups and information about each backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

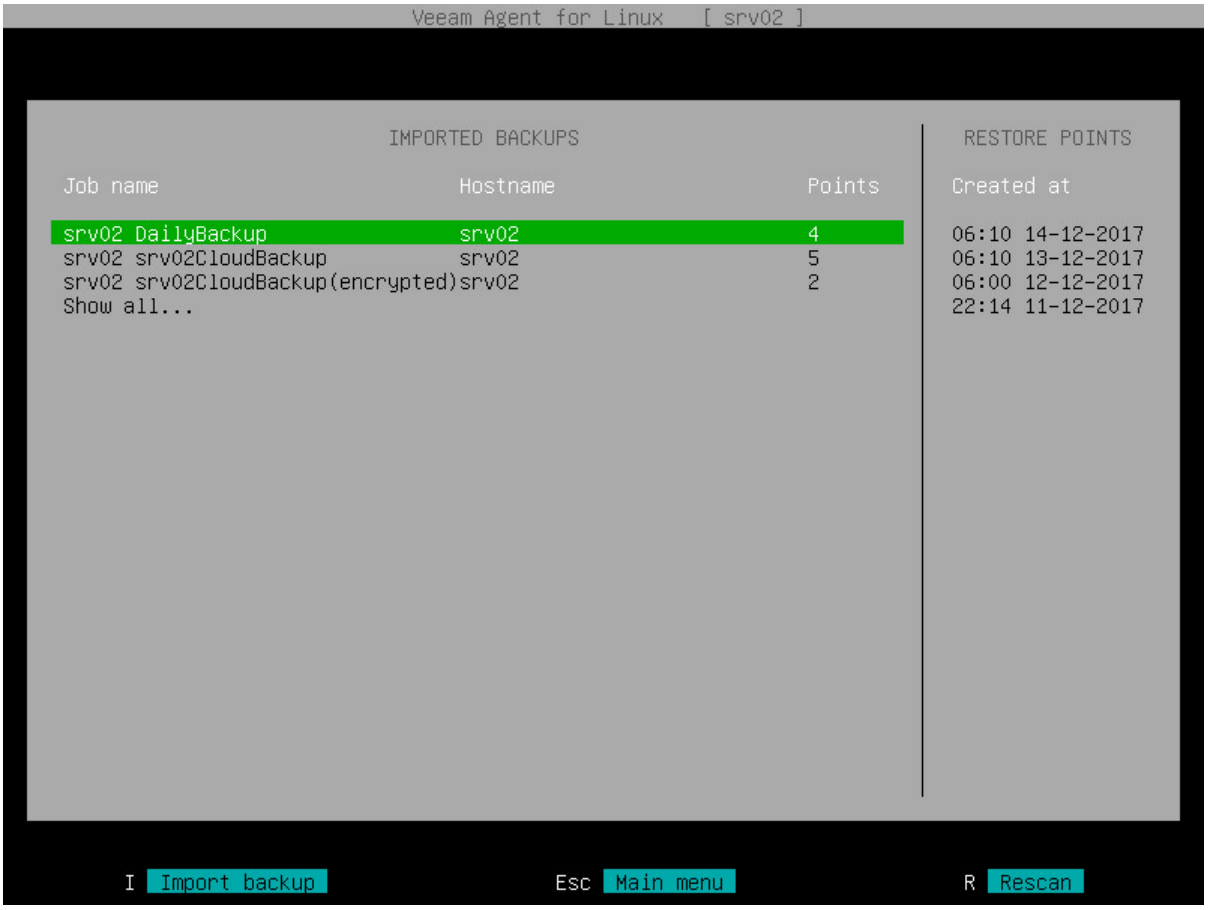
1. In the **Imported backups** pane, select with **Up** and **Down** keys the backup from which you want to recover data and press **Enter**.

   In the list of backups, Veeam Agent displays backups that were created by backup jobs configured with Veeam Agent on your computer. If Veeam Agent for Linux is connected to a Veeam Backup & Replication server or a Veeam Cloud Connect service provider, backups created in the Veeam backup repository or cloud repository also appear in the list.

   By default, Veeam Agent displays in the list only those backups in the Veeam backup repository that were created under your account. If you used an account to which the Veeam Backup Administrator role is assigned to connect to the Veeam backup server, you can also view all Veeam Agent backups that are stored in the Veeam backup repository to which Veeam Agent is connected. To view such backups, click the **Show all** link at the bottom of the list.

   If Veeam Agent fails to display backups stored in the Veeam backup repository for some reason, you can press the '**r**' key to rescan the backup repository. Veeam Agent will try to reconnect to the Veeam backup server and refresh the list of backups.

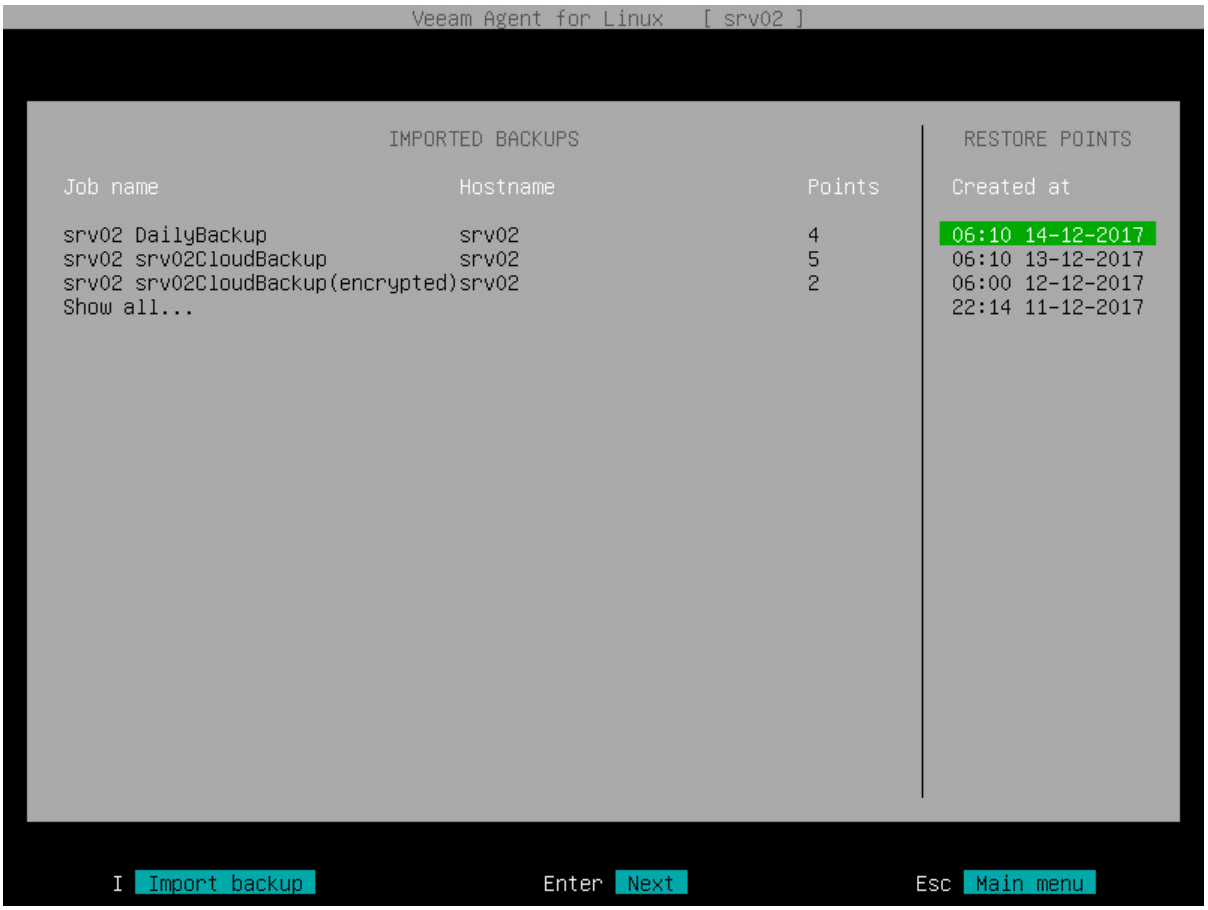
   If you want to recover data from a backup that is stored in another location, for example, a backup created with another instance of Veeam Agent in a network shared folder, you can import such backup. Press the '**i**' key, browse to the directory in which the backup file resides and select the necessary backup file. The selected backup file will be added to the list of backups.



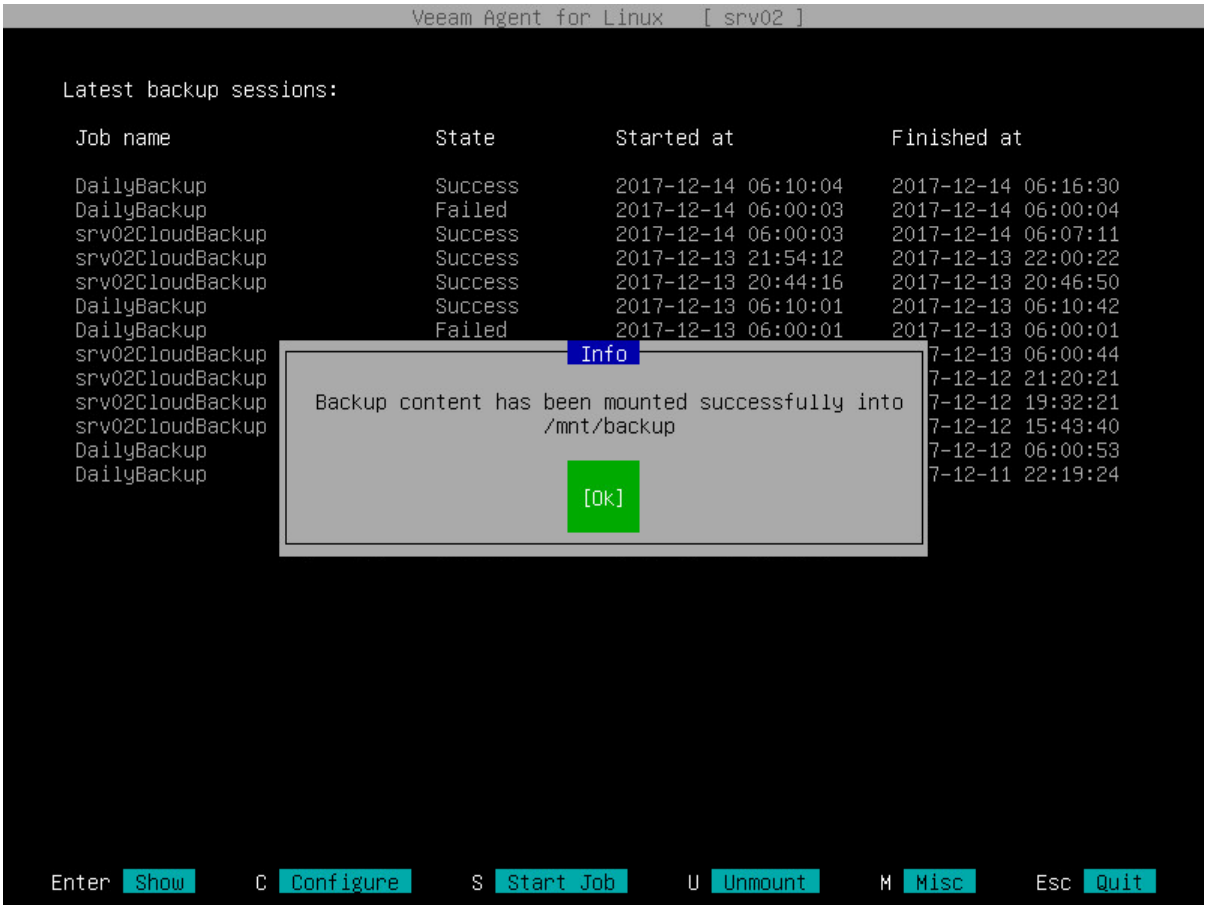
2. In the **Restore points** pane, select with **Up** and **Down** keys the restore point from which you want to recover data and press **Enter**.

NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To lean more, see Restoring Data from Encrypted Backups.



3. Veeam Agent will mount the content of the backup file to the `/mnt/backup` directory in the computer's file system and display a notification window with the corresponding message. Press **Enter** to close the window and return to the Veeam Agent control panel.



TIP

When you finish working with restored files and folders, you can unmount the backup from the `/mnt/backup` folder. To learn more, see Stop Backup Mount Session.

Step 3. Save Restored Files

When the backup file content is mounted to the `/mnt/backup` directory in the computer's file system, you can use Linux command line utilities or preferred file browser to work with restored files and directories. You can browse for files and directories in the mounted backup and copy files and directories that you want to restore to their initial location or to a new location.

In the following example, the restored file `Report1.pdf` is copied from the mounted backup to the new location with Linux command line utilities:

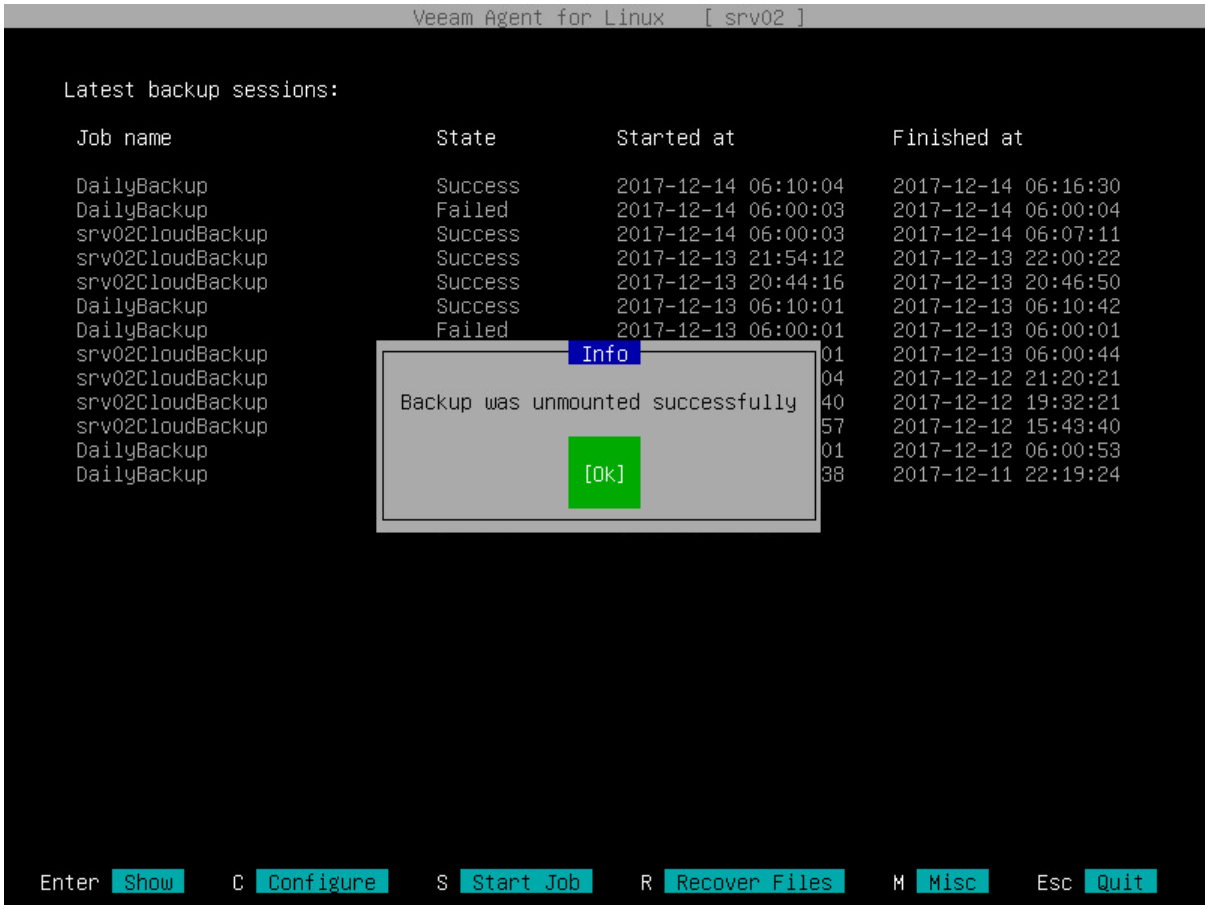
```
user@srv01:~$ ls Documents/
Reports
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/
Report1.pdf  Report2.pdf
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Repo
rt1.pdf Documents/
user@srv01:~$ ls Documents/
Report1.pdf  Reports
```

Step 4. Stop Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. To unmount a backup, you need to stop the backup mount session. This may be required, for example, if you want to stop working with files and folders in one backup and mount another backup for file-level restore. You can also stop the backup mount session to unmount a backup after you have finished working with restored files and folders.

To stop the backup mount session, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command:
2. In the Veeam Agent control panel, press the '**u**' key to unmount a backup.
3. Veeam Agent will stop the backup mount session and display a notification window. Press **Enter** to close the window and return to the Veeam Agent control panel.



Restoring Files and Folders with Command Line Interface

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent publishes the backup content directly into the computer file system. You can browse to files and folders in the backup, restore files and folders to their initial location, copy files and folders to a new location or simply target applications to restored files and work with them as usual.

With the Veeam Agent command line interface, you can restore files and folders in a more flexible way than with the use of the File Level Restore wizard. In particular, you can specify a directory in which Veeam Agent should mount the backup file content for file-level restore. You can also mount several backups to different directories to work with files and folders restored from different backups simultaneously.

You can use Veeam Agent commands to restore files and folders from backup or from specific restore point:

- Restore from backup

  When you restore files and folders from the backup, Veeam Agent will automatically select the latest restore point in the backup. You can restore files and folders to the state in which they were at the time when the latest restore point was created.

- Restore from a restore point

  When you restore files and folders from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

Before You Begin

Before you begin the file-level restore process, check the following prerequisites:

- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.
- [For backups of Btrfs file system] A machine on which you perform file-level restore must run the same or later Linux kernel version as the machine on which the backup was created.

  For example, you created a backup of a machine that runs Linux kernel version 4.14. If you perform file-level restore from this backup on another machine that runs Linux kernel 2.6, the file-level restore process will fail.

Restoring from Backup

With Veeam Agent command line interface, you can restore files and folders from the backup. When you perform file-level restore from the backup, Veeam Agent for Linux automatically selects the latest restore point in the backup. You can restore files and folders to the state in which they were at the time when the latest restore point was created.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list --all
```

In the list of backups, Veeam Agent displays information about all Veeam Agent for Linux backups stored in all backup repositories. If Veeam Agent is connected to a Veeam backup server, all Veeam Agent for Linux backups that are kept on Veeam backup repositories managed by this server also appear in this list.

The `--all` parameter is optional. If you do not use this option, Veeam Agent displays information about backups created by the current Veeam Agent computer only.

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as *Imported* in the **Repository** column. For information about the import procedure, see Importing Backups.
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name                   Backup ID                                  Repositor
y    Created at
srv01 SystemBackup         {45f074d2-d2d9-423d-84e9-8f1798b08d4c}  Repository_
1  2016-11-11 17:37
srv01 DocumentsBackup      {ea64a7e5-038a-4c86-970a-6d59d4cf3968}  Repository_
1  2016-11-11 18:30
srv01 HomePartitionBackup  {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b}  Repository_
2  2016-11-15 11:28
wrk01 SystemBackup         {951ac571-dd29-45ac-8624-79b8ccb45863}  Repository_
2  2016-11-13 15:26
wrk02 SystemBackup         {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167}  Repository_
3  2016-11-13 15:59
```

TIP

If you want to recover data from a backup that is stored in another location, for example, a backup created with another instance of Veeam Agent, you can import such backup into the Veeam Agent database on your computer. To learn more, see Importing Backups.

Step 2. Explore Backup Content

For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

To view detailed information about specific backup, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup_id> — ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Name	Name of the volume in the backup.
Device	Path to the block device that represents the volume.
FS UUID	File system ID.
Offset	Position of the volume on the computer disk.
Size	Size of the volume in the backup.

For a file-levwel backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Backed up	Backup scope for the file-level backup job.

For example:

```
user@srv01:~$ veeamconfig backup show --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
     Machine name: srv01 DocumentsBackup
        File-level backup
        Backed up:
           /home/user/Documents
```

Step 3. Mount Backup

To mount a backup for file-level restore, use the following command:

```
veeamconfig backup mount --id <backup_id> --mountdir <path>
```

where:

- <backup_id> — ID of the backup that you want to mount to the computer file system for file-level restore.
- <path> — path to the directory to which you want to mount the backup file content.

For example:

```
user@srv01:~$ veeamconfig backup mount --id ea64a7e5-038a-4c86-970a-6d59d4cf396
8 --mountdir /mnt/backup
Backup is mounted.
Session ID: [{2a313184-32d0-4d3a-a1b0-2eebac986047}].
Logs stored in: [/var/log/veeam/Mount/Session_{2a313184-32d0-4d3a-a1b0-2eebac98
6047}].
```

Step 4. Monitor Mount Process and Result

You can monitor the backup mount process by viewing the mount session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

`<session_id>` — ID of the backup mount session.

For example:

```
user@srv01:~$ veeamconfig session log --id 2a313184-32d0-4d3a-a1b0-2eebac986047
2016-11-22 17:30:34 UTC {30878c82-27d0-45dc-ab21-6f27d5082fd4} [info] Job start
ed at 2016-11-22 20:30:34
2016-11-22 17:30:34 UTC {714b21d0-0d20-486e-b1e5-22d5fb5a8ee9} [info] Mounting
restore point
2016-11-22 17:30:35 UTC {d331f038-5b7c-4549-85cf-5e1b54dbaf71} [info] Restore p
oint has been mounted
```

To ensure that the backup is successfully mounted, you can browse to the directory that you specified in the `veeamconfig backup mount` command. For example:

```
user@srv01:~$ ls /mnt/backup/
FileLevelBackup_0
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see Viewing Session Status.

Step 5. Save Restored Files

When the backup file content is mounted to the computer file system, you can use Linux command line utilities or preferred file browser to work with restored files and folders. You can browse for files and folders in the mounted backup and copy files and folders that you want to restore to their initial location or to a new location.

In the following example, the restored file `Report1.pdf` is copied from the mounted backup to a new location with the Linux command line utilities:

```
user@srv01:~$ ls Documents/
Reports
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/
Report1.pdf  Report2.pdf
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Repo
rt1.pdf Documents/
user@srv01:~$ ls Documents/
Report1.pdf  Reports
```

Step 6. Stop Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. After you have finished working with restored files and folders, you can stop the backup mount session to unmount the backup.

To stop the backup mount session, use the following command:

```
veeamconfig session stop --id <session_id>
```

where:

<session_id> — ID of the backup mount session that you want to stop.

Veeam Agent will stop the mount session and unmount the backup from the computer file system. For example:

```
user@srv01:~$ veeamconfig session stop --id 2a313184-32d0-4d3a-a1b0-2eebac98604
7
Session has stopped.
user@srv01:~$ ls /mnt
user@srv01:~$
```

Restoring from Restore Point

With Veeam Agent command line interface, you can restore files and folders from the specific restore point. When you restore files and folders from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list --all
```

In the list of backups, Veeam Agent displays information about all Veeam Agent for Linux backups stored in all backup repositories. If Veeam Agent is connected to a Veeam backup server, all Veeam Agent for Linux backups that are kept on Veeam backup repositories managed by this server also appear in this list.

The `--all` parameter is optional. If you do not use this option, Veeam Agent displays information about backups created by the current Veeam Agent computer only.

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as *Imported* in the **Repository** column. For information about the import procedure, see Importing Backups.
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name                   Backup ID                                  Repositor
y    Created at
srv01 SystemBackup         {45f074d2-d2d9-423d-84e9-8f1798b08d4c}  Repository_
1  2016-11-11 17:37
srv01 DocumentsBackup      {ea64a7e5-038a-4c86-970a-6d59d4cf3968}  Repository_
1  2016-11-11 18:30
srv01 HomePartitionBackup  {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b}  Repository_
2  2016-11-15 11:28
wrk01 SystemBackup         {951ac571-dd29-45ac-8624-79b8ccb45863}  Repository_
2  2016-11-13 15:26
wrk02 SystemBackup         {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167}  Repository_
3  2016-11-13 15:59
```

TIP

If you want to recover data from a backup that is stored in another location, for example, a backup created with another instance of Veeam Agent, you can import such backup into the Veeam Agent database on your computer. To learn more, see Importing Backups.

Step 2. Explore Restore Points

To view information about restore points in the backup, use the following command:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where:

`<backup_id>` — ID of the backup for which you want to view information on restore points.

You can view the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: • Full • Increment.
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: • True • False

For example:

```
user@srv01:~$ veeamconfig backup info --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
Job name               OIB ID                                      Type       Creat
ed at        Is corrupt
srv01 DocumentsBackup  {0f3c9f3e-3985-4dc9-8cd6-979dba810c2f}  Full       2016-
11-11 18:31    false
srv01 DocumentsBackup  {ff0c6969-8b9b-4865-b4f9-d686faf41d50}  Increment  2016-
11-14 13:35  false
srv01 DocumentsBackup  {a9e420df-d749-4b9a-b675-19d8e94c3bf1}  Increment  2016-
11-14 13:43  false
srv01 DocumentsBackup  {51e6056b-c0ae-40bc-bcf4-4ad7339f647a}  Increment  2016-
11-14 15:48  false
srv01 DocumentsBackup  {b127e64e-1f1c-4e0b-bb36-b087761267b3}  Increment  2016-
11-20 19:52  false
```

or

```
user@srv01:~$ veeamconfig point list --backupid ea64a7e5-038a-4c86-970a-6d59d4c
f3968
Job name               OIB ID                                      Type       Creat
ed at        Is corrupt
srv01 DocumentsBackup  {0f3c9f3e-3985-4dc9-8cd6-979dba810c2f}  Full       2016-
11-11 18:31  false
srv01 DocumentsBackup  {ff0c6969-8b9b-4865-b4f9-d686faf41d50}  Increment  2016-
11-14 13:35  false
srv01 DocumentsBackup  {a9e420df-d749-4b9a-b675-19d8e94c3bf1}  Increment  2016-
11-14 13:43  false
srv01 DocumentsBackup  {51e6056b-c0ae-40bc-bcf4-4ad7339f647a}  Increment  2016-
11-14 15:48  false
srv01 DocumentsBackup  {b127e64e-1f1c-4e0b-bb36-b087761267b3}  Increment  2016-
11-20 19:52  false
```

Step 3. Mount Restore Point

To mount a backup for file-level restore, use the following command:

```
veeamconfig point mount --id <point_id> --mountdir <path>
```

where:

- `<point_id>` — ID of the restore point that you want to mount to the computer file system for file-level restore.
- `<path>` — path to the directory to which you want to mount the backup file content.

For example:

```
user@srv01:~$ veeamconfig point mount --id b127e64e-1f1c-4e0b-bb36-b087761267b3
--mountdir /mnt/backup
Restore point is mounted.
Session ID: [{4d69dd85-ac60-4cff-883d-50f25f49a9c8}].
Logs stored in: [/var/log/veeam/Mount/Session_{4d69dd85-ac60-4cff-883d-50f25f49
a9c8}].
```

Step 4. Monitor Mount Process and Result

You can monitor the restore point mount process by viewing the mount session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

`<session_id>` — ID of the restore point mount session.

For example:

```
user@srv01:~$ veeamconfig session log --id 4d69dd85-ac60-4cff-883d-50f25f49a9c8
2016-11-23 12:44:55 UTC {9c5c8ece-cb88-4742-bb90-1f8ff79b4bdc} [info] Job start
ed at 2016-11-23 15:44:55
2016-11-23 12:44:55 UTC {4ac10045-a74b-4a41-9c5e-53521cba1045} [info] Mounting
restore point
2016-11-23 12:44:56 UTC {540a61f7-5d5c-47d5-a2b8-51daa694d5ec} [info] Restore p
oint has been mounted
```

To ensure that the restore point is successfully mounted, you can browse to the directory that you specified in the `veeamconfig point mount` command. For example:

```
user@srv01:~$ ls /mnt/backup/
FileLevelBackup_0
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see Viewing Session Status.

Step 5. Save Restored Files

When the restore point is mounted to the computer file system, you can use Linux command line utilities or preferred file browser to work with restored files and folders. You can browse for files and folders in the mounted backup and copy files and folders that you want to restore to their initial location or to a new location.

In the following example, the restored file `Report1.pdf` is copied from the mounted restore point to a new location with the Linux command line utilities:

```
user@srv01:~$ ls Documents/
Reports
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/
Report1.pdf  Report2.pdf
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Repo
rt1.pdf Documents/
user@srv01:~$ ls Documents/
Report1.pdf  Reports
```

Step 6. Stop Backup Mount Session

When Veeam Agent mounts a restore point for file-level restore, Veeam Agent starts a new restore point mount session. After you have finished working with restored files and folders, you can stop the mount session to unmount the restore point.

To stop the restore point mount session, use the following command:

```
veeamconfig session stop --id <session_id>
```

where:

<session_id> — ID of the restore point mount session that you want to stop.

Veeam Agent will stop the mount session and unmount the restore point from the computer file system. For example:

```
user@srv01:~$ veeamconfig session stop --id 4d69dd85-ac60-4cff-883d-50f25f49a9c
8
Session has stopped.
user@srv01:~$ ls /mnt
user@srv01:~$
```

Exporting Backup to Virtual Disk

You can export a backup to a virtual disk in the VHD format. You can then attach the created VHD disc to a virtual machine to recover your computer in a virtual environment.

- Exporting Backups
- Exporting Restore Points

Exporting Backups

You can export the backup file to a virtual disk in the VHD format. When you export a backup, you export to a virtual disk data pertaining to the latest restore point in the backup. The created VHD disk will reflect the state in which backed-up volumes were at the time when the latest restore point was created.

To export backup to a VHD disk:

1. Start the export process with the following command:

   ```
   veeamconfig backup export --id <backup_id> --outdir <path>
   ```

   where:

   - `<backup_id>` — ID of the backup that you want to export to a virtual disk.
   - `<path>` — full path to a directory in which you want to save the created virtual disk. Specifying relative paths is not supported.

   For example:

   ```
   user@srv01:~$ veeamconfig backup export --id 45f074d2-d2d9-423d-84e9-8f179
   8b08d4c  --outdir /home/user/disk
   Export has been started.
   Session ID: [{5f001367-8937-46e0-a756-449bf9f1a182}].
   Logs stored in: [/var/log/veeam/Export/Session_{5f001367-8937-46e0-a756-44
   9bf9f1a182}].
   ```

2. You can monitor the export process and result by viewing the export session log with the following command:

   ```
   veeamconfig session log --id <session_id>
   ```

   where:

   `<session_id>` — ID of the export session.

   For example:

   ```
   user@srv01:~$ veeamconfig session log --id 5f001367-8937-46e0-a756-449bf9f
   1a182
   2016-11-27 11:20:56 UTC {b54af37c-35a6-4807-80d2-0f070f024e69} [info] Job
   started at 2016-11-27 14:20:56
   2016-11-27 11:20:56 UTC {48d699d2-86cf-4a32-b9c8-ab51b8325f3c} [info] Expo
   rting virtual disks content
   2016-11-27 11:20:57 UTC {0e2e7d97-f067-4823-8dde-084c401eb62b} [processing
   ] Restoring device: [30460cb5].
   2016-11-27 11:22:59 UTC {0e2e7d97-f067-4823-8dde-084c401eb62b} [info] Devi
   ce [30460cb5] has been exported
   2016-11-27 11:23:00 UTC {36f0d0c5-2af7-48d8-abc2-c8ef9aaffc54} [info] Virt
   ual disks content has been exported
   ```

   You can also check the restore session status with the `veeamconfig session info` command. To learn more, see Viewing Session Status.

3. Exported backup will be saved as a virtual disk file in the specified directory. You can check this with a file browser or with the following command:

```
ls <path>
```

where:

<path> — path to the directory in which the virtual disk with the backup is saved.

For example:

```
user@srv01:~$ ls disk/
dev_30460cb5.vhd
```

Exporting Restore Points

You can export the specific restore point to a virtual disk in VHD format. When you export a restore point, you select the necessary restore point in the backup to recover data to a desired point in time. The created VHD disk will reflect the state in which backed-up volumes were at the time when the selected restore point was created.

To export restore point to a VHD disk:

1. Start the export process with the following command:

   ```
   veeamconfig point export --id <point_id> --outdir <path>
   ```

   where:

   - `<point_id>` — ID of the restore point that you want to export to a virtual disk.
   - `<path>` — full path to a directory in which you want to save the created virtual disk. Specifying relative paths is not supported.

   For example:

   ```
   user@srv01:~$ veeamconfig point export --id b319ea1f-59a2-41ea-9ca3-b668e8
   6ac941 --outdir /home/user/veeam/
   Export has been started.
   Session ID: [{aeb9c549-a660-4a0e-b89c-cb076b8bfa85}].
   Logs stored in: [/var/log/veeam/Export/Session_{aeb9c549-a660-4a0e-b89c-cb
   076b8bfa85}].
   ```

2. You can monitor the export process and result by viewing the export session log with the following command:

   ```
   veeamconfig session log --id <session_id>
   ```

   where:

   `<session_id>` — ID of the export session.

   For example:

   ```
   user@srv01:~$ veeamconfig session log --id aeb9c549-a660-4a0e-b89c-cb076b8
   bfa85
   2016-05-05 11:15:21 UTC {b950503d-55c9-435f-946e-1078184f5a86} [info] Job
   started at 2016-05-05 14:15:21
   2016-05-05 11:15:21 UTC {32d56391-9002-431e-ae6b-2285537a67e5} [info] Expo
   rting virtual disks content
   2016-05-05 11:15:22 UTC {ba3dabe0-0556-430c-9671-9448a6dc4bcb} [processing
   ] Restoring device: [30460cb5].
   2016-05-05 11:17:26 UTC {ba3dabe0-0556-430c-9671-9448a6dc4bcb} [info] Devi
   ce [30460cb5] has been exported
   2016-05-05 11:17:26 UTC {9e945c29-900e-4a07-9e3b-ccf7f156807d} [info] Virt
   ual disks content has been exported
   ```

   You can also check the restore session status with the `veeamconfig session info` command. To learn more, see Viewing Session Status.

3. Exported backup will be saved as a virtual disk file in the specified directory. You can check this with a file browser or with the following command:

```
ls <path>
```

where

<path> — path to the directory in which the virtual disk with the backup is saved.

For example:

```
user@srv01:~$ ls /home/user/veeam/
dev_30460cb5.vhd
```

Restoring Data from Encrypted Backups

When you restore data from an encrypted backup, Veeam Agent performs data decryption automatically in the background or requires you to specify a password.

- If encryption keys required to unlock the backup file are available in the Veeam Agent database, that is, if you encrypt and decrypt the backup file on the same Veeam Agent computer, you do not need to specify the password. Veeam Agent uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore from the encrypted backup does not differ from that from an unencrypted one.
- If encryption keys are not available in the Veeam Agent database, you need to provide a password to unlock the encrypted file. The password must be the same as the password that was used to encrypt the backup file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including restore points that were encrypted with an old password and restore points that were created before you have enabled the encryption option for the job.

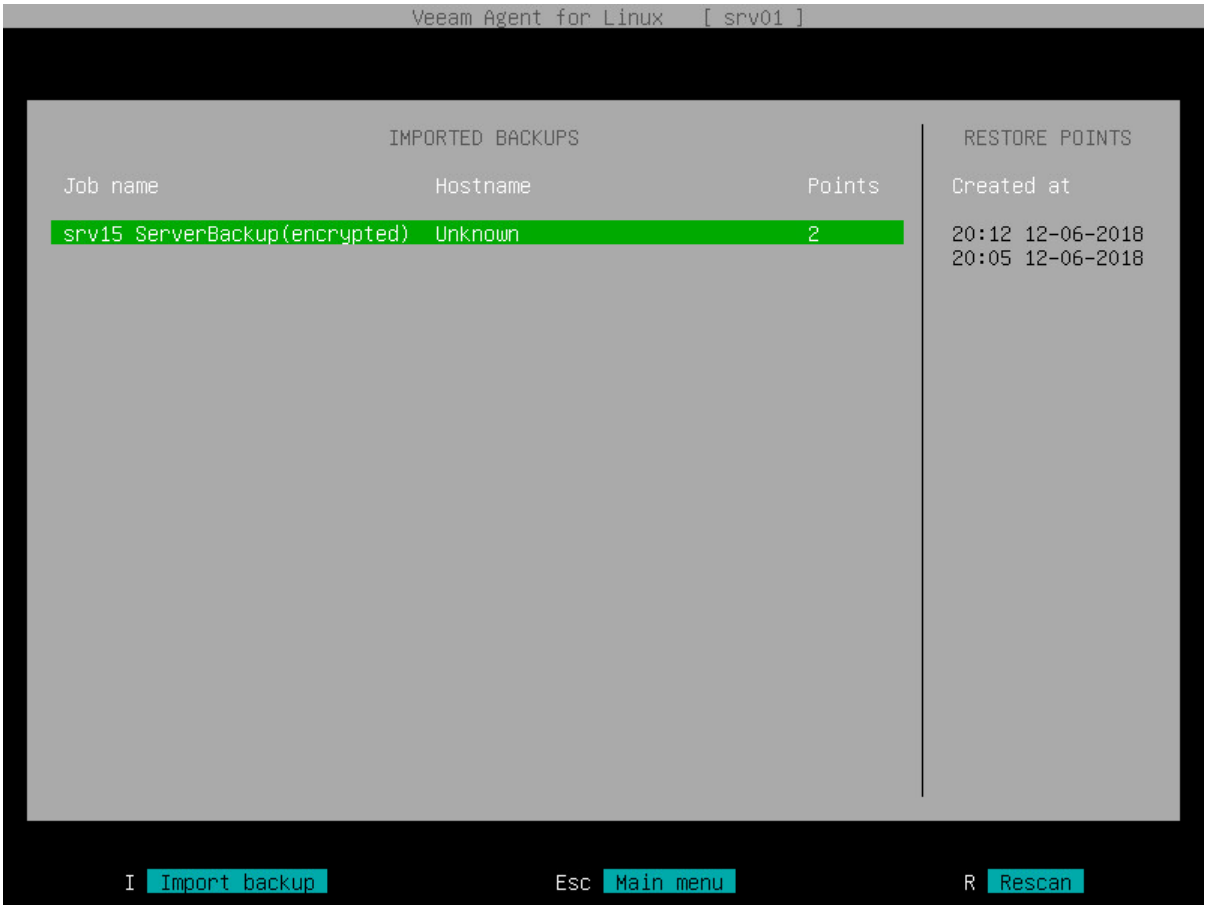
  The process of unlocking an encrypted backup file differs depending on what Veeam Agent user interface you use for data restore.

  - Veeam Agent graphical user interface
  - Veeam Agent command line interface

Restoring Data from Encrypted Backups Using GUI

To restore data from an encrypted backup using the Veeam Agent graphical user interface:

1. Launch the necessary data restore wizard:
   - If you want to perform file-level restore from an encrypted backup that was created on another Veeam Agent computer, launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command. To learn more, see Restoring Files and Folders.
   - If you want to perform volume-level restore or file-level restore recovery from an encrypted backup, boot from the Veeam Recovery Media and launch the necessary data restore wizard. To learn more, see Restoring from Veeam Recovery Media.
2. Follow the steps of the wizard to specify where the encrypted backup file that you plan to use for restore resides. If the backup file resides in a remote location, select the backup location type and specify settings to connect to the backup location.
3. Select the encrypted backup and restore point from which you want to restore data.

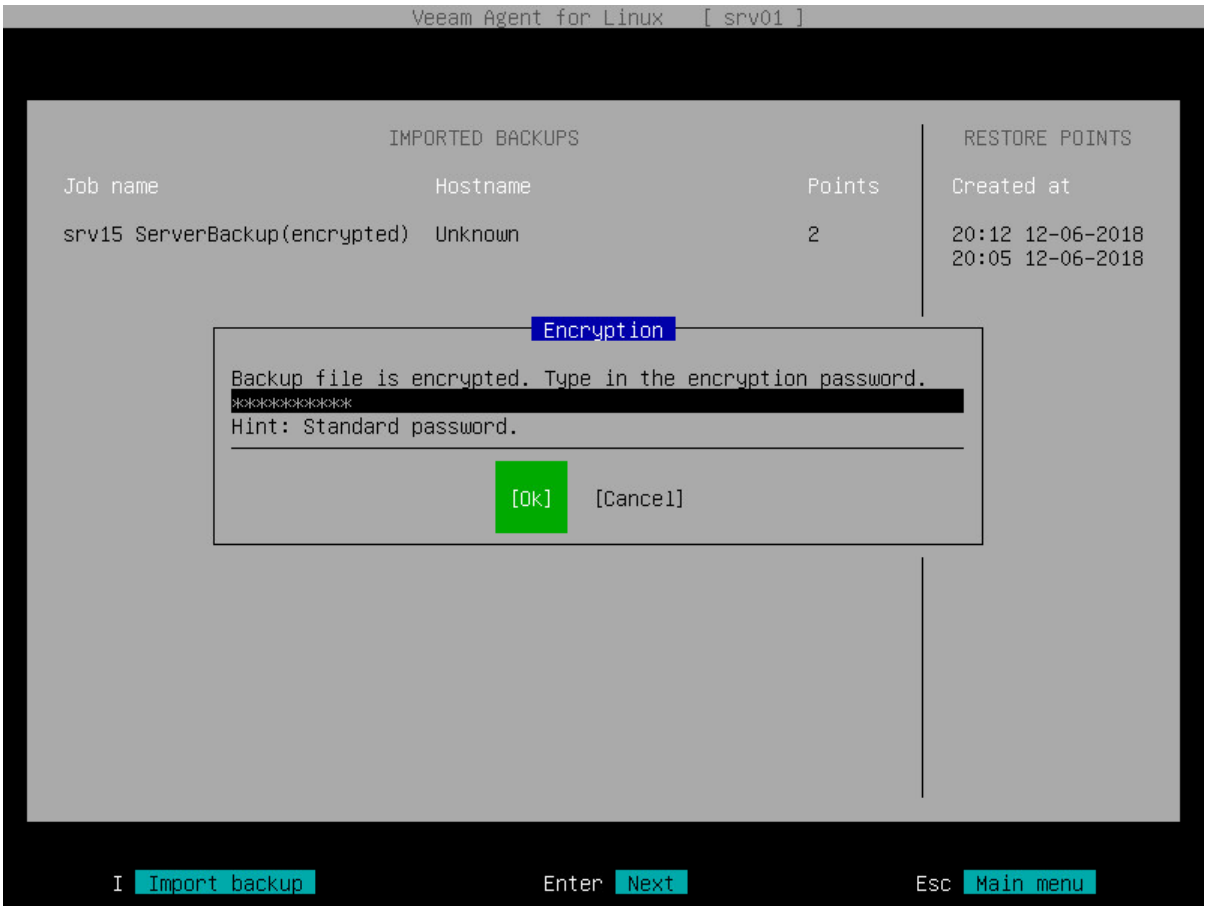


4. Veeam Agent will display the **Encryption** window. Enter the password for the backup file.

   In the **Hint** field of the **Encryption** window, Veeam Agent displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.

   If you changed the password one or several times while the backup chain was created, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

   If you enter correct password, Veeam Agent will decrypt the backup metadata. You will be able to continue the restore operation in a regular manner.



Restoring Data from Encrypted Backups Using Command Line Interface

To restore data from an encrypted backup using the Veeam Agent command line interface, complete the following steps:

1. Import the encrypted backup file to the Veeam Agent database. To learn more, see Importing Encrypted Backups.
2. Perform the necessary restore operation in a regular manner. To learn more, see Restoring Volumes with Command Line Interface and Restoring Files and Folders with Command Line Interface.

Reporting

Veeam Agent for Linux provides several ways to get information about performed operations:

- With the Veeam Agent control panel
- With the Veeam Agent command line interface

For every data transfer operation, for example data backup and restore, backup import and export, Veeam Agent starts a new session. You can monitor performance of sessions started by Veeam Agent in the following ways:

- Monitor backup job session progress with the control panel.
- View real-time backup job session statistics with the control panel.
- View backup job sessions results with the control panel.
- View the session status using the command line interface.
- View session logs.

Viewing Job Session Progress

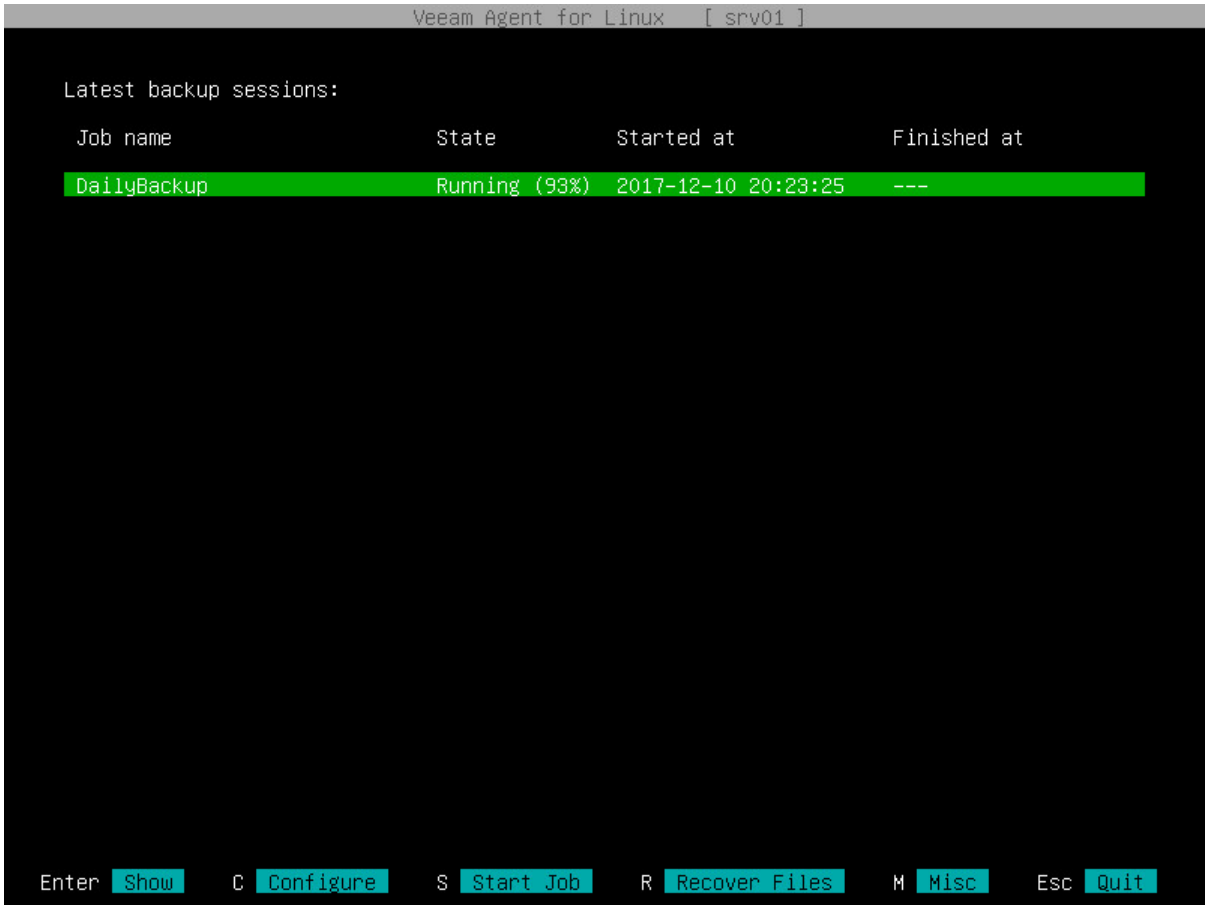
You can monitor the backup job session progress in the list of sessions in the Veeam Agent control panel. For the currently running backup job session, Veeam Agent shows session status and percentage of session completion in the **State** column of the list of sessions.

To view backup job session progress, do the following:

1. If you have started the backup job from the command line, launch the Veeam Agent control panel with the `veeam` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, monitor progress of the currently running session.

   If you have started the backup job from the Veeam Agent control panel, Veeam Agent will immediately display the list of backup job sessions with the currently running session.

> **TIP**
>
> You can stop the backup job session at any time. To stop the backup job session, press the '**s**' key.



Viewing Real-Time Job Session Statistics

You can view real-time statistics for a job session in the Veeam Agent control panel. Veeam Agent shows detailed data for every backup job session: 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.

To view detailed information on the currently running backup job session, do the following:

1. If you have started the backup job from the command line, launch the Veeam Agent control panel with the `veeam` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, select the currently running session with **Up** and **Down** keys and press **Enter**.

   If you have started the backup job from the Veeam Agent control panel, the current session will be already selected in the list of backup job sessions.

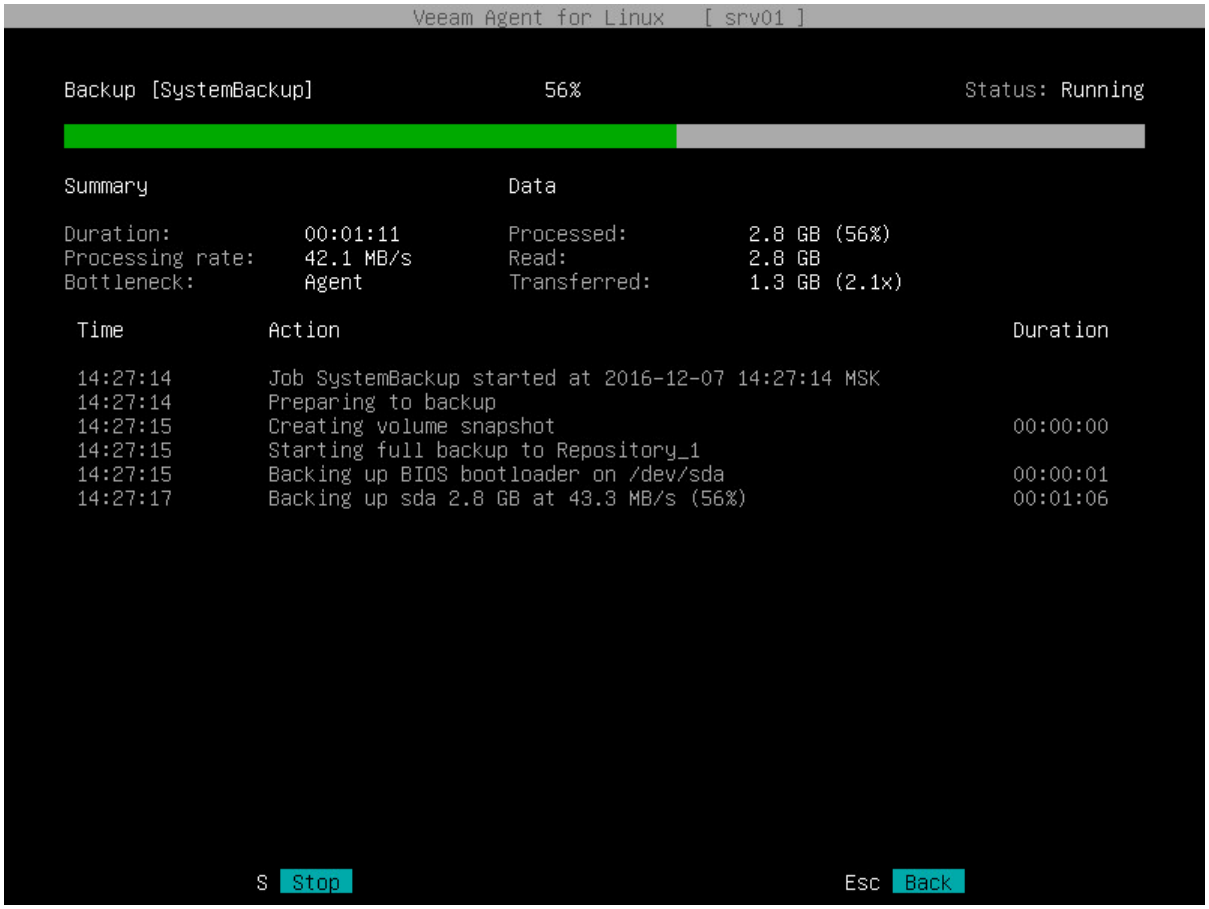
> **TIP**
>
> You can stop the backup job session at any time. To stop the backup job session, press the '**s**' keyboard key.

Statistics Counters

Veeam Agent for Linux displays jobs statistics for the following counters:

- The pane at the top of the control panel shows information on the job session type, percentage of the job completion and session status. If Veeam Agent operates in the Server edition and you have created more than one backup job, the job name also appears on the pane.
- The **Summary** box shows general information about the job:
  - **Duration** — time from the job start till the job end.
  - **Processing rate** — average speed of data processing. This counter is a ratio between the amount of processed data (**Processed** counter) and job duration (**Duration** counter).
  - **Bottleneck** — bottleneck in the data transmission process.
- The **Data** box shows information about processed data:
  - **Processed** — total size of all volumes processed by the job.
  - **Read** — amount of data read from the backed-up volume by Veeam Agent for Linux prior to applying compression. For incremental job runs, the value of this counter is typically lower than the value of the **Processed** counter. Veeam Agent reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target location.
  - **Transferred** — amount of data transferred from the backed-up volume to the backup location after applying compression. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Agent can perform additional activities with data, for example, decompress data prior to writing the file to disk. The activities can impact the size of the resulting file.
- The box in the center of the control panel shows a list of operations performed during the job session, their start time and duration time. To scroll the list of operations, use **Up** and **Down** arrow keys on the keyboard.

- The pane at the lower side of the control panel shows help information on how to navigate the control panel.



Viewing Job Session Result

You can view detailed statistics on every backup job session performed by Veeam Agent for Linux.

To view statistics for a specific job session:

1. Open the Veeam Agent control panel with one of the following commands:

   ```
   veeam
   ```

   or

   ```
   veeamconfig ui
   ```

   or

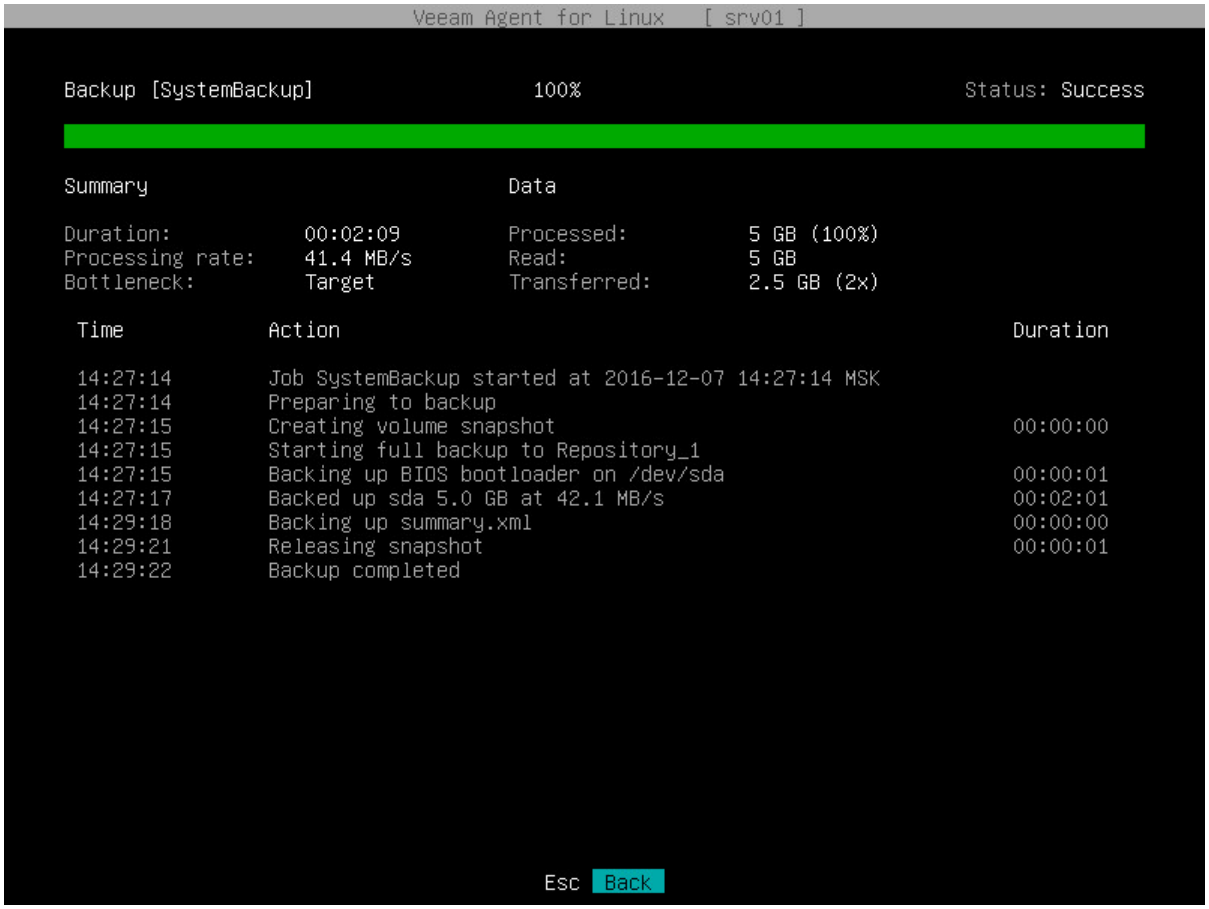
   ```
   veeamconfig session ui
   ```

2. In the **Latest backup sessions** list, select the necessary backup job session with **Up** and **Down** keys and press **Enter**.

TIP

To return to the list of backup job sessions, press **Esc**. You can then select another backup job session or exit the Veeam Agent control panel in one of the following ways:

- with the **Esc** key — if you opened the control panel with the `veeam` or `veeamconfig ui` command.
- with the '**q**' key — if you opened the control panel with the `veeamconfig session ui` command.



Viewing Session Status

You can view status of every session that was started by Veeam Agent for Linux. To view the session status, use the following command:

```
veeamconfig session info --id <session_id>
```

where:

<session_id> — ID of the session for which you want to check status.

Veeam Agent displays the following information about sessions:

Parameter	Description
ID	ID of the session.
Job name	Name of the backup job parent to the session. Veeam Agent displays value for this parameter only for backup job sessions.
Job ID	ID of the backup job parent to the session. Veeam Agent displays value for this parameter only for backup job sessions.
State	Current status of the session.
Start time	Date and time of the session start.
End time	Date and time of the session completion. Veeam Agent displays value for this parameter only for completed sessions.

The following example shows status information on the completed backup job session:

```
user@srv01:~$ veeamconfig session info --id 1592755d-3a2b-40a9-a036-5c81853b369
e
Backup session
  ID: {1592755d-3a2b-40a9-a036-5c81853b369e}
  Job name: SystemBackup
  Job ID: {2495911e-58db-4452-b4d1-f53dcfbc600e}
  State: Success
  Start time: 2016-11-11 14:37:21 UTC
  End time: 2016-11-11 14:40:02 UTC
```

The following example shows status information on the running volume restore session:

```
user@srv01:~$ veeamconfig session info --id 697d9348-9001-4845-8764-3cc4fb3f296
b
Restore session
  ID: {697d9348-9001-4845-8764-3cc4fb3f296b}
  State: Running
  Start time: 2016-11-27 10:35:47 UTC
  End time:
```

Viewing Session Logs

You can monitor the backup and restore process by viewing the backup job session and restore session logs in the Veeam Agent command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> — ID of the backup job or restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 0b72ef45-4c88-4639-b940-ad3828b1cd4e
2016-11-27 11:04:04 UTC {b141f32a-3e77-45a6-b55a-c100a1464d67} [info] Job start
ed at 2016-11-27 14:04:04
2016-11-27 11:04:04 UTC {9b60ac03-2de0-4fe2-a00e-bec556d98ee8} [info] Starting
volume restore
2016-11-27 11:04:07 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [processing] sdb
2016-11-27 11:04:15 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 512.0
kB at 58.6kB/s (0%)
2016-11-27 11:04:25 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 125.0
MB at 6.7MB/s (0%)
2016-11-27 11:04:35 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 238.5
MB at 8.3MB/s (1%)
...
2016-11-27 11:14:32 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB
at 10.7MB/s (92%)
2016-11-27 11:14:35 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB
at 10.6MB/s (97%)
2016-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB
at 10.6MB/s (100%)
2016-11-27 11:14:37 UTC {00add723-cbfa-4cc8-b299-d2349a051d6f} [warn] /dev/sdb
has a duplicate filesystem UUID
2016-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb resto
red 6.5GB at 10.6MB/s
2016-11-27 11:14:37 UTC {8b8742a2-1c80-4e14-bbf1-45a3612bc3a7} [info] Volume re
store completed
```

Managing Configuration Database

You can perform the following operations with the Veeam Agent configuration database:

- Export configuration database to a configuration file.
- Import configuration database to Veeam Agent.

Exporting Configuration Database

You can export the Veeam Agent configuration database to a configuration file in the XML format. This may be useful, for example, if you want to change Veeam Agent settings by editing a configuration file or copy the Veeam Agent configuration to another computer.

To export the Veeam Agent configuration database, use the following command:

```
veeamconfig config export --file <path>
```

where:

<path> — path to a configuration file to which you want to export the configuration database.

For example:

```
user@srv01:~$ veeamconfig config export --file veeam/config.xml
```

NOTE

A directory in which you want to save the configuration file must exist in the file system.

Importing Configuration Database

You import the Veeam Agent configuration from a file in the XML format to the configuration database. This may be useful, for example, if you have changed Veeam Agent for Linux settings by editing a configuration file or want to apply configuration of another instance of Veeam Agent to Veeam Agent installed on your computer.

> **NOTE**
>
> Veeam Agent for Linux 5.0 does not support import of XML configuration files generated by earlier versions of Veeam Agent.

To import the Veeam Agent configuration database, use the following command:

```
veeamconfig config import --file <path>
```

where:

<path> — path to a configuration file from which you want to import the configuration database.

For example:

```
user@srv01:~$ veeamconfig config import --file veeam/config.xml
```

Exporting Product Logs

Veeam Agent offers a simple and convenient way to collect product logs and export them to an archive file. This operation may be required if you want to report an issue and need to attach log files to the support case.

When you export logs, Veeam Agent collects its log files and configuration files, exports them to an archive file in the `tar.gz` format and saves this archive file to a directory on the Veeam Agent computer.

You can perform the export logs operation in one of the following ways:

- With the Veeam Agent control panel — in this case, you can specify a directory to which Veeam Agent should save the log archive.
- With the command line interface — in this case, Veeam Agent will save the log archive to the current working directory.

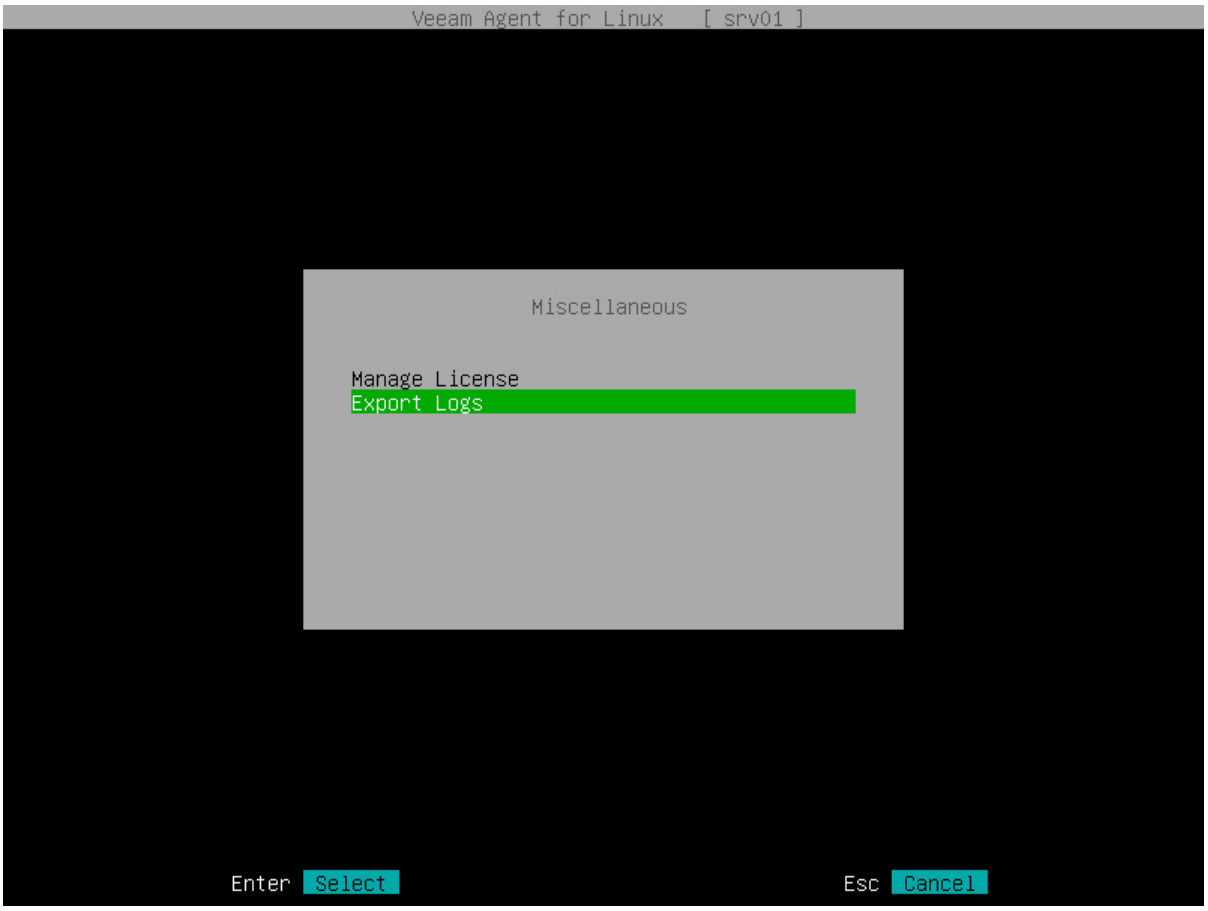
> **TIP**
>
> When you perform restore operations after booting from the Veeam Recovery Media, Veeam Agent also saves restore logs to the backup location. Restore logs are saved to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`. The archive is placed to the folder that contains the backup file from which you restored data.
>
> If you encounter problems after restoring from the Veeam Recovery Media, it is recommended that you attach restore logs, as well as product logs collected by Veeam Agent, to the support case.

Exporting Logs with Control Panel

You can use the Veeam Agent control panel to collect and export product logs. When you export logs with the control panel, you can choose where Veeam Agent should save the resulting log archive.

To export logs:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the '**m**' key to open the **Miscellaneous** menu.
3. In the menu, select the **Export Logs** option and press **Enter**.



4. In the **Choose logs directory** window, specify a directory to which you want to save the log archive:
   a. In the **Choose logs directory** window, select the necessary directory and press **Enter**.
   b. Repeat the step 'a' until a path to the directory in which you want to save exported logs appears in the **Current directory** field.
   c. To create a new directory, switch to the **Create Dir** button, press **Enter**, then type a name for the new directory and press **Enter**.

d. Switch to the **Ok** button and press **Enter**. Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`, and save the archive to the specified directory.



Exporting Logs with Command Line Interface

You can use the Veeam Agent command line interface to collect and export product logs. To export logs, use the following command:

```
veeamconfig grablogs
```

Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`, and save the archive to the current working directory.

For example:

```
user@srv01:~$ veeamconfig grablogs
Logs have been exported successfully.
```

Getting Support

If you have any questions or want to share your feedback about Veeam Agent, you can use one of the following options:

- You can search for the information on the necessary subject in the current Veeam Agent for Linux User Guide.
- You can visit Veeam R&D Forums and share your opinion or ask a question.
- If you use Veeam Agent with an active license installed, you can visit Veeam Customer Support Portal and submit a support case to the Veeam Customer Support Team.

Using with Veeam Backup & Replication

If you have the Veeam backup infrastructure deployed in the production environment, you can use Veeam Agent for Linux together with Veeam Backup & Replication.

IMPORTANT

If you plan to use Veeam Agent for Linux 5.0 with Veeam Backup & Replication, you must install Veeam Backup & Replication 11a on the Veeam backup server. Veeam Agent for Linux 4.0 supports integration with Veeam Backup & Replication 11 or later as well.

NOTE

The subsequent sections describe tasks available for Veeam Agent operating in the standalone mode. For information about Veeam Agent management in Veeam Backup & Replication, see Veeam Agent Management Guide.

Tasks with Veeam Backup & Replication

Veeam Backup & Replication lets you perform a number of additional disaster recovery tasks and administrative actions with Veeam Agent backups. You can:

Data protection tasks

- Create Veeam Agent backups on backup repositories
- Create Veeam Agent backups on a Veeam Cloud Connect repository
- Copy Veeam Agent backups to secondary backup repositories
- Archive Veeam Agent backups to tape

Restore tasks

- Restore a Veeam Agent backup to a VMware vSphere VM
- Restore a Veeam Agent Backup to a Microsoft Hyper-V VM
- Restore files and folders from Veeam Agent backups
- Restore application items from Veeam Agent backups with Veeam Explorer for Oracle
- Restore disks from Veeam Agent backups
- Export Veeam Agent backups to standalone full backup files
- Restore data from Veeam Agent backups to Microsoft Azure
- Restore data from Veeam Agent backups to Amazon EC2

Administrative tasks

- Import Veeam Agent backups
- Enable and disable Veeam Agent backup jobs
- View Veeam Agent backup job statistics
- Delete Veeam Agent backup jobs
- Remove Veeam Agent backups
- View Veeam Agent backup statistics
- Configure global settings
- Assign roles to users

Setting Up User Permissions on Backup Repositories

To be able to store backups on a backup repository managed by a Veeam backup server, the user must have access permissions on this backup repository.

> **NOTE**
>
> If you plan to create backups on 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.

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 on this backup repository.

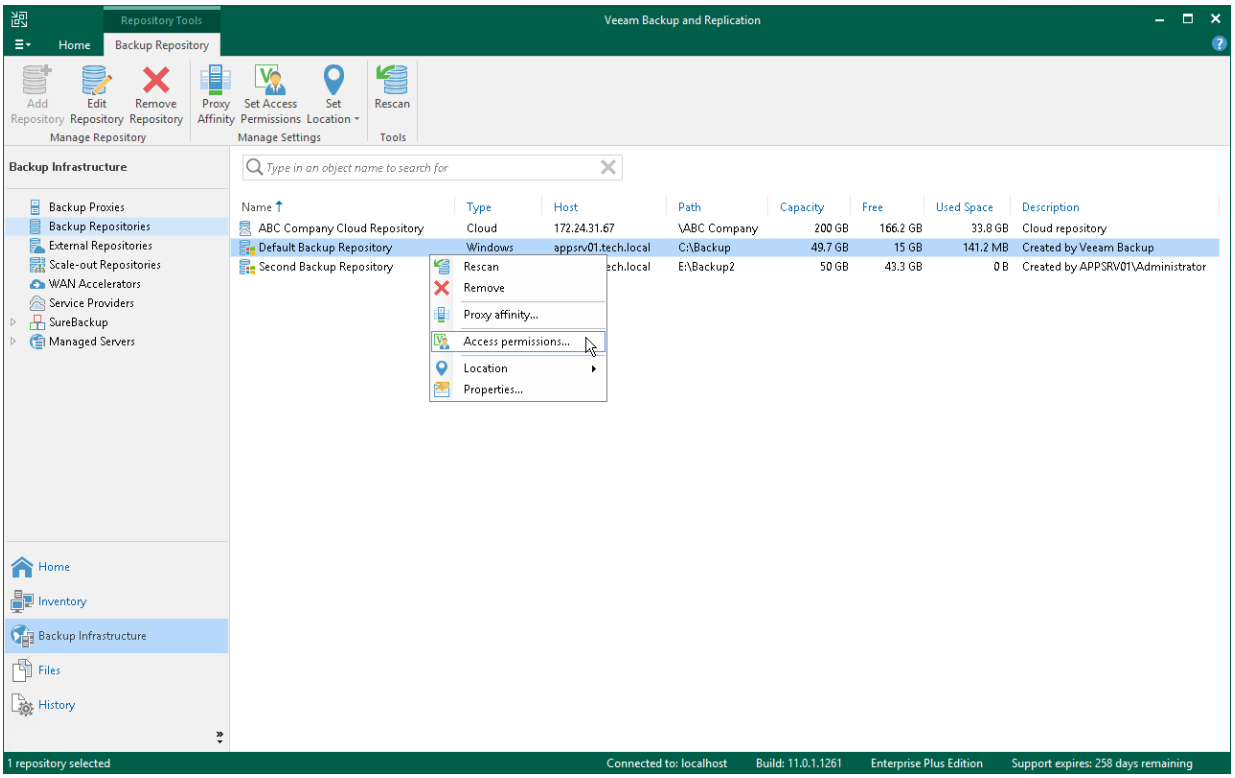
Right after installation, access permissions on the default backup repository are set to *Allow to everyone* for testing and evaluation purposes. If necessary, you can change these settings.

After you create a new backup repository, access permissions on this repository are set to *Deny to everyone*. To allow users to store backups on 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 to Veeam Agent users.
   - The **Scale-out Repositories** node — if you want to grant access permissions on a scale-out backup repository to Veeam Agent users.

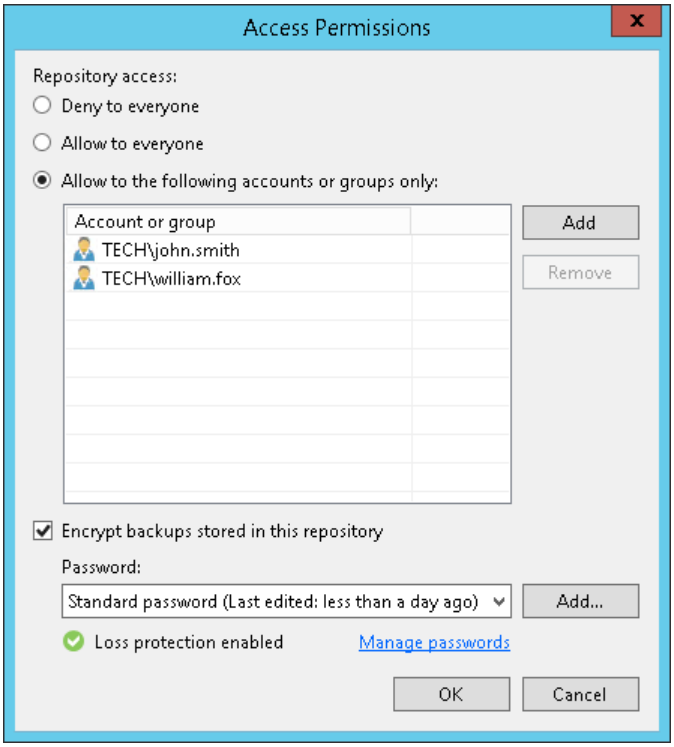
3. In the working area, select the necessary backup repository and click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**. If you do not see the **Set Access Permissions** button on the ribbon or the **Access permissions** command is not 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 on 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 on 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 on 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 on the backup repository side using its built-in encryption mechanism. To learn more, see Veeam Backup & Replication Documentation.



Managing License

If you plan to use Veeam Agent for Linux with Veeam Backup & Replication, you must have a license installed in Veeam Backup & Replication. The license must have a total number of instances that is sufficient to protect Linux-based machines (servers and workstations) on which you plan to install Veeam Agent. The number of backup jobs configured in Veeam Agent does not consume the license. To learn more about per-instance licensing in Veeam Backup & Replication, see the Licensed Objects section in the Veeam Backup & Replication User Guide.

After Veeam Agent for Linux connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming the license. By default, Veeam Agent operates in the Server edition. You can switch Veeam Agent to the Workstation edition manually if needed.

If one or more Veeam Agents operating in the Free edition are already connected to the backup server, they will start consuming the license immediately after the license is installed in Veeam Backup & Replication. Veeam Agents that exceed the license limit will not be able to back up data to the Veeam backup repository.

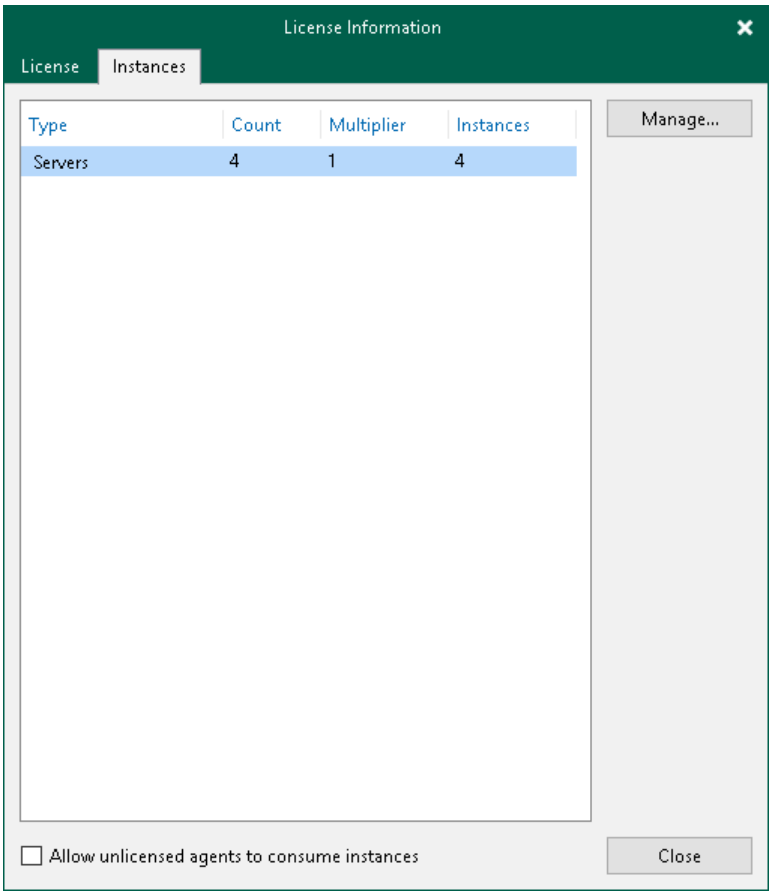
Veeam Agent keeps information about the license in its database. Information about the license is valid for 32 days. If Veeam Agent for Linux does not connect to Veeam Backup & Replication during this period, Veeam Backup & Replication will revoke its license.

Managing Instance Consumption by Veeam Agents

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, you can restrict instance consumption. After instance consumption is restricted, Veeam Agents will not be able to back up data to a Veeam backup repository.

To restrict instance consumption by Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, click the **Instances** tab.
3. On the **Instances** tab, clear the **Allow unlicensed agents to consume instances** check box.
4. Click **Close**.



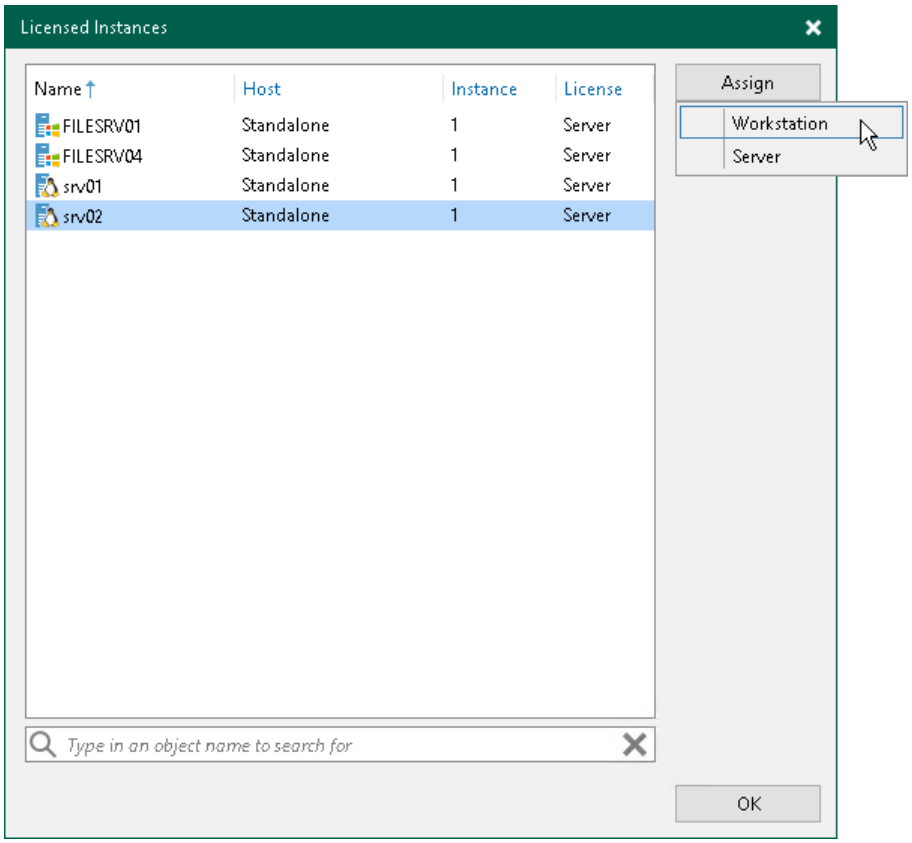
Assigning License to Veeam Agent

After Veeam Agent connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming the license. If the license allows Veeam Agent to operate in both the Workstation Server editions, Veeam Agent will operate in the Server edition.

You can also assign a license to Veeam Agent manually if needed. When you assign a license, you can select in which edition the licensed Veeam Agent will operate.

To assign a license:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the **Licensed Instances** window, select the Veeam Agent to which you want to assign the license, click **Assign** and select the desired edition: *Workstation* or *Server*.



Viewing Licensed Agents and Revoking License

When Veeam Backup & Replication connects to the backup server, Veeam Backup & Replication applies a license to the Veeam Agent. You can view to which Veeam Agents the license is currently applied.

To view a list of licensed Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.

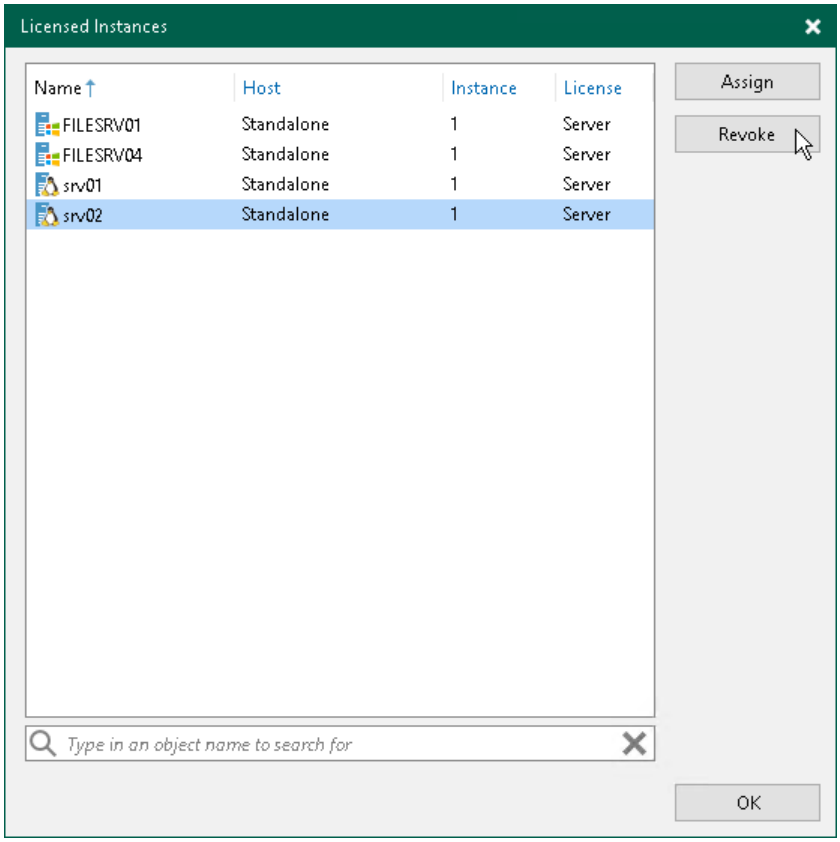
In the list of licensed instances, Veeam Backup & Replication displays Veeam Agents that have established a connection with the backup server during a Veeam Agent backup job session.

Revoking License from Veeam Agents

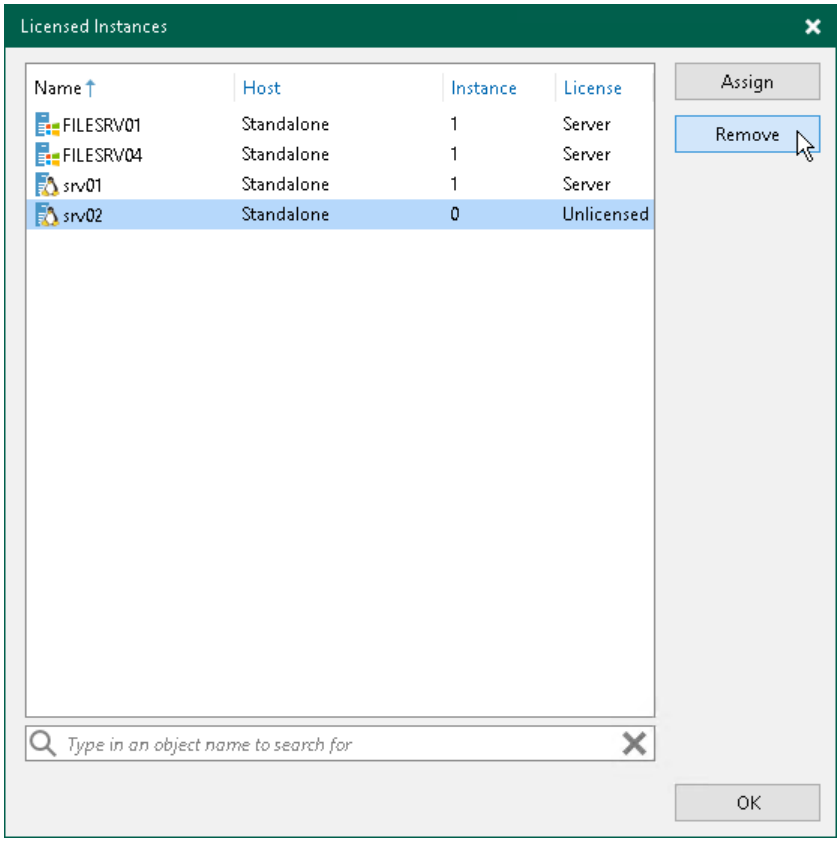
You can revoke the license from some Veeam Agents and re-apply it to other protected workloads. License revoking can be helpful, for example, if you do not want to use some Veeam Agents with Veeam Backup & Replication anymore.

To revoke a license from the Veeam Agent:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the **Licensed Instances** window, select the Veeam Agent from which you want to revoke the license and click **Revoke**. Veeam Backup & Replication will revoke the license from the Veeam Agent, and the license will be freed for other workloads protected with Veeam Backup & Replication.



The Veeam Agent from which you have revoked the license will become unable to connect to the Veeam backup server but will remain in the **Licensed Instances** list. To allow this Veeam Agent to create backups in the Veeam backup repository, select the Veeam Agent and click **Remove**. During the next backup job session, the Veeam Agent will connect to the Veeam backup server and start consuming the license.



Performing Data Protection Tasks

You can perform the following data protection tasks:

- Back up your data and store the resulting backup files on a backup repository managed by a Veeam backup server.
- Back up your data and store the resulting backup files on a Veeam Cloud Connect Repository.
- Copy Veeam Agent for Linux backups from the backup repository to a secondary backup repository with backup copy jobs.
- Archive Veeam Agent for Linux backups to tapes with backup to tape jobs.

Backing Up to Backup Repositories

You can store backups created with Veeam Agent on backup repositories connected to Veeam backup servers. To do this, you must perform the following actions:

1. Set up user permissions at the backup repository side.
2. Point the Veeam Agent for Linux backup job to the backup repository.

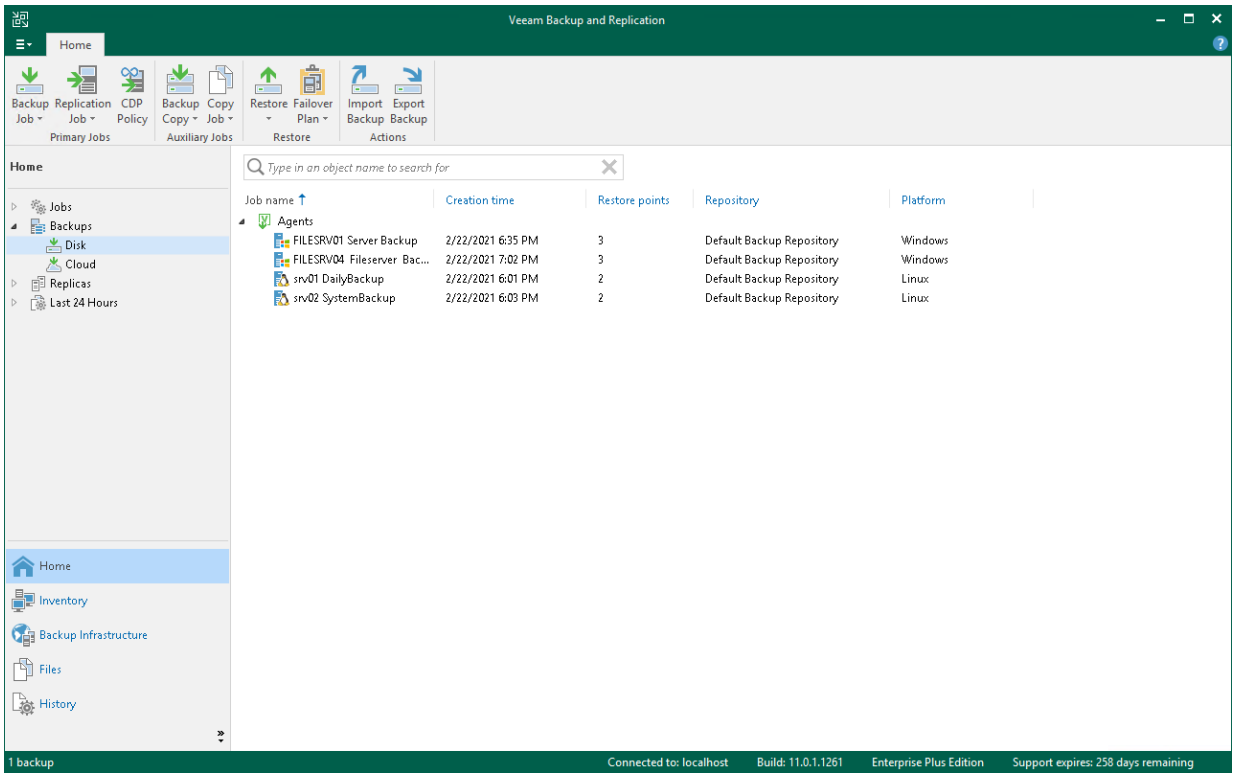
NOTE

Consider the following:

- A Veeam Agent backup job can be started automatically upon the defined schedule or manually from the Veeam Agent computer. You cannot start, stop, retry or edit Veeam Agent backup jobs in the Veeam Backup & Replication console.
- If the user is granted restore permissions on the Veeam backup server, the user will be able to see all backups on the backup repository.
- The user who creates a Veeam Agent backup on the backup repository is set as the owner of the backup file. The backup file owner can access this file and restore data from it. If the user who is not the backup file owner needs to perform operations with the backup file, the user must have the Veeam Backup & Replication role that allows to perform these operations. To learn more about roles, see the Roles and Users section in the Veeam Backup & Replication User Guide.

Backup jobs targeted at the backup repository become visible in Veeam Backup & Replication under the **Jobs** > **Backup** node in the **Home** view. Backups created with Veeam Agent are available under the **Backups** > **Disk** node in the **Home** view.

The Veeam Backup Administrator working with Veeam Backup & Replication can manage Veeam Agent backup jobs and restore data from these backups. To learn more, see Performing Restore Tasks and Performing Administration Tasks.



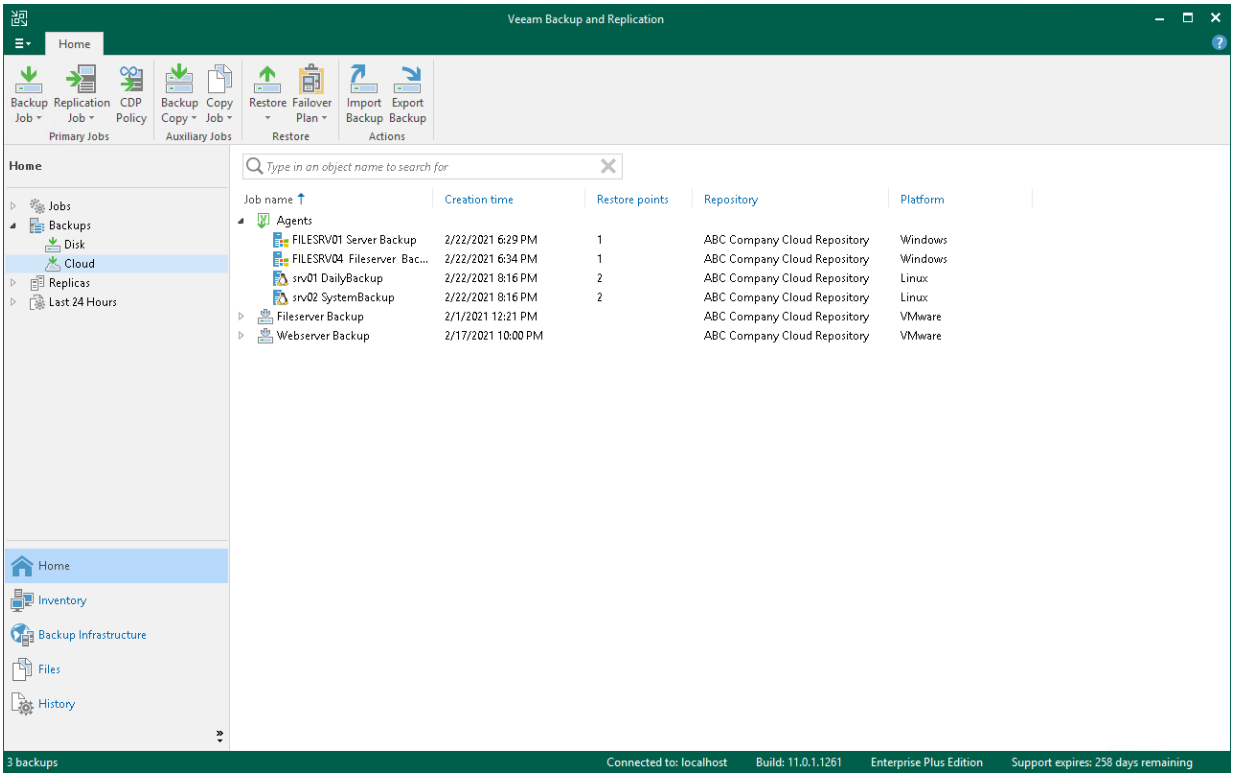
Backing Up to Cloud Repositories

You can use Veeam Agent to create backups on cloud repositories provided to you by a Veeam Cloud Connect service provider. To do this, you must connect to the service provider and point the backup job to the cloud repository. To connect to the service provider, you can use credentials of the tenant or subtenant account that your provider or backup administrator communicated to you. To learn more, see Veeam Cloud Connect Repository Settings.

Veeam Agent Backups on Tenant Side

Backups created with Veeam Agent are available under the **Backups** > **Cloud** node in the **Home** view of the Veeam Backup & Replication console deployed on the tenant side.

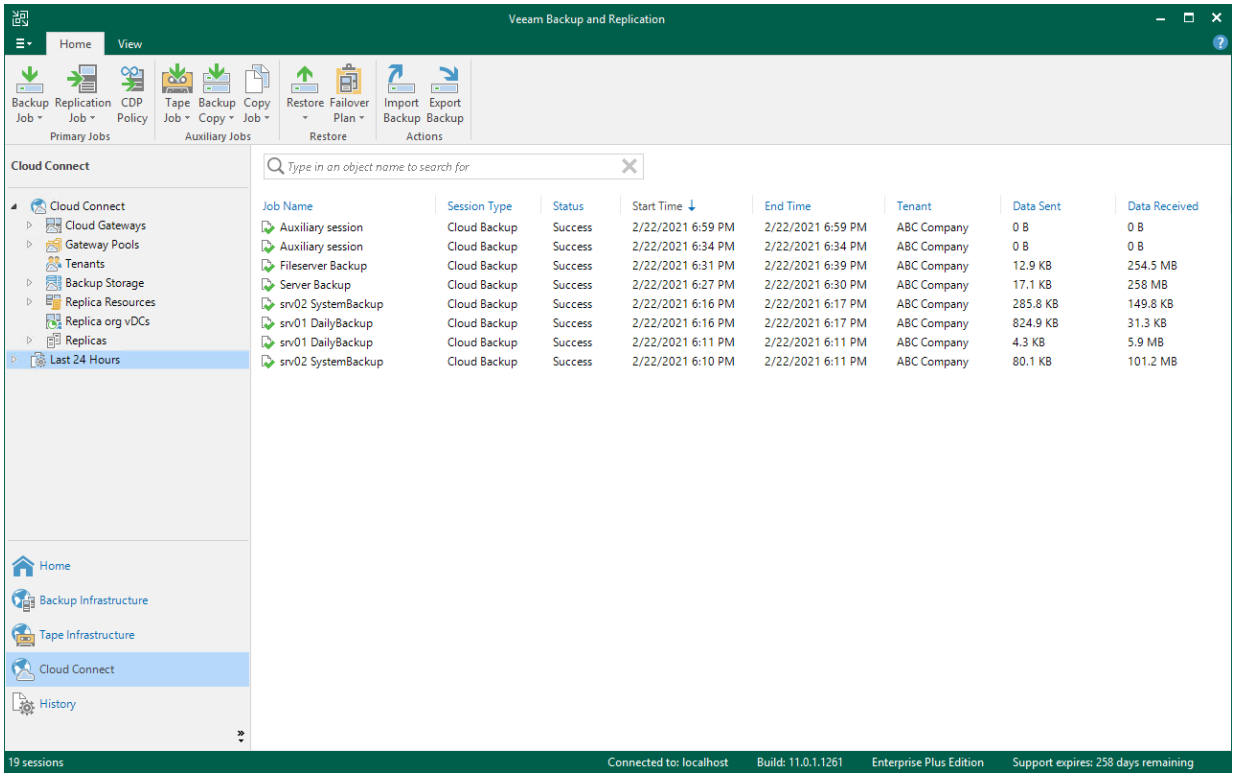
The backup administrator working with Veeam Backup & Replication on the tenant side can manage Veeam Agent backups created on the cloud repository and restore data from such backups. To recover data from a Veeam Agent for Linux backup, you can export disks of the Veeam Agent computer as virtual disks. To learn more, see Exporting Disks.



Veeam Agent Backups on Service Provider Side

The service provider can view information about backup and restore sessions performed by Veeam Agent users within the last 24 hours period. The list of sessions is available under the **Last 24 Hours** node in the **Cloud Connect** view of the Veeam Backup & Replication console deployed on the service provider side.

The service provider cannot perform restore tasks with Veeam Agent backups that are stored on the cloud repository.



Performing Backup Copy for Veeam Agent Backups

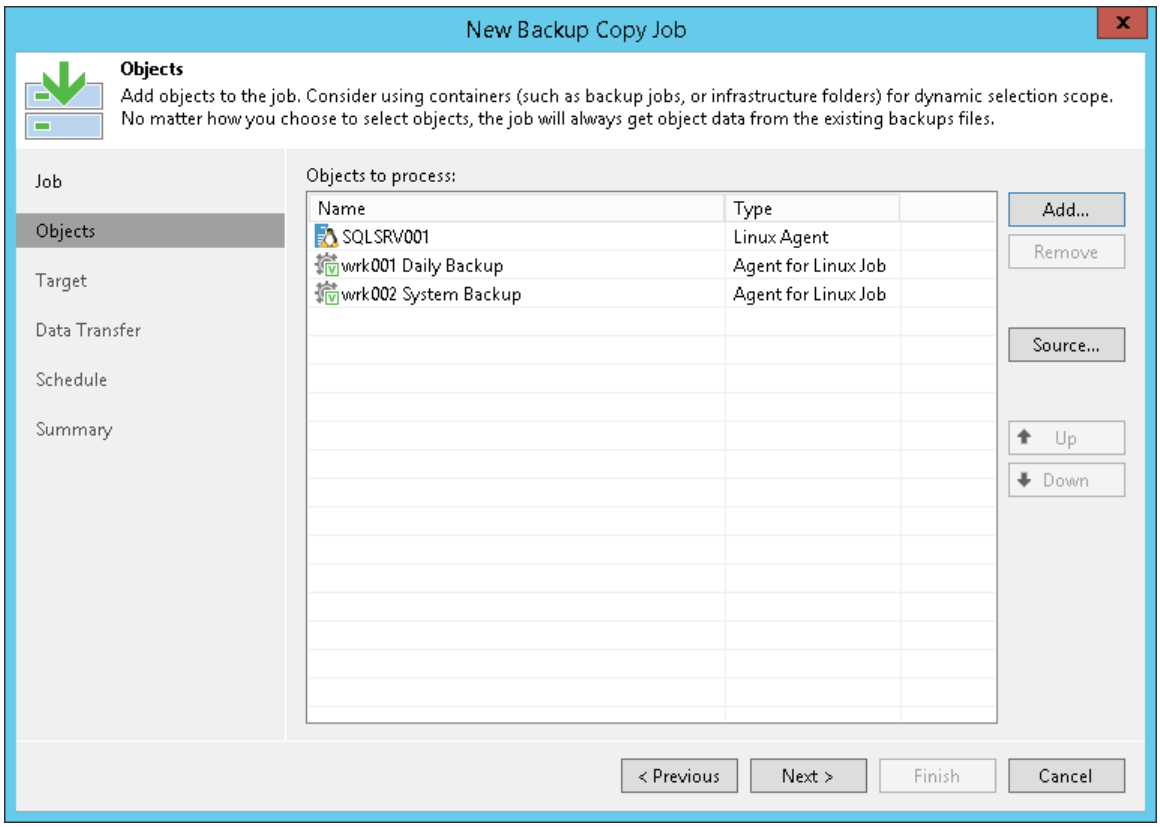
You can configure backup copy jobs that will copy backups created with Veeam Agent to a secondary backup repository.

Backup copy jobs treat Veeam Agent backups as usual backup files. The backup copy job setup and processing procedures practically do not differ from the same procedures for a backup copy job that processes VM backups. To learn more about backup copy jobs, see the Backup Copy section in the Veeam Backup & Replication User Guide.

Mind the following:

- You can process backups created by Veeam Agent only with backup copy jobs for Linux machine backups. You cannot add a Veeam Agent backup as an additional source of a backup copy job that processes VM backups.
- You can process backups created by Veeam Agent operating in the standalone mode only with backup copy jobs operating in the periodic copy mode. Immediate copy mode is not supported.
- You can process backups created by Veeam Agent operating in the managed mode with backup copy jobs operating in the periodic and immediate copy modes.

  Mind that for immediate copy mode, only Veeam Agent backups created by jobs managed by the backup server are supported.
- When mapping a backup copy job to a Veeam Agent backup, consider the limitations listed in the Backup File Mapping section in the Veeam Backup & Replication User Guide.



Restoring Data from Copies of Veeam Agent Backups

Backups copied to the secondary backup repository do not preserve user access permissions. At the same time, users who created backups do not have access permissions on these secondary repositories. For this reason, users cannot restore data from their backups residing in the secondary site.

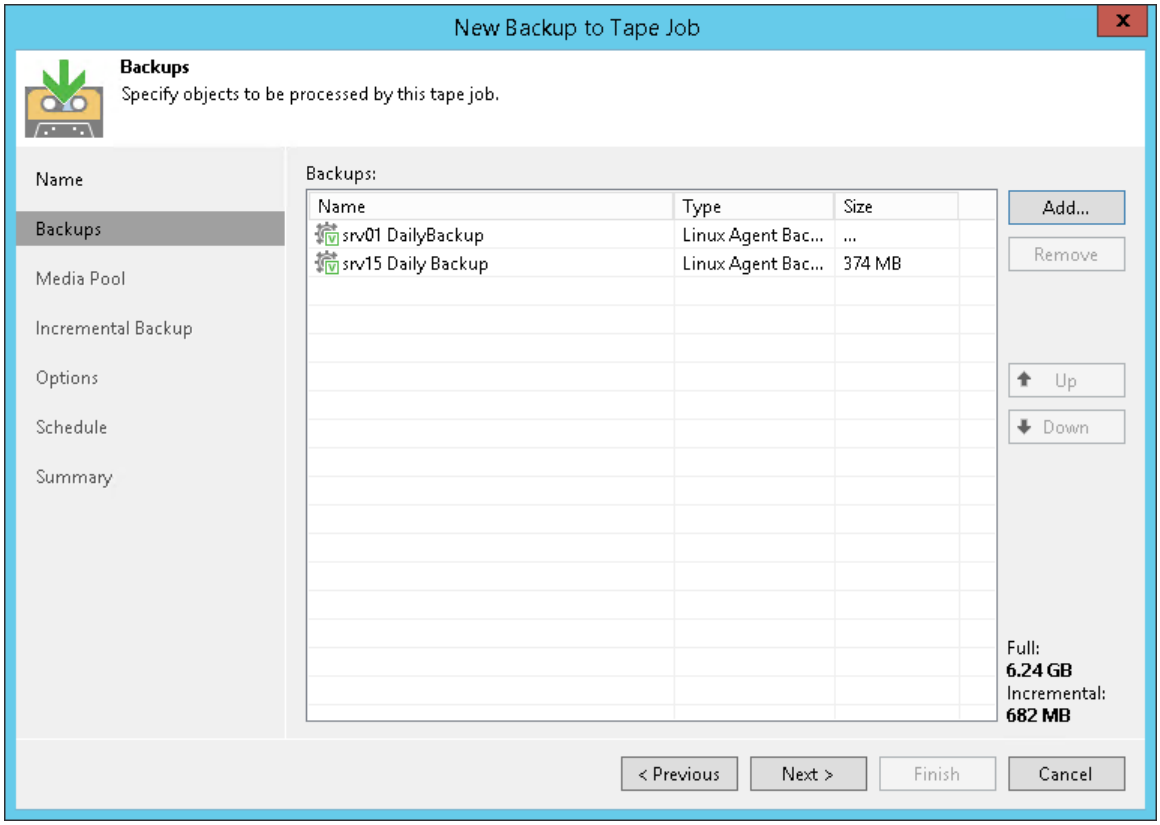
To overcome this limitation, you can delegate the restore task to backup administrators who work with Veeam Backup & Replication. Backup administrators can use Veeam Backup & Replication to export data contained in backup files as virtual disks.

Archiving Veeam Agent Backups to Tape

You can configure backup to tape jobs to archive Veeam Agent backups to tape.

Backup to tape jobs treat Veeam Agent backups as usual backup files. The archiving job setup and processing procedures practically do not differ from the regular ones. To learn more about backup to tape jobs, see Veeam Backup & Replication Documentation.

> **NOTE**
>
> Keep in mind that in backup to tape job schedule, you cannot select the **After this job option** for a Veeam Agent job configured directly on a Veeam Agent computer.



Performing Restore Tasks

You can perform the following restore operations:

- Restore a Veeam Agent backup to a VMware vSphere VM
- Restore a Veeam Agent Backup to a Microsoft Hyper-V VM
- Restore individual files and folders from Veeam Agent for Linux backups.
- Restore application items from Veeam Agent backups with Veeam Explorer for Oracle.
- Export computer disks as VMDK, VHD or VHDX disks.
- Export restore points of Veeam Agent backups to standalone full backup files.
- Restore data from Veeam Agent backups to Microsoft Azure.
- Restore data from Veeam Agent backups to Amazon EC2.
- Restore data from Veeam Agent backups to Google Cloud Platform.

Restoring Veeam Agent Backup to VMware vSphere VM

You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a VMware vSphere VM in your virtualization environment. For Instant Recovery to a VMware vSphere VM, you can use Veeam Agent backups created on the Veeam backup repository. You cannot perform this operation with Veeam Agent backups created on the Veeam Cloud Connect repository.

A restored VMware vSphere VM has the same settings as a backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves settings of the Veeam Agent computer from the backup and applies them to the target VM. These settings include:

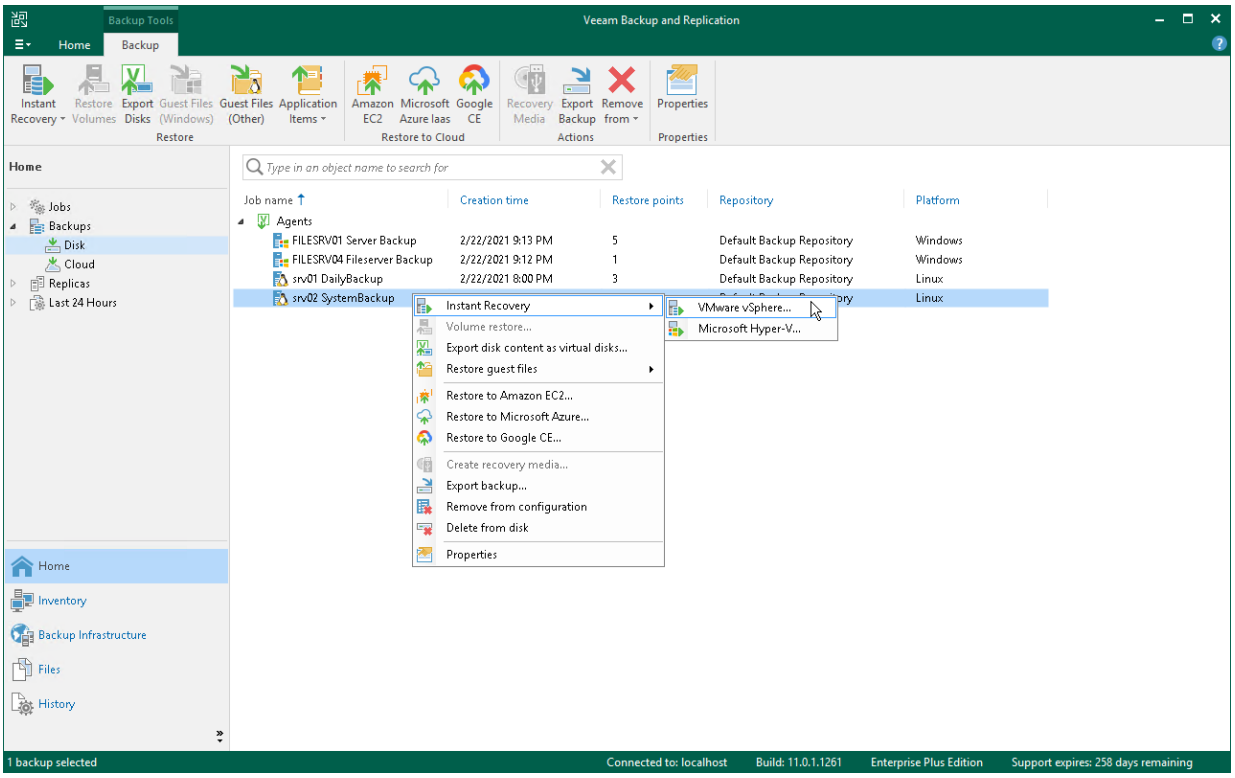
- Amount of RAM
- Number of CPU cores
- Number of network adapters
- Network adapter settings
- BIOS UUID

  If you do not want to preserve the backed-up machine UUID for a VMware vSphere VM, you can create a new UUID during the Instant Recovery configuration process.
- Number of disks and volumes
- Size of volumes

If you restore a Veeam Agent computer to a VMware vSphere VM, consider the following:

- Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.
- If you restore a workload to the production network, make sure that the original workload is powered off.
- If the disk you want to restore contains an LVM volume group, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the Performing Instant Recovery of Workloads to VMware vSphere VMs section in the Veeam Backup & Replication User Guide.



Restoring Veeam Agent Backup to Microsoft Hyper-V VM

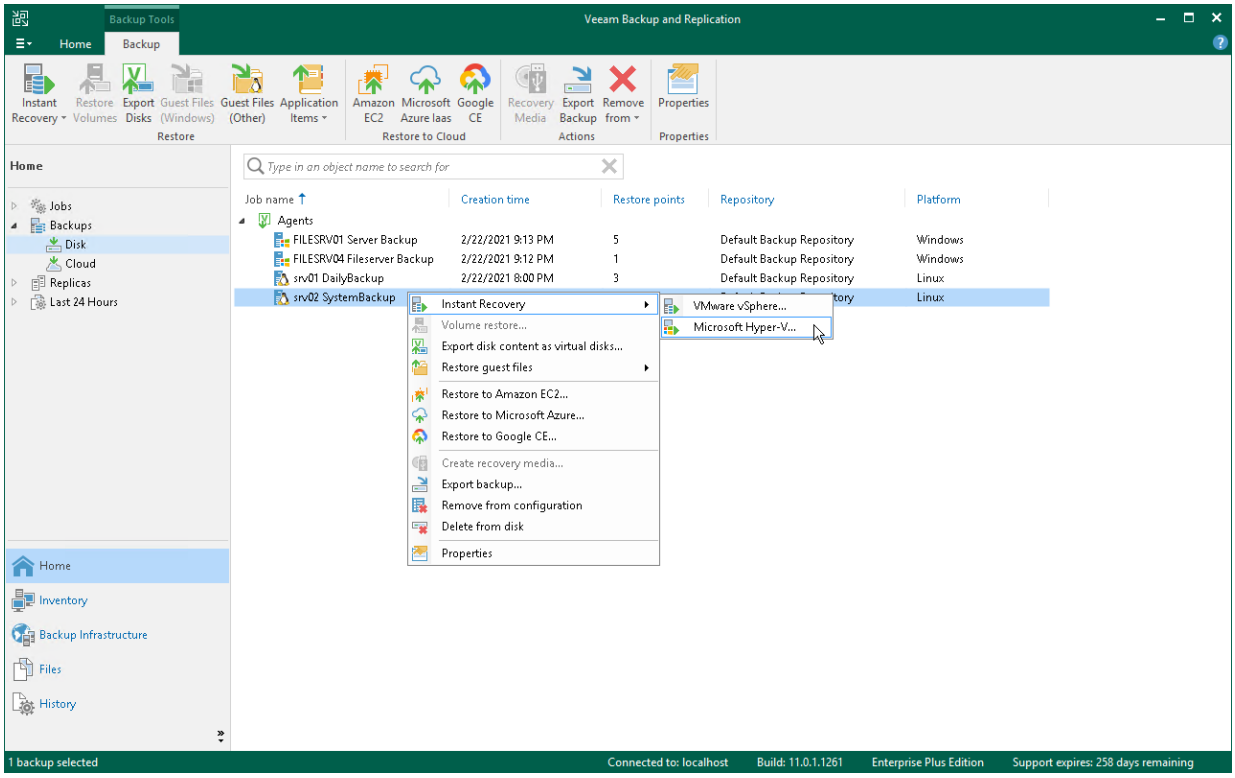
You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a Hyper-V VM in your virtualization environment. For Instant Recovery to a Hyper-V VM, you can use Veeam Agent backups created on the Veeam backup repository. You cannot perform this operation with Veeam Agent backups created on the Veeam Cloud Connect repository.

Mind that to restore to a Hyper-V VM, you must consider the Hyper-V limitations. To learn more, see this Microsoft article.

A restored Hyper-V VM has the same settings as a backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves settings of the Veeam Agent computer from the backup and applies them to the target VM.

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the Performing Instant Recovery of Workloads to Hyper-V VMs section in the Veeam Backup & Replication User Guide.

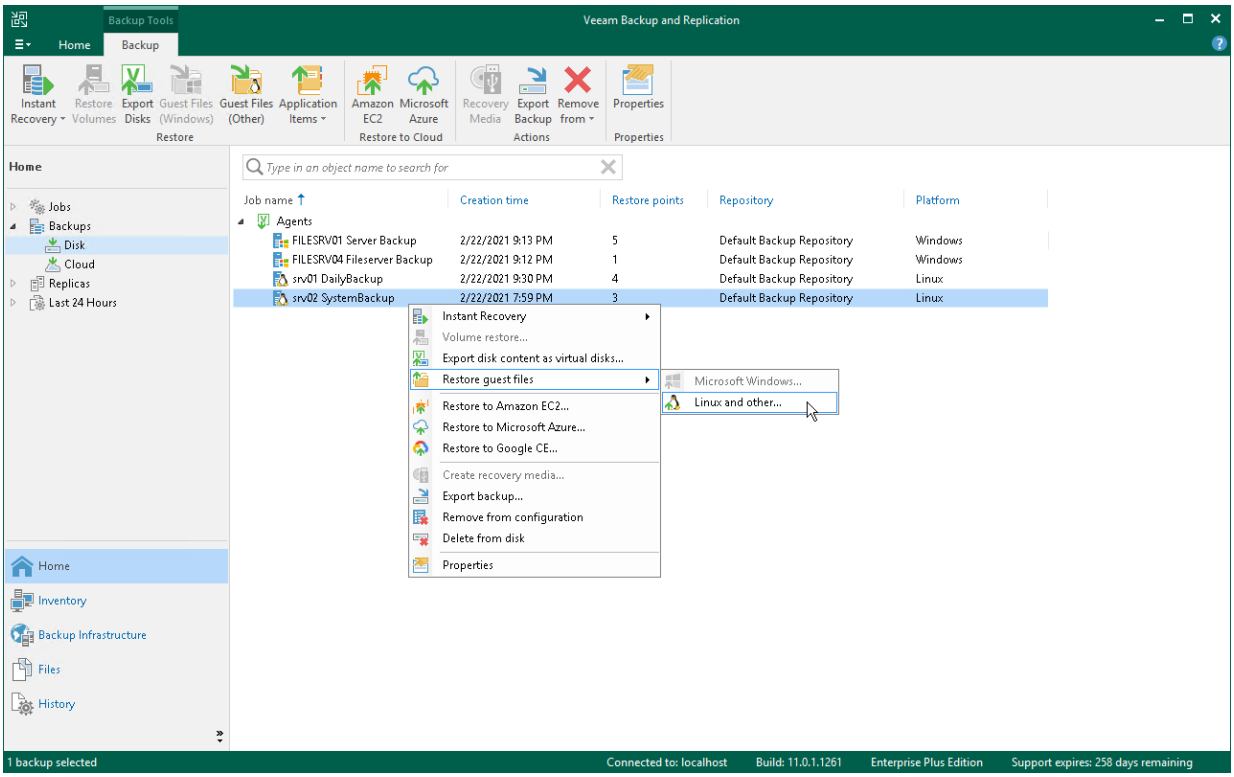
> **IMPORTANT**
>
> If the disk you want to restore contains an LVM volume group, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.



Restoring Files and Folders

You can restore individual files and folders from Veeam Agent backups.

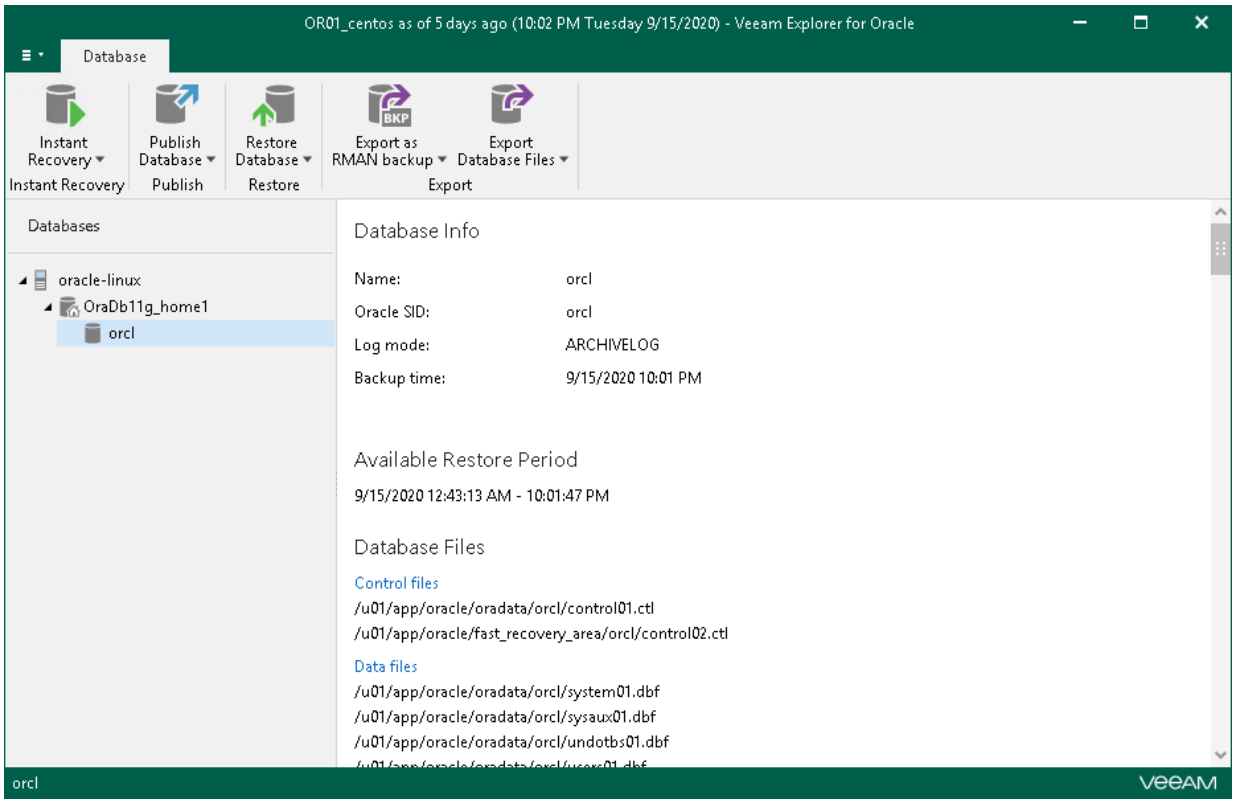
The procedure of file-level restore from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about file-level restore, see Veeam Backup & Replication Documentation.



Restoring Application Items

You can use Veeam Explorer for Oracle to restore application items from backups created with Veeam Agent for Linux.

The procedure of application-item restore from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about the application-items restore procedure, see the Veeam Explorer for Oracle section in the Veeam Explorers User Guide.



Exporting Disks

You can restore computer disks from volume-level backups and convert them to disks of the VMDK, VHD or VHDX format.

During disks restore, Veeam Agent creates standard virtual disks that can be used by VMware vSphere and Microsoft Hyper-V VMs.

- When you restore a disk in the VMDK format, Veeam Agent creates a pair of files that make up the VM virtual disk: a descriptor file and file with the virtual disk content.
- When you restore a disk in the VHD/VHDX format, Veeam Agent creates a file of the VHD or VHDX format.

You can save converted disks locally on any server added to the backup infrastructure or place disks on a datastore connected to an ESXi host (for VMDK disk format only). VMDK disks can be restored as thin provision and thick provision disks:

- Disks restored to a datastore are saved in the thin provisioned format.
- Disks restored to a server are saved in the thick provisioned format.

VHD/VHDX disks are always restored as dynamically expanding.

Veeam Agent supports batch disk restore. For example, if you choose to restore 2 computer disks, Veeam Agent will convert them to 2 virtual disks and store these disks in the specified location.

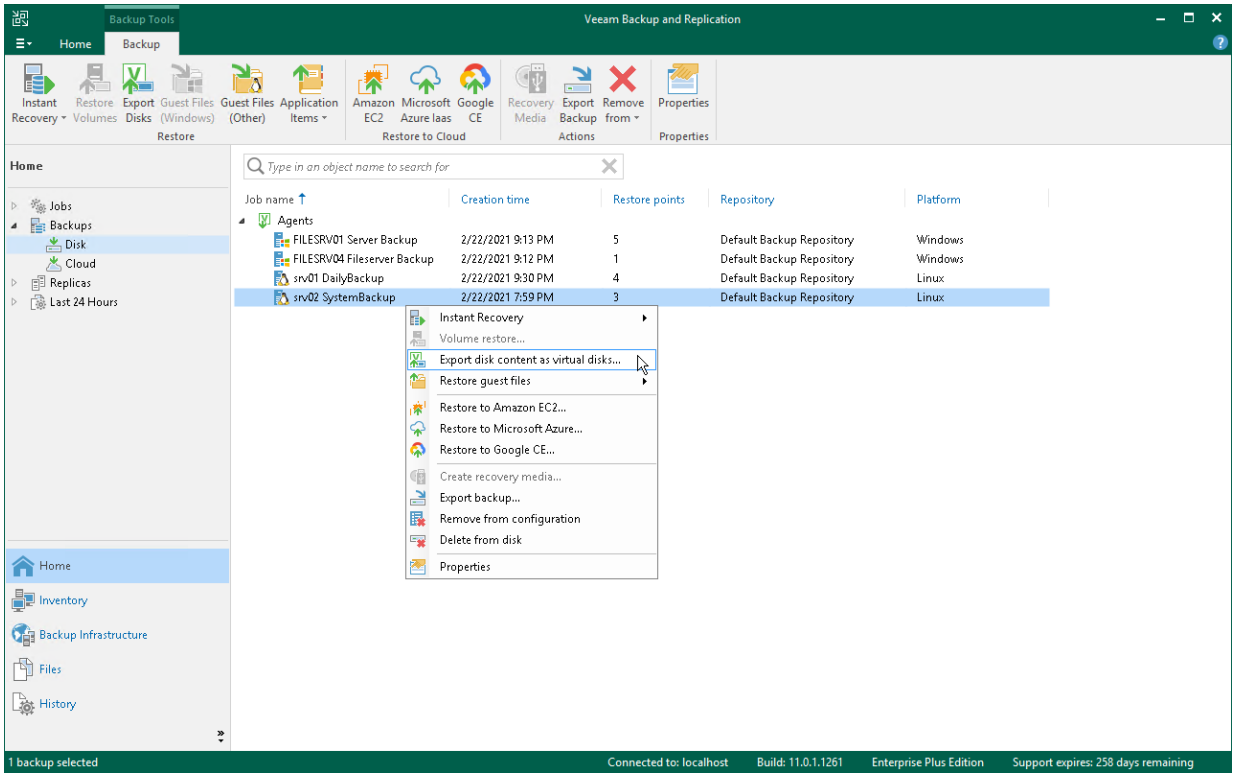
> **IMPORTANT**
>
> Consider the following:
>
> - If the backup from which you restore disks contains a Btrfs storage pool, during the disk restore process Veeam Backup & Replication will create a separate disk and restore the Btrfs pool to this disk.
> - If the disk you want to restore contains an LVM volume group, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.

To restore disks and convert them to the VMDK, VHD or VHDX format, use the **Export Disk** wizard.

Step 1. Launch Export Disk Wizard

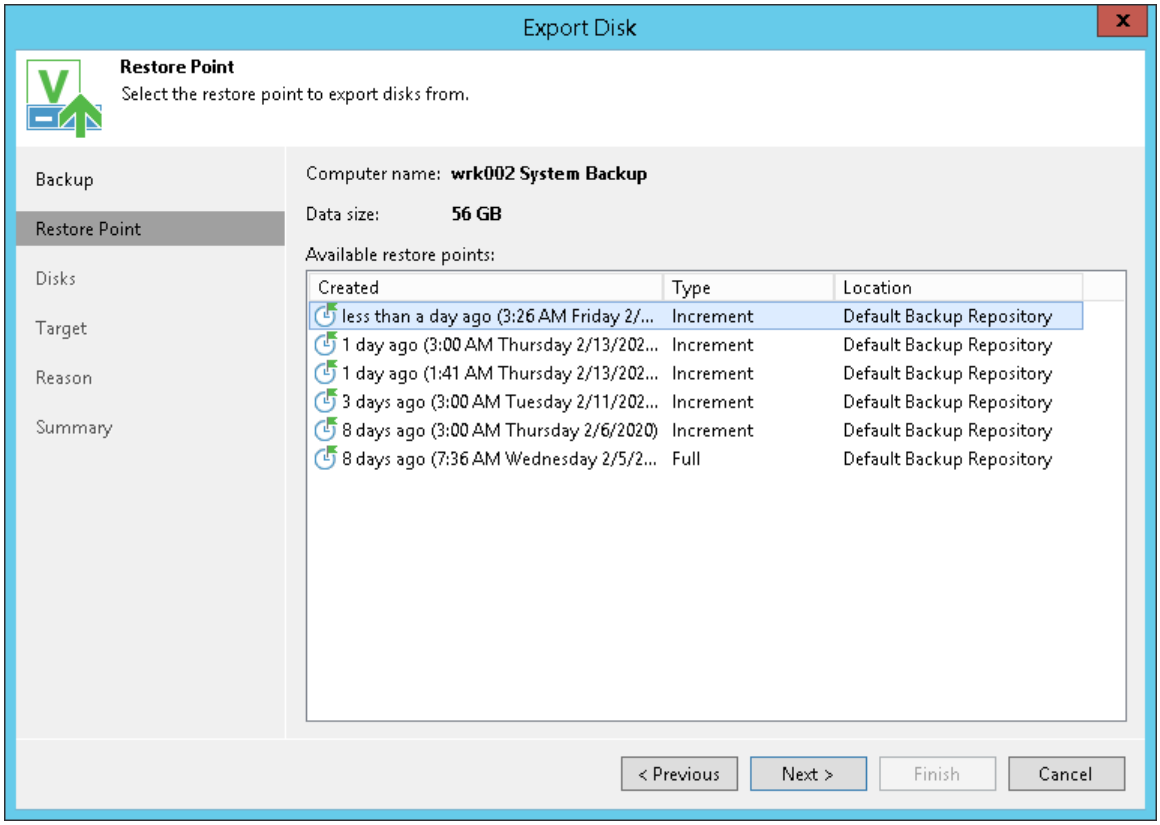
To launch the **Export Disk** wizard:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click **Disk** under the **Backups** node. In the working area, expand the **Agents** node, right-click the necessary backup and select **Export disk content as virtual disks**. You will pass to the Restore Point step of the wizard.



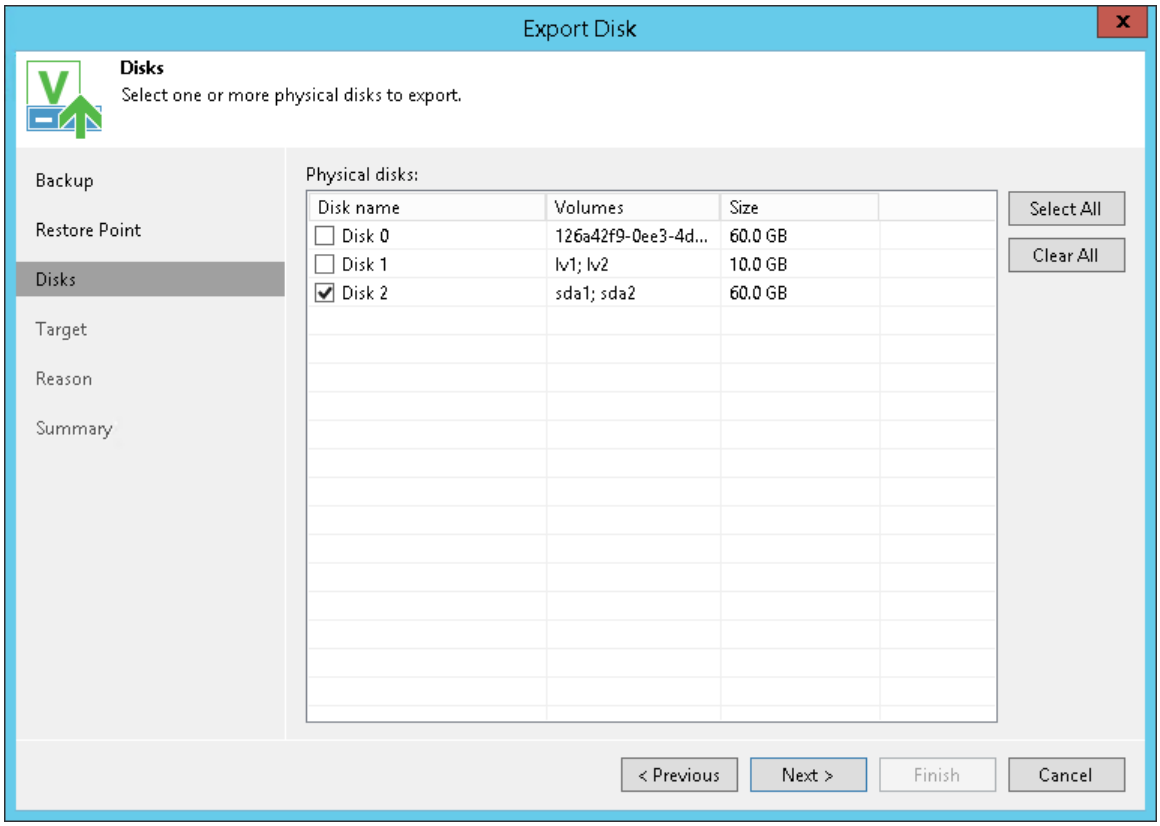
Step 2. Select Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point from which you want to restore disk. In the list of restore points, Veeam Agent displays all restore points that have been created. Make sure that you select a restore point that relates to a volume-level backup.



Step 3. Select Disks

At the **Disks** step of the wizard, select check boxes next to those disks that you want to export.



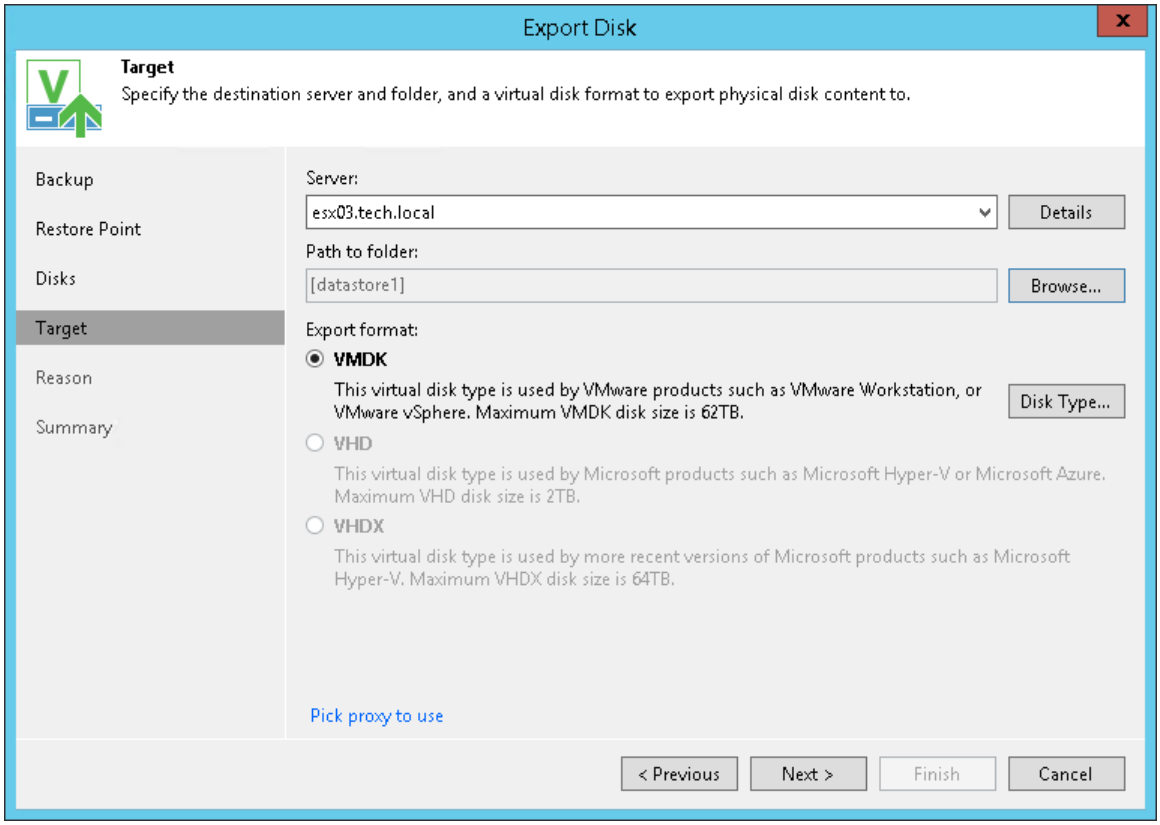
Step 4. Select Destination and Disk Format

At the **Target** step of the wizard, select the destination for disk export and format in which you want to save the resulting virtual disk.

1. From the **Server** list, select a server on which the resulting virtual disks must be saved. If you plan to save the disks in the VMDK format on a datastore, select an ESXi host to which this datastore is connected.
2. In the **Path to folder** field, specify a folder on the server or datastore where the virtual disks must be placed.
3. Select the export format for disks:
   - **VMDK** — select this option if you want to save the resulting virtual disk in the VMware VMDK format.
   - **VHD** — select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHD format.
   - **VHDX** — select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHDX format (supported by Microsoft Windows Server 2012 and later).

NOTE

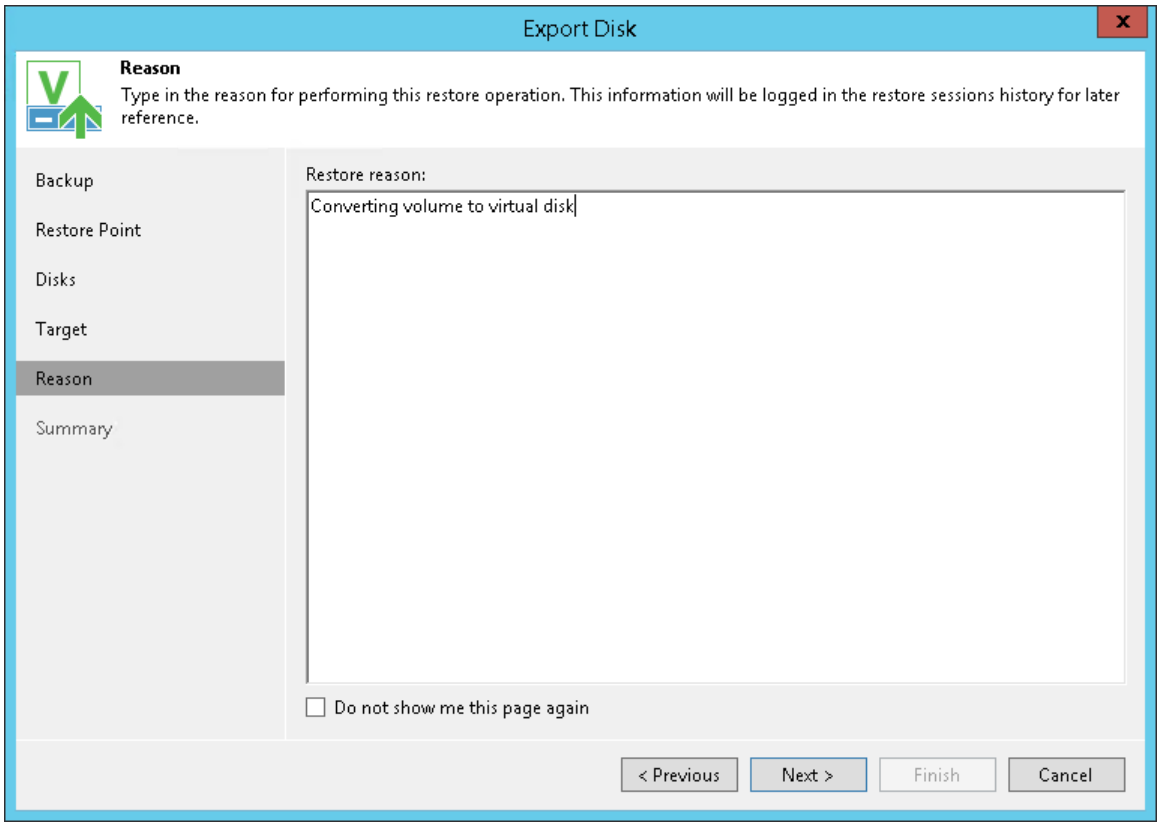
If you have selected to store the resulting virtual disk to a datastore, you will be able to save the virtual disk in the VMDK format only. Other options will be disabled.



Step 5. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the computer volume.

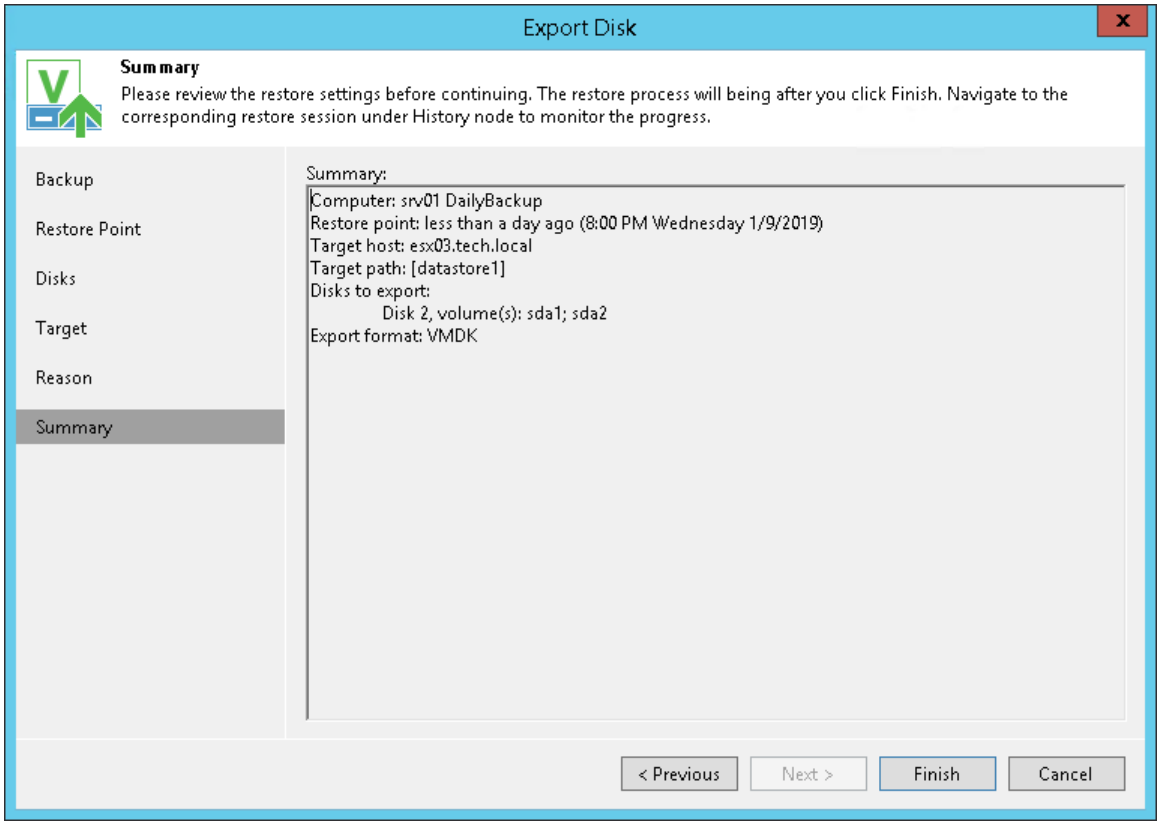
> **NOTE**
>
> If you do not want to display the **Restore Reason** step of the wizard in future, select the **Do not show me this page again** check box.



Step 6. Complete Restore Process

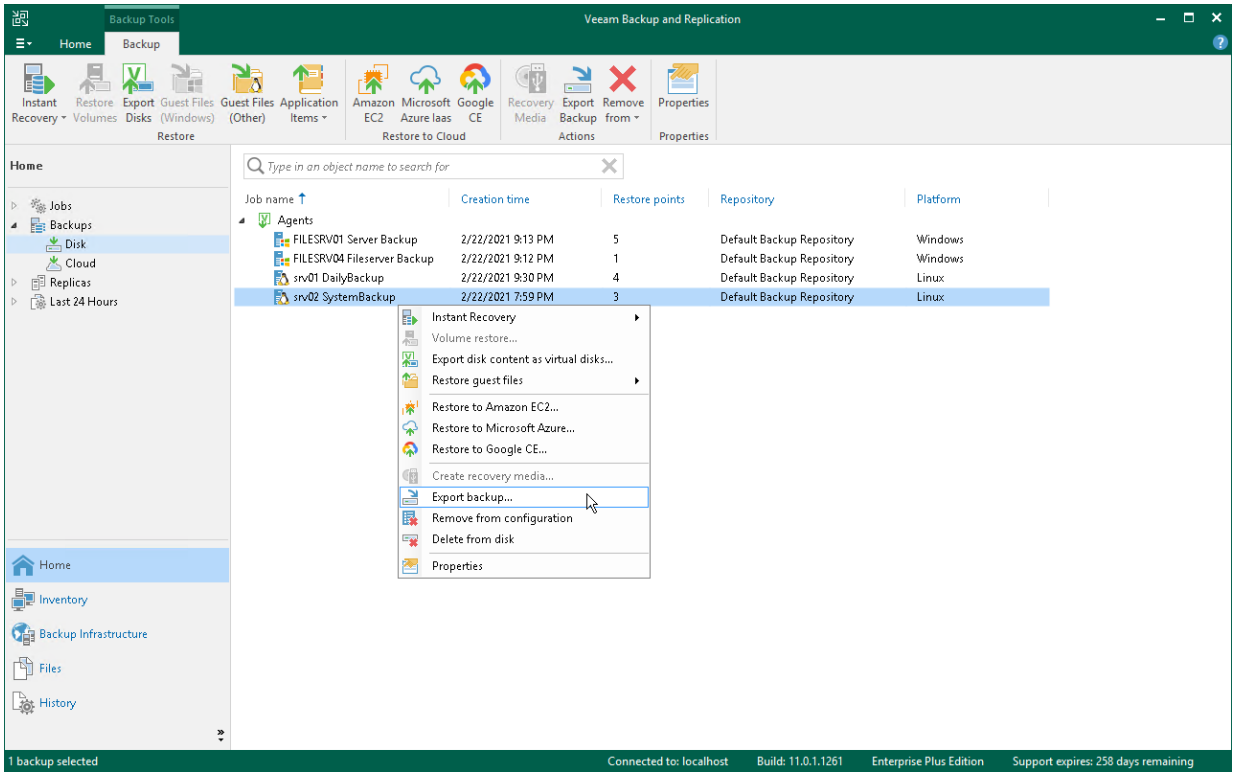
At the **Summary** step of the wizard, complete the procedure disk restore procedure.

1. Review details for the disk to be restored.
2. Click **Finish** to start the restore procedure and exit the wizard.



Exporting Restore Point to Full Backup File

You can restore data from a specific restore point in a Veeam Agent backup and export this data to a standalone full backup file. The procedure of Veeam Agent backup export does not differ from the same procedure for a VM. To learn more, see the Exporting Backups section in the Veeam Backup & Replication User Guide.



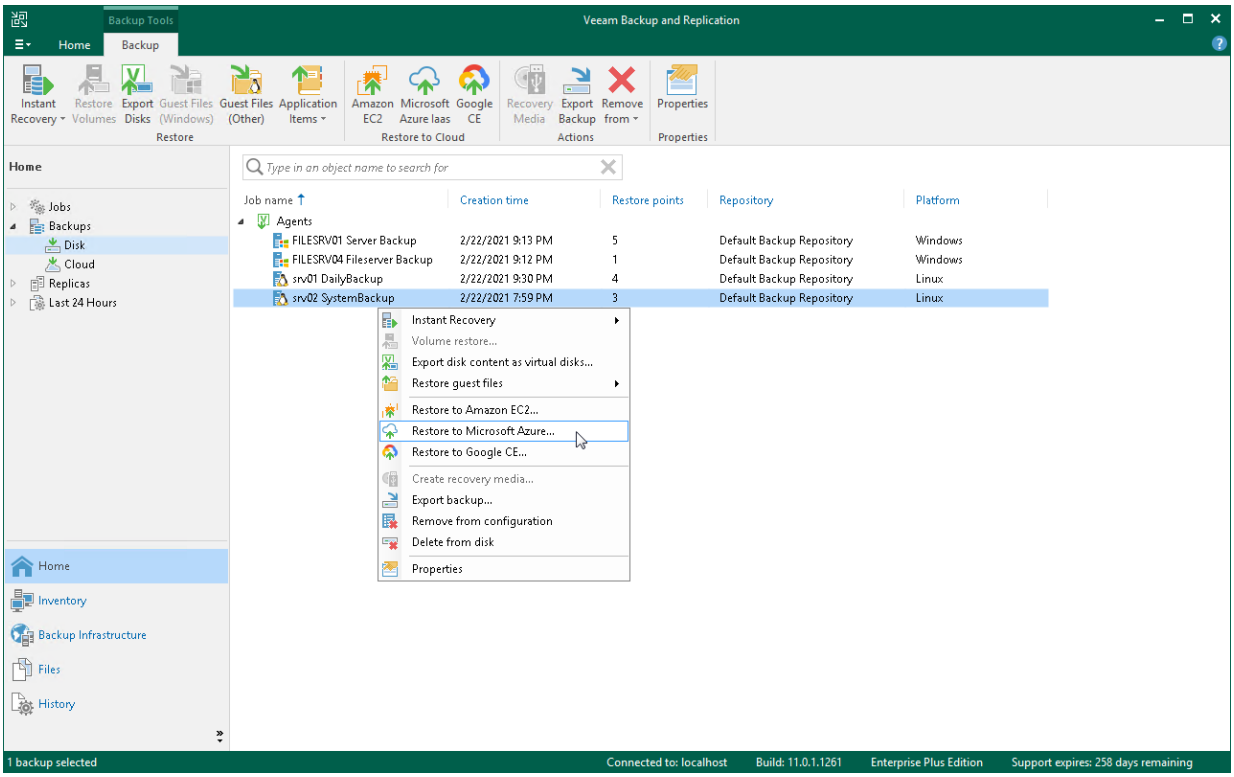
Restoring to Microsoft Azure

You can restore Linux machines from Veeam Agent backups to Microsoft Azure. Backups must be created at the entire machine level or volume level.

If you restore a Veeam Agent machine to Microsoft Azure, consider the following:

- If you recover a EFI-based system to Microsoft Azure, Veeam Agent will restore a BIOS-based Generation 1 VM.
- Veeam Backup & Replication offers experimental support for generation 2 VMs within restore to Microsoft Azure feature. To learn more, see the Generation 2 VM Support section in the Veeam Backup & Replication User Guide.
- If the disk you want to restore contains an LVM volume group, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.

The procedure of restore to Microsoft Azure from a Veeam Agent backup practically is similar to the same restore procedure from a VM backup. To learn more, see the Restore to Microsoft Azure section in the Veeam Backup & Replication User Guide.

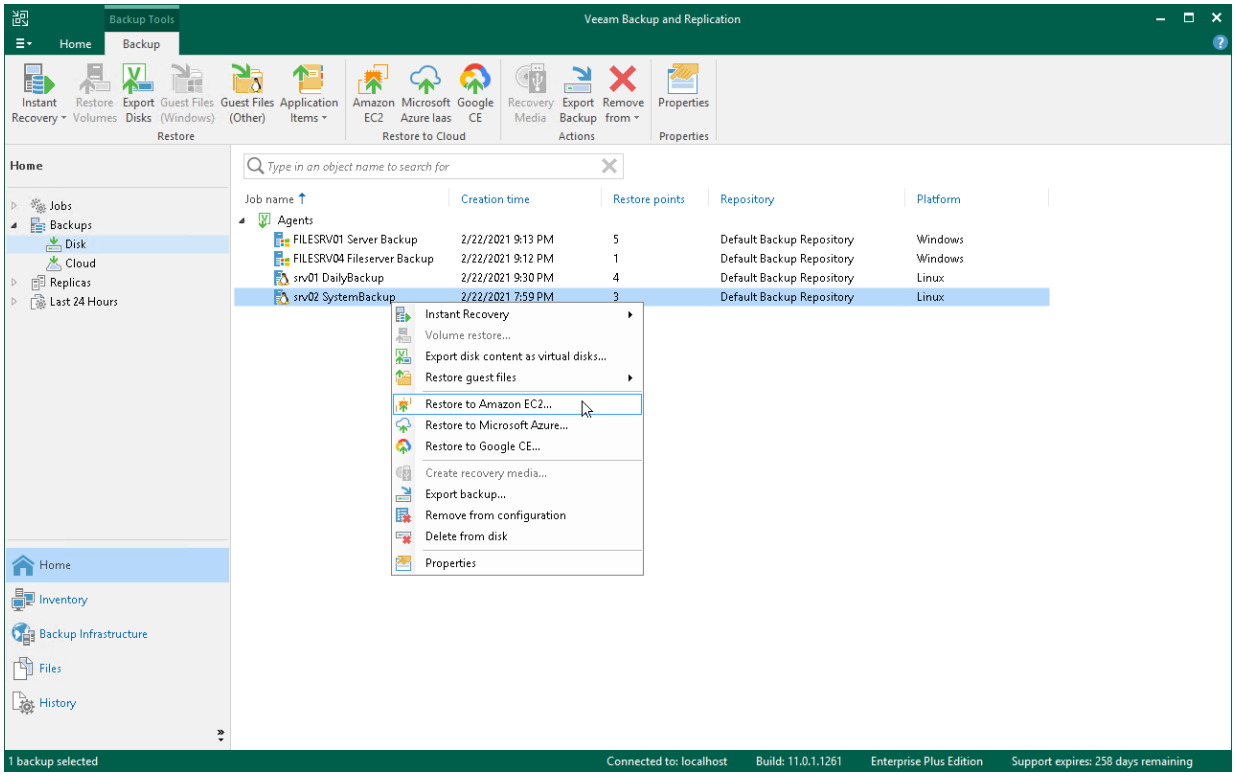


Restoring to Amazon EC2

You can restore Linux machines from Veeam Agent backups to Amazon EC2. Backups must be created at the entire machine level or volume level.

The procedure of restore to Amazon EC2 from a Veeam Agent backup practically is similar to the same restore procedure from a VM backup. To learn more, see the Restore to Amazon EC2 section in the Veeam Backup & Replication User Guide.

> **IMPORTANT**
>
> If the disk you want to restore contains an LVM volume group, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.

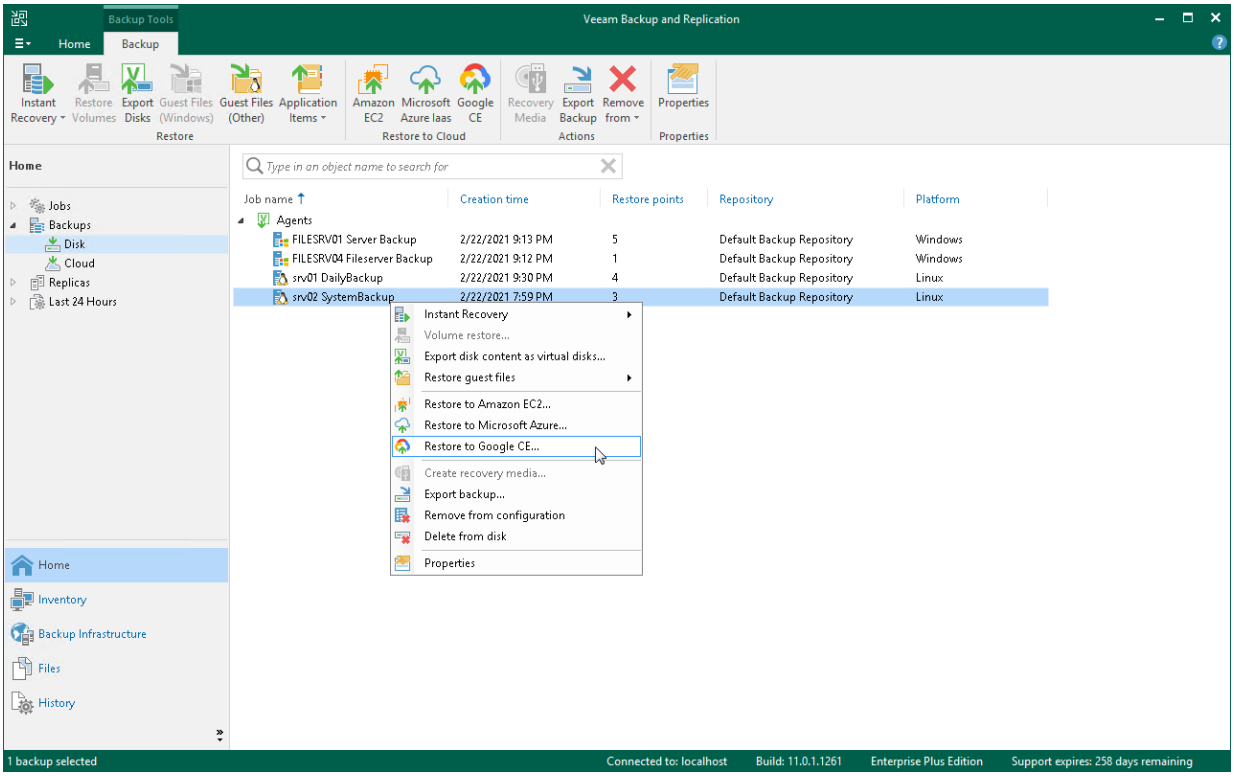


Restoring to Google Cloud Platform

You can restore Linux machines from Veeam Agent backups to Google Cloud Platform. Only volume-level backups are supported.

The procedure of restore to Google Cloud Platform practically does not differ from a regular one. To learn more about restore to Google Cloud Platform, see the Restore to Google Cloud Platform section in the Veeam Backup & Replication User Guide.

> **IMPORTANT**
>
> If the disk you want to restore contains an LVM volume group, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.



Performing Administration Tasks

You can manage Veeam Agent backup jobs and backups created with these jobs. Veeam Backup & Replication allows you to perform the following administration tasks:

- Import Veeam Agent backups
- Enable and disable Veeam Agent backup jobs
- Remove Veeam Agent backup jobs
- View Veeam Agent backup job statistics
- Remove Veeam Agent backups
- View Veeam Agent backup properties
- Configure global settings
- Assign roles to users

Importing Veeam Agent Backups

You may need to import a Veeam Agent backup in the Veeam Backup & Replication console in the following situations:

- The Veeam Agent backup is stored on a drive managed by another computer (not the Veeam backup server).
- The Veeam Agent backup is stored on a backup repository managed by another Veeam backup server.
- The Veeam Agent backup has been removed in the Veeam Backup & Replication console.

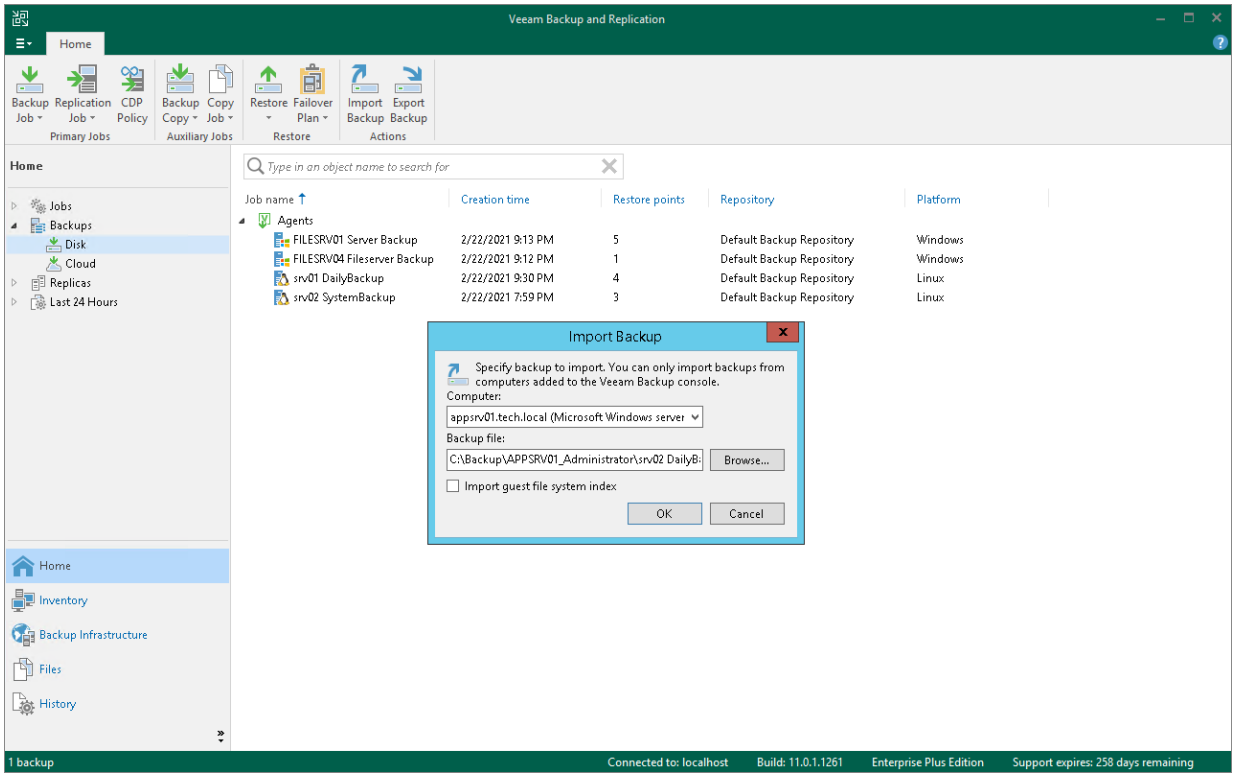
After importing, the Veeam Agent backup becomes available in the Veeam Backup & Replication console. You can restore data from such backup in a regular manner.

Before importing a backup, check the following prerequisites:

- The computer or server from which you plan to import the backup must be added to Veeam Backup & Replication. Otherwise you will not be able to access backup files.
- To be able to restore data from previous backup restore points, make sure that you have all incremental restore points in the same folder where the full backup file resides.

To import a Veeam Agent backup:

1. In Veeam Backup & Replication, click **Import Backup** on the **Home** tab.
2. From the **Computer** list, select the computer or 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 use the VBK files for import only if a corresponding VBM file is not available.
4. Click **OK**. The imported backup will become available in the **Home** view, under the **Backups** > **Disk (imported)** node in the inventory pane.



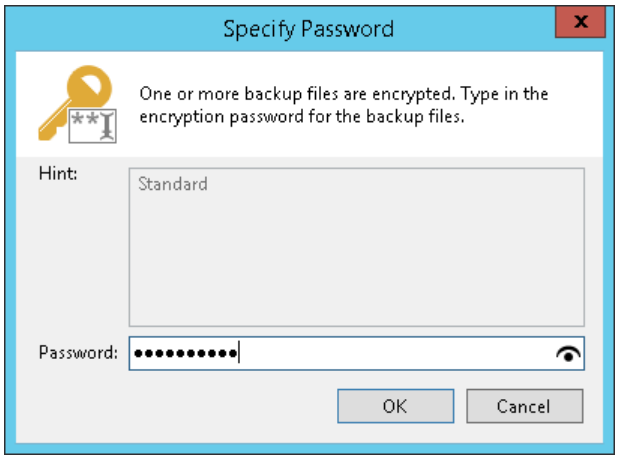
Importing Encrypted Backups

You can import Veeam Agent backups that were encrypted by Veeam Backup & Replication or Veeam Agent for Linux.

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 need to specify the latest password. For Veeam Agent backups, you can use the latest password to restore data form all restore points in the backup chain, including those restore points that were encrypted with an old password.

If you enter correct password, Veeam Backup & Replication will decrypt the backup file. The backup will be moved under the **Backups** > **Disk (imported)** node in the inventory pane.



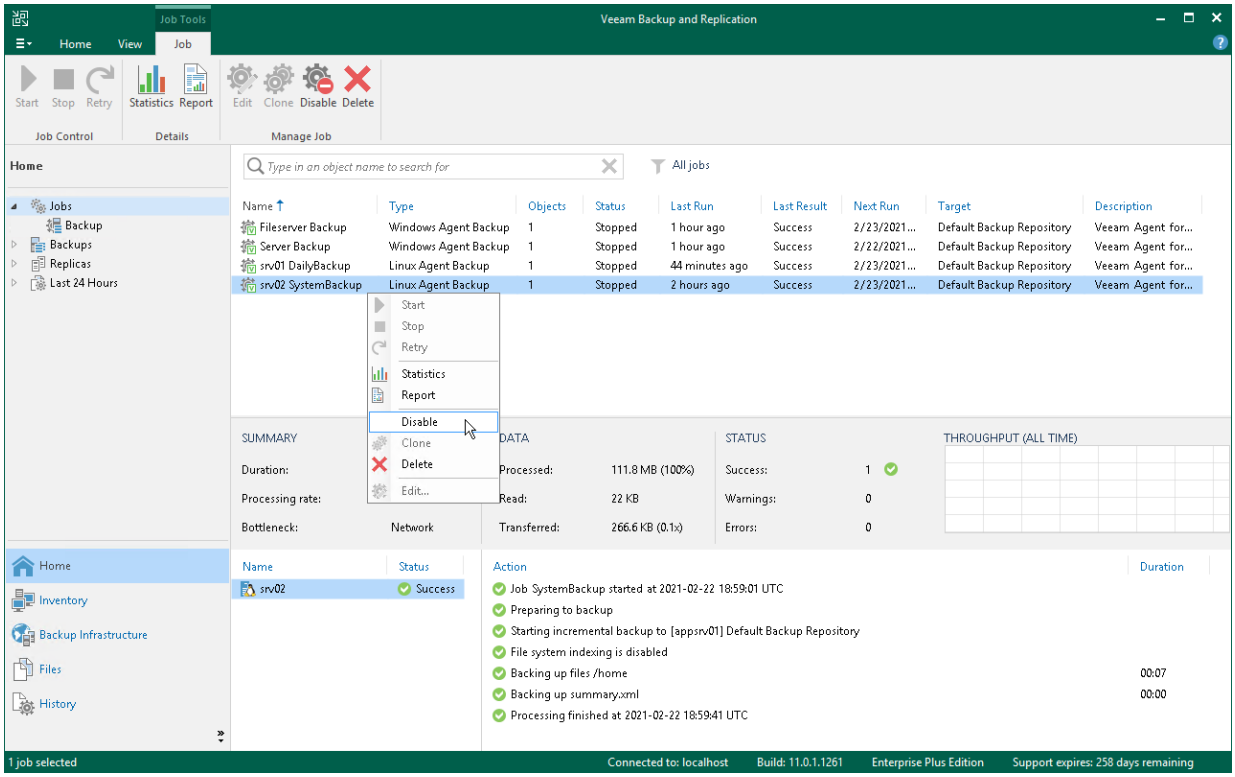
Enabling and Disabling Veeam Agent Backup Jobs

You can disable and enable Veeam Agent backup jobs in Veeam Backup & Replication.

When you disable the job, you prohibit the user to store the resulting backup to the backup repository. If the user starts a disabled job manually or the job starts by schedule, the job session will fail and report the "*Job is disabled on backup server*" error. To let Veeam Agent store backups to the backup repository again, you must enable the disabled job.

To disable or enable the scheduled backup job in Veeam Backup & Replication:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Disable** on the ribbon or right-click the necessary job in the working area and select **Disable**. To enable the disabled job, click **Disable** on the toolbar or right-click the job and select **Disable** once again.



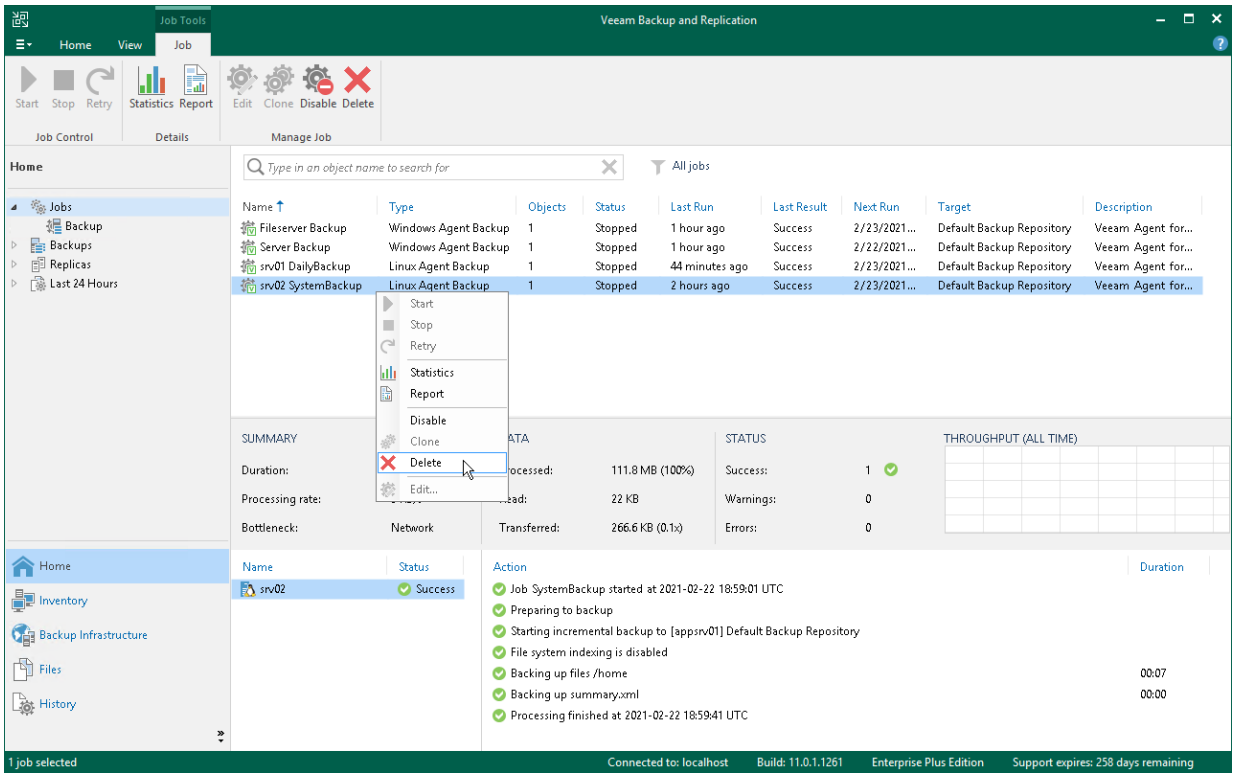
Deleting Veeam Agent Backup Jobs

You can delete Veeam Agent backup jobs.

When you delete a Veeam Agent backup job, Veeam Backup & Replication removes all records about the job from its database and console. When the user starts a new Veeam Agent backup job session manually or the job starts automatically by schedule, the job will appear in the Veeam Backup & Replication console again, and records about a new job session will be stored to the Veeam Backup & Replication database. To remove the job permanently, you must delete the job and unassign access rights permissions for this user from the backup repository.

To remove a job:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Delete** on the toolbar or right-click the necessary job in the working area and select **Delete**.

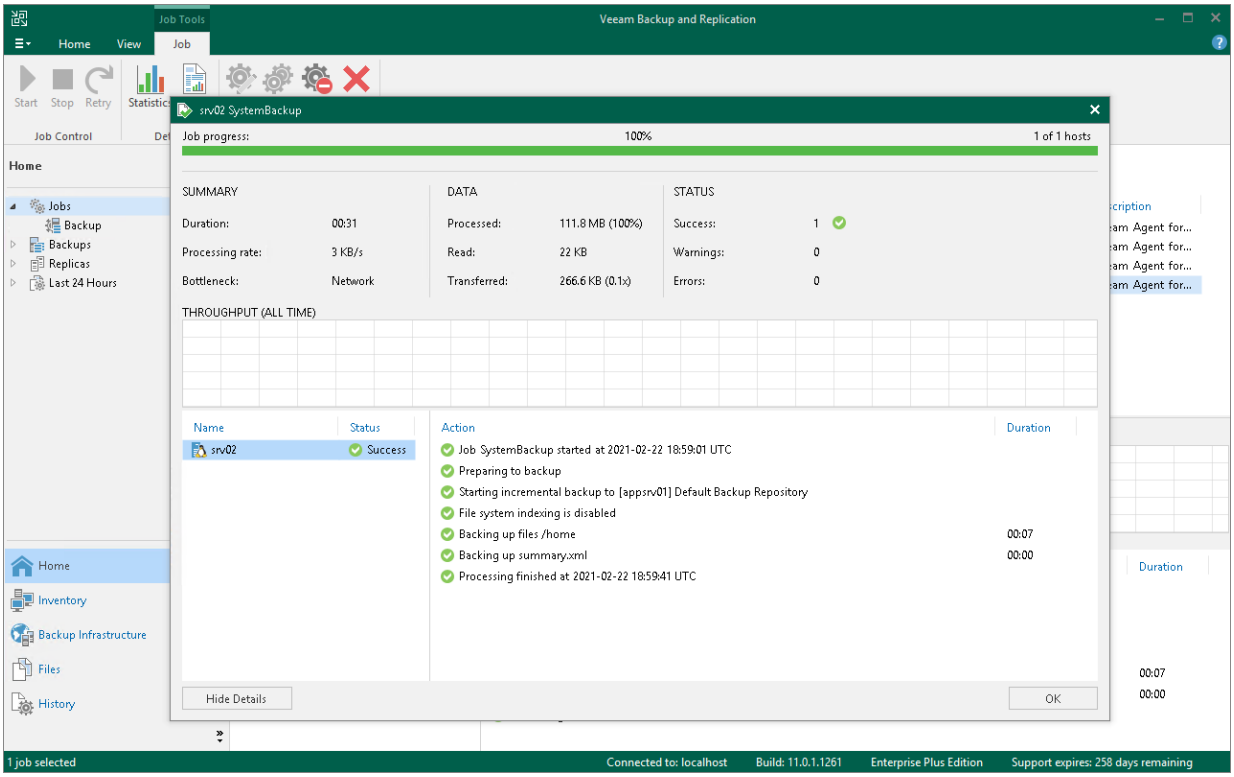


Viewing Veeam Agent Backup Job Statistics

You can view statistics about Veeam Agent backup jobs in the Veeam Backup & Replication console. Veeam Backup & Replication displays statistics for Veeam Agent backup jobs in the similar way as for regular backup jobs. The difference is that the list of objects included in the job contains a Veeam Agent machine instead of one or several VMs.

To view Veeam Agent backup job statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. In the working area, select the necessary Veeam Agent backup job and click **Statistics** on the ribbon or right-click the job and select **Statistics**.



Removing Veeam Agent Backups

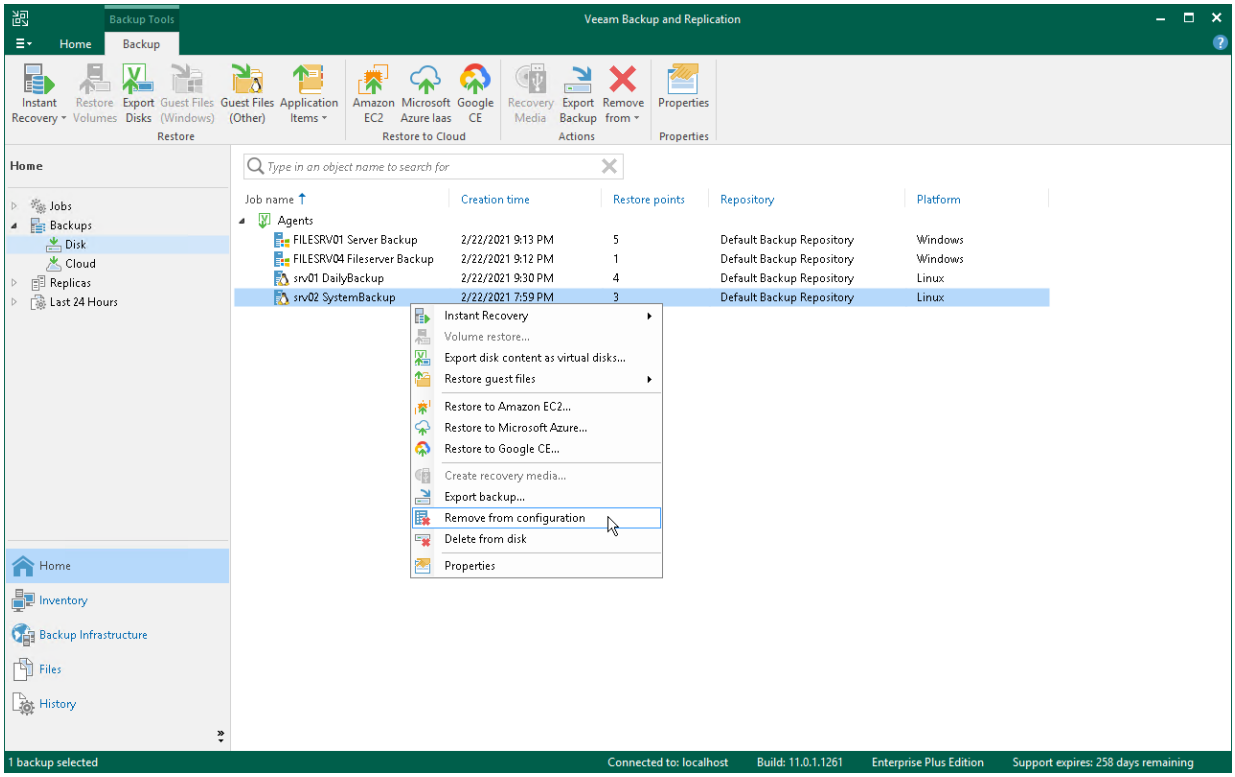
You can remove Veeam Agent backups from Veeam Backup & Replication or permanently delete Veeam Agent backups from the Veeam backup repository.

Removing from Configuration

When you remove a Veeam Agent backup from configuration, Veeam Backup & Replication deletes all records about the backup from its database and console. The actual backup files remain on the backup repository. You can import the backup to the Veeam Backup & Replication at any time later and restore data from it. To learn more, see Importing Veeam Agent for Linux Backups.

To remove a Veeam Agent backup from configuration:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click **Disk** under the **Backups** node.
3. In the working area, expand the **Agents** node, select the necessary backup and click **Remove from > Configuration** on the toolbar or right-click the backup and select **Remove from configuration**.



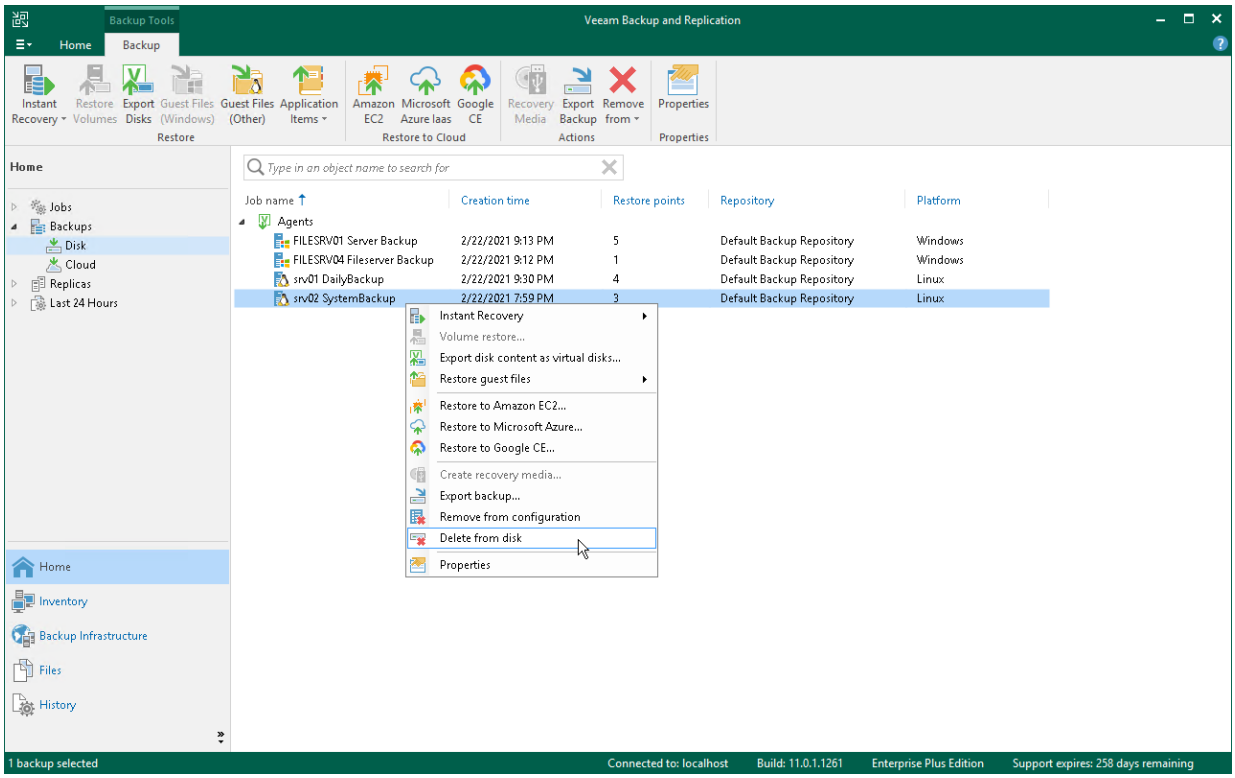
Removing from Veeam Backup Repository

When you remove a Veeam Agent backup from the Veeam backup repository, Veeam Backup & Replication deletes all records about the backup from its database and console. The actual backup files are removed from the backup repository, too.

To remove a Veeam Agent backup from the backup repository:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click **Disk** under the **Backups** node.

3. In the working area, expand the **Agents** node, select the necessary backup and click **Remove from > Disk** on the toolbar or right-click the backup and select **Delete from disk**.

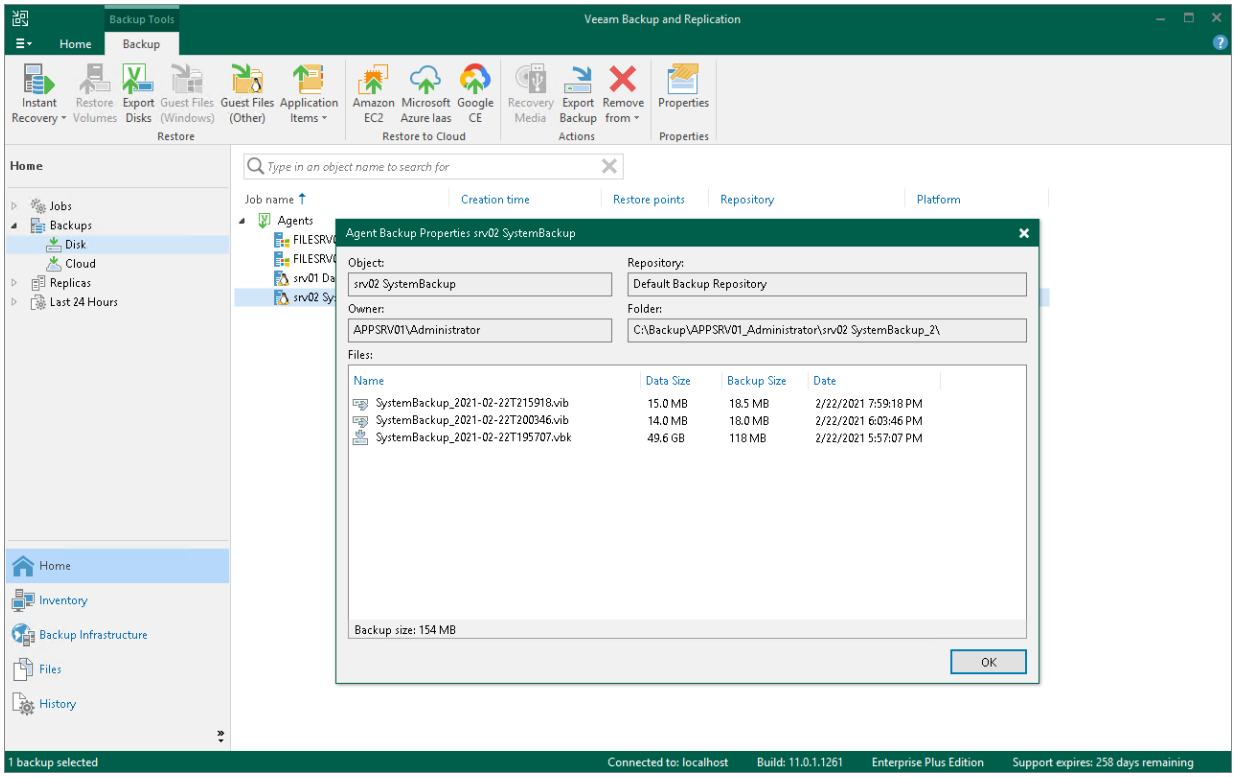


Viewing Veeam Agent Backup Properties

You can view statistics about Veeam Agent backups.

To view Veeam Agent backup statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click **Disk** under the **Backups** node.
3. In the working area, expand the **Agents** node, select the necessary backup and click **Properties** on the toolbar or right-click the backup and select **Properties**.



Configuring Global Settings

Global settings configured on the Veeam backup server apply to Veeam Agent backup jobs as well. You can:

- Configure network throttling settings so that Veeam Agent backup job does not consume all network resources.
- Configure global email settings to get alerted about the Veeam Agent backup job results.

To learn more, see Veeam Backup & Replication Documentation.

Assigning Roles to Users

User roles configured on the Veeam backup server apply to Veeam Agent backup jobs as well.

To learn more, see the Roles and Users section in the Veeam Backup & Replication User Guide.

Appendix A. Requirements for veeam-nosnap

Starting from build 5.0.1.4493, you can use a `veeam-nosnap` package to install Veeam Agent for Linux. This package allows Veeam Agent to operate without `veeamsnap` kernel module.

The `veeam-nosnap` package can be useful in the following cases:

- You do not want to install kernel sources and compilers on your computer.
- You want to use third-party tools to create data snapshots.
- You want to perform bare metal restore, but Veeam Recovery Media does not work with your computer. In this case you can install the `veeam-nosnap` package on LiveCD of your choice and access the Veeam recovery UI.

Before you install Veeam Agent using the `veeam-nosnap` package, consider the following limitations:

- The `veeam-nosnap` package has an experimental support status. For details about experimental support, see this Veeam KB article.
- The RAM-based changed block tracking (CBT) mechanism is not supported. As a result, if you plan to back up a significant amount of data, the backup will require greater time.
- Veeam Agent can create a snapshot of LVM logical volumes and Btrfs subvolumes. To back up data that resides on other file systems and volumes, you can use only file-level backup in the snapshot-less mode. For details, see Snapshot-Less File-Level Backup.
- For a successful backup, Veeam Agent requires unallocated extents on volume groups.
- For a successful bare metal restore, all disks of the Veeam Agent computer you want to restore must be available in the backup.

System Requirements for veeam-nosnap

If you plan to use the `veeam-nosnap` package to install Veeam Agent, the protected Linux computer meet requirements listed in the table below. To learn about system requirements for Veeam Agent installed using the `veeamsnap` kernel module, see System Requirements.

Specification	Requirement
Hardware	CPU: x86-64 processor (i386 or later). Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. Disk Space: 100 MB free disk space for product installation. Network: 10 Mbps or faster network connection to a backup target. System firmware: BIOS or UEFI. Disk layout: MBR or GPT. For virtual machines: Only full virtualization type is supported. Containers and paravirtualized instances are not supported. Oracle VM virtual machines are supported with limitations.
OS	**Important!** Check considerations and limitations that apply to the list of supported OSes. Linux kernel version 2.6.32 or later is supported. Both 64-bit and 32-bit versions (if applicable) of the following distributions are supported: • Debian 9.0 – 11.1 • Ubuntu 14.04, 16.04, 18.04, 20.04, 20.10, 21.04, 21.10 • RHEL 6.0 – 8.5 • CentOS 7.0 – 8.5 • 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)

Specification	Requirement
File System	**Important!** Check considerations and limitations that apply to the list of supported file systems. Veeam Agent for Linux supports consistent snapshot-based data backup for the following file systems: • LVM LVM volumes protected with encryption software such as dm-crypt are supported. • Btrfs (for OSes that run Linux kernel 3.16 or later) Btrfs is supported only if it resides directly on a physical device with no additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) below or above it. Data that resides on other file systems and volumes (including NFS and SMB shares) can be backed up using the snapshot-less mode. For details, see Snapshot-Less File-Level Backup.
Software	**Important!** Check considerations and limitations that apply to the list of supported components. Protected computer must have the following components installed: • perl • linux-headers (for Debian-based systems) • libudev • libacl • libattr • lvm2 • libfuse • dmidecode • libmysqlclient • libpq5 • efibootmgr (for UEFI-based systems) • isolinux (for Debian-based systems) • syslinux (for RedHat-based systems) • btrfs-progs (for backup of Btrfs file system) • mksquashfs (for custom Veeam Recovery Media) • unsquashfs (for custom Veeam Recovery Media) • wget (for custom Veeam Recovery Media) • xorriso (for custom Veeam Recovery Media with EFI support)

Considerations and Limitations

OS

- Linux kernel version 2.6.32 or later is supported as long as you use kernels supplied by your distribution with the following limitation: Linux kernel 2.6.32-754.6.3 in RHEL and Oracle Linux (RHCK) is not supported.
- Only GA versions of the supported distributions that have been released before the current version of Veeam Agent for Linux are supported.
- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.
- RHEL, CentOS and Oracle Linux (RHCK) are supported up to certain kernel versions. For details, see this Veeam KB article.

File System

- Veeam Agent for Linux does not back up volumes that reside on USB devices and SD cards
- Total size of all file systems included in a file-level backup must not exceed 218 TB. Size of a file included in a file-level backup must not exceed 16 TB.
- Veeam Agent supports backup of extended attributes with the following limitations:
  - Veeam Agent backs up extended attributes only with the following public namespaces: system, security, trusted, and user.
  - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

    For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the ea_inodes feature. Backups created using the ea_inodes feature cannot be mounted on kernel versions up to 4.12.
- Backup of file and directory attributes (for example, a — append only, c — compressed, and so on) is not supported.
- Each volume included in a backup must have a unique UUID.
- You cannot back up an entire system image or specific volumes of machines used as cluster nodes. Only snapshot-less file-level backup of cluster nodes is supported. That includes backup of machines that use shared disks, clustered file systems or clustered LVM.
- Certain limitations for EMC PowerPath configuration apply. To learn more, see this Veeam KB article.
- BFQ I/O scheduler is not supported.
- Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.

Software

> **IMPORTANT**
>
> Linux user account used to work with Veeam Agent for Linux must have the `/bin/bash` shell set as the default shell.

- Version of the `linux-headers` package (for Debian-based systems) varies according to the Linux kernel version that you use.

- The `dmidecode` package is required for Veeam Agent management — a valid BIOS UUID must be obtainable either from `dmidecode | grep -i uuid` or from `/sys/class/dmi/id/product_uuid`. Each Veeam Agent that consumes a license installed in Veeam Backup & Replication must have a unique BIOS UUID. If a valid UUID cannot be obtained, Veeam will generate it automatically.
- The `libmysqlclient` package is required to process MySQL database system located on the Veeam Agent server. For details, see Backup of MySQL Database. Package version varies according to the MySQL database system version that you use.
- The `libpq5` package is required to process PostgreSQL database system located on the Veeam Agent server. For details, see Backup of PostgreSQL Database.
- The `btrfs-progs` package version 3.16 or later is required.

Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see File System.

Backup Target

Backup can be performed to the following types of storage:

- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.
- Veeam Backup & Replication 11 or later backup repository (including deduplication appliances).
- Veeam Cloud Connect 11 or later cloud repository.

Network

Consider the following:

- Veeam Agent for Linux should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 addresses.

Veeam Agent for Microsoft Windows

Version 5.0

User Guide

January, 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 **8**

ABOUT THIS DOCUMENT **9**

OVERVIEW **10**

Solution Architecture 11

Standalone and Managed Operation Modes 12

Data Backup 13

Backup Types 14

How Backup Works 25

Backup Job 27

Backup Job Schedule 28

Ad-Hoc Backup 39

Parallel Disk Processing 43

Backup Chain 44

Changed Block Tracking 67

Data Compression 72

Guest Processing 74

Data Encryption 92

Backup Cache 105

Backup to Rotated Drives 109

Backup of External Drives 112

Data Restore 113

Volume-Level Restore 114

File-Level Restore 115

Volume Resize 116

Veeam Recovery Media 119

Drivers in Veeam Recovery Media 120

BitLocker Encrypted Volumes Support 121

Microsoft OneDrive Support 126

Integration with Veeam Backup & Replication 129

Backup to Veeam Cloud Connect Repository 131

MANAGING VEEAM AGENT IN VEEAM BACKUP & REPLICATION **132**

PLANNING AND PREPARATION **133**

System Requirements 134

Used Ports 138

LICENSING **141**

Product Editions 142

License Agreements 143
Installing License 144
Selecting Product Edition 146
Revoking License 147
INSTALLATION AND CONFIGURATION 148
Before You Begin 149
Installing Veeam Agent for Microsoft Windows 150
Installing Veeam Agent for Microsoft Windows in Unattended Mode 151
Using Sysprep and Veeam Agent for Microsoft Windows 152
Upgrading Veeam Agent for Microsoft Windows 153
Configuring Advanced Settings 154
Uninstalling Veeam Agent for Microsoft Windows 155
What You Do Next 156
GETTING STARTED 157
PERFORMING BACKUP 158
Creating Veeam Recovery Media 159
Before You Begin 160
Step 1. Launch Create Recovery Media Wizard 162
Step 2. Specify Recovery Media Options 163
Step 3. Specify Path to ISO 165
Step 4. Review Recovery Image Settings 166
What You Do Next 168
Creating Veeam Recovery Media with Command Line Interface 169
Auto-Configuring Scheduled Backup Jobs 170
Configuring Backup Job 171
Before You Begin 172
Step 1. Launch New Backup Job Wizard 173
Step 2. Specify Job Name and Description 174
Step 3. Select Backup Mode 175
Step 4. Specify Backup Scope Settings 177
Step 5. Select Backup Destination 186
Step 6. Specify Backup Storage Settings 188
Step 7. Specify GFS Retention Policy 209
Step 8. Specify Advanced Backup Settings 211
Step 9. Specify Backup Cache Settings 218
Step 10. Specify Guest Processing Settings 219
Step 11. Specify Backup Schedule 230
Step 12. Review Backup Job Settings 236
What You Do Next 237

Managing Backup Job 238
Editing Backup Job Settings 239
Disabling and Enabling Scheduled Backups 240
Removing Backup Job 242
Controlling Backup Post-Job Action 243
Performing Ad-Hoc Backups 244
Creating Incremental Backups 245
Creating Active Full Backups 247
Creating Standalone Full Backups 249
Performing Backup to Another Location 251
Performing Backup with Command Line Interface 253
Deleting Backups 257
Managing Backup Cache 259
Monitoring Backup Cache Activity 260
Pausing Backup Cache Synchronization 263
Deleting Restore Points from Backup Cache 264
PERFORMING RESTORE 265
Restoring from Veeam Recovery Media 266
Before You Begin 267
Step 1. Boot from Veeam Recovery Media 268
Step 2. Select Network Adapter or Wireless Network 270
Step 3. Launch Veeam Recovery Media Wizard 272
Step 4. Specify Backup File Location 273
Step 5. Select Remote Storage Type 275
Step 6. Specify Remote Storage Settings 276
Step 7. Select Backup 283
Step 8. Select Restore Point 285
Step 9. Select Data Restore Mode 286
Step 10. Map Restored Disks 287
Step 11. Resize Restored Volumes 293
Step 12. Start Restore Process 294
Using Veeam Agent and Microsoft Windows Tools 295
Using Microsoft Windows Recovery Environment 297
Restoring Volumes 298
Before You Begin 299
Step 1. Launch Volume Level Restore Wizard 300
Step 2. Specify Backup File Location 301
Step 3. Select Remote Storage Type 302
Step 4. Specify Remote Storage Settings 303

Step 5. Select Backup 310
Step 6. Select Restore Point 312
Step 7. Map Restored Disks 313
Step 8. Resize Restored Volumes 318
Step 9. Complete Restore Process 319
Restoring Files and Folders 320
Before You Begin 321
Step 1. Launch File Level Restore Wizard 322
Step 2. Specify Backup File Location 323
Step 3. Select Remote Storage Type 324
Step 4. Specify Remote Storage Settings 325
Step 5. Select Backup 333
Step 6. Select Restore Point 334
Step 7. Complete Restore Process 335
Step 8. Save Restored Files 336
Restoring Data from Encrypted Backups 342
REPORTING 344
Viewing Statistics in Control Panel 345
Viewing Statistics for Separate Restore Points 348
Viewing Information About Job Retries 352
Viewing Status of Restore Points in Backup Cache 354
Monitoring Backup State with Tray Agent 355
Monitoring Backup Process in Taskbar Button 357
Viewing and Dismissing Veeam Agent Events 358
Viewing Events with Windows Notification Center 362
Viewing Job Session Results in Email Reports 364
SPECIFYING SETTINGS 365
Throttling Backup Activities 366
Restricting Network Connections Usage 367
Limiting Bandwidth Consumption 368
Disabling Backup over Metered Connections 369
Disabling Backup over VPN Connections 371
Selecting Wireless Networks for Backup 373
Managing Rotated Drives 375
Disabling Control Panel Notifications 376
Enabling Email Notifications 377
Checking for New Product Versions and Updates 380
MANAGING VEEAM CBT DRIVER 381
Installing Veeam CBT Driver 382

Removing Veeam CBT Driver 384
Removing CBT Driver with Veeam Recovery Media 385
Resetting CBT 386
GETTING SUPPORT 387
Reporting Issues 388
USING WITH VEEAM BACKUP & REPLICATION 390
Setting Up User Permissions on Backup Repositories 392
Managing License 395
Managing Instance Consumption by Veeam Agents 396
Assigning License to Veeam Agent 397
Viewing Licensed Veeam Agents and Revoking License 398
Performing Data Protection Tasks 400
Backing Up to Backup Repositories 401
Backing Up to Cloud Repositories 402
Performing Backup Copy for Veeam Agent Backups 404
Archiving Veeam Agent Backups to Tape 406
Performing Restore Tasks 407
Restoring Veeam Agent Backup to Hyper-V VM 408
Restoring Veeam Agent Backup to VMware vSphere VM 410
Restoring Volumes 412
Restoring Files and Folders 413
Restoring Application Items 414
Exporting Disks 416
Restoring to Microsoft Azure 425
Restoring to Amazon EC2 426
Restoring to Google Cloud Platform 427
Exporting Restore Point to Full Backup File 428
Performing Administration Tasks 429
Importing Veeam Agent Backups 430
Enabling and Disabling Veeam Agent Backup Jobs 432
Deleting Veeam Agent Backup Jobs 433
Viewing Veeam Agent Backup Job Statistics 434
Removing Veeam Agent Backups 435
Viewing Veeam Agent Backup Properties 438
Configuring Global Settings 439
Assigning Roles to Users 440
AUTOMATING VEEAM AGENT FOR WINDOWS OPERATIONS 441
APPENDIX A. VEEAM AGENT EVENTS 442

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 Veeam Agent for Microsoft Windows.

Intended Audience

The user guide is intended for anyone who wants to use Veeam Agent for Microsoft Windows 5.0 or later to protect their computer.

Overview

Veeam Agent for Microsoft Windows is a data protection and disaster recovery solution for physical and virtual machines. Veeam Agent for Microsoft Windows can be used to protect different types of computers and devices: desktops, laptops and tablets. The solution can be installed on Windows-based workstations, physical servers and virtual machines. To learn about supported OSes, see System Requirements.

> **NOTE**
>
> Veeam Agent for Microsoft Windows can operate in either standalone or managed mode. Depending on the mode, Veeam Agent provides different features and limitations. To learn more, see Standalone and Managed Operation Modes.

Veeam Agent for Microsoft Windows offers a variety of features to protect your data. You can:

- Create a Veeam Recovery Media on an external hard drive, USB flash drive, CD/DVD/BD, or create an ISO file with the Veeam Recovery Media on disk.
- Create an entire system image backup, back up specific computer volumes or individual folders with files. Backups can be stored on an external hard drive, in a network shared folder, on a Veeam backup repository or Veeam Cloud Connect repository, as well as in Microsoft OneDrive.

In case of a disaster, you can perform the following restore operations:

- Start the OS from the Veeam Recovery Media and use Veeam Agent for Microsoft Windows and standard Microsoft Windows tools to diagnose and fix problems.
- Perform bare metal restore.
- Restore necessary data from backups to its original location or a new location.

Veeam Agent for Microsoft Windows integrates with Veeam Backup & Replication. Backup administrators who work with Veeam Backup & Replication can perform advanced tasks with Veeam Agent backups: perform data restore tasks with Veeam Agent backups, manage Veeam Agent backup jobs or backups created with these jobs.

Solution Architecture

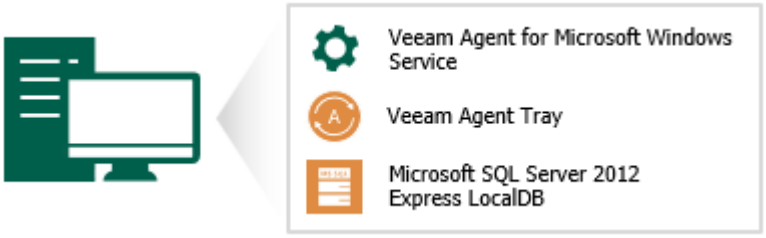
Veeam Agent for Microsoft Windows is set up on a computer whose data you want to protect.

Veeam Agent for Microsoft Windows has a one-service architecture. When you install the product, Veeam Agent deploys the following components on the computer:

- *Veeam Agent for Microsoft Windows Service* is a Microsoft Windows service responsible for performing all types of backup and restore tasks. The service is started automatically when you power on the computer, and runs in the background under the Local System account.
- *Veeam Agent Tray* is a tray agent that communicates with the Veeam Agent for Microsoft Windows Service to let you monitor the backup operation status and provide quick access to main functions of Veeam Agent for Microsoft Windows: starting backup and restore operations, viewing statistics for created backups and so on. The Veeam Agent Tray starts when you log on to the system and runs in the background.
- To store its configuration data, Veeam Agent uses the Microsoft SQL Server 2012 LocalDB Express. The LocalDB requires only few files to install and takes little resources to run a local on-demand Microsoft SQL Server instance. The LocalDB is executed as a subprocess launched by the Veeam Agent for Microsoft Windows Service. When the Veeam Agent for Microsoft Windows Service is stopped, the LocalDB subprocess is stopped, too.

NOTE

The account under which Veeam Agent for Microsoft Windows Service runs should not be changed. Configurations with custom account are not supported.



Standalone and Managed Operation Modes

Veeam Agent for Microsoft Windows can operate in two modes: *standalone mode* and *managed mode*.

- **Standalone mode** — in this mode, Veeam Agent for Microsoft Windows operates as a standalone product. To use Veeam Agent operating in the standalone mode, you must manually install the product directly on the computer whose data you want to protect.

  For Veeam Agent for Microsoft Windows operating in the standalone mode, data protection, disaster recovery, and administration tasks are performed in Veeam Agent by a user. You can also use Veeam Agent operating in the standalone mode with Veeam Backup & Replication. In this scenario, you can use a Veeam backup repository as a target location for Veeam Agent backups and use the Veeam Backup & Replication console to perform a number of tasks with Veeam Agent backup jobs and backups.

  The current User Guide covers subjects related to Veeam Agent for Microsoft Windows operating in the standalone mode only.

- **Managed mode** — in this mode, Veeam Agent for Microsoft Windows operates under control from Veeam Backup & Replication. To use Veeam Agent operating in the managed mode, you must deploy the product remotely from Veeam Backup & Replication.

  For Veeam Agent for Microsoft Windows operating in the managed mode, product deployment, data protection, disaster recovery, and administration tasks are performed by a backup administrator in the Veeam Backup & Replication console.

  Features and limitations of Veeam Agent for Microsoft Windows operating in the managed mode are different from those in the standalone mode. To learn about Veeam Agent operating in the managed mode, see the Veeam Agent Management Guide.

Data Backup

We recommend that you regularly back up data stored on your computer. Backup creates a safety copy of your data. If any kind of disaster strikes, you can restore your data from the backup and be sure that you will not lose the necessary information.

You can set up Veeam Agent for Microsoft Windows to perform automatic scheduled backups (triggered at specific time of the day or on specific events), or you can choose to back up data manually when needed. You can back up the entire computer image, specific computer volumes or individual folders with files.

Backups created with Veeam Agent for Microsoft Windows can be saved to one of the following locations:

- Removable storage device
- Local computer drive
- Network shared folder
- Backup repository managed by a Veeam backup server
- Cloud repository managed by a Veeam Cloud Connect service provider
- Microsoft OneDrive

Backup Types

Veeam Agent for Microsoft Windows lets you create the following backup types:

- Volume-level backup
- File-level backup

  In the file-level backup, you can select what personal data to back up. To learn more, see Personal Data Backup.

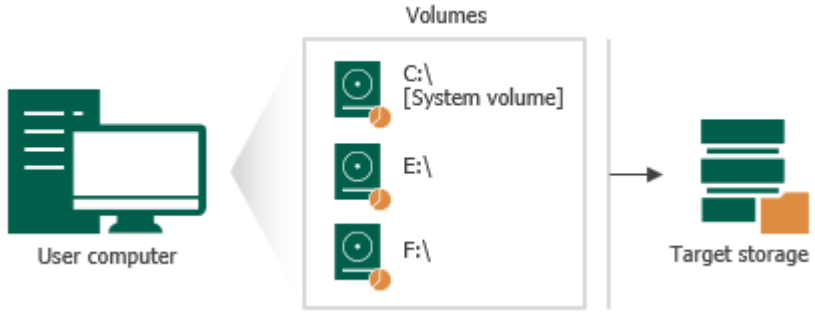
When backing up the operating system data, you can select the desired backup type. To learn more, see System State Data Backup.

Volume-Level Backup

You can set up Veeam Agent for Microsoft Windows to create a volume-level backup. The volume-level backup captures the whole image of a data volume (also called logical drive or partition) 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.

You can back up all computer volumes or specific computer volumes.

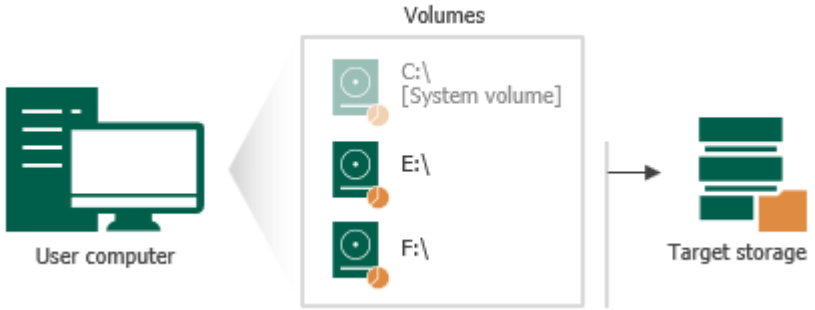
- When you back up the entire computer image, Veeam Agent for Microsoft Windows captures the content of all volumes on your computer. The resulting backup file contains all volume data and Microsoft Windows OS system data: system partition and boot partition. For GPT disks on Microsoft Windows 8.1, 10, 2012, 2012 R2, 2016 and 2019, Veeam Agent for Microsoft Windows additionally backs up the recovery partition.



- When you back up a specific computer volume, Veeam Agent for Microsoft Windows captures only that data that resides on this specific volume: files, folder, application data and so on.

  If you choose to back up the system volume (volume on which Microsoft Windows is installed), Veeam Agent for Microsoft Windows automatically includes the *System Reserved* partition into the backup scope. You can exclude the *System Reserved* partition from the backup if necessary. In this case, Veeam Agent for Microsoft Windows will capture only data on the system volume.

  To learn more, see System State Data Backup.

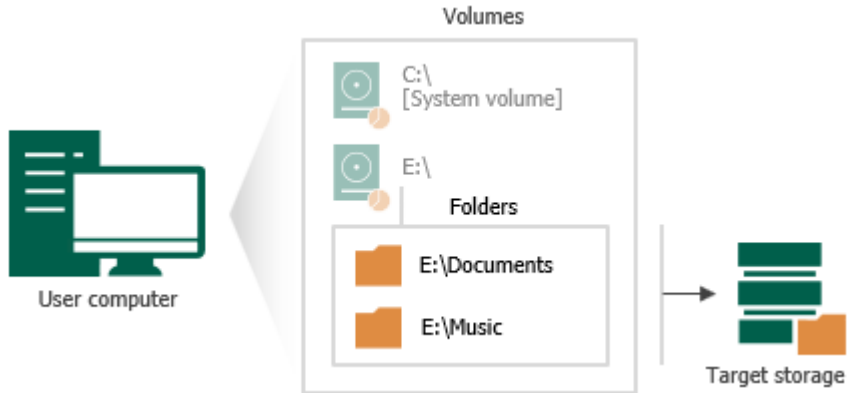


File-Level Backup

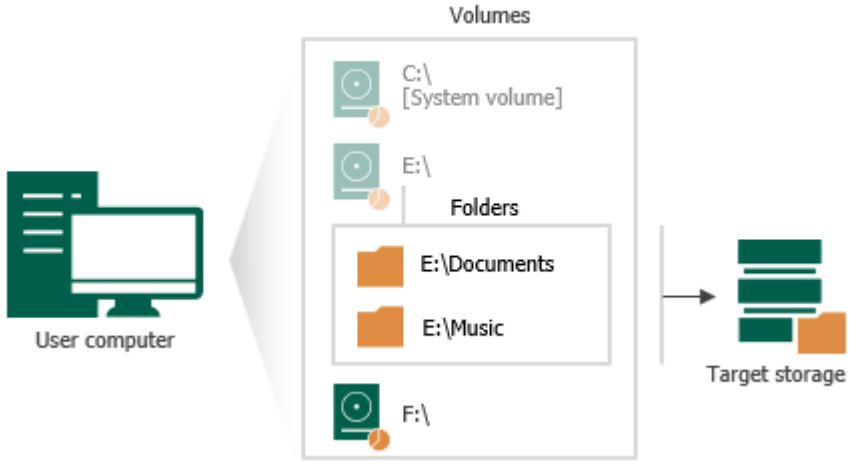
You can set up Veeam Agent for Microsoft Windows to create a file-level backup. The file-level backup captures only data of individual folders on the computer. You can use the file-level backup to restore files and folders that you have added to the backup scope.

Veeam Agent for Microsoft Windows lets you create two types of file-level backups:

- You can include individual folders into the backup. When you recover from such backup, you will be able to restore folders that you have selected to back up, and files in these folders.



- You can create a hybrid backup that will include folders and specific computer volumes. When you recover from such backup, you will be able to restore the following components:
  - For backed-up volume: the entire volume and individual files and folders on this volume.
  - For backed-up folders: folders that you have selected to back up, and files in these folders.

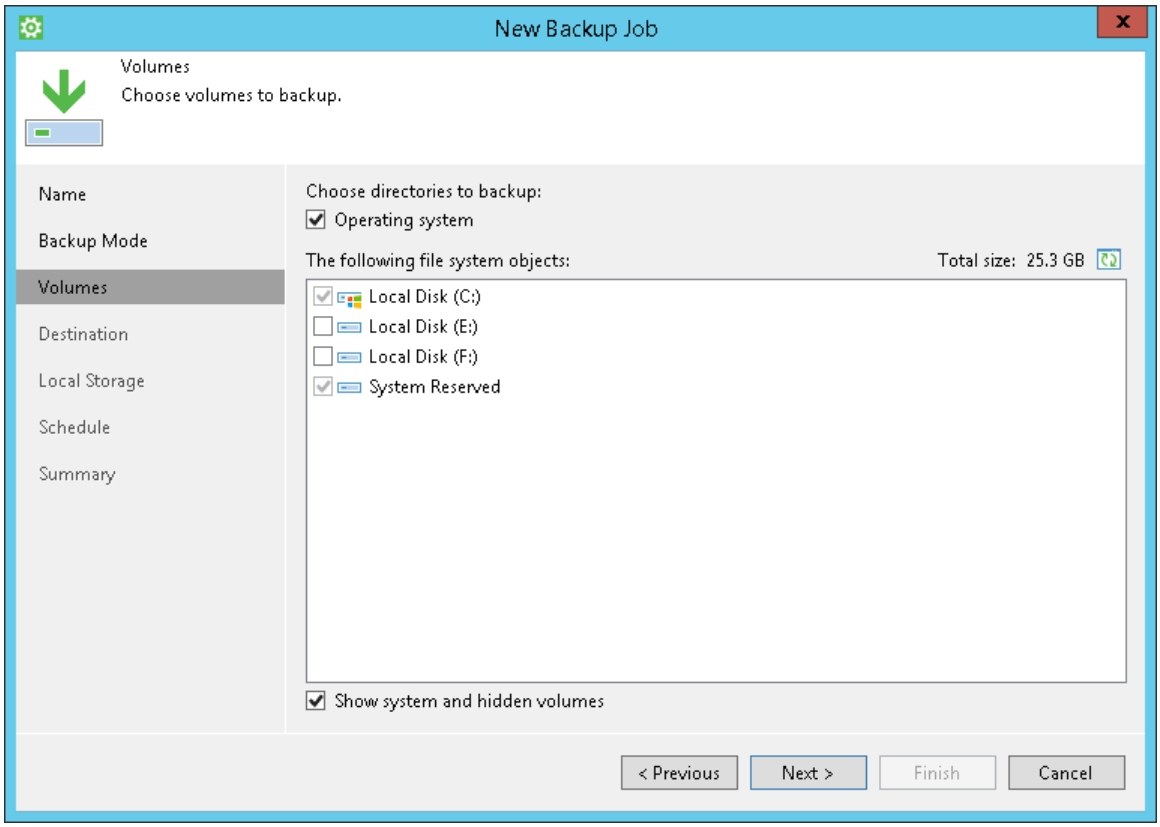


System State Data Backup

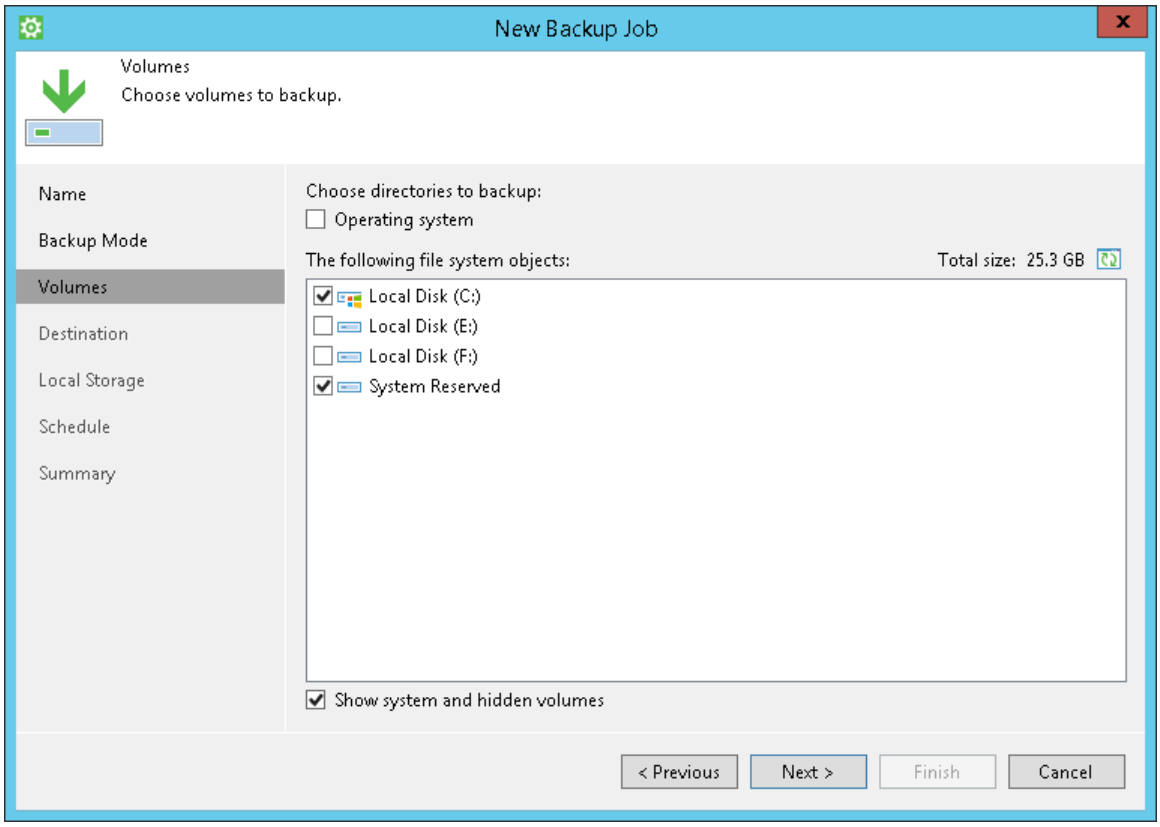
To be able to restore critical components related to the OS and start the OS after recovery, you must include in the backup the system volume (volume on which the OS is installed) and the System Reserved/UEFI or other system partitions. To simplify this process, Veeam Agent for Microsoft Windows offers you to add the *Operating system* component to the backup scope. When you select to back up the operating system data, Veeam Agent automatically includes in the backup all data related to the OS. This mechanism differs depending on the backup type: volume-level backup or file-level backup.

System State Backup in Volume-Level Backup

To include the operating system data in the volume-level backup, select *Operating system* data. When you select to back up the operating system data, Veeam Agent automatically includes in the backup the system volume and the System Reserved partition. If some additional system partitions appear on the Veeam Agent computer in the future, for example, after the OS upgrade, Veeam Agent will add these partitions to the backup scope, too.

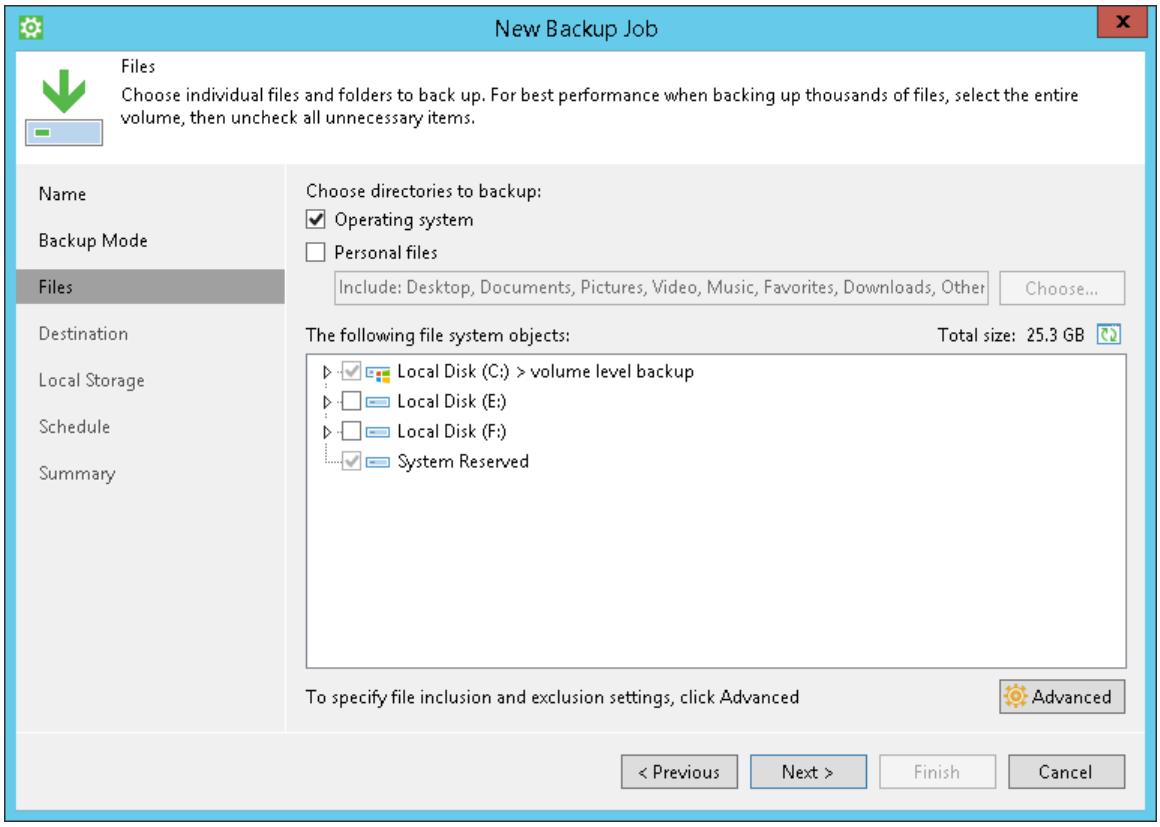


Alternatively, you can explicitly select to back up the system volume. When you select to back up the system volume, Veeam Agent automatically includes the System Reserved partition in the backup. However, if additional system partitions appear on the Veeam Agent computer, Veeam Agent may be unable to back up such volumes. Thus, we recommend that you use the *Operating system* option to create system state data backup.



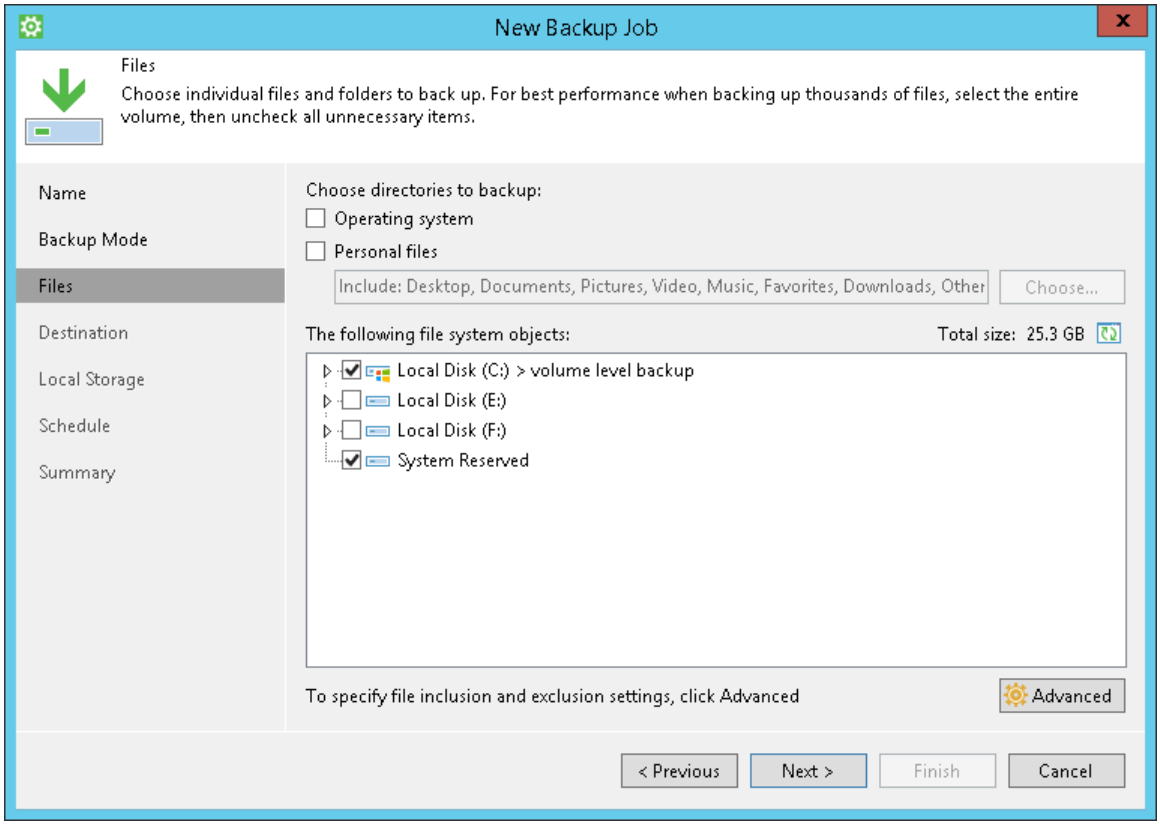
System State Backup in File-Level Backup

To include the operating system data in the file-level backup, select *Operating system* data. When you select to back up the operating system data, Veeam Agent automatically includes in the backup all data related to the OS: the system volume, personal files and the System Reserved partition.



Alternatively, you can select to back up the system volume and the System Reserved partition.

In this case, you will be able to exclude specific folders related to the OS from the backup (for example, the *Users* folder and *Documents* and *Settings* folder). When you select to back up the *Operating system* data, you cannot choose which components related to the OS must be backed up and which must be excluded.

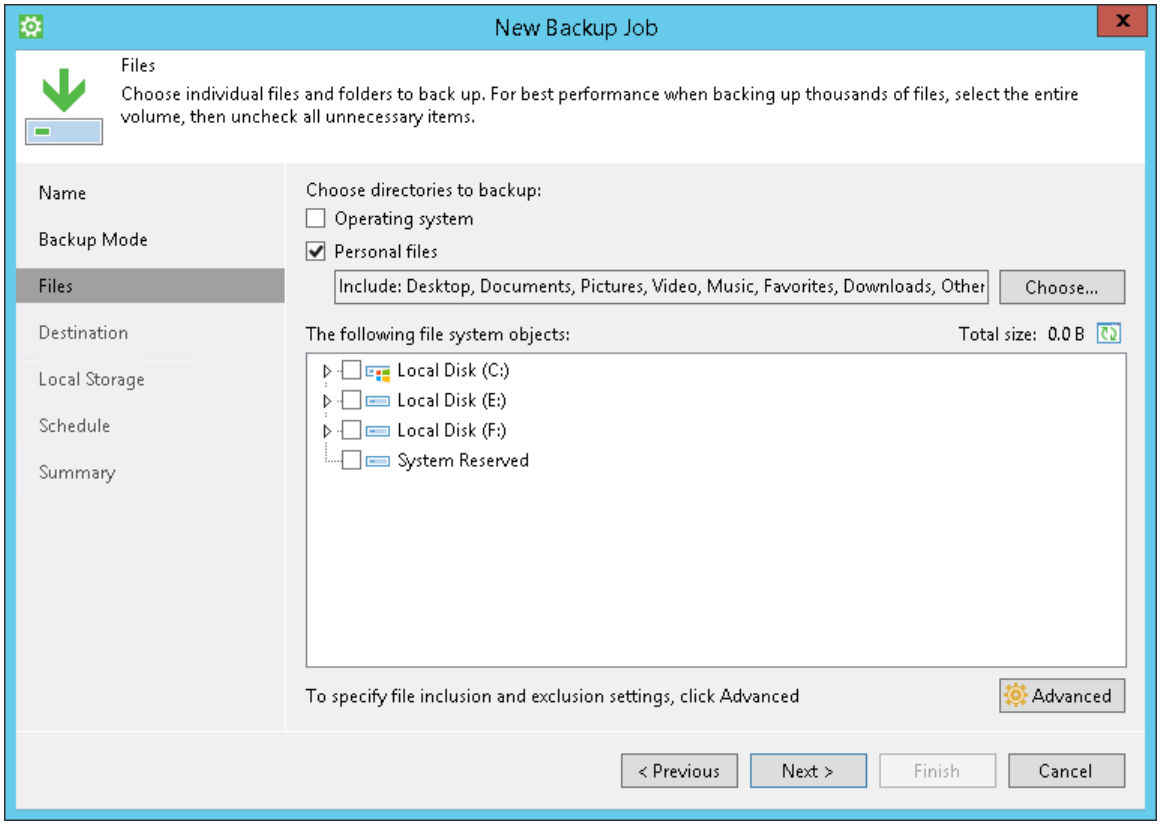


Personal Data Backup

You can include personal data in the backup. In the file-level backup, to simplify this process, Veeam Agent for Microsoft Windows offers you to add the *Personal files* component to the backup scope. When you select to back up the personal data, Veeam Agent automatically includes in the backup data related to Veeam Agent computer user profiles.

Personal Data Backup in File-Level Backup

To include personal data in the file-level backup, select the *Personal files* option. When you select to back up the personal data by default, Veeam Agent includes in the backup all data stored in the *Users* folder on the system volume excluding application data and data related to roaming user profiles. If you store personal folders of user profiles in custom locations, Veeam Agent backs up them too. If you do not want to back up default personal data, you can change the backup scope by selecting what folders Veeam Agent includes in the *Personal files* component.

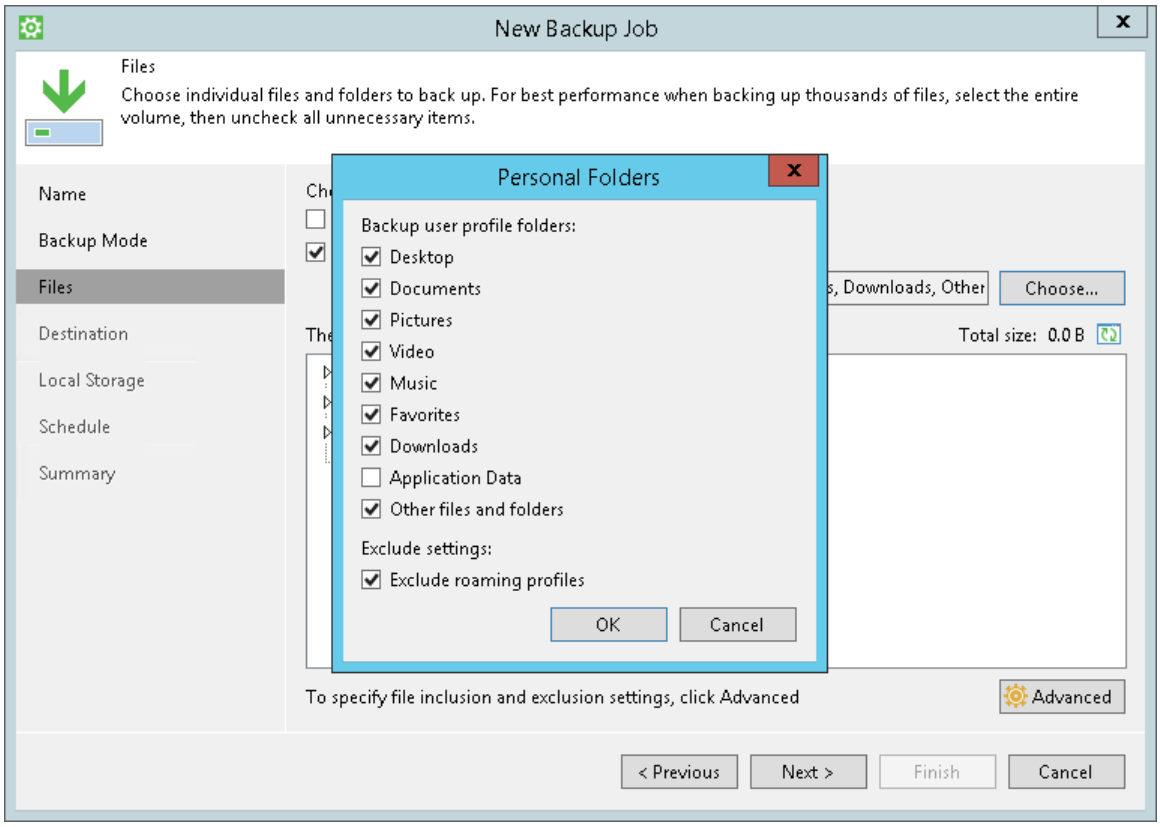


TIP

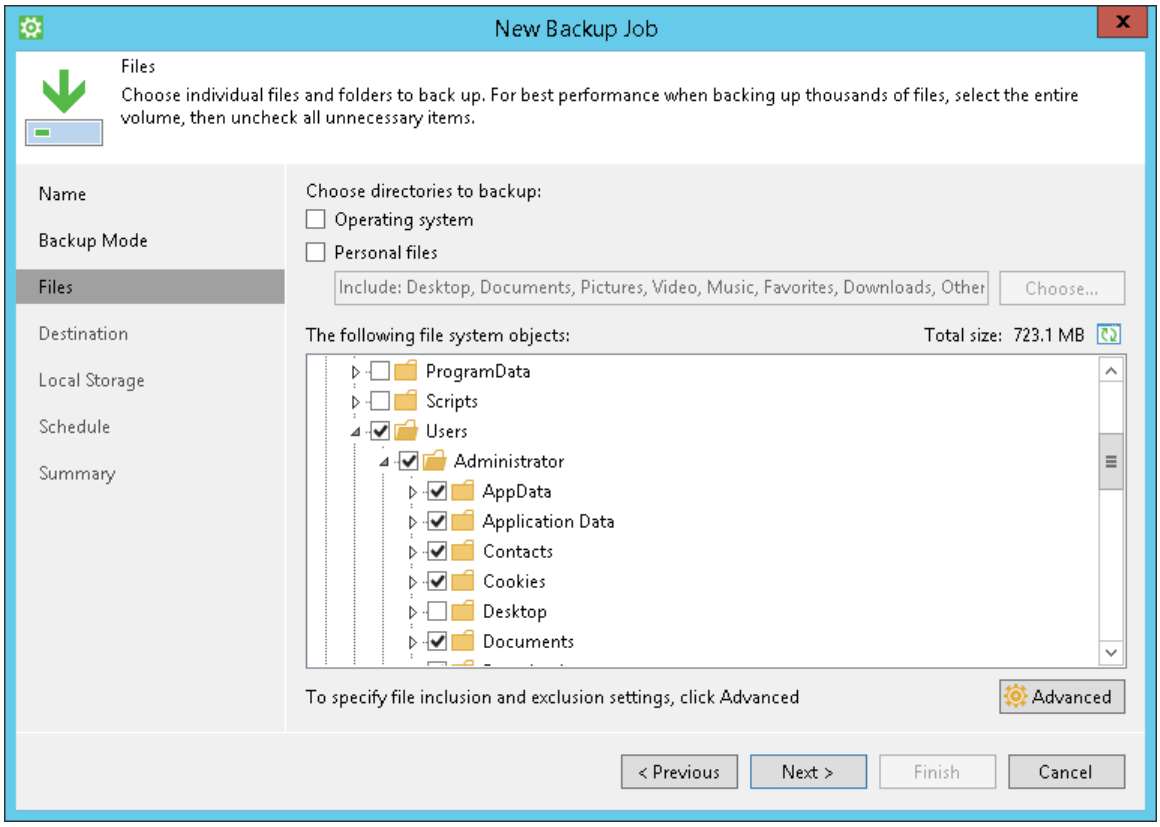
- By default, in the *Personal Folders* window, all options are selected in the **Backup user profile folders** list excluding the **Application Data** option and the **Exclude roaming profiles** option is selected.
- If you want Veeam Agent to back up personal data close to the way personal data was backed up in Veeam Agent 3.0 and 4.0, select all options in the **Backup user profile folders** list and clear the **Exclude roaming profiles** option. Note that in this case, the only difference is that Veeam Agent 5.0 also backs up selected folders of user profiles that are stored in custom locations.
- If Veeam Agent fails to back up personal data, Veeam Agent displays a warning message in the job session and a warning notification for each folder that Veeam Agent failed to back up. If you do not want to get these warnings, you can disable them with a registry value. For more information, contact Veeam Customer Support.

You can use the **Other files and folders** option to back up all folders and files that are located in the *Users* folder on the system volume, but are not available in the **Personal Folders** window. Mind that depending on the **Other files and folders** option, Veeam Agent behaves in one of the following ways:

- If the option is cleared, Veeam Agent backs up only selected personal folders in the *Users* folder on the system volume and personal folders stored in custom locations. Keep in mind, that in such case Veeam Agent backs up only those user profiles that are related to authenticated users. As a result, for example, default and system user profiles will be excluded. Besides that, temporary and corrupted user profiles will be excluded too.
- If the option is selected, Veeam Agent backs up all files and folders in the *Users* folder on the system volume and personal folders stored in custom locations, but excludes personal data according to the options that you cleared.

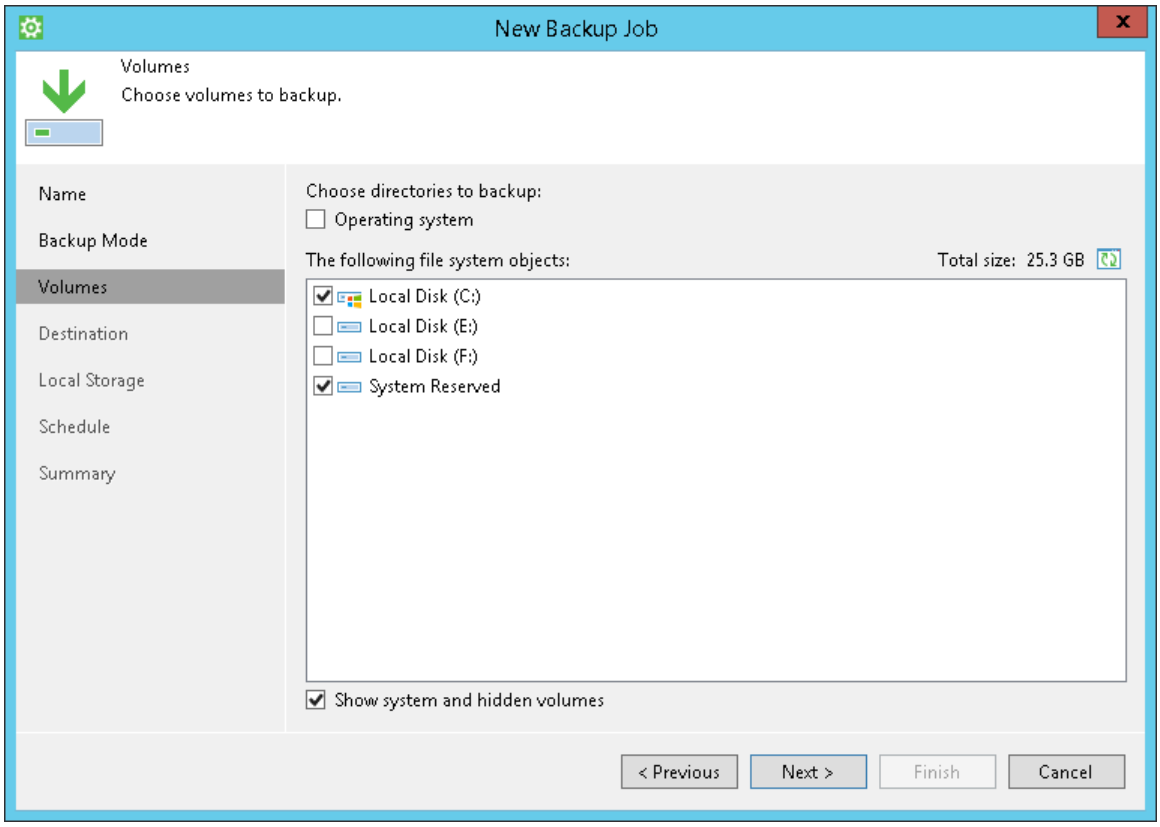


Alternatively, you can explicitly select to back up the *Users* folder. In this case, you still will be able to exclude specific subfolders of the *Users* folder from the backup. But you will need to exclude subfolders for each user manually. You will also need to find user folders stored in custom locations and include subfolders there if necessary. Thus, we recommend that you use the *Personal files* component to create personal data backup.



Personal Data Backup in Volume-Level Backup

Typically, data and settings of the Veeam Agent computer users are located in the Users folder on the system volume, for example, `C:\Users`. Thus, to include the operating system data in the volume-level backup, you must explicitly select to back up the system volume.



How Backup Works

During backup, Veeam Agent for Microsoft Windows performs the following operations:

1. Veeam Agent creates a Microsoft VSS snapshot of the volume whose data you want to back up. The VSS snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup.
   - In the Free and Workstation product editions, the VSS snapshot type depends on the OS of the Veeam Agent computer. On Microsoft Windows Client OSes, Veeam Agent creates a copy-only VSS snapshot. On Microsoft Windows Server OSes, Veeam Agent creates a full VSS snapshot.
   - In the Server product edition, the VSS snapshot type depends on the specified application-aware processing settings. To learn more, see Application-Aware Processing.

   Veeam Agent does not create a VSS snapshot for the EFI system partition on GPT disks, because its data does not change during backup. For the System Reserved and other system partitions, VSS snapshot can be created if there is enough free disk space on the partition.

   > **NOTE**
   >
   > By default, Microsoft Windows does not include offline Outlook Data Files (.ost) into a VSS snapshot. As a result, these files are not included into Veeam Agent backups, too.

2. Veeam Agent reads data from the created VSS snapshot, compresses it and copies it to the target location.
   - For volume-level backup, Veeam Agent copies data blocks of the whole volume.
   - For file-level backup, Veeam Agent creates a volume inside the backup file in the target location. The content of the volume in the backup file is synchronized with the volume on the source: Veeam Agent for Microsoft Windows copies only those data that you have selected to back up.

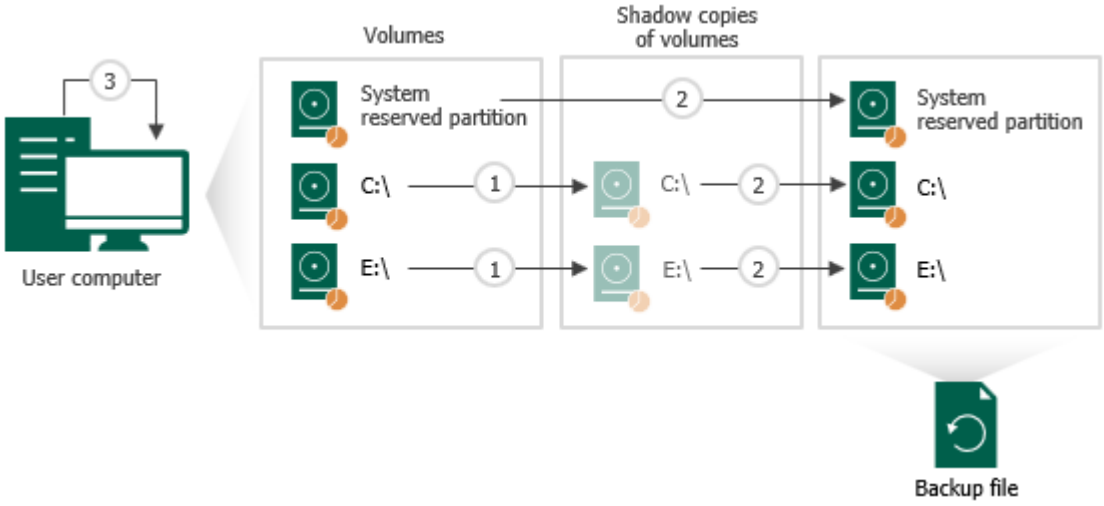
   During incremental backup, Veeam Agent uses Changed Block Tracking (CBT) to retrieve only those data blocks that have changed since the previous backup session. To learn more, see Change Block Tracking.

   In the target location, Veeam Agent stores copied data to the backup file.

3. If an application on the computer uses transaction logs to maintain the database consistency, Veeam Agent can truncate transaction logs upon successful backup.
    - In the Free and Workstation product editions, Veeam Agent truncates transaction logs depending on the OS of the Veeam Agent computer. On Microsoft Windows Client OSes, Veeam Agent does not truncate logs. On Microsoft Windows Server OSes, Veeam Agent always truncates logs.
    - In the Server product edition, Veeam Agent truncates logs depending on the specified application-aware processing settings. To learn more, see Application-Aware Processing.

IMPORTANT

The Veeam Agent Service runs under the LocalSystem account. On Microsoft SQL Server 2012, this account does not have necessary permissions to truncate transaction logs. If you want Veeam Agent to automatically truncate transaction logs, you need to manually add the LocalSystem account to a group that has the SQL Server System Administrator rights.



Backup Job

Veeam Agent for Microsoft Windows lets you configure a scheduled backup job that will perform backup automatically in a timely manner. You can set up the backup job once and forget about running the backup operation manually. Veeam Agent for Microsoft Windows will periodically launch the job to back up necessary data on your computer.

The backup job settings define what data you want to back up, what the target location and retention policy for created backups are and how often you want to back up your data. If necessary, you can re-configure the backup job at any time or remove the job.

Depending on the product edition, you can configure one or more backup jobs in Veeam Agent for Microsoft Windows:

- For the Free product edition, you can configure one backup job only.
- For the Workstation product edition, you can configure an unlimited number of backup jobs targeted at a Veeam Cloud Connect repository plus one backup job targeted at any other supported backup location (local or removable computer drive, network shared folder, Veeam backup repository or Microsoft OneDrive).
- For the Server product edition, you can configure an unlimited number of backup jobs targeted at any supported backup location.

For example, you can configure one backup job to create volume-level backup and another backup job to create file-level backup. Or you can configure backup jobs targeted at different backup locations to keep several copies of your backed-up data. You can also configure multiple backup jobs with individual schedule to fine-tune automatic backup creation process.

Settings of the backup job apply to ad-hoc backups as well: standalone full backups, active full backups and incremental backups.

Related Topic

Backup Job Schedule

Backup Job Schedule

Veeam Agent for Microsoft Windows launches the backup job according to the schedule you define. Scheduling options available for the backup job differ depending on the edition of Veeam Agent:

- For Free and Workstation product editions, you can schedule the job to start at specific time daily or on specific week days. You can also instruct Veeam Agent to automatically perform backup on specific events. To learn more, see Scheduling Options in Free and Workstation Editions.
- For the Server product edition, you can configure daily, monthly and periodic backup job schedule. You can also specify settings for automatic job retries and configure a backup window. To learn more, see Scheduling Options in Server Edition.

If you configure multiple backup jobs in Veeam Agent, you can specify individual schedule for each job. Veeam Agent processes backup jobs sequentially, one after another. If the backup job must start upon schedule at the time when another backup job is already running, the job will wait for the running job to complete, and start afterwards.

For portable devices, Veeam Agent does not start a backup job on the defined schedule if a device is working on battery and the battery level is below 20%.

If a backup job fails, Veeam Agent automatically retries the job. In Free and Workstation editions, Veeam Agent retries the job every 10 minutes within the next 23 hours. In the Server edition, you can specify retry settings along with other scheduling options. To learn more, see Automatic Job Retries and Job Retry.

In any edition, Veeam Agent can wake your computer from sleep at the time when the backup job must start. To learn more, see Computer Wake Up from Sleep.

Scheduling Options in Free and Workstation Editions

You can schedule the backup job to start at specific time daily or on specific week days. You can also instruct Veeam Agent for Microsoft Windows to automatically perform backup on specific events.

Missed Backup Schedule

Veeam Agent for Microsoft Windows does not perform scheduled backups if the computer is powered off. To handle situations of short power outage or computer restart, Veeam Agent for Microsoft Windows provides a tolerance window of 15 minutes for scheduled backups.

For example, you have configured the backup job to run daily at 10:00 PM. At 9:55 PM, there is a power outage that lasts for 10 minutes. When the computer is on again at 10:05, Veeam Agent for Microsoft Windows will automatically launch the scheduled job to back up your data.

Additionally, you can instruct Veeam Agent for Microsoft Windows to resume missed daily backup. If the computer is powered off at the time when the scheduled backup job must start, and you power on the computer later, Veeam Agent for Microsoft Windows will not wait for the next scheduled backup. Instead, Veeam Agent for Microsoft Windows will start the backup job right after the computer is powered on to ensure no necessary data is lost because of the missed backup.

Backup on Specific Events

In addition to the basic job schedule, you can instruct Veeam Agent for Microsoft Windows to launch the backup job on specific events. Veeam Agent for Microsoft Windows lets you trigger backup on the following events:

- Lock — the user locks the computer.
- Log off — the user performs a logout operation on the computer.
- When backup target is connected — the target backup location becomes available: the user attaches a known removable storage device to the computer or a network connection to the backup repository is established.

  You can instruct Veeam Agent for Microsoft Windows to eject the removable storage device after the backup job successfully completes. This helps to protect backup files in the target location from encrypting ransomware, such as CryptoLocker.

Backup on specific event helps you ensure that you capture all changes made within a specific time interval — for example, during a working day. When the necessary event occurs, Veeam Agent for Microsoft Windows automatically launches the scheduled backup job. As a result, you can be sure that all changes made within some period of time are backed up, and you do not lose your data.

If you choose to perform backup on specific events, you can restrict the frequency of backup job sessions. You can instruct Veeam Agent for Microsoft Windows not to start the backup job at specific events more often than once a specified time interval, for example, not more often than every 2 hours. This option does not affect daily schedule. Daily backups are performed according to the defined schedule regardless of the specified time interval.

Backup on specific events helps you fine-tune the backup job schedule. For example, you can specify the following scheduling settings for the backup job:

- The backup job must start automatically at 10:00 PM every day.
- The backup job must start at computer lock.
- The backup job must not run more often than every 2 hours.

Veeam Agent for Microsoft Windows will launch the backup job at the end of the working day, when you lock your computer. In addition, Veeam Agent for Microsoft Windows will perform backup at 10:00 PM regardless of the time interval between the computer lock and scheduled backup.

If you lock your computer later than at 10:00 PM, Veeam Agent for Microsoft Windows will perform backup in the following order. At 10:00 PM, Veeam Agent for Microsoft Windows will launch the backup job upon the daily schedule. If the time interval between the scheduled backup and computer lock is greater than 2 hours, Veeam Agent for Microsoft Windows will additionally perform backup at computer lock. If the time interval between the scheduled backup and computer lock is not greater than 2 hours, Veeam Agent for Microsoft Windows will not perform backup at computer lock.

Automatic Job Retries

Veeam Agent for Microsoft Windows supports automatic retries for the scheduled backup job. If the backup job is started on the defined daily schedule and fails for some reason, Veeam Agent automatically retries the job every 10 minutes within the next 23 hours.

If the backup job fails and Veeam Agent retries the job session, Veeam Agent does not transfer all the data again. Instead, Veeam Agent continues data transfer that was started before the backup job fail. To do so, Veeam Agent compares hash values for data blocks on source and target. After the hash comparison, Veeam Agent transfers only those data blocks that were not transferred before the job fail. If data blocks on source were changed before the automatic retry, Veeam Agent transfers these data blocks as well.

If you edit certain job configuration options before the automatic retry, Veeam Agent may require to transfer all the data to target again. These options include backup scope, data compression, storage optimization, data encryption. Mind that if you back up data to the Veeam backup repository and change data encryption options on the Veeam Backup & Replication side, Veeam Agent transfers all the data to the repository again.

Veeam Agent does not fail the backup job session in case Veeam Agent computer is put into sleep or hibernate mode during the job session, and the backup job is targeted at one of the following locations:

- Local computer drive
- Network shared folder
- Microsoft OneDrive

After the computer is turned on and the connection is restored, Veeam Agent operates in the same way as in case of automatic job retry. Veeam Agent does not transfer all the data again, but continues the job session from the point where it stopped.

Limitations for Automatic Job Retry

- Veeam Agent continues data transfer after the backup job fail only if you configured a job that creates volume-level backups.
- Veeam Agent does not continue data transfer after the backup job fail if the backup job is targeted at Microsoft OneDrive.
- Veeam Agent does not perform automatic retry for jobs started manually.
- Veeam Agent does not automatically retry the backup job if the job session is started when the computer is powered on after missed daily backup.
- Veeam Agent retries a job only if the previous job session has failed. Veeam Agent does not perform a retry if a job session has finished with the *Success* or *Warning* status.
- For portable devices, Veeam Agent does not automatically retry the backup job if a device is working on battery and the battery level is below 20%.

Scheduling Options in Server Edition

You can schedule the backup job to start automatically at specific time. Veeam Agent for Microsoft Windows lets you configure the following settings for the job:

- Scheduling settings
- Job retry settings
- Backup window settings

Automatic Startup Schedule

Veeam Agent for Microsoft Windows lets you configure the following scheduling settings for jobs:

- You can schedule the backup job to run at specific time every day or on selected days
- You can schedule the backup job to run periodically at specific time intervals
- You can schedule the backup job to run continuously

Job Started at Specific Time

You can schedule the backup job 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.



Job Started at Specific Time Intervals

You can schedule the backup job 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 Agent for Microsoft Windows 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.

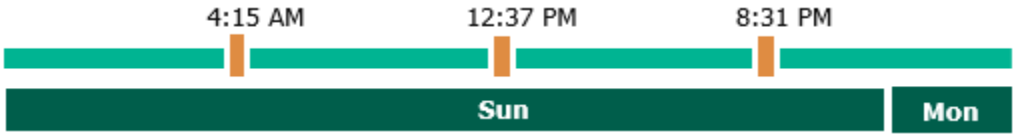


If a session of the periodically run job does not fit into the specified time interval and overlaps the next planned job session, Veeam Agent for Microsoft Windows starts the next backup job session at the nearest scheduled interval. For example, you set up the job to run with a 4-hour interval. The first job session starts at 12:00 AM, takes 5 hours and completes at 5:00 AM. In this case, Veeam Agent for Microsoft Windows will start a new job session at 8:00 AM.



Job Run Continuously

You can schedule the job to run continuously — that is, in a non-stop manner. A new session of the continuously running job starts as soon as the previous job session completes. Continuously run job can help you implement near-continuous data protection (near-CDP) for the most critical applications.



Job Retry

You can instruct Veeam Agent for Microsoft Windows to retry the backup job several times if the initial job session fails. By default, Veeam Agent for Microsoft Windows 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.

If a backup job fails and Veeam Agent retries the job session, Veeam Agent does not transfer all the data again. Instead, Veeam Agent continues data transfer that was started before the backup job fail. To do so, Veeam Agent compares hash values for data blocks on source and target. After the hash comparison, Veeam Agent also transfers only those data blocks that were not transferred before the job fail. If data blocks on source were changed before the retry, Veeam Agent transfers these data blocks as well.

If you edit certain job configuration options before the retry, Veeam Agent may require to transfer all the data to target again. These options include backup scope, data compression, storage optimization, data encryption. Mind that if you back up data to the Veeam backup repository and change data encryption options on the Veeam Backup & Replication side, Veeam Agent transfers all the data to the repository again.

Veeam Agent does not fail the backup job session in case Veeam Agent computer is put into sleep or hibernate mode during the job session, and the backup job is targeted at one of the following locations:

- Local computer drive
- Network shared folder
- Microsoft OneDrive

After the computer is turned on and the connection is restored, Veeam Agent operates in the same way as in case of job retry. Veeam Agent does not transfer all the data again, but continues the job session from the point where it stopped.

Limitations for Job Retry

- Veeam Agent continues data transfer after the backup job fail only if you configured a job that creates volume-level backups.
- Veeam Agent does not continue data transfer after the backup job fail if the backup job is targeted at Microsoft OneDrive.
- Veeam Agent does not perform retry for jobs started manually.
- Veeam Agent does not retry the backup job if the job session is started when the computer is powered on after missed daily backup.
- Veeam Agent retries a job only if the previous job session has failed. Veeam Agent does not perform a retry if a job session has finished with the *Success* or *Warning* status.
- For portable devices, Veeam Agent does not retry the backup job if a device is working on battery and the battery level is below 20%.

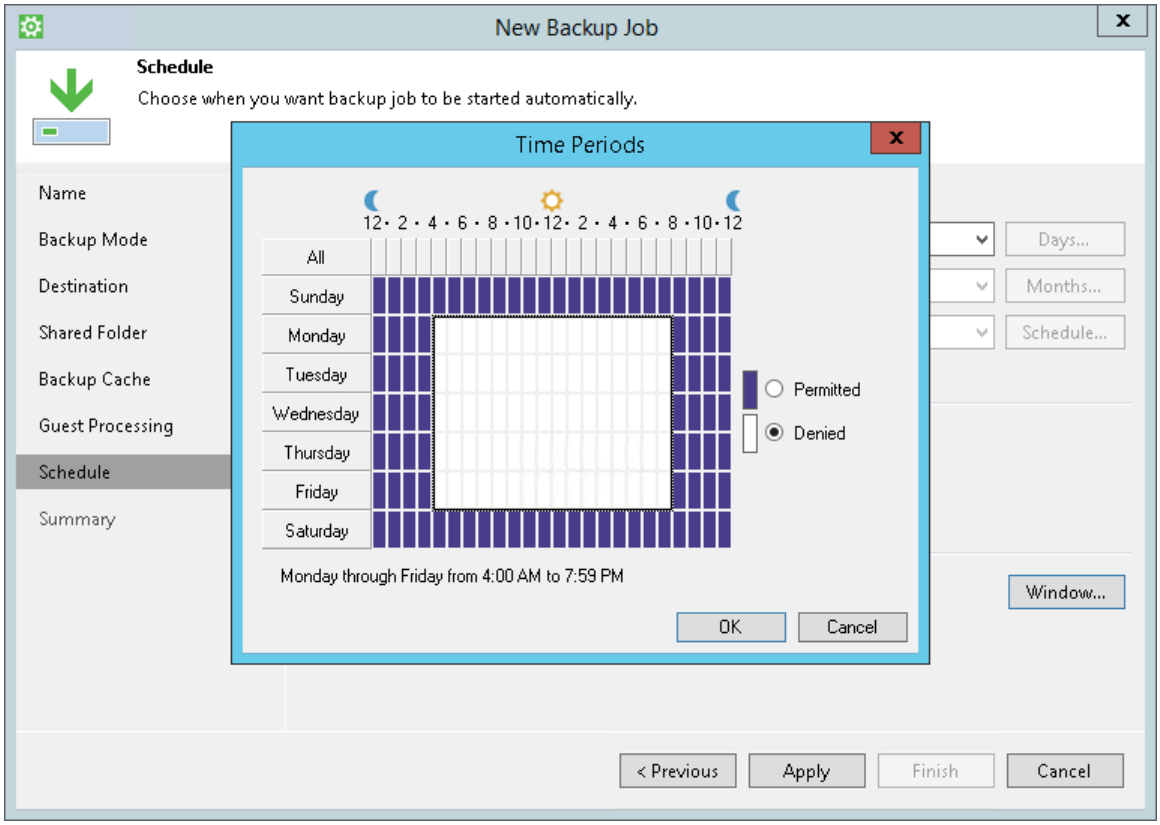
Backup Window

If necessary, you can specify a backup window for the backup job. The backup window is a period of time on week days when the job is permitted to run.

The backup window can be helpful if you do not want the data protection job to produce unwanted overhead for the production environment or do not want the job to overlap production hours. In this case, you can define the time interval during which the job must not run.

If the job exceeds the allowed window, it will be automatically terminated. In this case, data transport and backup chain transformation processes are stopped.

> **IMPORTANT**
>
> The backup window does not affect the process of uploading backup files from the backup cache to the target storage. If Veeam Agent creates one or more backup files in the backup cache, and then the backup target becomes available, Veeam Agent uploads backup files to the target location immediately, regardless of the specified backup window.

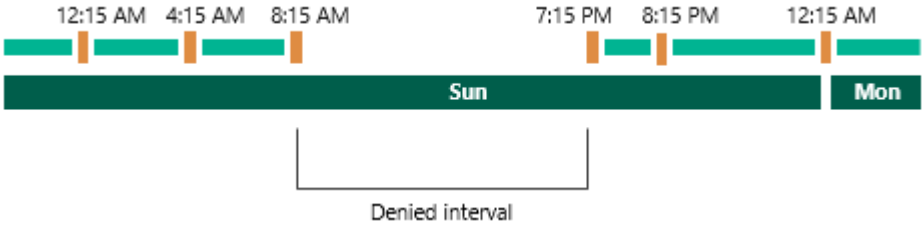


Backup Window for Periodically Run Job

If you define the backup window for the job that runs periodically at specific time intervals, Veeam Agent 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 the 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 Agent 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, because 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 Agent for Microsoft Windows will run the job by the defined schedule: at 8:15 PM, 12:15 AM and so on.



Computer Wake Up from Sleep

If your computer is in the standby mode at the time when the backup job must start, Veeam Agent for Microsoft Windows automatically wakes your computer from sleep. The wake-up feature lets you schedule your backup at night. At the defined time, Veeam Agent for Microsoft Windows will wake up the computer and perform a scheduled task. If necessary, you can additionally instruct Veeam Agent for Microsoft Windows to bring the computer back to the standby mode or power off the computer after the backup is finished.

Veeam Agent for Microsoft Windows wakes up the computer by default, unless the power saving settings on the computer prohibit this. If the wake up operation is not possible for some reason, the computer will remain in the standby mode, and the backup operation will not be performed. You can instruct Veeam Agent for Microsoft Windows to resume missed backup in such situations. To learn more, see Missed Backup Schedule.

> **IMPORTANT**
>
> [For tablets running Microsoft Windows 8.1 or later] If at the moment of backup a computer is in the Connected Standby power saving mode, Veeam Agent for Microsoft Windows will fail to wake it up due to limitations set by the OS itself.

Ad-Hoc Backup

You can create ad-hoc backups of your data when you need.

Ad-hoc backups let you capture your data at a specific point in time. You can create ad-hoc backups before you perform some alterations on your computer: install new software or enable a new feature. Ad-hoc backups help you protect your computer from potential data corruption or data loss that can be caused by these operations. If an error occurs, you can always restore data from the ad-hoc backup and bring your computer system to a state before the alteration was made.

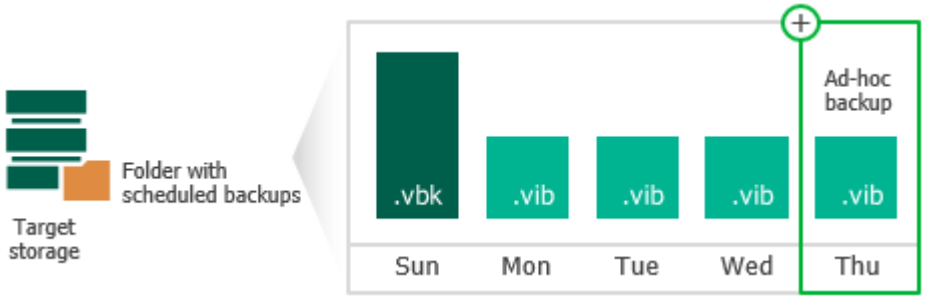
Veeam Agent for Microsoft Windows lets you create the following types of ad-hoc backups:

- Incremental backup
- Standalone full backup

Ad-Hoc Incremental Backup

If you want to create a new backup of your data in addition to backups created with the scheduled backup job, you can perform ad-hoc incremental backup. Ad-hoc incremental backup adds a new restore point to the backup chain. For example, you may want to back up your data before you install new software on your computer or enable a new feature.

For ad-hoc incremental backup, Veeam Agent for Microsoft Windows uses settings specified for the backup job. For example, if you have configured the backup job to perform backup of a specific volume, the ad-hoc incremental backup operation will create an incremental backup of this volume and save it in the target location, next to existing backup files in the backup chain.



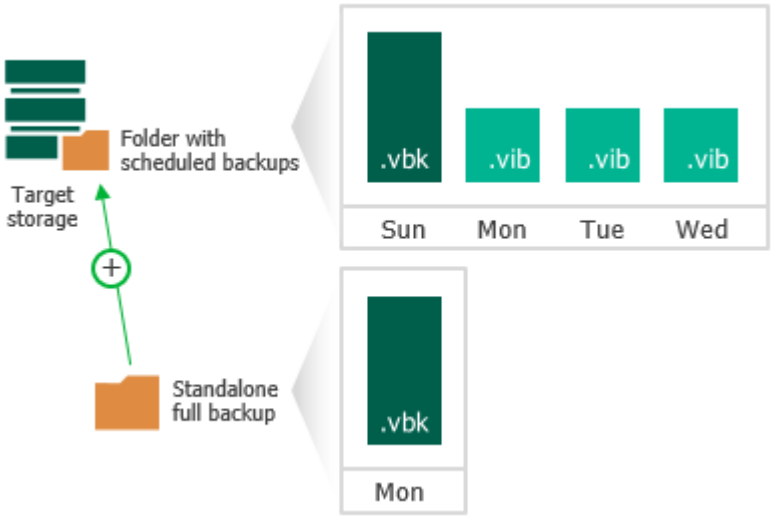
Unlike the backup job started upon schedule, the ad-hoc incremental backup task is not retried automatically. If the task fails for some reason, you will have to start it manually again.

Veeam Agent for Microsoft Windows treats restore points created by ad-hoc incremental backup as regular restore points, and applies to them retention policy settings specified for the backup job. To learn more, see Backup Retention Policy.

Standalone Full Backup

Sometimes you need to create a full backup of your data. For example, you may want to save a copy of your data on a CD or DVD or create a full backup of all data on your computer at some point in time. In these situations, you can perform standalone full backup.

When Veeam Agent for Microsoft Windows performs standalone full backup, it produces a full backup of your data in a separate folder in the target location. The standalone full backup is not associated with subsequent incremental backups. You can use it as an independent restore point for data recovery.



To create a standalone full backup, Veeam Agent for Microsoft Windows uses settings specified for the backup job. For example, if you have configured the backup job to perform backup of a specific volume, the standalone full backup will create a full backup of this volume in a separate folder in the target location.

Unlike the backup job started upon schedule, the standalone full backup task is not retried automatically. If standalone full backup fails for some reason, you will have to start the standalone full backup task manually again.

The standalone full backup is not removed by retention. To remove it, you must manually delete the full backup file from disk.

> **NOTE**
>
> You cannot perform standalone full backup if the backup job is targeted at a Veeam Cloud Connect repository or Microsoft OneDrive.

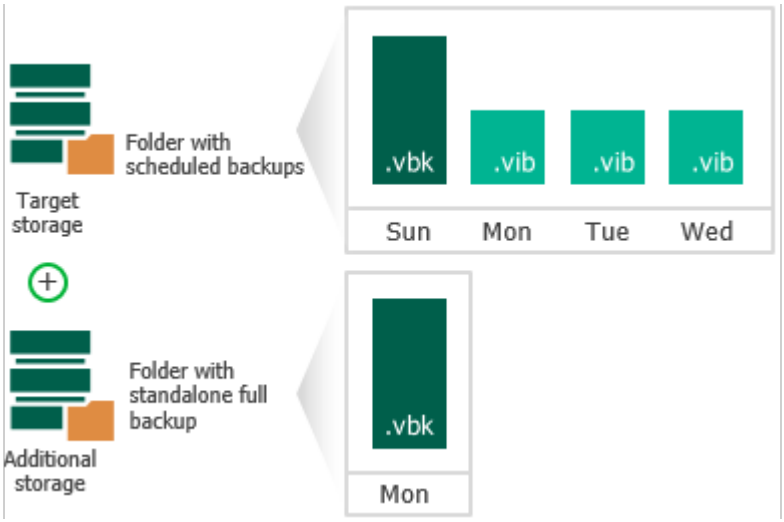
Standalone Full Backup to Another Location

You can create a standalone full backup in a separate location that is not specified as a target location in the backup job settings. For example, you may want to save a copy of your data on a removable storage device while your scheduled backup job is targeted at the network shared folder.

Backup to another location practically does not differ from regular standalone full backup. The only difference is that you must manually select a target location in which Veeam Agent for Microsoft Windows will save the backup file. You can save backup files to the following locations:

- Removable storage device
- Local computer drive
- Network shared folder

You cannot use a Veeam backup repository, Veeam Cloud Connect repository or Microsoft OneDrive storage as a target for backup to another location.



Parallel Disk Processing

Veeam Agent for Microsoft Windows supports parallel disk processing. If you included several disks of the Veeam Agent computer in the backup scope, Veeam Agent will process disks simultaneously.

How Parallel Disk Processing Works

During backup, Veeam Agent for Microsoft Windows creates a Microsoft VSS snapshot of the volume whose data you want to back up as described in How Backup Works. After the VSS snapshot is created, Veeam Agent transfers data blocks of disks to the target storage simultaneously.

The number of disks that are simultaneously processed by Veeam Agent is limited by repository settings on the Veeam Backup & Replication side. During the backup job session, Veeam Backup & Replication compares the number of currently processed disks to the limit specified in the properties of the backup repository. If the number of currently processed disks is greater than the limit, Veeam Agent will not transfer data of a new disk. Only after data transfer for at least one of the disks completes, Veeam Agent starts data transfer for the next disk. To learn more, see the Limitation of Concurrent Tasks section in the Veeam Backup & Replication User Guide.

You can also limit the number of disks that are processed simultaneously with a registry value. Keep in mind that if the number of simultaneously processed disks is limited by both repository settings and registry value, Veeam Agent uses the lower limit. For the registry value, see this Veeam KB article.

If processing of at least one disk fails, the entire backup job session completes with an error.

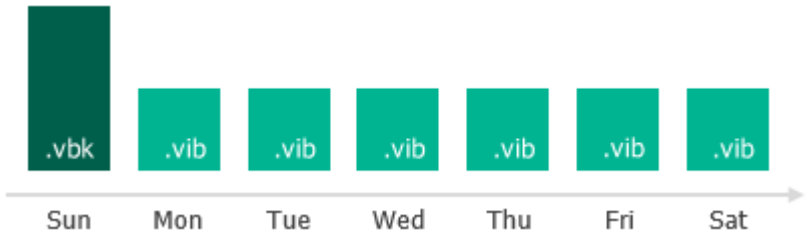
Limitations for Parallel Disk Processing

- Parallel disk processing is available only in the Server edition of Veeam Agent for Microsoft Windows.
- Veeam Agent offers parallel disk processing for backups that are created in the following types of target locations:
    - Veeam backup repository
    - Veeam Cloud Connect repository
- Parallel disk processing is not supported for dynamic disks.
- Parallel disk processing is not performed for file-level backups.

Backup Chain

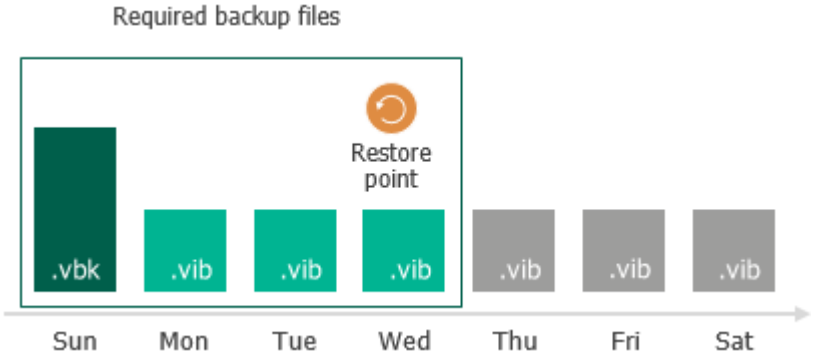
Every backup job session produces a new backup file in the target location. Backup files make up a backup chain. The backup chain can contain files of two types: full backup(s) and incremental backups.

- During the first backup job session, Veeam Agent performs full backup. Veeam Agent copies all data that you have chosen to back up (entire volumes and folders) and stores the resulting full backup file (VBK) in the target location. The full backup takes significant time to complete and produces a large backup file: you have to copy the whole amount of data.
- During subsequent backup job sessions, Veeam Agent performs incremental backups. It copies only new or changed data relatively to the last backup job session and saves this data as an incremental backup file (VIB) in the target location. Incremental backups typically take less time than full backup: you have to copy only changes, not the whole amount of data.



After several backup cycles, you have a chain of backup files in the target location: the first full backup file and subsequent incremental backup files. Every backup file contains a restore point for backed-up data. A restore point is a "snapshot" of your data at a specific point in time. You can use restore points to roll back your data to the necessary state.

To recover data to a specific restore point, you need a chain of backup files: a full backup file plus a set of incremental backup files following this full backup file. If some file from the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, we recommend that you do not delete separate backup files manually. To learn more, see Deleting Backups.



Types of Backup Files

Veeam Agent for Microsoft Windows produces backup files of the following types:

- VBK — full backup file.
- VIB — incremental backup file.
- VBM — backup metadata file. The backup metadata file is updated with every backup job session. It contains information about the computer on which the backup was created, every restore point in the backup chain, how restore points are linked to each other and so on. The backup metadata file is required for performing file-level and volume-level restore operations.

NOTE

For backup jobs with database log backup options enabled, Veeam Agent for Microsoft Windows additionally produces backup files of the following types:

- VLB, VSM and VLM files — for Microsoft SQL Server transaction log backups
- VLB, VOM and VLM files — for Oracle archived log backups

Keep in mind that Veeam Agent creates VLM files only in backup chains started in Veeam Agent for Microsoft Windows 5.0 and later. If you upgrade Veeam Agent and continue the backup chain that was started in the earlier Veeam Agent version, Veeam Agent will not create VLM files in this backup chain.

Short-Term Retention Policy

Restore points in the backup chain are not kept forever. They are removed according to the retention policy. The short-term retention policy helps maintain the life cycle of restore points and make sure that backup files do not consume the whole disk space.

Short-term retention policy depends on the edition of Veeam Agent for Microsoft Windows:

- In Free and Workstation editions, Veeam Agent retains restore points for a certain number of days. To learn more, see Retention Policy in Free and Workstation Editions.
- In the Server edition, you can select the short-term retention policy. Veeam Agent can retain restore points for a certain number of days or retains a certain number of latest restore points that is defined by the user. To learn more, see Retention Policy in Server Edition.

You can also specify the retention policy for outdated backups — backups for which Veeam Agent does not create new restore points within a specified time period. To learn more, see Retention Policy for Outdated Backups.

Retention Policy in Free and Workstation Editions

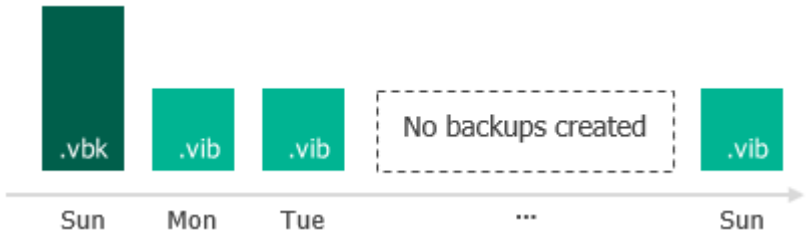
In Free and Workstation editions, Veeam Agent for Microsoft Windows retains restore points for the last N days; the number of days is defined by the user. During every backup job session, Veeam Agent for Microsoft Windows checks if there is any obsolete restore point in the backup chain. If some restore point is obsolete, it is removed from the chain.

For retention policy settings, Veeam Agent for Microsoft Windows takes into account not calendar days but days on which backup files were successfully created.

For example, you have configured the backup job in the following way:

- The backup job runs daily.
- The retention policy is set to 5 days.

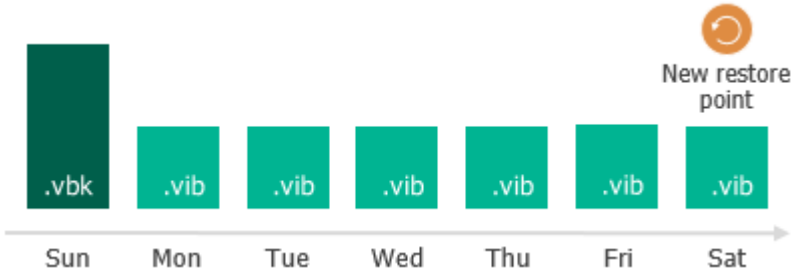
The backup job has successfully run 3 times and created 3 restore points in the backup chain. After that, you have turned off your computer for 10 days. When you turn on your computer, Veeam Agent for Microsoft Windows runs a backup job by schedule and creates a new restore point. The earliest restore point, however, is not removed from the backup chain. At the end of a new backup job session, the backup chain will have only 4 restore points created during 4 days when the backup job was successfully run.



Removing Backups by Retention

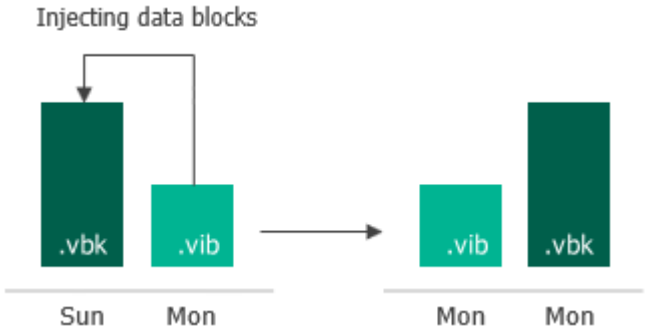
When removing obsolete restore points, Veeam Agent for Microsoft Windows does not simply delete backup files from disk. It transforms the backup chain so that the backup chain always contains a full backup file on which subsequent incremental backup files are dependent. To maintain the consistency of the backup chain, Veeam Agent for Microsoft Windows uses the following rotation scheme:

1. During every backup job session Veeam Agent for Microsoft Windows adds a backup file to the backup chain and checks if there is an obsolete restore point.

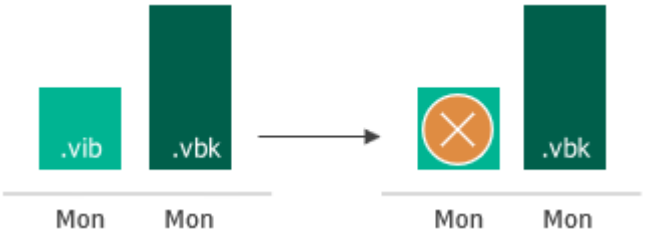


2. If an obsolete restore point exists, Veeam Agent for Microsoft Windows transforms the backup chain. As part of this process, it performs the following operations:

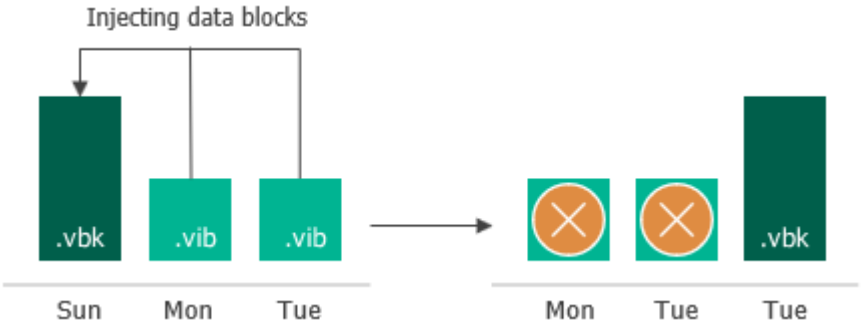
   a. Veeam Agent for Microsoft Windows re-builds the full backup file to include in it data of the incremental backup file that follows the full backup file. To do this, Veeam Agent for Microsoft Windows injects into the full backup file data blocks from the earliest incremental backup file in the chain. This way, a full backup 'moves' forward in the backup chain.



   b. The earliest incremental backup file is removed from the chain as redundant: its data has already been injected into the full backup file, and the full backup file includes data of this incremental backup file.



If the backup chain contains several obsolete restore points, the rebuild procedure is similar. Data from several restore points is injected to the re-built full backup file. This way, Veeam Agent for Microsoft Windows makes sure that the backup chain is not broken, and you will be able to recover your data to any restore point.



Retention Policy in Server Edition

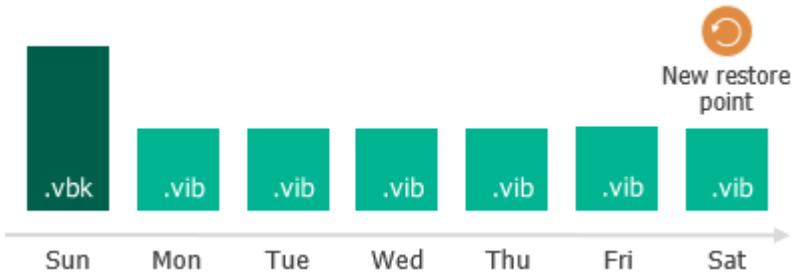
In the Server edition, you can select the logic behind the short-term retention policy. Veeam Agent offers the following options:

- Veeam Agent retains restore points for the last N days; the number of days is defined by the user. In this case, Veeam Agent works in a similar way as in the Free or Workstation editions. To learn more, see Retention Policy in Free and Workstation Editions.
- Veeam Agent retains the number of latest restore points defined by the user.

During every backup job session, Veeam Agent checks if there is any obsolete restore point in the backup chain. If some restore point is obsolete, it is removed from the chain.

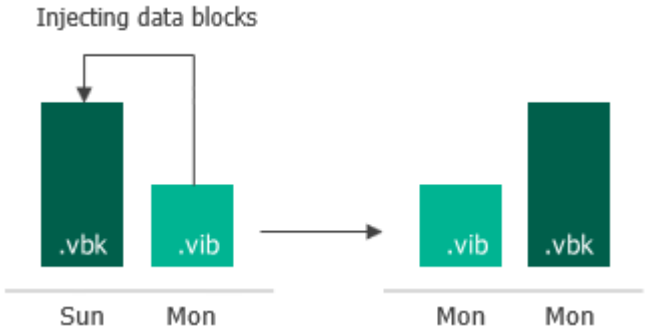
When removing obsolete restore points, Veeam Agent for Microsoft Windows does not simply delete backup files from disk. It transforms the backup chain so that the backup chain always contains a full backup file on which subsequent incremental backup files are dependent. To maintain the consistency of the backup chain, Veeam Agent for Microsoft Windows uses the following rotation scheme:

1. During every backup job session Veeam Agent for Microsoft Windows adds a backup file to the backup chain and checks if there is an obsolete restore point.

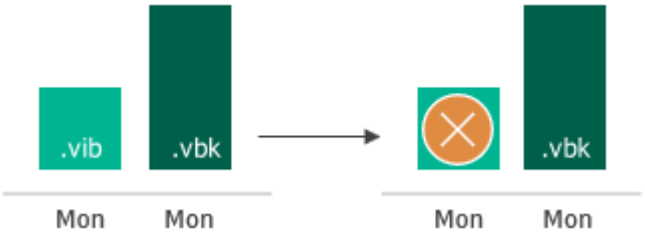


2. If an obsolete restore point exists, Veeam Agent for Microsoft Windows transforms the backup chain. As part of this process, it performs the following operations:

   a. Veeam Agent for Microsoft Windows re-builds the full backup file to include in it data of the incremental backup file that follows the full backup file. To do this, Veeam Agent for Microsoft Windows injects into the full backup file data blocks from the earliest incremental backup file in the chain. This way, a full backup 'moves' forward in the backup chain.



   b. The earliest incremental backup file is removed from the chain as redundant: its data has already been injected into the full backup file, and the full backup file includes data of this incremental backup file.



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.

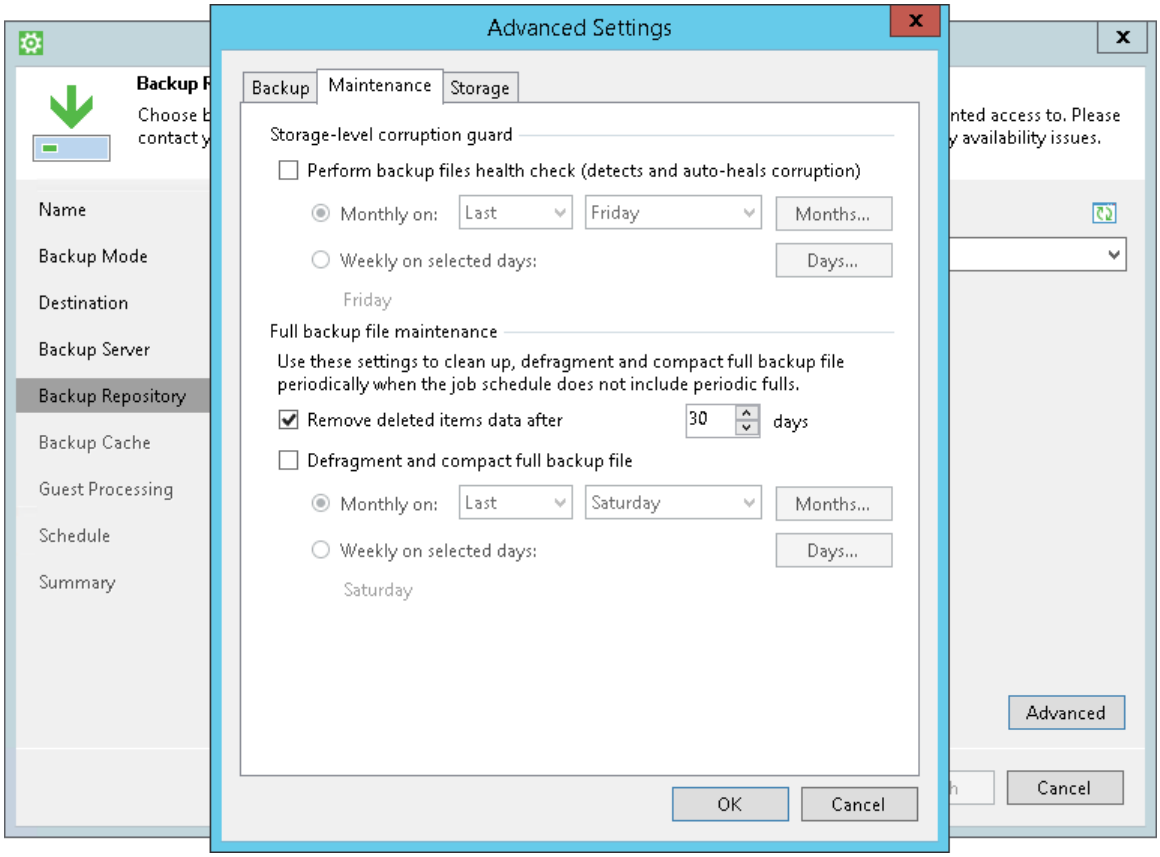
Veeam Agent for Microsoft Windows offers retention policy for outdated backups that are created in the following types of target locations:

- Veeam backup repository
- Veeam Cloud Connect repository

To specify retention policy for outdated backups, you must enable the **Remove deleted items data after <N> days** option in backup job settings and specify the number of days for which data of outdated backups must be retained on the backup repository. By default, Veeam Agent for Microsoft Windows retains outdated backups for 30 days. You can change the number of days to retain outdated backups if necessary.

Mind the following:

- You must use retention policy for outdated backups wisely. We strongly recommend that you set retention policy for outdated backups to 7 days or more to prevent unwanted data loss.
- The **Remove deleted items after <N> days** option lets you control data of backups for which Veeam Agent for Microsoft Windows does not produce new restore points for some time. In addition to it, Veeam Agent for Microsoft Windows applies general retention policy rules to maintain the necessary number of restore points in the backup chain. To learn more, see Backup Retention Policy.



How Retention Policy for Outdated Backups Works

Although you specify retention policy settings for outdated backups in Veeam Agent for Microsoft Windows, actions required to track and delete outdated Veeam Agent backups are taken on the Veeam Backup & Replication side.

To perform necessary actions, Veeam Backup & Replication uses the Veeam Agent backup retention job that runs on the backup server. The Veeam Agent backup retention job runs upon schedule, at 12:30 AM daily.

> **NOTE**
>
> You can change the schedule for the Veeam Agent backup retention job with registry values. For more information, contact Veeam Customer Support.

During the Veeam Agent backup retention job session, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication checks the configuration database to detect Veeam Agent backups that reside on the backup repository.
2. For each Veeam Agent backup, Veeam Backup & Replication checks whether the following conditions are met:
   a. The configuration database contains information about the backup job that created the Veeam Agent backup.

      The configuration database does not contain information about the backup job that created the backup in the following cases:

      - The backup job that created the backup was removed in Veeam Agent for Microsoft Windows.
      - The Veeam Agent backup was imported in the Veeam Backup & Replication console.
      - The Veeam Agent backup is a standalone full backup.

      If the configuration database does not contain information about the backup job that created the backup, Veeam Backup & Replication does not remove such a backup from the backup repository.

   b. The **Remove deleted items after <N> days** option is enabled in the backup job settings.

   c. No new restore points were created by the backup job for the last N days.

      Where N is the number of days specified for the **Remove deleted items data after <N> days** option.

   d. No new backup job sessions were detected by Veeam Backup & Replication for the last N days.

      Where N is the number of days specified for the **Remove deleted items data after <N> days** option.

   e. [For backup jobs with database log backup enabled] No database log copies (Microsoft SQL Server transaction log or Oracle archive log copies) were created by the database log backup job for the last N days.

      Where N is the number of days specified for the **Remove deleted items data after <N> days** option.
3. If all the conditions listed in the step 2 are met, Veeam Backup & Replication removes the Veeam Agent backup from the configuration database and deletes actual backup files from the backup repository.

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 Agent for Microsoft Windows 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 Agent 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 Agent assigns depend on the configured GFS retention policy settings.

> **NOTE**
>
> GFS flags can be assigned only to full backup files created during the time period specified in GFS policy settings.

If Veeam Agent assigns a GFS flag to a full backup file, this backup file can no longer be deleted or modified. Veeam Agent does not apply short-term retention policy settings to the full backup file. For example, Veeam Agent ignores the backup file when determining whether the number of allowed backup files is exceeded.

When the specified retention period ends, Veeam Agent 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.

Veeam Agent assigns GFS flags in the similar way as Veeam Backup & Replication does for VM backup files. To learn about logic behind GFS flags, see the Assignment of GFS Flags and Removal of GFS Flags sections in the Veeam Backup & Replication User Guide.

Limitations

When planning to use GFS retention policy, mind the following limitations:

- As Veeam Agent 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, the following applies:
  - Veeam Agent cannot transform the backup chain according to the short-term retention policy.
  - Veeam Agent is not able to merge data from incremental backup files into the full backup file.
- Veeam Agent 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 use GFS retention policy for a backup job targeted at Microsoft OneDrive.
- 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.

Related Topics

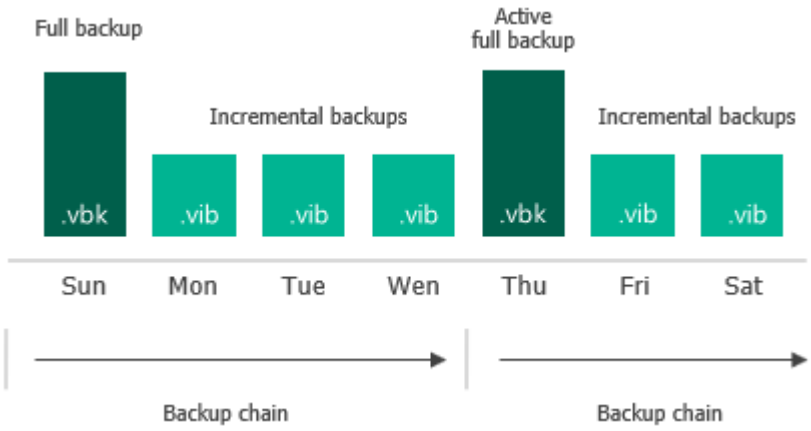
GFS Retention Policy Settings

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 Agent for Microsoft Windows lets you create active full backups.

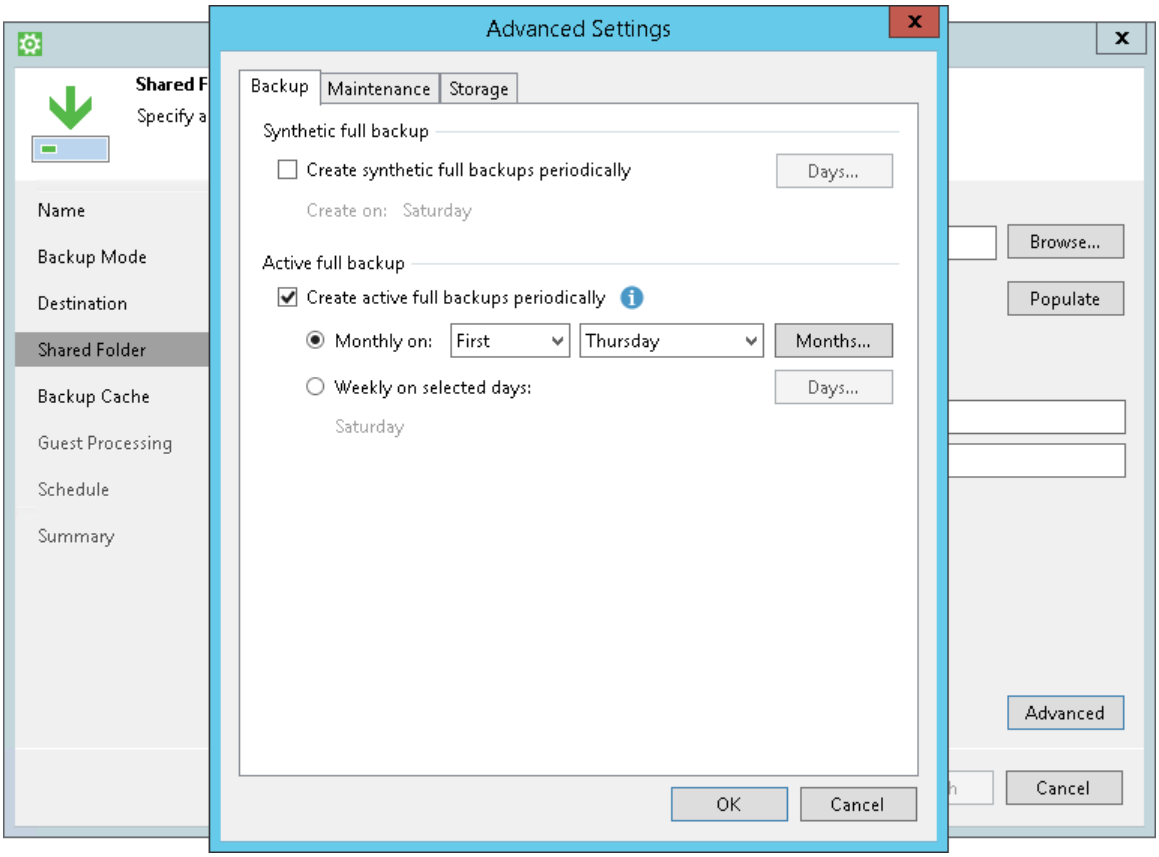
When Veeam Agent for Microsoft Windows performs active full backup, it produces a full backup file and adds this file to the backup chain.

The active full backup resets the backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file and its subsequent incremental backup files remain on the disk. After the last incremental backup file created prior to the active full backup becomes outdated, Veeam Agent for Microsoft Windows automatically deletes the previous backup chain. To learn more, see Retention Policy for Active Full Backups.



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 backup** command from the Veeam Agent Tray menu. To learn more, see Creating Active Full Backups.
- To schedule active full backups, specify scheduling settings in the **Advanced Settings** window of the **New Backup Job** wizard. You can schedule active full backups to run weekly, for example, every Saturday, or monthly, for example, every first Thursday of a month.

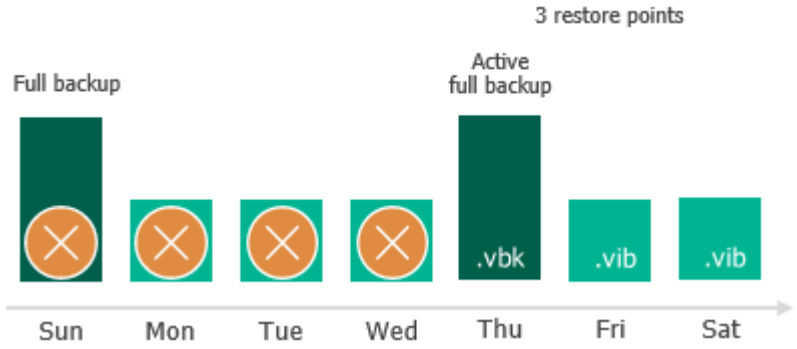


Retention Policy for Active Full Backups

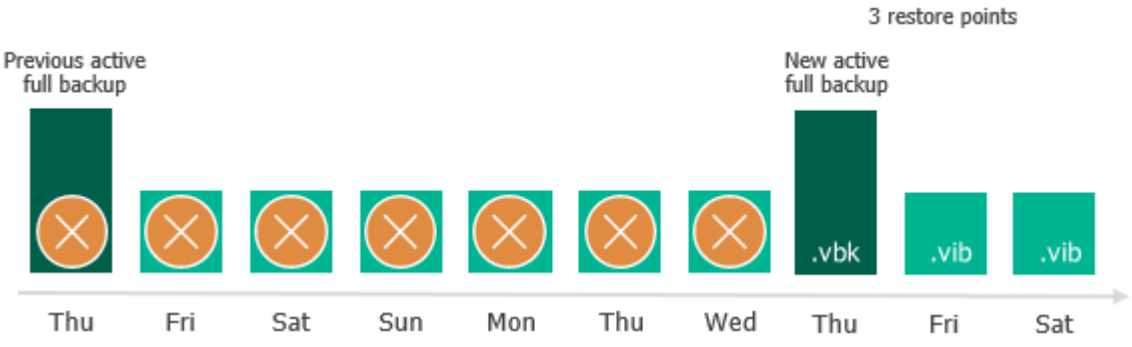
To be able to restore data from a Veeam Agent backup, you need to have a full backup file and a chain of subsequent incremental backup files on the disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you create an active full backup manually or set up the backup job to create active full backups, in some days there will be more restore points on the disk than specified by retention policy settings. Veeam Agent for Microsoft Windows will remove the full backup chain only after the last incremental backup file in the chain becomes outdated.

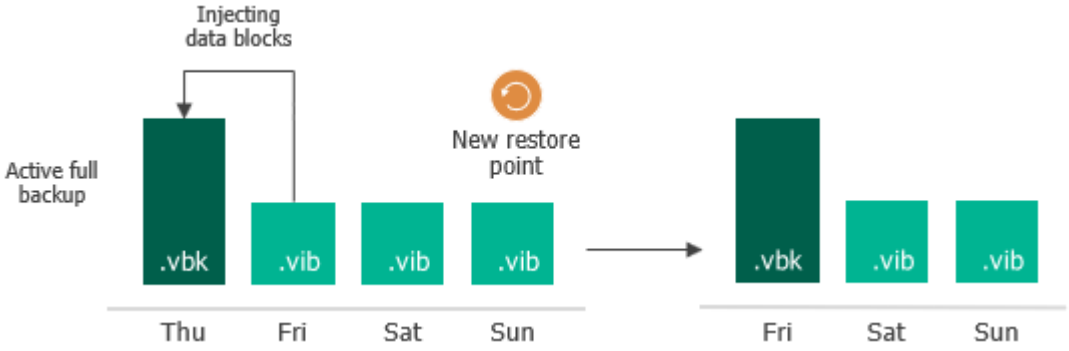
For example, the retention policy is set to 3 restore points. A full backup file is created on Sunday, incremental backup files are created Monday through Saturday, and an active full backup is scheduled on Thursday. Although the retention policy is already breached on Wednesday, the full backup is not deleted. Without the full backup, backup chain would be useless, leaving you without any restore point at all. Veeam Agent for Microsoft Windows will wait for the next 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.



Keep in mind that if the backup job is set up to create periodic active full backups, Veeam Agent for Microsoft Windows will never transform the backup chain. Instead, Veeam Agent for Microsoft Windows will always wait for the next full backup file and the necessary number of incremental backup files to be created, and only then will delete the whole previous chain. In the example above, Veeam Agent for Microsoft Windows will delete the previous chain every Saturday. As a result, although the retention policy is set to 3 restore points, the actual number of backup files on the disk will be greater most of the time.



In contrary, in a situation when you manually create a single active full backup, Veeam Agent for Microsoft Windows will treat the active full backup in the same way as a regular full backup. If some restore point becomes obsolete, Veeam Agent for Microsoft Windows will re-build the full backup file to include in it data of the incremental backup file that follows the full backup file. After that, Veeam Agent for Microsoft Windows will remove the earliest incremental backup file from the chain as redundant.



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.

> **NOTE**
>
> Synthetic full backup functionality is available only in Workstation and Server editions of Veeam Agent for Microsoft Windows.

In terms of data, the synthetic full backup is identical to a regular full backup. Synthetic full backup produces a VBK file that contains all data that you have chosen to back up. The difference between active and synthetic full backup lies in the way how backed-up data is retrieved:

- When you perform active full backup, Veeam Agent for Microsoft Windows reads backed-up data, compresses it and copies it to the target location.
- When you perform synthetic full backup, Veeam Agent for Microsoft Windows does not retrieve all backed-up data from the Veeam Agent computer. Instead, it creates a new incremental backup, then synthesizes a full backup from data you already have on the target location. Veeam Agent for Microsoft Windows accesses the previous full backup file and a chain of subsequent incremental backup files in the backup chain including a new incremental backup, consolidates 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.

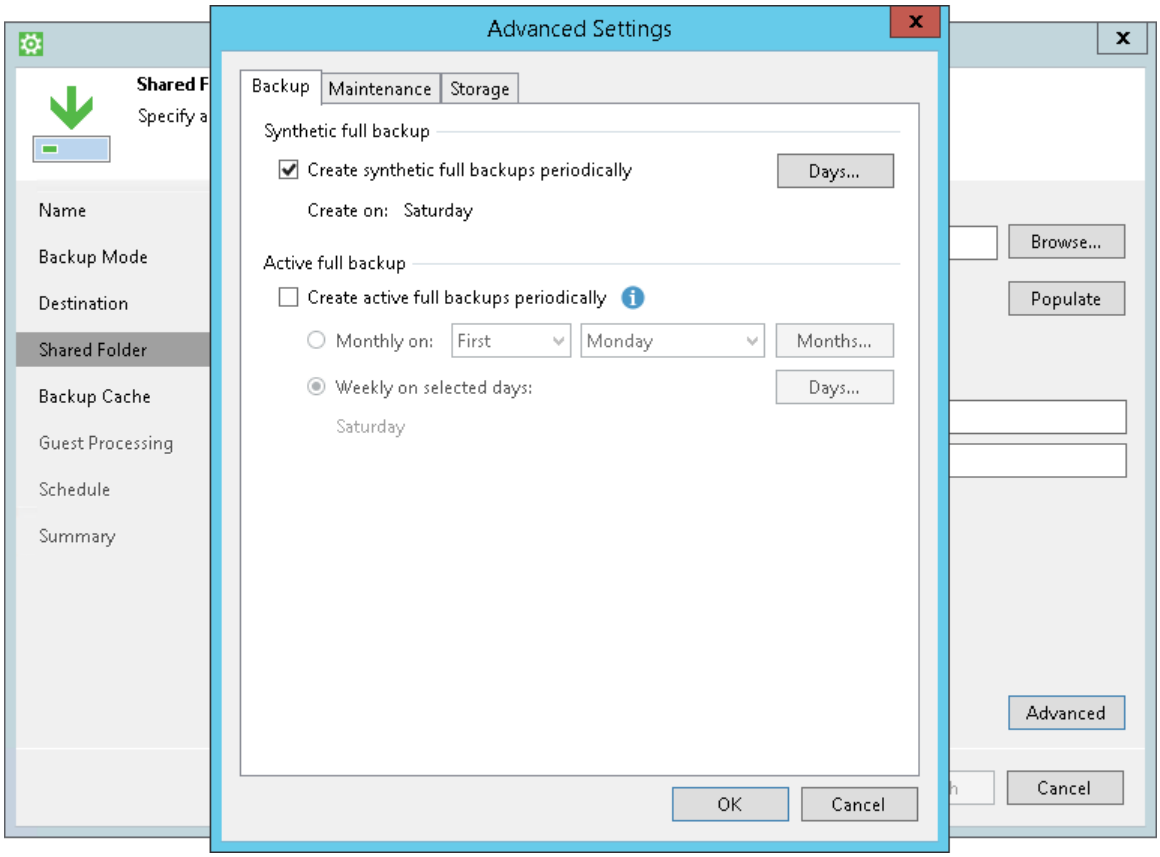
The synthetic full backup has a number of advantages:

- The synthetic full backup does not use network resources: it is created from backup files you already have on the target location.
- The synthetic full backup produces less load on the production environment: it is synthesized right on the target location.

Veeam Agent for Microsoft Windows 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 and its subsequent incremental backup files remain on the disk. After the last incremental backup file created prior to the synthetic full backup becomes outdated, Veeam Agent for Microsoft Windows automatically deletes the previous backup chain. To learn more, see Retention Policy for Synthetic Full Backups.

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.

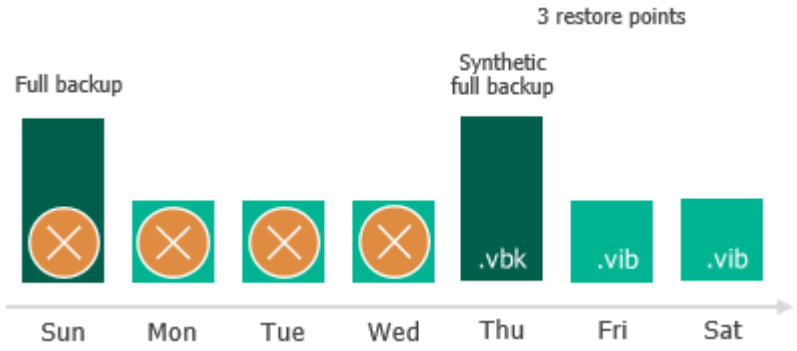


Retention Policy for Synthetic Full Backups

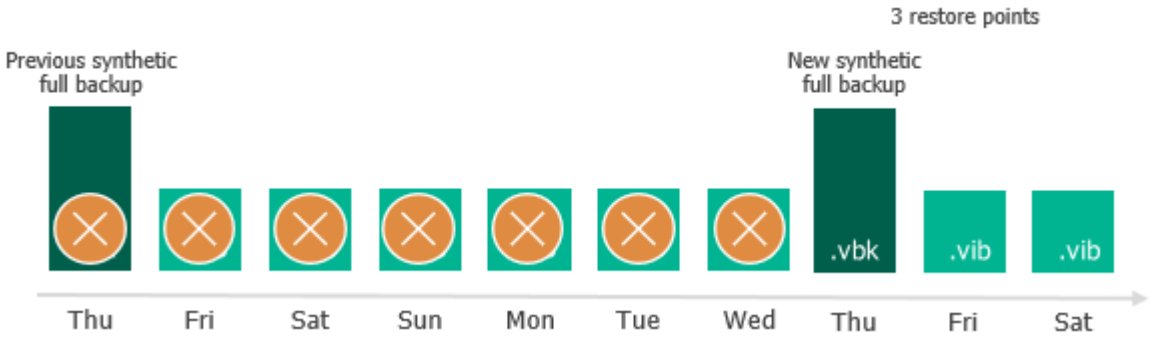
To be able to restore data from a Veeam Agent backup, you need to have a full backup file and a chain of subsequent incremental backup files on the disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you set up the backup job to create synthetic full backups, in some days there will be more restore points on the disk than specified by retention policy settings. Veeam Agent for Microsoft Windows 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. A full backup file is created on Sunday, incremental backup files are created Monday through Saturday, and synthetic full backup is scheduled on Thursday. Although the retention policy is already breached on Wednesday, the full backup is not deleted. Without the full backup, backup chain would be useless, leaving you without any restore point at all. Veeam Agent for Microsoft Windows will wait for the next 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.



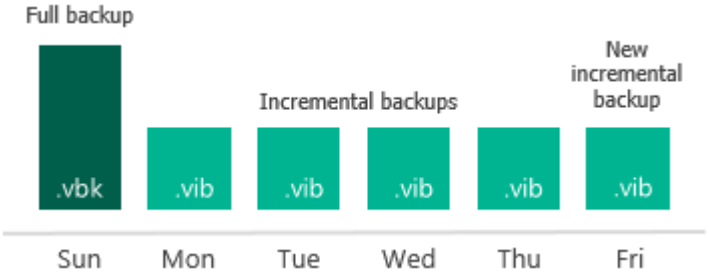
Keep in mind that if the backup job is set up to create synthetic full backups, Veeam Agent for Microsoft Windows will never transform the backup chain. Instead, Veeam Agent for Microsoft Windows will always wait for the next full backup file and the necessary number of incremental backup files to be created, and only then will delete the whole previous chain. In the example above, Veeam Agent for Microsoft Windows will delete the previous chain every Saturday. As a result, although the retention policy is set to 3 restore points, the actual number of backup files on the disk will be greater most of the time.



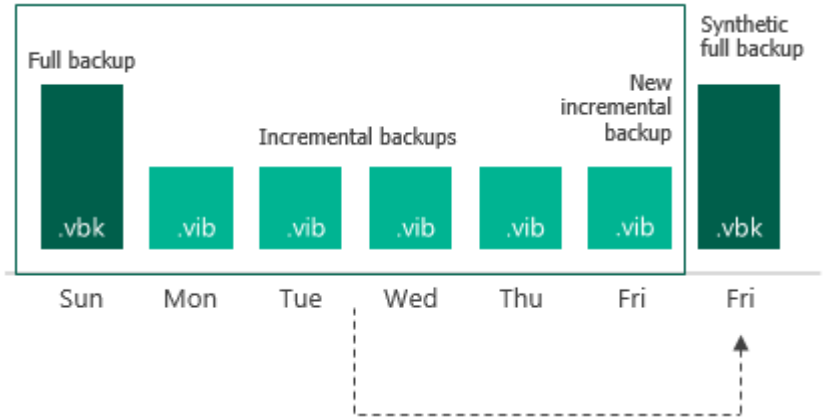
How Synthetic Full Backup Works

To create a synthetic full backup, Veeam Agent for Microsoft Windows performs the following steps:

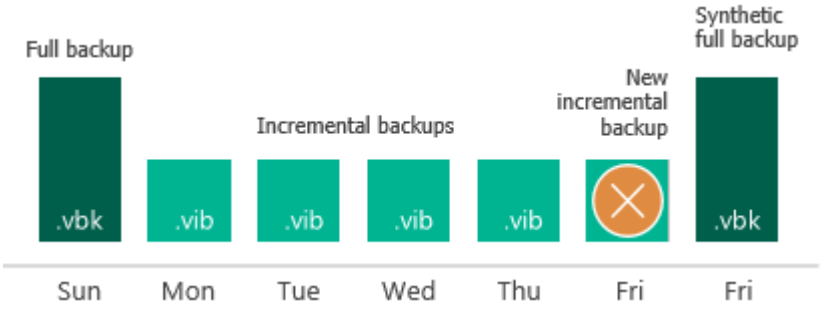
1. On a day when synthetic full backup is scheduled, Veeam Agent for Microsoft Windows triggers a new backup job session. During this session, Veeam Agent for Microsoft Windows first performs incremental backup in a regular manner and adds a new incremental backup file to the backup chain. Incremental backup helps Veeam Agent for Microsoft Windows ensure that the synthetic full backup includes the latest changes of the backed-up data.



2. At the end of the backup job session, Veeam Agent for Microsoft Windows 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 Agent for Microsoft Windows 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.



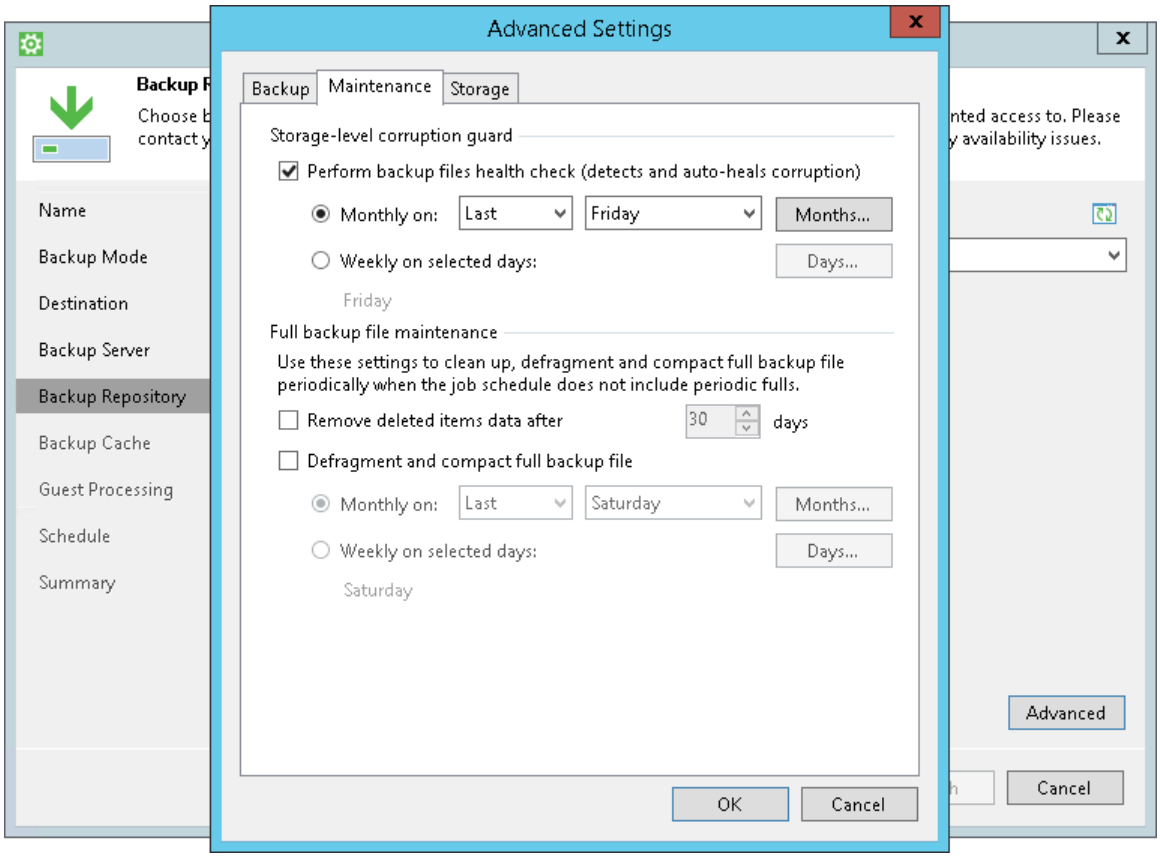
4. Every next job session creates a new incremental restore point starting from the synthetic full backup until the day on which synthetic full backup is scheduled. On this day, Veeam Agent for Microsoft Windows creates a new synthetic full backup.

Health Check for Backup Files

You can instruct Veeam Agent for Microsoft Windows to periodically perform a health check for the latest restore point in the backup chain. During the health check, Veeam Agent performs a CRC check for metadata and a hash check for 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.

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 Agent for Microsoft Windows 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 Agent 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.



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 supported for backup files created in Microsoft OneDrive.
- The health check is not performed during a backup job session that creates a restore point in the backup cache.

- The health check is not performed during the backup cache synchronization process.

How Health Check Works

When Veeam Agent for Microsoft Windows saves a new restore point to the backup location, it calculates CRC values for backup metadata and hash values for data blocks in the backup file, and saves these values in the metadata of the backup file, together with the backed-up data. During the health check session, Veeam Agent uses these values to make sure that a verified restore point is consistent.

> **NOTE**
>
> If the backup job is targeted at a Veeam backup repository or cloud repository, and you perform health check for the encrypted backup files, Veeam Backup & Replication will pass encryption keys to the Veeam backup repository or cloud repository. For more information on encryption, see Data Encryption.

Veeam Agent for Microsoft Windows performs the health check in the following way:

1. During the backup job session, Veeam Agent for Microsoft Windows creates a new restore point in the backup chain.
2. At the end of the backup job session, Veeam Agent for Microsoft Windows performs the health check. It calculates CRC values for backup metadata and hash values for data blocks in the backup file and compares them with the CRC and hash values that are already stored in the backup file.

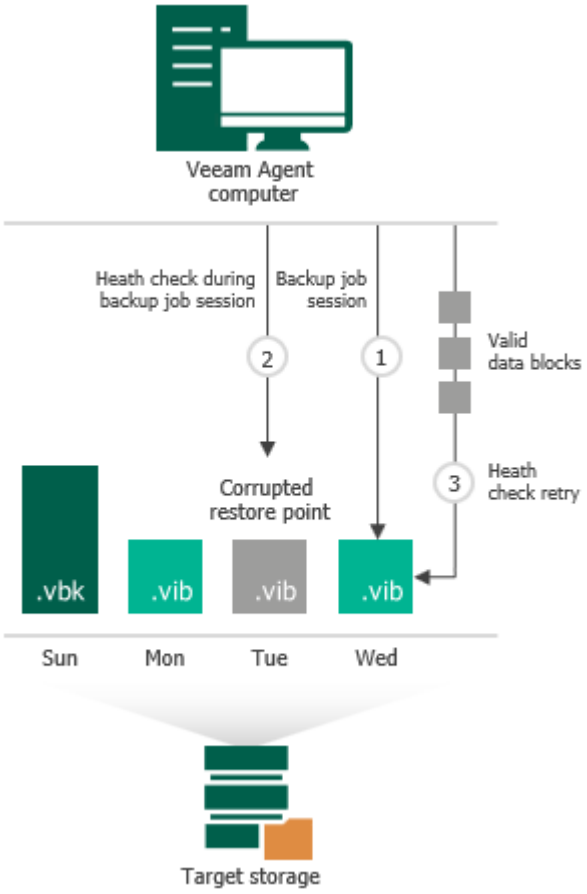
   During the health check, Veeam Agent for Microsoft Windows verifies the latest restore point in the backup chain (restore point created with the current backup job session — the session during which the health check is performed). If the latest restore point in the backup chain is incomplete (for example, if the current backup job session completed with an error), Veeam Agent checks the restore point preceding the latest one.

   If the health check does not detect data corruption, the backup job session completes in a regular way.
3. If the health check detects corrupted data, Veeam Agent for Microsoft Windows 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 Agent performs the following actions:

   - If the health check has detected corrupted backup metadata in the full backup file, Veeam Agent marks the backup chain starting from this full restore point as corrupted in the Veeam Agent for Microsoft Windows database. During the health check retry, Veeam Agent transports the entire data selected for backup from the Veeam Agent Computer, creates a new full backup file in the target location and saves transported data blocks to this backup file.
   - If the health check has detected corrupted backup metadata in the incremental backup file, Veeam Agent removes information about this incremental restore point and subsequent incremental restore points from the Veeam Agent for Microsoft Windows database. During the health check retry, Veeam Agent transports incremental data relatively the latest valid restore point in the backup chain from the Veeam Agent computer, creates a new incremental backup file in the target location and saves transported data blocks to this backup file.

- If the health check has detected corrupted data blocks in the full or incremental backup file, Veeam Agent marks the restore point that includes the corrupted data blocks and subsequent incremental restore points as corrupted in the Veeam Agent for Microsoft Windows database. During the health check retry, Veeam Agent transports data blocks from the Veeam Agent computer. In addition, Veeam Agent transports data blocks that have changed since the backup job session that has triggered the health check. Veeam Agent 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).



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 the attempts are not successful, Veeam Agent for Microsoft Windows will perform the health check during the last job retry in any case.

If the health check detects corrupted data, Veeam Agent for Microsoft Windows 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 Agent attempts to transport data blocks for the corrupted restore point from the Veeam Agent computer to the target location.

For a backup job started automatically upon schedule, the number of health check retries is equal to the number of job retries specified in the job settings. For a job started manually, Veeam Agent for Microsoft Windows performs 1 health check retry.

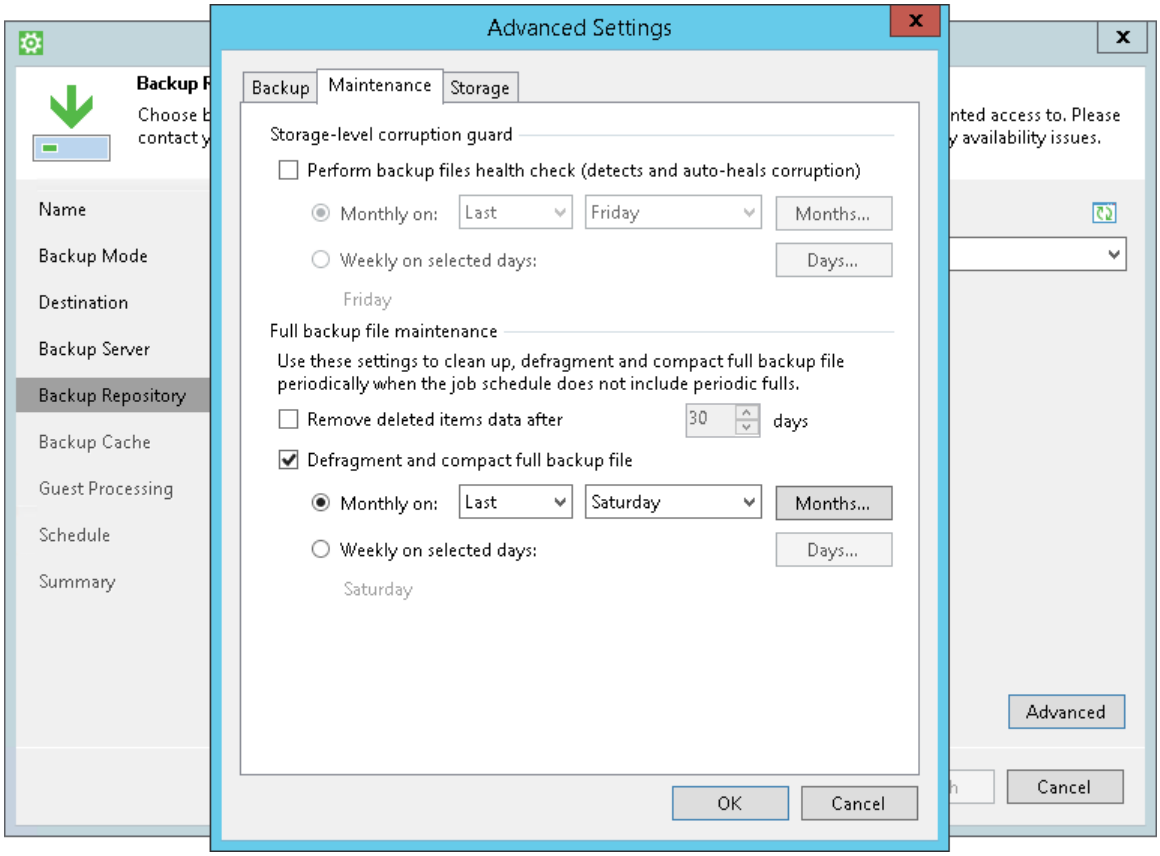
> **NOTE**
>
> If Veeam Agent for Microsoft Windows fails to fix the corrupted data during all health check retries, you must retry the job manually. In this case, Veeam Agent will transport the necessary data blocks from the Veeam Agent computer to fix the corrupted restore point.

Compact of Full Backup File

If you do not create periodic full backups, 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 badly 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 Agent for Microsoft Windows to compact the full backup file periodically. During the file compact operation, Veeam Agent creates a new empty file and copies to this file data blocks from the original 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 Agent for Microsoft Windows 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 is not available for backup jobs targeted at the Microsoft OneDrive cloud storage.
- The **Defragment and compact full backup file** option can be enabled only for backup jobs for which active full and synthetic full backups are not scheduled.
- The compact full backup file operation is not performed during backup job sessions that produce active full backups. If the backup job starts again on the same day when the active full backup was created, Veeam Agent for Microsoft Windows does not perform the compact full backup operation. This limitation helps reduce the number of backup operations — Veeam Agent considers that the full backup is recent and does not need to be rebuilt.

  If such situation occurs, Veeam Agent for Microsoft Windows triggers the full backup file compact operation during the next backup job session that produces an incremental backup file on another day.
- The target location must have enough space to store a file of the full backup size. During the compact process, Veeam Agent for Microsoft Windows creates auxiliary files that exist in the target location until the end of the compact operation.
- If you change the block size in backup job settings, Veeam Agent for Microsoft Windows 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 Agent for Microsoft Windows changes the compression level for the compacted backup file.
- The file compact operation is not performed during a backup job session that creates a restore point in the backup cache.
- The file compact operation is not performed during the backup cache synchronization process.

Removal of Deleted Drives Data

During the compact operation, Veeam Agent for Microsoft Windows does not simply copy all data blocks from the VBK file to the newly created file. It copies only data blocks of Veeam Agent computer drives whose information is stored in the Veeam Agent for Microsoft Windows database. For example, if a drive was removed from the Veeam Agent computer, 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.

Changed Block Tracking

To perform incremental backup, Veeam Agent for Microsoft Windows needs to know what data blocks have changed since the previous job session. To get the list of changed data blocks, Veeam Agent for Microsoft Windows uses the changed block tracking mechanism, or CBT. CBT increases the speed and efficiency of incremental backups.

To keep track of changed data blocks, Veeam Agent for Microsoft Windows can use the following mechanisms:

- **Default CBT mechanism** — this mechanism is enabled by default in all installations of Veeam Agent for Microsoft Windows. To learn more, see Default Changed Block Tracking Mechanism.
- **Veeam CBT driver** — this functionality is available if the Veeam Agent computer runs a Microsoft Windows Server OS and the Server edition of Veeam Agent for Microsoft Windows. To learn more, see Veeam Changed Block Tracking Driver.

Default Changed Block Tracking Mechanism

Veeam Agent for Microsoft Windows performs the changed block tracking differently depending on the backup type:

- CBT for volume-level backup
- CBT for file-level backup

CBT for Volume-Level Backup

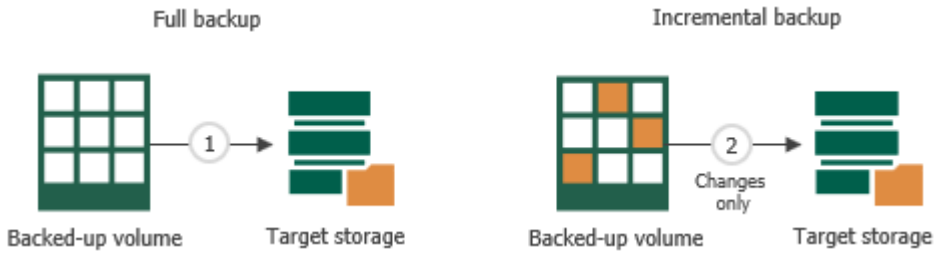
In case of the volume-level backup, Veeam Agent for Microsoft Windows performs changed block tracking in the following way:

1. During the full backup job session, Veeam Agent reads the Master File Table (MFT) of the backed-up volume. Veeam Agent uses MFT records to create digests with file system metadata, transfers the created digests to the target location and stores them to the resulting backup file.
2. During subsequent incremental job sessions, Veeam Agent performs the following operations:
   a. Reads the Master File Table (MFT) of the backed-up volume and creates the new digests with file system metadata.
   b. Interacts with the target backup location to obtain digests from the backup file that was created during the previous job session.
   c. Compares new and previous digests to detect files whose data blocks have changed on the volume since the previous job session.

      During incremental backup, Veeam Agent for Microsoft Windows reads from the VSS snapshot only data blocks pertaining to files that have changed since the previous job session. If Veeam Agent cannot calculate information about the changed files, for example, if it fails to retrieve digests from the backup file, Veeam Agent will need to read all data blocks from the VSS snapshot. As a result, the backup may take significantly more time.

NOTE

Veeam Agent for Microsoft Windows uses the default CBT mechanism for NTFS volumes only. As a result, for volumes that use other file systems, incremental backup will require greater time, because Veeam Agent for Microsoft Windows will read all data from the VSS snapshot to detect what blocks have changed since the last job session.



CBT for File-Level Backup

In case of the file-level backup, Veeam Agent for Microsoft Windows performs changed block tracking in the following way:

1. During the full backup job session, Veeam Agent creates a new NTFS partition in the backup file. Veeam Agent uses this partition to store all backed-up files and folders.
2. During subsequent incremental job sessions, Veeam Agent performs the following operations:
    a. Compares the last modification time attribute of files on the Veeam Agent computer and files in the backup. This operation allows Veeam Agent to detect files that have changed since the previous job session.
    b. Transfers data blocks that contain metadata related to each changed file and its parent folders from the target backup location to the Veeam Agent computer.

       If the backed-up file system has a complex folder structure with many hierarchy levels, and you back up data to a remote location, during incremental backup, the inbound network traffic on the Veeam Agent computer may exceed by far the outbound traffic. Significant amount of data can be transferred to the Veeam Agent computer even if few files are changed since the previous job session.
    c. Replaces file system metadata in the downloaded data blocks with the current file system metadata.
    d. Transfers data blocks with current metadata from the Veeam Agent computer to the target backup location along with changed files. In the target backup location, Veeam Agent stores data and metadata in the newly created backup file.

NOTE

Consider the following:

- Veeam Agent for Microsoft Windows detects changed files using the last modification time attribute. If you change the file in any way, but the last modification time attribute remains unchanged, Veeam Agent will not back up this file during an incremental job session.
- If the backed-up file has changes, Veeam Agent will copy to the target location not only changed data blocks of the file but the entire file. To save space on the target storage, consider using the volume-level backup. To learn more, see Backup Types.



Veeam Changed Block Tracking Driver

You can set up Veeam Agent for Microsoft Windows to use the Veeam CBT driver instead of the default CBT mechanism. The Veeam CBT driver is a class filter driver for volume devices that helps Veeam Agent for Microsoft Windows keep track of changed data blocks in a more efficient way. The driver is intended for servers running applications with large database files.

To use the Veeam CBT driver, the Veeam Agent computer must meet the following requirements:

- Run a Microsoft Windows Server OS
- Run the Server edition of Veeam Agent for Microsoft Windows

> **IMPORTANT**
>
> Consider the following:
>
> - Prior to installing the Veeam CBT driver on a computer running Microsoft Windows Server 2008 R2 SP1, make sure that update KB3033929 is installed in the OS.
>
> The update adds the SHA-2 code signing support that is required for verification of the Veeam CBT driver signature. Without this update installed, the OS running on a protected computer will fail to boot after you install the Veeam CBT driver. To learn more, see this Microsoft KB article.
> - Do not install the Veeam CBT driver on a computer running Microsoft Windows Server 2008 R2 SP1, 2012 or 2012 R2 if one or more volumes on this computer are encrypted with Microsoft BitLocker (or other encryption tool), or if you plan to use Microsoft BitLocker to encrypt volumes on this computer. Concurrent operation of Microsoft BitLocker and Veeam CBT driver may result in driver failures and may prevent the OS from starting.

To enable the advanced CBT mechanism provided by the Veeam CBT driver, you need to install the driver in the Veeam Agent control panel. You can perform this operation at any time you need. To activate the driver after installation, Veeam Agent for Microsoft Windows needs to reboot the computer. After computer reboot, the Veeam CBT driver will start keeping track of changed data blocks on computer volumes whose data you have selected for backup in the Veeam Agent backup job settings.

In contrast to the default CBT mechanism that supports NTFS volumes only, the Veeam CBT driver can keep track of changed data blocks on volumes that use the following file systems:

- NTFS
- FAT
- ReFS

Information about changed data blocks is registered in special VCT files. VCT files are stored in the `C:\ProgramData\Veeam\EndpointData\CtStore` folder on the Veeam Agent computer. When the backup job runs, Veeam Agent for Microsoft Windows uses VCT files to find out what data blocks have changed since the last run of the job, and copies only changed data blocks from the backed-up volume.

If you use the Veeam CBT driver, the `C:\ProgramData\Veeam\EndpointData\CtStore` folder accessibility is crucial for a successful backup. Consider the following:

- The folder and its content must not be compressed or encrypted.
- The folder and its parent folders must not be reparse points.

NOTE

Consider the following:

- If the Veeam Agent computer shuts down unexpectedly, the Veeam CBT driver may fail to register information about changed data blocks in a VCT file. In this case, during the next backup job session, Veeam Agent for Microsoft Windows will need to read all data from the backed-up volume to create incremental backup. As a result, incremental backup will require greater time.
- If data blocks are changed on a volume while this volume is mounted on another Windows-based machine, during the next backup job session, Veeam Agent for Microsoft Windows will also read all data from the volume to create incremental backup.
- The Veeam CBT driver cannot detect data block changes made on a volume that is mounted in a non-Windows OS. For example, such changes can be made when you boot your computer using a Linux-based antivirus rescue disc. To continue the backup chain after such changes, you need to create active full backup instead of incremental backup. Alternatively, you can reset CBT. To learn more, see Resetting CBT.

Data Compression

Veeam Agent for Microsoft Windows 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.

Data Compression

Data compression decreases the size of created backups but affects duration of the backup procedure. Veeam Agent for Microsoft Windows allows you to select one of the following compression levels:

- **None** compression level is recommended if you plan to store backup files on storage devices that support hardware compression and deduplication.
- **Dedupe-friendly** is an optimized compression level for very low CPU usage. You can select this compression level if you want to decrease the load on the CPU of the Veeam Agent computer.
- **Optimal** is the recommended compression level. It provides the best ratio between size of the backup file and time of the backup 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 backup file but reduces the backup performance. We recommend that you use the extreme compression level only on Veeam Agent computers with modern multi-core CPUs (6 cores recommended).

You can change data compression settings for existing backup jobs. New settings will not have any effect on previously created backup files in the backup chain. They will be applied to new backup files created after the settings were changed.

Compression settings are changed on the fly. You do not need to create a new full backup to use new settings — Veeam Agent for Microsoft Windows will automatically apply the new compression level to newly created backup files.

Storage Optimization

Depending on the type of storage you select as a backup target, Veeam Agent for Microsoft Windows uses data blocks of different size, which optimizes the size of a backup file and job performance. You can choose one of the following storage optimization options:

- The **Local target (large blocks)** option is recommended for backup jobs that can produce very large full backup files — larger than 16 TB. With this option selected, Veeam Agent for Microsoft Windows uses data block size of 4096 KB.

  Large data blocks produce a smaller metadata table that requires less memory and CPU resources to process. Keep in mind, however, that this storage optimization option will provide the largest size of incremental backup files.

- The **Local target** option is recommended for backup to SAN, DAS or local storage. With this option selected, Veeam Agent for Microsoft Windows uses data block size of 1024 KB.

  The SAN identifies larger blocks of data and therefore can process large amounts of data at a time. This option provides the fastest backup job performance.

- The **LAN target** option is recommended for backup to NAS and onsite backup. With this option selected, Veeam Agent for Microsoft Windows uses data block size of 512 KB. This option reduces the size of an incremental backup file because of reduced data block sizes.

- The **WAN target** option is recommended if you plan to use WAN for offsite backup. With this option selected, Veeam Agent for Microsoft Windows uses data block size of 256 KB. This results in the smallest size of backup files, allowing you to reduce the amount of traffic over WAN.

> **NOTE**
>
> Because of Microsoft OneDrive limitations, a larger number of backed-up data blocks results in lower write speed on a Microsoft OneDrive storage. As a result, we do not recommend to use data block size of less than 1024 KB for the backup job targeted at Microsoft OneDrive (as long as your network connection bandwidth allows you to transfer larger amounts of incremental backup data over the internet).

To apply new storage optimization settings, you must create an active full backup after you change storage optimization settings. Veeam Agent for Microsoft Windows will use the new block size for the active full backup and subsequent backup files in the backup chain.

Guest Processing

For Server edition of Veeam Agent for Microsoft Windows, you can specify guest processing options. Veeam Agent for Microsoft Windows offers the following guest processing options:

- Application-aware processing. You can create transactionally consistent backups of servers running applications that support Microsoft VSS. Application-aware processing guarantees that you can perform restore from Veeam Agent backups without data loss.
- Pre-freeze and post-thaw scripts. You can use pre-freeze and post-thaw scripts to quiesce applications that do not support Microsoft VSS.
- Transaction log truncation. You can set up the backup job to truncate transaction logs after the job successfully completes.
- Transaction log backup for Microsoft SQL Server and Oracle. You can set up the backup job to back up transaction logs from servers running Microsoft SQL Server and archived logs of Oracle database systems.
- File system indexing. You can set up the backup job to create a catalog of files and folders on the Veeam Agent computer OS. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you will be able to search for individual files in Veeam Agent backups and perform 1-click restore in Veeam Backup Enterprise Manager.

  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 Veeam Agent backup. For more information, see the Veeam Backup Enterprise Manager User Guide.

Supported Applications

Veeam Agent for Microsoft Windows supports VSS-aware processing for the following systems:

Specification	Requirement
Microsoft Active Directory Domain Controllers	The following versions of Microsoft Active Directory Domain Services servers (domain controllers) are supported: • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 Minimum supported domain and forest functional level is Windows 2003.
Microsoft Exchange	The following versions of Microsoft Exchange are supported: • Microsoft Exchange 2019 • Microsoft Exchange 2016 • Microsoft Exchange 2013 SP1 • Microsoft Exchange 2013 • Microsoft Exchange 2010 SP1, SP2, or SP3

Specification	Requirement
Microsoft SharePoint	The following versions of Microsoft SharePoint are supported: • Microsoft SharePoint 2019 • Microsoft SharePoint 2016 • Microsoft SharePoint 2013 • Microsoft SharePoint 2010 All editions are supported (Foundation, Standard, Enterprise).
Microsoft SQL Server	The following versions of Microsoft SQL Server are supported: • Microsoft SQL Server 2019 • Microsoft SQL Server 2017 • Microsoft SQL Server 2016 • Microsoft SQL Server 2014 • Microsoft SQL Server 2012 • Microsoft SQL Server 2008 R2 • Microsoft SQL Server 2008 • Microsoft SQL Server 2005 SP4 All editions of Microsoft SQL Server except LocalDB are supported.
Oracle	Oracle Database 11g and 21c are supported for the following operating systems (32-bit and 64-bit architecture): • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 **Important notes:** • Automatic Storage Management (ASM) is not supported. • Oracle Real Application Clusters (RAC) are not supported. • Oracle servers using Data Guard are not supported. • Oracle Database Express Edition is supported. • Configurations with different versions of Oracle Database deployed on the same server are not supported. • 32-bit Oracle running on 64-bit operating systems is not supported.

Application-Aware Processing

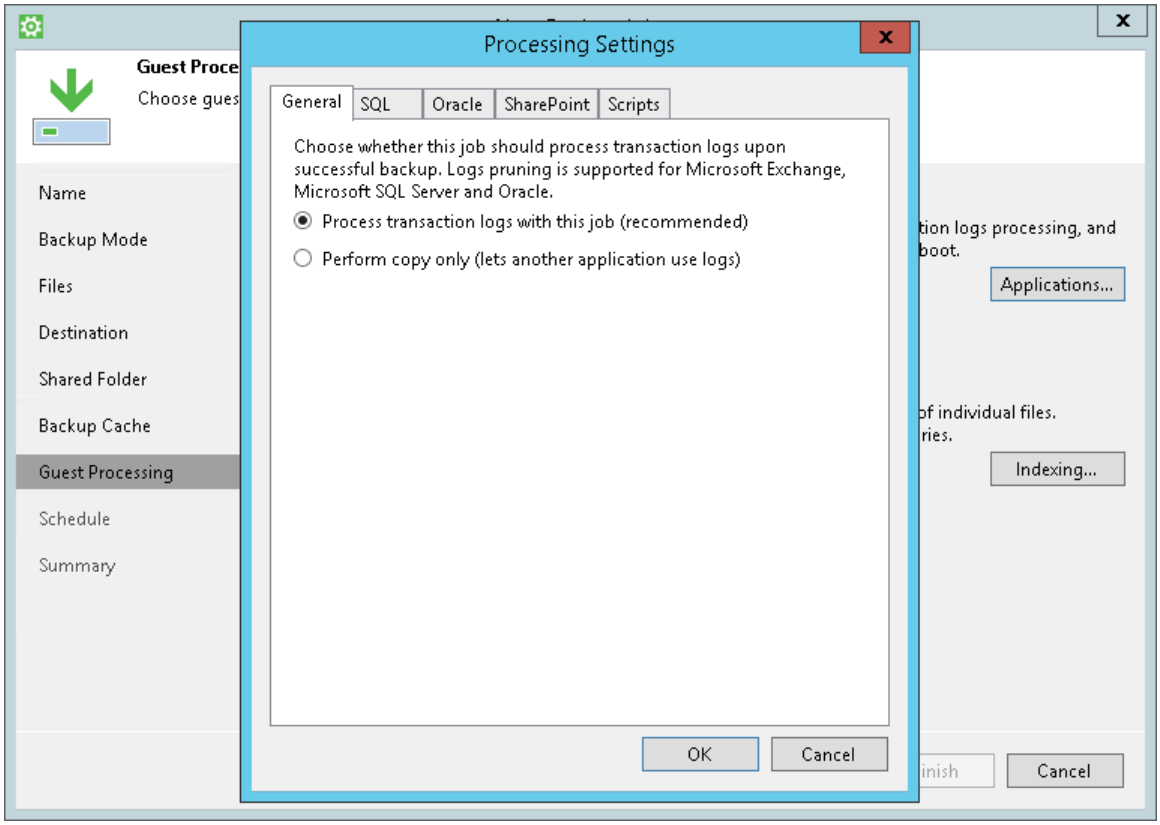
To create transactionally consistent backups of servers that run VSS-aware applications such as Microsoft SQL Server, Microsoft SharePoint, Microsoft Exchange or Oracle, you must enable application-aware processing for the backup job.

Application-aware processing is Veeam's proprietary technology based on Microsoft VSS. Microsoft VSS is responsible for quiescing applications and creating a consistent view of application data on the OS of the Veeam Agent computer. Use of Microsoft VSS ensures that there are no unfinished database transactions or incomplete application files when Veeam Agent creates a Microsoft VSS snapshot and starts copying backed-up data to the target location. For more information about Microsoft VSS, see this Microsoft article.

In the Server edition of Veeam Agent, the type of the VSS snapshot depends on application-aware processing settings:

- If application-aware processing is disabled for the backup job, Veeam Agent creates a copy-only VSS snapshot.
- If application-aware processing is enabled and the **Perform copy only** option is selected, Veeam Agent creates a copy-only VSS snapshot.
- If application-aware processing is enabled and the **Process transaction logs with this job** option is selected, Veeam Agent creates a full VSS snapshot.

To learn more, see Specify Application-Aware Processing Settings.



Application-aware processing is supported for Microsoft Windows 2008 R2 SP1 and later. To use application-aware processing, you must have the latest updates installed on the Veeam Agent computer OS. To learn more, see Supported Applications.

IMPORTANT

If your computer OS runs an application that does not support Microsoft VSS (there is no VSS writer for this particular type of application, for example, MySQL), Veeam Agent for Microsoft Windows will not be able to utilize Microsoft VSS and application-aware processing for this application. To process such applications, you can use pre-freeze and post-thaw scripts. For more information, see Pre-Freeze and Post-Thaw Scripts.

Pre-Freeze and Post-Thaw Scripts

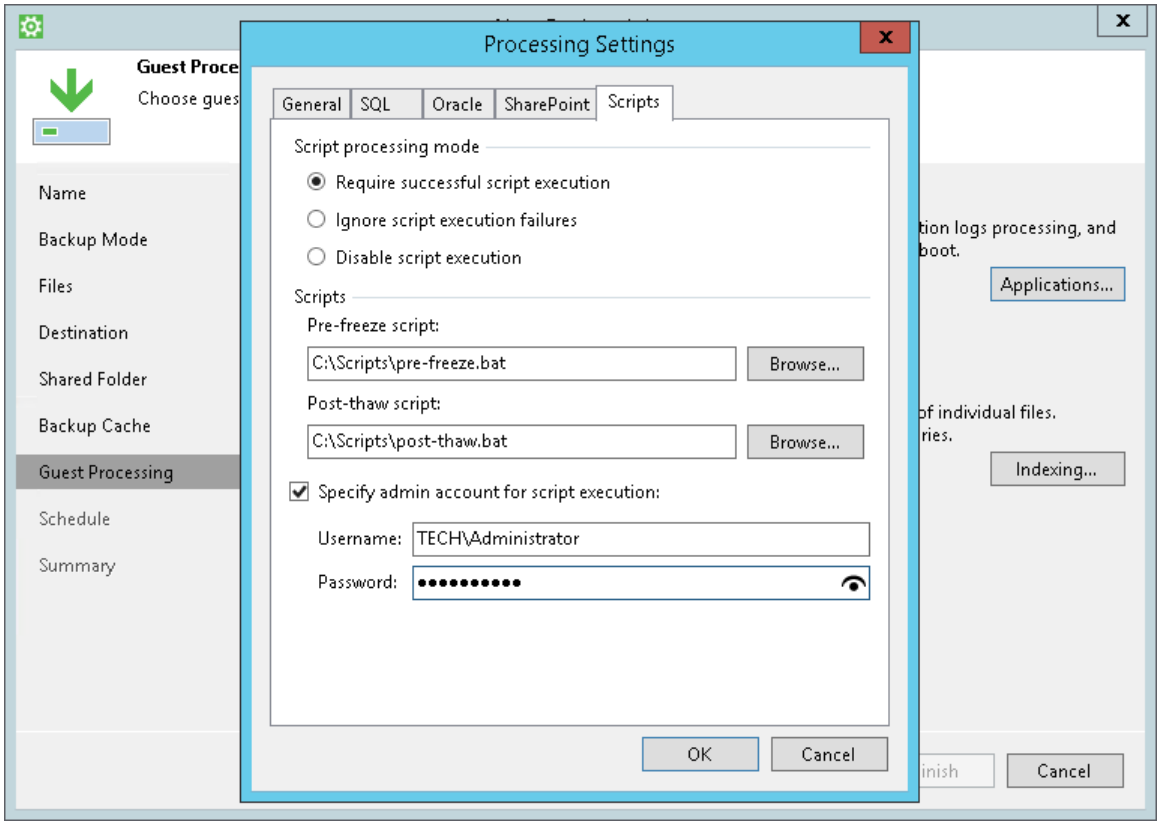
If Veeam Agent computer runs applications that do not support Microsoft VSS, you can instruct Veeam Agent for Microsoft Windows to run custom scripts during the backup job session. For example, the pre-freeze script may quiesce the file system and application data to bring the computer OS to a consistent state before Veeam Agent for Microsoft Windows creates a Microsoft VSS snapshot. After the VSS snapshot is created, the post-thaw script may bring the OS and applications to their initial state.

Veeam Agent for Microsoft Windows supports scripts in the EXE, BAT and CMD file format.

Scripts must be created beforehand. You must specify paths to them in the backup job settings. Scripts must reside on a local drive of the Veeam Agent computer.

When the backup job starts, Veeam Agent for Microsoft Windows executes scripts specified for the job. A 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 Agent for Microsoft Windows displays an error message in the job session and error or warning messages issued during script execution.



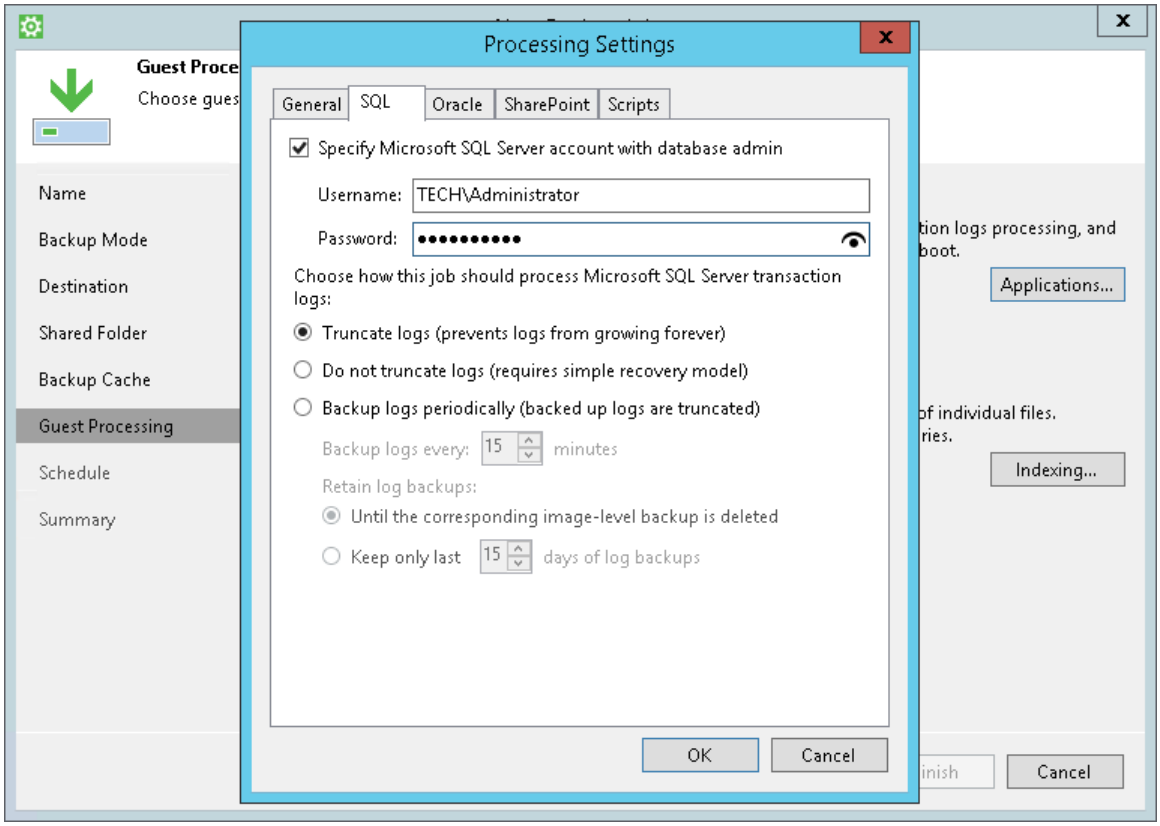
Limitations for Pre-Freeze and Post-Thaw Scripts

Veeam Agent for Microsoft Windows has one limitation 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 Agent for Microsoft Windows waits for 10 minutes and terminates the job.

Transaction Log Truncation

If you back up database systems that use transaction logs, for example, Microsoft SQL Server, you can instruct Veeam Agent for Microsoft Windows to truncate transaction logs so that logs do not overflow the storage space. Veeam Agent for Microsoft Windows provides the following options of transaction logs handling:

- Truncate logs
- Do not truncate logs
- Back up logs periodically



Truncate Logs

You can instruct Veeam Agent for Microsoft Windows to truncate logs after a backup is successfully created. With this option selected, Veeam Agent for Microsoft Windows behaves in the following way:

- If the backup job completes successfully, Veeam Agent for Microsoft Windows produces a backup file and truncates transaction logs on the Veeam Agent computer. As a result, you have the backup file that contains a computer image, image of a specific data volume or individual folders at a specific point in time.

  In this scenario, you can recover a database to the point in time when the backup file was created. As transaction logs on the Veeam Agent computer are truncated, you cannot use them to get the restored database to some point in time between backup job sessions.

- If the backup job fails, Veeam Agent for Microsoft Windows does not truncate transaction logs on the Veeam Agent computer. In this scenario, you can restore computer data from the most recent point in the backup 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. We recommend this option if you use another backup tool together with Veeam Agent for Microsoft Windows.

For example, you can use Veeam Agent for Microsoft Windows to create a computer image backup and instruct the native Microsoft SQL Server log backup job to back up transaction logs. If you truncate transaction logs with Veeam Agent for Microsoft Windows, 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 Agent for Microsoft Windows produces a backup file and does not trigger transaction log truncation. As a result, you have a backup file that contains a computer image, image of a specific data volume or individual folders captured at a specific point in time, and transaction logs. You can use transaction logs to restore the Veeam Agent computer to any point in time between job sessions. To do this, you must recover data from the backup file 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 or Oracle database system.

You can choose to back up database logs with Veeam Agent for Microsoft Windows. With this option selected, Veeam Agent for Microsoft Windows creates a backup and additionally copies Microsoft SQL Server transaction logs or Oracle archived logs and saves them to the backup location next to the backup files. To learn more, see Microsoft SQL Server and Oracle Logs Backup.

In this scenario, you can use transaction logs to restore the Veeam Agent computer to any point in time between backup job sessions. To do that, you must recover data from the Veeam Agent backup and use Veeam Explorer for Microsoft SQL Server or Veeam Explorer for Oracle to perform transaction log replay and get the database system to a necessary point in time.

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.

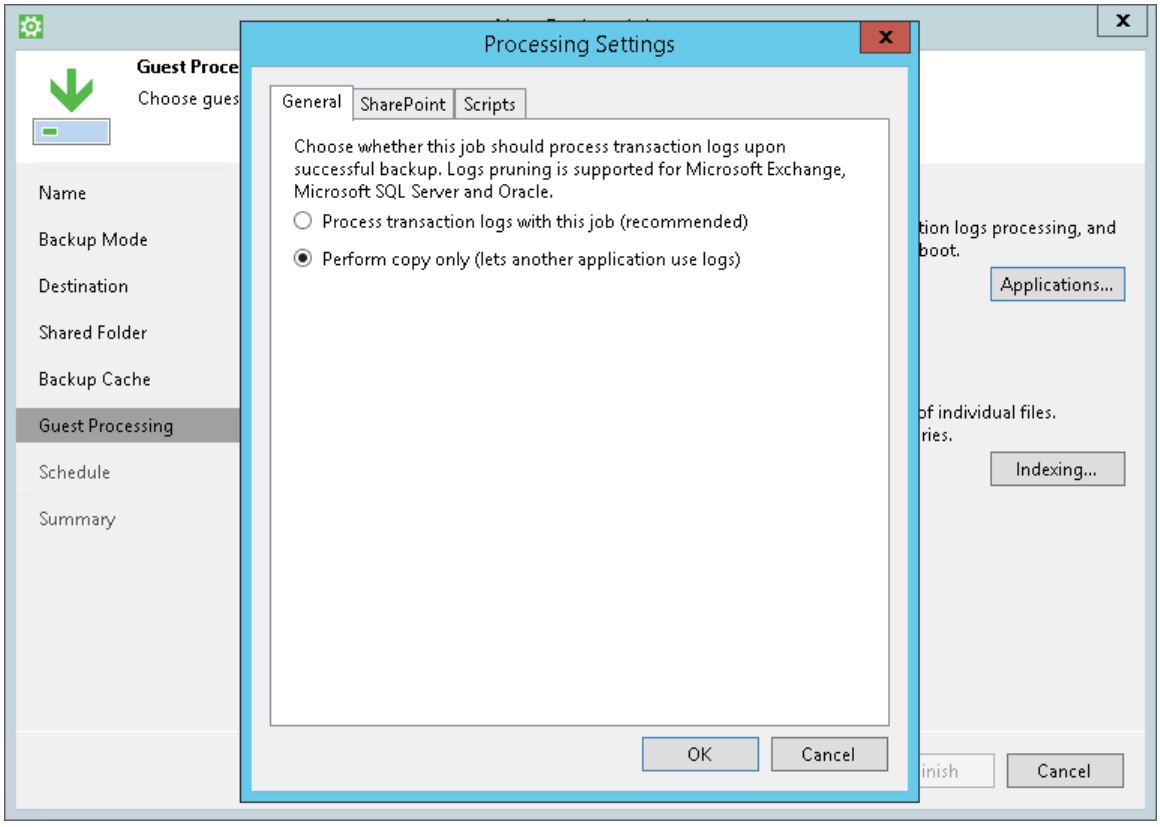
If you use native Microsoft SQL Server tools or 3rd party backup tools and also want to back up a machine that runs Microsoft SQL Server with Veeam Agent for Microsoft Windows, you must enable the **Perform copy only** option in the backup job settings.

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 Agent to preserve this chain (backup history). If the **Perform copy only** option is enabled, Veeam Agent produces the backup of the Copy Backup type that is independent of the existing chain of database backups and does not contain transaction log data. As a result, the backup does not change the log sequence number and transaction log backup time. To learn more about Copy Backup and other VSS backup types, see this Microsoft article.

IMPORTANT

Consider the following:

- Veeam Agent for Microsoft Windows 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.
- Veeam Agent for Microsoft Windows supports one transaction log backup job per Veeam Agent computer. If you plan to configure several backup jobs using Veeam Agent for Microsoft Windows, you can enable database log backup settings for one job only.



Microsoft SQL Server and Oracle Logs Backup

NOTE

This and subsequent sections describe application-aware processing of Microsoft SQL Server and Oracle database systems in Veeam Agent for Microsoft Windows. You can perform item-level recovery of Microsoft SQL Server and Oracle systems if you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication. For more information, see Veeam Backup & Replication Documentation.

Microsoft SQL Server Log Backup

You can instruct the Veeam Agent backup job to create volume level or file-level backups and also periodically back up database transaction logs. If Microsoft SQL Server fails, you can restore Microsoft SQL Server from the necessary restore point of the Veeam Agent backup. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you can also 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.

Requirements for Microsoft SQL Server Transaction Log Backup

- Veeam Agent for Microsoft Windows supports transaction log backups for the following systems:
  - Microsoft SQL Server 2019
  - Microsoft SQL Server 2017
  - Microsoft SQL Server 2016
  - Microsoft SQL Server 2014 SP2
  - Microsoft SQL Server 2012 SP3
  - Microsoft SQL Server 2008 R2 SP3
  - Microsoft SQL Server 2008 SP4
- 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 this Microsoft article.

Oracle Log Backup

Veeam Agent for Microsoft Windows supports backup of Oracle database archived logs. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you can use Veeam Explorer for Oracle to apply archived logs and get Oracle databases to the necessary state between backups.

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 archived logs are saved. We do not recommend this mode, because it does not provide for proper disaster recovery.

With ARCHIVELOG turned on, the Oracle system stores database archived logs to a certain location on the machine that runs the database system, as specified by the database administrator. Veeam Agent for Microsoft Windows allows you to set up the following ways of log handling:

- Instruct the backup job to collect log files from the Oracle system and ship them to the backup location where they are stored next to regular backup files created by Veeam Agent for Microsoft Windows.
- Skip log processing — log files remain untouched and are preserved within the Veeam Agent backup.

If you enable application-aware processing for Oracle, during the job session Veeam Agent for Microsoft Windows collects information about the database and processes archived logs according to job settings. Application-specific settings are configured at the **Guest Processing** step of the **New Backup Job** wizard — you can specify how logs should be backed up and/or truncated for Oracle databases.

Requirements for Oracle Archived Log Backup

- Veeam Agent for Microsoft Windows supports archived log backup and restore for Oracle database version 11.2 and later.
- Automatic Storage Management (ASM) is not supported.
- The database must run in the ARCHIVELOG mode.

Database Log Backup Job

To back up database logs (Microsoft SQL Server transaction logs and Oracle archived logs), you must specify advanced settings for transaction log backup in the Veeam Agent backup job settings. The resulting job will comprise two jobs:

- Parent backup job — the backup job that creates a volume-level or file-level backup. The backup job becomes the parent job after you enable database log backup options at the **Guest Processing** step of the **New Backup Job** wizard.
- Child job — a transaction log backup job. Veeam Agent for Microsoft Windows automatically creates the child job if transaction log backup is enabled for the backup job. Session data of the transaction log backup job is stored in the Veeam Agent for Microsoft Windows database and displayed in the Veeam Agent control panel. To learn more, see Transaction Log Backup Statistics.

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 restore point is present when it comes to transaction log replay.

> **NOTE**
>
> Veeam Agent for Microsoft Windows supports one transaction log backup job per Veeam Agent computer. If you plan to configure several backup jobs using Veeam Agent for Microsoft Windows, you can enable database log backup settings for one job only.

Sessions of Transaction Log Backup Jobs

The transaction log backup job runs permanently in the background, shipping transaction logs to the backup location 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, and/or when this parent backup job is disabled.

How Microsoft SQL Server Log Backup Works

The transaction log backup for Microsoft SQL Server is performed in the following way:

1. Veeam Agent for Microsoft Windows launches the parent backup job by schedule.
2. The parent backup job creates a volume-level or file-level backup and stores it to the backup location.
3. A new session of the transaction log backup job starts. Veeam Agent for Microsoft Windows copies transaction log files from the log archive destination (set by the Microsoft SQL Server administrator) to a temporary folder on the Veeam Agent computer file system.
4. Veeam Agent for Microsoft Windows detects what databases currently exist on the Microsoft SQL Server and maps this data with the information kept in the Veeam Agent for Microsoft Windows database. This periodic mapping reveals the databases for which Veeam Agent for Microsoft Windows must process transaction logs during this time interval.
5. Veeam Agent for Microsoft Windows transports transaction log backup copies from the temporary folder to the backup location and saves them as VLB files. As soon as copies of transaction log backups are saved to the backup location, transaction log backups in the temporary folder on the Veeam Agent computer are removed.

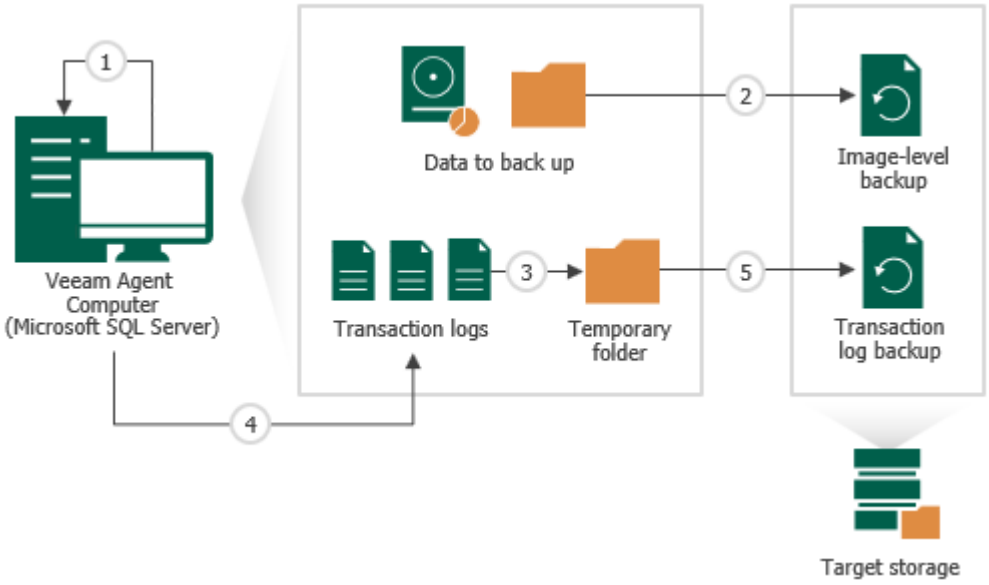
The session of the transaction log backup job remains working until the next start of the parent backup job. When a new session of the parent job starts, the transaction log backup job stops the current session and then starts a new session, performing steps 1–5.

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 Agent for Microsoft Windows enumerates log files in the temporary folder.

NOTE

Consider the following:

- 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.
- Default location of the temporary folder where Veeam Agent copies transaction log files is `C:\ProgramData\Veeam\SqlLogBackup`. If it is necessary to change the location, you can use a registry value. For more information, contact Veeam Customer Support.



How Oracle Archived Log Backup Works

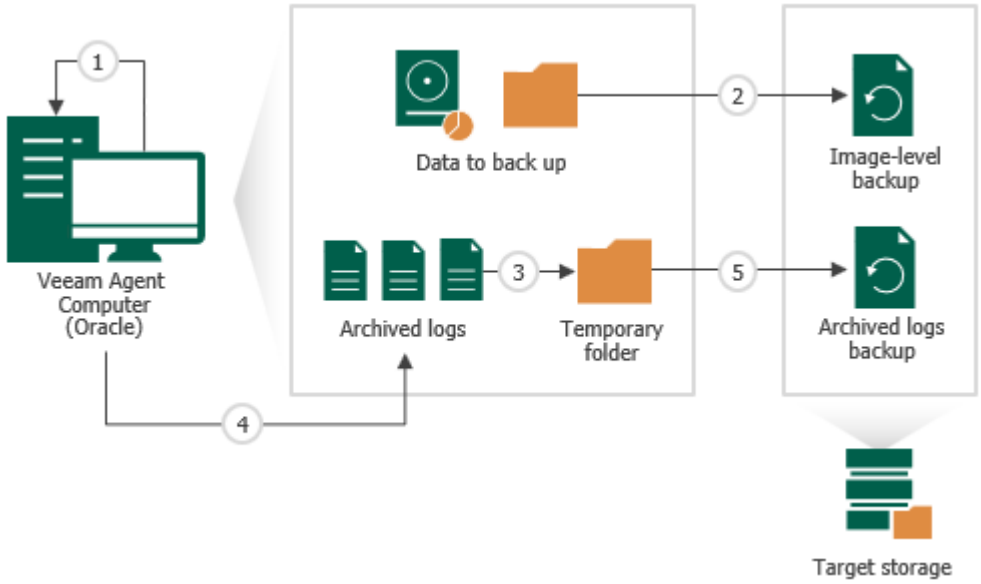
The archived log backup for Oracle is performed in the following way:

1. Veeam Agent for Microsoft Windows launches the parent backup job by schedule.
2. The parent backup job creates a volume-level or file-level backup and stores it to the backup location.
3. A new session of the archived log backup starts. Veeam Agent for Microsoft Windows scans the Oracle system and collects 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 Agent for Microsoft Windows copies archived log files from the log archive destination (set by the Oracle administrator) to a temporary folder on the Veeam Agent computer.
4. Veeam Agent for Microsoft Windows maps information about the Oracle system collected at the step 3 with information kept in the Veeam Agent for Microsoft Windows database. This periodic mapping helps reveal databases for which Veeam Agent for Microsoft Windows must ship archived logs to the backup location during this time interval.
5. Archived log backup files are transferred from the temporary folder on the Veeam Agent computer to the backup location.

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 Agent for Microsoft Windows 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 transaction log backup job will remain in the idle state, waiting for a new restore point to be created.



Retention for Database Log Backups

Transaction log backups are stored in files of the proprietary Veeam format — VLB. Veeam Agent for Microsoft Windows keeps transaction log backups together with the chain of backup files on the target location.

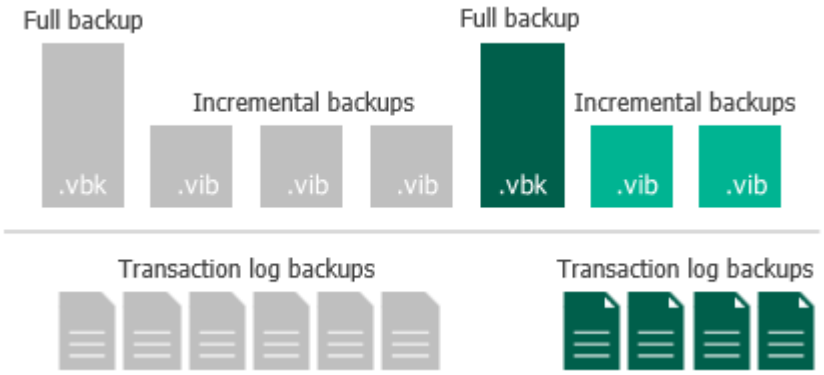
Veeam Agent for Microsoft Windows removes transaction log backups by retention. You can choose one of the following retention methods:

- 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 Agent for Microsoft Windows retains transaction log backups together with the corresponding backup file. When Veeam Agent for Microsoft Windows removes a restore point from the backup chain, it also removes a chain of transaction logs relating to this restore point.

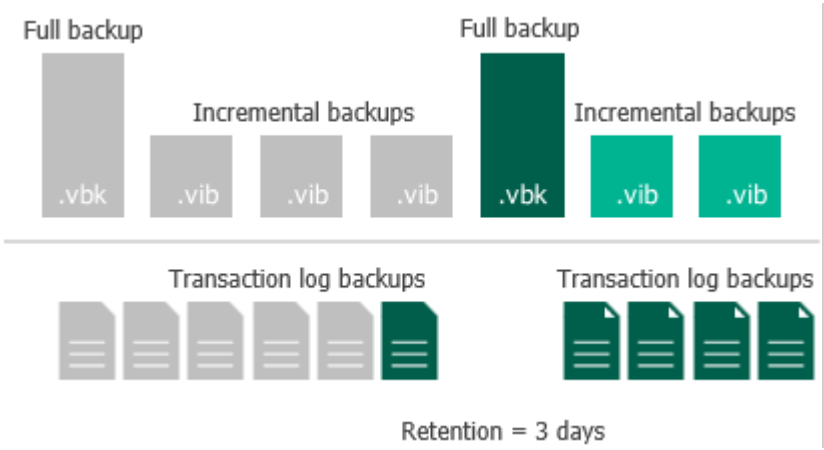
This method allows you to have both the file-level or volume-level backup and necessary transaction log backups at hand. If you need to recover a database to some state, you can restore a machine running Microsoft SQL Server or Oracle from the necessary restore point and perform transaction log replay to bring the database to the desired state.



Retain Logs for a Number of Days

You can instruct Veeam Agent for Microsoft Windows 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 Veeam Agent backup and transaction log backup are consistent. The restore point of the volume-level or file-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.

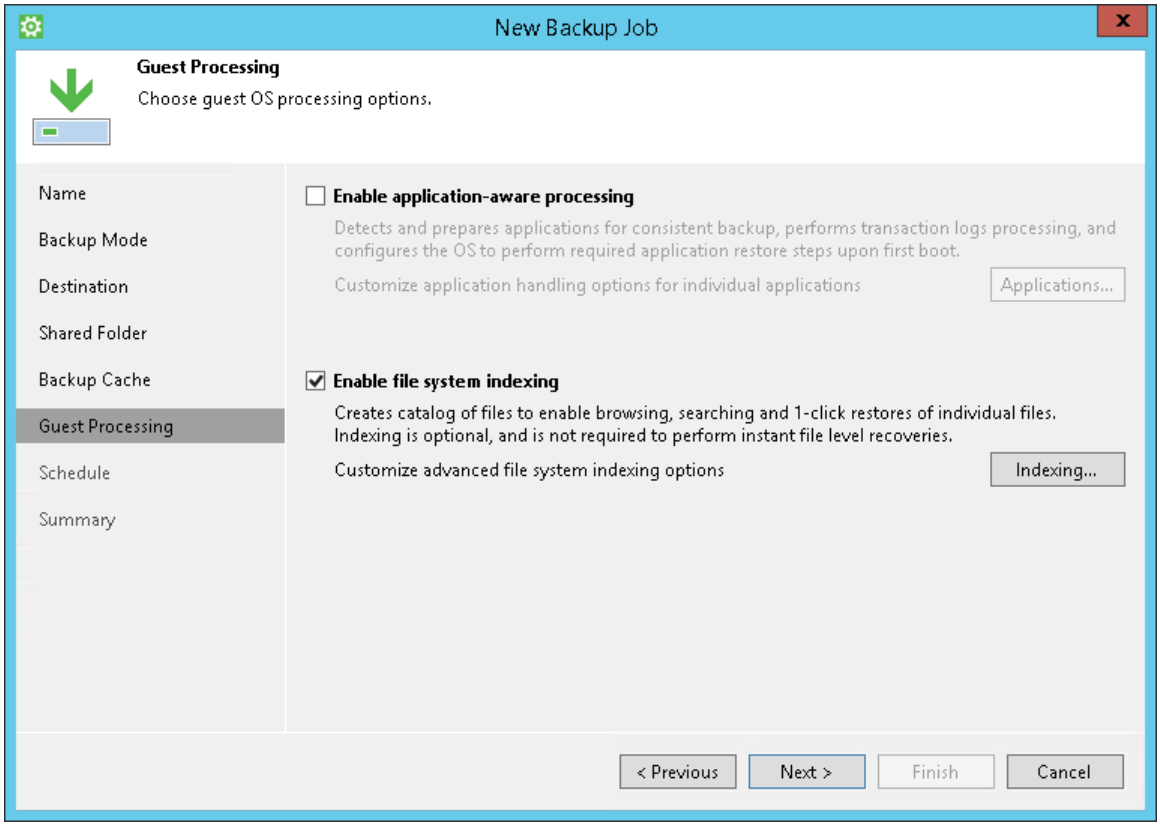


File System Indexing

You can instruct Veeam Agent for Microsoft Windows to create an index of files and folders on the Veeam Agent computer OS during backup. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you will be able to search for individual files inside Veeam Agent backups and perform 1-click restore in Veeam Backup Enterprise Manager.

NOTE

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 Veeam Backup Enterprise Manager User Guide.



How File Indexing Works

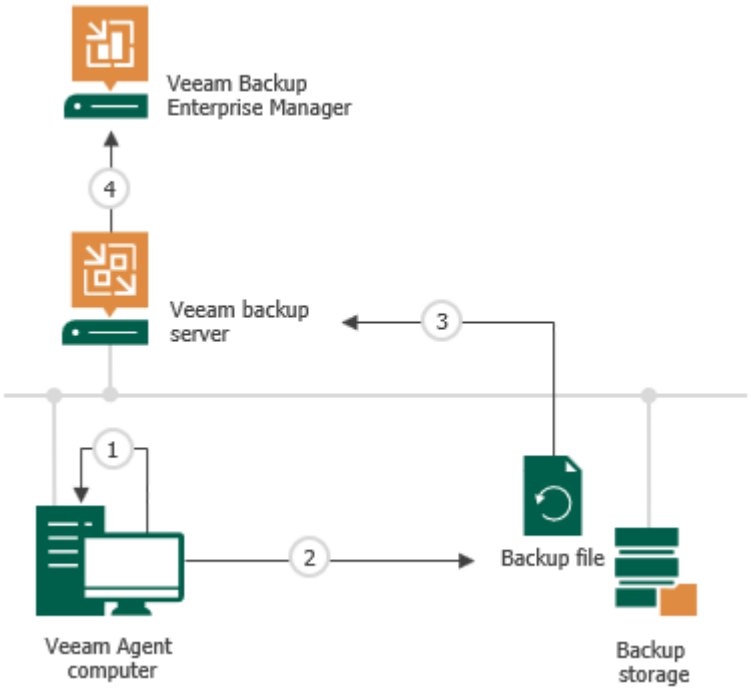
When you run a backup job with the file indexing option enabled, Veeam Agent for Microsoft Windows performs the following operations:

1. When the backup job starts, Veeam Agent for Microsoft Windows starts indexing the file system. The indexing procedure is carried out in parallel with the backup procedure. If indexing takes long, Veeam Agent for Microsoft Windows will not wait for the indexing procedure to complete. It will start copying data to the target location and continue file indexing.
2. When file indexing is complete, Veeam Agent for Microsoft Windows collects indexing data, writes it to an archive file and stores this archive file to the backup file along with the backed-up data.
3. If the backup job is set up to create backups in a Veeam backup repository, when the job completes, Veeam Guest Catalog Service running on the backup server also saves indexing data in the Veeam Catalog folder on the backup server.

   To learn more about the Veeam Guest Catalog Service, see the Veeam Backup Catalog section in the Veeam Backup & Replication User Guide.
4. 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.

NOTE

If the backup job is set up to create backups in a Veeam Cloud Connect repository, Veeam Backup & Replication running on the SP backup server does not save indexing data in the Veeam Catalog folder.



Data Encryption

Data security is an important part of the backup strategy. You must protect your information from unauthorized access, especially if you back up sensitive data to remote locations. To keep your data safe, you can use data encryption.

Data encryption transforms data to an unreadable, scrambled format with the help of a cryptographic algorithm and a secret key. If encrypted data is intercepted, it cannot be unlocked and read by the eavesdropper. Only intended recipients who know the secret key can reverse encrypted information back to a readable format.

In Veeam Agent for Microsoft Windows, encryption works at the backup job level.

Veeam Agent for Microsoft Windows uses the block cypher encryption algorithm. Encryption works at the source side. Veeam Agent for Microsoft Windows reads backed-up data, encodes data blocks, transfers them to the target location in the encrypted format and stores the data to a backup file. Data decryption is also performed on the source side: Veeam Agent for Microsoft Windows transfers encrypted data back to the source side and decrypts it there.

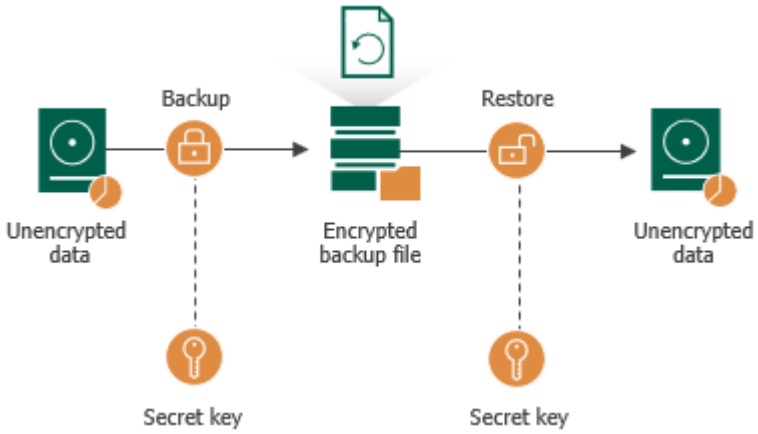
To create encrypted backups, you must enable the **Enable backup file encryption** option and specify a password that will be used for data encryption.

> **NOTE**
>
> You cannot enable encryption options in the properties of the Veeam Agent backup job if you have chosen to create Veeam Agent backups on a Veeam backup repository. For such jobs, encryption options are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the Data Encryption section in the Veeam Backup & Replication User Guide.

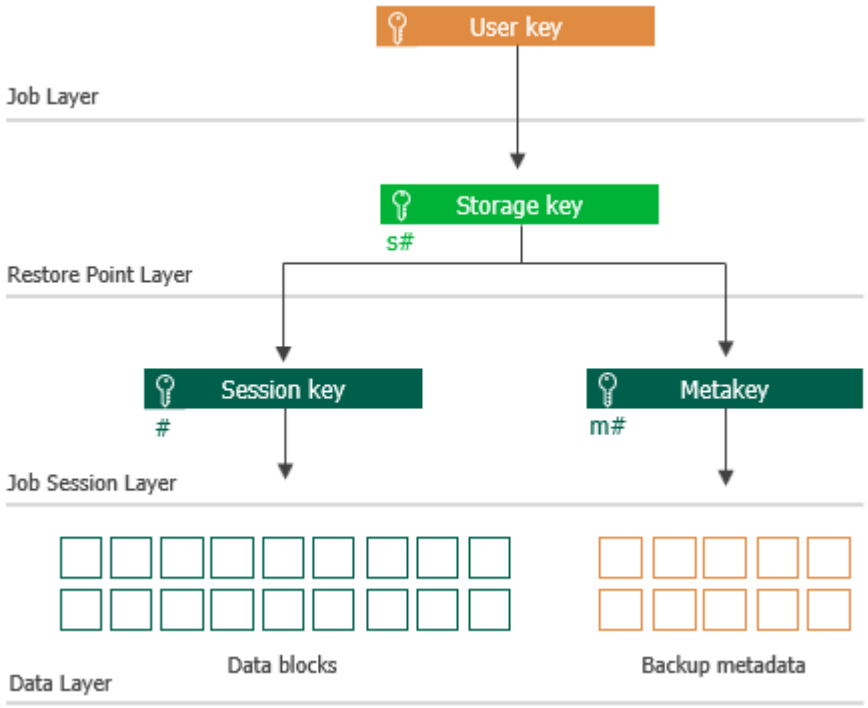
Encryption Algorithms

To encrypt data in backups and files, Veeam Agent for Microsoft Windows employs a symmetric key encryption algorithm.

The symmetric, or single-key encryption algorithm, uses a single, common secret key to encrypt and decrypt data. Before data is sent to the target location, it is encoded with a secret key. To restore encrypted data, you must have the same secret key. Users who do not have the secret key cannot decrypt data and get access to it.



Veeam Agent for Microsoft Windows relies on a hierarchical encryption scheme. Each layer in the hierarchy encrypts the layer below with a key of specific type.

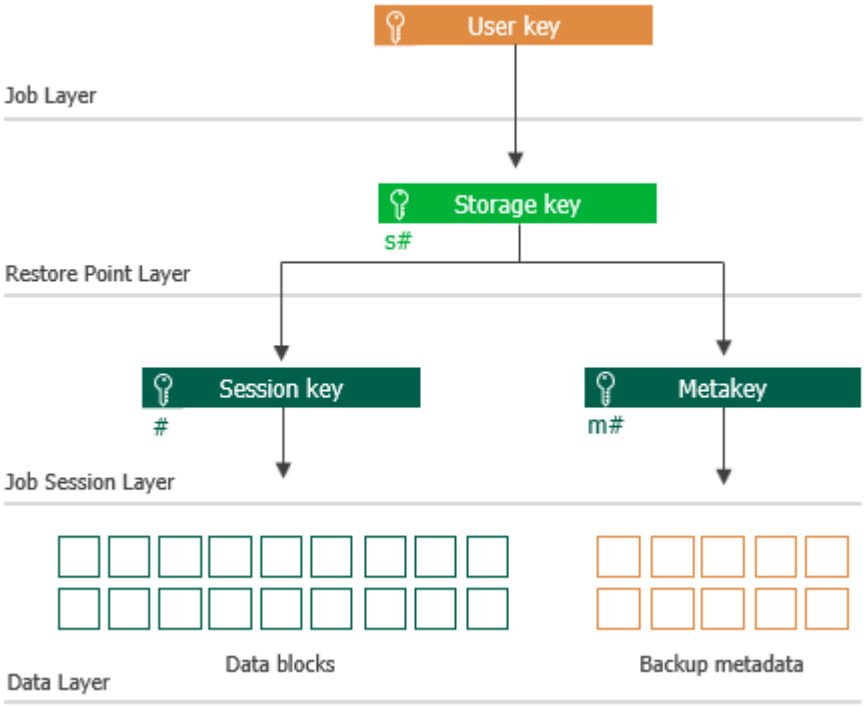


Encryption Keys

An encryption key is a string of random characters that is used to bring data to a scrambled format and back to unscrambled. Encryption keys encode and decode initial data blocks or underlying keys in the key hierarchy.

Veeam Agent for Microsoft Windows uses 4 types of keys:

- 3 service keys generated by Veeam Agent for Microsoft Windows:
  - Session key
  - Metakey
  - Storage key
- 1 key generated based on a user password: a user key.

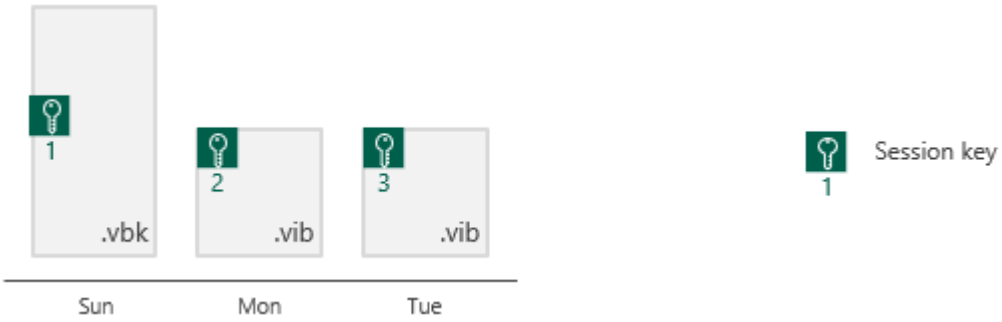


Session Keys and Metakeys

The session key is the lowest layer in the encryption key hierarchy. When Veeam Agent for Microsoft Windows encrypts data, it first encodes every data block in a file with a session key. For session keys, Veeam Agent for Microsoft Windows uses the AES algorithm with a 256-bit key length in the CBC-mode.

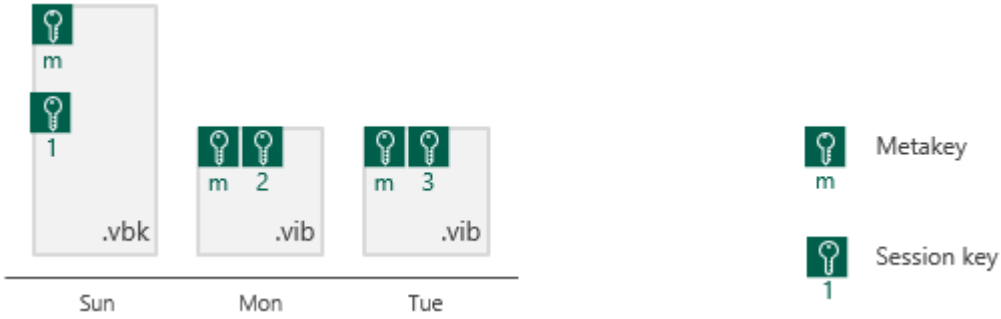
Veeam Agent for Microsoft Windows generates a new session key for every backup job session. For example, if you have created an encrypted backup job and run 3 job sessions, Veeam Agent for Microsoft Windows will produce 3 backup files that will be encrypted with 3 different session keys:

- Full backup file encrypted with session key 1
- Incremental backup file encrypted with session key 2
- Incremental backup file encrypted with session key 3



The session key is used to encrypt only data blocks in backup files. To encrypt backup metadata, Veeam Agent for Microsoft Windows applies a separate key — metakey. Use of a metakey for metadata raises the security level of encrypted backups.

For every job session, Veeam Agent for Microsoft Windows generates a new metakey. For example, if you have run 3 job sessions, Veeam Agent for Microsoft Windows will encrypt metadata with 3 metakeys.

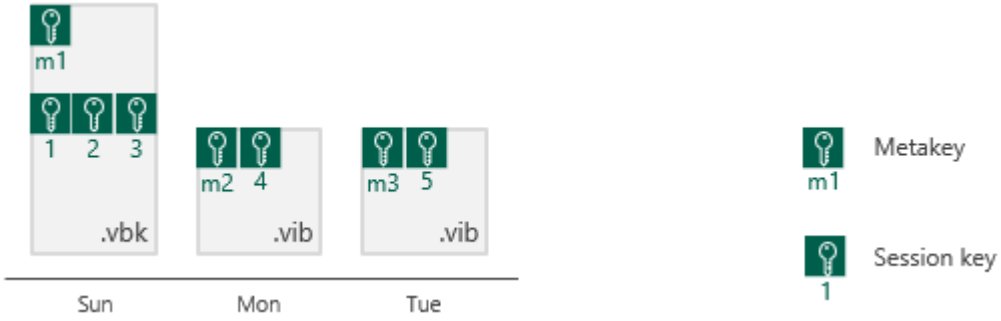


In the encryption process, session keys and metakeys are encrypted with keys of a higher layer — storage keys. Cryptograms of session keys and metakeys are stored to the resulting file next to encrypted data blocks. Metakeys are additionally kept in the Veeam Agent for Microsoft Windows database.

Storage Keys

Backup files in the backup chain often need to be transformed, for example, when the earliest incremental backup file in the chain becomes obsolete and its data should be included into the full backup file. When Veeam Agent for Microsoft Windows transforms a full backup file, it writes data blocks from several restore points to the full backup file. As a result, the full backup file contains data blocks that are encrypted in different job sessions with different session keys.

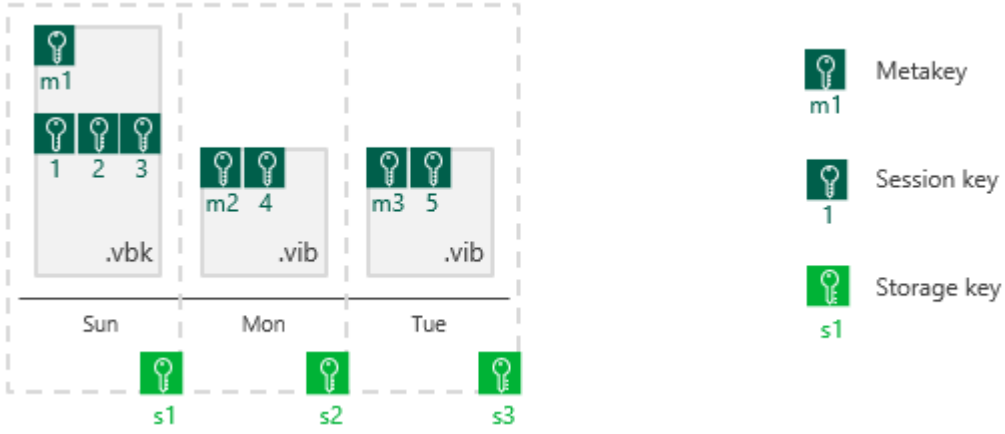
To restore data from such "composed" backup file, Veeam Agent for Microsoft Windows would require a bunch of session keys. For example, if the backup chain contains restore points for 2 months, Veeam Agent for Microsoft Windows would have to keep session keys for a 2-month period.



In such situation, storing and handling session keys would be resource consuming and complicated. To facilitate the encryption process, Veeam Agent for Microsoft Windows uses another type of service key — a storage key.

For storage keys, Veeam Agent for Microsoft Windows uses the AES algorithm with a 256-bit key length in the CBC-mode. A storage key is directly associated with one restore point in the backup chain. The storage key is used to encrypt the following keys in the encryption hierarchy:

- All session keys for all data blocks in one restore point
- Metakey encrypting backup metadata



During the restore process, Veeam Agent for Microsoft Windows uses one storage key to decrypt all session keys for one restore point, no matter how many session keys were used to encrypt data blocks in this restore point. As a result, Veeam Agent for Microsoft Windows does not need to keep the session keys history in the Veeam Agent for Microsoft Windows database. Instead, it requires only one storage key to restore data from one file.

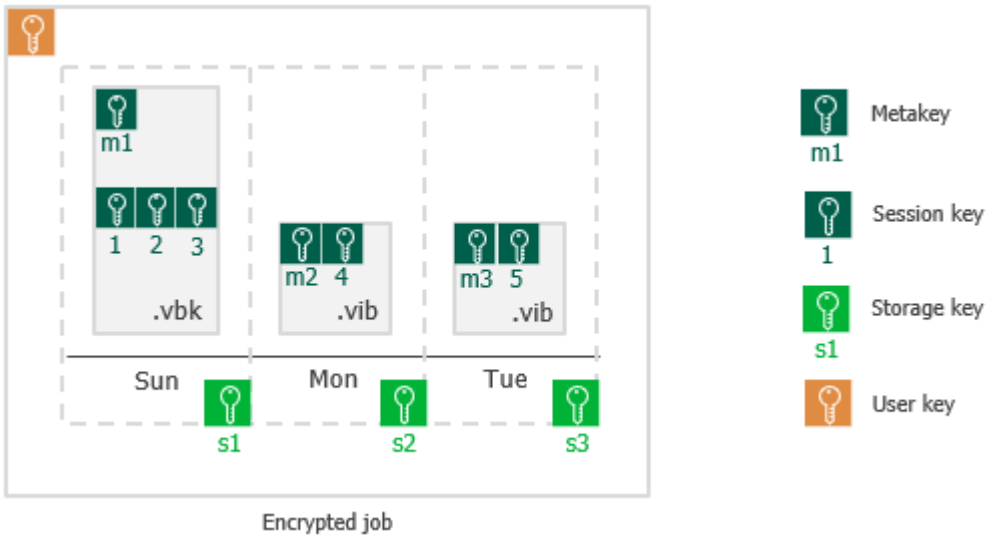
In the encryption process, storage keys are encrypted with a key of a higher layer — a user key. Cryptograms of storage keys are stored to the resulting file next to encrypted data blocks, and cryptograms of session keys and metakeys.

Storage keys are also kept in the Veeam Agent for Microsoft Windows database. To maintain a set of valid storage keys in the database, Veeam Agent for Microsoft Windows uses retention policy settings specified for the job. When some restore point is removed from the backup chain by retention, the storage key corresponding to this restore point is also removed from the Veeam Agent for Microsoft Windows database.

User Keys

When you enable encryption for a job, you must define a password to protect data processed by this job, and define a hint for the password. The password and the hint are saved in the job settings. Based on this password, Veeam Agent for Microsoft Windows generates a user key.

The user key protects data at the job level. In the encryption hierarchy, the user key encrypts storage keys for all restore points in the backup chain.



Encrypted job

Veeam Agent for Microsoft Windows saves a hint for the password to its database and to the backup metadata file (VBM). When you decrypt a file, Veeam Agent for Microsoft Windows displays a hint for the password that you must provide. After you enter a password, Veeam Agent for Microsoft Windows derives a user key from the password and uses it to unlock the storage key for the encrypted file.

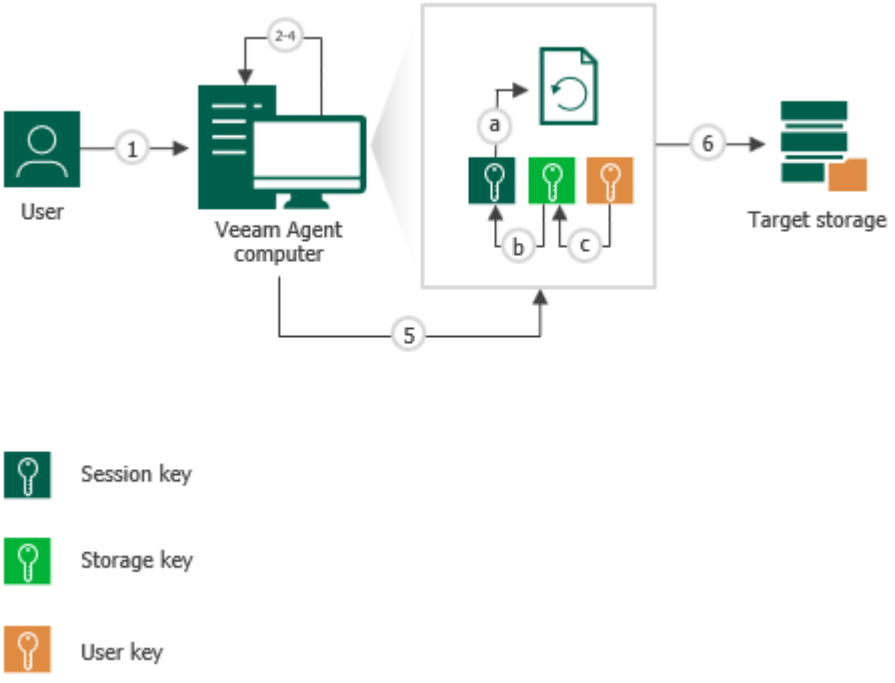
According to the security best practices, you should change passwords for encrypted jobs regularly. When you change a password for the job, Veeam Agent for Microsoft Windows creates a new user key and uses it to encrypt new restore points in the backup chain. If you lose a password that was specified for encryption, you can change the password in the encryption settings. You can use the new password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

How Data Encryption Works

Data encryption is performed as part of the backup process. Encryption works at the source side, before data is transported to the target location. As a result, encryption keys are not passed to the target side, which helps to avoid data interception.

The encryption process includes the following steps:

1. When you create a backup job, you enable the encryption option for the job and enter a password to protect data at the job level.
2. Veeam Agent for Microsoft Windows generates a user key based on the entered password.
3. When you start an encrypted job, Veeam Agent for Microsoft Windows creates a storage key and stores this key to its database.
4. Veeam Agent for Microsoft Windows creates a session key and a metakey. The metakey is stored to the Veeam Agent for Microsoft Windows database.
5. Veeam Agent for Microsoft Windows processes job data in the following way:
    a. The session key encrypts data blocks in the backup file. The metakey encrypts backup metadata.
    b. The storage key encrypts the session key and the metakey.
    c. The user key encrypts the storage key.
6. Encrypted data blocks are passed to the target. The cryptograms of the user key, storage key, session key and metakey are stored to the resulting file next to encrypted data blocks.



How Data Decryption Works

When you restore data from an encrypted backup file, Veeam Agent for Microsoft Windows performs data decryption automatically in the background or requires you to provide a password.

- If encryption keys required to unlock the backup file are available in the Veeam Agent for Microsoft Windows database, you do not need to enter the password. Veeam Agent for Microsoft Windows uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore does not differ from that from an unencrypted one.

  Automatic data decryption can be performed in one of the following situations:

  - You encrypt and decrypt the backup file on the same Veeam Agent computer using the same Veeam Agent for Microsoft Windows database.
  - You have included encryption keys into the Veeam Recovery Media and perform bare metal recovery after booting from this Veeam Recovery Media. To learn more, see Specify Recovery Media Options.

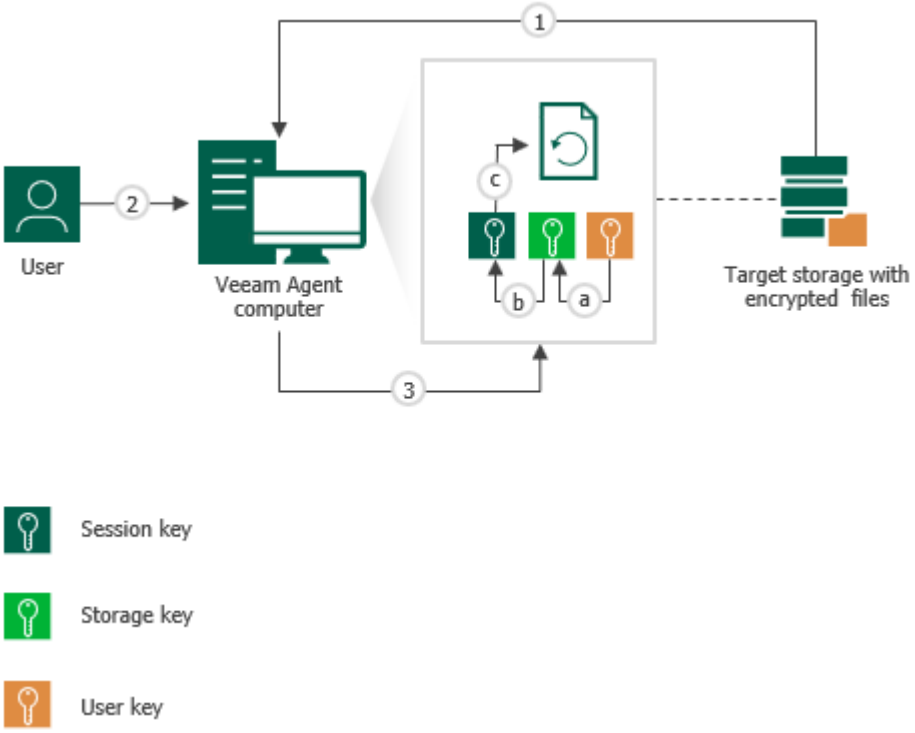
- If encryption keys are not available in the Veeam Agent for Microsoft Windows database, you need to provide a password to unlock the encrypted file.

Data decryption is performed at the source side, after data is transported back from the target side. As a result, encryption keys are not passed to the target side, which helps avoid data interception.

The decryption process includes the following steps. Keep in mind that steps 1 and 2 are required only if you decrypt the file on the Veeam Agent computer other than the computer where the file was encrypted.

1. You select the backup from which you want to restore data. Veeam Agent for Microsoft Windows notifies you that one or more files in the backup chain are encrypted and requires a password.
2. You specify a password for the imported file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent for Microsoft Windows, you can use the latest password to restore data form all restore points in the backup chain, including those restore points that were encrypted with an old password.
3. Veeam Agent for Microsoft Windows reads the entered password and generates the user key based on this password. With the user key available, Veeam Agent for Microsoft Windows performs decryption in the following way:
   a. Veeam Agent for Microsoft Windows applies the user key to decrypt the storage key.
   b. The storage key, in its turn, unlocks underlying session keys and a metakey.
   c. Session keys decrypt data blocks in the encrypted file.

After the encrypted file is unlocked, you can work with it as usual.

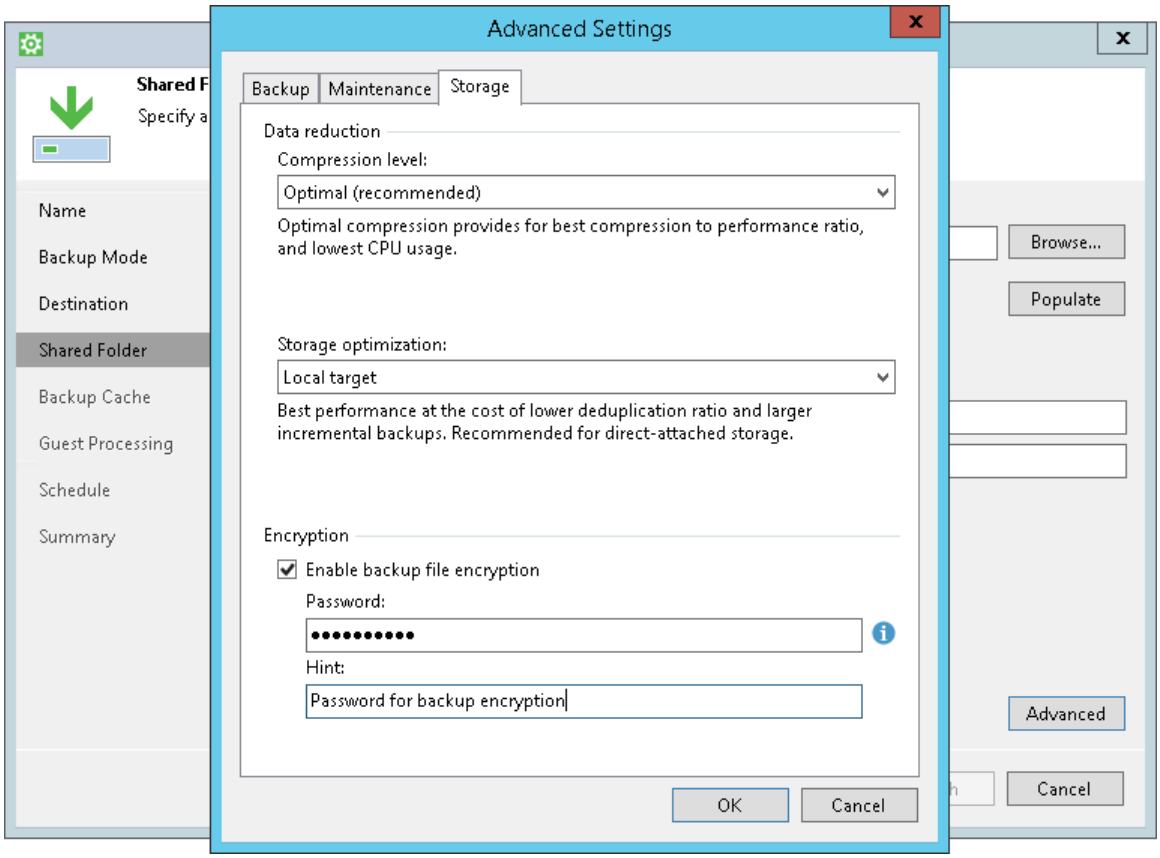


Backup Job Encryption

Encryption for the backup job is configured in the advanced job settings. You should enable the encryption option and specify a password to protect data in backup files produced by the backup job.

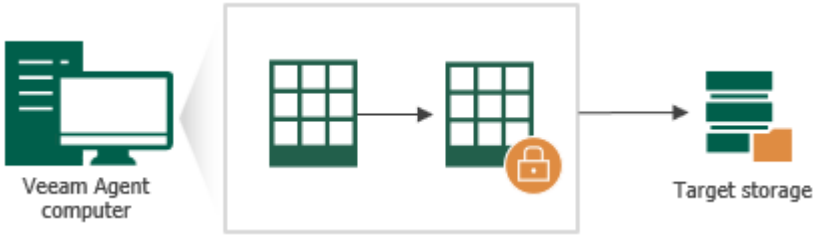
NOTE

You cannot specify encryption options for the backup job if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more, see Veeam Backup & Replication Documentation.



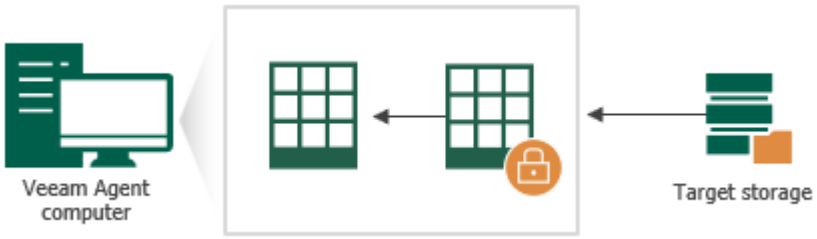
The backup job processing with encryption enabled includes the following steps:

1. You enable encryption for a backup job and specify a password.
2. Veeam Agent generates the necessary keys to protect backup data.
3. Veeam Agent encrypts data blocks and transfers them to the target location already encrypted.
4. On the target storage, encrypted data blocks are stored to a resulting backup file.



Restore of an encrypted backup file includes the following steps:

1. You select an encrypted backup file and define a password to decrypt the backup file. If the password has changed once or several times, you need to specify the latest password that was used to encrypt files in the backup chain.
2. Veeam Agent uses the provided password to generate user key and unlock the subsequent keys for backup file decryption.
3. Veeam Agent retrieves data blocks from the backup file, sends them to the target volume and decrypts them on the target volume.



Resuming Encrypted Backup Chain

In some situations, encryption keys may be unavailable in the Veeam Agent for Microsoft Windows database, and Veeam Agent for Microsoft Windows cannot create a new encrypted restore point in the backup chain. For example, this may happen after you change the password for encryption and then recover the entire Veeam Agent computer to a restore point that was created before you have changed the password. In this case, information about backup in the Veeam Agent for Microsoft Windows database will become outdated and will not match backup metadata residing on the target location. To continue the existing encrypted backup chain, you need to provide the latest password in the Veeam Agent control panel.

When the backup job is started (either manually or upon the defined schedule), Veeam Agent for Microsoft Windows detects the latest encrypted backup created by this job in the target location and displays a window in the Veeam Agent control panel offering to enter a password and continue the backup chain. You can choose to perform one of the following operations:

- To continue the existing encrypted backup chain, you can enter the password specified for encryption and click **OK**. If the password has changed more than once, you need to specify the latest password.

  After you provide the correct password, Veeam Agent for Microsoft Windows will use this password to decrypt backup metadata on the target location and update information about backup in its database. After that, Veeam Agent for Microsoft Windows will create the new incremental backup file in the existing encrypted backup chain. To encrypt this backup file and subsequent backup files, Veeam Agent for Microsoft Windows will use the password that is kept in its database at the time when continue the backup chain.

  You will be able to use this password to restore data from any restore point in the backup chain, including restore points that were encrypted with an older password and restore points that were created before you have enabled the encryption option for the job.

- If you do not remember the password or do not want to continue the existing backup chain for some reason, you can click **Start New** immediately without providing a password.

  When you start a new backup chain, Veeam Agent for Microsoft Windows creates the active full backup file. All incremental backup files use this full backup file as a new starting point. Depending on the defined encryption settings, Veeam Agent for Microsoft Windows creates a full backup file according to the following rules:

  - If the encryption option is disabled in Veeam Agent for Microsoft Windows at the time when you start a new backup chain, Veeam Agent for Microsoft Windows creates an unencrypted full backup file. Subsequent incremental backup files will be unencrypted, too.

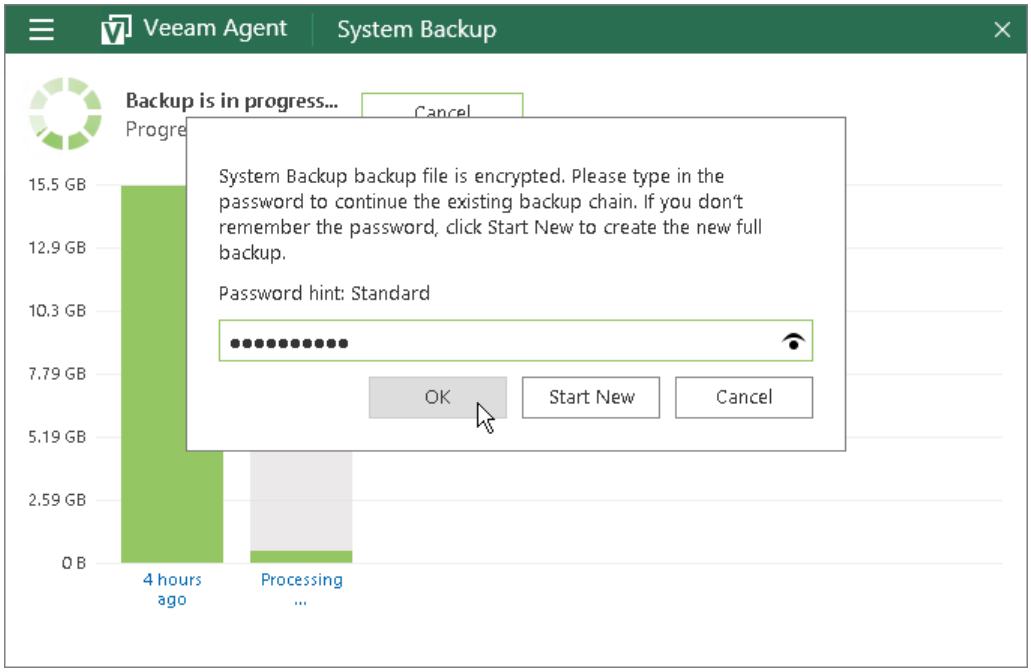
- If the encryption option is enabled in Veeam Agent for Microsoft Windows at the time when you start a new backup chain, Veeam Agent for Microsoft Windows creates an encrypted full backup file. To encrypt this backup file, Veeam Agent for Microsoft Windows uses the password that exists in its database at the time when you start a new backup chain.

- You can click **Cancel** to close the notification window and cancel the job. In this case, the next time the backup job is started, Veeam Agent for Microsoft Windows will again prompt to enter the password.

NOTE

Veeam Agent for Microsoft Windows displays the notification window for 1 hour. If you do not choose to continue the existing backup chain, start a new backup chain or cancel the backup job, after this time period expires, Veeam Agent for Microsoft Windows will automatically close the window and cancel the job. If email notification settings are enabled for the backup job, Veeam Agent for Microsoft Windows will also send an email report informing that the job was canceled.

Veeam Agent for Microsoft Windows will display this window every time the backup job is started (manually or upon schedule) until you choose to continue the existing backup chain or start a new backup chain.



Encryption Best Practices

To guarantee the flawless process of data encryption and decryption, consider the following advice.

Password

1. Use strong passwords that are hard to crack or guess. Consider the following recommendations:
    a. The password must be at least 8 characters long.
    b. The password must contain uppercase and lowercase characters.
    c. The password must be a mixture of alphabetic, numeric and punctuation characters.
    d. The password must significantly differ from the password you used previously.
    e. The password must not contain any real information related to you, for example, date of birth, your pet's name, your logon name and so on.
2. Provide a meaningful hint for the password that will help you recall the password. The hint for the password must significantly differ from the password itself. The hint for the password is displayed when you select an encrypted backup server and attempt to unlock it.
3. Change passwords for encrypted jobs regularly. Use of different passwords helps increase the encryption security level.

Encryption for Existing Job

If you enable encryption for an existing job, during the next job session Veeam Agent for Microsoft Windows will create active full backup. 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 backup job, Veeam Agent for Microsoft Windows does not encrypt the previous backup chain created with this job. However, Veeam Agent for Microsoft Windows encrypts backup metadata. As a result, you need to enter the password to restore data from unencrypted backup files in the backup chain as well as from encrypted backup files in this chain.

Backup Cache

A remote storage specified as a target location for backup files may be unavailable at the time when the backup job must start. In this case, Veeam Agent cannot create a regular restore point upon the defined schedule. As a result, the backup chain on the remote storage will not contain a sequence of restore points that precisely complies with the backup schedule.

To overcome this limitation, Veeam Agent offers the concept of the backup cache. The backup cache is a temporary local storage in which Veeam Agent creates backup files in case the backed-up data cannot be transferred to a remote location. When the target location becomes available, Veeam Agent uploads backup files from the backup cache to the remote storage, adding regular restore points to the backup chain.

The backup cache lets you perform scheduled backup in due time ensuring that the resulting backup chain will contain "snapshots" of your data at desired points in time. This may be helpful, for example, for laptop users who go on business trips with no or limited access to the corporate network in which the backup location resides.

Technically, the backup cache is a local folder on the computer on which Veeam Agent is installed. A user can define a folder for the backup cache and the size of the backup cache in the backup job settings.

The backup cache is available if the following types of storage are chosen as a target location:

- Network shared folder
- Microsoft OneDrive storage
- Backup repository managed by a Veeam backup server
- Cloud repository managed by a Veeam Cloud Connect service provider

How Backup Cache Works

When you create a backup job targeted at a remote storage, you can select to use the backup cache in its properties. The procedure of data backup with the backup cache enabled is performed in the following way:

1. The user enables the backup cache in the properties of the backup job targeted at a remote location.
2. During a regular backup job session, Veeam Agent for Microsoft Windows creates in the backup cache a map of data blocks on the remote location. Information about data blocks on the remote location is saved in a file with digests and stored in the folder `C:\ProgramData\Veeam\EndpointData\CacheDigests`.

   Veeam Agent for Microsoft Windows will use the created digests to create incremental backup files in the backup cache when the remote storage itself is unavailable at the time of scheduled backup.
3. If the target location is unavailable at the time when the scheduled backup job must start, Veeam Agent for Microsoft Windows creates the new restore point in the backup cache.

   The target location is considered as unavailable in the following conditions:

   - [For network shared folder] Veeam Agent for Microsoft Windows Service that runs on the protected computer cannot connect to the network shared folder. In this case, Veeam Agent for Microsoft Windows will immediately start creating the new restore point in the backup cache.
   - [For Veeam backup repository] Veeam Agent for Microsoft Windows Service cannot connect to the Veeam Backup Service that runs on the backup server to which the backup repository specified as a target location for Veeam Agent backups is connected.
   - [For Veeam Cloud Connect repository] Veeam Agent for Microsoft Windows Service cannot connect to one of the following services on the Veeam Cloud Connect provider side:
     - *Veeam Backup Service* that runs on the backup server used for managing the Veeam Cloud Connect infrastructure.
     - *Veeam Cloud Connect Service* that runs on the backup server used for managing the Veeam Cloud Connect infrastructure.
     - *Veeam Cloud Gateway Service* that runs on a cloud gateway deployed in the Veeam Cloud Connect infrastructure.
   - [For Microsoft OneDrive] Veeam Agent for Microsoft Windows Service cannot connect to Microsoft OneDrive API.

   In case the connection to the target location is lost when the backup job is already running, Veeam Agent for Microsoft Windows performs backup based on the following rules:

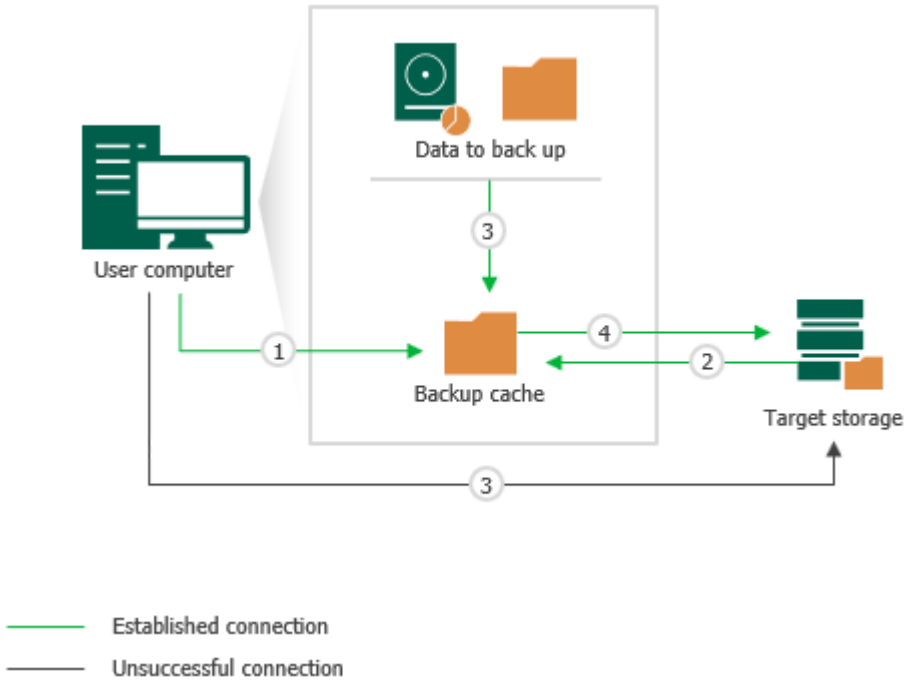
   - [For network shared folder] Veeam Agent for Microsoft Windows immediately switches to the backup cache and writes to the backup cache all data that must be backed up within the current backup job session. If some data has already been transferred to the target location, Veeam Agent for Microsoft Windows starts the data transfer operation from the beginning and transfers all backed-up data to the backup cache.
   - [For Veeam backup repository, Veeam Cloud Connect repository and Microsoft OneDrive] Veeam Agent for Microsoft Windows tries to reconnect to the target location. The reconnection process may last 30 minutes or more. After the reconnection period expires, Veeam Agent for Microsoft Windows switches to the backup cache and writes to the backup cache only remaining data that has not been transferred to the target location yet. Data that has been already transferred to the target location remains on the backup repository.

   The process of data backup to the backup cache practically does not differ from the regular one. The difference is that the resulting backup file is saved to the local folder instead of the remote storage.

If the target location becomes available after Veeam Agent for Microsoft Windows has started creating a restore point in the backup cache, Veeam Agent for Microsoft Windows will not switch back to the target location. Instead, Veeam Agent for Microsoft Windows will create a restore point in the backup cache and then upload this restore point to the target location.

4. After the remote storage becomes available, Veeam Agent for Microsoft Windows uploads backup files that were created in the backup cache to the target location. If more than one restore point was created in the backup cache, these restore points are uploaded to the target location sequentially, one by one.

   Until all restore points are uploaded from the backup cache to the target location, Veeam Agent for Microsoft Windows will continue to create new restore points in the backup cache even if the target location is available at the time when the backup job is running.



Limitations for Backup Cache

The backup cache has the following limitations:

- The backup cache functionality is available only in Workstation and Server editions of Veeam Agent for Microsoft Windows.
- You cannot use the backup cache for a file-level backup job.
- You cannot restore data from backup files that reside in the backup cache. Until restore points are uploaded from the backup cache to the target location, these restore points are not considered as a fully valid part of the backup chain.
- Veeam Agent for Microsoft Windows does not support creating full backups (including active full backups and synthetic full backups) in the backup cache except for the very first full backup file in the backup chain.
- Veeam Agent for Microsoft Windows does not support creating encrypted backups in the backup cache. If encryption options are specified for the backup job, Veeam Agent will create unencrypted backup files in the backup cache. When the target location becomes available, Veeam Agent will encrypt data prior to uploading it to the remote storage.
- Veeam Agent for Microsoft Windows does not support creating transaction log backups in the backup cache. You cannot enable transaction log backup and the backup cache for the backup job simultaneously.
- Restore points created in the backup cache cannot be uploaded to the target location after you perform the following operations:
    - Change target location for backup files in the properties of the backup job
    - Change a folder defined for the backup cache
    - Move or delete backup files on the target location
    - [For Veeam backup repository] Enable or disable data encryption settings in Veeam Backup & Replication.

    If the backup cache contains one or more restore points at the time when you perform one of these operations, you need to delete restore points from the backup cache.
- You cannot delete restore points from the backup cache while Veeam Agent for Microsoft Windows is performing the following operations:
    - Creating a new restore point in the backup cache
    - Uploading a restore point from the backup cache to the target location

Backup to Rotated Drives

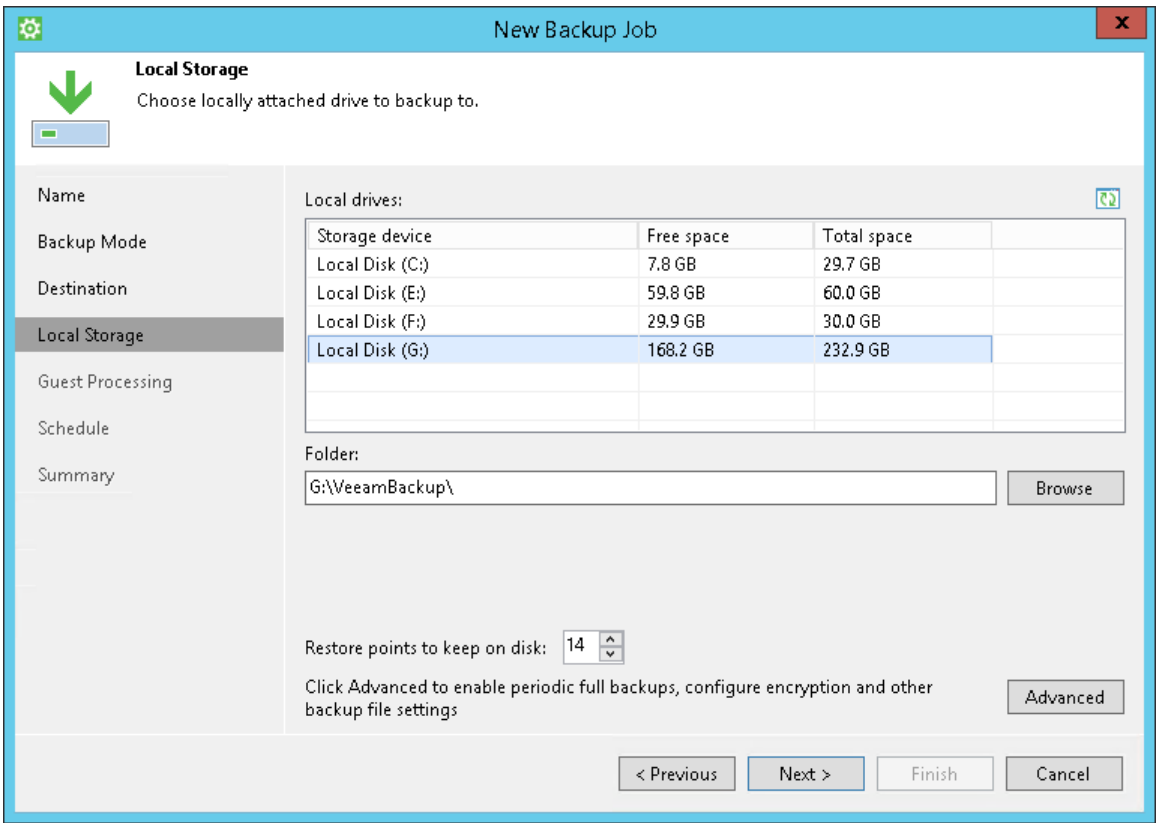
You can use rotated drives as a target location for backups. This scenario can be helpful if you want to store backups on several external hard drives (for example, USB or FireWire) and plan to swap these drives between different locations regularly.

Backup on rotated drives is performed in the following way:

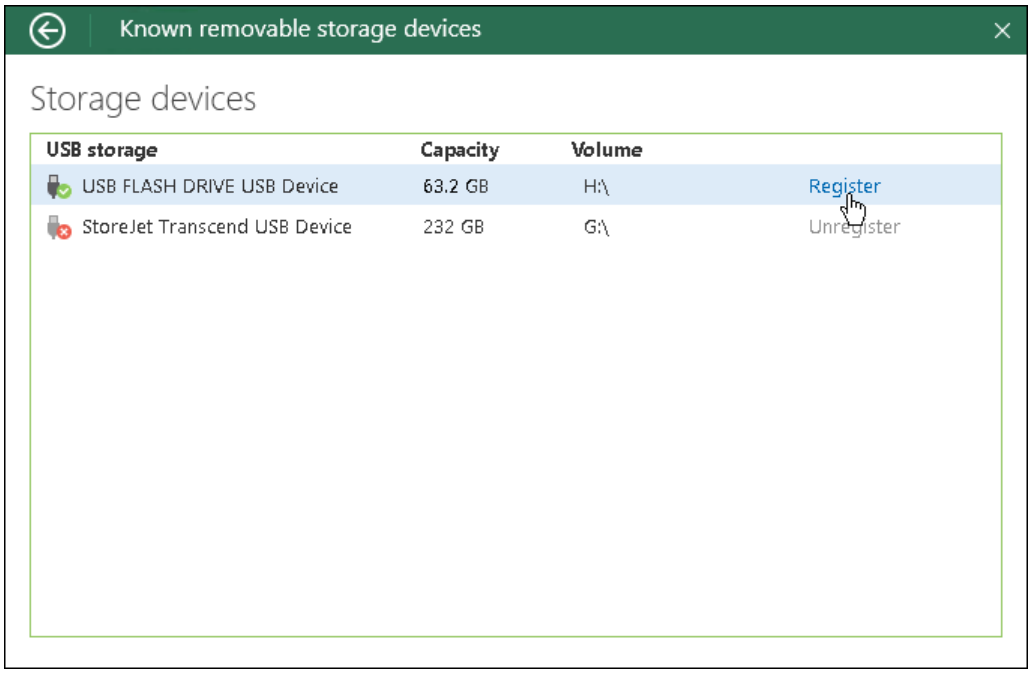
1. Veeam Agent for Microsoft Windows creates a backup chain on an external drive that you use as a backup target. The backup chain consists of the first full backup and a set of subsequent incremental backups.
2. When you swap drives and attach a new external drive, Veeam Agent creates a separate backup chain on the new drive.
3. After you swap drives again, Veeam Agent detects if there is a backup chain on the currently attached drive. If the backup chain exists, Veeam Agent continues the existing chain: it creates a new incremental backup file and adds it to the existing backup files.

To use rotated drives for backup, you must perform the following actions:

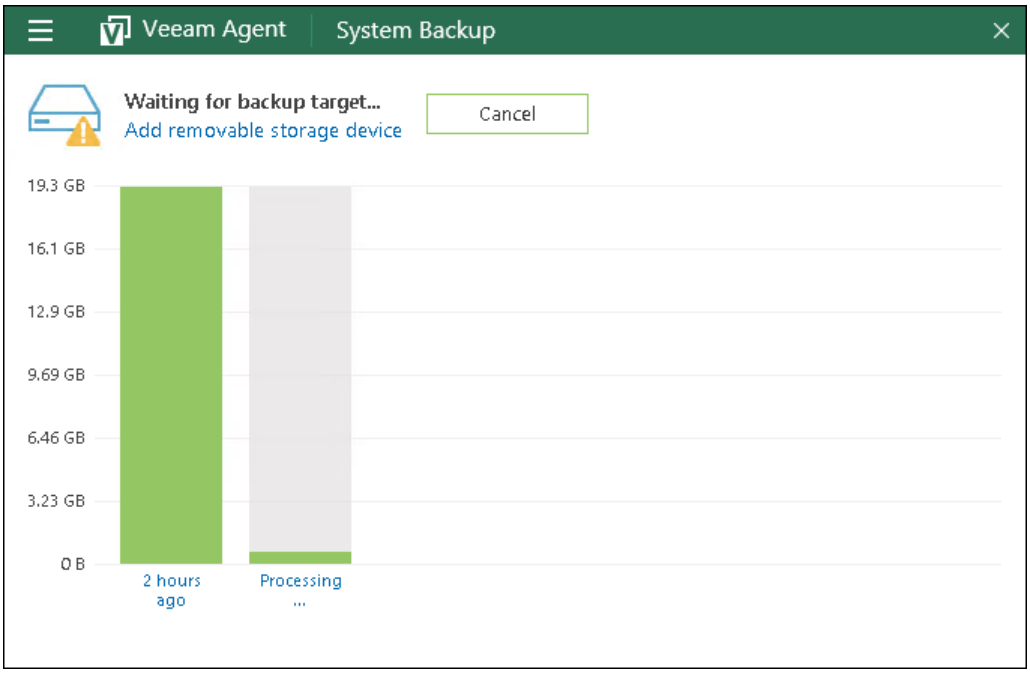
1. Attach one of external drives from the set to your computer.
2. Configure the backup job to store backups on the currently connected external drive. To do this:
   a. At the **Local Storage** step of the wizard, select the connected drive.
   b. From the **Local drives** list, select the necessary volume on the connected drive and specify a folder where backups must be stored.
   c. Save the job settings.



3. When you need to swap files, disconnect the drive that was used previously and attach a new drive to your computer.
4. Register a newly connected drive as a known removable storage in Veeam Agent for Microsoft Windows. To do this:
   a. Double-click the Veeam Agent icon in the system tray to open the control panel.
   b. From the main menu, select **Settings.**
   c. Click the **Manage registered storage devices** link.
   d. Click **Register** next to the newly connected drive.



If you do not register the newly connected drive before the backup job starts, Veeam Agent will be unable to detect the backup target and launch the backup job. Veeam Agent will display a warning in the system tray and in the control panel. To register a new device, click the **Add removable storage device** link in the **Status** view of the control panel and register the newly connected drive as described above. To learn more, see Managing Rotated Drives.



5. After you register the newly connected drive, you can start a new backup session manually or wait Veeam Agent to start a new session.

NOTE

Keep in mind that you cannot target a backup job at a certain set of rotated drives. If the backup job to rotated drive starts, and the drive that was used previously is not connected, Veeam Agent creates a backup file on any registered drive that is currently attached to the Veeam Agent computer.

Backup of External Drives

You can use Veeam Agent for Microsoft Windows to back up data that resides on external USB/eSATA or FireWire drives connected to the Veeam Agent computer.

> **IMPORTANT**
>
> Veeam Agent supports backup of external drives that support Microsoft VSS: HDD, SSD, and so on. USB flash drives (USB sticks) are not supported, because it is impossible to create a VSS snapshot on a device of this type. To check if your drive is supported, see this Veeam KB article.

You can include an external drive in any type of backup: entire computer backup, volume-level backup or file-level backup.

- For entire computer backup, you can instruct Veeam Agent for Microsoft Windows to back up data of all external drives the are currently connected to the Veeam Agent computer. To do this, you must enable the **Include external USB drives** option at the **Backup Mode** step of the **New Backup Job** wizard.

  An external drive whose data you want to back up must be connected to the Veeam Agent computer at the time when the backup job must start.

  If an external drive that was backed up within a previous backup job session is not connected to the Veeam Agent computer when the backup job must start, Veeam Agent for Microsoft Windows will back up data of all existing drives and skip the absent one. The backup job session will complete successfully.

- For volume-level backup, you can include in the backup a specific volume that resides on an external drive. To do this, you must select the necessary volume at the **Volumes** step of the **New Backup Job** wizard.

  If a USB drive that contains the volume included in the backup is not connected to the Veeam Agent computer when the backup job must start, the backup job will complete with an error.

- For file-level backup, you can include in the backup data of a specific volume that resides on an external drive (entire volume or specific folders with files). To do this, you must select the necessary object at the **Files** step of the **New Backup Job** wizard.

  If a USB drive that contains a volume or folder included in the backup is not connected to the Veeam Agent computer when the backup job must start, the backup job will complete with an error.

> **NOTE**
>
> Veeam Agent for Microsoft Windows does not back up external drives used as rotated drives. USB drives registered as a known removable storage in Veeam Agent for Microsoft Windows are excluded from the backup.

Data Restore

Veeam Agent for Microsoft Windows offers two data restore scenarios:

- You can perform volume-level restore to recover the entire system image of your computer or specific computer volumes. To learn more, see Volume-Level Restore.
- You can perform file-level restore to recover individual files and folders. To learn more, see File-Level Restore.

When performing volume-level restore, you can resize restored volumes to fit available space on target location. To learn more, see Volume Resize.

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 Microsoft Windows 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.

Keep in mind that you cannot browse the volume in the backup and select individual application items, files and folders for restore. For granular file-level restore, you can use the File-Level Restore option.

A volume can be restored to its original location or new location. If you restore the volume to its original location, Veeam Agent for Microsoft Windows overwrites data on the original volume. If you restore the volume to a new location, and the target disk contains any data, Veeam Agent for Microsoft Windows overwrites data in the target location with data retrieved from the backup.

A volume can be restored to a new location that has greater or less space than the size of the volume in the backup. Depending on the amount of free disk space on target location, you can select either to shrink or to extend the volume during restore. To learn more, see Volume Resize.

Limitations for Volume-Level Restore

Volume restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the swap file is currently hosted.
- You cannot restore a volume to the volume where the backup file used for restore is located.

To overcome the first two limitations, you can create a Veeam Recovery Media and use the **Veeam Bare Metal Recovery** wizard for volume-level restore. To learn more, see Veeam Recovery Media.

File-Level Restore

If you have lost or modified files and folders 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 Microsoft Windows does not extract files and folders from the backup file. Instead, it uses Veeam's proprietary driver to publish the backup content directly into the computer file system, under `C:\VeeamFLR\<Volume N>`. For accessing the backup file content, Veeam Agent for Microsoft Windows uses a separate program — Virtual Disk Driver (VDK) that is provided with the product.

After the backup content is mounted, you can use a built-in Veeam Backup browser or Microsoft Windows Explorer to browse and copy necessary files and folders to your local machine drive, save them in a network shared folder or point applications to files and work with them in a regular way.

Volume Resize

With Veeam Agent for Microsoft Windows, you can resize backup volumes during Volume-Level Restore. When you select to resize a volume, Veeam Agent for Microsoft Windows restores data from the backup and resizes the restored volume to the specified size.

There are two ways to resize a volume depending on the amount of free disk space on the target location:

- **Volume shrink** — you can shrink a volume when you restore it to a new location that has less space than the size of the volume in the backup. You can also shrink a volume that is restored to its original location to free disk space on the target location. To learn more, see How Volume Shrink Works.
- **Volume extend** — you can extend a backup volume when you restore it to a new location that has more available disk space than the size of the backup volume. To learn more, see How Volume Extend Works.

Volume resize may be also helpful when you need to restore data after hardware upgrade. For example, you may want to resize volumes in the following situations:

- Shrink backup data to restore system volumes of your computer to a smaller disk after you replace an old HDD drive with a faster but less capacitive SSD drive.
- Extend the backup volume during volume-level restore to a new, more capacitive HDD drive.

You can restore and resize volumes:

- With the **Volume Restore** wizard when Restoring Volumes under Microsoft Windows system.
- With the **Veeam Bare Metal Recovery** wizard when Restoring from Veeam Recovery Media.

The volume resize option is available only in the **Manual restore** mode at the **Disk Mapping** step of the wizard.

Limitations for volume resize

Volume resize has the following limitations:

- You cannot restore a volume to the volume of the smaller size if the amount of data stored on the backup volume exceeds the free space on the target disk.
- You can only resize basic volumes that use the NTFS file system.
- If you resize a BitLocker encrypted volume during restore, the restored volume will be unencrypted.

How Volume Shrink Works

When you restore a volume to a target location of the smaller size, Veeam Agent for Microsoft Windows performs the following operations:

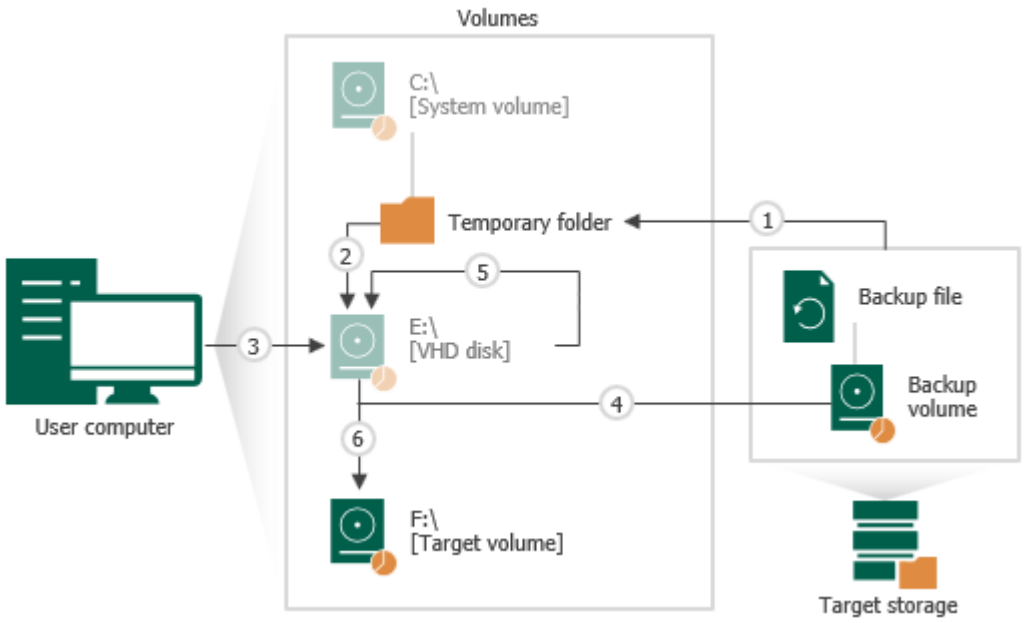
1. When you select the **Resize** option to shrink a volume, Veeam Agent for Microsoft Windows mounts the backup volume to a temporary NTFS folder on the system drive, for example: `C:\Users\Username\AppData\Local\Temp`.
2. Veeam Agent for Microsoft Windows mounts the created NTFS folder as a VHD disk next to other disks that are present on the computer.

   Mounting VBK file content as a VHD disk makes it possible for Veeam Agent for Microsoft Windows to use Microsoft Windows system's disk management tools to measure current size of the backup volume and maximum and minimum size for the restored volume.
3. Veeam Agent for Microsoft Windows sends a query request to the mounted VHD disk to calculate its size, amount of stored data and free disk space by which the volume can be shrunk.

   This step may take some time depending on the size of the backup volume and its data fragmentation ratio.

   When the query is complete and you specify the desired size for the restored volume, Veeam Agent for Microsoft Windows unmounts the VHD disk.
4. When you start the restore process, Veeam Agent for Microsoft Windows creates on the target disk a volume of the specified size and restores to that volume the amount of backed-up data that fits the specified size.
5. Veeam Agent for Microsoft Windows mounts the backup volume as a VHD disk as described in steps 1 and 2 and starts to shrink it to the size of the target volume. During the process of volume shrink, empty data blocks from the part of the mounted VHD disk that does not fit the size of the target volume are moved to the part of the disk that contains actual data.
6. Veeam Agent for Microsoft Windows captures on the VHD disk data blocks that are moved during shrink and writes them to the target volume.

   When all data blocks are written to the target volume, Veeam Agent for Microsoft Windows unmounts the VHD disk.



How Volume Extend Works

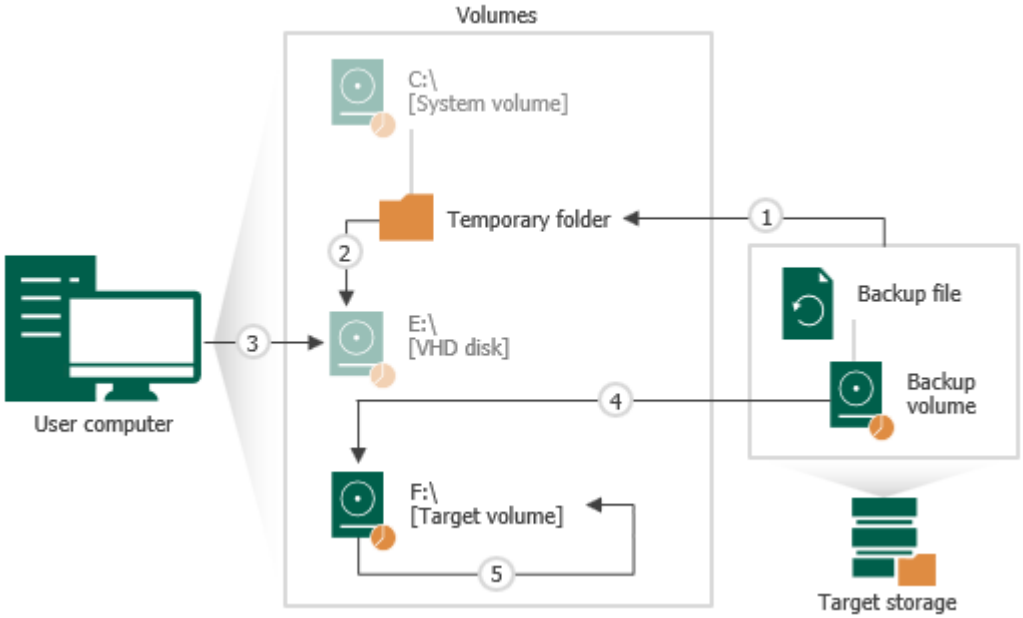
When you restore a volume to a target location of the larger size, Veeam Agent for Microsoft Windows performs the following operations:

1. When you select the **Resize** option to extend a volume, Veeam Agent for Microsoft Windows mounts the backup volume to a temporary NTFS folder on the system drive, for example: `C:\Users\Username\AppData\Local\Temp`.
2. Veeam Agent for Microsoft Windows mounts the created NTFS folder as a VHD disk next to other disks that are present on the computer.

   Mounting VBK file content as a VHD disk makes it possible for Veeam Agent for Microsoft Windows to use Microsoft Windows system's disk management tools to measure current size of the backup volume and maximum and minimum size for the restored volume.
3. Veeam Agent for Microsoft Windows sends a query request to the mounted VHD disk to calculate its size, amount of stored data and free disk space by which the volume can be extended.

   This step may take some time depending on the size of the backup volume and its data fragmentation ratio.

   When the query is complete and you specify the desired size for the restored volume, Veeam Agent for Microsoft Windows unmounts the VHD disk.
4. When you start the restore process, Veeam Agent for Microsoft Windows creates on the target disk a volume of the same size as the backup volume and restores to that volume all data blocks from the backup volume.
5. When all data blocks are written to the target location, Veeam Agent for Microsoft Windows extends the size of the target volume to the specified size.



Veeam Recovery Media

Veeam Agent for Microsoft Windows lets you create a Veeam Recovery Media — a recovery image of your computer.

The recovery image is a "copy" of your OS with the limited functionality — it contains all data required to run Microsoft Windows Recovery Environment (Windows RE) and provides an alternative way to boot your computer. If the OS installed on the computer fails to start for some reason, you can boot the Windows RE from the recovery image. After booting, you can do the following:

- You can use Veeam Agent for Microsoft Windows and Microsoft Windows tools to diagnose problems and fix errors on your computer.
- You can restore data from a backup to your computer. For this scenario, you must have a backup created with Veeam Agent for Microsoft Windows or system image created with Microsoft Windows.

The recovery image can be helpful if one of the following errors occur:

- The OS on the computer fails to start.
- The computer is blocked with malware and you cannot get access to your data.
- You want to perform bare metal restore from the backup on the computer without the OS and other software installed.
- You want to restore the system volume of the computer and so on.

You can create a recovery image on different kinds of media:

- Removable storage devices such as USB drives or SD cards
- CD/DVD/BD
- ISO images on local or external computer drives

When you boot from the Veeam Recovery Media, you can use the Veeam Agent for Microsoft Windows recovery environment to fix the OS system errors on your computer or restore data from the backup. Veeam Agent for Microsoft Windows offers a set of tools for the computer system image and data recovery:

- Bare Metal Recovery — the Veeam Agent for Microsoft Windows wizard to recover data on the original computer or a new computer.
- Windows Recovery Environment — a built-in Microsoft Windows tool to recover the computer system image.
- Tools — Veeam Agent for Microsoft Windows and Microsoft Windows utilities for advanced computer administration.

Limitations for Veeam Recovery Media

- You cannot restore dynamic volumes using a Veeam Recovery Media. To restore dynamic volumes, you can recover data from the volume-level backup on a working computer system. To learn more, see Restoring Volumes.
- The Veeam Recovery Media is based on the Microsoft Windows RE. Due to Microsoft limitations, Microsoft Windows RE automatically reboots after 72 hours of continuous use. All data that has not been saved before reboot will be lost.

Drivers in Veeam Recovery Media

The Veeam Recovery Media created with Veeam Agent for Microsoft Windows contains the following data:

- Set of files required to start your computer OS from the recovery media.
- Diagnostic tools from Microsoft and Veeam.
- Drivers required to run hardware and devices on your computer in a regular way. When you boot your computer from the Veeam Recovery Media, drivers included into the Veeam Recovery Media are automatically loaded on the recovered OS.
- Network connection settings from your computer. When you boot your computer from the Veeam Recovery Media, network settings included into the Veeam Recovery Media are automatically applied and can be used to connect to the remote backup storage.
- If you have enabled data encryption options for the backup job, you can also include a decryption key into the Veeam Recovery Media. To learn more, see Creating Veeam Recovery Media.
- Web browser required to log in to a Microsoft OneDrive account. If your backed-up data resides in Microsoft OneDrive, you can boot your computer from the Veeam Recovery Media and recover the necessary data from Microsoft OneDrive. To learn more, see Microsoft OneDrive Support.

> **NOTE**
>
> The Veeam Recovery Media contains all locales (languages) that are included in the OS of the Veeam Agent Computer.

You can include the following drivers in the Veeam Recovery Media:

- **Drivers that are currently installed on your computer**. Veeam Agent for Microsoft Windows detects hard disk controller drivers, network adapter drivers and USB controller drivers and includes them into the Veeam Recovery Media.
- **Additional storage and network drivers**. If you use non-standard drivers, you can include them in the created Veeam Recovery Media manually. For example, you can include drivers for a discrete network card, third-party USB 3.0 controllers and non-standard hard disk controllers.

> **TIP**
>
> If you do not include some drivers in the Veeam Recovery Media, you can load them from the computer drive when you perform bare metal recovery. To learn more, see Restoring from Veeam Recovery Media.

BitLocker Encrypted Volumes Support

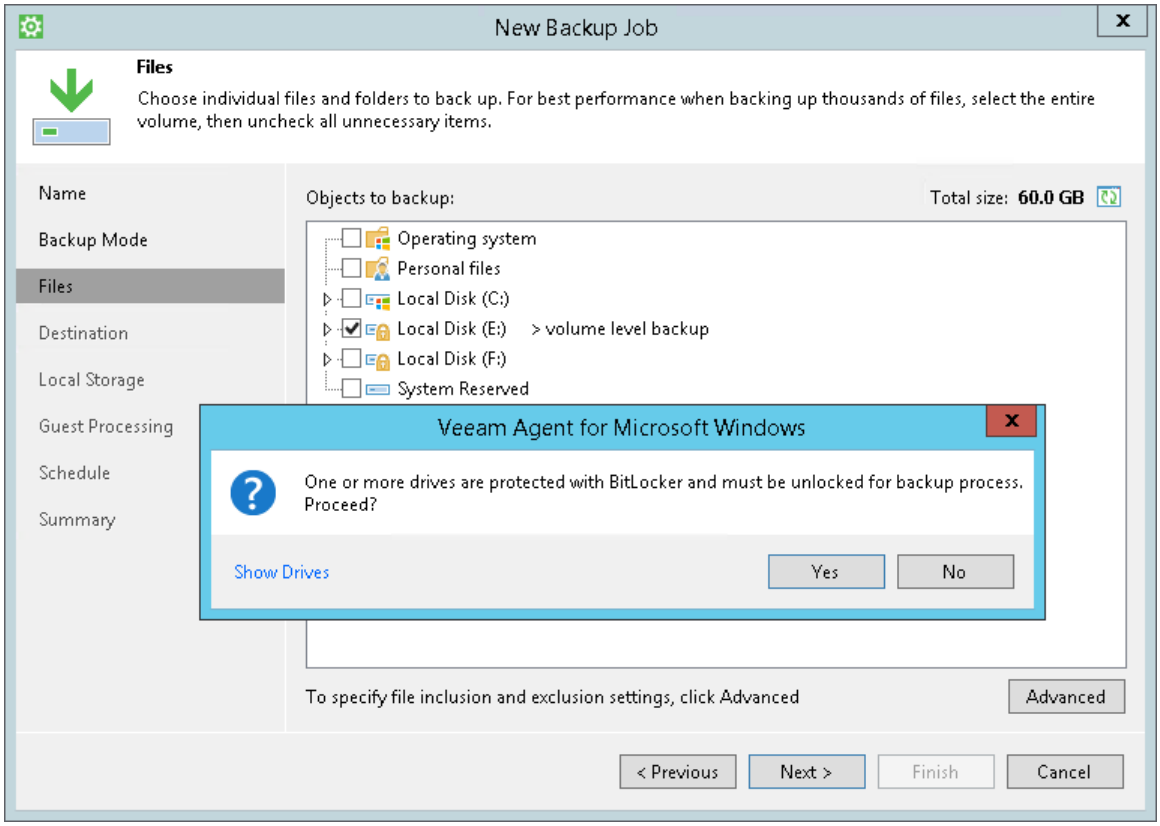
Veeam Agent for Microsoft Windows supports scenarios of data backup and restore to/from volumes encrypted with Microsoft Windows BitLocker.

Data Backup

You can create backups of BitLocker encrypted volumes and store backups created with Veeam Agent for Microsoft Windows on BitLocker encrypted volumes.

BitLocker encrypted volumes (both source and target) must be unlocked at the moment when Veeam Agent for Microsoft Windows starts the backup operation.

- If the volume added to the backup scope is locked at the moment of backup, the backup job will be unable to process it and will fail.
- If the volume to which the backup file must be stored is locked at the moment of backup, the backup job will be unable to save the resulting file, and the job will fail.



Data Restore

You can restore data from backups stored on BitLocker encrypted volumes and restore data to BitLocker encrypted volumes.

Veeam Agent for Microsoft Windows restores volumes in their initial state:

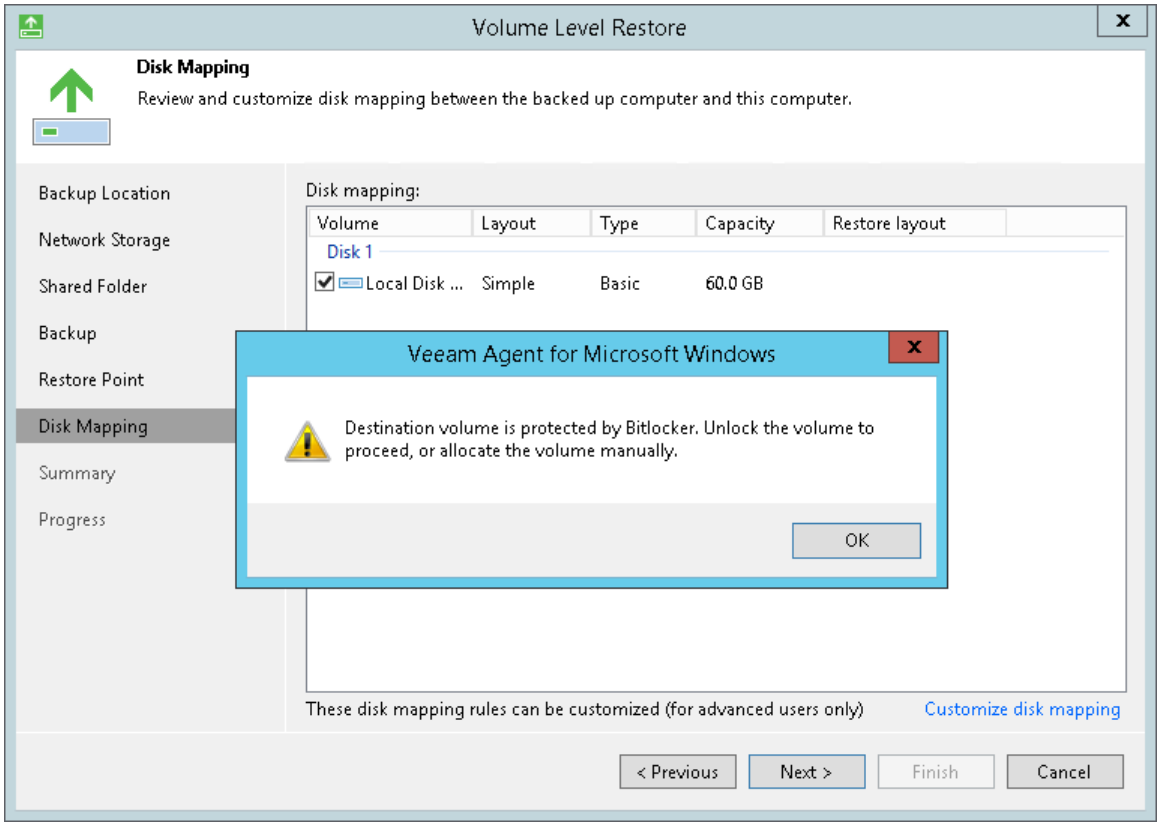
- If you restore an encrypted volume to its original location, the restored volume will be encrypted.
- If you restore an unencrypted volume to an encrypted volume, the restored volume will be unencrypted.

IMPORTANT

If you resize a BitLocker encrypted volume during restore, the restored volume will be unencrypted. To learn more about volume resize, see Volume Resize.

BitLocker encrypted volumes must be unlocked at the moment when you perform the restore operation.

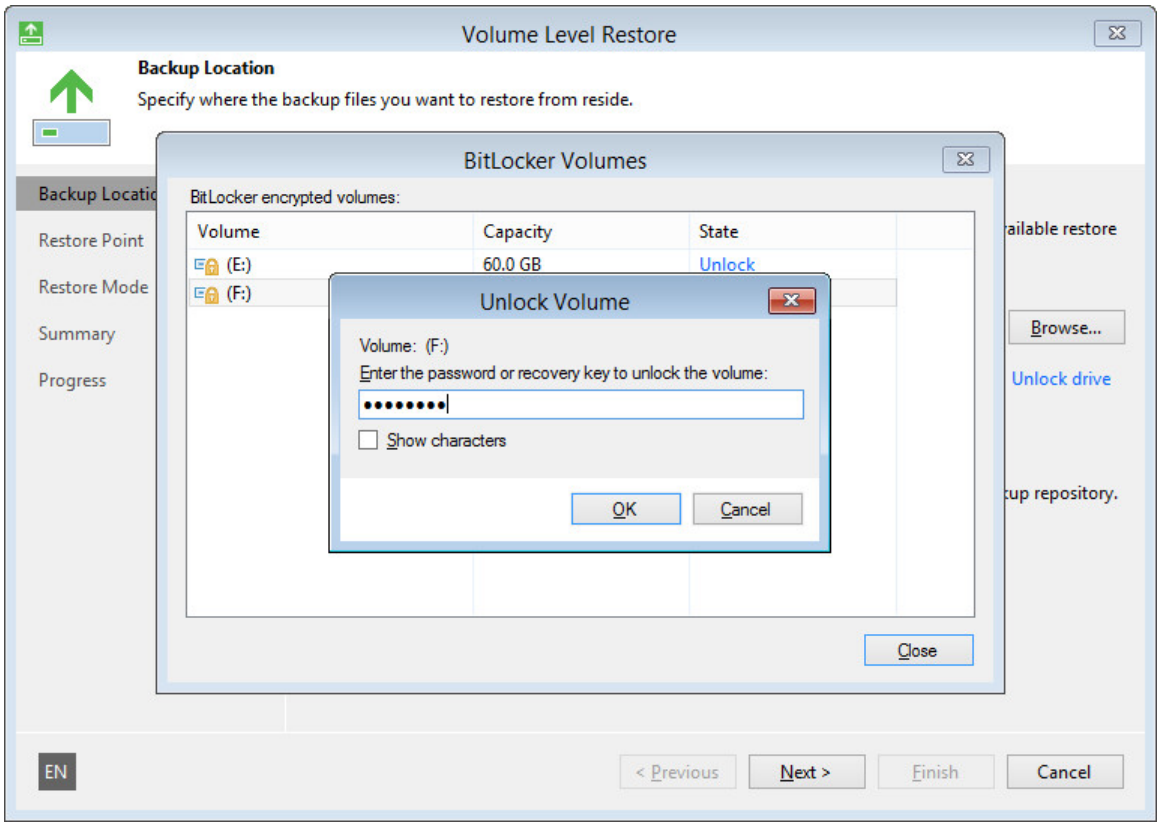
- If the backup file is stored on a locked volume, Veeam Agent for Microsoft Windows will fail to access it, and you will not be able to restore data from it.
- If you perform volume-level restore, and the target volume is locked, Veeam Agent for Microsoft Windows will display a warning and will ask you to unlock the volume. You can do this using the Microsoft Windows UI.



Veeam Recovery Media

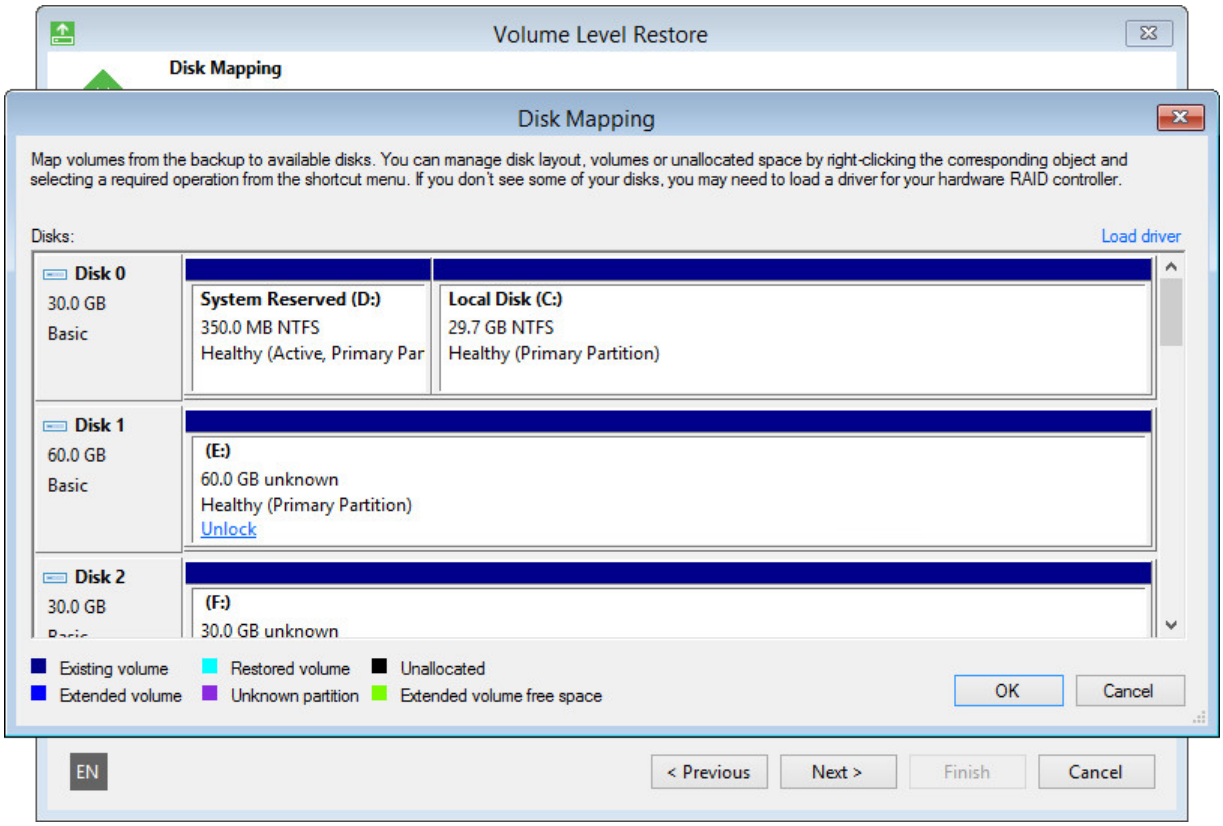
If you boot from the Veeam Recovery Media, you can restore data from backups stored on BitLocker encrypted volumes and restore data to BitLocker encrypted volumes.

- If the backup file that you want to use for data restore resides on a locked volume, Veeam Agent for Microsoft Windows cannot access this backup file. To unlock the volume with the backup file, click **Unlock drive** under the **Backup file** field and enter a password for the volume.

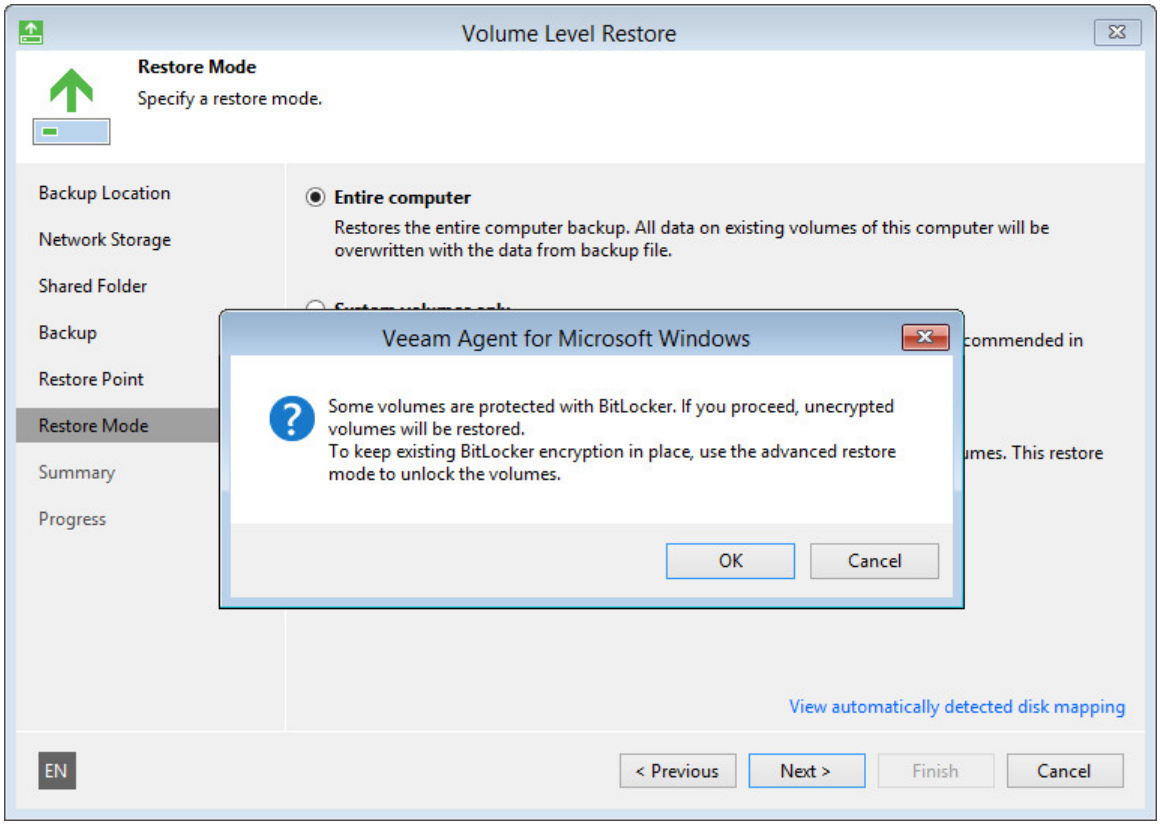


- If the target volume is BitLocker encrypted and locked, at the **Restore Mode** step of the wizard, Veeam Agent for Microsoft Windows displays a warning informing about it. You can use one of the following scenarios:
  - You can restore data to the target volume and keep BitLocker encryption enabled for the volume. To do this, you must unlock the volume before you start data restore.

    To unlock the volume, click **Cancel** in the warning window. At the **Restore Mode** step of the wizard, select **Manual Restore**. At the **Disk Mapping** step of the wizard, click **Customize disk mapping** and click **Unlock** under the necessary volume.



- You can restore data to the target volume and disable BitLocker encryption for the volume. To do this, click **OK** in the warning window. Veeam Agent for Microsoft Windows will delete existing BitLocker encrypted partitions on the volume, format the disk and restore data from the backup as unencrypted.



IMPORTANT

Veeam Agent for Microsoft Windows cannot back up volumes formatted as FAT32 and encrypted with BitLocker. In general, FAT32 does not allow storing VSS snapshots on the same volume. When Veeam Agent for Microsoft Windows triggers a VSS snapshot of a FAT32 formatted volume, the VSS snapshot is stored on another, non-FAT32 volume on the computer.

If BitLocker is enabled, the VSS cannot save the snapshot on another volume due to Microsoft limitations, and the backup process fails.

Microsoft OneDrive Support

Veeam Agent for Microsoft Windows lets you back up your data to Microsoft OneDrive. This functionality is intended for users of home PCs, workstations and laptops who want to keep backups off-site without the need to set up the backup infrastructure.

Backup to Microsoft OneDrive is supported for Veeam Agent computers running the following operating systems:

- Microsoft Windows 7 SP1
- Microsoft Windows 8.1
- Microsoft Windows 10
- Microsoft Windows 11

To create backups in Microsoft OneDrive, you need to select Microsoft OneDrive as a target for backup files in the properties of the Veeam Agent backup job. When you configure the backup job, you must provide credentials of the Microsoft OneDrive account that has access to the Microsoft OneDrive storage. You can use the following types of Microsoft OneDrive accounts to back up your data:

- OneDrive (personal)
- OneDrive for Business

When the backup job session starts, Veeam Agent for Microsoft Windows copies the backed-up data directly to Microsoft OneDrive. In contrast to other types of backup location, during backup to Microsoft OneDrive, Veeam Agent for Microsoft Windows does not store copied data blocks to backups files of the VBK or VIB type. Instead, Veeam Agent for Microsoft Windows stores the entire data blocks 'as is'. This helps Veeam Agent for Microsoft Windows upload backed-up data, perform transformations with restore points in the backup chain and delete obsolete restore points by retention in a more efficient way.

For backups residing in Microsoft OneDrive, backup metadata is stored alongside the backed-up data itself. In addition, Veeam Agent for Microsoft Windows saves a copy of backup metadata to the *C:\ProgramData\Veeam\OneDriveBackup* folder on the Veeam Agent computer.

You can use backups that reside in Microsoft OneDrive to perform the same data recovery operations that are available for Veeam Agent backups in other types of target location. You can restore entire computer volumes or individual files and folders. In case your computer fails to start, you can also boot your computer from the Veeam Recovery Media and restore data from a volume-level backup residing in Microsoft OneDrive.

> **NOTE**
>
> You can use Microsoft OneDrive as a target for Veeam Agent backup jobs configured directly on a Veeam Agent computer only. Backup to Microsoft OneDrive for Veeam Agent for Microsoft Windows managed within the Veeam Agent management scenario is not supported.

Authorization in Microsoft OneDrive

Veeam Agent for Microsoft Windows communicates with Microsoft OneDrive over the OneDrive API. To let Veeam Agent for Microsoft Windows access the API, you need to sign in to Microsoft OneDrive in the **New Backup Job** wizard. When you sign in to Microsoft OneDrive, Veeam Agent for Microsoft Windows obtains from Microsoft OneDrive authorization tokens that provide access to the OneDrive API.

To work with the OneDrive API, Veeam Agent for Microsoft Windows uses the following types of tokens:

- **Access token.** Veeam Agent for Microsoft Windows uses this token to send requests to the OneDrive API. The access token is valid for 2 hours. To obtain a new access token after this period expires without the need to sign in to Microsoft OneDrive once again, Veeam Agent for Microsoft Windows uses the refresh token.
- **Refresh token.** Veeam Agent for Microsoft Windows uses this token to request a new access token from Microsoft OneDrive after the initial access token expires. In contrast to the access token, the refresh token is saved to Veeam Agent for Microsoft Windows database. The refresh token can be valid for 14 to 180 days depending on the type of account used to connect to Microsoft OneDrive:
  - 180 days — for a personal OneDrive account.
  - 14 to 90 days — for a OneDrive for Business account. Veeam Agent for Microsoft Windows can use the refresh token for more that 14 days if one of the following operations is performed within a 14-day period:
    - Sign-in to Microsoft OneDrive in the **New Backup Job** wizard
    - Backup to Microsoft OneDrive
    - Backup chain transformation in Microsoft OneDrive
    - Backup cache synchronization with Microsoft OneDrive
    - Data restore from Microsoft OneDrive

    After a 90-day period expires, you need to sign in to your OneDrive for Business account once again in the **New Backup Job** wizard. Otherwise, the backup job will fail.

Limitations for Backup to Microsoft OneDrive

Backup to Microsoft OneDrive has the following limitations:

- You cannot create a standalone full backup in Microsoft OneDrive.
- Veeam Agent does not support health check for backup files created in Microsoft OneDrive.
- Veeam Agent does not support creating Microsoft SQL Server transaction log backups and Oracle archived log backups in Microsoft OneDrive.
- Veeam Agent does not support long-term retention policy (GFS) for a backup job targeted at Microsoft OneDrive.
- In the Free edition of Veeam Agent, if you want to save backup files in Microsoft OneDrive, you can create backups that contain operating system data only.
- You cannot move backup files residing in Microsoft OneDrive to a backup storage of a different type and continue the Veeam Agent backup chain in another target location.
- If you use Veeam Agent with Veeam Backup & Replication, you cannot import a backup residing in Microsoft OneDrive in the Veeam Backup & Replication console.
- Due to limitations set in Microsoft OneDrive for third-party applications, backup to Microsoft OneDrive and data restore from Microsoft OneDrive requires longer time even in fast networks.
- Backup to Microsoft OneDrive is not supported for the following OneDrive storage plan regions:
    - US Government
    - China
    - Germany

Integration with Veeam Backup & Replication

If you plan to use Veeam Agent for Microsoft Windows 5.0 with Veeam Backup & Replication, you must install Veeam Backup & Replication 11a on the Veeam backup server. Veeam Agent for Microsoft Windows 4.0 supports integration with Veeam Backup & Replication 11 or later as well.

> **NOTE**
>
> You can also use Veeam Backup & Replication to manage Veeam Agent for Microsoft Windows on computers in your infrastructure. Within the Veeam Agent management scenario, you can remotely deploy Veeam Agent for Microsoft Windows to your computers, as well as configure and manage Veeam Agent backup jobs in Veeam Backup & Replication. To learn more, see the Veeam Agent Management Guide.
>
> The current guide covers subjects related to Veeam Agent for Microsoft Windows operating in the standalone mode.

You can store backup files created with Veeam Agent for Microsoft Windows on backup repositories managed by Veeam Backup & Replication. To do this, you must select a backup repository as a target location in the properties of the Veeam Agent backup job. To store Veeam Agent backups, you can use a simple backup repository or a scale-out backup repository.

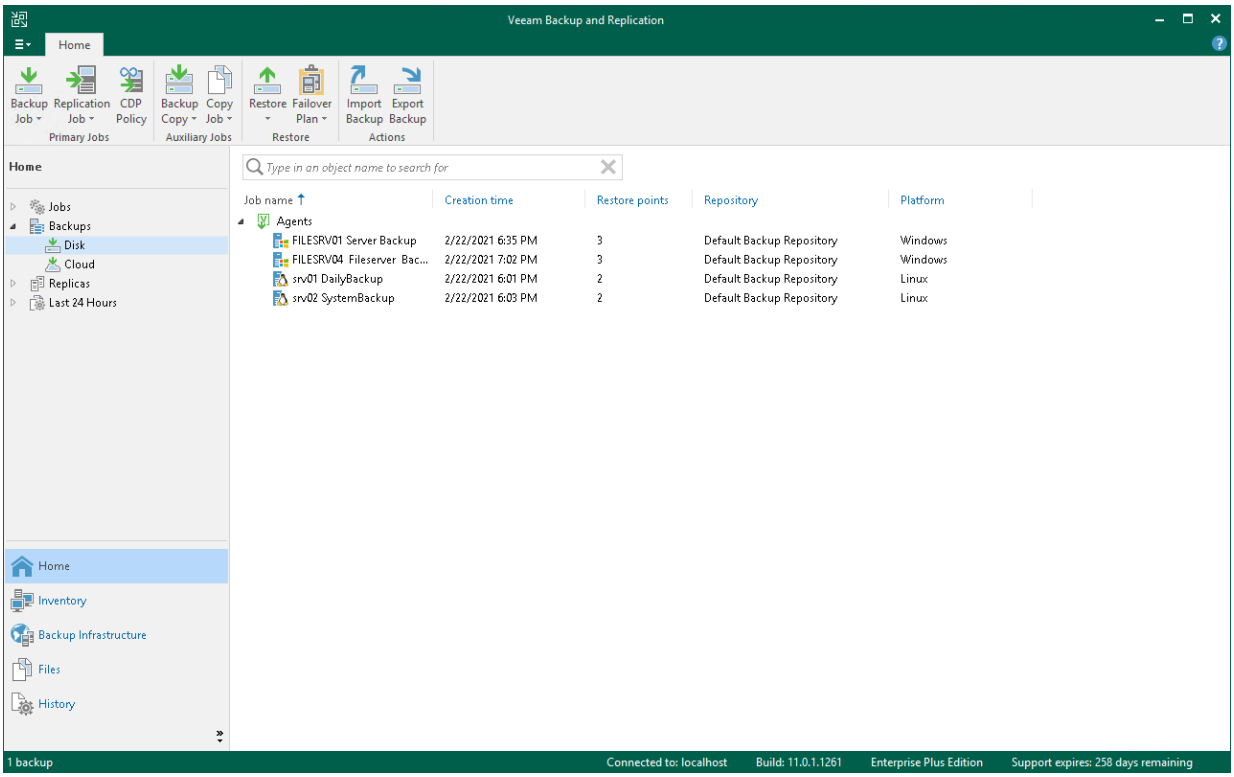
Veeam Agent for Microsoft Windows works with the backup repository as with any other target location. Backup files are stored to a separate folder; you can perform standard restore operations using these files.

Information about Veeam Agent backups stored on the backup repositories, backup jobs and sessions becomes available in the Veeam Backup & Replication console:

- The Veeam Agent for Microsoft Windows backup job is displayed in the list of jobs in Veeam Backup & Replication.
- Backup files created with Veeam Agent for Microsoft Windows are displayed in the list of backups, under the **Backups** > **Disk** node.
- Performed job sessions are available in the **History** view of Veeam Backup & Replication.

Backup administrators working with Veeam Backup & Replication can perform a set of operations with Veeam Agent backups:

- Perform data protection operations: copy Veeam Agent backups to secondary backup repositories and archive these backups to tape.
- Perform restore operations: restore individual files and folders, application items; restore computer disks and convert them to the VMDK, VHD or VHDX format; restore Veeam Agent backups to Microsoft Azure, Amazon Elastic Compute Cloud or to Hyper-V VMs.
- Perform administrative tasks: disable and delete Veeam Agent backup jobs, remove Veeam Agent backups and so on.



Backup to Veeam Cloud Connect Repository

If you want to store your data in the cloud, you can connect to a Veeam Cloud Connect service provider (SP) and create Veeam Agent backups in a cloud repository.

Limitations for Backup to Cloud Repository

Backup to a Veeam Cloud Connect repository has the following limitations:

- You can store backups in a Veeam Cloud Connect repository if you use Workstation or Server edition of Veeam Agent for Microsoft Windows only.
- You cannot use a Veeam Cloud Connect repository as a target for standalone full backup. For details, see Standalone Full Backup.
- Veeam Agent for Microsoft Windows does not support creating transaction log backups in the cloud repository. You cannot enable transaction log backup options in the properties of the backup job targeted at the cloud repository.

Managing Veeam Agent in Veeam Backup & Replication

Veeam Backup & Replication lets you automate management of Veeam Agent for Microsoft Windows on multiple computers in your infrastructure. You can deploy Veeam Agent for Microsoft Windows, configure Veeam Agent backup jobs and perform other data protection and administration tasks on remote computers. To use the Veeam Agent management functionality in Veeam Backup & Replication, you must install Veeam Backup & Replication on the Veeam backup server.

To learn more, see the Veeam Agent Management Guide.

Planning and Preparation

Before you install Veeam Agent for Microsoft Windows, make sure that the target computer meets the system requirements and all required ports are open.

System Requirements

The protected computer must meet requirements listed in the table below.

NOTE

The following system requirements apply to the following Veeam Agent for Microsoft Windows configuration:

- Veeam Agent version is 5.0.2.

  To learn system requirements for Veeam Agent previous versions, see Release Notes for Veeam Agent for Microsoft Windows 5.0 and Release Notes for Veeam Agent for Microsoft Windows 5.0.1.
- Veeam Agent is operating in the standalone mode.

  To learn about system requirements for Veeam Agent managed by Veeam Backup & Replication, see the System Requirements section in the Veeam Agent Management Guide.

Specification	Requirement
Hardware	CPU: x86-64 processor. Memory: 2 GB RAM or more. Memory consumption varies depending on number and size of processed disks. Disk Space: 200 MB for product installation. Network: 1 Mbps or faster. High latency and reasonably unstable WAN links are supported. System firmware: BIOS or UEFI. Drive encryption: Microsoft BitLocker (optional). BitLocker encrypted volumes must be unlocked at the moment when Veeam Agent starts the backup or restore operation. Only Microsoft BitLocker is supported for drive encryption. Other drive encryption products are not supported.

Specification	Requirement
OS	Both 64-bit and 32-bit (where applicable) versions of the following operating systems are supported[1]: • 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 SP1[2] • Microsoft Windows 11 (version 21H2) • Microsoft Windows 10 Semi-Annual Channel (from version 1803 up to version 21H2)[3] • Microsoft Windows 10 Long-Term Servicing Channel (versions 1507, 1607 and 1809) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1 [1] Consider the following: • Small Business Server, Server Essentials, and Server Storage editions of Microsoft Windows Server OSes are supported. • Server Core installations of Microsoft Windows Server OSes are supported for Veeam Agent for Microsoft Windows managed by Veeam Backup & Replication only. To learn more, see Veeam Agent Management Guide. Veeam Agent operating in the standalone mode does not support Server Core installations. • Windows Embedded / Windows IoT OSes are supported (except for custom builds by certain vendors that do not have components required for Veeam Agent operation). • Microsoft Failover Clusters are supported for Veeam Agent for Microsoft Windows managed by Veeam Backup & Replication only. To learn more, see Veeam Agent Management Guide. Veeam Agent operating in the standalone mode does not support Microsoft Failover Clusters. [2] Veeam CBT driver is supported only if Microsoft Windows update KB3033929 is installed on the Veeam Agent computer. [3] Microsoft Windows 10 Education is supported starting from build 10586 and later.

Specification	Requirement
File System	Microsoft Windows FAT32/exFAT, NTFS, ReFS file systems are supported. The supported file system must reside on a volume that is 64 TB or smaller, because Veeam Agent uses the Microsoft Software Shadow Copy Provider to create a volume shadow copy during the backup. To learn more about the limitation, see this Microsoft article.
Software	The following required 3rd party software is included in the setup program: • Microsoft .NET Framework 4.5.2 • Microsoft SQL Server 2012 Management Objects • Microsoft SQL Server System CLR Types • Windows Universal C Run-time Library When installing the product, the setup program checks whether all prerequisite software is available on the target computer. If some of the required software components are missing, the missing software is installed automatically. If you plan to use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you must install Veeam Backup & Replication 11 on the Veeam backup server.
Microsoft SQL Database	Microsoft SQL Server 2012 SP4 Express LocalDB Edition (installed with the product).

Veeam Agent for Microsoft Windows works with only those hard drive types that are supported by the Microsoft Windows OS. Thus, Veeam Agent supports the 512 bytes and 4 KB sector hard drives only. Other hard drive types are not supported. To learn more, see this Microsoft article.

Backup Source

Veeam Agent for Microsoft Windows supports backup of data that resides in the following types of storage:

- Local (internal) storage of the protected computer.
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives (USB sticks and SD cards are not supported).
- Storage Area Network (SAN), such as iSCSI connected volumes.

Keep in mind that Veeam Agent cannot work with third-party volume managers installed on the protected computer. Such managers may not allow Veeam Agent to interact with necessary interfaces and services properly.

Backup Target

Backup can be performed to the following types of storage:

Disk-based storage

- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives, and raw device mapping (RDM) volumes.

  > **IMPORTANT**
  >
  > We do not recommend targeting a backup job at the storage device with the exFAT file system. If the protected computer runs Microsoft Windows 10 or Microsoft Windows Server 2019 and later, this configuration may lead to the backup data corruption caused by the exFAT file system issue.

- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) share.
- Storage Area Network (SAN), such as iSCSI connected volumes.
- Veeam Backup & Replication 11 or later backup repository.

Cloud storage

- Veeam Cloud Connect 11 or later cloud repository.
- Microsoft OneDrive storage (for Microsoft Windows 7 SP1, 8.1, 10 and 11).

Network

Consider the following:

- Veeam Agent for Microsoft Windows should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Agent cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 addresses.

Used Ports

Make sure that you open ports listed below to enable proper work of Veeam Agent for Microsoft Windows.

From	To	Protocol	Port	Notes
Veeam Agent Computer	Veeam Backup Server	TCP	10001	Default port used by Veeam Agent for Microsoft Windows operating in the standalone mode for communication with the Veeam Backup server. Data between the Veeam Agent computer and backup repositories is transferred directly, bypassing Veeam backup servers.
	Veeam Agent Computer	TCP	9395+, 6183+	Ports used locally on the Veeam Agent computer for communication between Veeam Agent components and Veeam Agent for Microsoft Windows Service. If the default port number is already in use, Veeam Agent for Microsoft Windows Service will try to use the next port number.
	Veeam Update Notification Server (autolk.veeam.com)	TCP	443	Default port used to download information about available updates from the Veeam Update Notification Server over the Internet.
Communication with Veeam Backup & Replication Repositories				
Veeam Agent Computer	Linux server performing the role of a backup repository	TCP	22	Port used as a control channel from the Veeam Agent computer to the target Linux host.
		TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
	Microsoft Windows server performing the role of a backup repository	TCP	49152-65535	Dynamic RPC port range. For more information, see this Microsoft article.
		TCP	2500 to 3300	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
	Shared folder CIFS (SMB) share	TCP UDP	137 to 139, 443, 445	Ports used as a transmission channel from the Veeam Agent computer to the target CIFS (SMB) share. Ports 137 to 139 are used by backup infrastructure components to communicate using NetBIOS. Port 443 is used to connect the target CIFS (SMB) share over QUIC. For more information, see this Microsoft article.
	Gateway Microsoft Windows server	TCP UDP	137 to 139, 445	If a CIFS (SMB) share is used as a backup repository and a Microsoft Windows server is selected as a gateway server for this CIFS share, these ports must be opened on the gateway Microsoft Windows server. Ports 137 to 139 are used by backup infrastructure components to communicate using NetBIOS.
		TCP	49152-65535	Dynamic RPC port range. For more information, see this Microsoft article.
		TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
Communication with Veeam Cloud Connect Repositories				
Veeam Agent Computer	Cloud gateway	TCP	6180	Port on the cloud gateway used to transport Veeam Agent data to the Veeam Cloud Connect repository.
	Certificate Revocation Lists	TCP	80 or 443 (most popular)	Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider. Generally, information about CRL locations can be found on the CA website.

IMPORTANT

The list of ports required for computers booted from the Veeam Recovery Media is the same as the list of ports required for Veeam Agent computers.

Licensing

You can use Veeam Agent for Microsoft Windows as a free product. In this case, you do not need to obtain and install any license.

To work with a commercial version of Veeam Agent for Microsoft Windows, you must obtain a license and install it on the protected computer. If you do not install a license, you will be able to use the Free edition of the product only.

If you plan to use a commercial version of the product with Veeam Backup & Replication, you must install and manage the license in the Veeam Backup & Replication console or in Veeam Backup Enterprise Manager. To learn more, see Managing License.

Product Editions

Veeam Agent for Microsoft Windows offers three product editions that define product functionality and operation modes:

- *Server* — a commercial edition that provides access to all product functions. The Server edition is intended for performing data protection tasks on servers that run the Microsoft Windows OS. To use the Server edition of Veeam Agent for Microsoft Windows, you must obtain and install a license on the protected computer. The license must have a number of instances that is enough to protect a machine with the Server product edition.
- *Workstation* — a commercial edition that offers capabilities for performing data protection tasks on desktop computers and laptops that run the Microsoft Windows OS. To use the Workstation edition of Veeam Agent for Microsoft Windows, you must obtain and install a license on the protected computer. The license must have a number of instances that is enough to protect a machine with the Workstation product edition.
- *Free* — a free edition that offers limited capabilities sufficient for personal use of the product. In contrast to Workstation and Server editions, the Free edition does not require a license.

For more information about product editions, pricing and features available for them, see this Veeam webpage.

To learn more about instance licensing, see this Veeam webpage.

When you install a license on the protected computer, you can select the product edition of Veeam Agent for Microsoft Windows: Workstation or Server (if both editions are supported by the license). To learn more, see Selecting Product Edition.

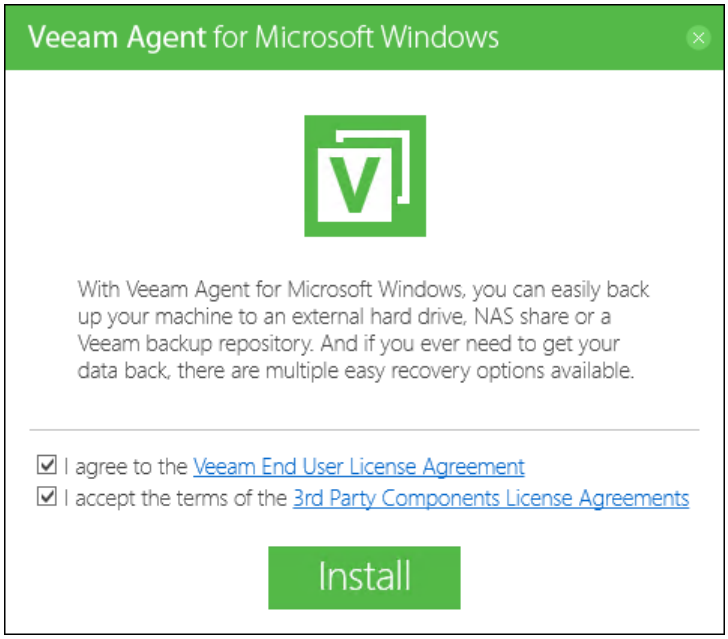
If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you should manage product licenses and editions from the Veeam Backup & Replication console. To learn more, see Managing License.

If the license has expired and you have used backup job options available for Workstation and Server editions, you must disable these options in the properties of the backup job. Otherwise, the backup job will fail.

License Agreements

When you install Veeam Agent for Microsoft Windows, you must accept the terms of the product license agreement and terms of the 3rd party components license agreements.

- To view the product license agreement, click the **Veeam End User License Agreement** link in the installation window or visit this Veeam webpage.
- To view 3rd party components license agreements, click the **3rd Party Components License Agreements** link in the installation window.



Installing License

When you launch the Veeam Agent control panel for the first time, Veeam Agent displays a notification window offering to install a license. You can choose to install the license immediately or postpone this operation.

- If you choose to install the license, you can immediately browse for the license key on your computer and complete the license installation process.
- If you choose to postpone the license installation process, you will be able to install a license later at any time you need.

Until you install a license, you can use the Free edition of the product. To switch to a commercial version of Veeam Agent for Microsoft Windows, you need to obtain and install a license.

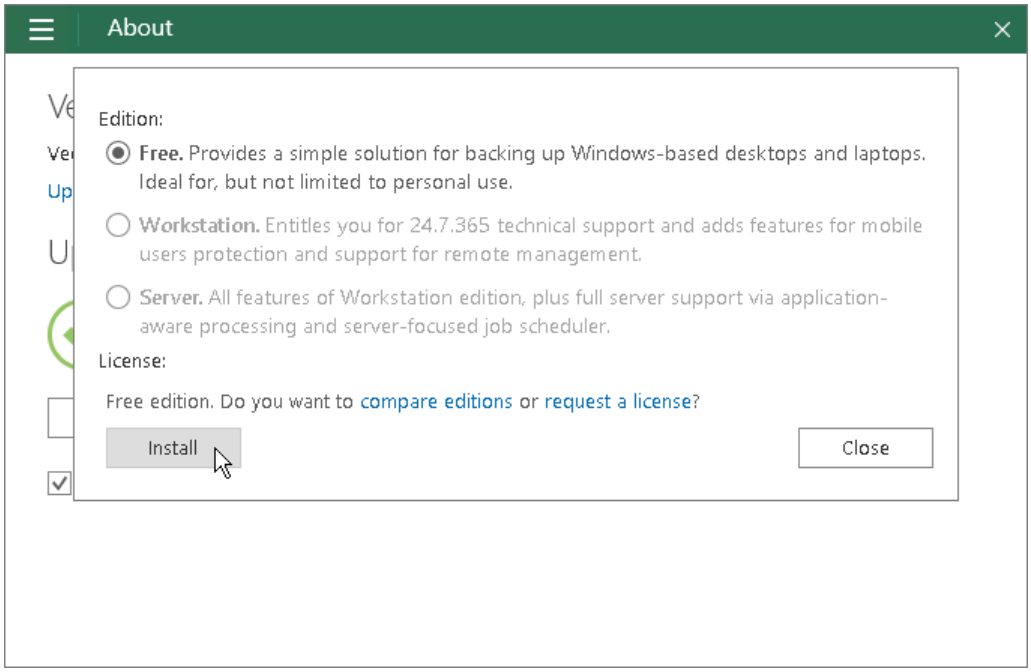
> **NOTE**
>
> If you plan to use a Veeam Backup & Replication repository as a target location for Veeam Agent backups, you must install a license in Veeam Backup & Replication. The license must have enough instances to protect machines with Veeam Agents that back up data to the Veeam Backup & Replication repository. To learn more, see Managing License.

To install a license:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **About**.
3. In the **Version** section, click the following link:
   a. **Update license to get additional features** — if the license is not installed yet, and you run the Free edition of Veeam Agent for Microsoft Windows.
   b. **Manage license and edition** — if the license is already installed on the Veeam Agent computer, and you want to change the license or select the product edition.
4. In the dialog window, click **Install** and browse for the LIC file.

   Veeam Agent for Microsoft Windows will install the license and select the product edition that is allowed by the license. If a license supports both the Workstation and Server editions, Veeam Agent for Microsoft Windows will select the product edition based on the type of the Microsoft Windows OS installed on the Veeam Agent computer.

   You can change the product edition manually if needed. To learn more, see Selecting Product Edition.



Selecting Product Edition

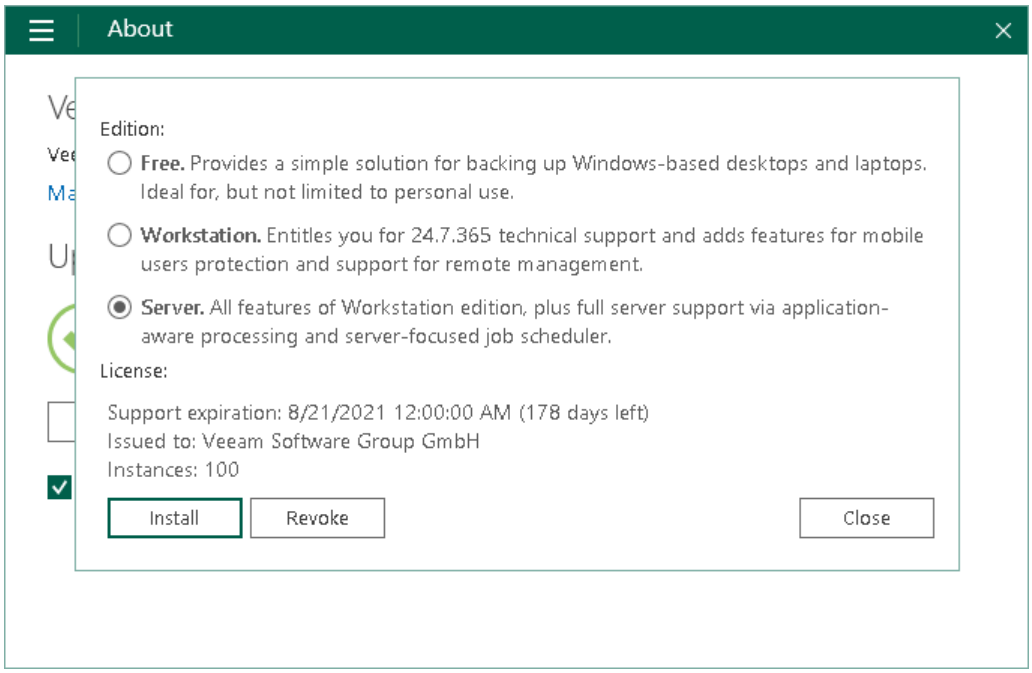
When you install a license, Veeam Agent for Microsoft Windows automatically selects the product edition that is allowed by the license. If a license supports both the Workstation and Server editions, Veeam Agent for Microsoft Windows will select the product edition based on the type of the Microsoft Windows OS installed on the Veeam Agent computer.

You can change the product edition manually if needed. To select the product edition:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **About**.
3. In the **Version** section, click **Manage license and edition**.
4. In the dialog window, in the **Edition** section, select the desired product edition. To learn more about editions of Veeam Agent for Microsoft Windows, see Product Editions.

NOTE

After you switch from the Server edition to the Workstation edition, or vice versa, Veeam Agent for Microsoft Windows will disable the backup job. This operation is required, because backup retention policies and available backup job options differ in Workstation and Server editions. To enable the job, you must edit the backup job settings in accordance with the selected edition.

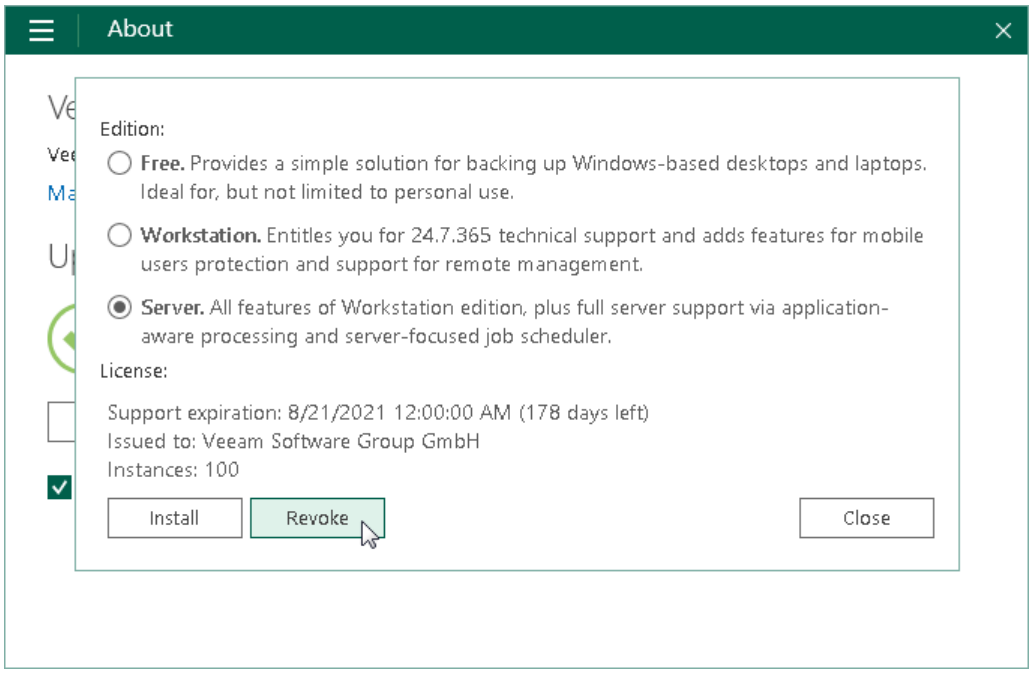


Revoking License

You can revoke a license at any time if needed, for example, after the license is expired, and you want to continue using the Free edition of Veeam Agent for Microsoft Windows.

To revoke a license:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **About**.
3. In the **Version** section, click **Manage license and edition**.
4. In the dialog window, click **Revoke**.



Installation and Configuration

You can install Veeam Agent for Microsoft Windows on any computer whose data you plan to protect — server, desktop, laptop or tablet.

Before You Begin

Before you start the installation process, check the following prerequisites:

- The computer on which you plan to install Veeam Agent must satisfy system requirements specified in this document. To learn more, see System Requirements.
- Make sure the operating system on the computer you want to protect is up-to-date. Otherwise, the correct functioning of Veeam Agent is not guaranteed.
- You must run the Veeam Agent setup file under the Administrator account or any user account that has Administrator privileges on the computer where you plan to install the product.
- Veeam Agent for Microsoft Windows requires the following components:
  - Microsoft SQL Server System CLR Types
  - Microsoft SQL Server 2012 Management Objects
  - Microsoft SQL Server 2012 (LocalDB)
  - Microsoft .NET Framework 4.5.2 or later
  - Windows Universal C Run-time Library

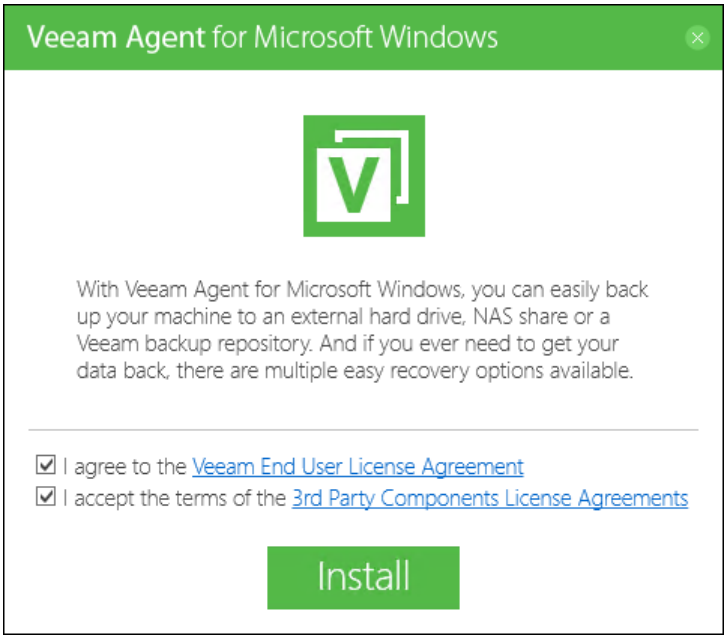
  If these components are not pre-installed on the computer, the setup will install them during the product installation process.
- The product program files are placed to the `%ProgramFiles%\Veeam\Endpoint Backup` folder on the system volume. Make sure that you have enough free space on the system volume to install the product. Veeam Agent requires at least 200 MB.
- [Recommended] If you want to configure a scheduled backup job with default settings after the installation, you must prepare a USB storage device.
- [Recommended] If you want to create a recovery image of your computer on a USB storage device, CD/DVD/BD or make an ISO image, prepare the necessary device/media or make sure that you have enough free disk space in the target location. On average, the size of the created recovery image is 500 MB.

  During the recovery image creation, Veeam Agent formats the removable storage device. If you have important information on the device, create a copy of this data in some other location.

Installing Veeam Agent for Microsoft Windows

To install Veeam Agent for Microsoft Windows:

1. Download the Veeam Agent for Microsoft Windows setup archive from this Veeam webpage, and save the downloaded archive on the computer where you plan to install the product.
2. Double-click the downloaded setup archive.
3. To install Veeam Agent for Microsoft Windows, you must accept the license agreements:
   a. Select the **I agree to the Veeam End User License Agreement** check box.
   b. Select the **I accept the terms of the 3rd Party Components License Agreements** check box.
4. Click **Install**.



5. After the installation process is complete, you can instruct Veeam Agent for Microsoft Windows to perform the following advanced actions:
   - Auto-configure settings for the backup job. To learn more, see Auto-Configuring Scheduled Backup Job.
   - Create a recovery image for your computer. To learn more, see Creating Veeam Recovery Media.

Installing Veeam Agent for Microsoft Windows in Unattended Mode

You can install Veeam Agent for Microsoft Windows in the unattended mode using the command line interface. The unattended installation mode does not require user interaction — the installation runs automatically in the background, and you do not have to respond to the installation wizard prompts. You can use the unattended installation mode to automate the Veeam Agent installation process in large-scale environments.

Prerequisite Software

During the product installation, Veeam Agent for Microsoft Windows automatically sets up the following required prerequisite components:

- Microsoft SQL Server System CLR Types for SQL Server 2012
- Microsoft SQL Server 2012 Management Objects
- Microsoft SQL Server 2012 Express LocalDB Edition
- Windows Universal C Run-time Library

Veeam Agent for Microsoft Windows will also set up Microsoft .NET Framework 4.5.2 if it does not detect this component on the computer during the product installation.

In some cases, installation of prerequisite software requires computer reboot. This can happen, for example, if you have an earlier version of a prerequisite component installed on the computer and during the installation process this component is used by third-party software.

In this situation, unattended setup will install Veeam Agent for Microsoft Windows but will not start the Veeam Agent for Microsoft Windows service. After you reboot the computer, the Veeam Agent for Microsoft Windows service will be started and Veeam Agent for Microsoft Windows will be fully functioning.

Installation Syntax

To install Veeam Agent for Microsoft Windows in the unattended mode, use a command with the following syntax:

```
<path_to_exe> /silent /accepteula /acceptthirdpartylicenses
```

where `<path_to_exe>` — path to the Veeam Agent for Microsoft Windows installation file.

Veeam Agent for Microsoft Windows uses the following codes to report about the installation results:

- 1000 — Veeam Agent for Microsoft Windows has been successfully installed.
- 1001 — prerequisite components required for Veeam Agent for Microsoft Windows have been installed on the machine. Veeam Agent for Microsoft Windows has not been installed. The machine needs to be rebooted.
- 1002 — Veeam Agent for Microsoft Windows installation has failed.
- 1101 — Veeam Agent for Microsoft Windows has been installed. The machine needs to be rebooted.

Using Sysprep and Veeam Agent for Microsoft Windows

You can pre-install Veeam Agent in a custom Microsoft Windows system image that will be used for deployment on different computers. To do this, you should perform a set of configuration steps in the reference Microsoft Windows system installation that will be included in a deployment image.

To configure a custom Microsoft Windows system image with Veeam Agent:

1. Install Veeam Agent in a Microsoft Windows system image. To learn more, see Installing Veeam Agent for Microsoft Windows.
2. Configure the backup job in the way you want it to work on computers with pre-installed Veeam Agent. To learn more, see Configuring Backup Job.

> **NOTE**
>
> It is advised to configure the backup job for the entire computer backup. In case of volume-level backup, it may be necessary to reconfigure the backup job after Microsoft Windows is deployed to the target computer and include the necessary volumes in the backup once again. This may happen if volumes' GUIDs were changed at the stage of Microsoft Windows generalization with Sysprep.

3. Create a registry value: `HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Endpoint Backup\SysprepMode (DWORD)=1`.

   This registry value is used to regenerate the job ID when Veeam Agent starts for the first time on the new computer. If you do not create the registry value, the backup job may fail as soon as it is started on the new computer.
4. Run the Sysprep tool in the *Generalize* mode to remove any system-specific data. If you need to run the Sysprep tool in the *Audit* mode, do not forget to re-create the registry value afterwards.
5. Deploy the image on the necessary computers in any convenient way. To learn more about deployment of Microsoft Windows system to new computers, see this Microsoft article.

When you deploy the created image on the computer, Veeam Agent will re-generate internal IDs of the backup jobs. As a result, the backup job will be fully functional.

Upgrading Veeam Agent for Microsoft Windows

For Veeam Agent for Microsoft Windows, upgrade to newer versions is supported. You can start the upgrade process from the Veeam Agent control panel when the new version becomes available. To learn how to check for product updates, see Checking for New Product Versions and Updates.

During the upgrade process, configuration and backup files that were created with the previous version of Veeam Agent are not impacted in any way.

To upgrade Veeam Agent for Microsoft Windows:

1. Double-click the Veeam Agent icon in the system tray, or right-click the Veeam Agent icon in the system tray and select **Control Panel**.
2. Open the **About** tab.
3. If the new version of Veeam Agent is available, click **Download**.
4. When the download is complete, click **Install** to run the setup archive.
5. To upgrade Veeam Agent, you must accept the license agreement. Read the license agreement, select the **I agree to the Veeam End User License Agreement** check box and click **Update**.

> **NOTE**
>
> Consider the following:
>
> - In some cases, upgrade to the new version of Veeam Agent may require computer reboot.
> - You can also download the Veeam Agent setup archive from this Veeam webpage. Save the downloaded archive on the computer where you plan to install the new version of the product and double-click the setup archive to start the upgrade.

Unattended Upgrade

You can upgrade Veeam Agent to a newer version in the unattended mode using the same command that is used for unattended installation. To learn more, see Installing Veeam Agent for Microsoft Windows in Unattended Mode.

Configuring Advanced Settings

Connection Settings for Veeam Backup Server

If you want to connect Veeam Agent computer to Veeam backup server as a member of the protection group for pre-installed Veeam Agents, you must apply connection settings from the configuration file. The configuration file is one of the Veeam Agent for Microsoft Windows setup files that you must obtain from your System Administrator. To learn more about protection group for pre-installed Veeam Agents, see the Protection Group Types section in the Veeam Agent Management Guide.

To connect Veeam Agent for Microsoft Windows to Veeam backup server, use Veeam Agent Configurator. To learn more, see the SetVBRSettings section in the Veeam Agent Configurator Reference.

Uninstalling Veeam Agent for Microsoft Windows

To uninstall Veeam Agent for Microsoft Windows:

1. From the **Start** menu, select **Control Panel** > **Programs and Features**.
2. In the programs list, right-click Veeam Agent and select **Uninstall**. Wait for the process to complete.

The LocalDB and other prerequisite components installed and used by Veeam Agent are not removed during the uninstall process. To remove each of the remaining components, right-click it in the programs list and select **Uninstall**.

What You Do Next

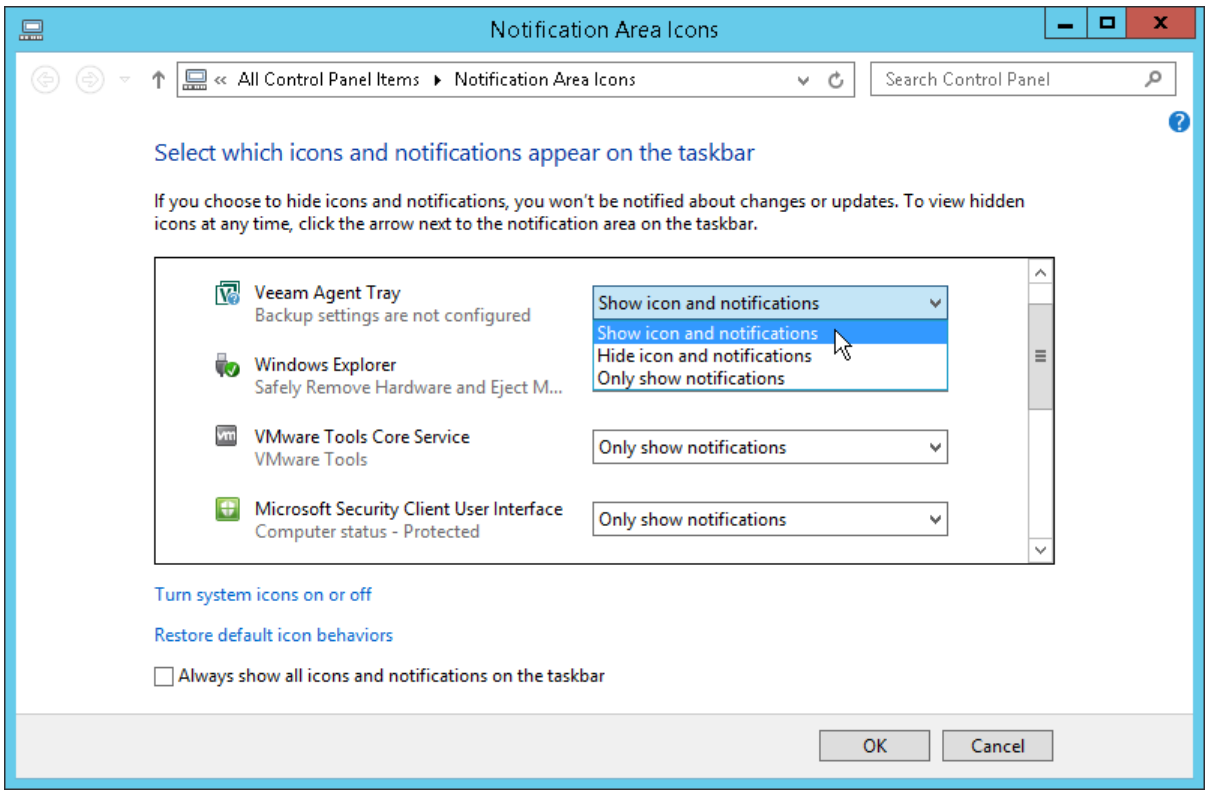
After the product installation, Veeam Agent for Microsoft Windows displays its icon in the system tray. You can use the system tray icon to perform main operations in Veeam Agent for Microsoft Windows:

- Configure the backup job and start ad-hoc backup operations
- Launch restore wizards
- Open the Veeam Agent control panel
- Monitor the state of backup tasks and so on

Depending on the current settings of your Microsoft Windows OS, the Veeam Agent for Microsoft Windows icon may not be displayed in the system tray.

To bring the icon to the system tray:

1. In Microsoft Windows, open the **Notification Area Icons** view. To do this, do either of the following:
   - Click the arrow in the system tray and click the **Customize** link.
   - From the Microsoft Windows main menu, select **Control Panel** and navigate to **Appearance and Personalization**. In the **Taskbar** section, select **Customize icons on the taskbar**.
2. In the **Notification Area Icons** window, find Veeam Agent Tray.
3. In the **Behaviors** column, set the **Show icon and notification** setting for it.
4. Click **OK**.



Getting Started

To protect your computer from a disaster of any kind, you must perform the following operations in Veeam Agent for Microsoft Windows:

1. **Create a Veeam Recovery Media**.

   The Veeam Recovery Media provides an alternate way to boot the Microsoft Windows RE. If your computer fails to start or the hard disk gets corrupted, you can boot the Windows RE from the Veeam Recovery Media and restore your data.

   To learn more, see Creating Veeam Recovery Media.

2. **Define what data you want to back up and configure the backup job**.

   Before you configure a backup job, you should decide on the following backup details:

   - Backup scope: entire computer image, individual computer volumes or specific computer folders.
   - Backup destination: where you want to store created backups.
   - Backup schedule: how often you want to back up your data.

   After that, you can configure the backup job. The backup job runs automatically by the defined schedule, captures the data that you have added to the backup scope and creates a chain of restore points in the target location. If your data gets lost or corrupted, you can restore it from the required restore point.

   To learn more, see Configuring Backup Job.

3. **Specify Veeam Agent settings**.

   You can define resource usage settings during backup, instruct Veeam Agent to automatically check for new product versions and so on. To learn more, see Specifying Settings.

4. **Monitor backup task performance**.

   You can use the Veeam Agent control panel to check how backup tasks are being performed, what errors have occurred during backup job sessions and so on. To learn more, see Reporting.

5. In case of a disaster, you can **restore the entire computer image or specific data** on the computer. To learn more, see Performing Restore.

Performing Backup

To protect your computer and data, you can perform the following operations:

- Create a Veeam Recovery Media. You can use the Veeam Recovery Media to boot the Microsoft Windows RE from the recovery media in case the OS on your computer fails to start. To learn more, see Creating Veeam Recovery Media and Creating Veeam Recovery Media with Command Line Interface.
- Back up your data. You can back up your data to protect the entire computer image, individual volumes or folders on your computer. Veeam Agent lets you configure a scheduled backup job with the default settings right after the product installation, configure a backup job with custom settings or create ad-hoc backups at any time you need. You can use data backup to restore necessary information if data on your computer gets corrupted or you delete some files and folders by mistake. To learn more, see Configuring Backup Job.

Creating Veeam Recovery Media

You can create a Veeam Recovery Media — a recovery media for your computer. The Veeam Recovery Media contains all data that is required to run the Microsoft Windows RE. If your computer stops working or the hard disk fails, you can boot from the Veeam Recovery Media, instead of booting from the hard drive. After booting, you can use Veeam and Microsoft tools to fix errors, recover the system image of your computer and your data.

NOTE

In some cases, Windows Recovery Environment components may be missing on the system, and Veeam Agent for Microsoft Windows will not find them during the Veeam Recovery Media creation. In such case you will be prompted to do one of the following:

- Insert the Windows Installation Media so that Veeam Agent for Microsoft Windows can load the necessary components from it.
- Download and install Windows Assessment and Deployment Kit (MS ADK).

Keep in mind that the Veeam Recovery Media created with MS ADK components will not contain the following tools:

- Windows Recovery Environment
- Memory Diagnostic
- Startup Repair

Before You Begin

Before you create a Veeam Recovery Media, check the following prerequisites:

Removable Storage Device Scenario (USB, SD Card and Other)

- The removable storage device must be inserted into a corresponding slot on the computer or connected to the computer.
- The removable storage device must have enough capacity to store the created recovery image. On average, the size of the created recovery image without manually loaded drivers is 500 MB.
- During the recovery image creation, Veeam Agent for Microsoft Windows formats the removable storage device. If you have important information on the device, create a copy of this data in some other location.
- If you want to include specific storage and network drivers into the recovery image, place them to a local folder on your computer or in a network shared folder to which you have access and read permissions. During the recovery image creation process, you will be able to define a path to this folder, and Veeam Agent for Microsoft Windows will include the drivers into the recovery image.
- [For Microsoft Windows 2008 R2 and later] If you want your computer to detect a Wi-Fi network and connect to it after you boot from the recovery image, enable the Wireless LAN Service feature on your computer. In this case, Veeam Agent for Microsoft Windows will add wireless networking support files to the Veeam Recovery Media. To learn more about the Wireless LAN Service, see this Microsoft article.

CD/DVD/BD Scenario

- An empty or re-writable CD/DVD/BD must be inserted into a CD/DVD/BD drive on the computer.
- The CD/DVD/BD must have enough capacity to store the created recovery image. On average, the size of the created recovery image without manually loaded drivers is 500 MB.
- If you want to include specific storage and network drivers into the recovery image, place them to a local folder on your computer or in a network shared folder on which you have read permissions. During the recovery image creation, you will be able to define a path to this folder, and Veeam Agent for Microsoft Windows will include the drivers into the recovery image.
- [For RW CD/DVD/BD] During the recovery image creation, Veeam Agent for Microsoft Windows erases information on the CD/DVD/BD. If you have important information on the CD/DVD/BD, create a copy of this data in some other location.
- [For Microsoft Windows 2008 R2 and later] If you want your computer to detect a Wi-Fi network and connect to it after you boot from the recovery image, enable the Wireless LAN Service feature on your computer. In this case, Veeam Agent for Microsoft Windows will add wireless networking support files to the Veeam Recovery Media. To learn more about the Wireless LAN Service, see this Microsoft article.

Local Target and Shared Folder Scenario (ISO)

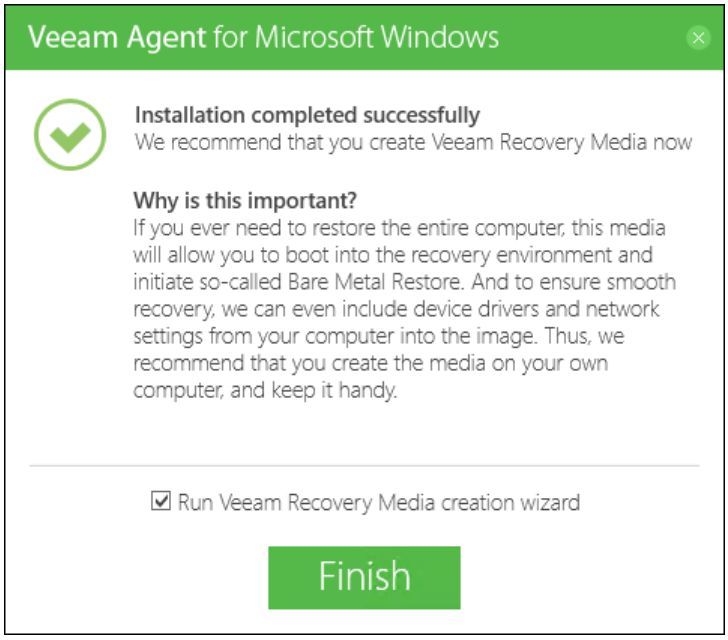
- If you want to include specific storage and network drivers into the recovery image, place them to a local folder on your computer or in a network shared folder on which you have read permissions. During the recovery image creation, you will be able to define a path to this folder, and Veeam Agent for Microsoft Windows will include the drivers into the recovery image.
- [For shared folders] If you plan to save the created ISO file in a network shared folder, make sure that you have access to this folder and write permissions on it.
- [For Microsoft Windows 2008 R2 and later] If you want your computer to detect a Wi-Fi network and connect to it after you boot from the recovery image, enable the Wireless LAN Service feature on your computer. In this case, Veeam Agent for Microsoft Windows will add wireless networking support files to the Veeam Recovery Media. To learn more about the Wireless LAN Service, see this Microsoft article.

Step 1. Launch Create Recovery Media Wizard

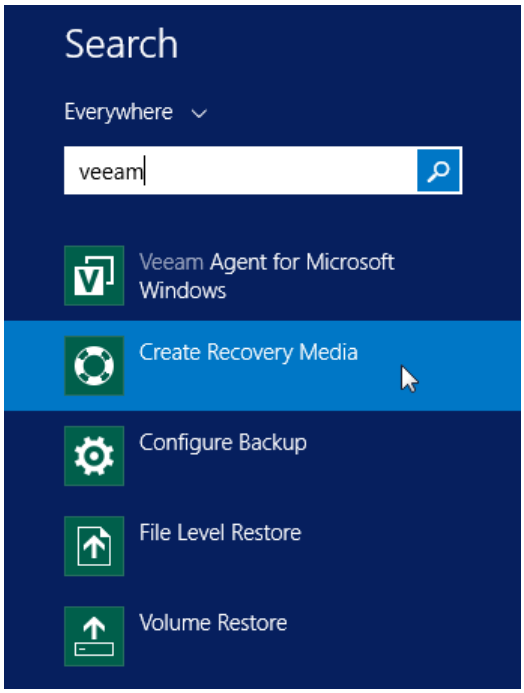
You can launch the **Create Recovery Media** wizard right after the product installation process or at any time later.

To launch the **Create Recovery Media** wizard after installation:

1. At the last step of the installation wizard, select the **Run Veeam Recovery Media creation wizard** check box.
2. Click **Finish**. Veeam Agent for Microsoft Windows will automatically launch the **Create Recovery Media** wizard.



To launch the **Create Recovery Media** wizard at any time, from the Microsoft Windows Start menu, select **All Programs** > **Veeam** > **Tools** > **Create Recovery Media** or use the Microsoft Windows search to find the **Create Recovery Media** option on your computer.



Step 2. Specify Recovery Media Options

At the **Recovery Media** step of the wizard, specify on which type of media you want to create a recovery image and what drivers you want to include in the recovery image.

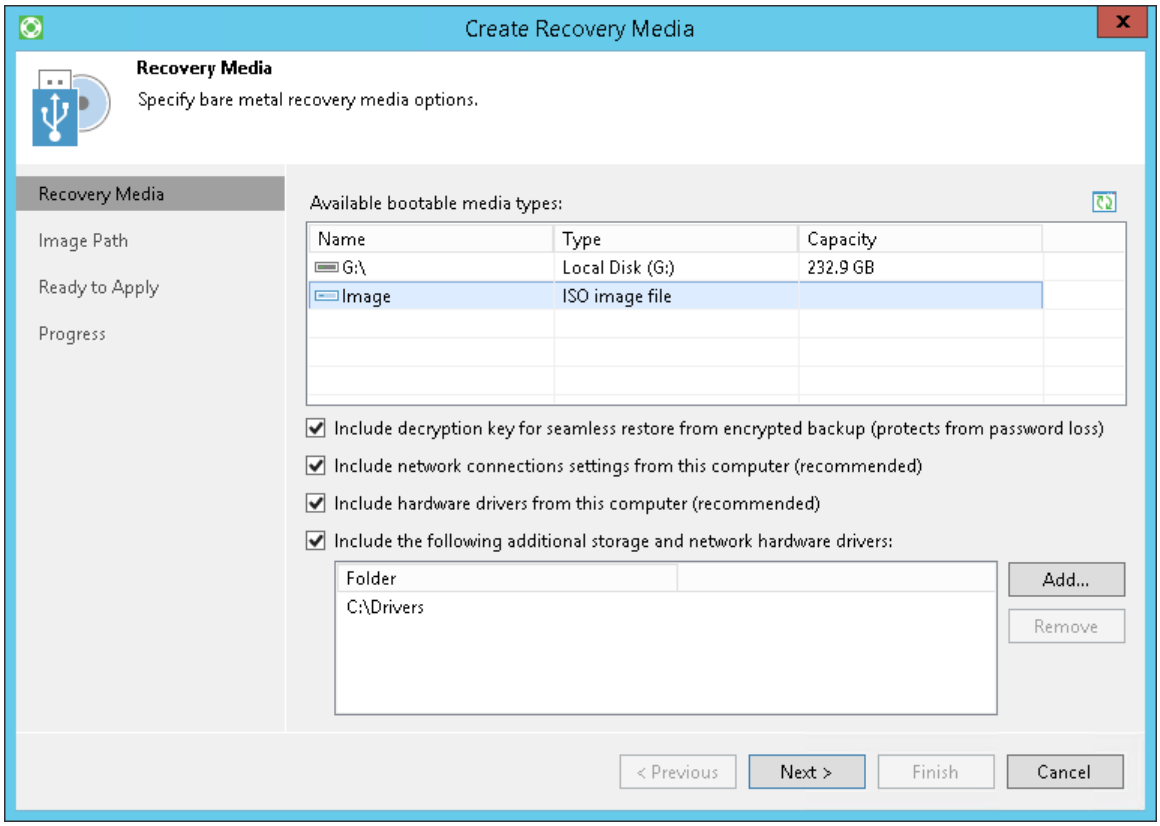
1. In the **Available bootable media types** list, select a media for the recovery image. You can create the following types of recovery images:
   - Recovery image on a removable storage device. You can create a recovery image on a USB drive, SD card and so on. Veeam Agent for Microsoft Windows displays all removable storage devices currently attached to your computer. Select the necessary one in the list.
   - Recovery image on an optical disk. You can create a recovery image on a CD, DVD or BD. Veeam Agent for Microsoft Windows displays all CD, DVD and BD drives available on your computer. Select the necessary one in the list.
   - ISO file with the recovery image. You can create a recovery image in the ISO file format and save the resulting file locally on your computer or in a network shared folder.
2. If you have enabled data encryption options for the backup job and want to include the decryption key in the recovery image, select the **Include decryption key for seamless restore from encrypted backup** check box. In this case, when you use the created recovery image to perform bare metal recovery, you will not have to enter the password used for encryption.
3. If you want to include in the recovery image current network settings, make sure that the **Include network connections settings from this computer** check box is selected. When you use the created Veeam Recovery Media to boot your computer, these settings will be automatically applied and will be used to connect to the remote backup storage.
4. If you want to include in the recovery image storage and network drivers that are currently installed on your computer, make sure that the **Include hardware drivers from this computer** check box is selected. Veeam Agent for Microsoft Windows will detect hard disk controller drivers, network adapter drivers and USB controller drivers and include them into the recovery image. When you use the created Veeam Recovery Media to boot your computer, these drivers will be automatically injected into Windows RE.
5. If you want to include in the recovery image additional storage and network drivers that you may need when booting from the recovery image, select the **Include the following additional storage and network hardware drivers** check box, click **Add** and select a folder containing necessary drivers. The folder that you select must contain all files of the driver package (files in CAT, INF and SYS formats).

   We strongly recommend that you enable this option if you use drivers that are not included into the Microsoft Windows installation DVD. For example, you can include drivers for a discrete network card, third party USB 3.0 controllers and non-standard hard disk controllers.

IMPORTANT

Mind the following:

- When you boot your computer from the Veeam Recovery Media, Veeam Agent for Microsoft Windows does not automatically install additional drivers included in the recovery image. You need to install such drivers manually using the **Load Driver** tool. To learn more, see Using Veeam Agent and Microsoft Windows Tools.
- We do not recommend to include large amount of additional drivers (1 GB and more) in the Veeam Recovery Media. When you boot your computer from the Veeam Recovery Media, Veeam Agent for Microsoft Windows loads all additional drivers stored in the Veeam Recovery Media into your computer RAM. If the total size of the recovery environment is approximately equal to or greater than the amount of RAM, Windows RE will fail to load.



Step 3. Specify Path to ISO

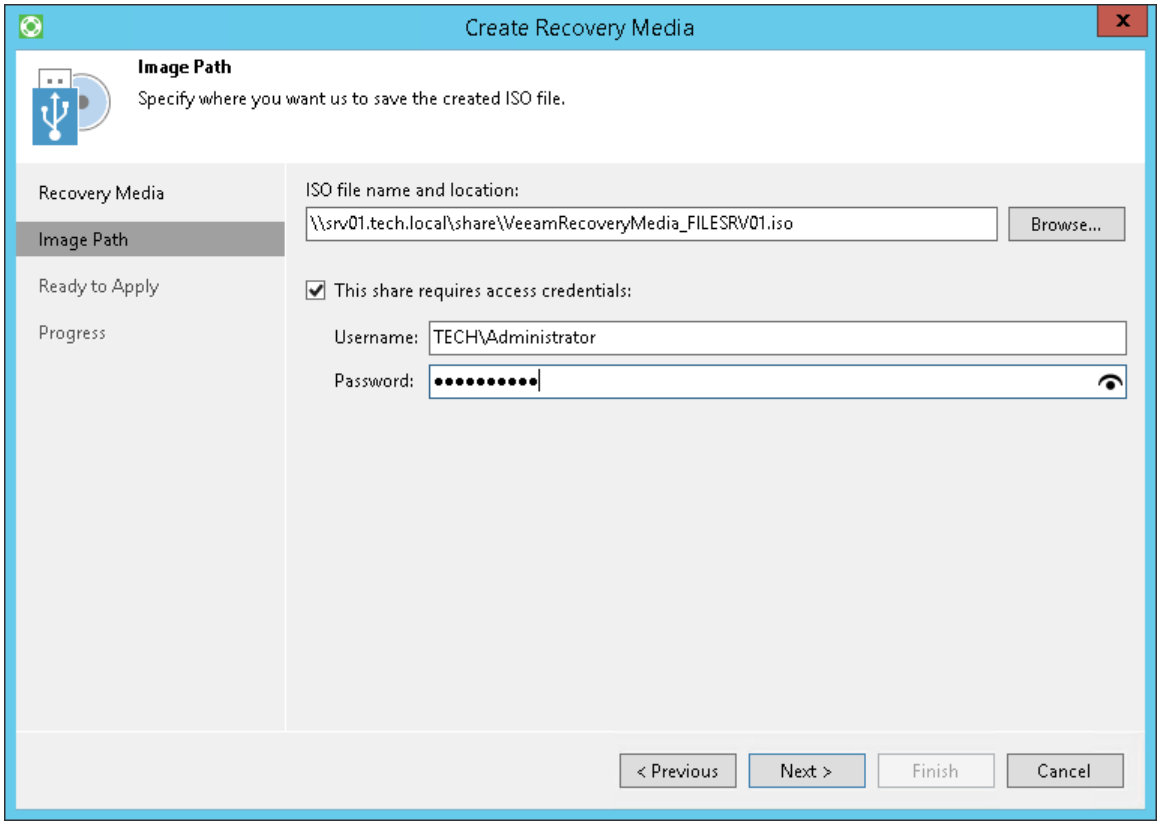
The **Image Path** step of the wizard is available if you have selected to create an ISO file with the recovery image.

Select a location where you want to save the ISO file.

1. In the **ISO file name and location** field, specify a real path to the folder where you want to save the created recovery image and the ISO file name. You can save the ISO file in the following locations:
   - Local folder: select the necessary folder on your computer.
   - Network shared folder: specify a UNC path to a network shared folder. Keep in mind that a UNC path always starts with two back slashes (\\).

   We strongly recommend that you store the recovery image in a location other than a local computer drive. If you choose to save the recovery image in a local folder on your computer, you can copy it to an external location afterwards. In this case, the recovery image will always be available should computer volumes get corrupted or the computer fail to start.
2. If you chose to save the ISO file in a network shared folder and this folder requires authentication, select the **This share requires access credentials** check box and enter the user name and password in the **Username** and **Password** fields. The user name must be specified in the *DOMAIN\UserName* format.

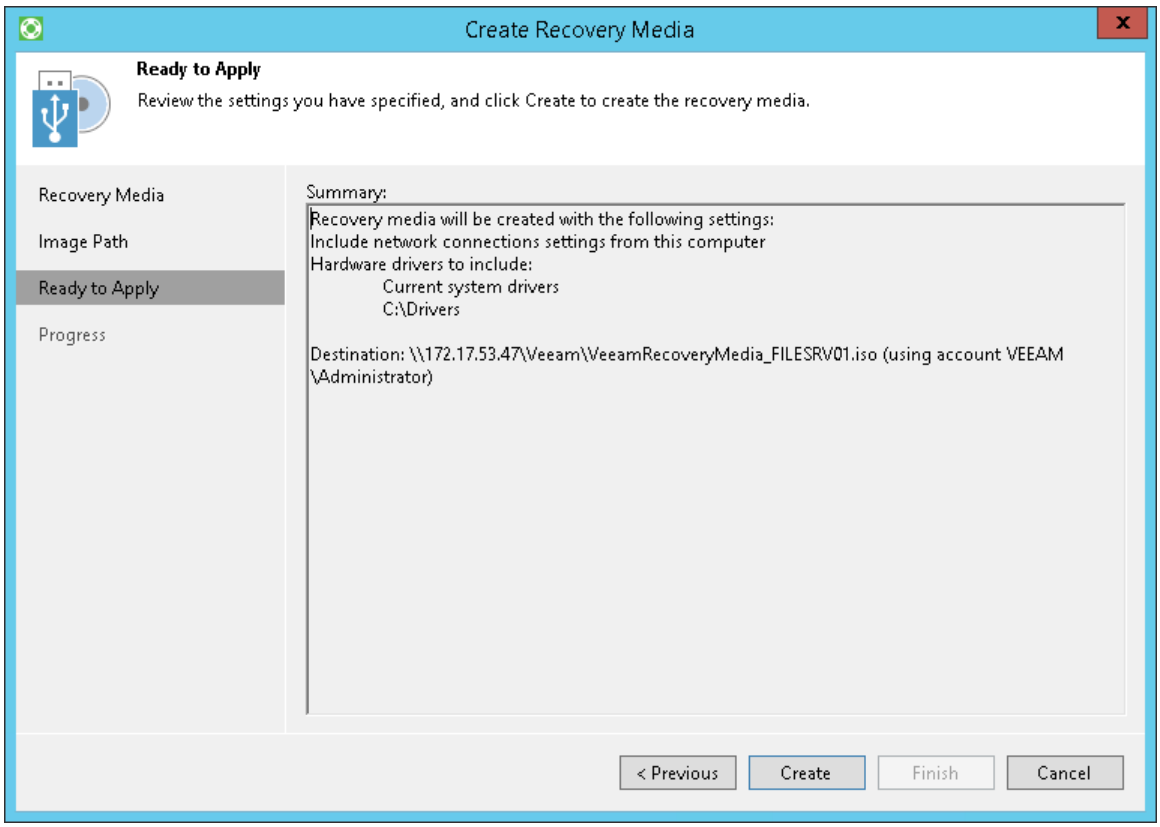
   To view the entered password, click and hold the eye icon on the right of the **Password** field.



Step 4. Review Recovery Image Settings

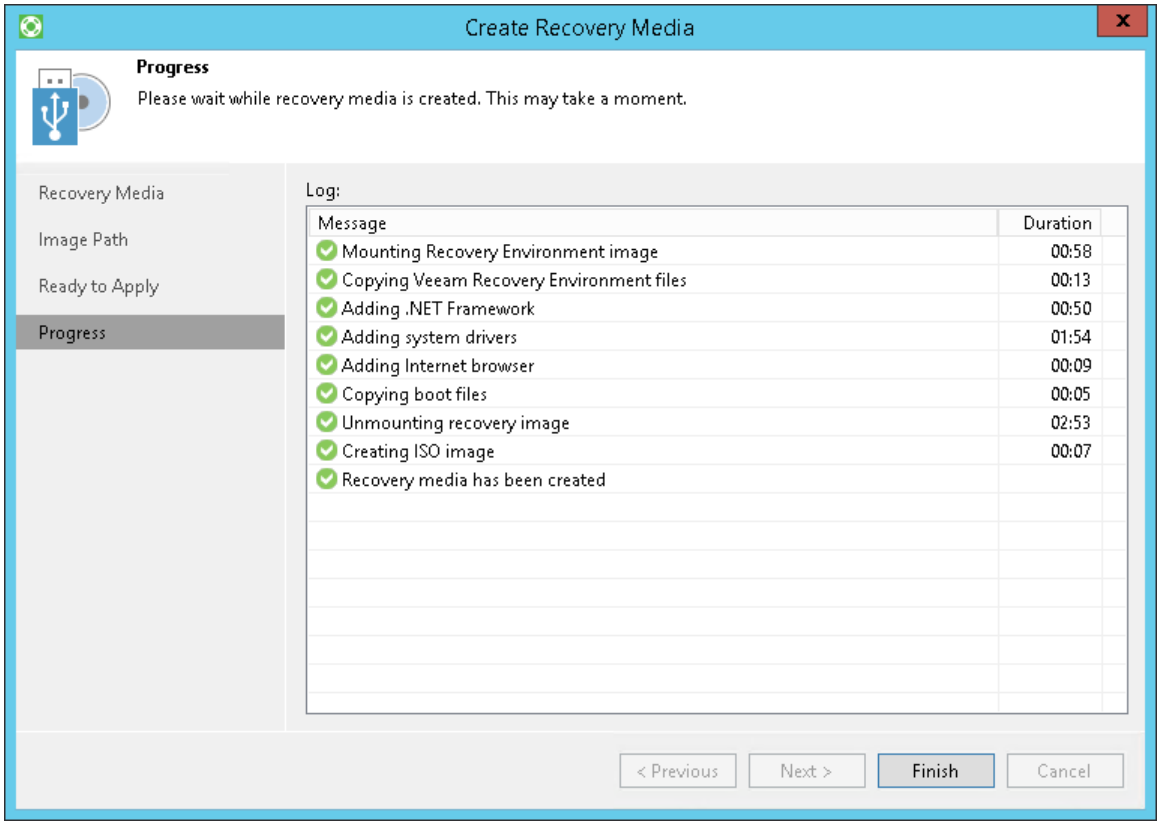
At the **Ready to Apply** step of the wizard, review settings of the recovery image that you plan to create and click **Create**.

Veeam Agent for Microsoft Windows will collect files necessary for recovery image creation and write the resulting recovery image to the specified target or burn it to CD/DVD/BD.



The process of recovery image creation may take some time. Wait for the process to complete and click **Finish** to exit the wizard.

If you want to interrupt the process of recovery image creation, click **Cancel** or close the wizard window.



What You Do Next

[For ISO] After the recovery image is created, you can burn the created ISO to a CD/DVD/BD. To do this, you can use native Microsoft Windows tools or third-party software.

Creating Veeam Recovery Media with Command Line Interface

In addition to creating the Veeam Recovery Media with the **Create Recovery Media** wizard, you can also use the command line interface to create the recovery image and save it as an ISO file. To create the Veeam Recovery Media, use a command with the following syntax:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /createreco
verymediaiso /f:<location>\<name>.iso
```

where:

- `<location>` — path to the folder in which the recovery image ISO file will be created. All folders that are mentioned in the path must be created in advance. Names of folders in the path must not contain spaces.
- `<name>` — name of the resulted ISO file. The file name must not contain spaces.

NOTE

Consider the following:

- When you create the Veeam Recovery Media using the command line interface, Veeam Agent for Microsoft Windows includes in the recovery image storage and network drivers that are currently installed on the Veeam Agent computer, and current network settings of this computer. If you want to include additional drivers and/or data decryption key in the recovery image, you must create the Veeam Recovery Media using the **Create Recovery Media** wizard. To learn more, see Creating Veeam Recovery Media.
- If you do not use the `/f:` option to specify a path to the folder in which to save the resulting ISO file, Veeam Agent for Microsoft Windows will save the ISO file to the `C:\Users\%UserProfile%\Documents` folder with the `VeeamRecoveryMedia_<machine>.iso` name, where `<machine>` — name of your Veeam Agent computer.
- You cannot save the ISO file to a network shared folder using the command line interface.

You can use the last exit code to verify if the backup job has completed successfully. To check the last exit code, use the *%ERRORLEVEL%* variable in `cmd.exe`. Veeam Agent can provide the following exit codes:

- 0 — recovery image successfully created
- 1 — the create recovery image operation failed

Example of Use

The example below creates an ISO file with the `VeeamRE` name in the `E:\Veeam\RecoveryImage` folder.

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /createreco
verymediaiso /f:E:\Veeam\RecoveryImage\VeeamRE.iso
```

Auto-Configuring Scheduled Backup Jobs

After the product installation, you can instruct Veeam Agent for Microsoft Windows to auto-configure the scheduled backup job with the default settings. The backup job will have the following settings:

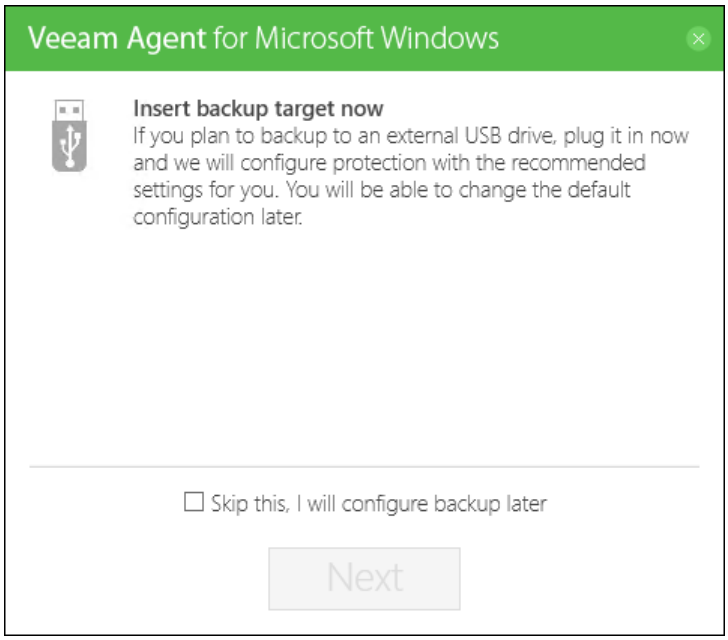
- Backup scope: entire computer
- Target destination: USB drive connected to the computer
- Schedule: 12:30 AM nightly

That is, the scheduled backup job will run regularly to create an entire computer backup at 12:30 AM and save this backup on the USB drive.

To auto-configure the scheduled backup job:

1. At the **Insert backup target now** step of the setup wizard, make sure that the **Skip this, I will configure backup later** check box is not selected.
2. Insert a USB drive to a USB slot on your computer.
3. Follow the steps of the installation wizard. At the last step of the wizard, click **Finish**.

> **IMPORTANT**
>
> USB storage devices formatted as FAT32 do not allow storing files larger than 4 GB in size. For this reason, we recommend that you do not use such USB storage devices as a backup target.



Configuring Backup Job

To back up your data, you must configure a backup job. The backup job defines how, where and when to back up data. You can choose one of the following backup types:

- Backup of an entire computer image
- Backup of specific computer volumes, for example, a system volume or secondary volume
- Backup of individual folders, for example, documents folder or folder with music

Before You Begin

Before you configure the backup job, check the following prerequisites:

- The target location where you plan to store backup files must have enough free space.
- Available backup job options depend on the edition of Veeam Agent for Microsoft Windows. Make sure that you have obtained and installed a license that has the desired number of instances. To learn more, see Licensing.
- [For Veeam Backup & Replication repository targets] You can store created backups in a backup repository only if the backup server runs Veeam Backup & Replication 11 or later. If you plan to use a commercial version of Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you must install a license in Veeam Backup & Replication that has a sufficient number of instances to protect machines with Veeam Agent. To learn more, see Managing License.
- [For Veeam Backup & Replication repository targets] If you plan to use a Veeam Backup & Replication repository as a target for backups, you must preconfigure user access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.
- [For Veeam Cloud Connect repository targets] If you plan to use a cloud repository as a target for backups, make sure that the service provider has communicated to you the necessary data: cloud gateway settings, user account settings and certificate thumbprint.
- [For Microsoft OneDrive targets] To log in to Microsoft OneDrive, Veeam Agent for Microsoft Windows uses the Internet Explorer web browser. ActiveX and JavaScript must be enabled in the web browser settings.
- A user account under which you launch the **New Backup Job** wizard must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

Backup has the following limitations:

- You cannot save the backup of entire computer on the local computer disk. Use an external hard drive or USB drive, network shared folder or backup repository as a target location.
- Veeam Agent for Microsoft Windows does not back up data to which symbolic links are targeted. It only backs up the path information that the symbolic links contain. After restore, identical symbolic links are created in the restore destination.
- Keep in mind that Veeam Agent stops running the backup job after 3 weeks.

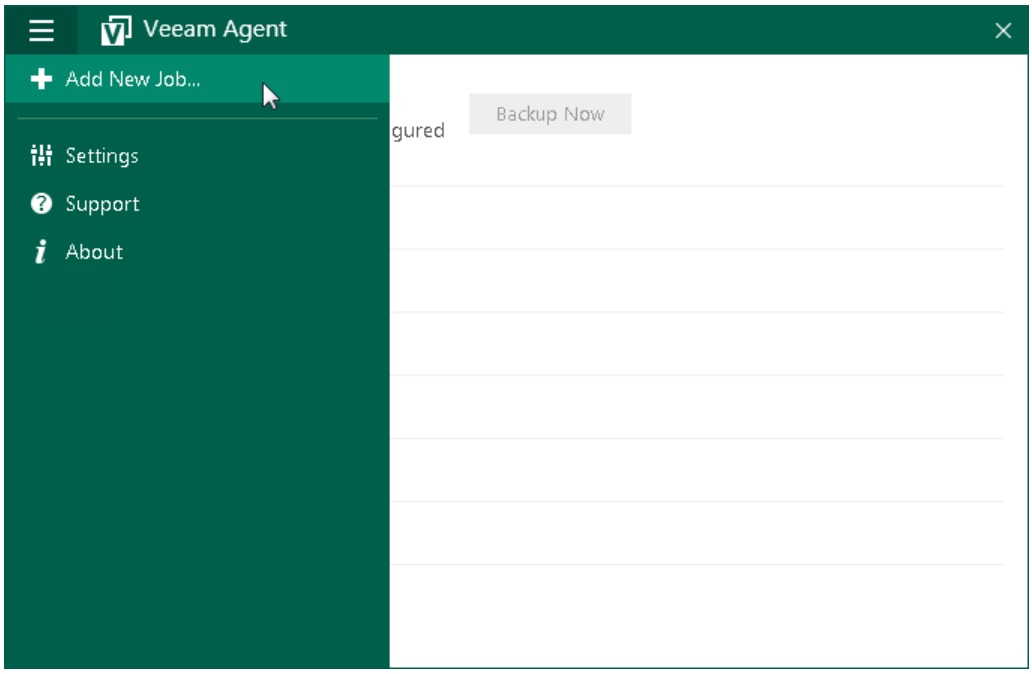
Step 1. Launch New Backup Job Wizard

To launch the **New Backup Job** wizard, do either of the following:

- [If no backup jobs are configured] Double-click the Veeam Agent for Microsoft Windows icon in the system tray.
- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Backup** > **Configure backup**.

> **NOTE**
>
> The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

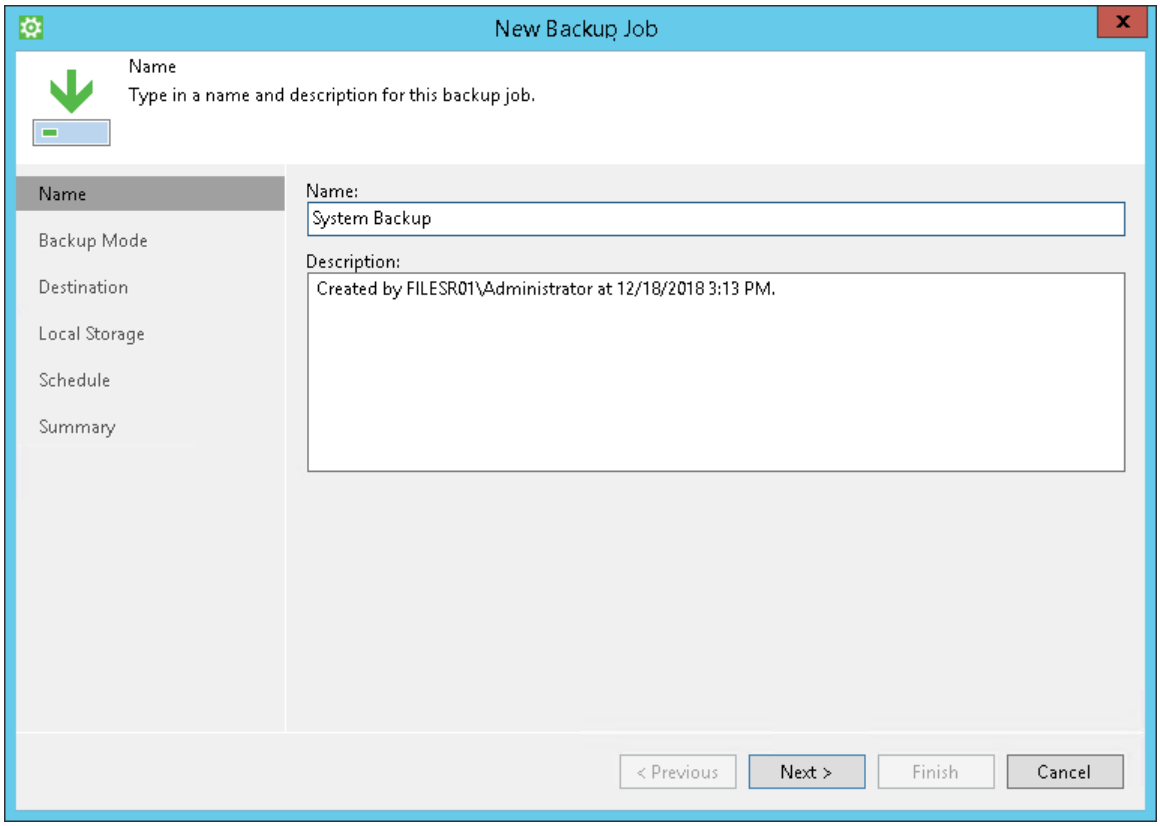
- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**. Then, from the main menu, select **Add New Job**.
- From the Microsoft Windows **Start** menu, select **All Programs** > **Veeam** > **Tools** > **Configure Backup** or use the Microsoft Windows search to find the **Configure Backup** option on your computer.



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.



Step 3. Select Backup Mode

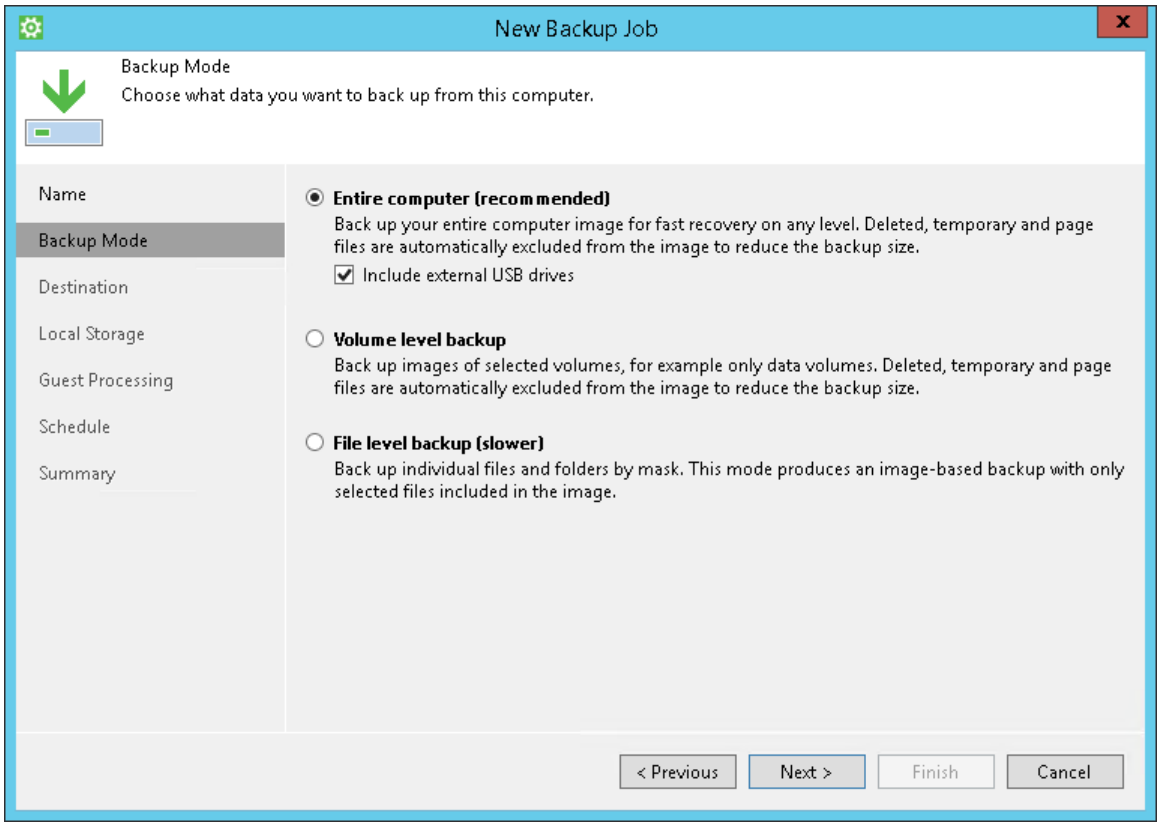
At the **Backup Mode** step of the wizard, select the mode in which you want to create a backup:

1. In the **Backup mode** section, select the backup mode. You can select one of the following options:
   - **Entire computer** — select this option if you want to create a backup of the entire computer image. When you restore data from such backup, you will be able to recover the entire computer image as well as data on specific computer volumes: files, folders, application data and so on. With this option selected, you will pass to the Destination step of the wizard.
   - **Volume level backup** — select this option if you want to create a backup of specific computer volumes, for example, all volumes except the system one. When you restore data from such backup, you will be able to recover data on these volumes only: files, folders, application data and so on. With this option selected, you will pass to the Volumes step of the wizard.
   - **File level backup** — select this option if you want to create a backup of individual folders on your computer. When you backup data in this mode, you can configure additional rules to include/exclude files of the specific type in/from your backup. When you restore data from such backup, you will be able to recover backed-up files and folders. With this option selected, you will pass to the Files step of the wizard.

2. [For entire computer backup] If you want to include in the backup one or more external USB drives, select the **Include external USB drives** check box. With this option selected, Veeam Agent for Microsoft Windows will include in the backup all supported external drives that are connected to the Veeam Agent computer at the time when the backup job starts. Veeam Agent for Microsoft Windows supports backup of external drives that support Microsoft VSS: HDD, SSD, and so on. USB flash drives (USB sticks) are not supported. To learn more, see Backup of External Drives.

TIP

File-level backup is typically slower and requires more network traffic than volume-level backup. Depending on the performance capabilities of your computer and backup environment, the difference between file-level and volume-level backup job performance may increase significantly. If you plan to back up all folders with files on a specific volume or back up large amount of data, we recommend that you configure volume-level backup instead of file-level backup.



Step 4. Specify Backup Scope Settings

Specify backup scope for the backup job:

- Select volumes to back up — if you have selected the **Volume level backup** option at the Backup Mode step of the wizard.
- Select folders to back up — if you have selected the **File level backup** option at the Backup Mode step of the wizard.

Selecting Volumes to Back Up

The **Volumes** step of the wizard is available if you have chosen to create a volume-level backup.

At this step of the wizard, you must specify the backup scope — define what volumes you want to include in the backup.

To specify the backup scope, choose volumes and objects that you want to include in the backup. You can include the following data in the backup:

- Computer volumes (located on local drives or external USB drives). To include individual volumes of your computer to the backup scope, select check boxes next to necessary volumes.

  If you select all volumes at this step of the wizard, Veeam Agent will create a backup of the entire computer image. However, if a new volume appears on a Veeam Agent computer, Veeam Agent will not include this volume in the backup scope automatically. Veeam Agent will back up only those volumes that are explicitly selected for backup.

- Mount points. To include individual volumes mounted on your computer to the backup scope, select check boxes next to necessary mount points.

  In the list of objects to back up, Veeam Agent displays available mount points after volumes.

- System state data. To include system state data into the backup, do one of the following:

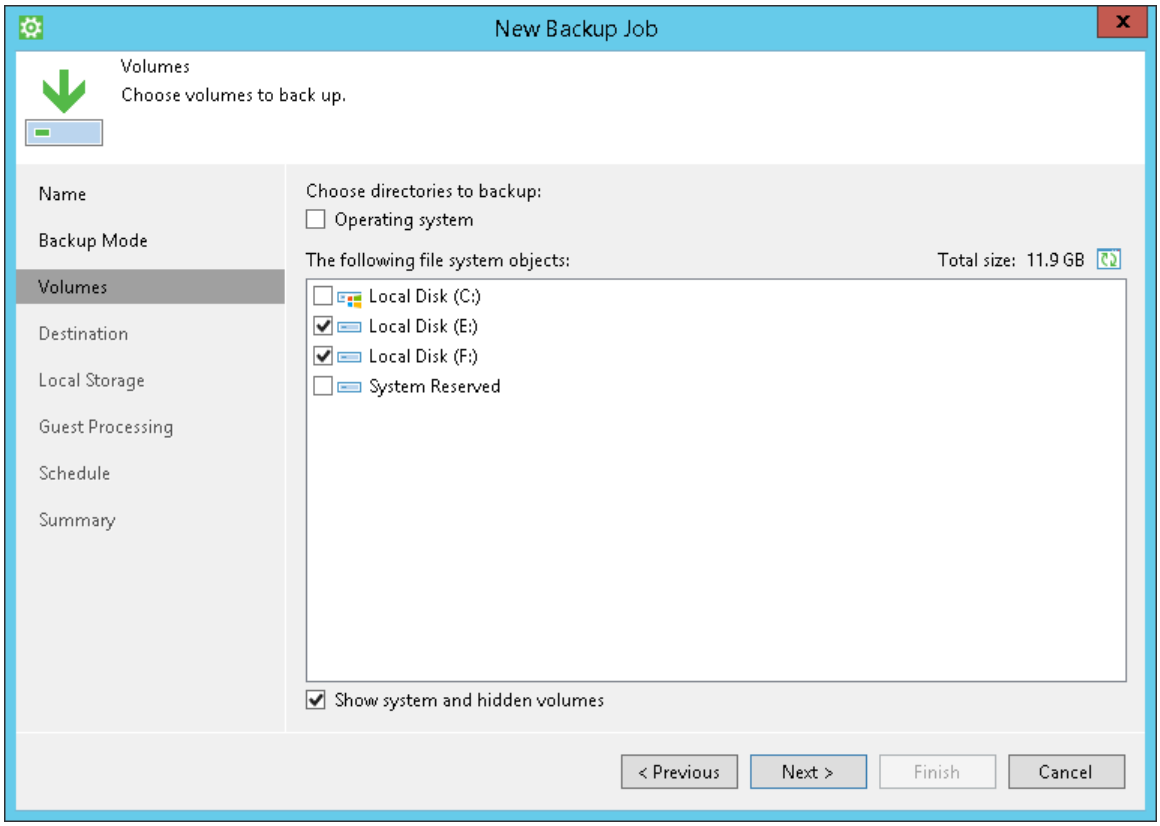
  - In the **Choose directories to backup** list, select the **Operating system** check box.
  - In the **The following file system objects** list, select the check box next to the system volume. Typically, the system volume is the **Local Disk (C:)** volume
  - Select the **Show system and hidden volumes** check box at the bottom of the window. In the **The following file system objects** list, select the **System Reserved** check box and the check box next to the system volume. Typically, the system volume is the **Local Disk (C:)** volume.

  Veeam Agent for Microsoft Windows will include in the backup scope the Microsoft Windows system partition and boot partition of your computer. For GPT disks on Microsoft Windows 8.1, 10, 2012, 2012 R2, 2016 and 2019, Veeam Agent will additionally back up the recovery partition. To learn more, see System State Data Backup.

When you include a system volume in the backup, Veeam Agent automatically includes the System Reserved/UEFI or other system partitions in the backup too. If you do not want to back up the system state data, you can clear the **System Reserved** check box. However, in this case Veeam Agent does not guarantee that the OS will boot properly when you attempt to recover from such backup. To learn more, see System State Data Backup.

NOTE

Veeam Agent automatically adds to the list of exclusions the following Microsoft Windows objects for all computer users: temporary files folder, Recycle Bin, Microsoft Windows pagefile, hibernate file and VSS snapshot files from the System Volume Information folder.



Selecting Folders to Back Up

The **Files** step of the wizard is available if you have chosen to create a file-level backup.

In the file-level backup mode, you can create two types of backups:

- File-level backup that includes individual folders on your computer.
- Hybrid backup that contains individual folders and specific volumes of your computer.

At this step of the wizard, you must specify the backup scope — define what folders with files or entire volumes you want to include in the backup.

To specify the backup scope, select check boxes next to necessary objects. You can include the following data in the backup:

- Operating system data — data related to the OS installed on your computer.
- Personal files — data related to user profiles. With this option enabled, Veeam Agent for Microsoft Windows will include in the backup scope settings and data related to Veeam Agent computer user profiles.

  You can specify what personal data to include in the backup and choose whether to exclude roaming profiles from the backup. To do this, click **Choose** next to the **Personal files** field and select the necessary options in the **Personal Folders** window. To learn more, see Personal Data Backup.
- System reserved data — system data required to boot the OS installed on your computer. With this option enabled, Veeam Agent for Microsoft Windows will include in the backup scope Microsoft Windows system partition and boot partition of your computer. For GPT disks on Microsoft Windows 8.1, 10, 2012, 2012 R2, 2016 and 2019, Veeam Agent for Microsoft Windows will additionally back up the recovery partition. To learn more, see System State Data Backup.
- Individual folders. To learn more, see Selecting Folders.
- Individual computer volumes and mount points (located on local drives or external USB drives). To learn more, see Selecting Volumes.

IMPORTANT

Consider the following:

- Veeam Agent for Microsoft Windows does not include the following Microsoft Windows objects in the backup: temporary files folder, Recycle Bin, Microsoft Windows pagefile, hibernate file and VSS snapshot files from the System Volume Information folder.
- If you add Microsoft OneDrive folders to the backup scope, Veeam Agent includes only local copies of files that are available on the Veeam Agent computer. To be backed-up, these files must be available even when Veeam Agent computer is not connected to the Internet. To make sure that all the content of Microsoft OneDrive folders is backed-up, download all Microsoft OneDrive files and folders to the Veeam Agent computer. For details, see this Veeam KB article.
- If you have a file-level backup job configured and need to extend the volume where backed-up files reside, we strongly recommend to create an active full backup after the volume is extended. Otherwise, Veeam Agent may skip files during the job run even if these files are added to the backup scope. To learn more, see Creating Active Full Backups.

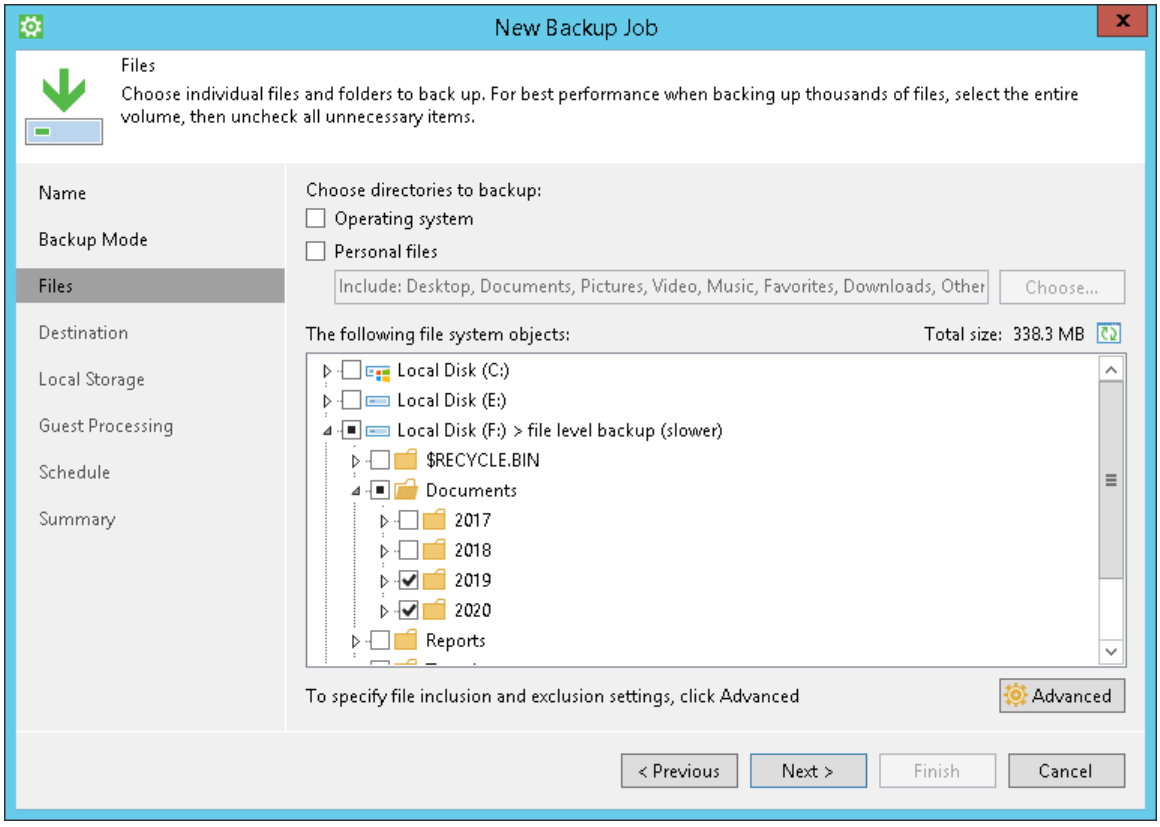
Selecting Folders

To include a folder in the file-level backup:

1. In the **The following file system objects** list, browse to the folder that you want to back up and select the check box next to this folder.

   After you select a folder, Veeam Agent will display the *file level backup* mark next to the volume that contains this folder.

2. If you want to include or exclude files of a specific type in/from the backup, click **Advanced**. To learn more, see Configuring Filters.



Alternatively, if you want to quickly include a large number of folders in the backup, you can select folders in the following way:

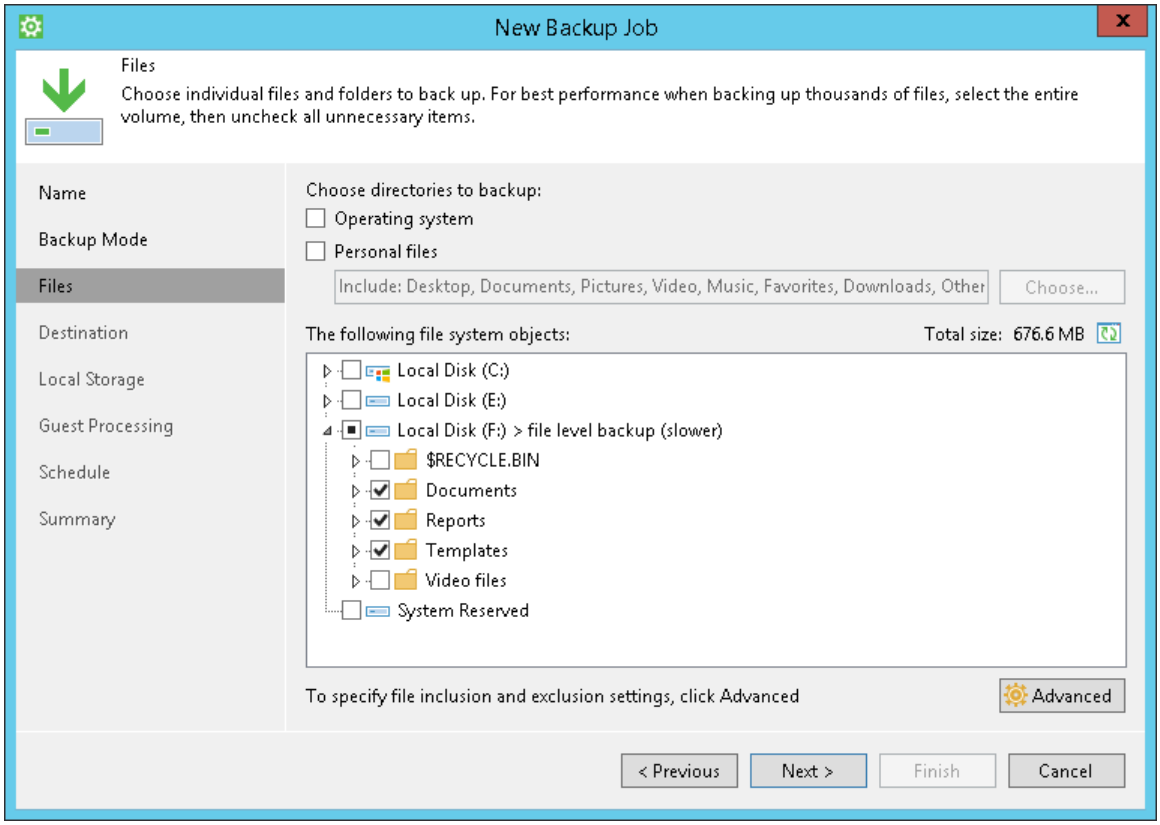
1. Select the check box next to the volume that contains folders you want to back up.

   Veeam Agent will display the *volume level backup* mark next to the selected volume.

2. Expand the volume and clear the check boxes next to one or more folders that you want to exclude from the backup.

3. Select the check box next to the volume once again.

   Veeam Agent will display the *file level backup* mark next to the volume. You can include or exclude files that reside on this volume in/from the backup. To learn more, see Configuring Filters.



Selecting Volumes

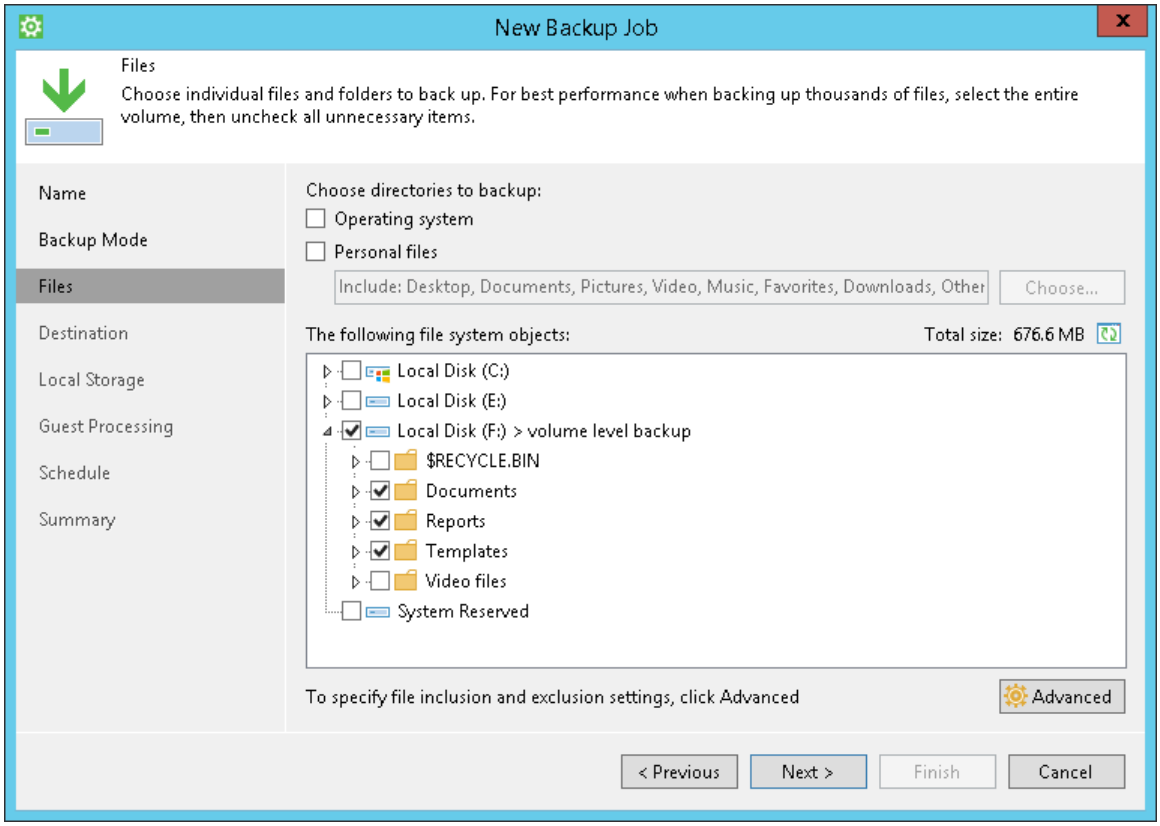
You can include individual volumes in the backup at the **Files** step of the wizard. This may be useful if you want to have a volume-level backup of the volume with specific folders excluded. Volume-level backup also offers better performance in case you back up a complex folder structure with many files.

You can also include individual mount points in the backup at the **Files** step of the wizard, but you cannot exclude folders from mount points.

To include a volume or a mount point in the file-level backup:

1. In the **The following file system objects** list, select the check box next to the item that you want to back up.
2. If you want to exclude a folder on the selected volume from the backup, expand the volume and clear the check box next to the necessary folder.

> **NOTE**
>
> If you select a volume for backup, you cannot use filters to include or exclude files of a specific type in/from the backup. You can only exclude specific folders that reside on the volume. If you want to include or exclude files, you must select individual folders for backup. To learn more, see Selecting Folders.



Configuring Filters

To include or exclude files of a specific type in/from the file-level backup, you can configure filters.

> **NOTE**
>
> Veeam Agent for Microsoft Windows applies filters to files in specific folders that you include in the backup. Filters are not applied to computer volumes and mount points selected for backup. If you plan to create a hybrid backup that will contain volumes, mount points, and folders, filters will be applied to files in folders only.

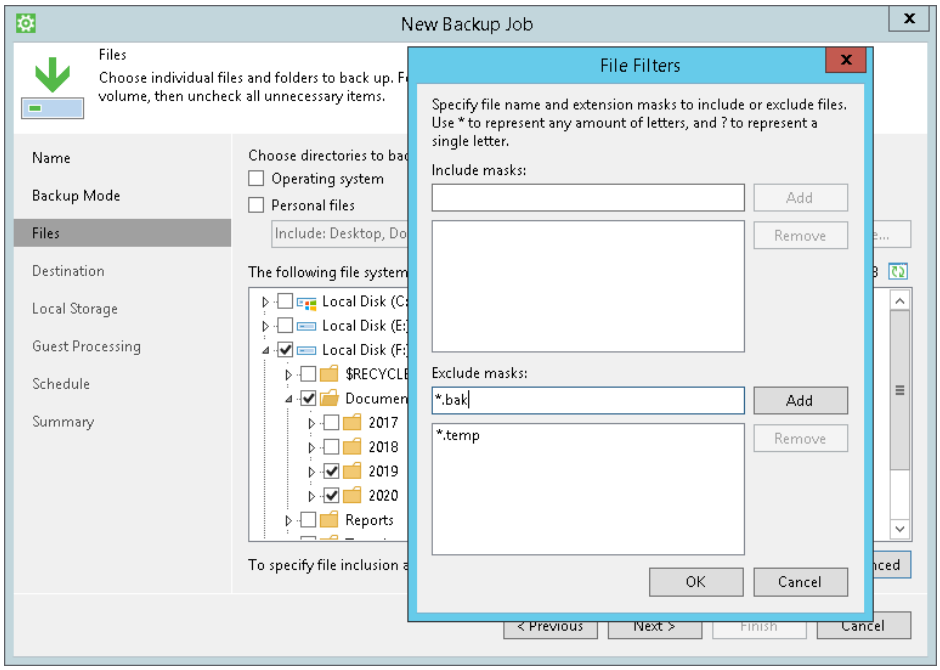
To configure a filter:

1. At the **Files** step of the wizard, click **Advanced**.
2. Specify what files you want to back up:
   - In the **Include masks** field, specify file names and/or masks for file types that you want to back up, for example, `MyMovie.avi`, `*filename*`, `*.docx`, `*.mp3`. Veeam Agent will create a backup only for selected files. Other files will not be backed up.
   - In the **Exclude masks** field, specify file names and/or masks for file types that you do not want to back up, for example, `OldPhotos.rar`, `*.temp`, `*.tmp`, `*.back`. Veeam Agent will back up all files except files of the specified type.
3. Click **Add**.
4. Repeat steps 2–3 for each mask that you want to add.

You can use a combination of include and exclude masks. Keep in mind that exclude masks have a higher priority than include masks. For example, you can specify masks in the following way:

- Include mask: `*.avi`
- Exclude mask: `*movie*`

Veeam Agent will include in the backup all files of the AVI format that do not contain *movie* in their names.

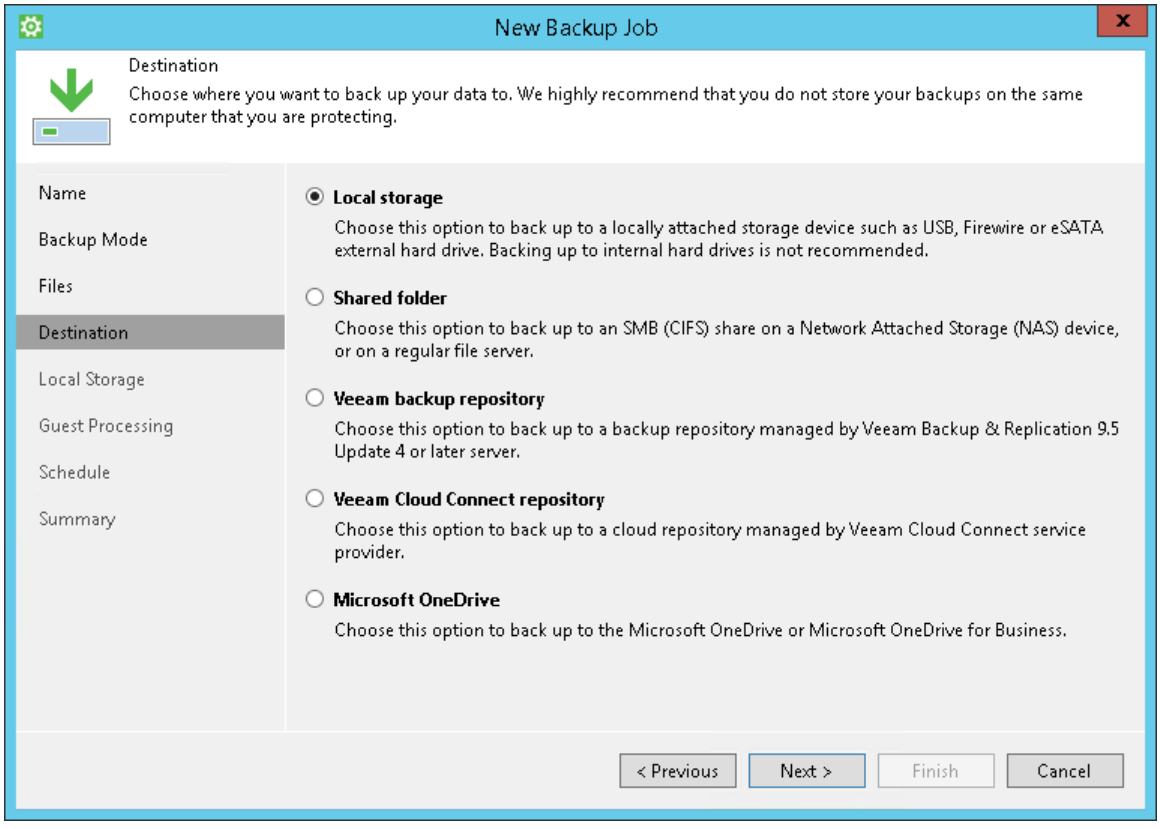

New Backup Job
Files
Choose individual files and folders to back up.
volume, then uncheck all unnecessary items.
Name
Backup Mode
Files
Destination
Local Storage
Guest Processing
Schedule
Summary
Choose directories to ba
Operating system
Personal files
Include: Desktop, Do
The following file system
Local Disk (C:
Local Disk (E:
Local Disk (F:
$RECYCLE
Documen
2017
2018
2019
2020
Reports
To specify file inclusion a
File Filters
Specify file name and extension masks to include or exclude files. Use * to represent any amount of letters, and ? to represent a single letter.
Include masks:
Add
Remove
Exclude masks:
*.bak
Add
*.temp
Remove
OK
Cancel
< Previous
Next >
Finish
Cancel

Step 5. Select Backup Destination

At the **Destination** step of the wizard, select a target location for the created backup.

You can store backup files in one of the following locations:

- **Local storage** — select this option if you want to save the backup on a removable storage device attached to the computer or on a local computer drive. With this option selected, you will pass to the Local Storage step of the wizard.
- **Shared folder** — select this option if you want to save the backup in a network shared folder. With this option selected, you will pass to the Shared folder step of the wizard.
- **Veeam backup repository** — select this option if you want to save the backup on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the Backup Server step of the wizard.
- **Veeam Cloud Connect repository** — select this option if you want to save the backup on a cloud repository exposed to you by the Veeam Cloud Connect service provider. With this option selected, you will pass to the Service Provider step of the wizard.
- **Microsoft OneDrive** — select this option if you want to save the backup in the Microsoft OneDrive cloud storage. With this option selected, you will pass to the Microsoft OneDrive step of the wizard.



IMPORTANT

Consider the following:

- We strongly recommend that you store backups in the external location like USB storage device or shared network folder. You can also keep your backup files on the separate non-system local drive.
- If you select to store the backup on a local folder included in the backup scope, Veeam Agent will automatically exclude this folder from the backup.
- The **Microsoft OneDrive** option is not available if the Veeam Agent computer runs a Microsoft Windows Server OS.
- For the Free product edition, the **Veeam Cloud Connect repository** option is not available.
- For the Free product edition, if you want to save backup file in Microsoft OneDrive, you can create backups that contain operating system data only. To learn more, see System State Data Backup.

Step 6. Specify Backup Storage Settings

Specify backup storage settings for the backup job:

- Local storage settings — if you have selected the **Local storage** option at the Destination step of the wizard.
- Shared folder settings — if you have selected the **Shared folder** option at the Destination step of the wizard.
- Veeam backup repository settings — if you have selected the **Veeam backup repository** option at the Destination step of the wizard.
- Veeam Cloud Connect repository settings — if you have selected the **Veeam Cloud Connect repository** option at the Destination step of the wizard.
- Microsoft OneDrive settings — if you have selected the **Microsoft OneDrive** option at the Destination step of the wizard.

Local Storage Settings

The **Local Storage** step of the wizard is available if you have chosen to save the backup on a local drive of your computer.

Specify local storage settings:

1. In the **Local drives** list, select a drive where you want to store the backup.
2. In the **Folder** field, specify a path to the folder where backup files must be saved. By default, Veeam Agent saves files in the `VeeamBackup` folder.
3. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

   To learn more, see Mapping Backup Job.
4. Specify short-term backup retention policy settings:

   - [For Free and Workstation product editions] In the **Keep backups for <N> days (excluding days with no backup)** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days. After this period is over, Veeam Agent will remove the earliest restore points from the backup chain.

     To learn more, see Backup Retention Policy in Free and Workstation Editions.
   - [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

     - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
     - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

     To learn more, see Backup Retention Policy in Server Edition.
5. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the GFS Retention Policy section in the Veeam Backup & Replication User Guide.
6. Click **Advanced** to specify advanced settings for the backup job. To learn more, see Specify Advanced Backup Settings.

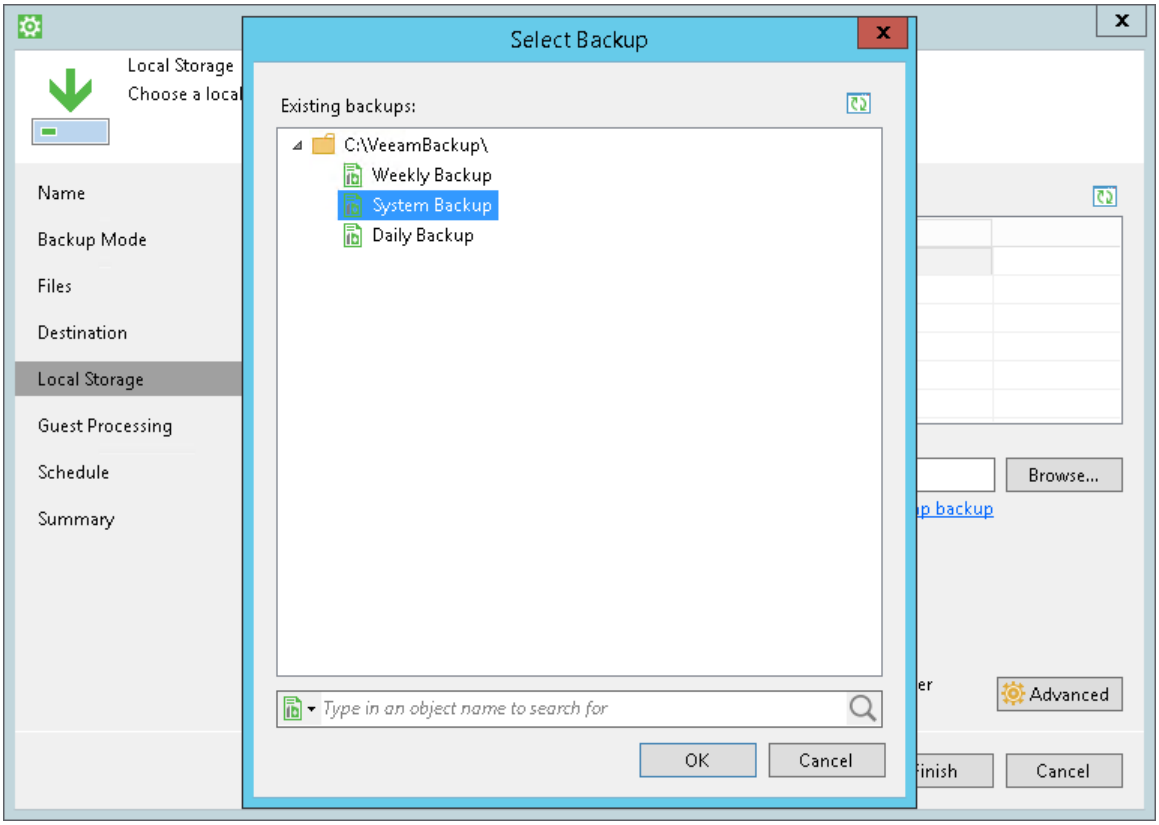
> **IMPORTANT**
>
> USB storage devices formatted as FAT32 do not allow storing files larger than 4 GB in size. For this reason, we recommend that you do not use such USB storage devices as a backup target.

Mapping Backup Job

If you have selected to map the job to the already created backup that is stored on a local drive, perform the following steps:

1. In the **Folder** field, specify the same path as was specified in the **Folder** field for the job that was used to create the backup.
2. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.

   Keep in mind that Veeam Agent displays backups stored in the folder that is specified in the **Folder** field and its first-level subfolders.



3. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.

Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have chosen to save the backup in a network shared folder.

Specify shared folder settings:

1. In the **Shared folder** field, type a UNC name of the network shared folder in which you want to store backup files. Keep in mind that the UNC name always starts with two back slashes (\\).
1. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

   To learn more, see Mapping Backup Job.
2. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the down-level logon name format. For example, *DOMAIN\UserName* or *HOSTNAME\UserName*.

   To view the specified password, click and hold the eye icon on the right of the **Password** field.

   If you do not select the **This share requires access credentials** check box, Veeam Agent for Microsoft Windows will connect to the shared folder using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can grant *Full Control* access permissions on the shared folder and underlying file system to the computer account *(DOMAIN\COMPUTERNAME$)*.
3. To view how much free space is available in the selected shared folder, click **Populate**.
4. Specify short-term backup retention policy settings:

   - [For Free and Workstation product editions] In the **Keep backups for <N> days (excluding days with no backup)** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent for Microsoft Windows keeps backup files for 7 days. After this period is over, Veeam Agent for Microsoft Windows will remove the earliest restore points from the backup chain.

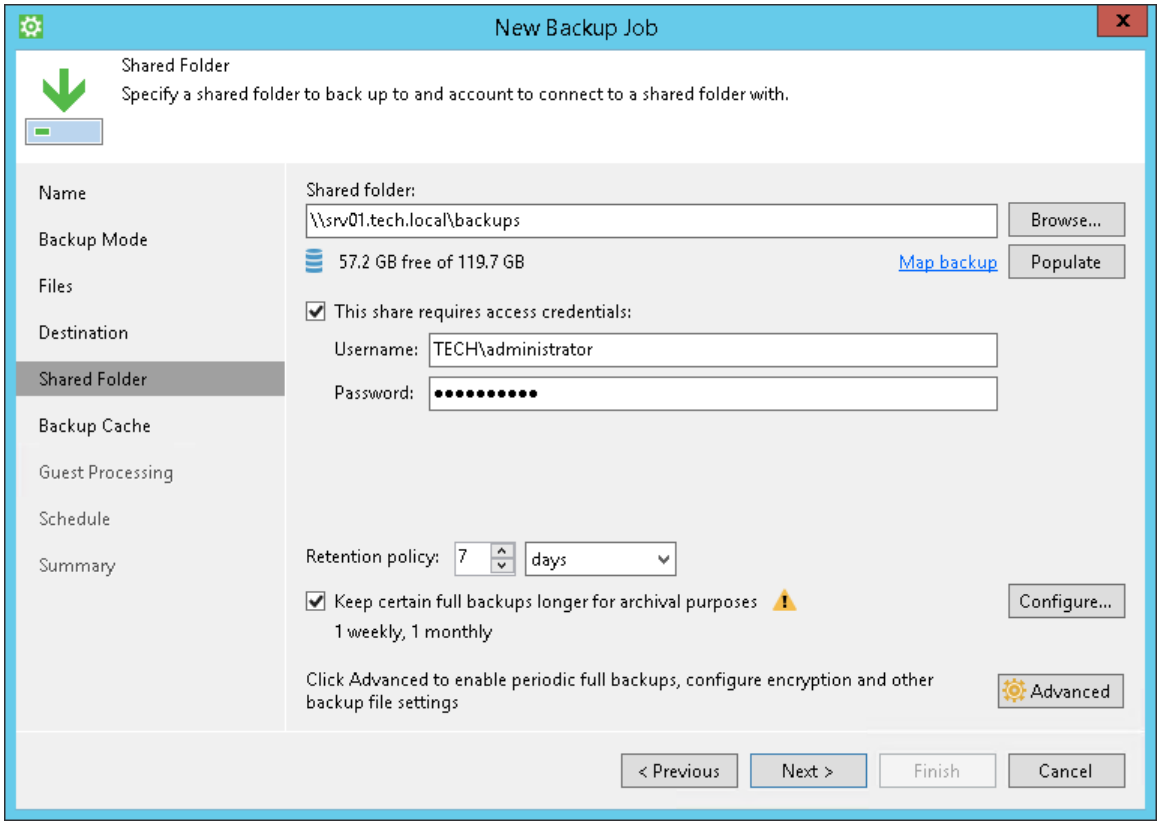
     To learn more, see Backup Retention Policy in Free and Workstation Editions.
   - [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

     - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
     - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

     To learn more, see Backup Retention Policy in Server Edition.

5. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the GFS Retention Policy section in the Veeam Backup & Replication User Guide.

6. Click **Advanced** to specify advanced settings for the backup job. To learn more, see Specify Advanced Backup Settings.

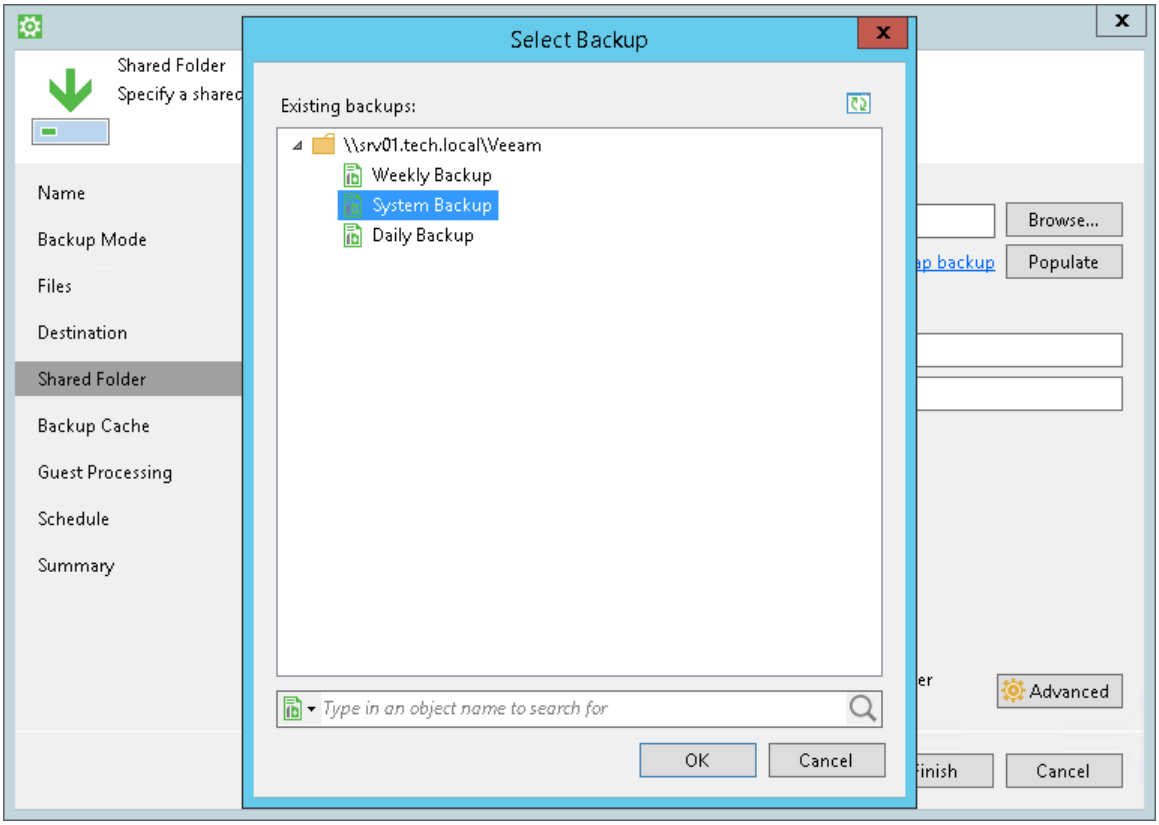


Mapping Backup Job

If you have selected to map the job to the already created backup that is stored on a network shared folder, perform the following steps:

1. In the **Shared folder** field, specify the same path as was specified in the **Shared folder** field for the job that was used to create the backup.
2. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.

   Keep in mind that Veeam Agent displays backups stored in the folder that is specified in the **Shared folder** field and its first-level subfolders.



3. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.

Veeam Backup Repository Settings

If you have selected to store backup files on a Veeam backup repository, specify settings to connect to the backup repository:

1. At the Backup Server step of the wizard, specify backup server settings.
2. At the Backup Repository step of the wizard, select the Veeam backup repository.

Specifying Backup Server Settings

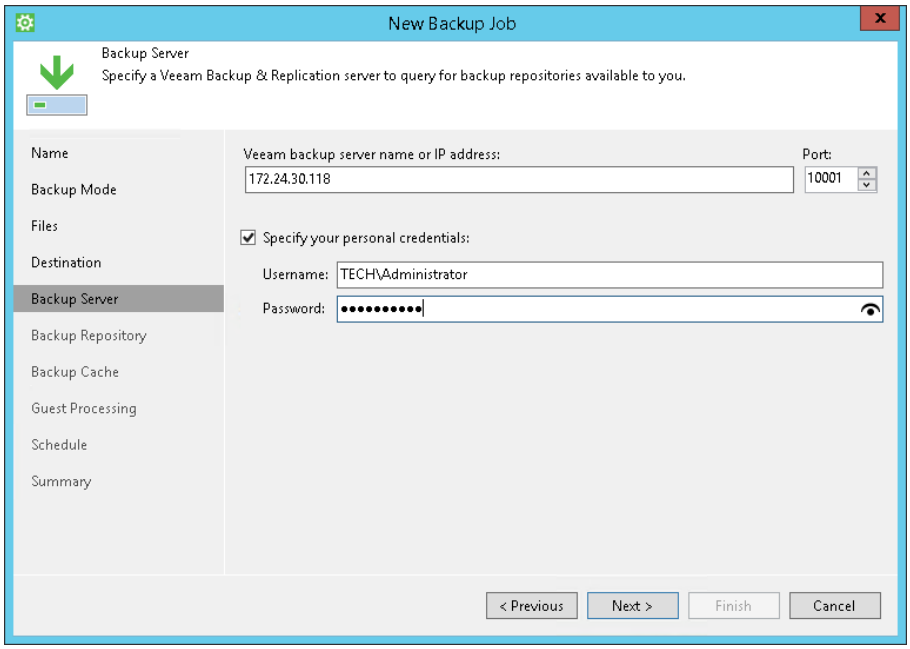
The **Backup Server** step of the wizard is available if you have chosen to store backup files on a Veeam backup repository.

Specify settings for the Veeam backup server that manages the target backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup repository. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. Select the **Specify your personal credentials** check box. In the **Username** and **Password** fields, specify a user name and password of the account that has access to this backup repository. Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see Setting Up User Permissions on Backup Repositories.

   If you do not select the **Specify your personal credentials** check box, Veeam Agent for Microsoft Windows will connect to the backup repository using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can add the computer account (*DOMAIN\COMPUTERNAME$*) to an AD group and grant access rights on the backup repository to this group.

   Setting access permissions on the backup repository to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). If you have set such permissions on the backup repository, you can omit specifying credentials. However, we recommend this scenario for demo environments only.



IMPORTANT

Mind the following:

- If you plan to use a commercial version of Veeam Agent with Veeam Backup & Replication, you must install a license in Veeam Backup & Replication in advance, before connecting to the backup server.
- If you specify a DNS name of the Veeam backup server, make sure that the Veeam backup server name is resolved into IPv4 address on the machine where Veeam Agent is installed. The Veeam Backup Service in Veeam Backup & Replication listens on IPv4 addresses only. If the Veeam backup server name is resolved into IPv6 address, Veeam Agent will fail to connect to the Veeam backup server.
- If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain on the backup repository.

Selecting Backup Repository

The **Backup Repository** step of the wizard is available if you have chosen to save backup files on a Veeam backup repository.

Specify settings for the target backup repository:

1. From the **Backup repository** list, select a backup repository where you want to store created backups. The **Backup repository** list displays only those backup repositories on which you have permissions to store data.

   To store Veeam Agent backups, you can use a simple backup repository or a scale-out backup repository.

   To refresh the list of backup repositories, click the **Refresh** button at the top right corner of the **Backup repository** field. Backup repositories list refresh may be required if you change permission settings for a specific backup repository on the Veeam backup server and want to display this backup repository in the **New Backup Job** wizard. To learn more, see Setting Up User Permissions on Backup Repositories.

1. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

   To learn more, see Mapping Backup Job.

3. Specify short-term backup retention policy settings:

   - [For Free and Workstation product editions] In the **Keep backups for <N> days (excluding days with no backup)** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days. After this period is over, Veeam Agent will remove the earliest restore points from the backup chain.

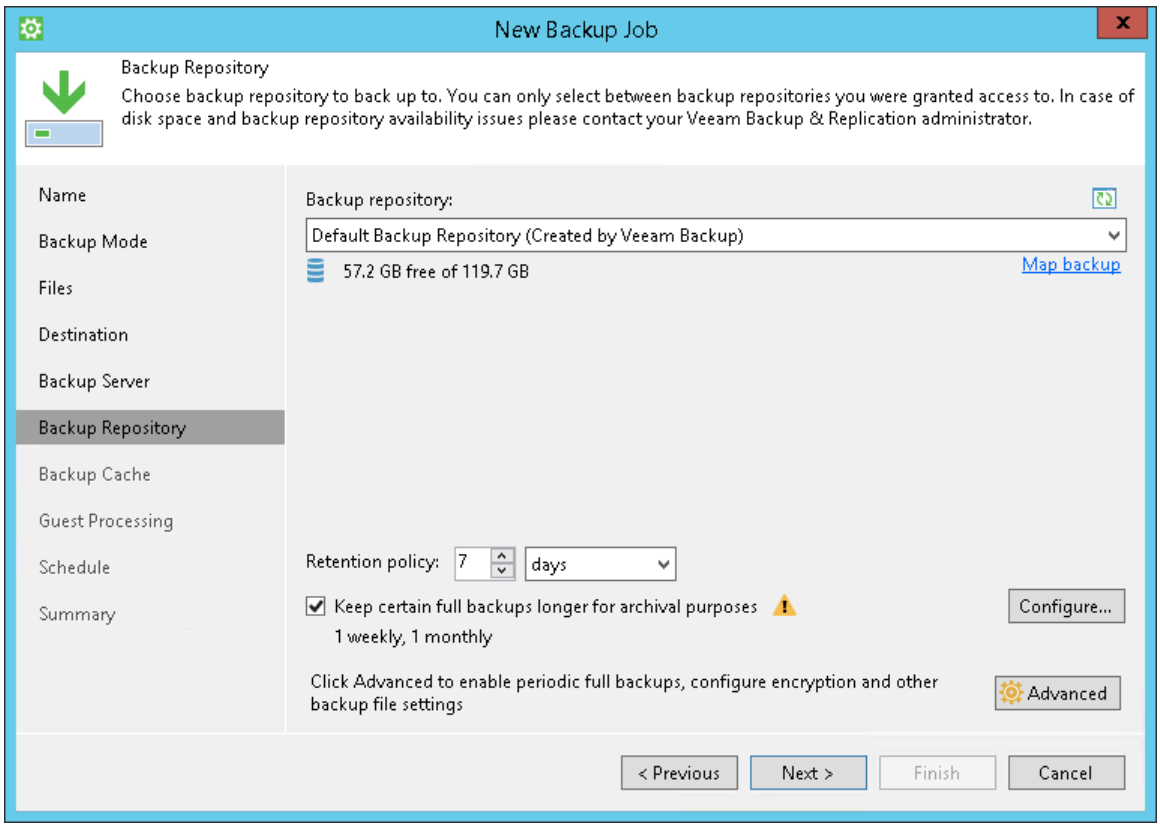
     To learn more, see Backup Retention Policy in Free and Workstation Editions.

   - [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

     - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
     - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

     To learn more, see Backup Retention Policy in Server Edition.

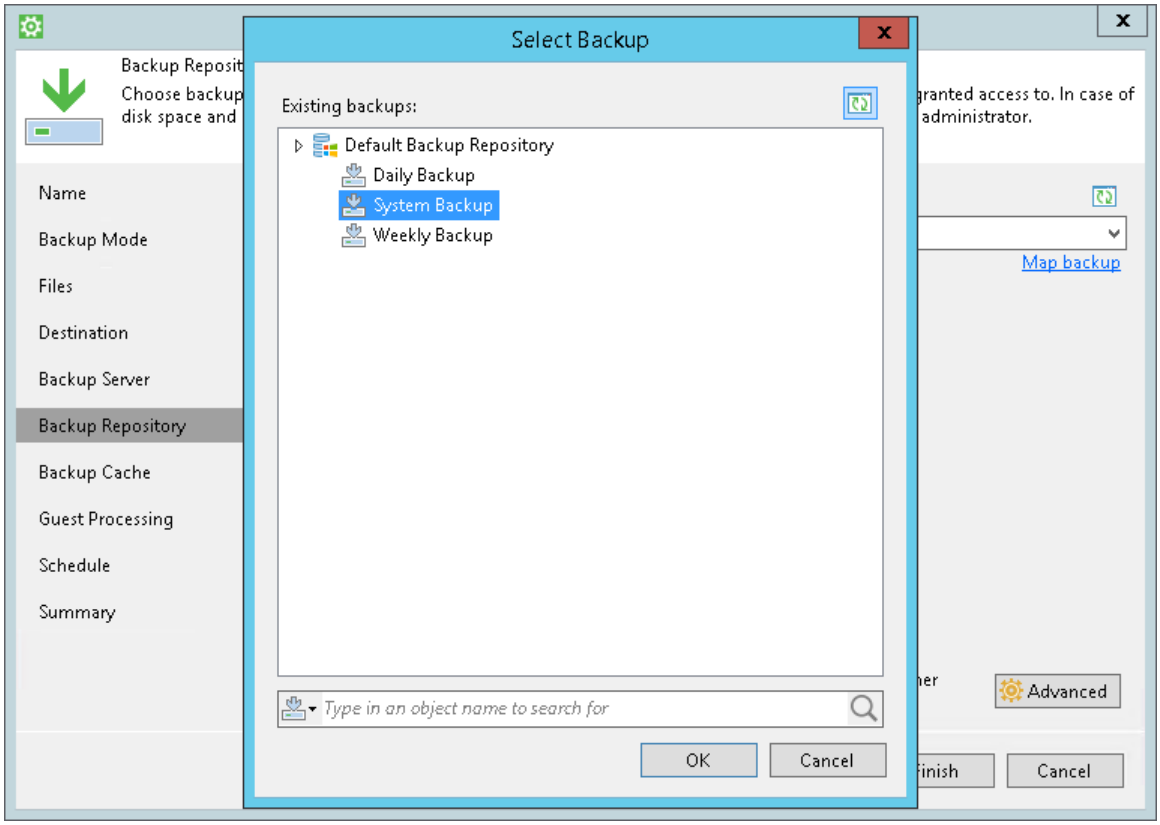
4. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the GFS Retention Policy section in the Veeam Backup & Replication User Guide.
5. Click **Advanced** to specify advanced settings for the backup job. To learn more, see Specify Advanced Backup Settings.



Mapping Backup Job

If you have selected to map the job to the already created backup that is stored on a Veeam backup repository, perform the following steps:

1. If you map the job to an encrypted backup, you must decrypt the backup on the Veeam Backup & Replication side before mapping. To learn more, see the Restoring Data from Encrypted Backups section in the Veeam Backup & Replication User Guide.
2. Click the **Map backup** link and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window. Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same account used to connect to the repository.



The backup job mapping can be helpful in case of the backup seeding. If you want to seed backup files, you must place these files in a specific folder on the target location. On the Veeam backup repository, Veeam Agent stores backups in the folders with paths of the following format:

```
<path_to_folder>\<domain_name>_<user_name>\<backup_job_name>
```

where:

- `<path_to_folder>` — path to the folder on the Veeam backup repository.
- `<domain_name>` — domain name specified to connect to the repository.
- `<user_name>` — user name specified to connect to the repository.
- `<backup_job_name>` — name of the backup job.

For example:

```
C:\Backup\TECH_Administrator\System Backup
```

To learn more about backup seeding scenarios for Veeam backup repositories, see this Veeam KB article.

Veeam Cloud Connect Repository Settings

If you have selected to store backup files on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. Specify service provider settings.
2. Verify the TLS certificate and specify user account settings.
3. Select the cloud repository.

Specifying Service Provider Settings

The **Service Provider** step of the wizard is available if you have chosen to save backup files on a Veeam Cloud Connect repository.

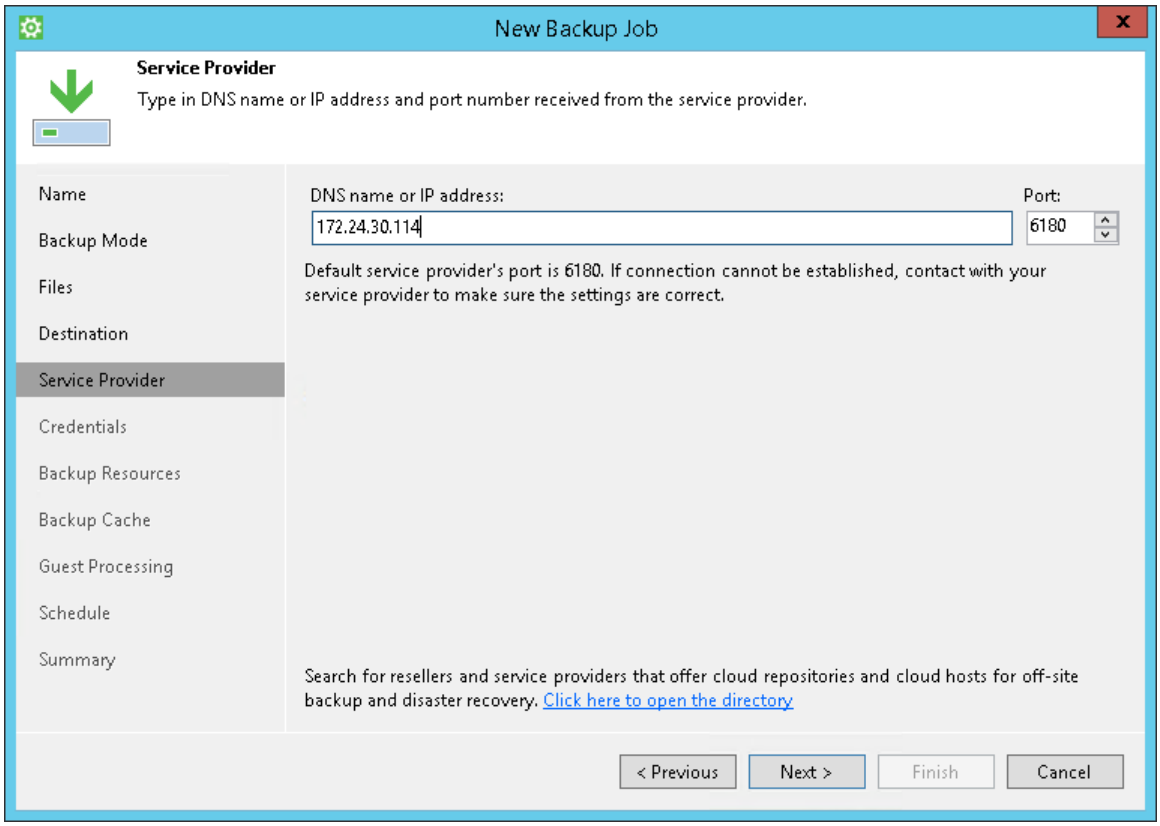
Specify settings for the cloud gateway that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.

TIP

You can look for service providers who offer Repository as a Service using Veeam Backup & Replication. The list of service providers is published on the Veeam website and constantly updated. You can select the necessary service provider from the list and contact this service provider to get the cloud repository service.

To find a service provider, click the *Click here to open the directory* link. Veeam Agent for Microsoft Windows will open a webpage on the Veeam website. Use the filter on the webpage to find the necessary service provider by the type of provided cloud services, service provider datacenter location or service area.



Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to save backup files on a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the cloud repository.

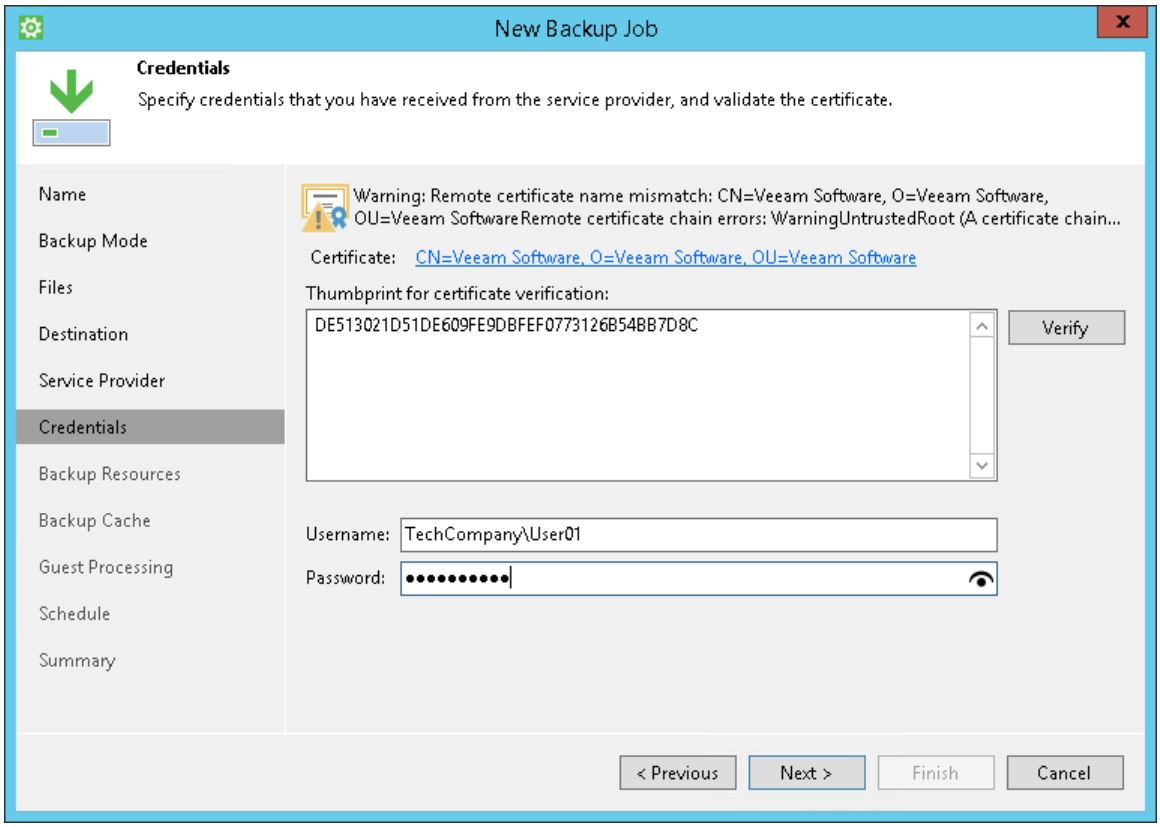
1. At the top of the wizard window, Veeam Agent for Microsoft Windows displays information about the TLS certificate obtained from the SP side. You can view the certificate settings and verify the TLS certificate.

   TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

   - To view the TLS certificate, click the certificate link.
   - To verify the TLS certificate with a thumbprint, copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Thumbprint for certificate verification** field. Click **Verify**. Veeam Agent for Microsoft Windows will check if the thumbprint you enter matches the thumbprint of the obtained TLS certificate.

2. In the **Username** field, enter the user name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The user name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.



Selecting Cloud Repository

The **Backup Resources** step of the wizard is available if you have chosen to save backup files on a cloud repository and specified settings to connect to the SP.

Specify settings for the cloud repository:

1. From the **Available cloud repositories** list, select a cloud repository where you want to store created backups. The **Available cloud repositories** list displays only those backup repositories on which you have permissions to store data.

   To refresh the list of cloud repositories, click the **Refresh** button at the top right corner of the **Available cloud repositories** field. Cloud repositories list refresh may be required if you change permission settings for a specific cloud repository and want to display this cloud repository in the **New Backup Job** wizard.

2. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

   To learn more, see Mapping Backup Job.

3. Specify short-term backup retention policy settings:

   - [For Free and Workstation product editions] In the **Keep restore points for the last <N> days when computer was used** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent for Microsoft Windows keeps backup files for 7 days. After this period is over, Veeam Agent for Microsoft Windows will remove the earliest restore points from the backup chain.

     To learn more, see Backup Retention Policy in Free and Workstation Editions.

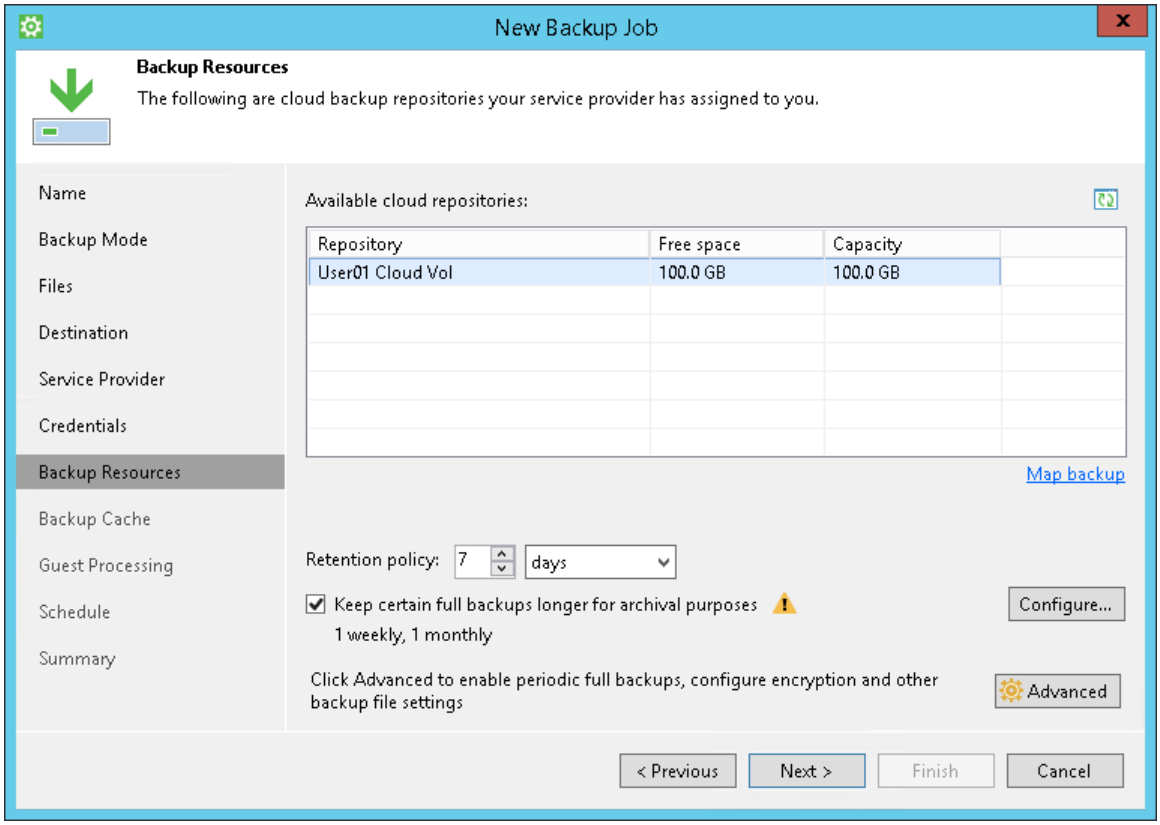
   - [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

     - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
     - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

     To learn more, see Backup Retention Policy in Server Mode.

4. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the GFS Retention Policy section in the Veeam Backup & Replication User Guide.

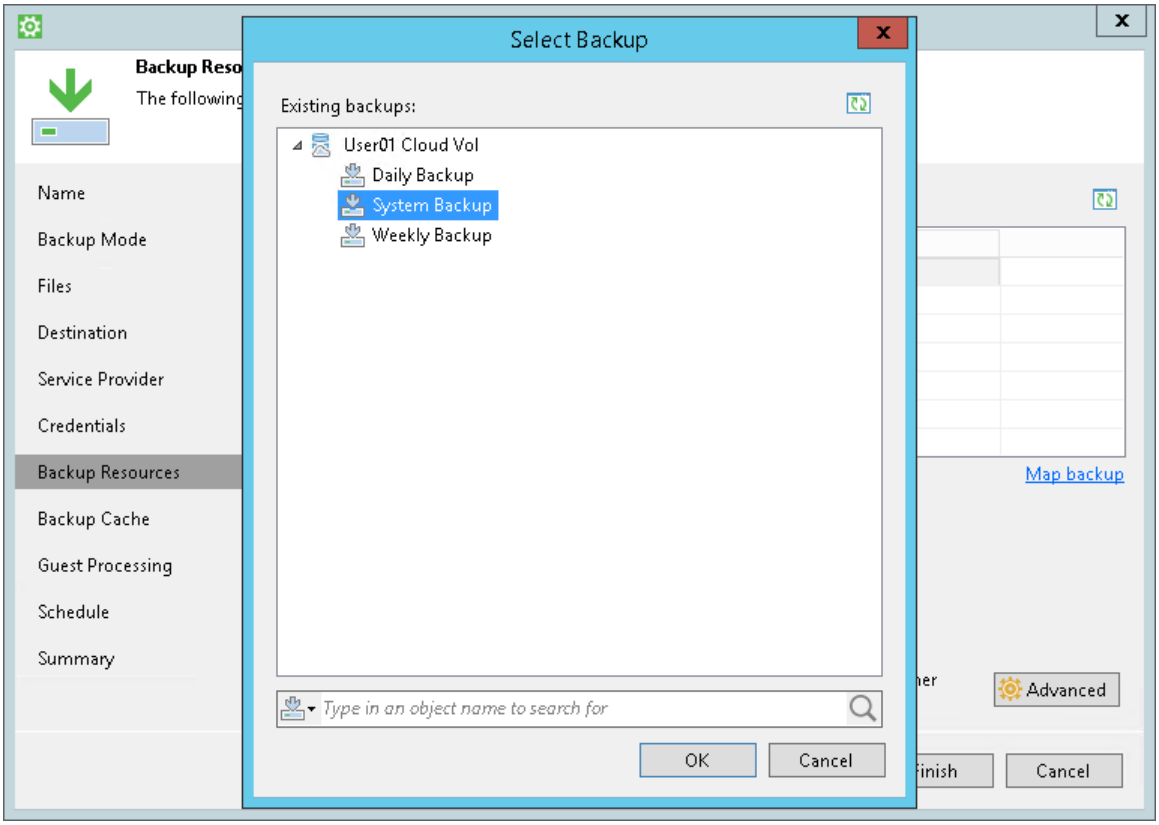
5. Click **Advanced** to specify advanced settings for the backup job. To learn more, see Specify Advanced Backup Settings.



Mapping Backup Job

If you have selected to map the job to the already created backup, perform the following steps:

1. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window. Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same tenant or subtenant account used to connect to the repository.



2. If you map the job to an encrypted backup, and encryption keys in the cloud repository database do not coincide with the encryption keys in the Veeam Agent database, you must provide the password specified for encryption after the job start. In the dialog window, enter the password and click **OK**. To learn more, see Resuming Encrypted Backup Chain.

The backup job mapping can also be helpful in case of the backup seeding. If you want to seed backup files, you must place these files in a specific folder on the target location. On the cloud repository, Veeam Agent stores backups in the folders with paths of the following format:

- If a tenant account is specified in the **Username** field at the **Credentials** step of the wizard:

  `<path_to_repository>\<tenant_name>\<backup_job_name>`

- If a subtentant account is specified in the **Username** field at the **Credentials** step of the wizard:

  `<path_to_repository>\<tenant_name>\Users\<subtenant_name>\<backup_job_name>`

where:

- `<path_to_repository>` — path to the folder on the cloud repository.
- `<tenant_name>` — name of the tenant account.
- `<subtenant_name>` — name of the subtenant account.
- `<backup_job_name>` — name of the backup job.

For example:

- `C:\Backup\TechCompany\System Backup`

  or

- `C:\Backup\TechCompany\Users\User01\System Backup`

To learn more about backup seeding scenarios for cloud repositories, see this Veeam KB article.

Microsoft OneDrive Settings

The **Microsoft OneDrive** step of the wizard is available if you have chosen to save the backup in the Microsoft OneDrive cloud storage.

Specify settings to connect to Microsoft OneDrive:

1. Click **Click to sign in to Microsoft OneDrive**.
2. In the **Microsoft OneDrive** web browser window, follow instructions to specify credentials of the Microsoft OneDrive account and click **Sign in**.

> **NOTE**
>
> You need to sign in to Microsoft OneDrive in the backup job wizard once in every 90 days. If no Veeam Agent backup job sessions are performed for 14 days within a 90-day period, you also need to sign in to Microsoft OneDrive in the backup job wizard after the 14-day period expires. To learn more, see Authorization in Microsoft OneDrive.

3. Specify backup retention policy settings:
   - [For Free and Workstation product editions] In the **Keep restore points for the last <N> days when computer was used** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent for Microsoft Windows keeps backup files for 7 days. After this period is over, Veeam Agent for Microsoft Windows will remove the earliest restore points from the backup chain.

     To learn more, see Backup Retention Policy in Free and Workstation Editions.
   - [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.
     - If the **restore points** option is selected and the number of restore points is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.
     - If the **days** option is selected and the specified period is over, Veeam Agent will remove the earliest restore points from the backup chain.

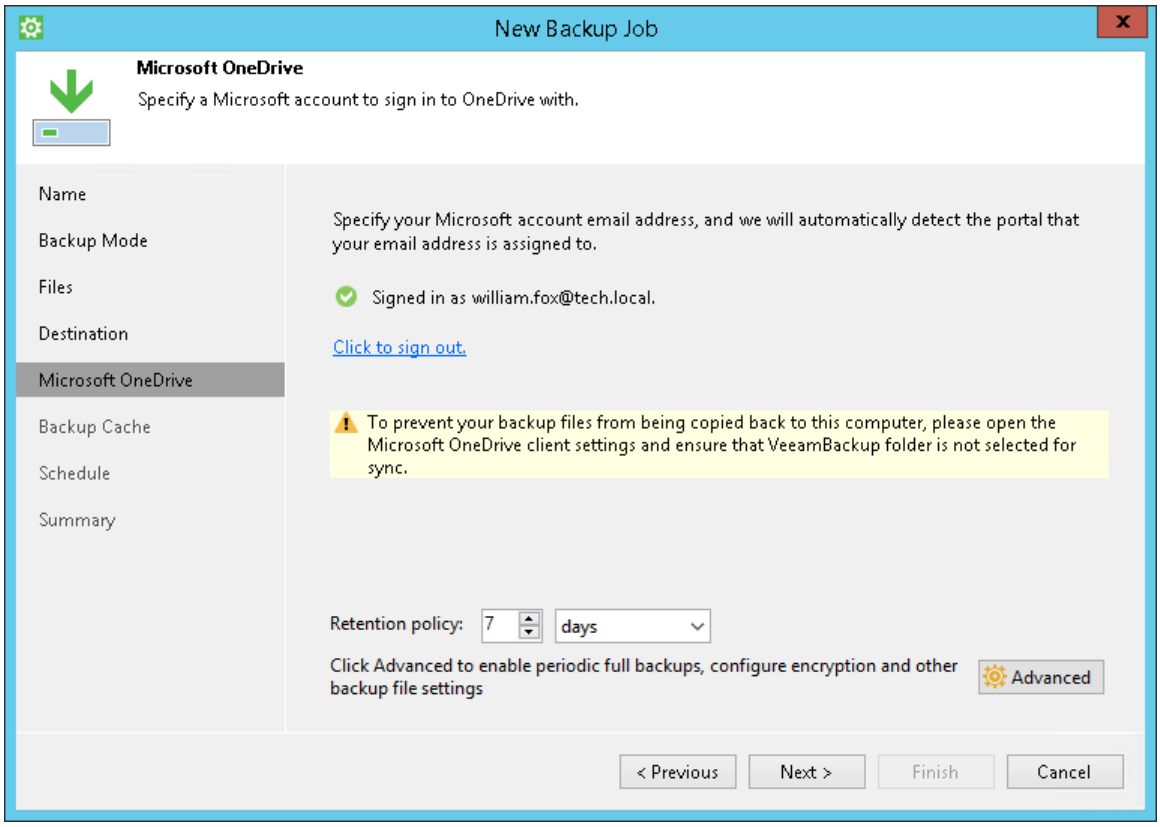
     To learn more, see Backup Retention Policy in Server Edition.

4. Click **Advanced** to specify advanced settings for the backup job. To learn more, see Specify Advanced Backup Settings.

TIP

Consider the following:

- If you want to change settings to connect to Microsoft OneDrive, click the **Click to sign out** link and repeat steps 1–2 to specify another account.
- If you use the Microsoft OneDrive client to sync your OneDrive storage and a local folder on your Veeam Agent computer, you should disable synchronization for the *VeeamBackup* folder in the client settings. Otherwise, Veeam Agent backups created on the Microsoft OneDrive storage will also appear on a computer whose data you back up consuming space on a local disk of this computer.



Step 7. Specify GFS Retention Policy

This step of the wizard is available if the following conditions are met:

- You use the Workstation or Server edition of Veeam Agent for Microsoft Windows.
- You have chosen to save backup files on a local storage, in a network shared folder, on a Veeam backup repository, or on a Veeam Cloud Connect repository.

To configure GFS retention policy settings for the backup job:

1. Select the Click **Configure** at one of the following steps of the wizard:
   - Local Storage — if you have selected the **Local storage** option at the Destination step of the wizard.
   - Shared Folder — if you have selected the **Shared folder** option at the Destination step of the wizard.
   - Backup Repository — if you have selected the **Veeam backup repository** option at the Destination step of the wizard.
   - Backup Resources — if you have selected the **Veeam Cloud Connect repository** option at the Destination step of the wizard.
2. In the **Configure GFS** window, do the following:
   a. If you want to create weekly restore points for archival purposes, 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.

      In the **If multiple full backups exist, use the one from** list, select a week day when Veeam Agent must assign the weekly GFS flag to a full restore point.

   b. If you want to create monthly restore points for archival purposes, 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.

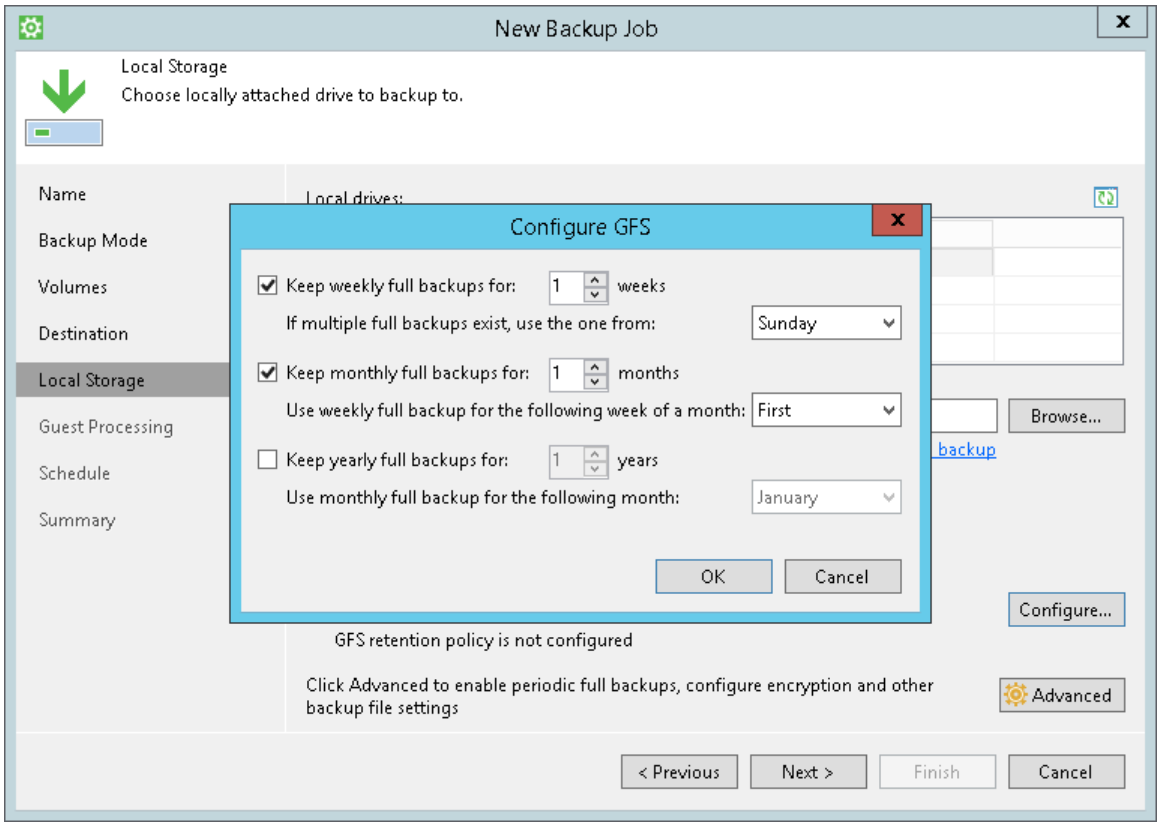
      In the **Use weekly full backup for the following week of a month** list, select a week when Veeam Agent must assign the monthly GFS flag to a full restore point. A week equals 7 calendar days; for example, the first week of May is days 1–7, and the last week of May is days 25–31.

   c. If you want to create yearly restore points for archival purposes, 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.

      In the **Use monthly full backup for the following month** list, select a month when Veeam Agent 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 GFS Flags section in the Veeam Backup & Replication User Guide.



Related Topic

Long-Term Retention Policy (GFS).

Step 8. Specify Advanced Backup Settings

In the **Advanced Settings** window, specify advanced settings for the backup job:

- Backup settings
- Maintenance settings
- Storage settings

Backup Settings

To specify settings for a backup chain created with the backup job:

1. Click **Advanced** at one of the following steps of the wizard:
   - Local Storage — if you have selected the **Local storage** option at the Destination step of the wizard.
   - Shared Folder — if you have selected the **Shared folder** option at the Destination step of the wizard.
   - Backup Repository — if you have selected the **Veeam backup repository** option at the Destination step of the wizard.
   - Backup Resources — if you have selected the **Veeam Cloud Connect repository** option at the Destination step of the wizard.
   - Microsoft OneDrive — if you have selected the **Microsoft OneDrive** option at the Destination step of the wizard.

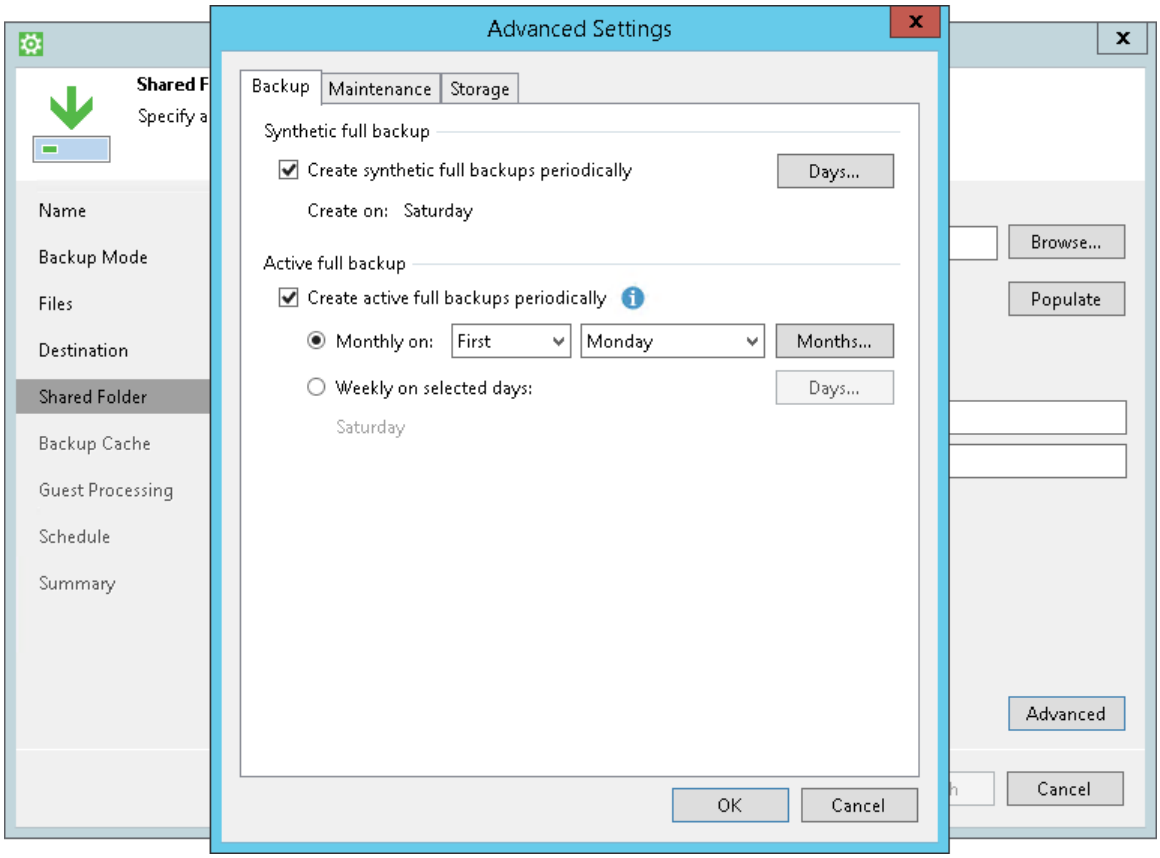
2. If you want to periodically create synthetic full backups, on the **Backup** tab, select the **Create synthetic full backups periodically** check box and click **Days** to schedule synthetic full backups on the necessary week days.
3. If you want to periodically create active full backups, select the **Create active full backups periodically** check box. Use the **Monthly on** or **Weekly on selected days** options to define scheduling settings.

Keep in mind that to create active and synthetic full backups on selected days, Veeam Agent must run the backup job on these days. For details about backup job schedule, see Scheduling Settings.

NOTE

Consider the following:

- Synthetic full backup functionality is available only in Workstation and Server editions of Veeam Agent for Microsoft Windows.
- Before scheduling periodic full backups, you must make sure that you have scheduled the backup job to run on the enough free space on the target location. As an alternative, you can create active full backups manually when needed. For more information, see Creating Active Full Backups.
- If you schedule the active full backup and synthetic full backup on the same day, Veeam Agent for Microsoft Windows will perform only active full backup. Synthetic full backup will be skipped.



Maintenance Settings

To specify maintenance settings for the backup chain created with the backup job:

1. Click **Advanced** at one of the following steps of the wizard:
   - Local Storage — if you have selected the **Local storage** option at the Destination step of the wizard.
   - Shared Folder — if you have selected the **Shared folder** option at the Destination step of the wizard.
   - Backup Repository — if you have selected the **Veeam backup repository** option at the Destination step of the wizard.
   - Backup Resources — if you have selected the **Veeam Cloud Connect repository** option at the Destination step of the wizard.
   - Microsoft OneDrive — if you have selected the **Microsoft OneDrive** option at the Destination step of the wizard.
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 schedule for the health check.

   An automatic health check can help you avoid a situation when a restore point gets corrupted, making all dependent restore points corrupted, too. If during the health check Veeam Agent for Microsoft Windows detects corrupted data blocks in the latest restore point in the backup chain (or 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 Veeam Agent computer to the target location. 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.

   > **NOTE**
   >
   > The **Perform backup files health check** option is not available if you have selected the **Microsoft OneDrive** option at the Destination step of the wizard.

4. [For Veeam backup repository and cloud repository targets] Select the **Remove deleted items data after** check box and specify the number of days for which you want to keep the backup created with the backup job in the target location. If Veeam Agent for Microsoft Windows does not create new restore points for the backup, the backup will remain in the target location for the period that you have specified. When this period is over, the backup will be removed from the target location.

   The retention period for outdated backups is useful if the job is configured to create active full backups or synthetic full backups with subsequent transformation, and you want to make sure that the full backup does not include redundant data.

   By default, the retention period for outdated backups is 30 days. Do not set this retention period to 1 day or a similar short interval. In the opposite case, the backup job may work not as expected and remove data that you still require. For more information, see Retention Policy for Outdated Backups.

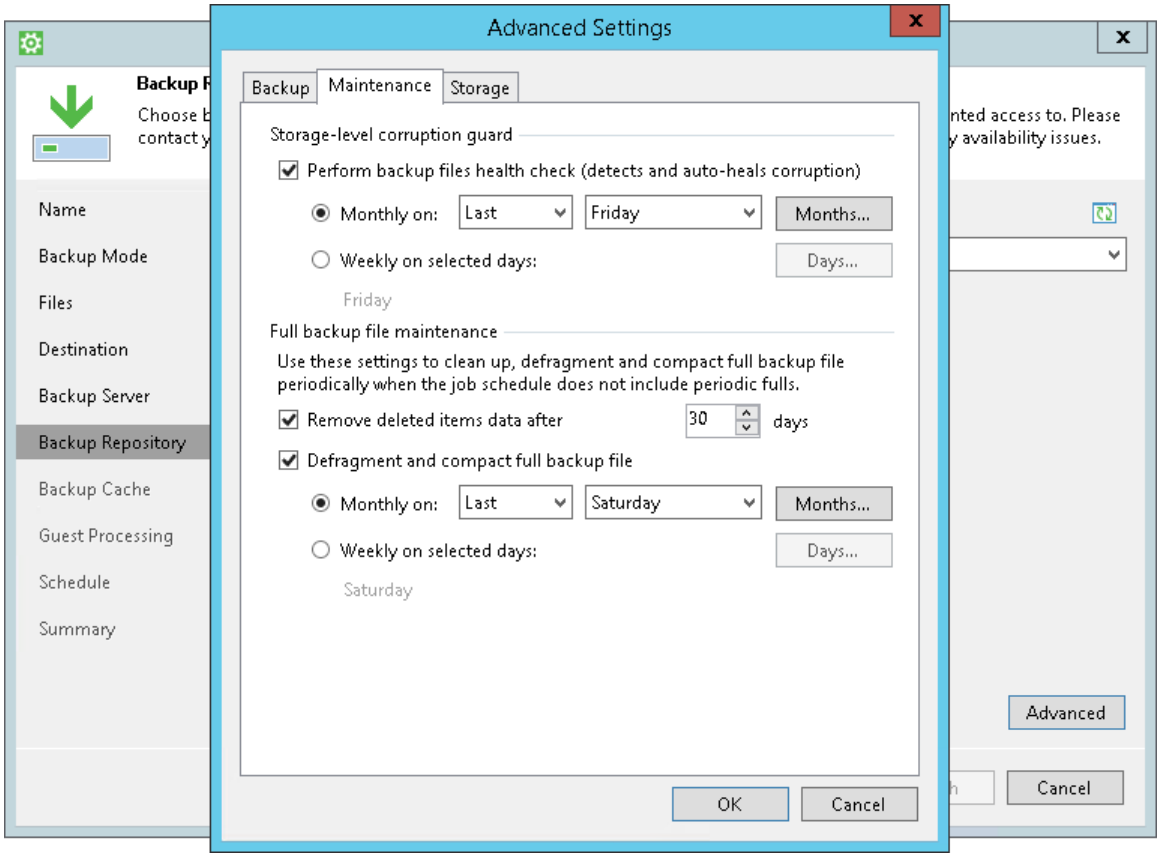
   > **NOTE**
   >
   > The **Remove deleted items data after** option is available only if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the Destination step of the wizard.

4. 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 Agent for Microsoft Windows creates a new empty file and copies to this file data blocks from the full backup file. As a result, the full backup file gets defragmented, and the speed of reading from and writing to the backup file increases.

   If the full backup file contains data blocks for deleted drives, Veeam Agent for Microsoft Windows will remove these data blocks. For more information, see Compact of Full Backup File.

   > **NOTE**
   >
   > The **Defragment and compact full backup file** option is not available if you have selected the **Microsoft OneDrive** option at the Destination step of the wizard.



Storage Settings

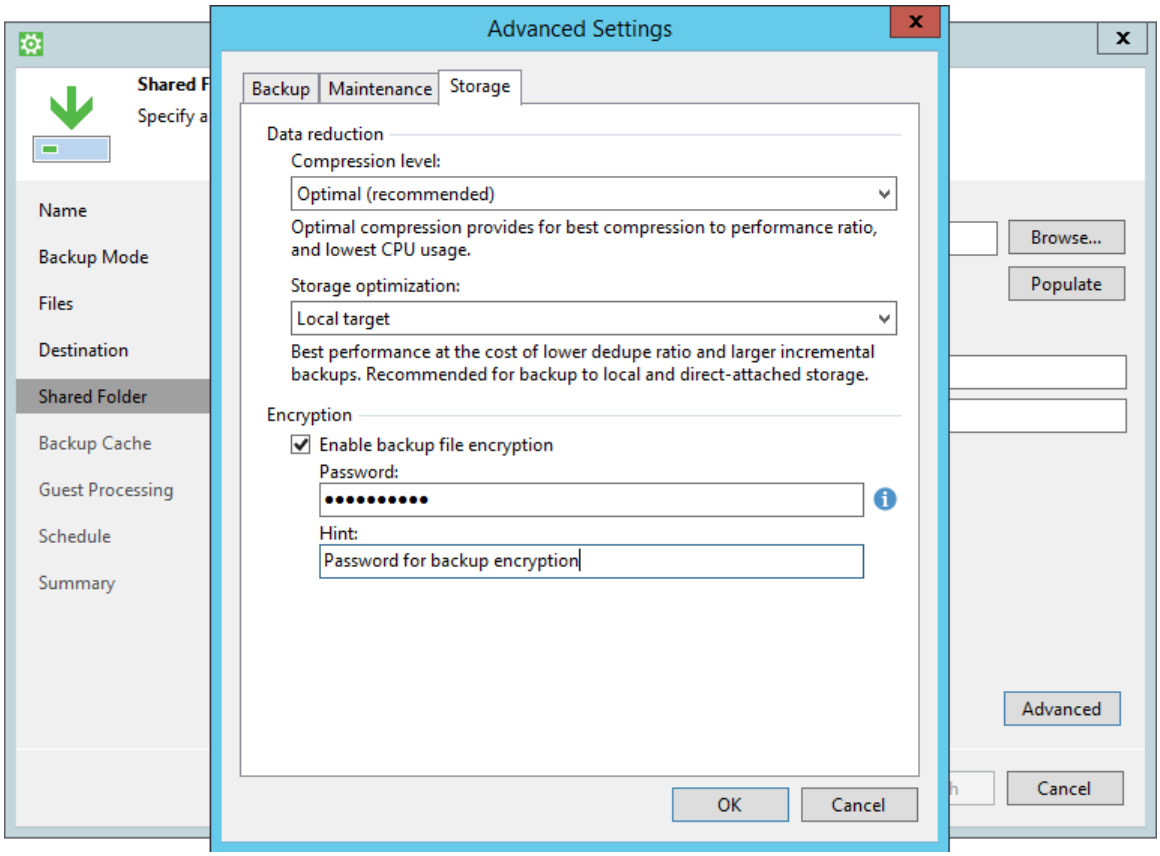
To specify storage settings for the backup job:

1. Click **Advanced** at one of the following steps of the wizard:
   - Local Storage — if you have selected the **Local storage** option at the Destination step of the wizard.
   - Shared Folder — if you have selected the **Shared folder** option at the Destination step of the wizard.
   - Backup Repository — if you have selected the **Veeam backup repository** option at the Destination step of the wizard.
   - Backup Resources — if you have selected the **Veeam Cloud Connect repository** option at the Destination step of the wizard.
   - Microsoft OneDrive — if you have selected the **Microsoft OneDrive** option at the Destination step of the wizard.
2. Click the **Storage** tab.
3. From the **Compression level** list, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.
4. In the **Storage optimization** section, select what type of backup target you plan to use: *Local target (large blocks)*, *Local target*, *LAN target* or *WAN target*. Depending on the chosen storage type, Veeam Agent for Microsoft Windows will use data blocks of different size to optimize the size of backup files and job performance.
5. If you want to encrypt the content of backup files, in the **Encryption** section, specify encryption settings for the backup job:
   a. Select the **Enable backup file encryption** check box.
   b. In the **Password** field, type a password that you want to use for encryption.
   c. In the **Hint** field, type a hint for the password. In case you lose the password, the specified hint will help you to remember the lost password.

NOTE

Consider the following:

- You cannot specify encryption options for the backup job if you have chosen to save backup files on a Veeam backup repository. Encryption of Veeam Agent backups stored on the backup repository are managed per repository by a backup administrator working with Veeam Backup & Replication. To learn more, see the Data Encryption and Access Permissions sections in the Veeam Backup & Replication User Guide.
- If you lose a password that was specified for encryption, you can change the password in the encryption settings. After the backup job creates a new restore point encrypted with the new password, you will be able to use this password to restore data form all restore points in the backup chain, including those restore points that were encrypted with an old password.
- If you enable encryption for the existing backup job that has already created one or more restore points, during the next job session, Veeam Agent for Microsoft Windows will create active full backup. 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 the existing backup job, Veeam Agent for Microsoft Windows does not encrypt the previous backup chain created with this job.



Step 9. Specify Backup Cache Settings

The **Backup Cache** step of the wizard is available if you have chosen to save backup files on a remote storage: in a network shared folder, on a Veeam backup repository, on a Veeam Cloud Connect repository or in Microsoft OneDrive.

Specify backup cache settings:

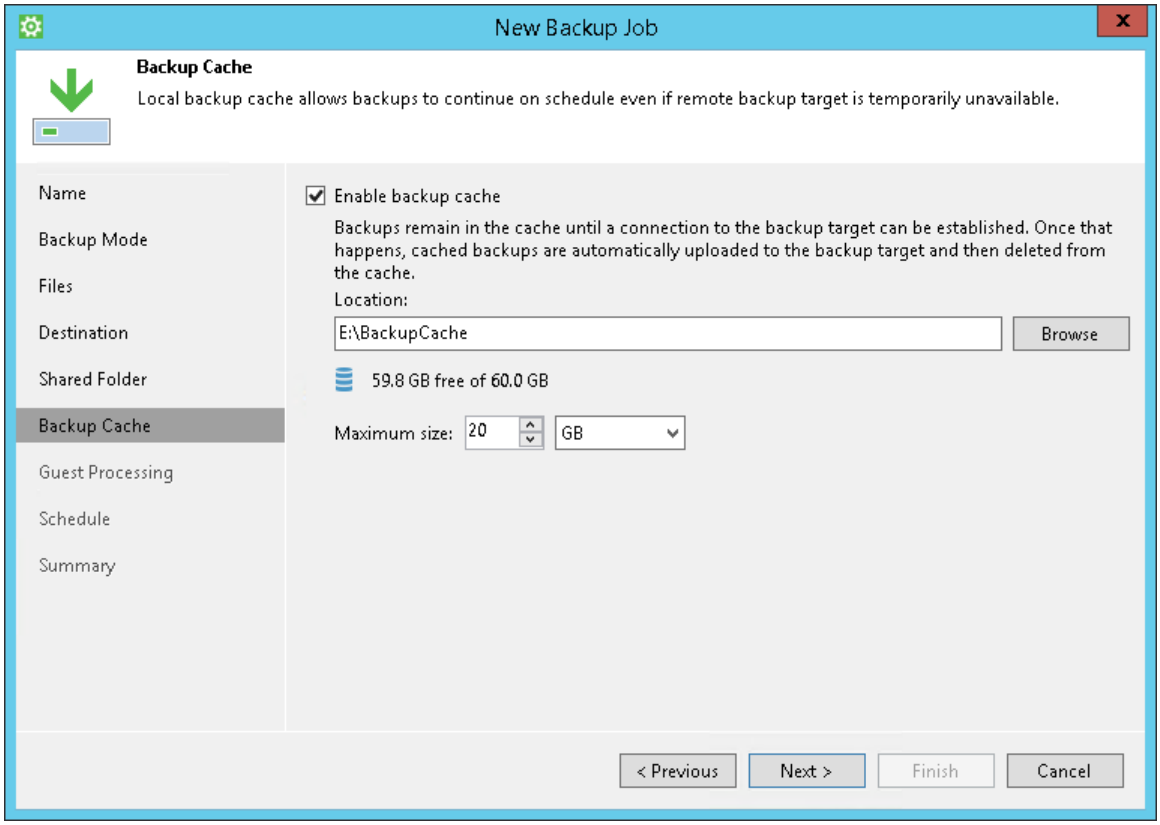
1. Select the **Enable backup cache** check box.
2. In the **Location** field, specify a path to the folder on your computer in which backup files must be stored.
3. In the **Maximum size** field, specify the size for the backup cache.

   When defining the size of the backup cache, assume the following:

   - Each full backup file may consume about 50% of the backed-up data size.
   - Each incremental backup file may consume about 10% of the backed-up data size.

TIP

For the backup cache, you can use a dedicated removable storage device, for example, a USB key or an SD card. In this case, the backup cache will not consume disk space on the local drive of the Veeam Agent computer.

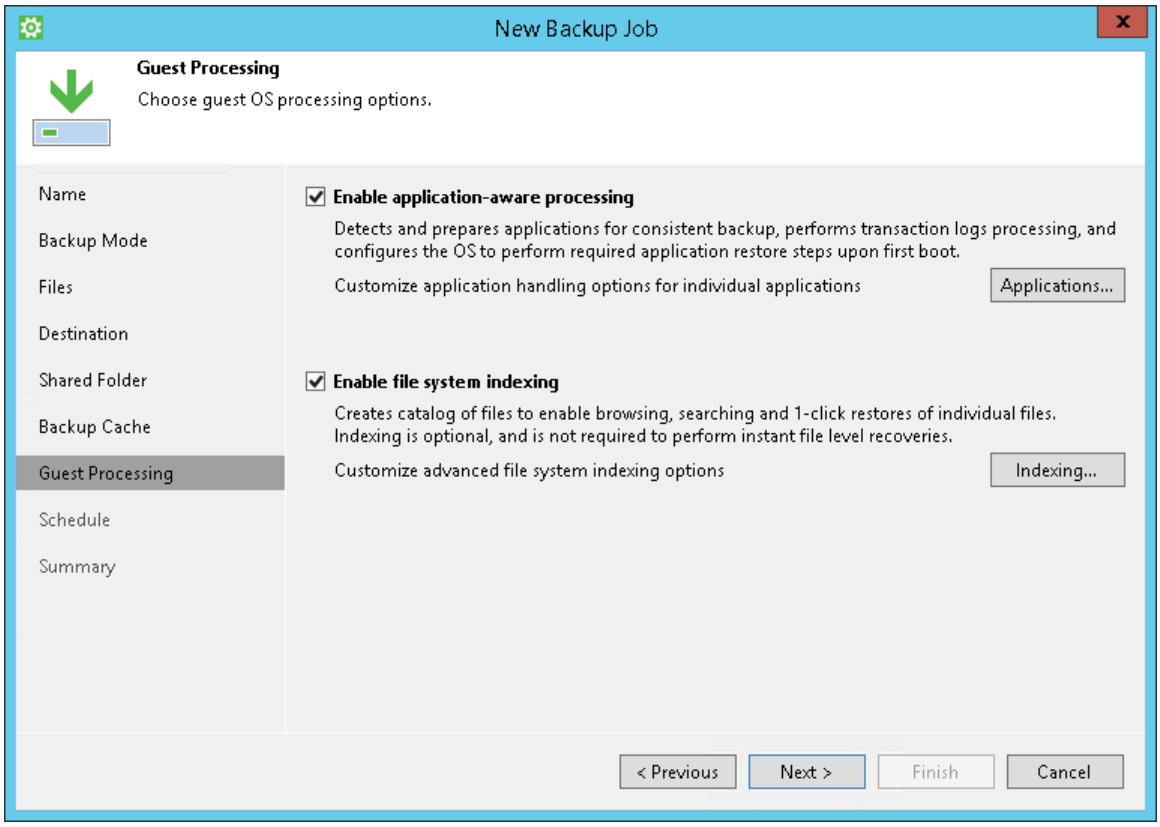


Step 10. Specify Guest Processing Settings

The **Guest Processing** step of the wizard is available in the Server edition of Veeam Agent for Microsoft Windows.

You can enable the following settings for guest OS processing:

- Application-aware processing
- Transaction log handling for Microsoft SQL Server
- Archived log handling for Oracle databases
- SharePoint account settings
- Use of pre-freeze and post-thaw scripts
- File indexing



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 proper recovery of applications without data loss.

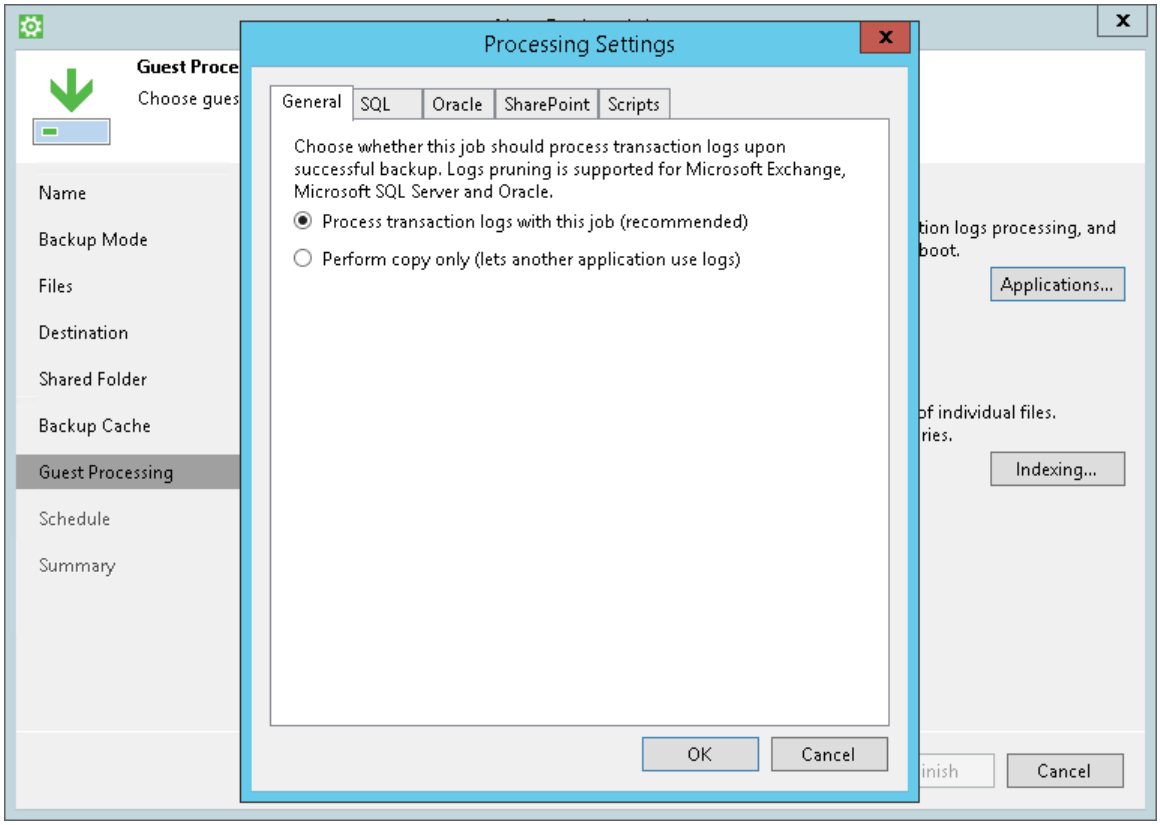
To enable application-aware processing:

1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.
3. In the **Processing Settings** window, on the **General** tab, specify if Veeam Agent must process transaction logs [For Microsoft Exchange, Microsoft SQL and Oracle] or copy-only backups must be created.
   a. Select **Process transaction logs with this job** if you want Veeam Agent to process transaction logs.
      - [For Microsoft Exchange] With this option selected, Veeam Agent will wait for backup to complete successfully and then trigger truncation of transaction logs. If the backup job fails, the logs will remain untouched until the next backup job session.
      - [For Microsoft SQL Server and Oracle] You will have to specify settings for transaction log handling on the **SQL** and **Oracle** tabs of the **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 tool to maintain consistency of the database state. Veeam Agent will create a copy-only backup. The copy only backup preserves the chain of full/differential backup files and transaction logs. After a copy-only backup, Veeam Agent does not trigger truncation of transaction logs. For more information, see this Microsoft article.

IMPORTANT

Consider the following:

- [For Microsoft Exchange] Veeam Agent for Microsoft Windows performs truncation of Microsoft Exchange transaction logs only if all disks that contain Microsoft Exchange databases and logs are included in a volume-level backup job.
- [For Microsoft SQL Server and Oracle] If both Microsoft SQL Server and Oracle Server are installed on one guest OS, you can enable log backup settings for one application only: either Microsoft SQL Server or Oracle.



Microsoft SQL Server Transaction Log Settings

If you back up Microsoft SQL Server, you can specify how Veeam Agent for Microsoft Windows must process database transaction logs:

1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.
3. In the **Processing Settings** window, on the **General** tab, select **Process transaction logs with this job**.
4. In the **Processing Settings** window, click the **SQL** tab.
5. To specify a user account that Veeam Agent for Microsoft Windows will use to connect to the Microsoft SQL Server, select the **Specify Windows account with sysadmin role on SQL Server** check box and enter a user name and password for the Microsoft Windows user account.

   The Microsoft Windows account used to connect to the Microsoft SQL Server must be assigned the following roles:

   - *Administrator* role on the target computer with Veeam Agent.
   - *Sysadmin* role on the target Microsoft SQL Server.

   To provide minimal permissions, the account must be assigned the following roles and permissions:

   - SQL Server instance-level roles: *dbcreator*, *public*.
   - Database-level roles: *db_backupoperator*, *db_denydatareader*, *public*; for system databases (master, model, msdb) — *db_backupoperator*, *db_datareader*, *public*; for system database (msdb) — *db_datawriter*.
   - Securables: *view any definition*, *view server state*.
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. Veeam Agent for Microsoft Windows will wait for the backup to complete successfully and then truncate transaction logs. If the backup job fails, the logs will remain untouched until the next backup job session.
   - Select **Do not truncate logs** to preserve transaction logs. When the backup job completes, Veeam Agent for Microsoft Windows will not truncate transaction logs.

     We recommend 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 may grow large and consume all disk space. In this case, the database administrator must take care of transaction logs him-/herself.
   - Select **Backup logs periodically** to back up transaction logs with Veeam Agent for Microsoft Windows. Veeam Agent for Microsoft Windows will periodically copy transaction logs to the backup location and store them together with the image-level backup. During the backup job session, transaction logs will be truncated.

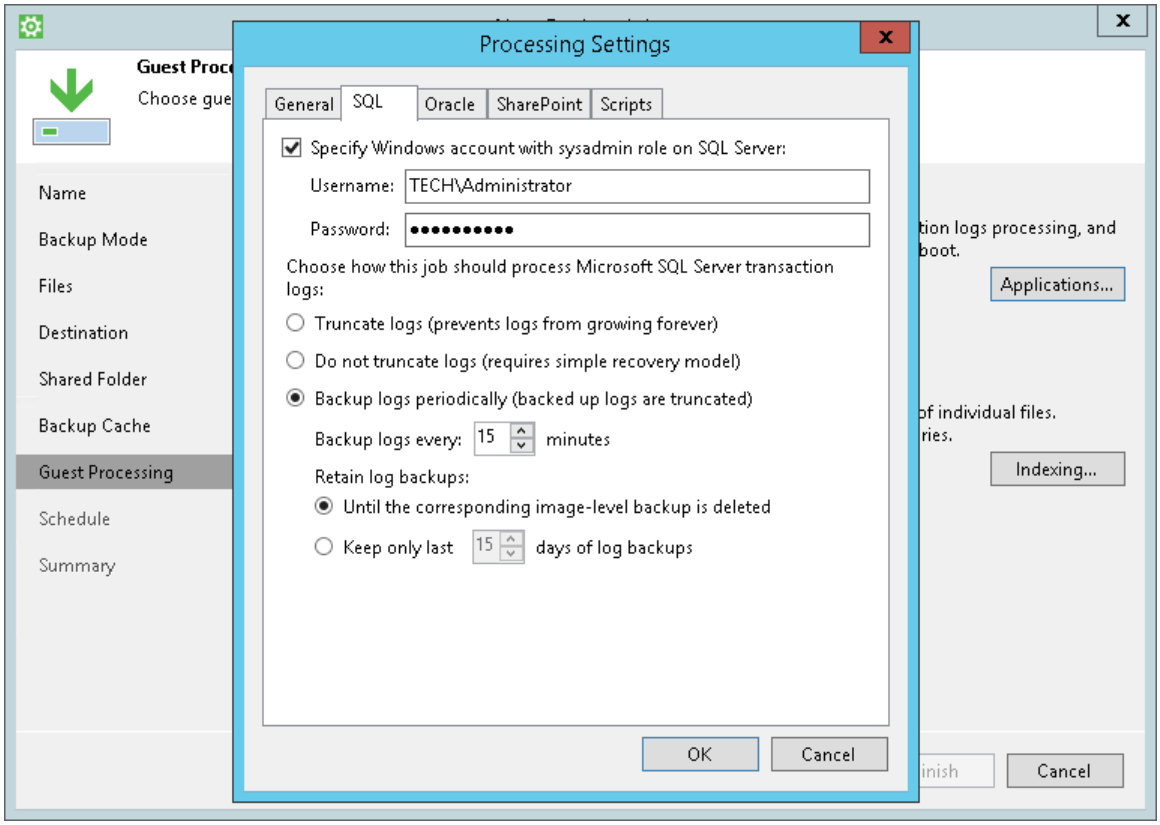
     For more information, see Microsoft SQL Server and Oracle Logs Backup and Restore.

> **NOTE**
>
> You can enable transaction log backup in one backup job only. If you have already configured a job that is set up to back up transaction logs, you will not be able to select the **Backup logs periodically** option in the properties of another backup job.

If you have selected to back up transaction logs with Veeam Agent for Microsoft Windows, 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 location.
   - 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 backup. For more information, see Retention for Transaction Log Backups.

> **IMPORTANT**
>
> Veeam Agent for Microsoft Windows automatically excludes its configuration database from application-aware processing during backup. Transaction logs for the configuration database are not backed up.



Oracle Archived Log Settings

If you back up an Oracle database, you can specify how Veeam Agent for Microsoft Windows must process archived logs:

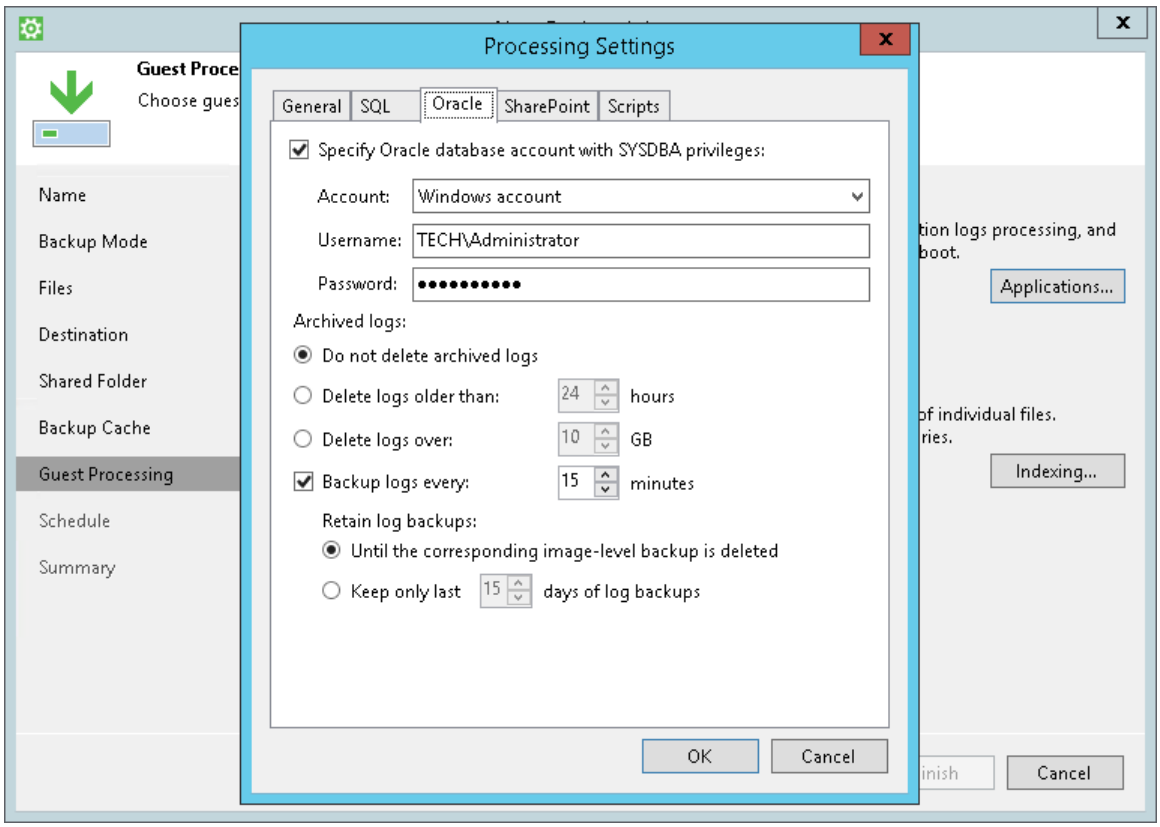
1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.
3. In the **Processing Settings** window, on the **General** tab, select **Process transaction logs with this job**.
4. In the **Processing Settings** window, click the **Oracle** tab.
5. Specify a user account that Veeam Agent for Microsoft Windows will use to connect to the Oracle database:
   a. Select the **Specify Oracle database account with SYSDBA privileges** check box
   b. In the **Account** field, select what type of user account you plan to use: *Oracle account* or *Windows account*.
   c. In the **Username** and **Password** fields, enter a username and password for the account.

   The specified account must have SYSDBA rights on the Oracle database.
6. In the **Archived logs** section, specify if Veeam Agent for Microsoft Windows must delete archived logs on the Oracle database:
   - Select **Do not delete archived logs** if you want Veeam Agent for Microsoft Windows to preserve archived logs. When the backup job completes, Veeam Agent for Microsoft Windows will not delete archived logs.

     We recommend that you select this option for databases for which the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space. In this case, the database administrator must take care of archived logs him-/herself.
   - Select **Delete logs older than <N> hours** or **Delete logs over <N> GB** if you want Veeam Agent for Microsoft Windows to delete archived logs that are older than <N> hours or larger than <N> GB. Veeam Agent for Microsoft Windows will wait for the backup to complete successfully and then trigger archived logs truncation via Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next backup job session.
7. To back up Oracle archived logs with Veeam Agent for Microsoft Windows, 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.

   > **NOTE**
   >
   > You can enable archived log backup in one backup job only. If you have already configured a job that is set up to back up archived logs, you will not be able to select the **Backup logs every <N> minutes** option in the properties of another backup job.

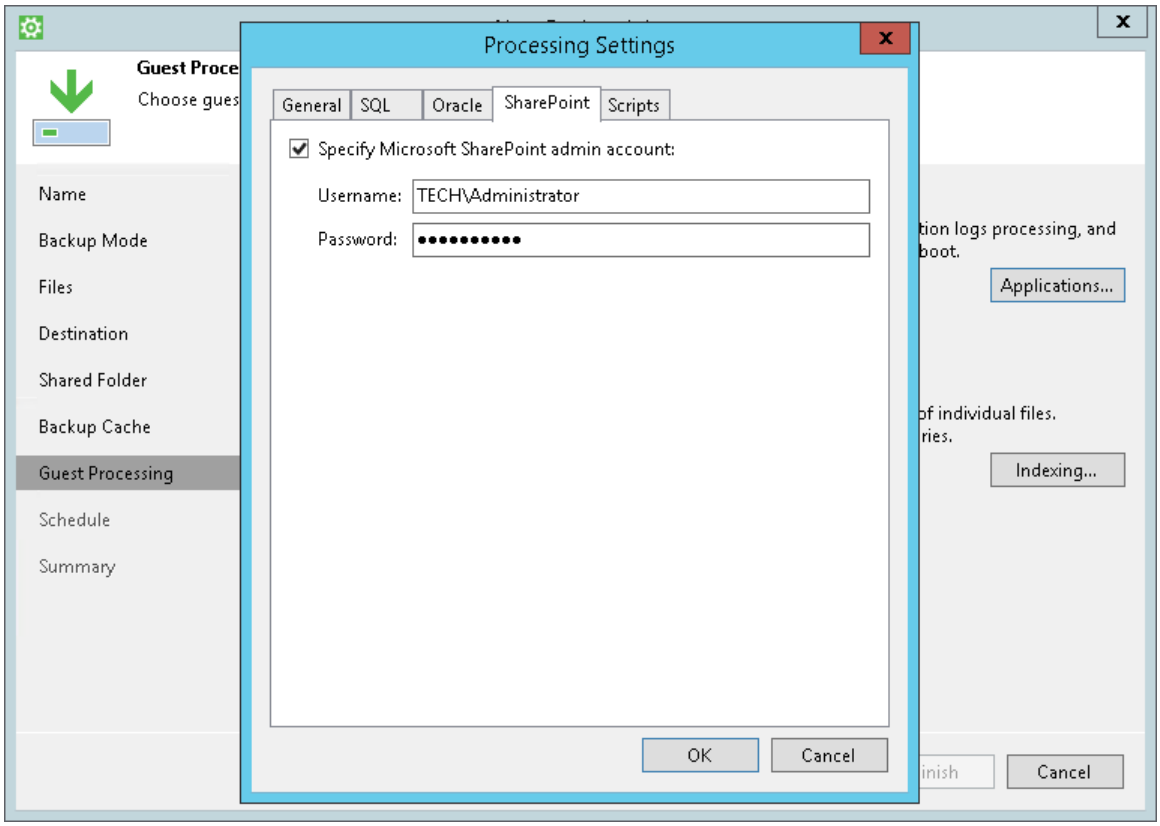
8. In the **Retain log backups** section, specify retention policy for archived logs stored in the backup location:
   - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for Veeam Agent backups and archived log backups.
   - Select **Keep only last <n> days t**o 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 Veeam Agent backups. For more information, see Retention for Archived Log Backups.



Microsoft SharePoint Account Settings

If you back up Microsoft SharePoint, you must specify a user account that has enough permissions on the application:

1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.
3. In the **Processing Settings** window, click the **SharePoint** tab.
4. Select the **Specify Microsoft SharePoint admin account** check box and enter a user name and password for the user account.



Pre-Freeze and Post-Thaw Scripts

If you plan to back up data of applications that do not support VSS, you can specify what scripts Veeam Agent for Microsoft Windows must use to quiesce the OS on your computer. The pre-freeze script quiesces the file system and application data to bring the OS to a consistent state before Veeam Agent for Microsoft Windows creates a VSS snapshot. After the VSS snapshot is created, the post-thaw script brings the file system and applications to their initial state.

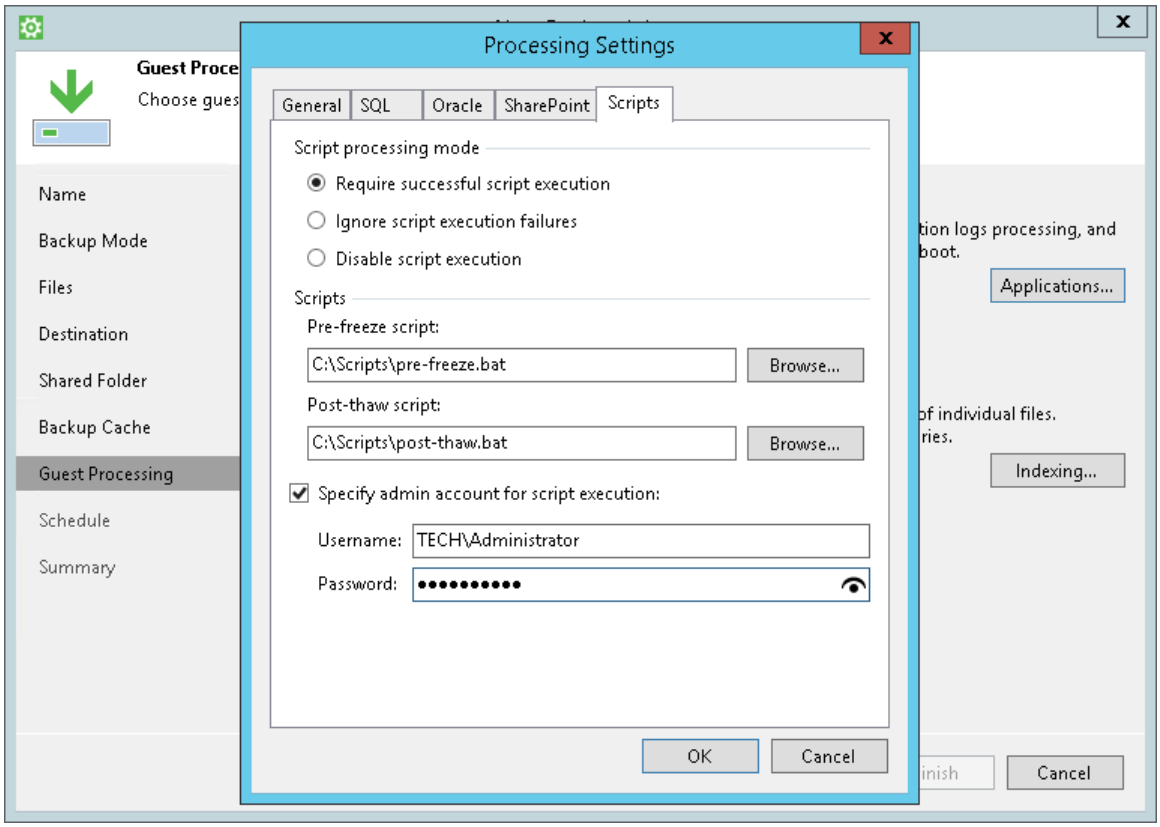
To specify pre-freeze and post-thaw scripts for the job:

1. At the **Guest Processing** step, click **Applications**.
2. In the **Processing Settings** window, click the **Scripts** tab.
3. In the **Script processing mode** section, specify the scenario for scripts execution:
   - Select **Require successful script execution** if you want Veeam Agent for Microsoft Windows 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.
4. In the **Scripts** section, specify paths to pre-freeze and post-thaw scripts. Scripts must reside on a local drive of the Veeam Agent computer.

   Veeam Agent for Microsoft Windows supports the following types of scripts:
   - Program files in the EXE, BAT and CMD format
   - Windows script files in the JS, VBS and WSF format
   - PowerShell script files in the PS1 format

   You can use scripts of other formats as well, but we cannot guarantee correct processing of such scripts.

5. By default, Veeam Agent for Microsoft Windows performs guest processing activities under the Local System account. To specify a user account that Veeam Agent for Microsoft Windows will use to run pre-freeze and post-thaw scripts, select the **Specify admin account for script execution** check box and enter a user name and password for the user account.

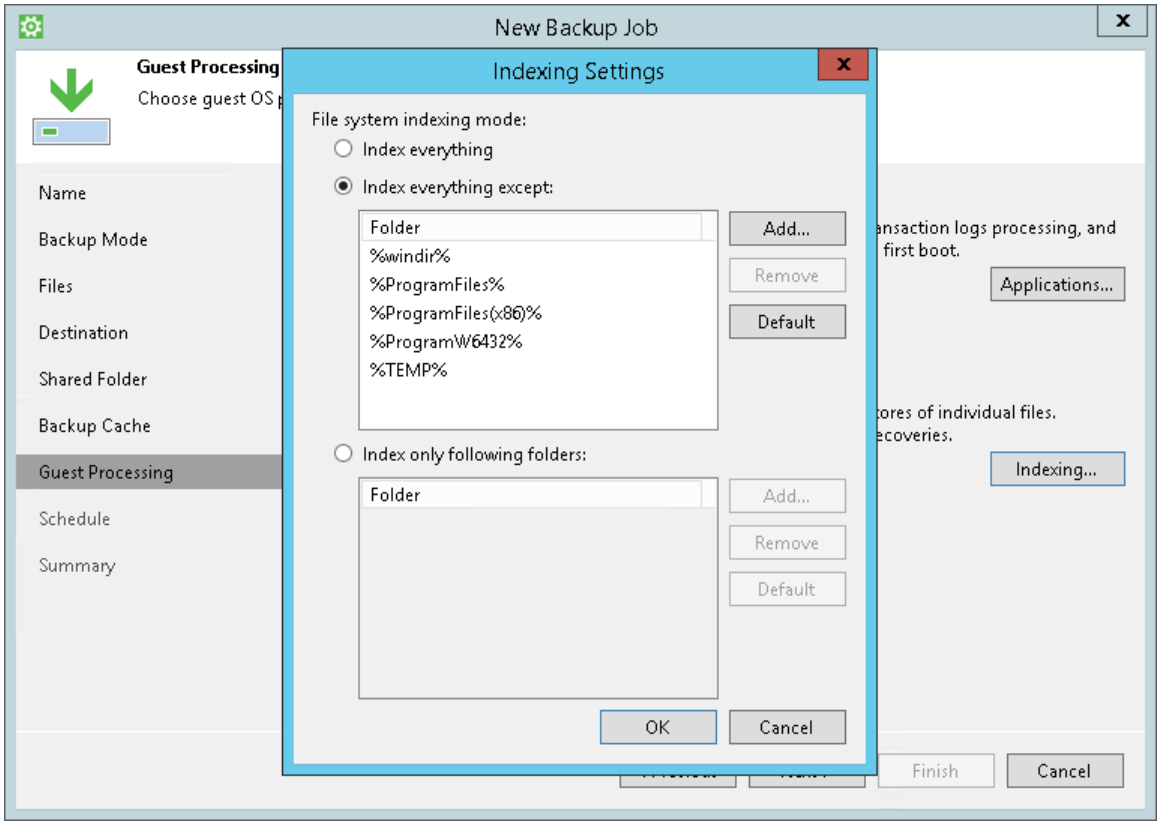


File Indexing

To specify file indexing options:

1. At the **Guest Processing** step of the wizard, click **Indexing**.
2. In the **Indexing Settings** window, specify the indexing scope:
   - Select **Index everything** if you want to index all files within the backup scope that you have specified at the Backup mode step of the wizard. Veeam Agent for Microsoft Windows will index all files that reside:
     - on your computer OS (for entire computer backup)
     - on the volumes that you have selected for backup (for volume-level backup)
     - in the directories that you have selected for backup (for file-level backup)
   - Select **Index everything except** if you want to index all files on your computer OS 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%*.



Step 11. Specify Backup Schedule

At the **Schedule** step of the wizard, specify the schedule according to which you want to perform backup. Backup job scheduling options differ depending on the edition of Veeam Agent for Microsoft Windows:

- Scheduling Settings in Free and Workstation Editions
- Scheduling Settings in Server Edition

Scheduling Settings in Free and Workstation Editions

At the **Schedule** step of the wizard, set up Veeam Agent for Microsoft Windows to run the backup job periodically at specific time or after specific computer events:

- Select the time to launch the backup job.
- Select events that trigger the backup job launch.

> **IMPORTANT**
>
> If the power scheme on your computer does not allow using wake up timers, Veeam Agent for Microsoft Windows will ask you to change the power scheme settings. Click **Yes** to allow Veeam Agent for Microsoft Windows to wake your computer from sleep for backup.
>
> You can manually change the power scheme settings on your computer. To do this, navigate to **Control Panel** > **All Control Panel Items** > **Power Options** > **Edit Plan Settings**.

Selecting Time to Launch Backup Job

In the **Periodically** section, specify time to launch the backup job:

1. Select the **Daily at** check box and use the fields on the right to specify time and days when the backup job must start:
   - *Everyday* — select this option to start the job at specific time daily.
   - *On week-days* — select this option to start the job at specific time on week-days.
   - *On these days* — select this option to start the job at specific time on selected days.

   You can leave the **Daily at** check box unchecked to configure the backup job without daily schedule. In this case, you will be able to use the configured backup job to perform backup automatically at specific events. You can also use the configured backup job to create ad-hoc incremental and standalone full backups. To learn more, see Performing Ad-Hoc Backups.
2. If you have selected the *On these days* option, click the **Days** button and clear check boxes for the days when the job must not start.
3. Select the action that Veeam Agent for Microsoft Windows must perform in case your computer is powered off at the time when the scheduled backup job must start.
   - *Backup once powered on* — select this option if you want Veeam Agent for Microsoft Windows to start the scheduled backup job when you power on the computer.
   - *Skip backup* — select this option if you want Veeam Agent for Microsoft Windows not to start the scheduled backup job when the computer is powered on. Veeam Agent for Microsoft Windows will perform backup at the next scheduled time.
4. If you want Veeam Agent for Microsoft Windows to perform a finalizing action after the backup job completes successfully, select the necessary action:
   - *Keep running* — select this option if the computer must keep on working.
   - *Sleep* — select this option if you want Veeam Agent for Microsoft Windows to bring your computer to the standby mode.
   - *Shutdown* — select this option if you want Veeam Agent for Microsoft Windows to shut down your computer.
   - *Hibernate* — select this option if you want Veeam Agent for Microsoft Windows to bring your computer to the hibernate mode. This option is available if the hibernate mode is enabled on your computer. To learn more, see this Microsoft article.

   Veeam Agent for Microsoft Windows applies this setting only to scheduled backups. If you start standalone full backup or incremental backup manually, Veeam Agent for Microsoft Windows will ignore this setting, and the computer will not be shut down or brought to the standby mode when the backup job completes.

   When the backup job completes, Veeam Agent for Microsoft Windows will prompt a dialog with a countdown to the selected post-job action. You can select to proceed to the action immediately or to cancel the action. To learn more, see Controlling Backup Post-Job Action.

Selecting Events that Trigger Backup Job Launch

In the **At the following events** section, specify computer events that trigger the backup job launch:

- Select the **Lock** check box if you want to start the scheduled backup job when the user locks the computer.

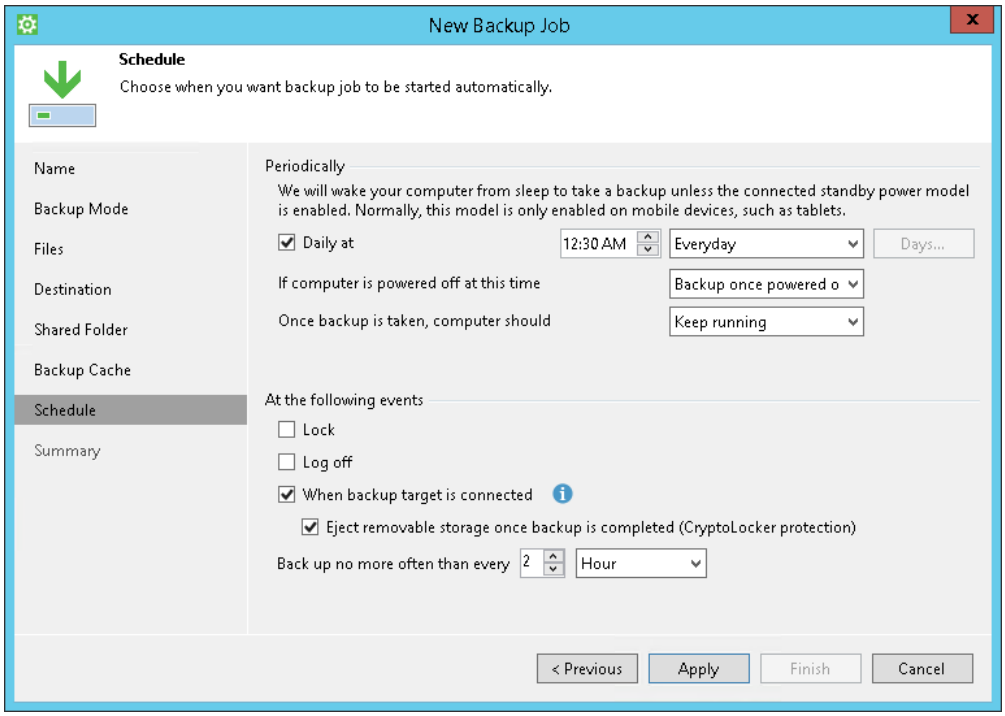
- Select the **Log off** check box if you want to start the scheduled backup job when the user working with the computer performs a logout operation.
- Select the **When backup target is connected** check box if you want to start the scheduled backup job when the backup storage becomes available (for example, when the computer connects to a local network and the target shared folder is accessible).
- Select the **Eject removable storage once backup is completed** check box if you want Veeam Agent for Microsoft Windows to unmount the storage device after the backup job completes successfully. With this option selected, backup files on the removable storage will be protected from encrypting ransomware, such as CryptoLocker.

  Veeam Agent applies this setting only to backup jobs triggered by computer events. In case of backup jobs scheduled to run at specific time, Veeam Agent will ignore this setting, and the storage device will not be unmounted after the backup job completes successfully.

  > **IMPORTANT**
  >
  > The *Eject removable storage once backup is completed* option does not guarantee a bulletproof protection against ransomware. To ensure your backups are safe, keep the OS up to date and regularly scan your backup repository for virus threats using modern antivirus software.

- Use the **Back up no more often than every <N> <time units>** field to restrict the frequency of backup job sessions. Specify a minutely, hourly or daily interval between the backup job sessions.

  The *Back up no more often than every <N> <time units>* option is applied only to job sessions started at specific events. Daily backups are performed according to defined schedule regardless of the time interval specified for this setting.



Scheduling Settings in Server Edition

At the **Schedule** step of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.

> **IMPORTANT**
>
> If the power scheme on your computer does not allow using wake up timers, Veeam Agent for Microsoft Windows will ask you to change the power scheme settings. Click **Yes** to allow Veeam Agent for Microsoft Windows to wake your computer from sleep for backup.
>
> You can manually change the power scheme settings on your computer. To do this, navigate to **Control Panel** > **All Control Panel Items** > **Power Options** > **Edit Plan Settings**.

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 backup job manually to create 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.
   - 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:

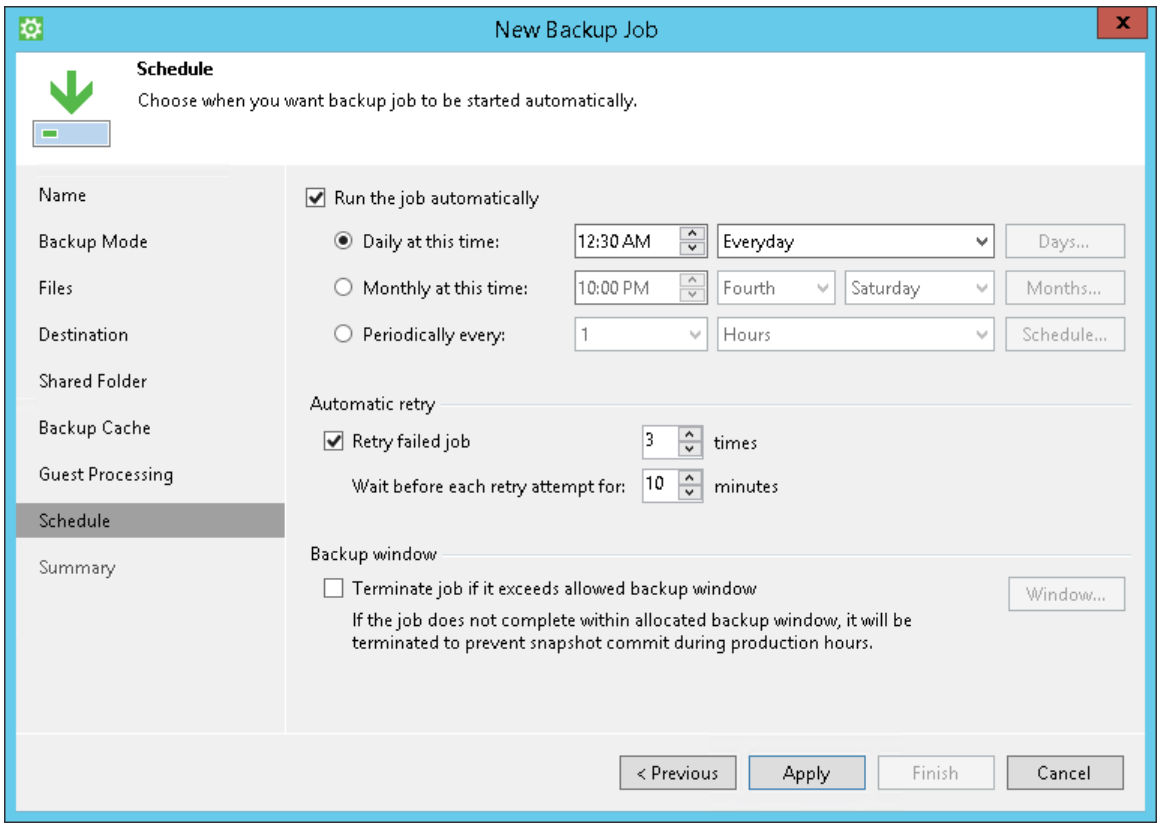
     - Veeam Agent for Microsoft Windows always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Agent for Microsoft Windows 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.
3. In the **Automatic retry** section, define whether Veeam Agent for Microsoft Windows must attempt to run the backup job again if the job fails for some reason. Enter the number of attempts to run the job and define time intervals between them. If you select continuous backup, Veeam Agent for Microsoft Windows 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 impact performance of your server. 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.



Step 12. Review Backup Job Settings

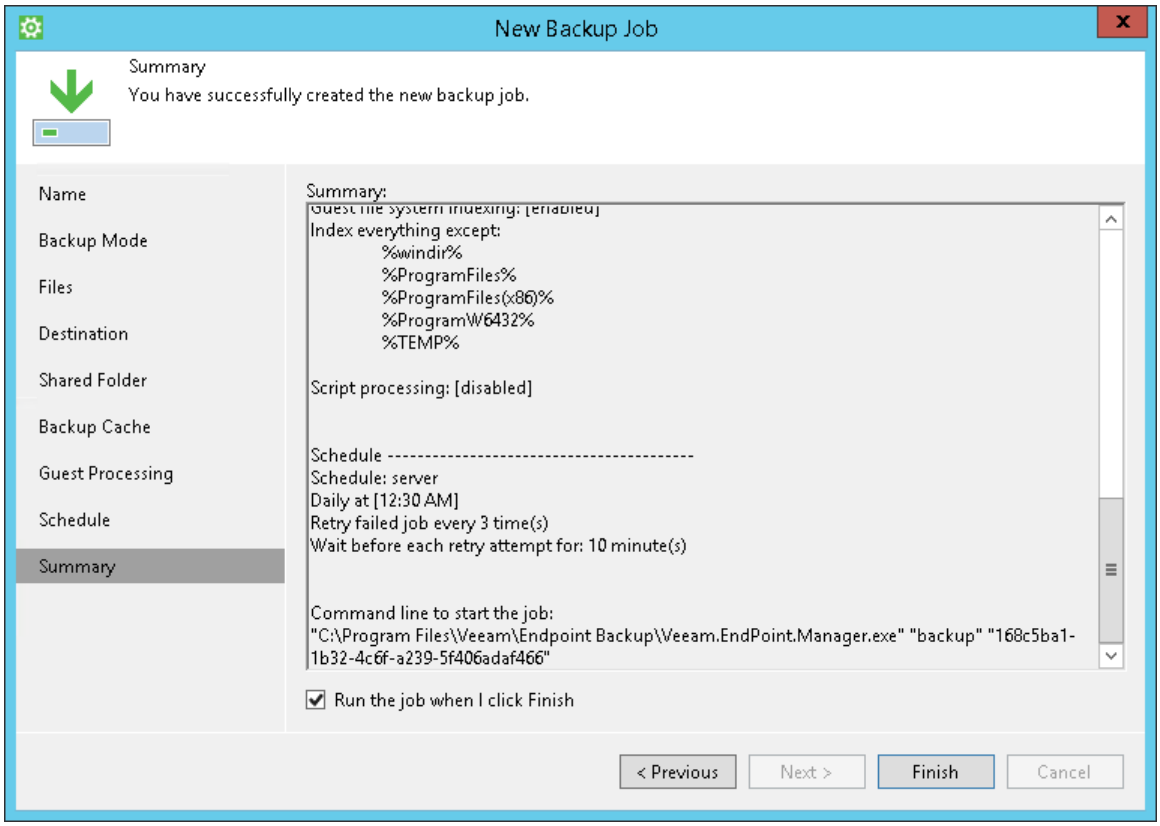
At the **Summary** step of the wizard, complete the backup job configuration process.

1. Review settings of the configured backup job.

   In addition to backup job settings, Veeam Agent for Microsoft Windows displays a command to start the job using the command-line interface.

2. To start the job after you close the wizard, select the **Run the job when I click Finish** check box.

3. Click **Finish**.



What You Do Next

After you configure the scheduled backup job, Veeam Agent for Microsoft Windows displays a clock over its icon in the system tray. The clock identifies that your computer is protected with the scheduled backup job. Veeam Agent for Microsoft Windows will periodically start the scheduled backup job to back up selected data and add a new restore point to the backup chain in the target location.

If necessary, you can also perform the following backup operations when you need it:

- Create a standalone full backup
- Create an incremental backup
- Create an active full backup

If some of your data gets lost or corrupted, you can do the following:

- Recover all computer volumes or specific volumes from the backup
- Recover individual files and folders from the backup

Managing Backup Job

After you configure the scheduled backup job, you can perform the following actions with it:

- Edit the backup job settings
- Disable and enable the backup job
- Remove the backup job

Editing Backup Job Settings

If you want to change settings of a backup job, you can edit it at any time. For example, you may want to edit the backup job to add a new folder to the backup scope, change the target location or job scheduling settings.

To access backup job settings, do one of the following:

- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**. Then hover over the name of the backup job whose settings you want to change, and select **Edit job**.
- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Backup** > **Configure backup**.

> **NOTE**
>
> The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

Then edit the job settings as required. To learn more about available job settings, see Configuring Backup Job.

If you change the target location for the backup job, during the next backup job session Veeam Agent for Microsoft Windows will perform full data backup. All subsequent backup sessions will produce incremental backups — Veeam Agent for Microsoft Windows will copy only changed data to the target location and add a new incremental backup file to the backup chain.

If you change the backup scope for the backup job, during the next backup job session Veeam Agent for Microsoft Windows will create a new incremental backup. The backup will contain the following data:

- All data blocks pertaining to new data added to the backup scope.
- Changed data blocks pertaining to original data in the backup scope (data that was processed by the job at the time before you changed the backup scope).

> **TIP**
>
> Full backup takes much more time than incremental backup. If you change the target location, you can copy an existing backup chain to the new location manually. In this case, the new backup job session will produce an incremental backup file and add it to the backup chain.

Editing Encryption Settings

If you change encryption settings for the backup job, during the next backup job session Veeam Agent for Microsoft Windows will create active full backup — encrypted (if encryption was enabled) or unencrypted (if encryption was disabled). All subsequent backup sessions will produce incremental backups.

Enabling or disabling encryption does not affect backup files that were created before you have changed encryption settings.

If the backup chain contains encrypted and unencrypted backup files, you need to provide a password to restore data from any restore point in this chain. After all encrypted backup files are removed from the backup chain according to retention policy, you will be able to restore data from remaining unencrypted restore points without providing a password.

Disabling and Enabling Scheduled Backups

You can disable scheduled backups if you do not want to run automatic backups for some period of time. For example, you may want to put backup activities on hold if you plan to perform resource consuming operations on your computer at the time when the backup job is scheduled. After the operations are completed, you can enable scheduled backups again.

The disabling option is applicable to backup job sessions started upon schedule. You can create standalone full backups and perform ad-hoc incremental backup even if the scheduled backups are disabled.

The disabling option is applicable to all backup jobs configured in Veeam Agent for Microsoft Windows. If you enable this option, all jobs that you configured will not start automatically upon the defined schedule. If you want to prevent a specific job from starting automatically, disable scheduling options in the properties of this job.

If you configured a backup job that is set up to create database log backups, after you disable scheduled backups, the database log backup job will be disabled, too.

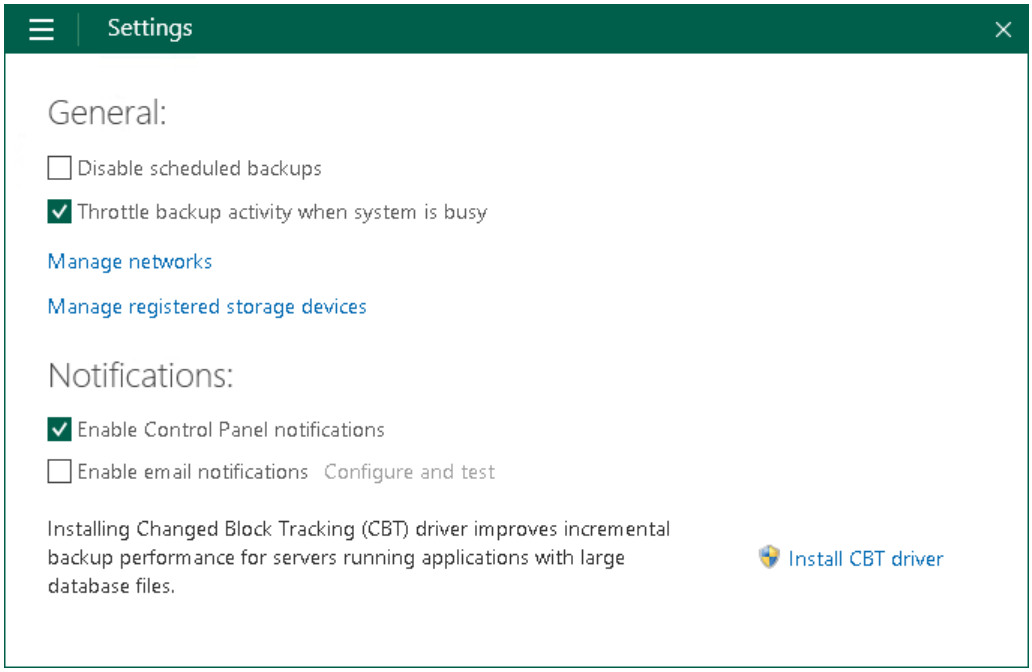
The disabling option does not put on hold the backup cache synchronization process. If Veeam Agent has created one or more backup files in the backup cache, and then the backup target becomes available, Veeam Agent will immediately upload backup files to the target location.

To disable scheduled backups:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Select the **Disable scheduled backups** check box.

To enable scheduled backups:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Clear the **Disable scheduled backups** check box.



Removing Backup Job

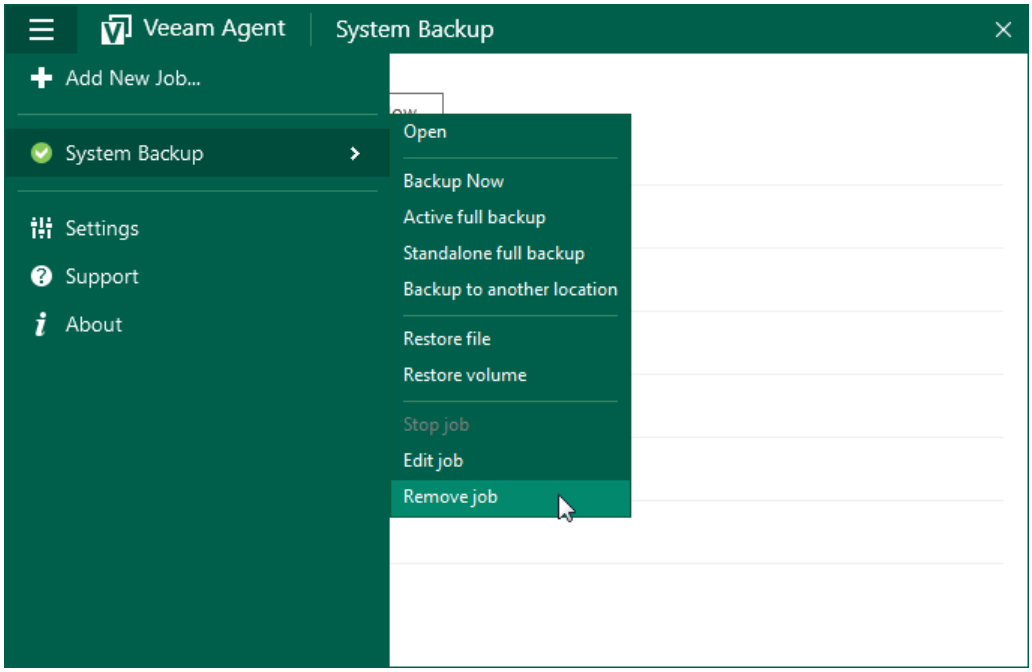
You can remove a backup job configured in Veeam Agent for Microsoft Windows. When you remove a backup job, Veeam Agent for Microsoft Windows also removes information about backup job sessions from the Veeam Agent database. After you remove the job, information about sessions of this job is not displayed in the **Status** view of the control panel any longer.

> **NOTE**
>
> You cannot remove a backup job that is currently running. You must stop the job first.

Backup files created by the deleted backup job remain intact in the target location.

To remove a backup job:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to remove, and select **Remove job**.
3. In the displayed window, click **Yes**.



Controlling Backup Post-Job Action

In the Free and Workstation product editions, you can set up Veeam Agent to perform a finalizing action after the backup job completes successfully:

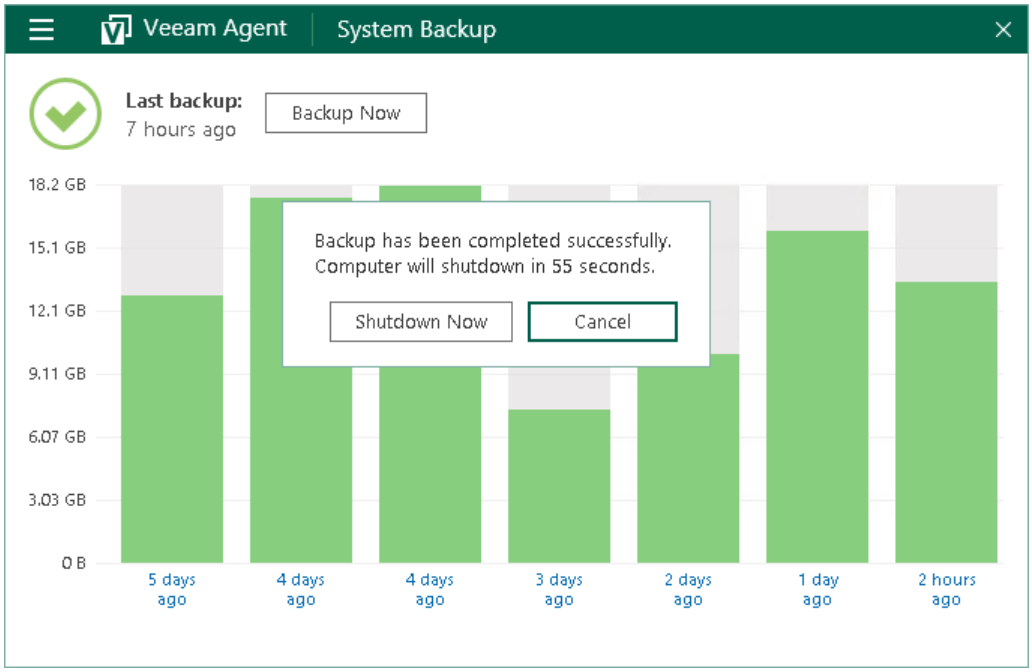
- *Sleep* — bring your computer to the standby mode.
- *Hibernate* — bring your computer to the hibernate mode.
- *Shutdown* — shut down your computer.

To learn more, see Specify Backup Schedule.

When the backup job completes, Veeam Agent opens control panel and prompts a dialog with a countdown to the specified action. Timeout between the backup job completion and the backup post-job action is 60 seconds.

- To proceed to the backup post-job action immediately, click **Sleep/Hibernate/Shutdown Now**.
- To cancel the action (for example, if you want to continue working or to save your data before turning off the computer), click **Cancel**.

If you do not select any option, Veeam Agent will perform the specified action when timeout expires.



Performing Ad-Hoc Backups

In addition to running scheduled backups, you can create ad-hoc backups of your data at any time you need. Veeam Agent for Microsoft Windows lets you perform the following types of ad-hoc backups:

- Incremental ad-hoc backup
- Active full backup
- Standalone full backup
- Backup to another location

Creating Incremental Backups

You can create an ad-hoc incremental backup of your data in addition to the scheduled backup. Ad-hoc incremental backup may be necessary if you want to capture your data at a specific point in time, for example, before you install new software on your computer. Ad-hoc incremental backup lets you produce an additional restore point in the backup chain at any time and does not require you to reconfigure the scheduling settings in the backup job.

Before you perform ad-hoc incremental backup, check the following prerequisites:

- The backup job that you want to use to perform ad-hoc incremental backup must be configured and successfully run at least once.
- You cannot perform ad-hoc incremental backup if a backup task of any type is currently running. These include a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.

To perform ad-hoc incremental backup:

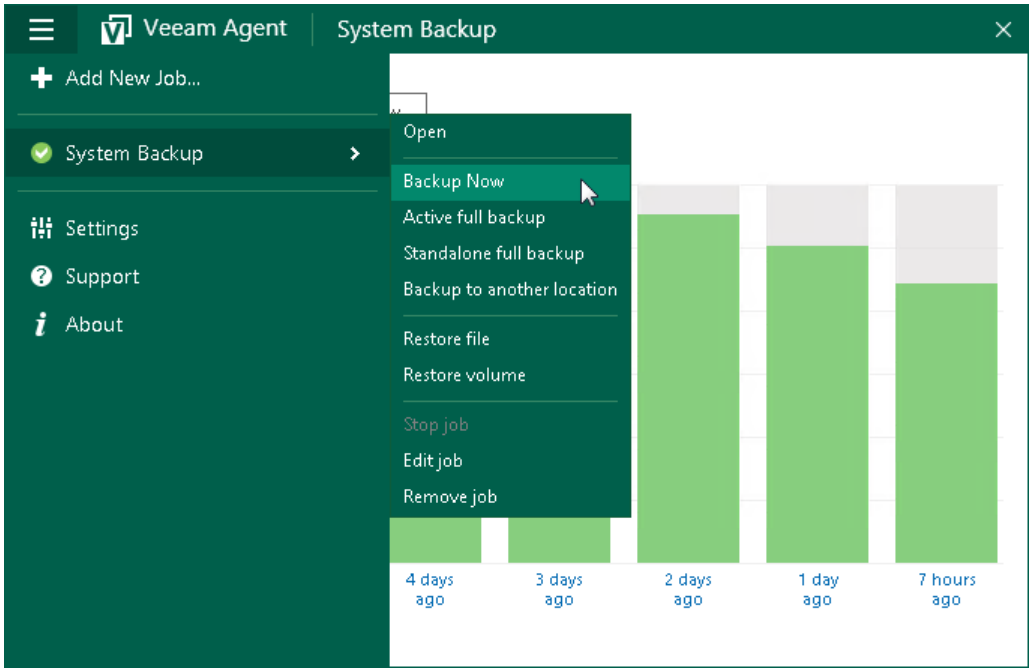
1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. Do either of the following:
    - If the Veeam Agent control panel currently displays statistics of the job that you want to use to perform ad-hoc incremental backup, in the job statistics window, click **Backup Now**.
    - If you want to perform ad-hoc incremental backup using another backup job configured in Veeam Agent, in the main menu, hover over the name of the necessary job and select **Backup now**.

If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the ad-hoc incremental backup task from the system tray menu:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select **Backup** > **Backup now**.

> **NOTE**
>
> The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

Veeam Agent for Microsoft Windows will perform incremental backup using settings of the scheduled backup job and add a new restore point to the backup chain in the target location.



Creating Active Full Backups

You can create an ad-hoc full backup — active full backup, and add it to the backup chain on the target storage. 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 on the target storage until it is removed from the backup chain according to the retention policy.

> **IMPORTANT**
>
> If you have a file-level backup job configured in Veeam Agent and you need to extend the volume where backed-up files reside, we strongly recommend to create an active full backup after the volume is extended. Otherwise, Veeam Agent may skip files during the job run even if these files are added to the backup scope.

Before you create an active full backup, check the following prerequisites:

- The backup job must be configured.
- You cannot create an active full backup if a backup task of any type is currently running. This includes a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.
- A user account under which you start the **Active full backup** operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

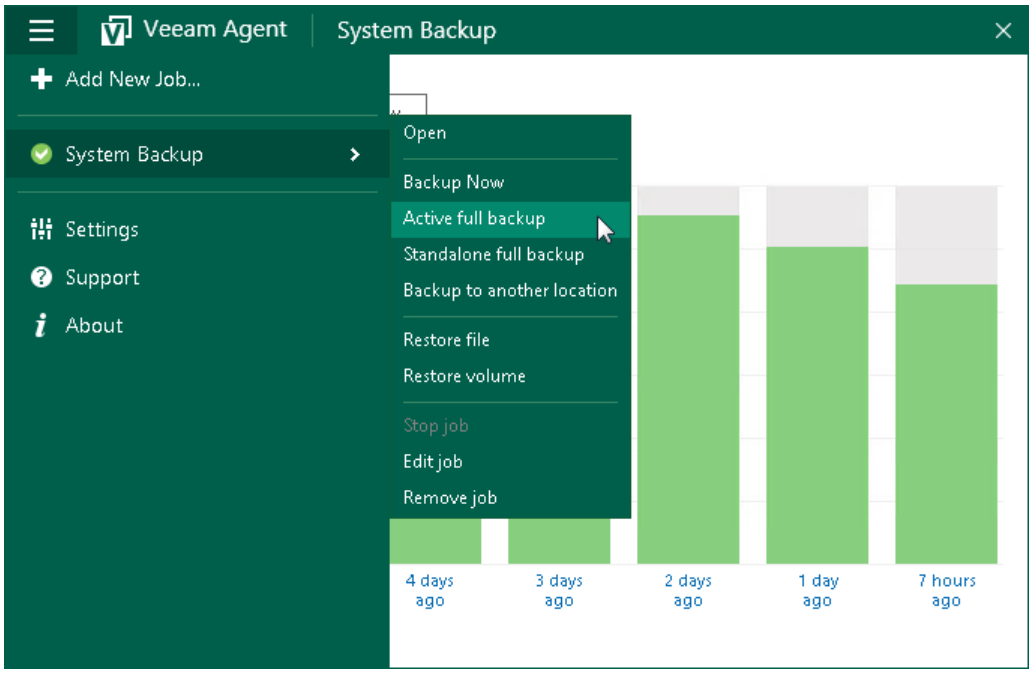
To perform active full backup:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to use to create an active full backup, and select **Active full backup**. Veeam Agent for Microsoft Windows will create a full backup file using settings of the backup job and add this backup file to the backup chain.

If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the active full backup task from the system tray menu:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select **Backup** > **Active full backup**.

> **NOTE**
>
> The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.



Creating Standalone Full Backups

If you want to back up your data at a specific point in time, you can create a standalone full backup. The standalone full backup is independent: it is not followed by subsequent incremental backups and is not removed by retention. You can use the standalone full backup to create an additional restore point from which you can recover your data.

Before you create a standalone full backup, check the following prerequisites:

- The backup job must be configured.
- You cannot create a standalone full backup if a backup task of any type is currently running. This includes a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.
- You cannot create a standalone full backup if the backup job is targeted at the cloud repository. If you want to create a full backup file not associated with the backup chain, you can perform standalone full backup to another location. To learn more, see Performing Backup to Another Location.
- A user account under which you start the **Standalone full backup** operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

To create a standalone full backup:

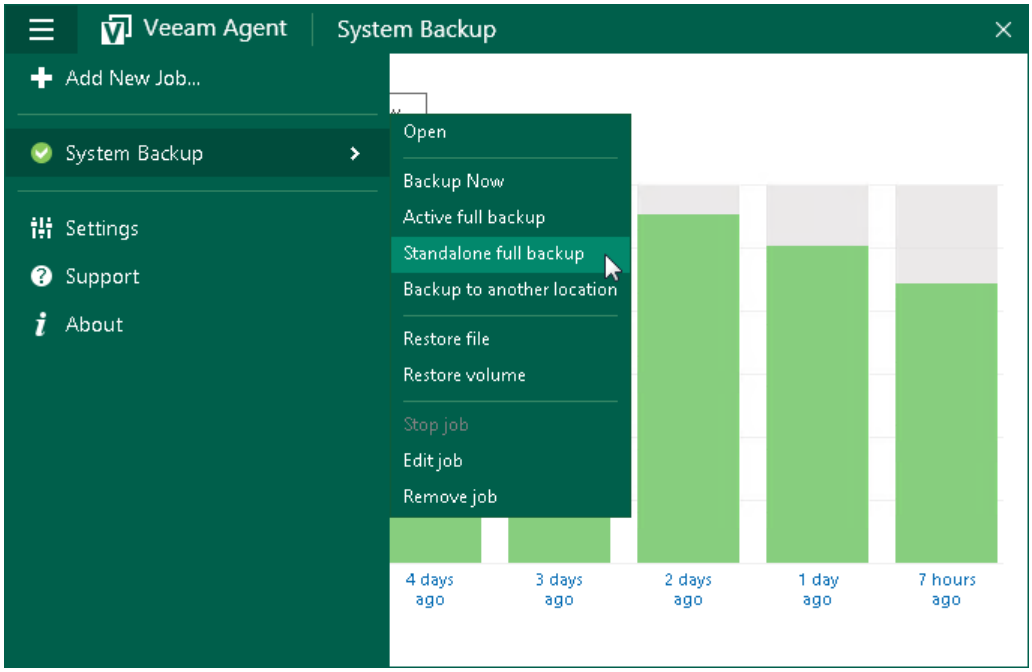
1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to use to create a standalone full backup, and select **Standalone full backup**. Veeam Agent for Microsoft Windows will create a full backup file using settings of the backup job. The resulting full backup file will be saved in the target location specified in the job settings, and placed to a separate folder. The folder is named in the following way: `<BackupJobName>.adhoc.<DateandTime>`.

If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the standalone full backup task from the system tray menu:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select **Backup** > **Standalone full backup**.

> **NOTE**
>
> The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

You can also create a standalone full backup in a location that is not specified in the backup job settings. To learn more, see Performing Backup to Another Location.



Performing Backup to Another Location

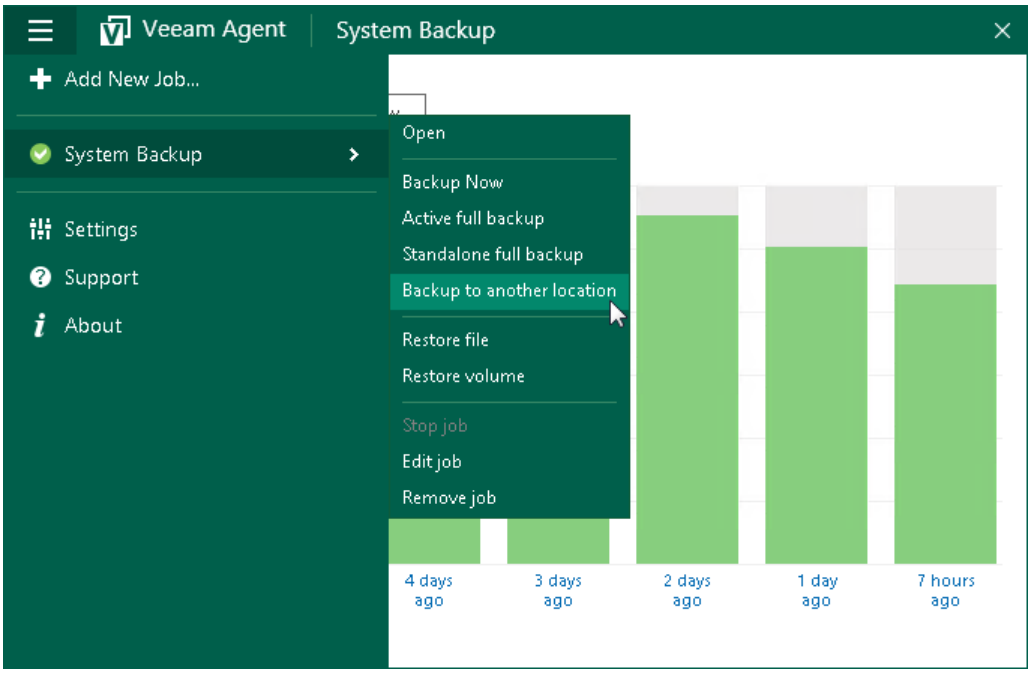
You can create a standalone full backup in a separate location that is not specified as a target location in the backup job settings. Performing backup to another location is similar to creating regular standalone full backups. The main difference is that you must manually select a target location in which Veeam Agent for Microsoft Windows will save the backup file.

Before you perform backup to another location, check the following prerequisites:

- The backup job must be configured.
- You cannot perform backup to another location if a backup task of any type is currently running. This includes a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.
- A user account under which you start the **Backup to another location** operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

To perform backup to another location:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to use to perform backup, and select **Backup to another location**.



NOTE

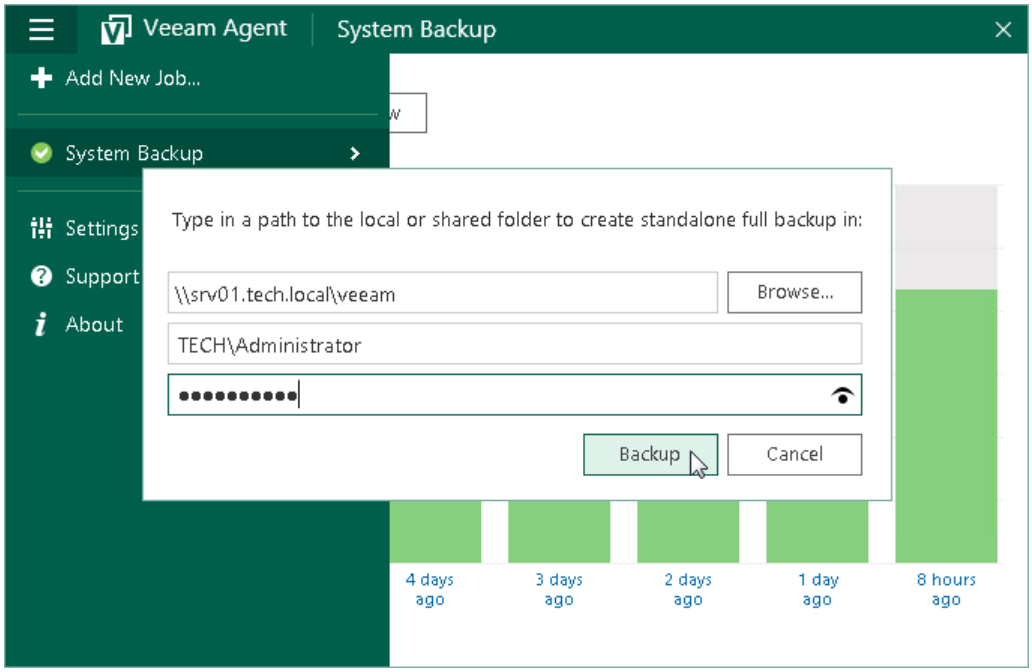
If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the standalone full backup task from the system tray menu:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select Backup > Backup to another location.

Note that the **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent.

3. In the standalone full backup dialog window, specify the target location for the backup file:
   - If you want to save the backup file in a folder on a local drive or a removable storage device, click **Browse** and select the necessary folder or type a path to the folder where backup file must be saved.
   - If you want to save the backup file in a network shared folder, type a UNC name of the network shared folder. Keep in mind that the UNC name always starts with two back slashes (\\). If the network shared folder requires authentication, specify a user name and password of the account that has access permissions on this shared folder. The user name must be specified in the *DOMAIN\UserName* format.
4. Click **Backup**.

   Veeam Agent for Microsoft Windows will create a full backup file using settings of the selected backup job. The resulting full backup file will be saved in a separate folder in the specified location. The folder is named in the following way: `<BackupJobName>.adhoc.<DateandTime>`.



Performing Backup with Command Line Interface

In addition to running scheduled backup jobs and performing ad-hoc backups from the Veeam Agent tray agent or control panel, you can create backups with the command line interface. For example, you can use commands for running a backup job in custom scripts to set up more detailed backup schedule than the daily schedule configured with the control panel.

You can run a backup job from the command line interface to create the following types of backups:

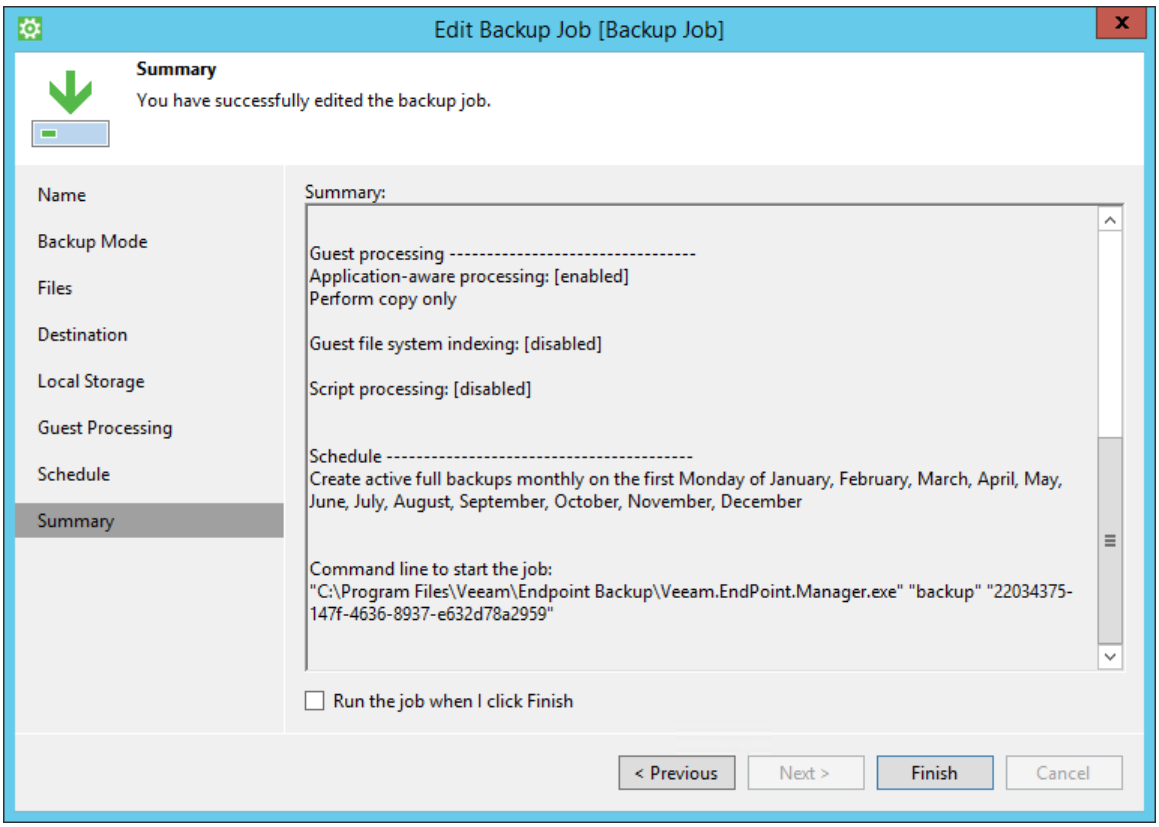
- Full or incremental backup (regular restore point in the backup chain)
- Active full backup
- Standalone full backup
- Backup to another location

> **IMPORTANT**
>
> If you create backups with the command line interface, Veeam Agent will ignore the following network connection restrictions:
>
> - Disabled backup over metered connections
> - Disabled backup over VPN connections
> - Backup over selected wireless networks
>
> To learn more about network usage settings, see Restricting Network Connections Usage.

Before you create a backup from the command line interface, check the following prerequisites:

- The backup job must be configured.
- You cannot run a backup job from the command line interface if a backup task for this job is currently running. This includes a scheduled backup, standalone full backup or ad-hoc incremental backup.

  If more than one backup job is configured, you can run multiple backup jobs at the same time.
- To create a standalone full backup or backup to another location, you must run the command line interface with administrative privileges.
- If more than one backup job is configured, you must specify the job ID in the *"00000000-0000-0000-0000-000000000000"* format. You can find the job ID at the **Summary** step of the backup job wizard.



Creating Backups

To perform backup, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /backup
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /backup "<ID>"
```

where `<ID>` — unique job ID generated by Veeam Agent for Microsoft Windows.

Creating Active Full Backups

To create an active full backup, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /activefull
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /activefull "<ID>"
```

where <ID> — unique job ID generated by Veeam Agent for Microsoft Windows.

Creating Standalone Full Backups

To create a standalone full backup, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone "<ID>"
```

where <ID> — unique job ID generated by Veeam Agent for Microsoft Windows.

Performing Backup to Another Location

To create a standalone full backup to a different location than a location that is specified in the backup job settings, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone <location>
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone <location> "<ID>"
```

where:

- `<location>` — path to a folder in which the backup should be created.
- `<ID>` — unique job ID generated by Veeam Agent for Microsoft Windows.

IMPORTANT

You can specify a network shared folder as a target location for standalone full backup only if read and write permissions on this folder are granted to *Everyone* or to the *LocalSystem* account of the Veeam Agent computer. You cannot specify credentials to access the network shared folder in the command.

Monitoring Backup Job Status

When you start a backup job from the command line interface, it runs automatically in the background. You can view information about the backup job session or the created restore point in the control panel. To learn more, see Viewing Statistics in Control Panel.

You can also use the last exit code to verify if the backup job has completed successfully. To check the last exit code, use the *%ERRORLEVEL%* variable in `cmd.exe`.

Veeam Agent for Microsoft Windows can provide the following exit codes:

- 0 — backup successfully created
- -1 — backup job failed to start or completed with error
- 5 — backup job is currently running and cannot be started from the command line interface

Deleting Backups

Backup files created with Veeam Agent for Microsoft Windows are removed automatically according to the retention policy settings. If necessary, you can remove a backup file manually with a file manager, for example, Microsoft Windows Explorer.

Always delete the whole backup chain from the target location. If you delete a full backup file or individual incremental backup file from the backup chain, the chain will be broken, and Veeam Agent for Microsoft Windows will fail to perform the scheduled backup next time.

If you remove the whole backup chain from the target location, during the next backup job session, Veeam Agent for Microsoft Windows will produce a new full backup. All subsequent backups will be incremental.

Deleting Backups with Command Line Interface

You can use the command line interface to delete Veeam Agent backups from a Veeam Cloud Connect repository or Microsoft OneDrive storage.

When you delete a backup from a cloud repository, Veeam Agent for Microsoft Windows deletes actual backup files from the cloud repository and removes records about the backup from the Veeam Backup & Replication database on the SP backup server. After information about the backup is removed from the SP backup server, Veeam Backup & Replication removes this information from its database and console on the tenant backup server, too.

When you delete a backup from Microsoft OneDrive, Veeam Agent for Microsoft Windows deletes actual backup files from the target location and removes records about the backup from the Veeam Agent for Microsoft Windows database.

Before you delete a backup, check the following prerequisites:

- To perform the delete backup operation, you must run the command line interface with administrative privileges.
- The delete backup operation cannot be performed if a backup or a restore task is currently running.
- [For Veeam Cloud Connect repository] The cloud repository from which you want to delete a backup must be specified as a target location for backup files in the backup job settings. To learn more, see Selecting Cloud Repository.
- [For Veeam Cloud Connect repository] Credentials of the user account (tenant account or subtenant account) whose backup you want to delete must be specified in the backup job settings. To learn more, see Specifying User Account Settings.
- [For Microsoft OneDrive] You must be signed in to the Microsoft OneDrive account whose backup you want to delete in the **New Backup Job** wizard. To learn more, see Microsoft OneDrive Settings.

To delete a Veeam Agent backup, use a command with the following syntax:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /deleteback
up
```

> **TIP**
>
> If more than one backup job is configured, Veeam Agent for Microsoft Windows will provide a list of backup jobs. To delete a backup, type the number that is displayed next to the backup job whose backup you want to delete.

Veeam Agent for Microsoft Windows can provide the following exit codes:

- 0 — backup job was successfully deleted
- -1 — the delete backup job operation failed

Managing Backup Cache

You can perform the following operations with the backup cache:

- Monitor backup cache activity
- Pause backup cache synchronization
- Delete restore points from the backup cache

Monitoring Backup Cache Activity

You can use the Veeam Agent control panel to view information about backup cache activity. In the **Backup Cache** view, Veeam Agent for Microsoft Windows displays a list of restore points that were created in the backup cache, their status and size of the resulting backup file. For restore points that are being uploaded or have been already uploaded to the target location, Veeam Agent for Microsoft Windows also displays the upload speed.

Viewing Restore Points in Backup Cache

To view information about backup cache activity, open the **Backup Cache** view in one of the following ways:

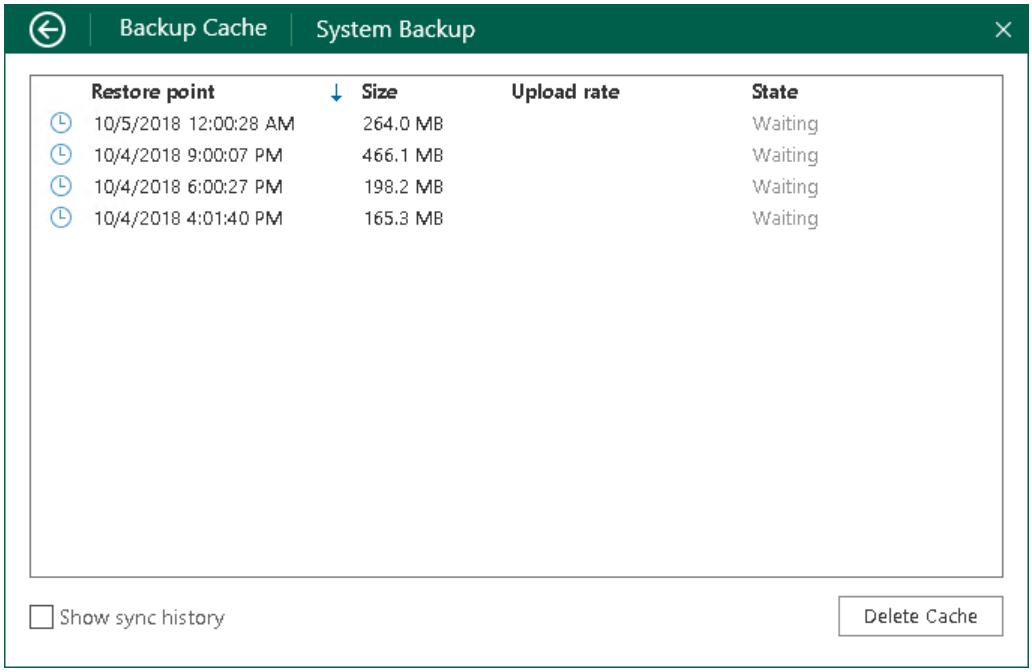
- If one backup job is configured:

  Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **View cache**. The Veeam Agent control panel will open, and you will pass immediately to the **Backup Cache** view.

  > **NOTE**
  >
  > The **View cache** option is not available in the system tray menu if more than one backup job is configured in Veeam Agent for Microsoft Windows.

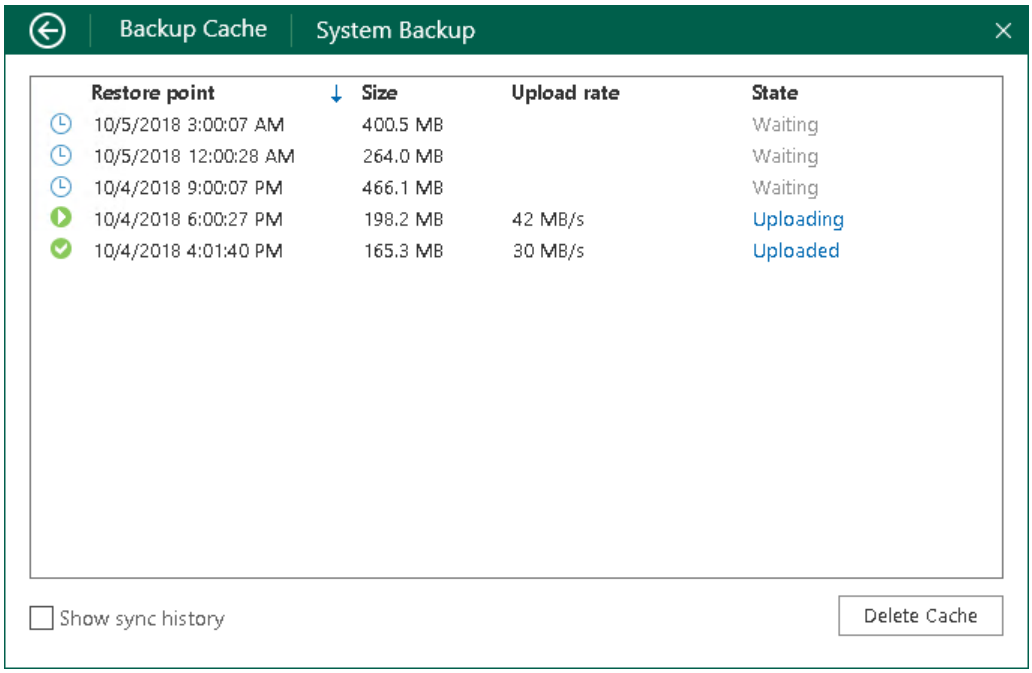
- If more than one backup job is configured:

  Double-click the Veeam Agent for Microsoft Windows icon, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**. In the main menu, hover over the name of the job that created restore points in the backup cache and select **Open backup cache**.



Viewing Backup Cache History

By default, the **Backup Cache** view contains a list of restore points that are waiting for upload or currently being uploaded to the target location. Restore points that have already been uploaded to the target location are not displayed in the list. To view such restore points, in the **Backup Cache** view, select the **Show sync history** check box.

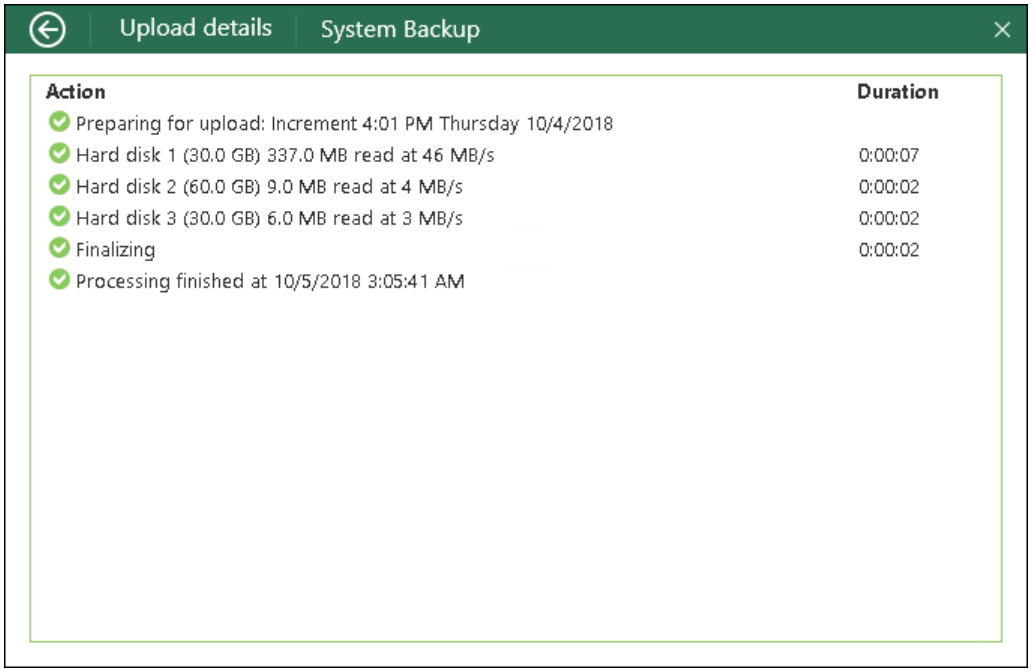


Viewing Upload Details for Restore Points

For every restore point that is being uploaded or has been uploaded to the target location, you can also view detailed information on the upload process:

1. Open the **Backup Cache** view in one of the following ways:
   - [If the backup cache is enabled in the properties of one job] Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **View cache**.
   - [If the backup cache is enabled in the properties of multiple jobs] Double-click the Veeam Agent for Microsoft Windows icon, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**. In the main menu, hover over the name of the job that created restore points in the backup cache and select **Open backup cache**.
2. In the **Backup Cache** view, click one of the following links next to the necessary restore point:
   - **Uploading** — for a restore point that is currently being uploaded to the target location.
   - **Uploaded** — for a restore point that has been already uploaded to the target location.

In the **Upload details** view, Veeam Agent for Microsoft Windows will provide detailed information about operations performed as part of the restore point upload process.



Pausing Backup Cache Synchronization

After at least one restore point is created in the backup cache, Veeam Agent for Microsoft Windows starts monitoring availability of the target location. To perform this operation, Veeam Agent for Microsoft Windows starts the backup cache synchronization job that runs in the background.

The backup cache synchronization job can be paused automatically by Veeam Agent for Microsoft Windows or manually.

Veeam Agent pauses the backup cache synchronization job automatically in case Veeam Agent computer is put into sleep or hibernate mode during the job session, and the backup job is targeted at one of the following locations:

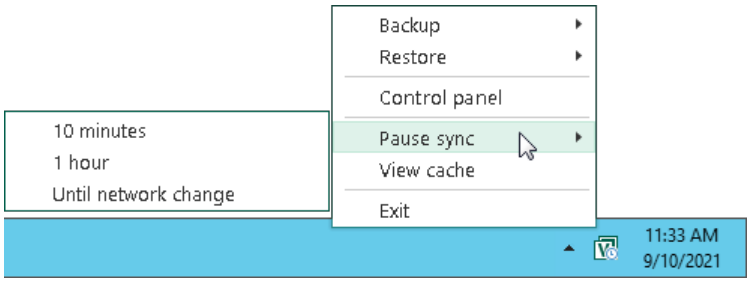
- Local computer drive
- Network shared folder
- Microsoft OneDrive

If the connection was lost during data transfer, Veeam Agent does not transfer all the data again. Instead, Veeam Agent continues data transfer that was started before the connection loss. After the connection is restored, Veeam Agent transfers only those data blocks that were not transferred before. Mind that if you store backups on Microsoft OneDrive, Veeam Agent transfers all the data again after the connection loss.

You can pause the backup cache synchronization job manually, for example, if you know that the target location will not become available for a while and want to reduce impact of Veeam Agent for Microsoft Windows on your OS performance.

To pause backup cache synchronization manually:

- Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
- Select one of the following options:
  - **Pause sync** > **10 minutes** — to pause backup cache synchronization for 10 minutes.
  - **Pause sync** > **1 hour** — to pause backup cache synchronization for 1 hour.
  - **Pause sync** > **Until network change** — to pause backup cache synchronization until new network settings are applied to the network adapter of the Veeam Agent computer.



Deleting Restore Points from Backup Cache

You can delete restore points from the backup cache manually if needed. For example, it is required to delete restore points after Veeam Agent for Microsoft Windows creates one or more restore points at the backup cache, and then you change the target location for backup files in the backup job settings.

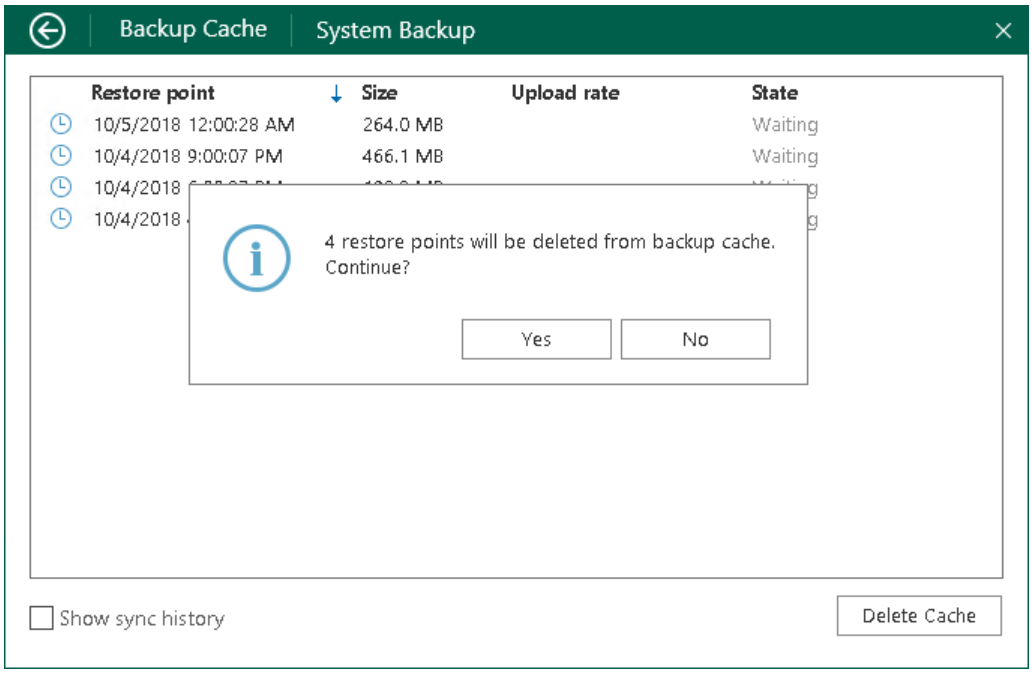
To delete restore points from the backup cache:

1. Open the **Backup Cache** view in one of the following ways:
   - [If the backup cache is enabled in the properties of one job] Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **View cache**.

     > **NOTE**
     >
     > The **View cache** option is not available in the system tray menu if the backup cache is enabled for more than one backup job configured in Veeam Agent for Microsoft Windows.

   - [If the backup cache is enabled in the properties of multiple jobs] Double-click the Veeam Agent for Microsoft Windows icon, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**. In the main menu, hover over the name of the job that created restore points in the backup cache and select **Open backup cache**.
2. In the **Backup Cache** view, click **Delete Cache**.
3. In the window notifying that restore points will be deleted from the backup cache, click **Yes**.

> **NOTE**
>
> The first backup job session following the deletion of restore points from the backup cache must complete successfully and create backup files on the target location. During this session, Veeam Agent for Microsoft Windows will create a new map of target location data blocks in the backup cache. If you delete restore points from the backup cache, and then run the backup job when the target location is unavailable, the backup job will fail.



Performing Restore

If you experience a problem with your computer, your data gets lost or corrupted, you can use one of the following options to recover your data or bring the computer back to work:

- Restoring from Veeam Recovery Media
- Using Veeam Agent and Microsoft Windows Tools
- Using Microsoft Windows Recovery Environment
- Restoring Volumes
- Restoring Files and Folders
- Restoring Data from Encrypted Backups

Restoring from Veeam Recovery Media

If the OS on your computer fails to start, you can use the Veeam Recovery Media to recover your computer. The Veeam Recovery Media will help you boot the computer in the limited mode. After booting, you can use Veeam Agent for Microsoft Windows or standard Microsoft Windows tools to diagnose problems and fix errors. You can also use a backup created with Veeam Agent for Microsoft Windows to restore the whole system image of your computer or specific volumes on your computer.

Before You Begin

Before you boot from the recovery image and recover your data, check the following prerequisites:

- You must have a successfully created recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For data recovery, you can use a volume-level backup created with Veeam Agent for Microsoft Windows or system image created with Microsoft Windows. Make sure that the backup or system image is available on the computer drive (local or external), in a network shared folder, on the backup repository managed by a Veeam backup server or on the cloud repository.
- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- Recovery images for Microsoft Windows 32-bit OSes can be booted in the BIOS system only. Recovery images for Microsoft Windows 64-bit OSes can be booted in the BIOS and UEFI systems.

Mind the following:

- When you create a Veeam Recovery Media, Veeam Agent for Microsoft Windows stores settings for languages added to the list of input languages on your computer. If necessary, you can switch between languages using a hotkey combination when working with the **Veeam Recovery Media** wizard. The default key combination is typically **[Shift]** + **[Alt]**.
- NIC Teaming is not supported by Veeam Recovery Media due to Windows Recovery Environment limitations. To avoid network connection issues in Veeam Recovery Media, we recommend to disable NIC Teaming in the network switch configuration.
- You can open the Command Prompt at any moment. To do this, press **[Shift]** + **[F10]** on the keyboard.
- If you perform restore on a tablet, you can use a virtual keyboard to enter necessary restore settings in the **Veeam Recovery Media** wizard.

Step 1. Boot from Veeam Recovery Media

To boot from the Veeam Recovery Media:

1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.

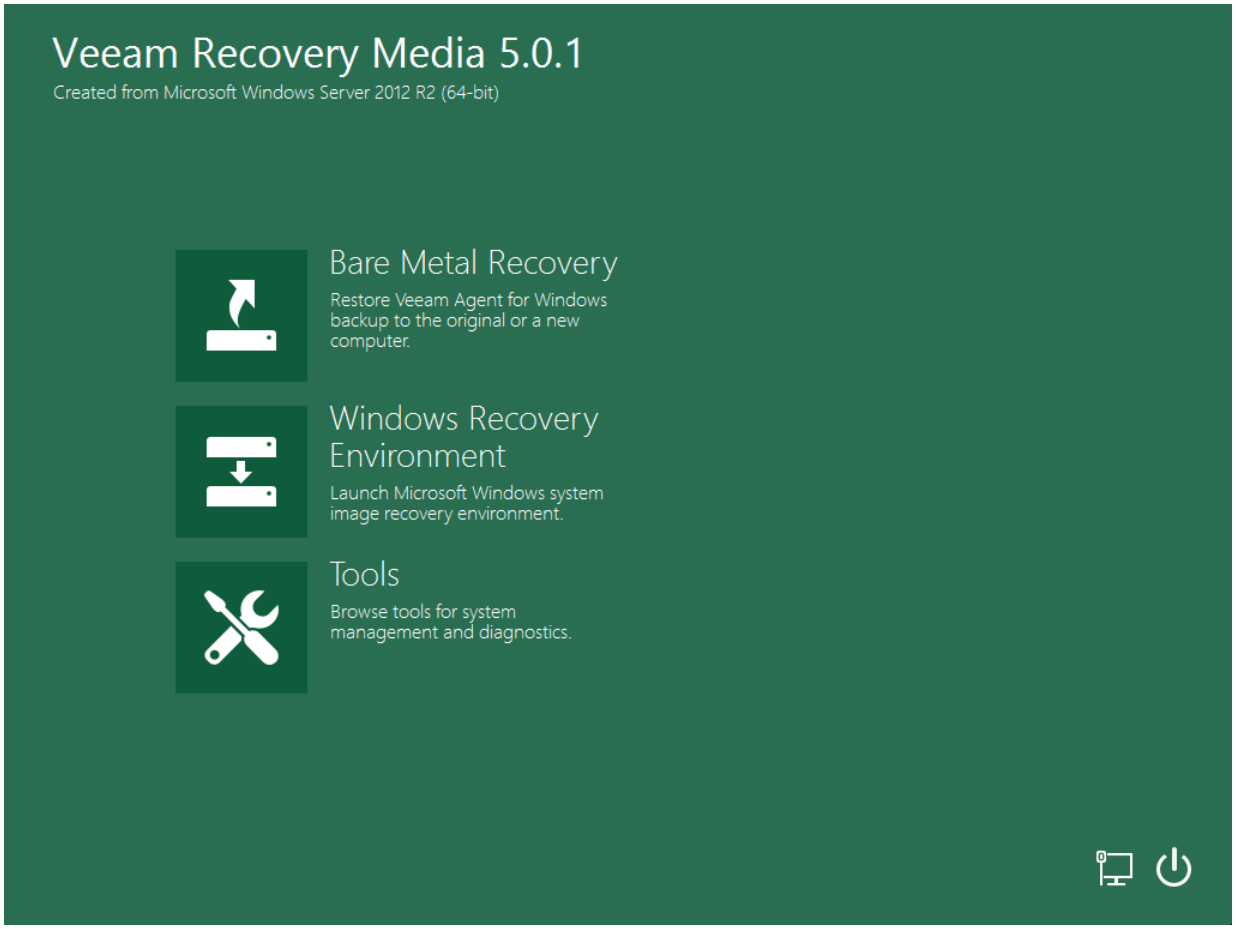
   [For removable storage device] Attach the removable storage device with the recovery image to your computer.

2. Start your computer.
3. You will be offered to boot the OS from the CD/DVD/BD or attached removable storage. Press any key on the keyboard to continue.
4. Wait for Veeam Agent for Microsoft Windows to load files from the Veeam Recovery Media. Loading the OS from the Veeam Recovery Media usually takes more time than loading the OS from the local computer drive.
5. After the OS has loaded, make sure network settings are specified correctly and configure network if necessary. To learn more, see Select Network Adapter or Wireless Network.
6. Choose the necessary recovery tool to use. Veeam Agent for Microsoft Windows offers the following tools:
   - Bare Metal Recovery — the Veeam Recovery Media wizard to recover data on the original computer or perform bare metal recovery.
   - Windows Recovery Environment — built-in Microsoft Windows tools to recover the computer system image.
   - Tools — Veeam Agent for Microsoft Windows and Microsoft Windows utilities for advanced computer administration.

TIP

Consider the following:

- To shut down or restart your computer, click the **Power Options** button at the bottom right corner of the Veeam Recovery Media screen and select the necessary option: **Shut down** or **Restart**.
- You can use the Veeam Recovery Media to start the Microsoft Windows OS in the safe mode. To do this, click the **Power Options** button at the bottom right corner of the Veeam Recovery Media screen and select **Boot in safe mode**. In the displayed window, select the necessary safe boot mode (*Minimal*, *Network* or *Repair*), and click **OK**.



Step 2. Select Network Adapter or Wireless Network

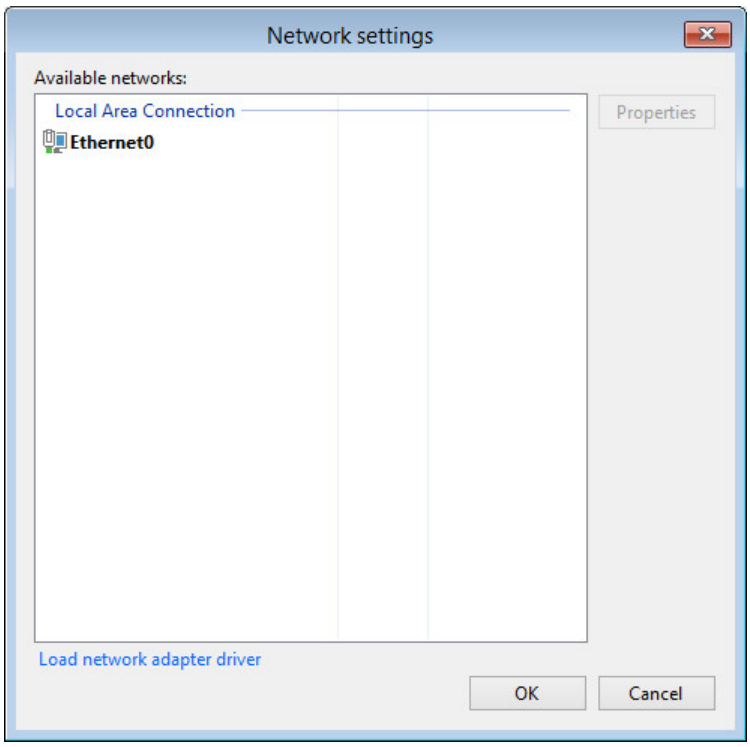
To open the **Network settings** window, click the **Network Settings** button at the bottom right corner of the **Veeam Recovery Media** screen.

> **TIP**
>
> The *Network Settings* button appearance may vary depending on the detected network connection: Ethernet or wireless. If your computer is connected to a wireless network, the *Network Settings* button will indicate Wi-Fi signal strength.

Select a network adapter or wireless network that you want to use to connect to the network shared folder or Veeam backup repository where the backup resides.

- If network connection settings are included in the Veeam Recovery Media, or if there is a DHCP server in your network, Veeam Agent for Microsoft Windows will configure the network settings automatically and display available network adapters in the list.
- If you want to access the network shared folder or Veeam backup repository using a wireless network, select the necessary network in the list and click **Next**. If the wireless network is password protected, you will be prompted to specify a password for this network.
- You can manually configure TPC/IP v4 settings for adapters if necessary. To do this, select an adapter in the list and click **Properties**.

> **NOTE**
>
> You will be prompted to configure network settings manually if Veeam Agent for Microsoft Windows does not detect available networks and there are no network settings included in the Veeam Recovery Media.



Installing Network Adapter Drivers

The list of networks can be empty. This can happen in two situations:

- The driver for the network adapter is included in the Veeam Recovery Media but failed to be installed automatically for some reason.
- The driver for the network adapter is not included in the Veeam Recovery Media.

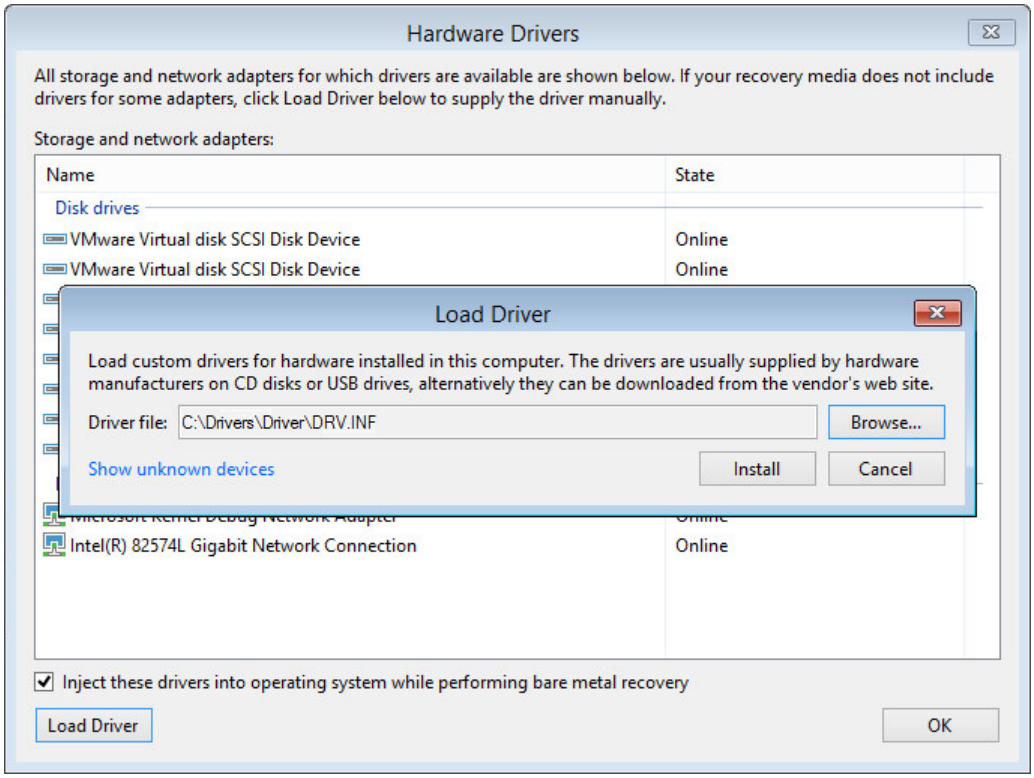
To install drivers that were included in the Veeam Recovery Media:

1. At the **Network settings** window, click **Load network adapter driver**.
2. In the **Hardware Drivers** window, select the necessary device.

   If you want to include in the restored operating system all the drivers that were saved to the Veeam Recovery Media, select the **Inject these drivers into operating system while performing bare metal recovery** option. In case the option is not selected, the restored operating system will include only default Windows hardware drivers.
3. Click the **Install** link next to the selected device.

To install drivers that were not included in the Veeam Recovery Media:

1. At the **Network settings** window, click **Load network adapter driver**.
2. At the bottom of the **Hardware Drivers** window, click the **Load Driver** button and select the INF file in the driver package folder. You can also click the **Show unknown devices** link to see a list of all existing devices without drivers. This information may help you to identify the exact device for which you need to install the driver.
3. Click **Install**.



Step 3. Launch Veeam Recovery Media Wizard

To launch the **Veeam Recovery Media** wizard, on the **Veeam Recovery Media** screen, click **Bare Metal Recovery**.

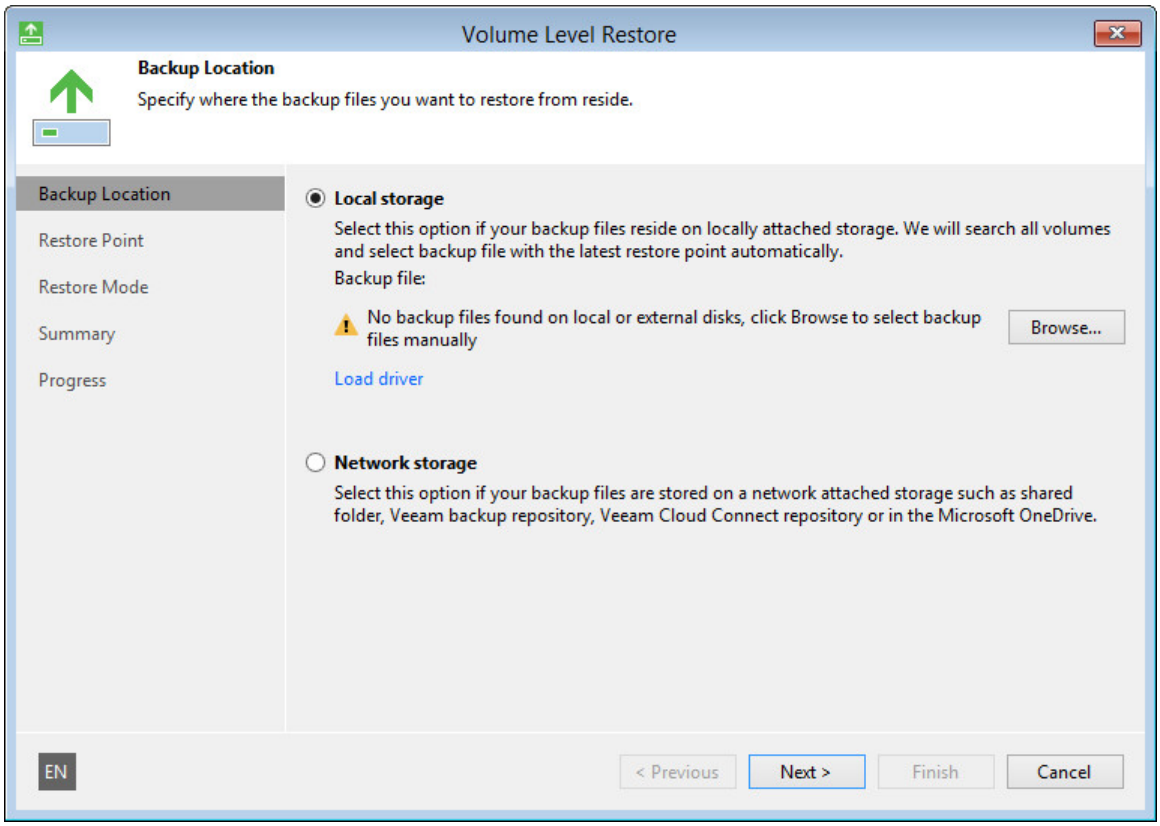


Step 4. Specify Backup File Location

At the **Backup Location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

By default, Veeam Agent for Microsoft Windows automatically locates the latest backup on the computer drive and you pass immediately to the Restore Point step of the wizard. If Veeam Agent for Microsoft Windows fails to locate the backup on the local computer drive for some reason, or the backup file is located in a network shared folder, on a backup repository or cloud repository, select where the backup file resides:

- **Local storage** — select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. Click **Browse** and select a backup metadata file (VBM).
- **Network storage** — select this option if the backup file resides in a network shared folder, in a Microsoft OneDrive cloud storage, on a backup repository managed by a Veeam backup server or on a cloud repository exposed to you by a Veeam Cloud Connect service provider. In this case, the Veeam Recovery Media wizard will include additional steps for specifying the backup file location settings.



Installing Drivers for Remote Storage Devices

A removable storage device with the backup file may not be displayed in the list of devices. This can happen in two situations:

- The driver for the remote storage device is included in the Veeam Recovery Media but failed to be installed automatically for some reason.
- The driver for the remote storage device is not included in the Veeam Recovery Media.

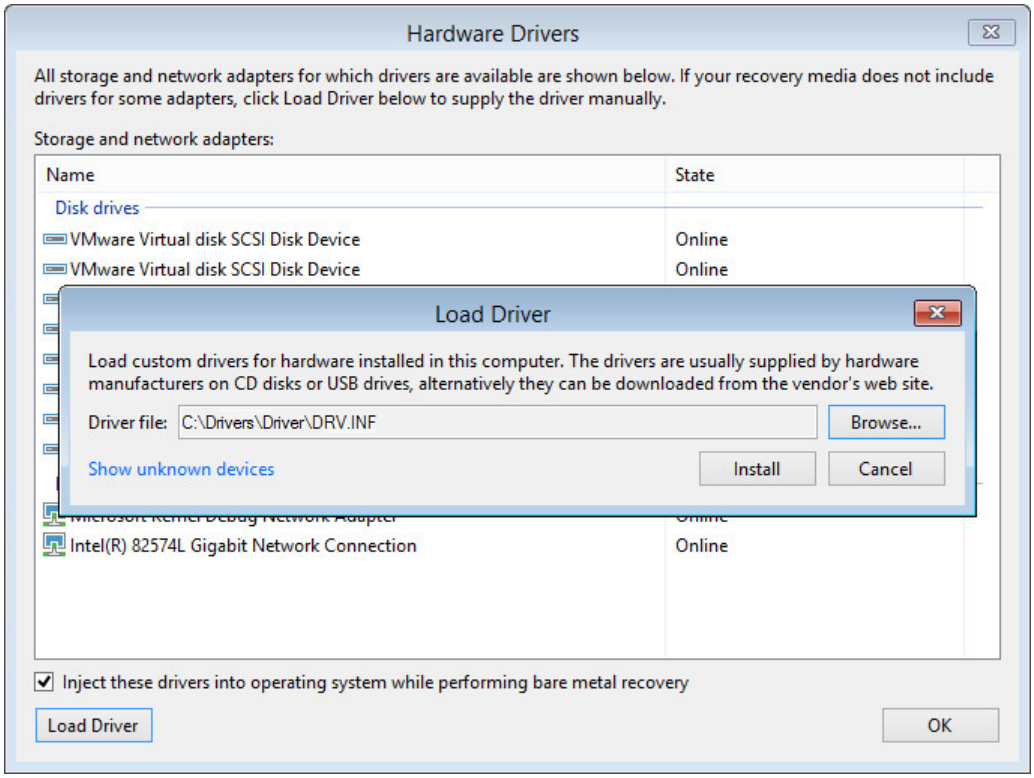
To install drivers that were included in the Veeam Recovery Media:

1. At the **Backup Location** step of the wizard, click **Load driver**.
2. In the **Hardware Drivers** window, select the necessary device.

   If you want to include in the restored operating system all the drivers that were saved to the Veeam Recovery Media, select the **Inject these drivers into operating system while performing bare metal recovery** option. In case the option is not selected, the restored operating system will include only default Windows hardware drivers.
3. Click the **Install** link next to the selected device.

To install drivers that were not included in the Veeam Recovery Media:

1. At the **Backup Location** step of the wizard, click **Load driver**.
2. At the bottom of the **Hardware Drivers** window, click the **Load Driver** button and select the INF file in the driver package folder. You can also click the **Show unknown devices** link to see a list of all existing devices without drivers. This information may help you to identify the exact device for which you need to install the driver.
3. Click **Install**.

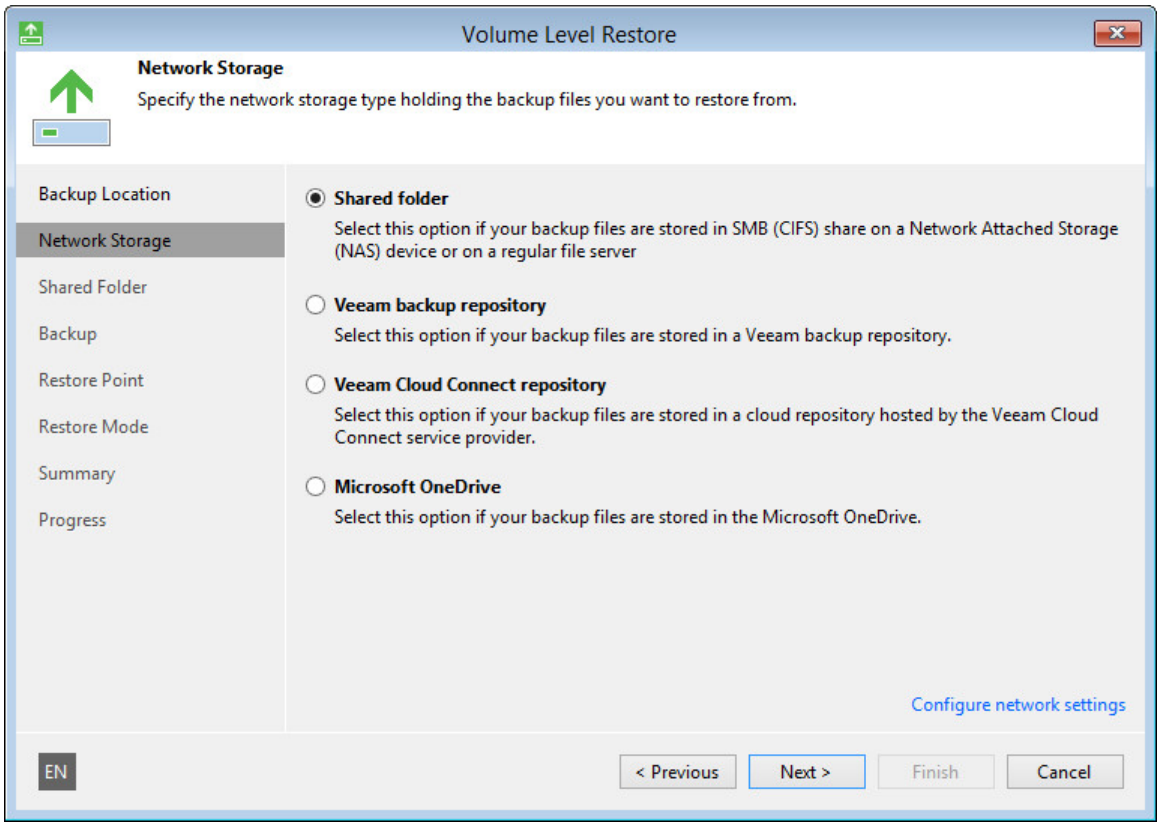


Step 5. Select Remote Storage Type

The **Network Storage** step of the wizard is available if you have selected to restore data from a backup file that resides in a remote location — in a network shared folder, on a backup repository or a cloud repository.

Select where the backup file resides:

- **Shared folder** — select this option if the backup file resides in a network shared folder. With this option selected, you will pass to the Shared Folder step of the wizard.
- **Veeam backup repository** — select this option if the backup file resides on a backup repository managed by a Veeam backup server. With this option selected, you will pass to the Backup Server step of the wizard.
- **Veeam Cloud Connect repository** — select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the Service Provider step of the wizard.
- **Microsoft OneDrive** — select this option if the backup file resides in a Microsoft OneDrive cloud storage. With this option selected, you will pass to the Microsoft OneDrive step of the wizard.



Step 6. Specify Remote Storage Settings

Specify settings for the remote storage that contains a backup file from which you plan to restore data:

- Shared folder settings — if you have selected the **Shared folder** option at the Network Storage step of the wizard.
- Veeam backup repository settings — if you have selected the **Veeam backup repository** option at the Network Storage step of the wizard.
- Veeam Cloud Connect repository settings — if you have selected the **Veeam Cloud Connect repository** option at the Network Storage step of the wizard.
- Microsoft OneDrive settings — if you have selected the **Microsoft OneDrive** option at the Network Storage step of the wizard.

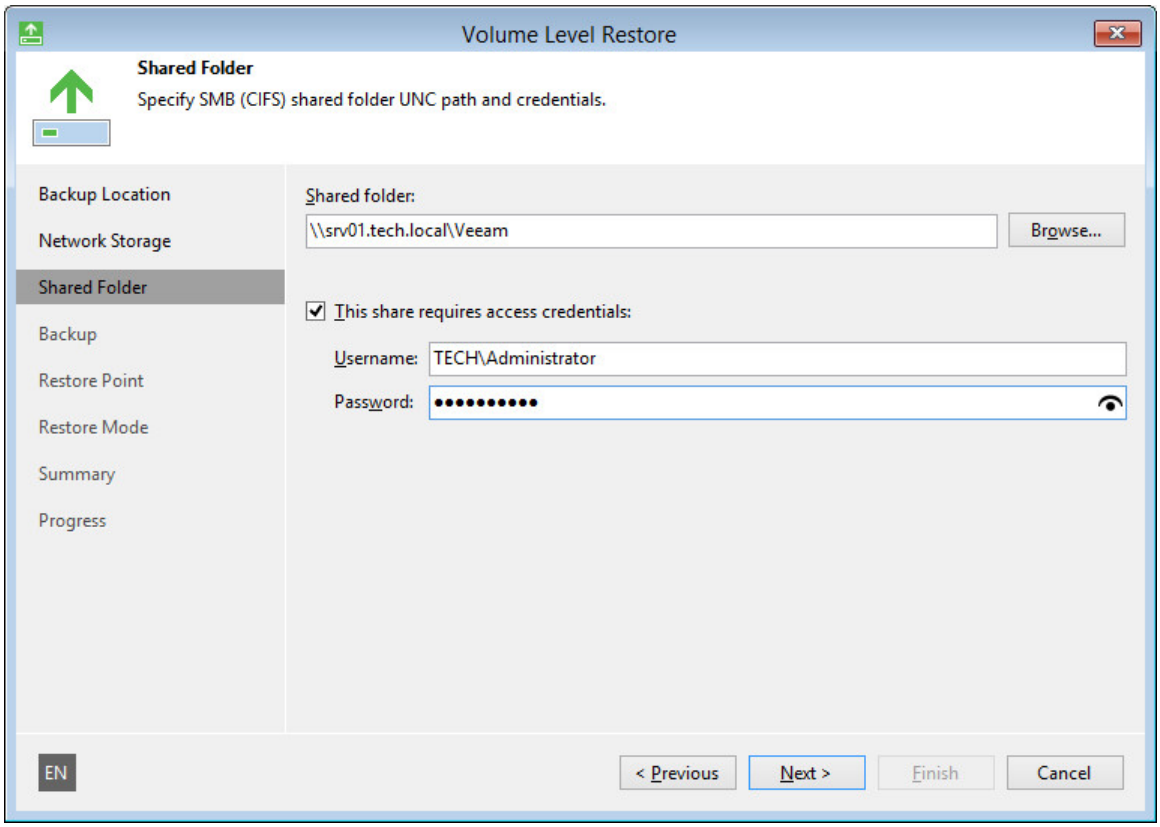
Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. In the **Shared folder** field, enter a UNC name of the network shared folder with a backup file. Keep in mind that the UNC name always starts with two back slashes (\\).
2. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the down-level logon name format. For example, *DOMAIN\UserName* or *HOSTNAME\UserName*.

   To view the specified password, click and hold the eye icon on the right of the **Password** field.



Backup Server Settings

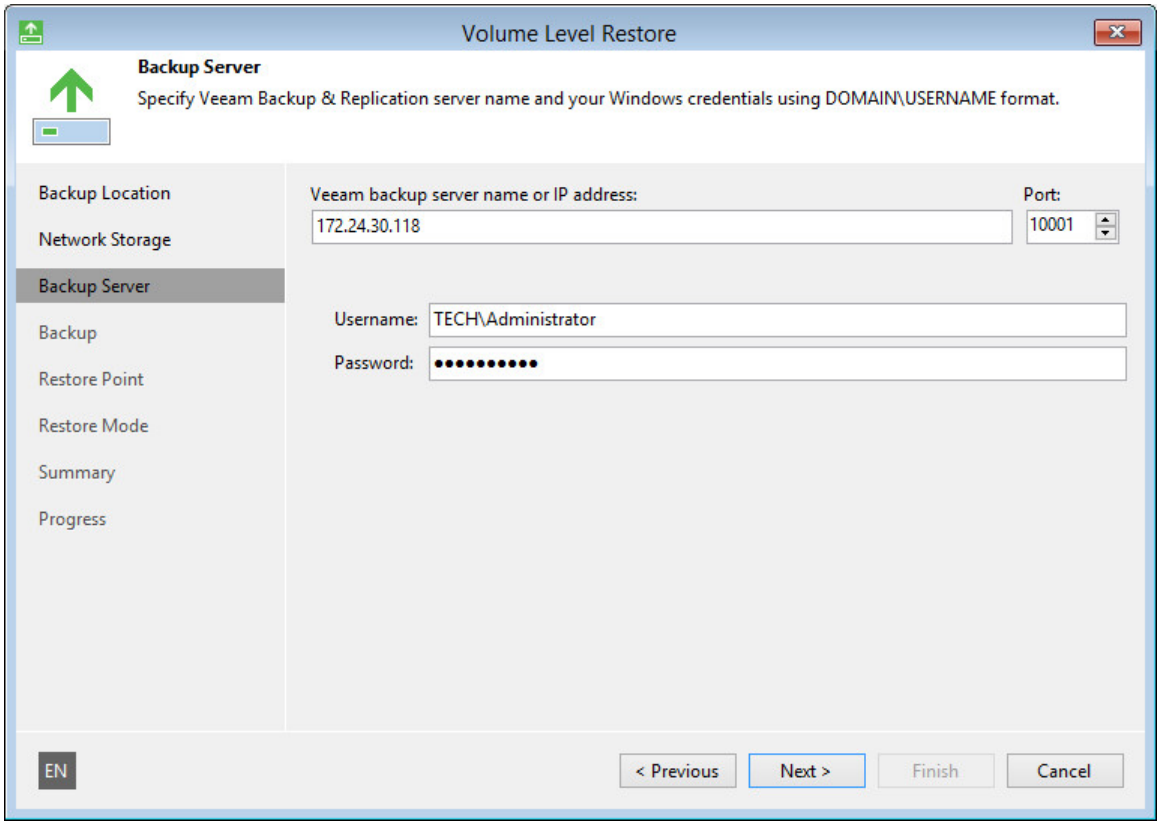
The **Backup Server** step of the wizard is available if you have chosen to restore data from a backup file located on a backup repository.

Specify settings for the Veeam backup server that manages the backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup repository. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. In the **Username** and **Password** fields, enter a user name and password of the account that has access to this backup repository. Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see Setting Up User Permissions on Backup Repositories.

IMPORTANT

If you specify a DNS name of the Veeam backup server, make sure that the Veeam backup server name is resolved into IPv4 address on the machine where Veeam Agent for Microsoft Windows is installed. The Veeam Backup Service in Veeam Backup & Replication listens on IPv4 addresses only. If the Veeam backup server name is resolved into IPv6 address, Veeam Agent for Microsoft Windows will fail to connect to the Veeam backup server.



Service Provider Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

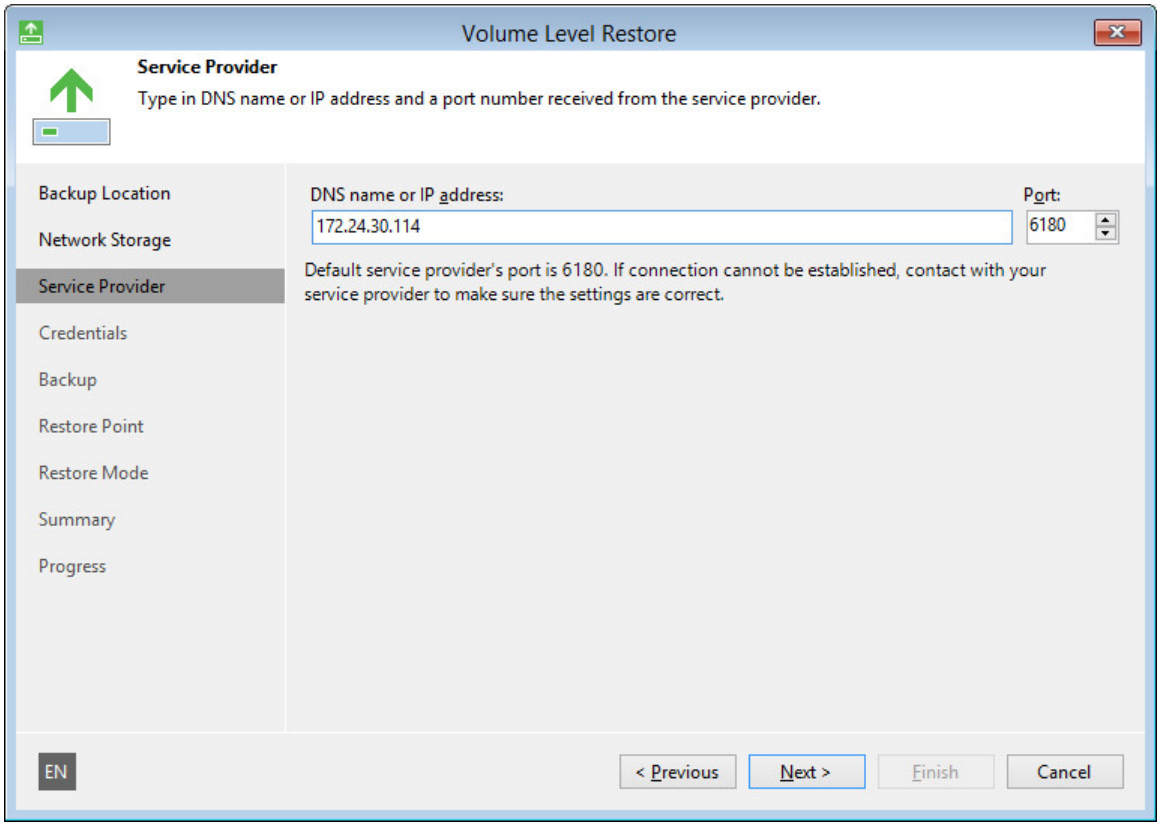
1. Specify service provider settings.
2. Verify the TLS certificate and specify user account settings.

Specifying Service Provider Settings

The **Service provider** step of the wizard is available if you have chosen to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.



Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the service provider.

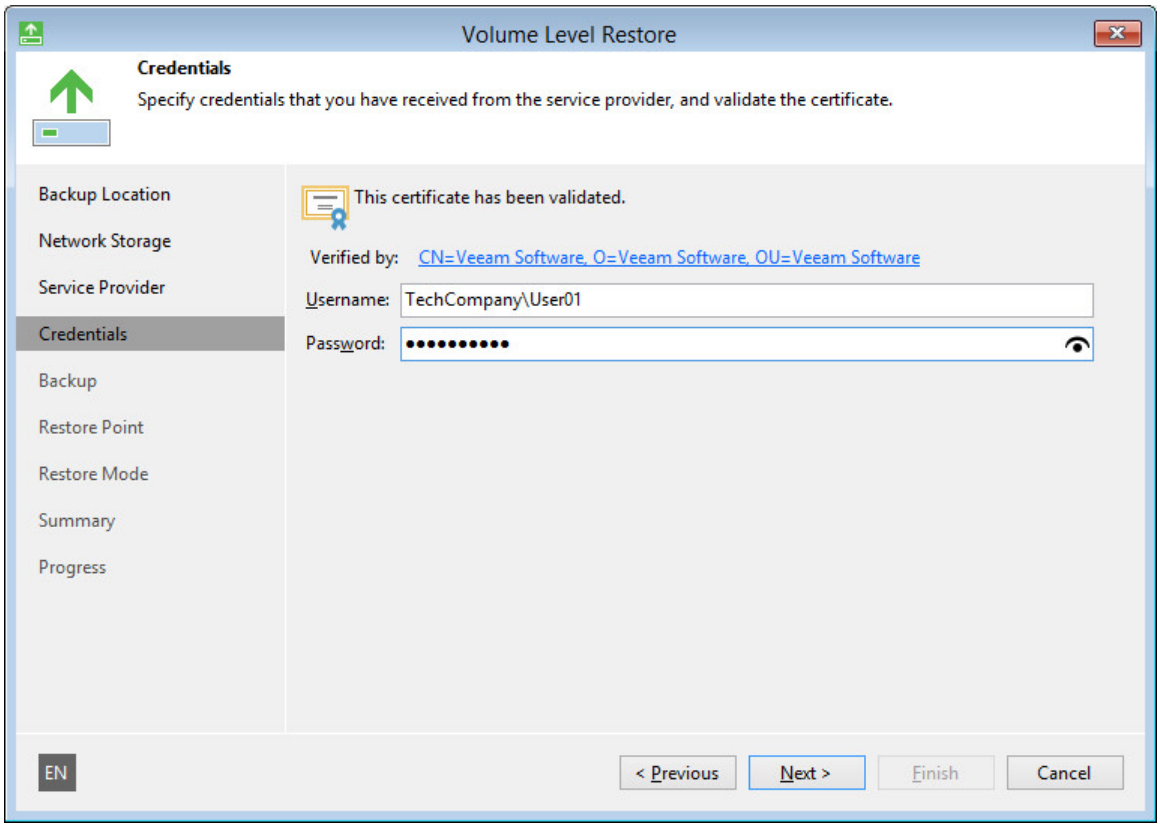
1. At the top of the wizard window, Veeam Agent for Microsoft Windows displays information about the TLS certificate obtained from the SP side. You can view the certificate settings and verify the TLS certificate.

   TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

   - To view the TLS certificate, click the certificate link.
   - To verify if the TLS certificate with a thumbprint, copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Fingerprint for certificate verification** field. Click **Verify**. Veeam Agent for Microsoft Windows will check if the thumbprint you enter matches the thumbprint of the obtained TLS certificate.

2. In the **Username** field, enter the user name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The user name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.



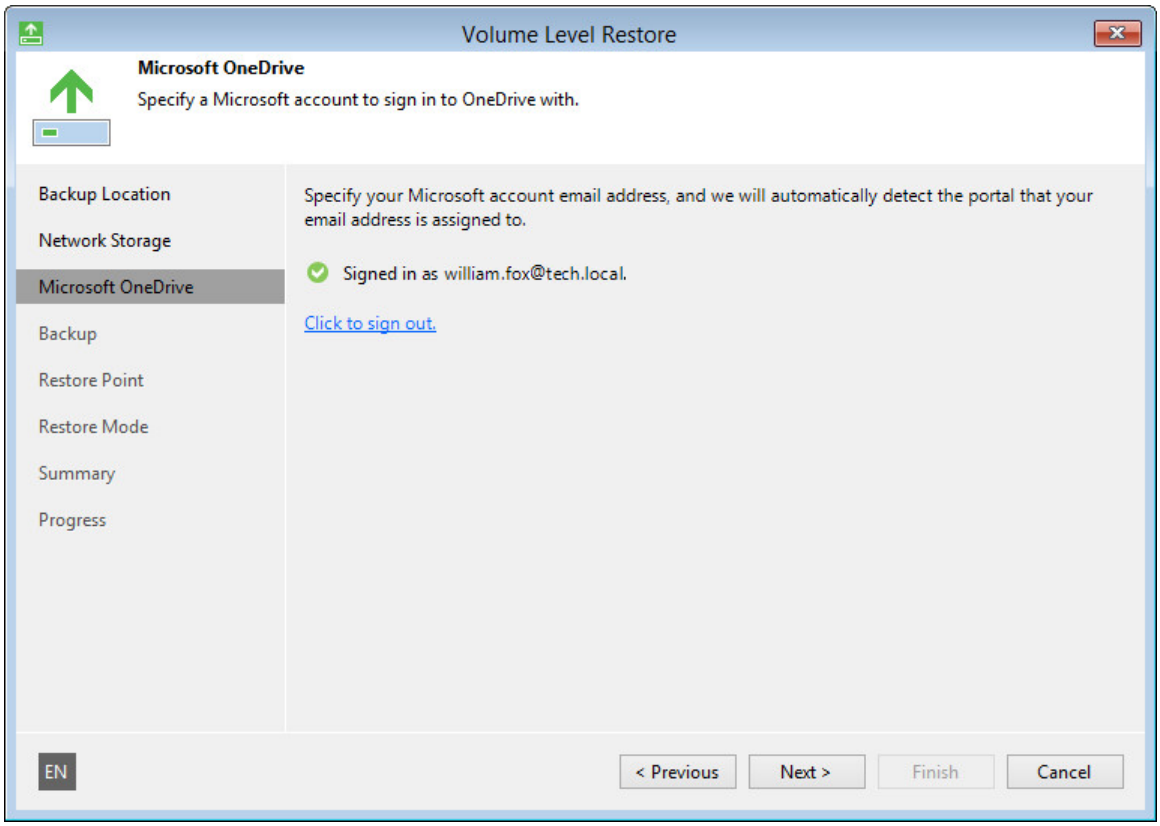
Microsoft OneDrive Settings

The **Microsoft OneDrive** step of the wizard is available if you have chosen to restore data from a backup file located in the Microsoft OneDrive cloud storage.

Specify settings to connect to Microsoft OneDrive:

1. Click **Click to sign in to Microsoft OneDrive**.
2. In the **Sign in to your account** window, follow instructions to specify credentials of the Microsoft OneDrive account that has access to the storage where the backup file resides and click **Sign in**.

If you want to change settings to connect to Microsoft OneDrive, click the **Click to sign out** link and repeat steps 1–2 to specify another account.



Step 7. Select Backup

The **Backup** step of the wizard is available if you have chosen to restore data from a backup file that resides in a remote location — in a network shared folder, on a Veeam backup repository or Veeam Cloud Connect Repository.

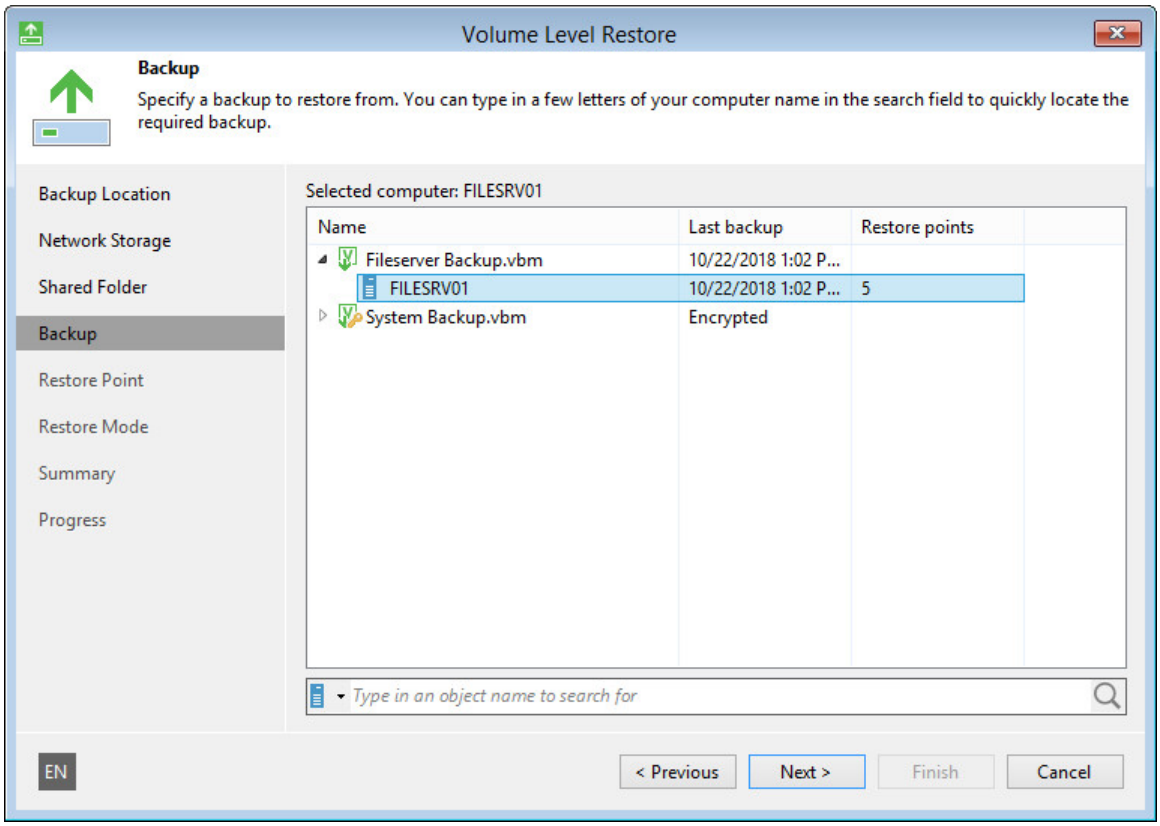
From the list of backups, select a backup from which you want to recover data. To quickly find the necessary backup, use the search field at the bottom of the window: enter a backup name or a part of it in the search field and click the **Start search** button on the right or press **[ENTER]**.

In the list of backups, Veeam Agent for Microsoft Windows displays only those backups that meet the following criteria:

1. Backups created at the volume level. File-level backups are not displayed.
2. [For backup repository target] Backups accessible by the user whose credentials are specified at the Backup Server step of the wizard:
    - If you specify credentials for the user who has access to the backup repository, the list of backups will include only backups created by this user.
    - If you specify credentials for the Backup Administrator on the backup server, the list of backups will include all Veeam Agent backups stored on the backup repository.
3. [For cloud repository target] Backups accessible by the user whose credentials are specified at the Credentials step of the wizard:
    - If you specify credentials for the tenant account, the list of backups will include backups created by all users who create backups under this tenant account and its subtenant accounts.
    - If you specify credentials for the subtenant account, the list of backups will include only those Veeam Agent backups that were created under this subtenant account.

NOTE

If you want to restore data from an encrypted backup, and the Veeam Recovery Media from which you booted your computer does not contain encryption keys required to unlock the backup file, you need to provide a password to unlock the encrypted file. To learn more, see Restoring Data from Encrypted Backups.

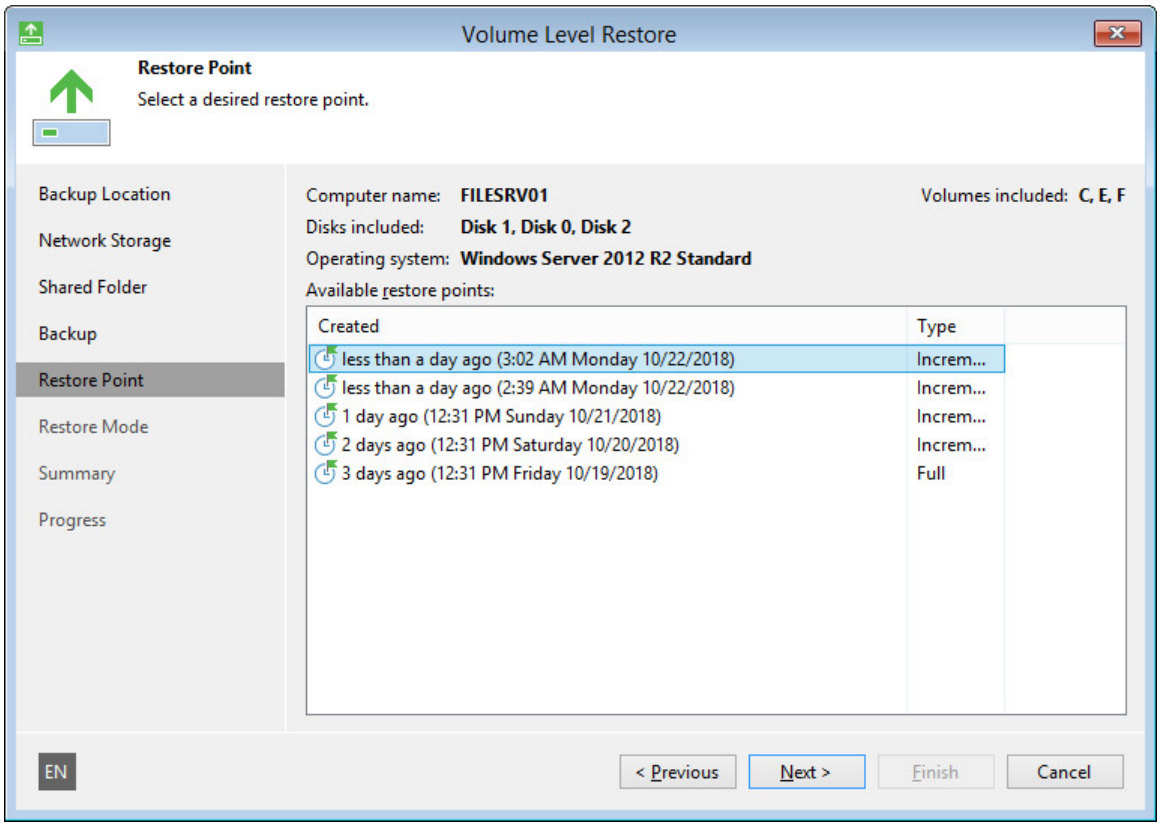


Step 8. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to recover data.

By default, Veeam Agent for Microsoft Windows uses the latest restore point. However, you can select any valid restore point to recover files and folders to a specific point in time.

Veeam Agent displays only restore points of volume-level backups. For example, if you have run 2 job sessions to create a backup of all computer volumes and then changed the backup scope to file-level backup, Veeam Agent will display only 2 restore points in the list.



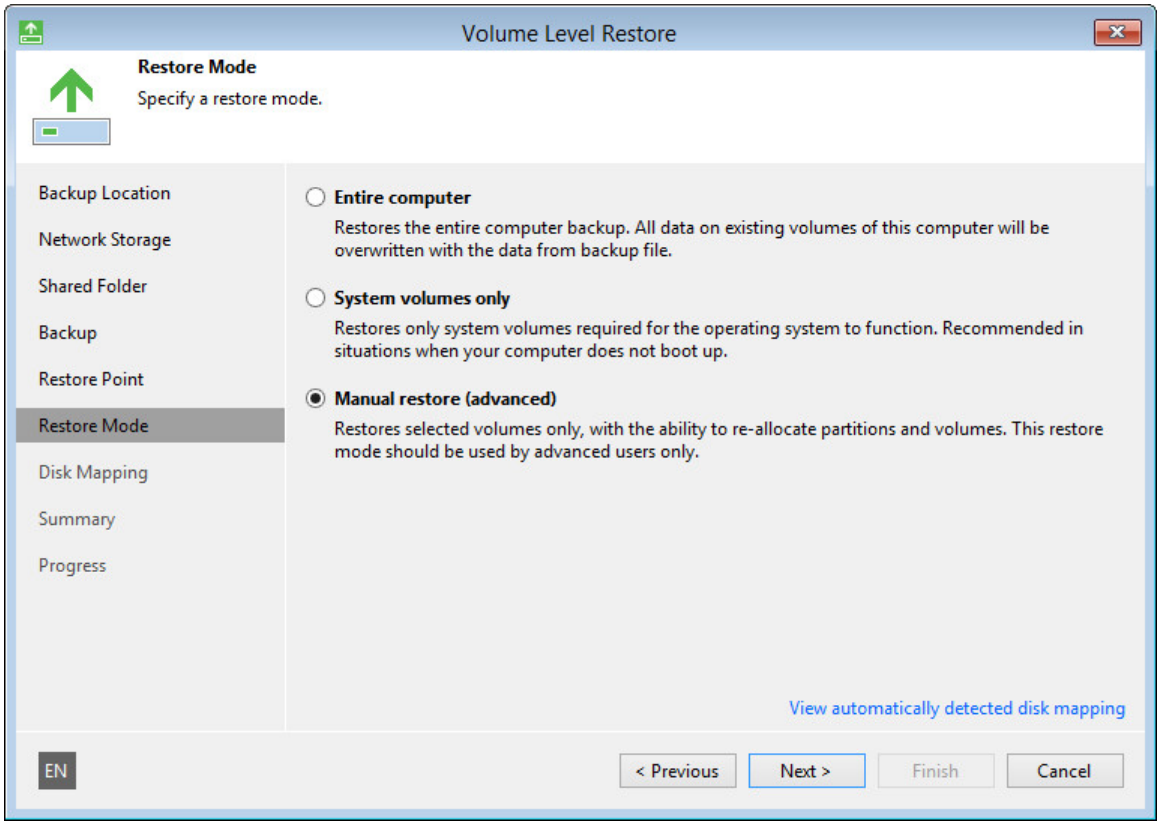
Step 9. Select Data Restore Mode

At the **Restore Mode** step of the wizard, select the data restore mode:

- **Entire computer** — select this option if you want to restore the whole system image of your computer. In this case, Veeam Agent for Microsoft Windows will attempt to map volumes from the backup to existing computer volumes and will overwrite existing data with data restored from the backup.
- **System volumes only** — select this option if you want to restore only system state data and the system volume (volume on which the Microsoft OS is installed). In this case, Veeam Agent for Microsoft Windows will restore the Microsoft Windows system partition and boot partition from the backup to your computer. For GPT disks on Microsoft Windows 8.1, 10, 2012 2012 R2, 2016 and 2019 Veeam Agent for Microsoft Windows will additionally restore the recovery partition.
- **Manual restore** — select this option if you want to choose what computer volumes you want to restore and manually allocate disk space on restored volumes. This option is recommended for users who have experience in working with Microsoft Windows disks and partitions.

To view the current disk allocations settings on your computer, at the bottom of the wizard click **View automatically detected disk mapping**.

> **IMPORTANT**
>
> You will not be able to restore data in the *Entire computer* or *System volumes only* mode, if disks on a computer have not enough space to embed volume data from the backup. In this situation, you will be prompted to use the *Manual restore* mode.



Step 10. Map Restored Disks

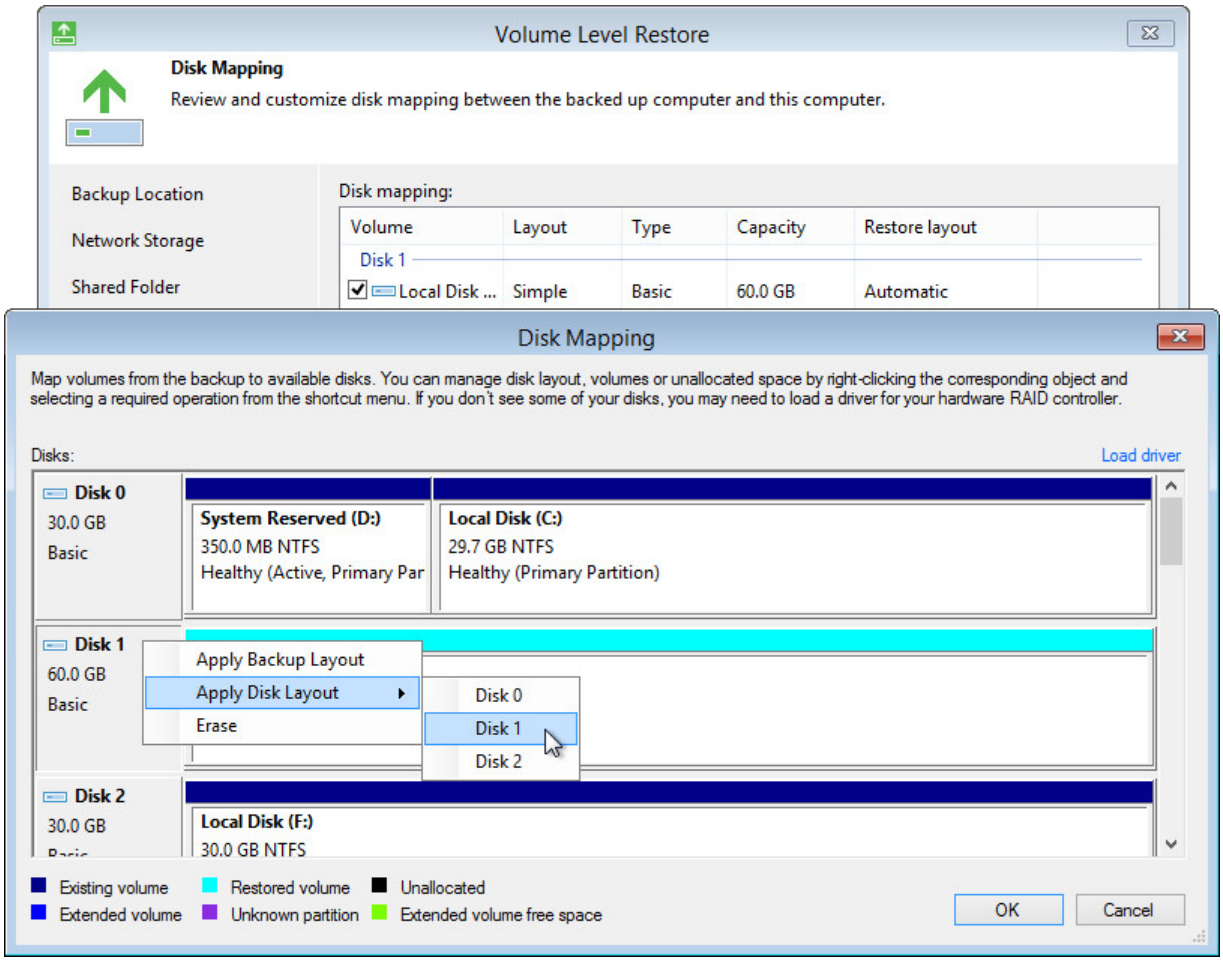
The **Disk Mapping** step of the wizard is available if at the Restore Mode step of the wizard you have chosen to restore data in the *Manual* mode.

You can map volumes that you want to restore from the backup to disks on the target computer.

> **IMPORTANT**
>
> We strongly recommend that you change disk mapping settings only if you have experience in working with Microsoft Windows disks and partitions. If you make a mistake, your computer data may get corrupted.

To map volumes:

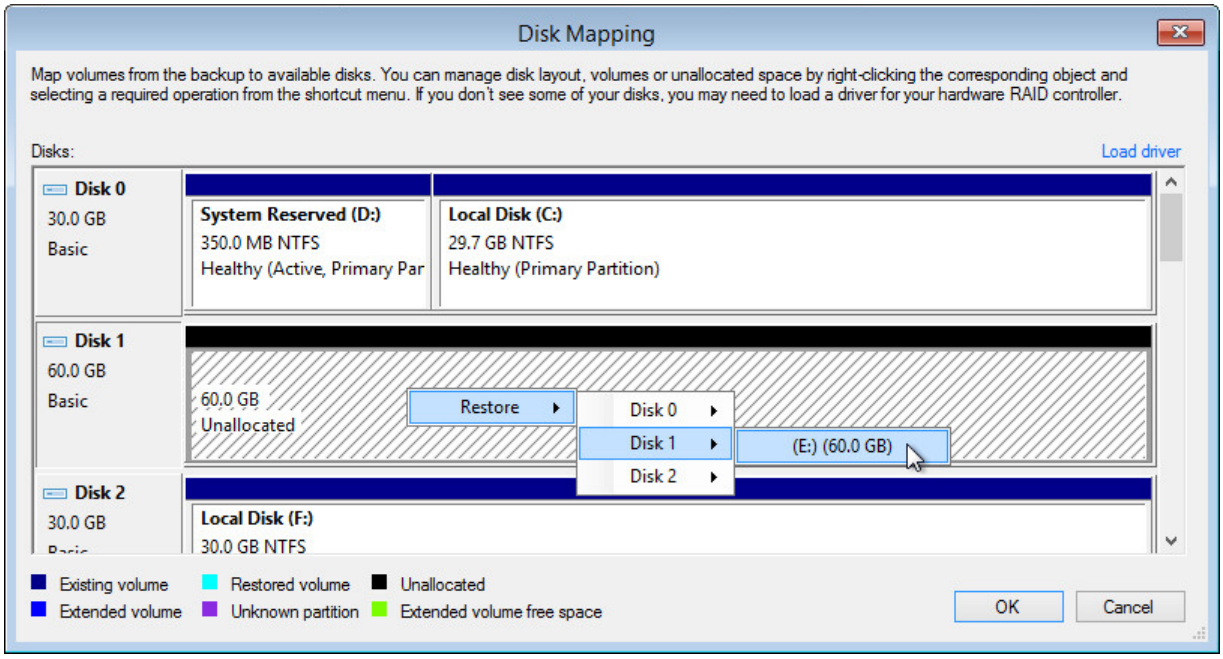
1. Select check boxes next to volumes that you want to restore from the backup.
2. [For restore to a new location] By default, Veeam Agent for Microsoft Windows restores all volumes to their initial location. If the initial location is unavailable, a volume is restored to a disk of the same or larger size. To map the restored volume to another computer disk, at the bottom of the wizard click **Customize disk mapping**. In the **Disk Mapping** window, specify how volumes must be restored:
   - Right-click the target disk on the left and select the necessary disk layout:
     - **Apply Backup Layout** — select this option if you want to apply to disk the settings that were used on your computer at the moment when you performed backup.
     - **Apply Disk Layout** — select this option if you want to apply to the current disk settings of another disk.
     - **Erase** — select this option if you want to discard the current disk settings.



- Right-click unallocated disk space in the disk area on the right and select what volume from the backup you want to place on this computer disk.

  If you restore to a dynamic disk, you can specify the type and size of the target volume before mapping. After you select the volume to place on the disk, you will pass to the Allocate Volume window.

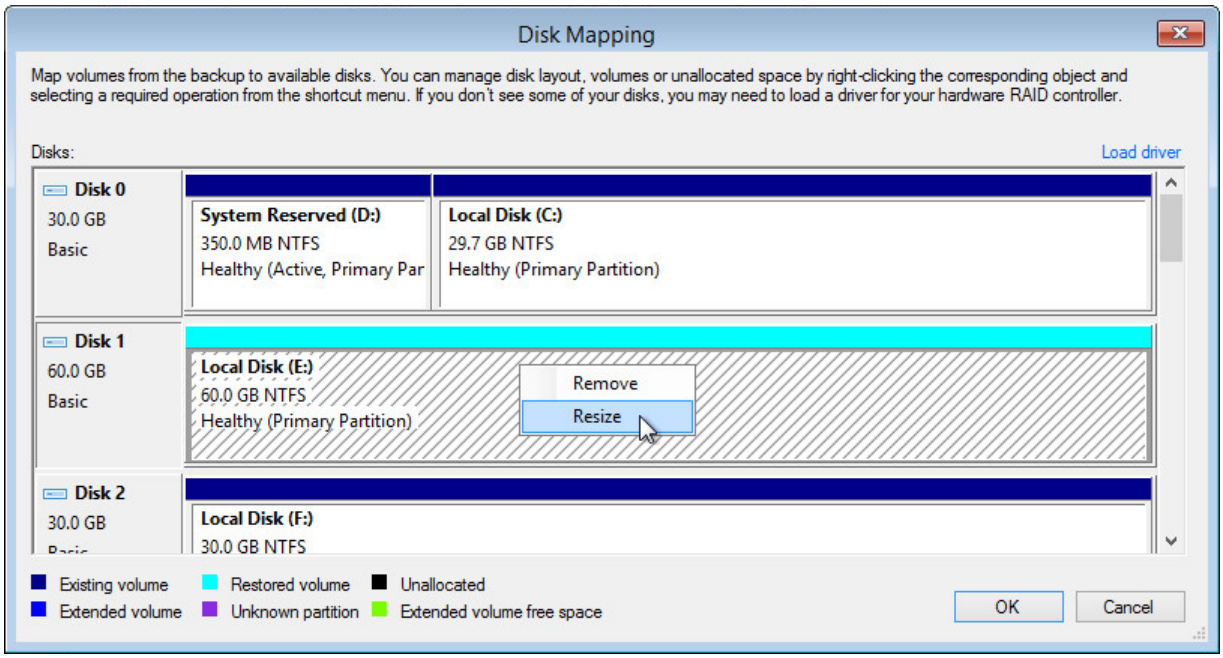
  If you want to change disk layout configured by Veeam Agent for Microsoft Windows, right-click an automatically mapped volume and select **Remove**. You will be able to use the released space for mapping volumes in your own order.



3. [For restore with volume resize] You can resize a volume mapped by Veeam Agent for Microsoft Windows to a target computer disk. To resize a volume, right-click it in the **Disk Mapping** window and select **Resize**. With this option selected, you will pass to the Volume Resize window.

NOTE

If you map a backup volume that is larger than the amount of available space on the target disk, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume. After you agree and click **OK**, Veeam Agent for Microsoft Windows will prepare to shrink the volume to the size of available disk space.



Installing Storage Adapter Drivers

A computer disk may not be available in the list of disks. This can happen in two situations:

- The driver for the storage adapter is included in the Veeam Recovery Media but failed to be installed automatically for some reason.
- The driver for the storage adapter is not included in the Veeam Recovery Media.

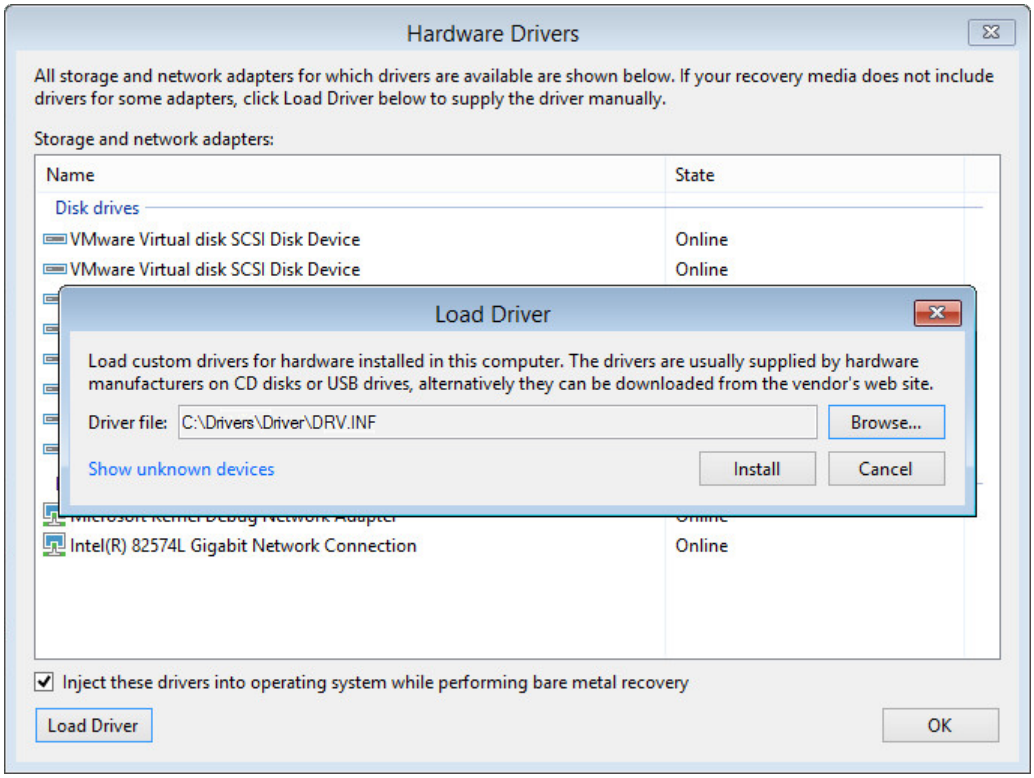
To install drivers that were included in the Veeam Recovery Media:

1. At the **Disk Mapping** step of the wizard, click **Load driver**.
2. In the **Hardware Drivers** window, select the necessary device.

   If you do not want to save drivers for listed devices to the restored operating system, clear the **Inject these drivers into operating system while performing bare metal recovery** check box.
3. Click the **Install** link next to the selected device.

To install drivers that were not included in the Veeam Recovery Media:

1. At the **Disk Mapping** step of the wizard, click **Load driver**.
2. At the bottom of the **Hardware Drivers** window, click the **Load Driver** button and select the INF file in the driver package folder. You can also click the **Show unknown devices** link to see a list of all existing devices without drivers. This information may help you to identify the exact device for which you need to install the driver.
3. Click **Install**.



Allocating Restored Volumes to Dynamic Disks

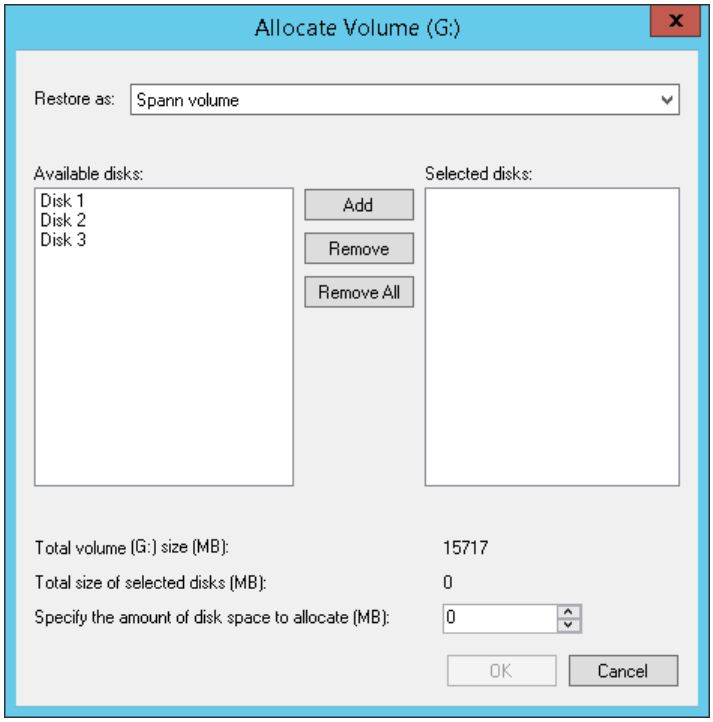
At the **Disk Mapping** step of the wizard, you can set the necessary type and size for the restored volumes. A volume will get the specified type and size during the process of data restore.

You can do this only if all the following conditions are met:

- You restore data to a dynamic disk. To learn more about dynamic volume types, see this Microsoft article.
- At the **Restore Mode** step of the wizard, you have selected the **Manual restore** option.
- At the **Disk Mapping** step of the wizard, you have clicked the **Customize disk mapping** link.

To set volume type and size:

1. Right-click unallocated space of a dynamic disk **and** select what volume from the backup you want to place on this disk. Veeam Agent for Microsoft Windows will prompt you to set the restored volume type and size.
2. In the **Allocate** window, specify the restored volume type:
   - If you want to restore the volume as a simple volume, from the **Restore as** drop-down list, select the **Simple volume** option and click **OK**.
   - If you want to restore the volume as a spanned, striped, mirrored, or parity volume, from the **Restore as** drop-down list, select the desired option and specify the following settings:
     i. In the **Available disks** list, select volume extents you want to restore. You can add or delete volume extents using the **Add**, **Remove**, and **Remove All** buttons.
     ii. In the **Specify the amount of disk space to allocate** field, specify the total size of restored extents.
     iii. Click **OK**.



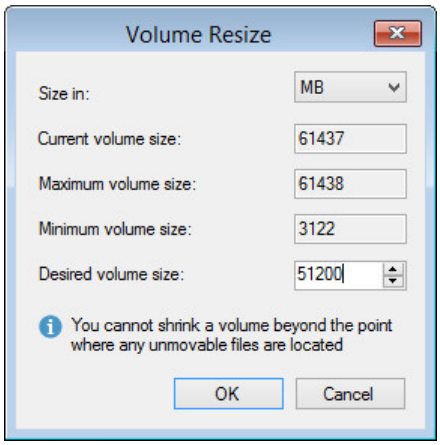
Step 11. Resize Restored Volumes

At the **Disk Mapping** step of the wizard you can set the necessary size for the restored volumes. You can resize a volume if you have chosen to restore data in the *Manual* mode and customize disk layout. A volume will be shrunk or extended to the specified size during the process of data restore.

> **NOTE**
>
> By default, Veeam Agent for Microsoft Windows displays volume size in megabytes (MB). This allows you to specify the desired size for the volume precisely. You can also choose to display volume size in gigabytes (GB). This may be helpful when you need to resize volumes on larger computer disks and want to simplify disk size calculations.
>
> When you use GB as a volume size unit, you can specify volume size with integral numbers, for example, 1 GB, 60 GB or 200 GB, but not 0,8 GB, 60,5 GB or 200,7 GB. However, if the maximum volume size is in fact greater than the displayed value for less than 1 GB, Veeam Agent for Microsoft Windows will automatically add the exceeding amount of disk space to the extended volume. For example, if the maximum volume size is 60,2 GB, Veeam Agent for Microsoft Windows will display this size as 60 GB. When you specify 60 GB as a desired volume size, Veeam Agent for Microsoft Windows will extend the volume to 60,2 GB.

To resize a volume:

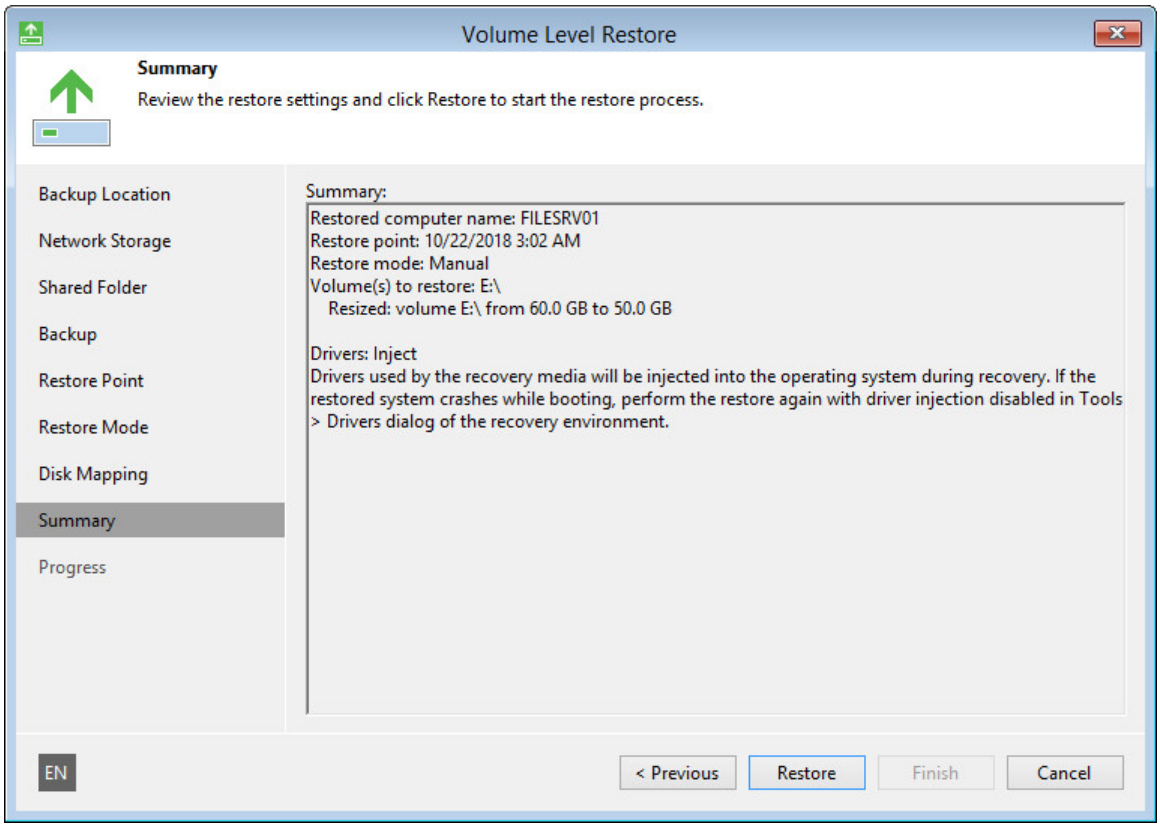
1. Specify a volume you want to resize:
   a. Right-click a restored volume mapped to a target disk and select **Resize**.
   b. [For volume shrink] Right-click unallocated disk space and select what volume from the backup you want to place on the computer disk. If the selected volume is larger than the amount of unallocated disk space, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume.
2. In the **Volume Resize** window, select the volume size unit and specify the desired size for the restored volume.



Step 12. Start Restore Process

At the **Summary** step of the wizard, finalize the recovery process.

1. Review the specified recovery settings.
2. Click **Restore** to start the recovery process. Veeam Agent for Microsoft Windows will perform partition re-allocation operations if necessary, restore the necessary data from the backup and overwrite data on your computer with it.



Using Veeam Agent and Microsoft Windows Tools

When you boot from the Veeam Recovery Media, you can use a set of tools to repair typical causes of unbootable OS, diagnose your computer and perform advanced administration tasks. Veeam Agent for Microsoft Windows offers its native tools and standard Microsoft Windows recovery tools.

> **IMPORTANT**
>
> Veeam Agent for Microsoft Windows includes Microsoft Windows Tools in the Veeam Recovery Media. If some of Microsoft Windows Tools components are missing on the computer, some of Microsoft Windows Tools may not be available when you boot from the Veeam Recovery Media.

To open the tools view, on the **Veeam Recovery Media** screen, click **Tools**. Then choose the necessary tool from the list:

- **Command Prompt** — use this option to start the Microsoft Windows command prompt (`cmd.exe`).
- **Reset Password** — use this option to reset a password for the built-in Administrator account to none. The next time you boot your computer from the hard disk under the Administrator account, you will not have to specify any password. Mind the following:
  - The password reset option does not function on domain controller machines.
  - If the built-in Administrator account is disabled, this account will be enabled by the password reset option.
- **Load Driver** — use this option to load from external sources drivers that are not available on the Veeam Recovery Media. Drivers can be loaded from the computer drive or from a network shared folder.
- **Memory Diagnostic** (Microsoft utility) — use this option to check the system memory of your computer and detect potential problems. The utility can be started during the current work session or when you boot your computer the next time. To learn more, see this Microsoft article.
- **Startup Repair** (Microsoft utility) — use this option to fix system problems that may prevent Microsoft Windows from starting, for example, missing and damaged system files or the corrupted boot sector. To learn more, see this Microsoft article.
- **Export Logs** — use this option to export the Veeam Agent for Microsoft Windows debug logs to a ZIP file and save this file on a removable storage appliance attached to your computer.

NOTE

Do not save the archive file with debug logs on the local disk `X:` of the recovery image OS. This local disk is a temporary storage that will be automatically deleted after you finish working with the Veeam Recovery Media.



Using Microsoft Windows Recovery Environment

If you have a Microsoft system image on the computer drive or a DVD archive with Microsoft system images, you can recover your computer using the Microsoft Windows System Image Recovery tool.

To access the Microsoft Windows System Image Recovery tool, on the **Veeam Recovery Media** screen, click **Windows Recovery Environment**.

The process of recovery does not differ from the process performed in Microsoft Windows. To learn more, see this Microsoft article.



Restoring Volumes

You can restore a specific computer volume or all volumes from the volume-level backup.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent for Microsoft Windows will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent for Microsoft Windows will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

Before You Begin

Before you begin the volume-level restore process, check the following prerequisites and limitations.

General

- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders and on backup repositories] You must have access to the target location where the backup file resides.
- [For backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.
- A user account under which you start the restore operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

ReFS Restore

If you restore a ReFS volume from a backup that was made on another machine, the Veeam Agent computer where you perform restore must run the OS that supports the specific ReFS version. For example, you can restore a ReFS 3.x volume only to a computer that runs one of the following OSes:

- Microsoft Windows 10 version 1803 or later
- Microsoft Windows Server 2016 or later

NTFS Restore

- If you restore an NTFS volume from a backup that was made on another machine, and Data Deduplication was enabled for the volume, the Veeam Agent computer where you perform restore must run the same OS version or later as the backed-up machine OS.
- If you restore an NTFS volume from a backup that was made on another machine, and the allocation unit size of the volume was set to 128K or greater, the Veeam Agent computer where you perform restore must run one of the following OSes:
    - Microsoft Windows 10 version 1803 or later
    - Microsoft Windows Server 2019 or later

Volume-Level Restore Limitations

Volume-level restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Microsoft Windows swap file is hosted.
- You cannot restore a volume to the volume where the backup file that you use for restore is located.

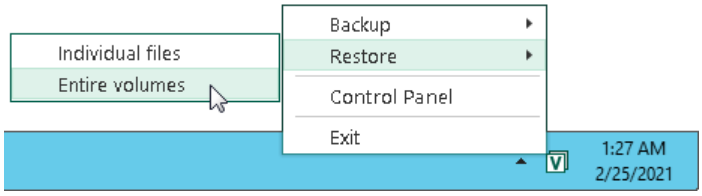
To overcome the first two limitations, you can boot from the recovery image and use the **Volume Level Restore** wizard for volume-level restore. To learn more, see Restoring from Veeam Recovery Media.

Step 1. Launch Volume Level Restore Wizard

To launch the **Volume Level Restore** wizard, do either of the following:

- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Restore** > **Entire volumes**.
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the control panel, click a bar of the necessary backup job session. Click **Restore Volumes** at the bottom of the window.
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the main menu, hover over the name of the job that that created the backup from which you want to restore data, and select **Restore volume**.
- From the Microsoft Windows **Start** menu, select **All Programs** > **Veeam** > **Tools** > **Volume Restore**.

If Veeam Agent for Microsoft Windows automatically detects backups of your computer in the target location, you will pass immediately to the Restore Point step of the wizard.

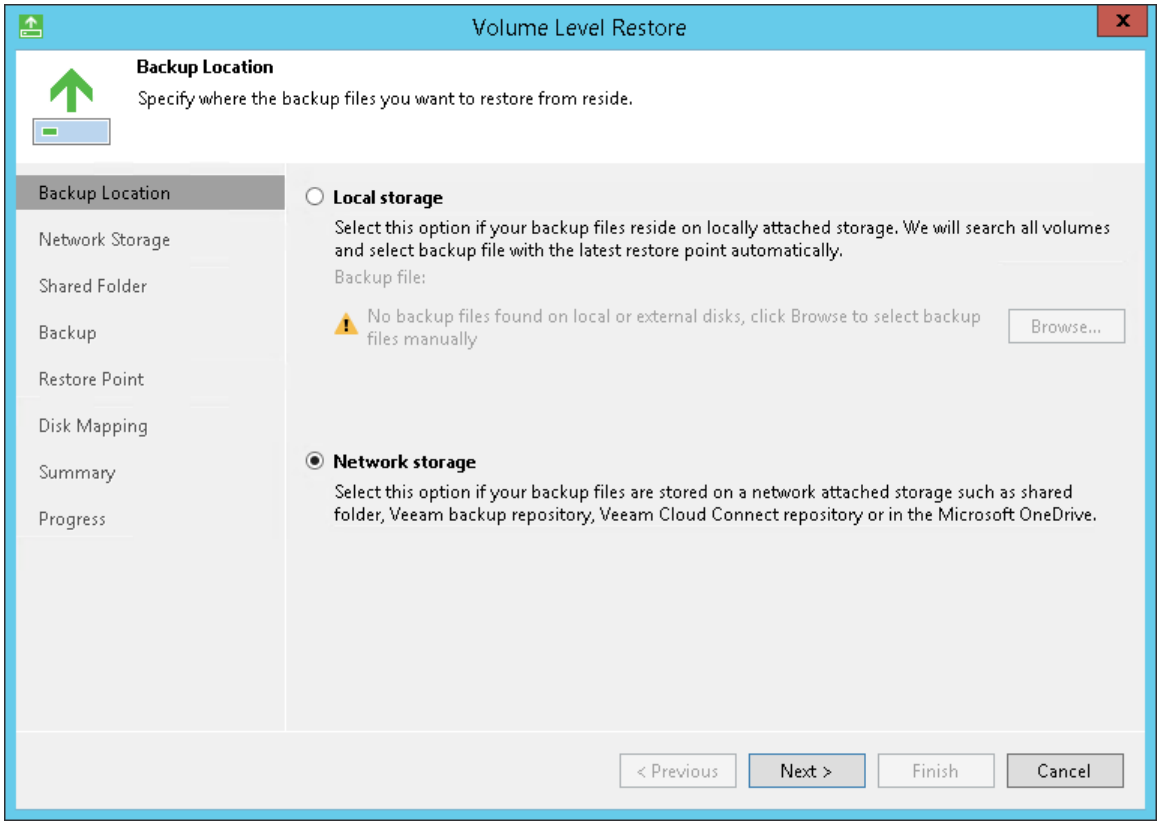


Step 2. Specify Backup File Location

At the **Backup Location** step of the wizard, specify where the backup file that you plan to use for restore resides.

If you configured only one backup job, Veeam Agent automatically locates the latest backup of this backup job on the target storage, and you pass immediately to the Restore Point step of the wizard. If there are several backup jobs configured or Veeam Agent fails to locate the backup for some reason, specify where the backup file resides:

- **Local storage** — select this option if the backup file resides on the computer drive, external drive or removable storage device that is currently connected to your computer. Click **Browse** and select a backup metadata file (VBM).
- **Network storage** — select this option if the backup file resides in a network shared folder, in a Microsoft OneDrive cloud storage, on a backup repository managed by a Veeam backup server or on a cloud repository exposed to you by a Veeam Cloud Connect service provider. In this case, the Volume Level Restore wizard will include additional steps for specifying file location settings.



Step 3. Select Remote Storage Type

The **Network Storage** step of the wizard is available if you have chosen to restore data from a backup file that resides in a remote location — in a network shared folder, on a backup repository or a cloud repository.

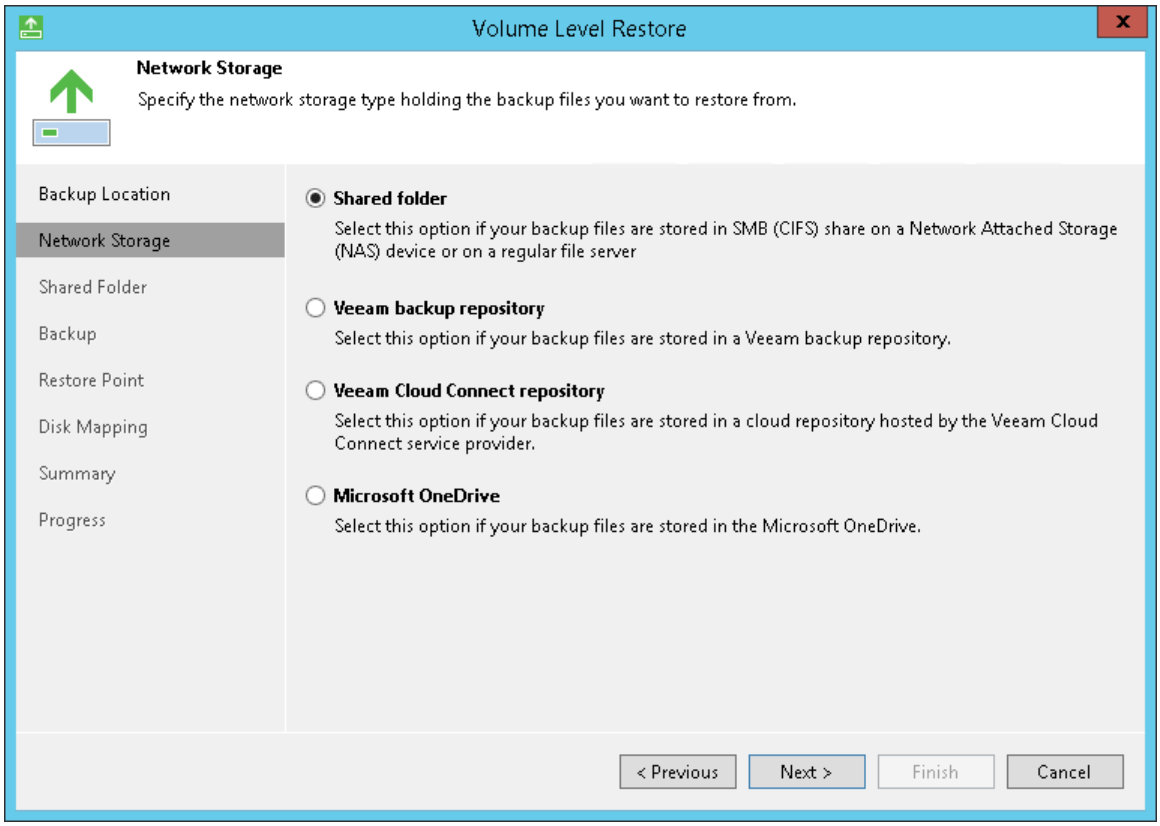
Select where the backup file is located:

- **Shared folder** — select this option if the backup file resides in a network shared folder. With this option selected, you will pass to the Shared Folder step of the wizard.
- **Veeam backup repository** — select this option if the backup file resides on a backup repository managed by a Veeam backup server. With this option selected, you will pass to the Backup Server step of the wizard.
- **Veeam Cloud Connect repository** — select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the Service Provider step of the wizard.
- **Microsoft OneDrive** — select this option if the backup file resides in the Microsoft OneDrive cloud storage. With this option selected, you will pass to the Microsoft OneDrive step of the wizard.

NOTE

Consider the following:

- The Microsoft OneDrive option is not available if the Veeam Agent computer runs a Microsoft Windows Server OS.
- The **Veeam Cloud Connect repository** option is also available at the **Network Storage** step of the wizard. However, for the Free product edition, you cannot store backups on a Veeam Cloud Connect repository. As a result, the restore process for the backup stored on a cloud repository is not described in this guide.



Step 4. Specify Remote Storage Settings

Specify settings for the remote storage that contains a backup file from which you plan to restore data:

- Shared folder settings — if you have selected the **Shared folder** option at the Network Storage step of the wizard.
- Veeam backup repository settings — if you have selected the **Veeam backup repository** option at the Network Storage step of the wizard.
- Veeam Cloud Connect repository settings — if you have selected the **Veeam Cloud Connect repository** option at the Network Storage step of the wizard.
- Microsoft OneDrive settings — if you have selected the **Microsoft OneDrive** option at the Network Storage step of the wizard.

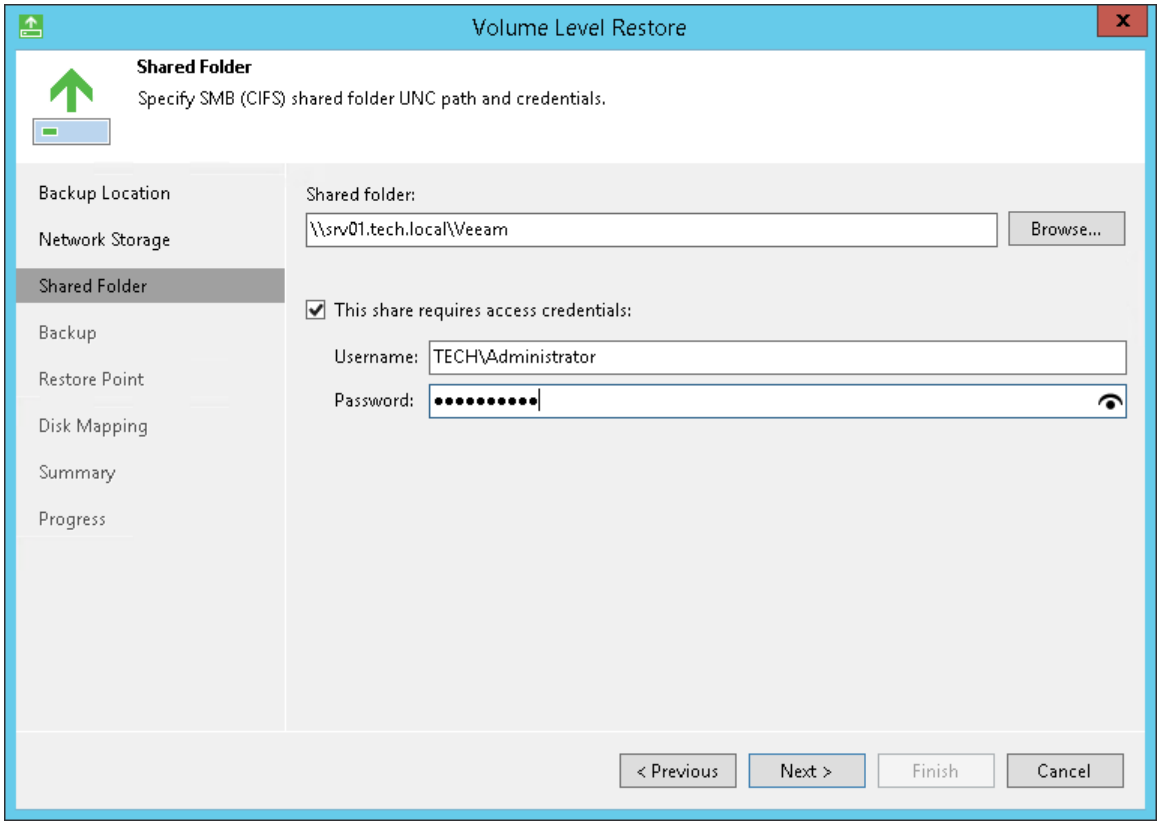
Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have chosen to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. In the **Shared folder** field, enter a UNC name of the network shared folder with the backup file. Keep in mind that the UNC name always starts with two back slashes (\\).
2. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the down-level logon name format. For example, *DOMAIN\UserName* or *HOSTNAME\UserName*.

   If you do not select the **This share requires access credentials** check box, Veeam Agent for Microsoft Windows will connect to the shared folder using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed.
3. To view the entered password, click and hold the eye icon on the right of the **Password** field.



Backup Server Settings

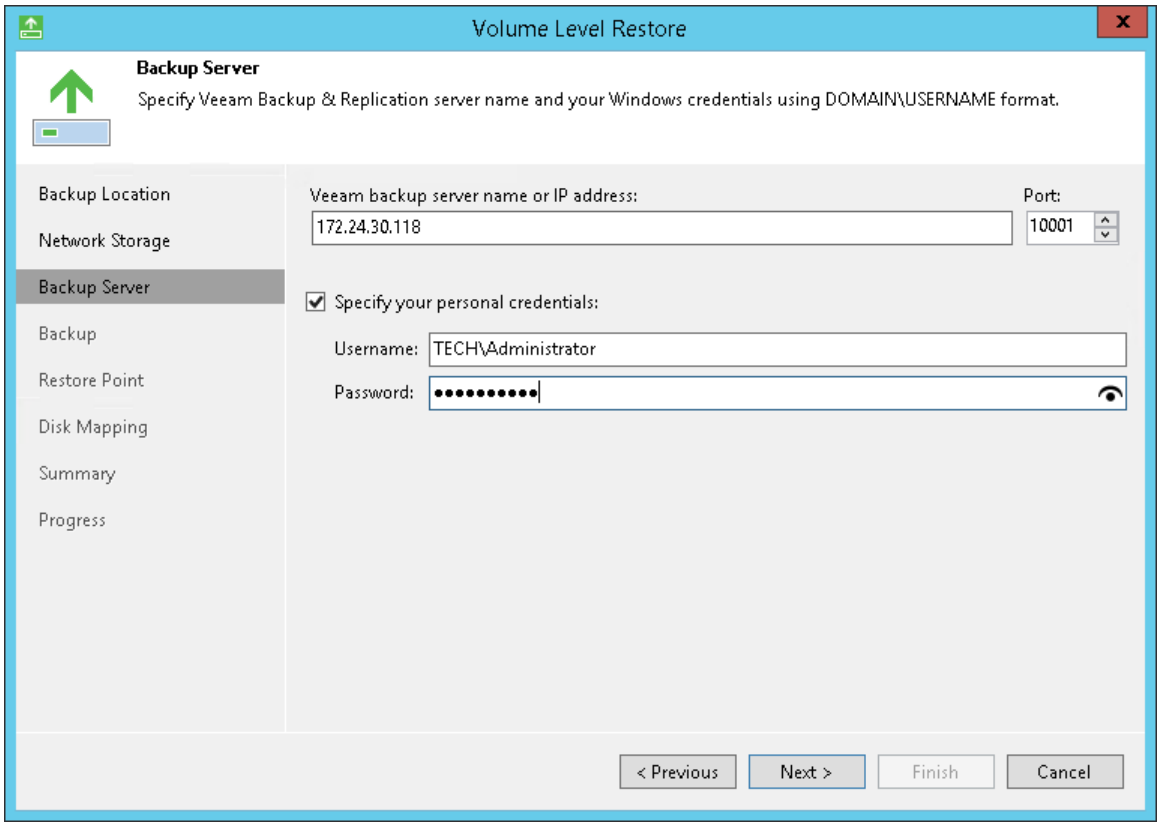
The **Backup Server** step of the wizard is available if you have chosen to restore data from a backup file located on a backup repository.

Specify settings for the Veeam backup server that manages the backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup repository. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. Select the **Specify your personal credentials** check box. In the **Username** and **Password** fields, enter a user name and password of the account that has access to this backup repository. Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see Setting Up User Permissions on Backup Repositories.

   If you do not select the **Specify your personal credentials** check box, Veeam Agent for Microsoft Windows will connect to the backup repository using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can add the computer account (*DOMAIN\COMPUTERNAME$*) to an AD group and grant access rights on the backup repository to this group.

   Setting access permissions on the backup repository to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). If you have set such permissions on the backup repository, you can omit specifying credentials. However, we recommend this scenario for demo environments only.



IMPORTANT

If you specify a DNS name of the Veeam backup server, make sure that the Veeam backup server name is resolved into IPv4 address on the machine where Veeam Agent for Microsoft Windows is installed. The Veeam Backup Service in Veeam Backup & Replication listens on IPv4 addresses only. If the Veeam backup server name is resolved into IPv6 address, Veeam Agent for Microsoft Windows will fail to connect to the Veeam backup server.

Service Provider Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

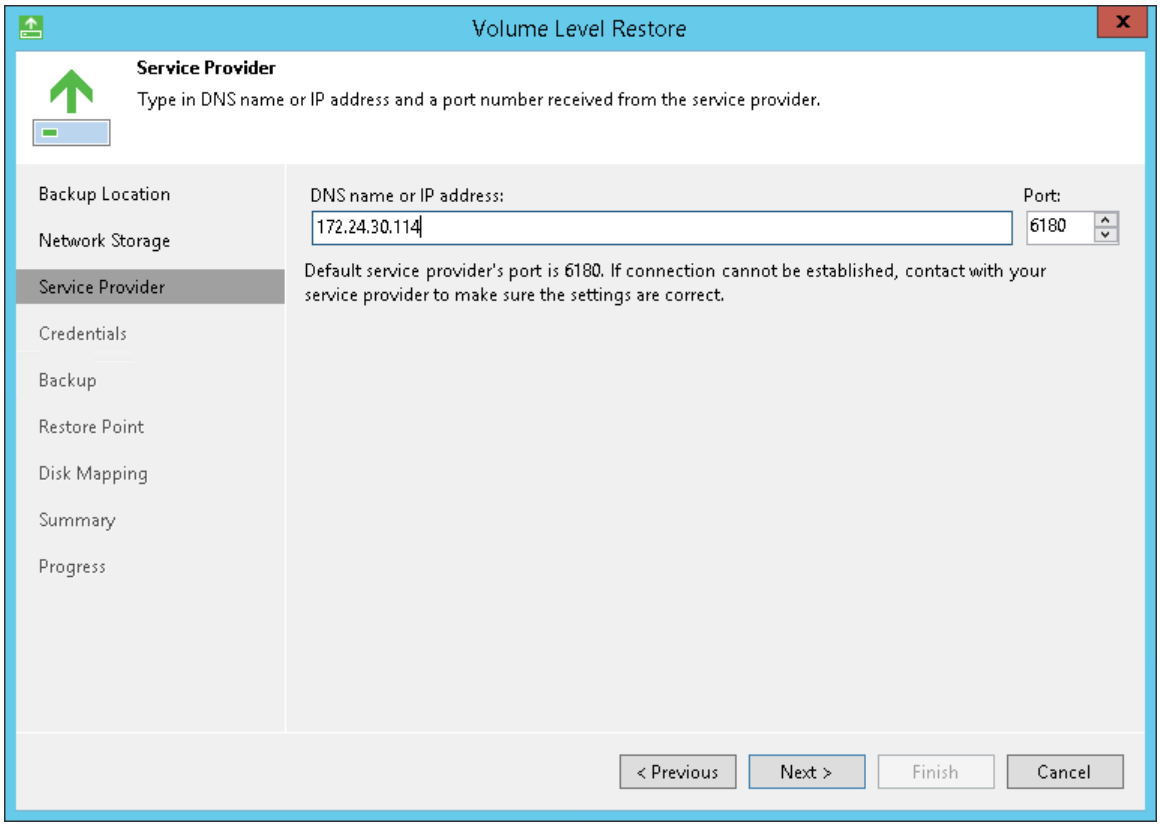
1. Specify service provider settings.
2. Verify the TLS certificate and specify user account settings.

Specifying Service Provider Settings

The **Service provider** step of the wizard is available if you have chosen to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.



Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the service provider.

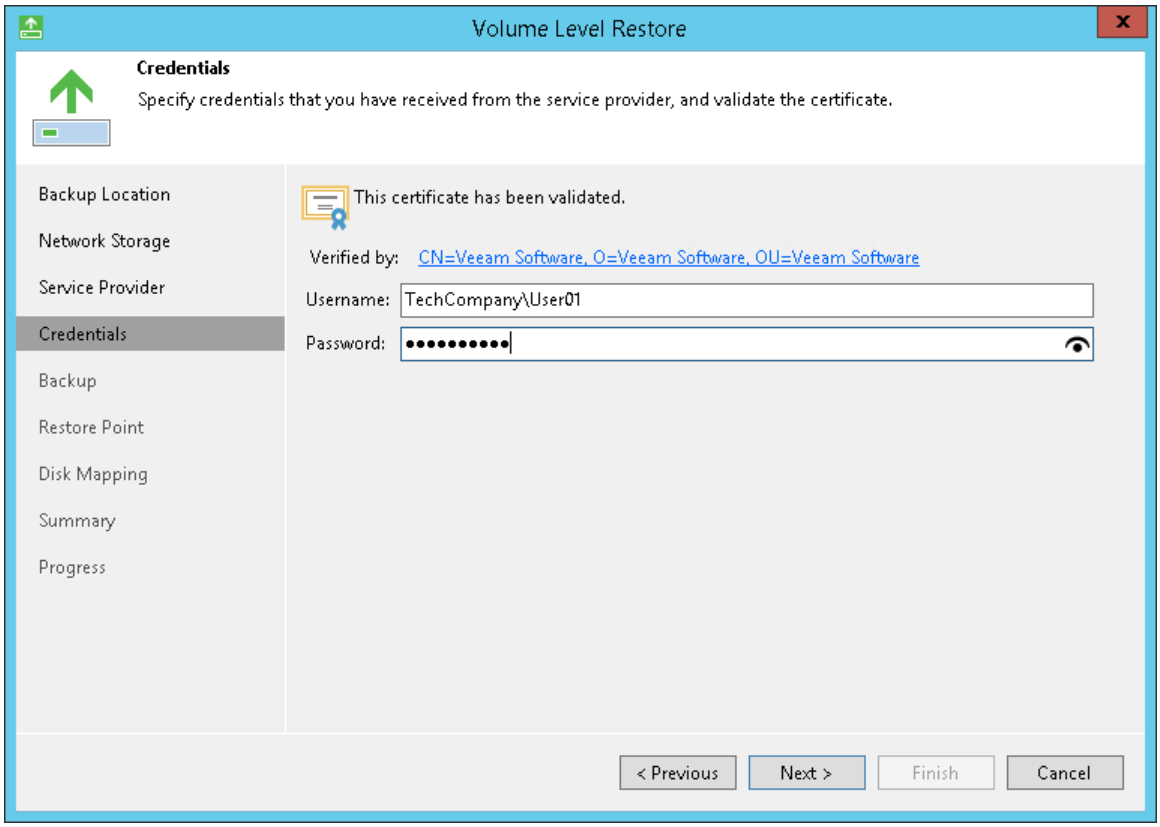
1. At the top of the wizard window, Veeam Agent for Microsoft Windows displays information about the TLS certificate obtained from the SP side. You can view the certificate settings and verify the TLS certificate.

   TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

   - To view the TLS certificate, click the certificate link.
   - To verify if the TLS certificate with a thumbprint, copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Fingerprint for certificate verification** field. Click **Verify**. Veeam Agent for Microsoft Windows will check if the thumbprint you enter matches the thumbprint of the obtained TLS certificate.

2. In the **Username** field, enter the user name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The user name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.



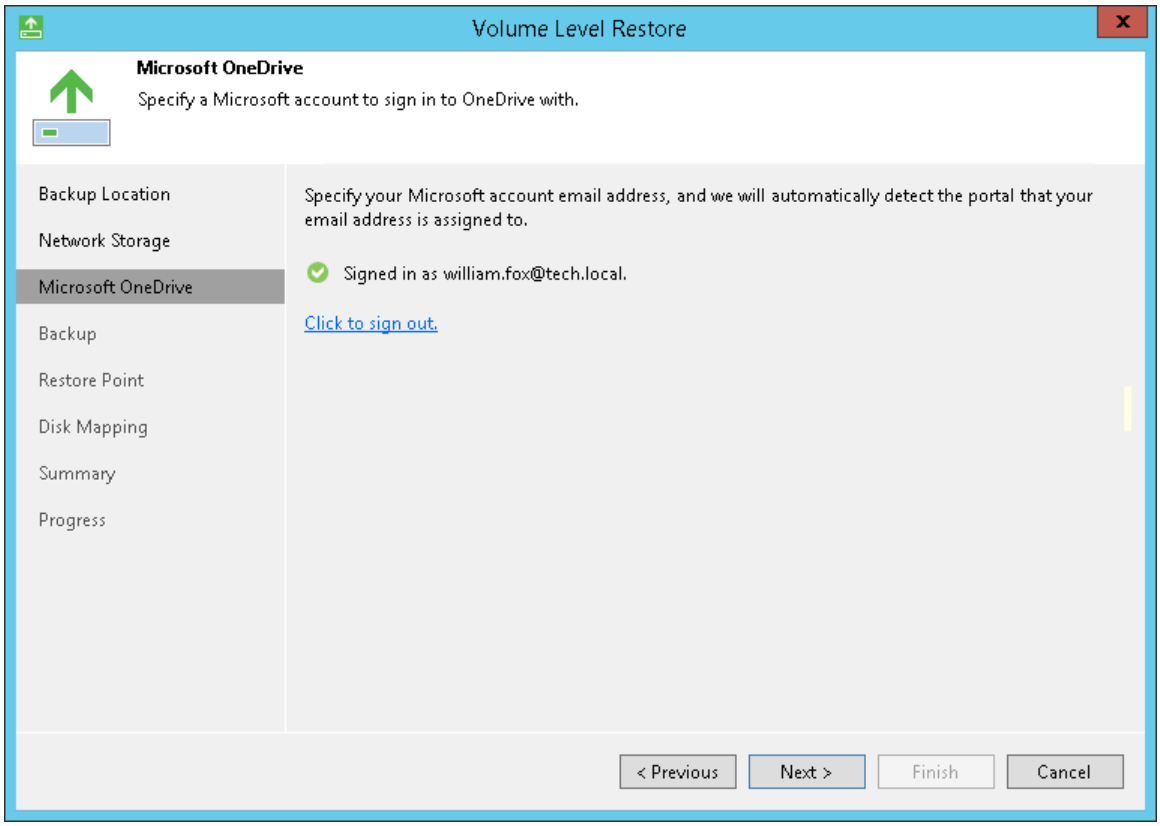
Microsoft OneDrive Settings

The **Microsoft OneDrive** step of the wizard is available if you have chosen to restore data from a backup file located in the Microsoft OneDrive cloud storage.

Specify settings to connect to Microsoft OneDrive:

1. Click **Click to sign in to Microsoft OneDrive**.
2. In the **Microsoft OneDrive** window, follow instructions to specify credentials of the Microsoft OneDrive account that has access to the storage where the backup file resides and click **Sign in**.

If you want to change settings to connect to Microsoft OneDrive, click the **Click to sign out** link and repeat steps 1-2 to specify another account.



Step 5. Select Backup

The **Backup** step of the wizard is available if you have chosen to restore data from a backup file located in a network shared folder or on a backup repository.

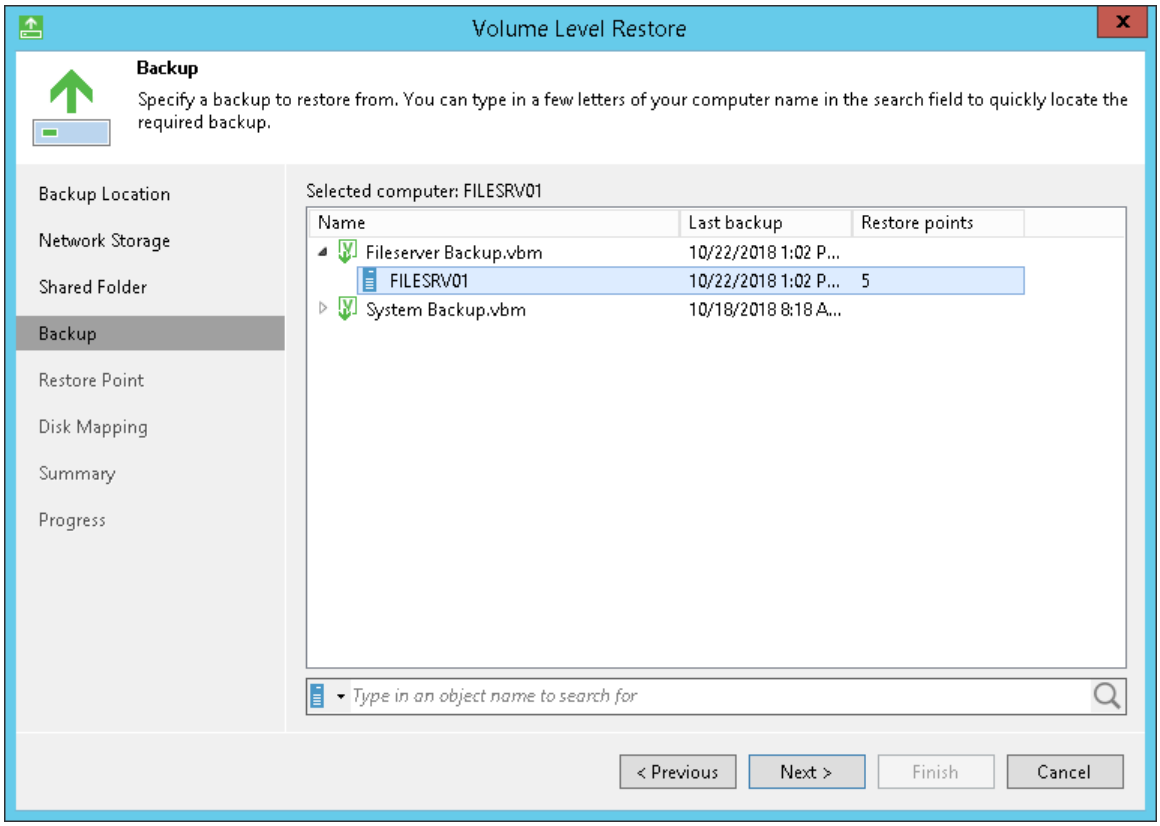
From the list of backups, select a backup from which you want to recover data. To quickly find the necessary backup, use the search field at the bottom of the window: enter a backup name or a part of it in the search field and click the **Start search** button on the right or press **[ENTER]**.

In the list of backups, Veeam Agent for Microsoft Windows displays only those backups that meet the following criteria:

1. Backups created at the volume level. File-level backups are not displayed.
2. [For backup repository target] Backups accessible by the user whose credentials are specified at the Backup Server step of the wizard:
   - If you specify credentials for the user who has access to the backup repository, the list of backups will include only backups created by this user.
   - If you specify credentials for the user who is assigned the *Backup Administrator* or *Restore Operator* role on the backup server, the list of backups will include all Veeam Agent backups stored on the backup repository.
3. [For cloud repository target] Backups accessible by the user whose credentials are specified at the Credentials step of the wizard:
   - If you specify credentials for the tenant account, the list of backups will include backups created by all users who create backups under this account.
   - If you specify credentials for the subtenant account, the list of backups will include only those Veeam Agent backups that were created under this subtenant account.

NOTE

If you restore data from an encrypted backup that was created on another Veeam Agent computer, you need to provide a password to unlock the encrypted file. To learn more, see Restoring Data from Encrypted Backups.

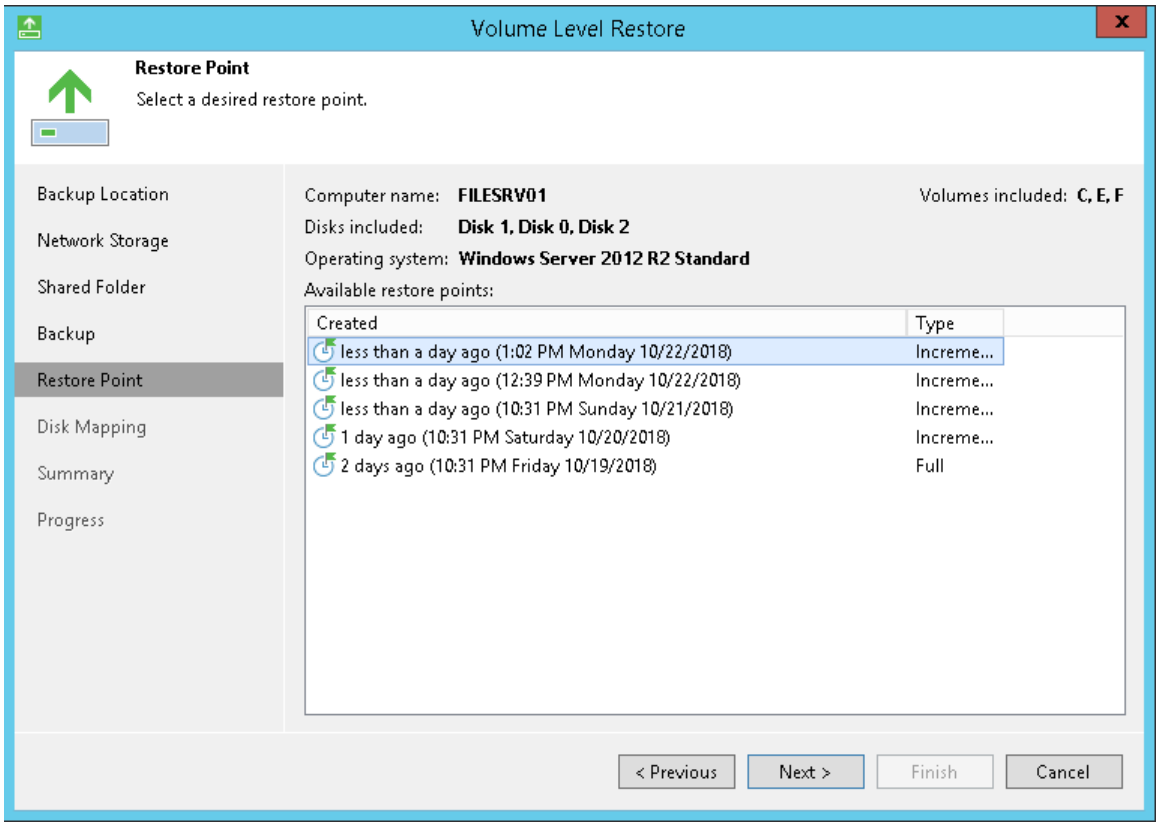


Step 6. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to recover data.

By default, Veeam Agent for Microsoft Windows uses the latest restore point. However, you can select any valid restore point to recover volumes to a specific point in time.

Veeam Agent for Microsoft Windows displays restore points for volume-level backups only. For example, if you have run 3 job sessions to create a backup of all computer volumes and then changed the backup scope to file-level backup, Veeam Agent for Microsoft Windows will display only 3 restore points in the list.



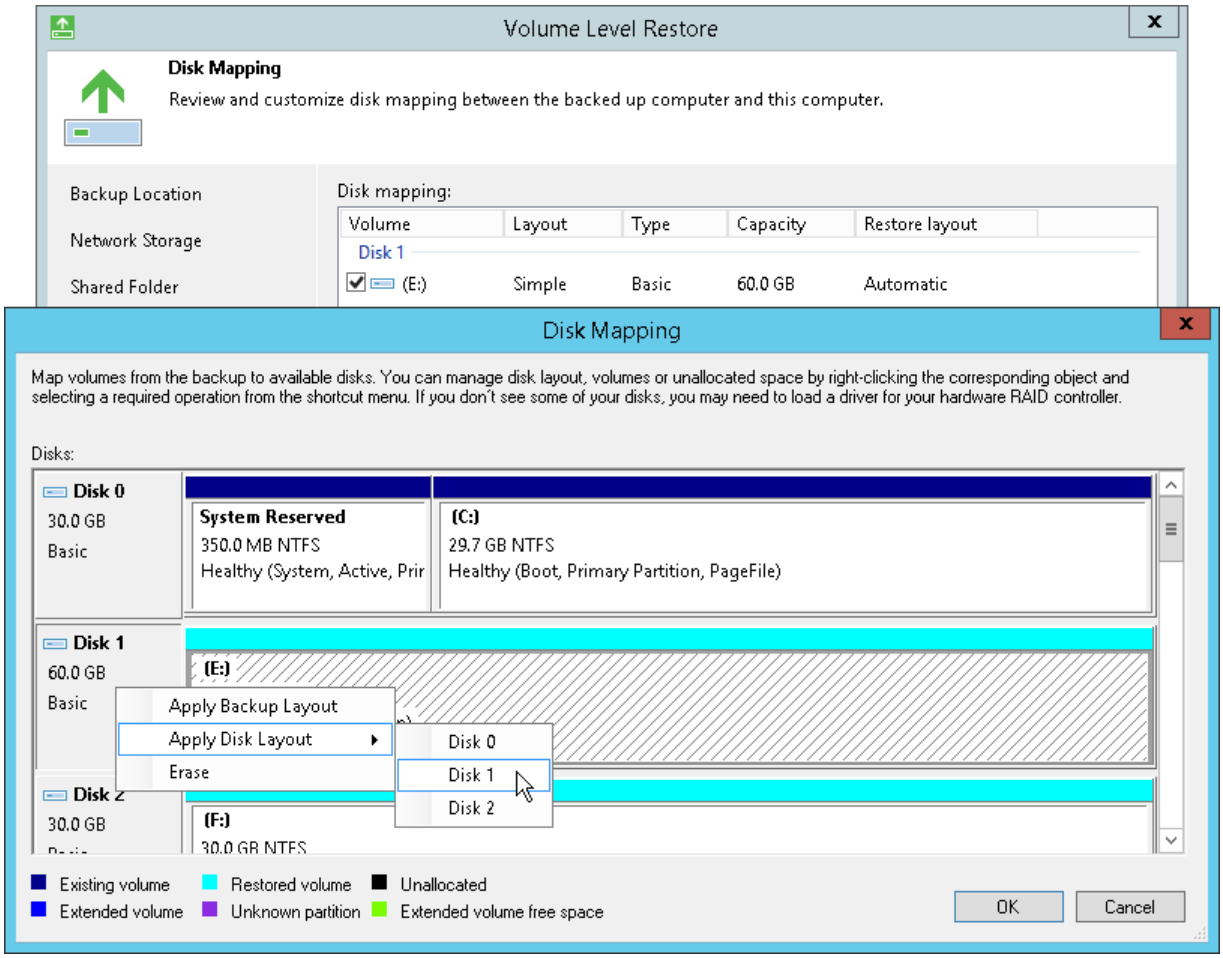
Step 7. Map Restored Disks

At the **Disk Mapping** step of the wizard, select what volume(s) you want to restore and map volumes from the backup to volumes on your computer.

> **IMPORTANT**
>
> We strongly recommend that you change disk mapping settings only if you have experience in working with Microsoft Windows disks and partitions. If you make a mistake, your computer data may get corrupted.

To select volumes to restore:

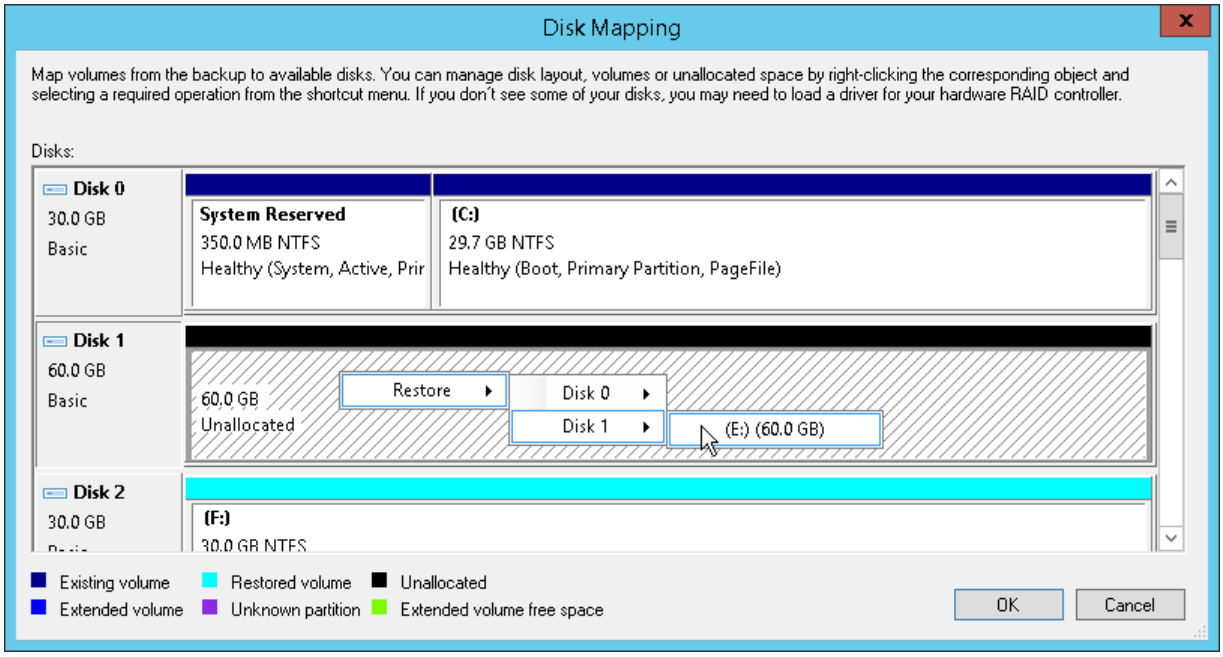
1. Select check boxes next to volumes that you want to restore from the backup.
2. [For restore to a new location] By default, Veeam Agent for Microsoft Windows restores volumes to their initial location. If the initial location is unavailable, a volume is restored to a disk of the same or larger size. To map the restored volume to another computer disk, at the bottom of the wizard click **Customize disk mapping**. In the **Disk Mapping** window, specify how volumes must be restored:
   - Right-click the target disk on the left and select the necessary disk layout:
     - **Apply Backup Layout** — select this option if you want to apply to disk the settings that were used on your computer at the moment when you performed backup.
     - **Apply Disk Layout** — select this option if you want to apply to the current disk settings of another disk.
     - **Erase** — select this option if you want to discard the current disk settings.



- Right-click unallocated disk space in the disk area on the right and select what volume from the backup you want to place on this computer disk.

  If you restore to a dynamic disk, you can specify the type and size of the target volume before mapping. After you select the volume to place on the disk, you will pass to the Allocate Volume window.

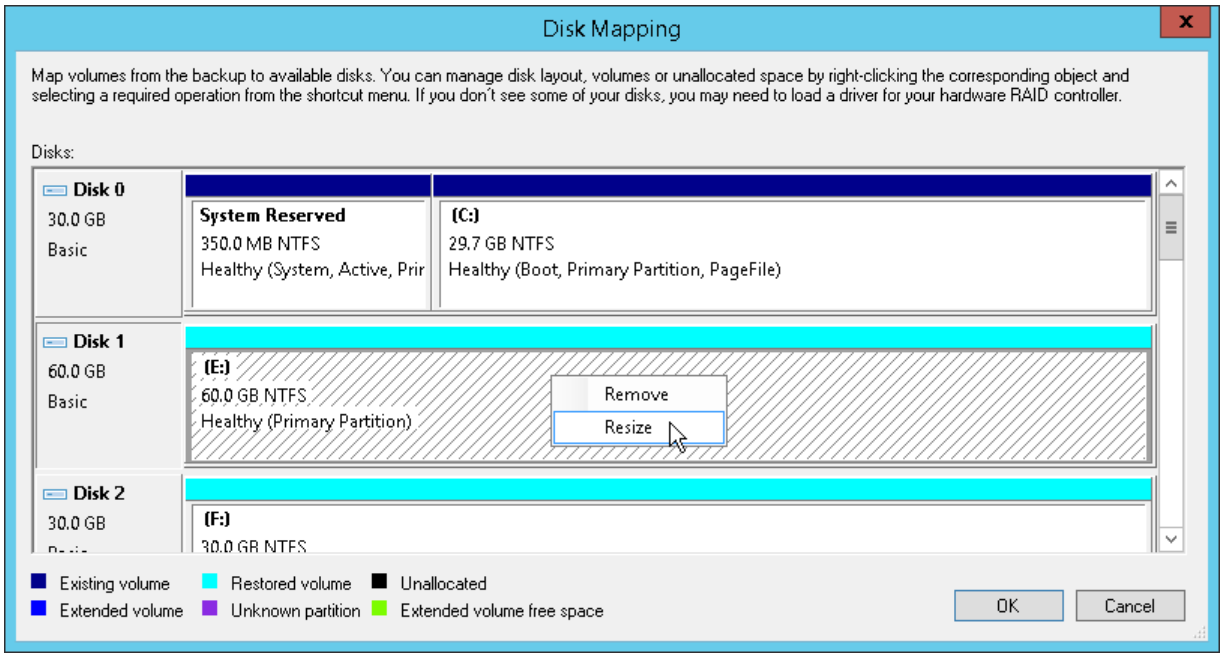
  If you want to change disk layout configured by Veeam Agent for Microsoft Windows, right-click an automatically mapped volume and select **Remove**. You will be able to use the released space for mapping volumes in your own order.



3. [For restore with volume resize] You can resize a volume mapped by Veeam Agent for Microsoft Windows to a target computer disk. To resize a volume, right-click it in the **Disk Mapping** window and select **Resize**. With this option selected, you will pass to the Volume Resize window.

NOTE

If you map a backup volume that is larger than the amount of available space on the target disk, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume. After you agree and click **OK**, Veeam Agent for Microsoft Windows will prepare to shrink the volume to the size of available disk space.



Allocating Restored Volumes to Dynamic Disks

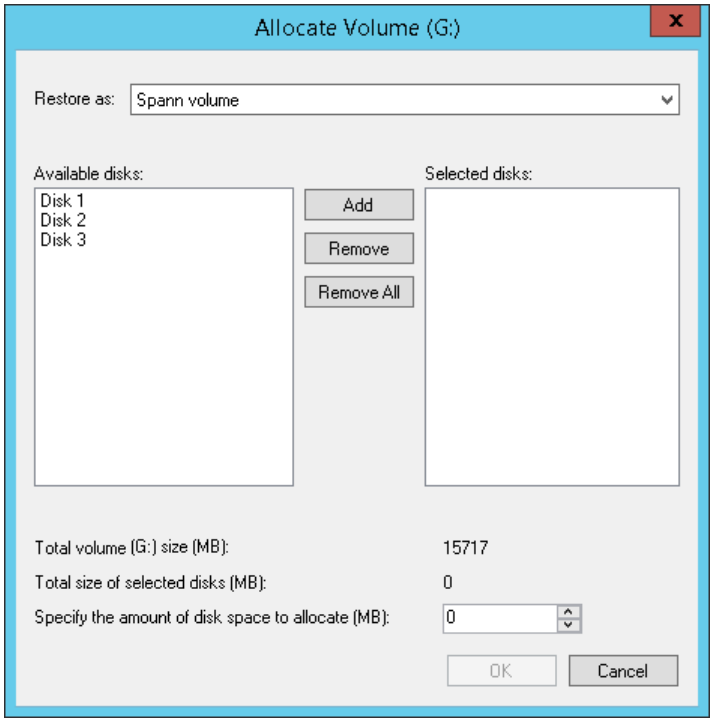
At the **Disk Mapping** step of the wizard, you can set the necessary type and size for the restored volumes. A volume will get the specified type and size during the process of data restore.

You can do this only if all the following conditions are met:

- You restore data to a dynamic disk. To learn more about dynamic volume types, see this Microsoft article.
- At the **Restore Mode** step of the wizard, you have selected the **Manual restore** option.
- At the **Disk Mapping** step of the wizard, you have clicked the **Customize disk mapping** link.

To set volume type and size:

1. Right-click unallocated space of a dynamic disk **and** select what volume from the backup you want to place on this disk. Veeam Agent for Microsoft Windows will prompt you to set the restored volume type and size.
2. In the **Allocate** window, specify the restored volume type:
    - If you want to restore the volume as a simple volume, from the **Restore as** drop-down list, select the **Simple volume** option and click **OK**.
    - If you want to restore the volume as a spanned, striped, mirrored, or parity volume, from the **Restore as** drop-down list, select the desired option and specify the following settings:
        i. In the **Available disks** list, select volume extents you want to restore. You can add or delete volume extents using the **Add**, **Remove**, and **Remove All** buttons.
        ii. In the **Specify the amount of disk space to allocate** field, specify the total size of restored extents.
        iii. Click **OK**.



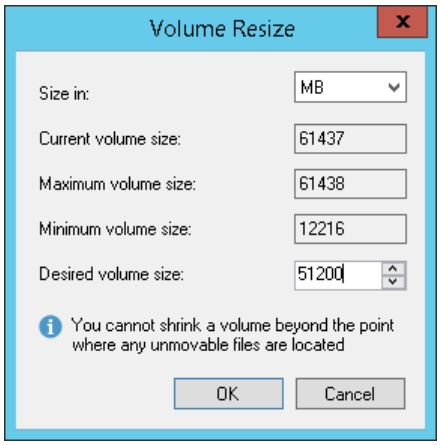
Step 8. Resize Restored Volumes

At the **Disk Mapping** step of the wizard you can set the necessary size for the restored volumes. You can resize a volume if you have chosen to restore data in the *Manual* mode and customize disk layout. A volume will be shrunk or extended to the specified size during the process of data restore.

> **NOTE**
>
> By default, Veeam Agent for Microsoft Windows displays volume size in megabytes (MB). This allows you to specify the desired size for the volume precisely. You can also choose to display volume size in gigabytes (GB). This may be helpful when you need to resize volumes on larger computer disks and want to simplify disk size calculations.
>
> When you use GB as a volume size unit, you can specify volume size with integral numbers, for example, 1 GB, 60 GB or 200 GB, but not 0,8 GB, 60,5 GB or 200,7 GB. However, if the maximum volume size is in fact greater than the displayed value for less than 1 GB, Veeam Agent for Microsoft Windows will automatically add the exceeding amount of disk space to the extended volume. For example, if the maximum volume size is 60,2 GB, Veeam Agent for Microsoft Windows will display this size as 60 GB. When you specify 60 GB as a desired volume size, Veeam Agent for Microsoft Windows will extend the volume to 60,2 GB.

To resize a volume:

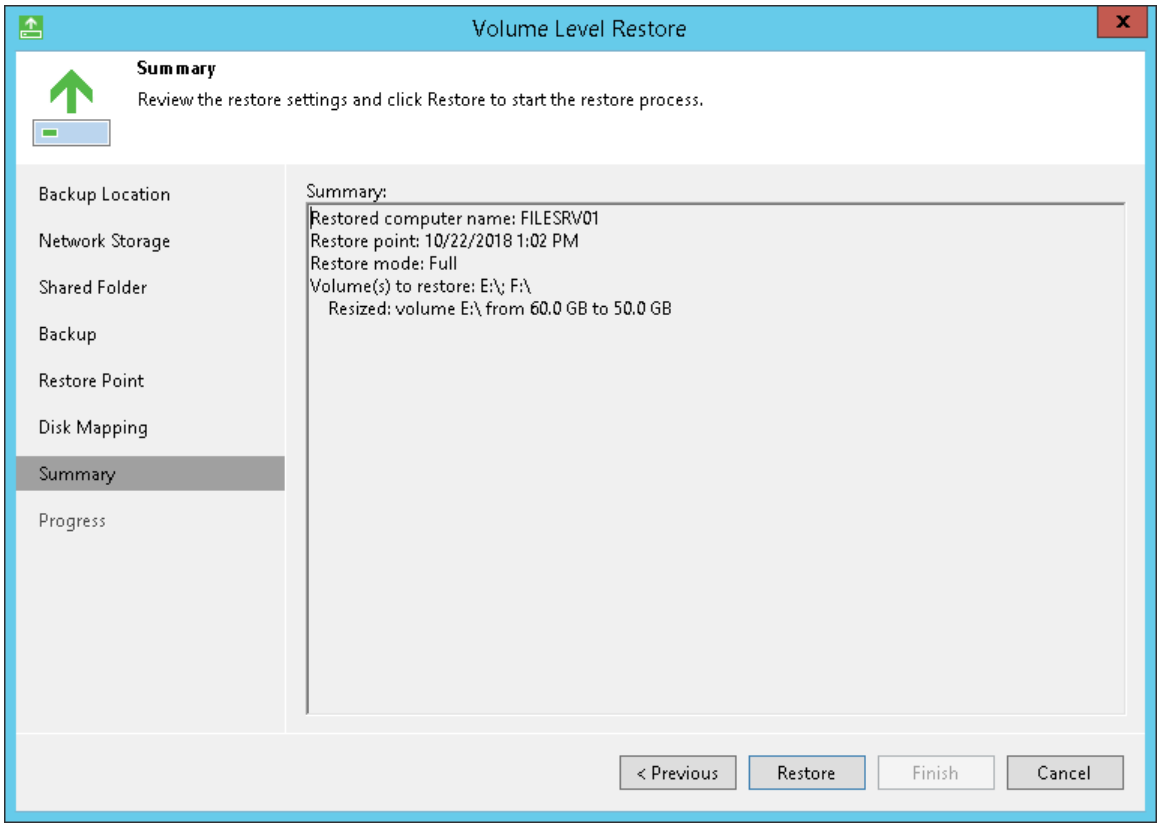
1. Specify a volume you want to resize:
   a. Right-click a restored volume mapped to a target disk and select **Resize**.
   b. [For volume shrink] Right-click unallocated disk space and select what volume from the backup you want to place on the computer disk. If the selected volume is larger than the amount of unallocated disk space, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume.
2. In the **Volume Resize** window, select the volume size unit and specify the desired size for the restored volume.



Step 9. Complete Restore Process

At the **Summary** step of the wizard, complete the procedure of volume-level restore.

1. Review settings of the restore process.
2. Click **Restore** to start the recovery process. Veeam Agent for Microsoft Windows will perform partition re-allocation operations if necessary, restore the necessary volume data from the backup and overwrite volume data on your computer with the restored data.



Restoring Files and Folders

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent for Microsoft Windows publishes the backup content directly into the computer file system and displays it in the Veeam Backup browser. You can restore files and folders to their initial location, copy files and folders to a new location or target applications to restored files and work with them as usual.

Before You Begin

Before you begin the file-level restore process, check the following prerequisites.

General

- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders and on backup repositories] You must have access to the target location where the backup file resides.
- [For backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see Setting Up User Permissions on Backup Repositories.
- A user account under which you start the restore operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

ReFS Restore

If you restore files or folders from a ReFS volume-level backup that was made on another machine, the Veeam Agent computer where you perform restore must run the OS that supports the specific ReFS version. For example, you can restore files from a ReFS 3.x volume only to a computer that runs one of the following OSes:

- Microsoft Windows 10 version 1803 or later
- Microsoft Windows Server 2016 or later

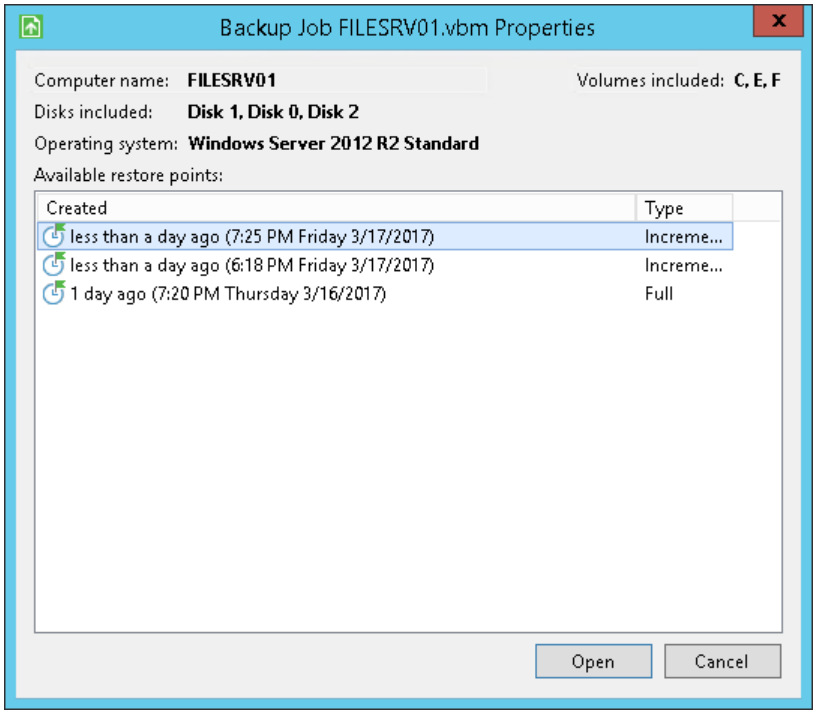
NTFS Restore

- If you restore files or folders from an NTFS volume-level backup that was made on another machine, and Data Deduplication is enabled for some volumes of the backup, the Veeam Agent computer where you perform restore must run the same OS version or later as the backed-up machine OS.
- If you restore files or folders from an NTFS volume-level backup that was made on another machine, and the allocation unit size of volumes is set to 128K or greater, the Veeam Agent computer where you perform restore must run one of the following OSes:
    - Microsoft Windows 10 version 1803 or later
    - Microsoft Windows Server 2019 or later

Step 1. Launch File Level Restore Wizard

To launch the **File Level Restore** wizard, do either of the following:

- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Restore** > **Individual files**.
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the control panel, click a bar of the necessary backup job session. Click **Restore Files** at the bottom of the window. Veeam Agent for Microsoft Windows will automatically publish the backup content into the computer file system and open the Veeam Backup browser.
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the main menu, hover over the name of the job that that created the backup from which you want to restore data, and select **Restore file**.
- From the Microsoft Windows **Start** menu, select **All Programs** > **Veeam** > **File Level Restore**.
- In Microsoft Windows Explorer, double-click the necessary VBK or VBM file or right-click the file and select **Extract**. In the displayed window, select the restore point from which you want to recover files and click **Open**. Veeam Agent for Microsoft Windows will automatically publish the backup content into the computer file system and open the Veeam Backup browser.

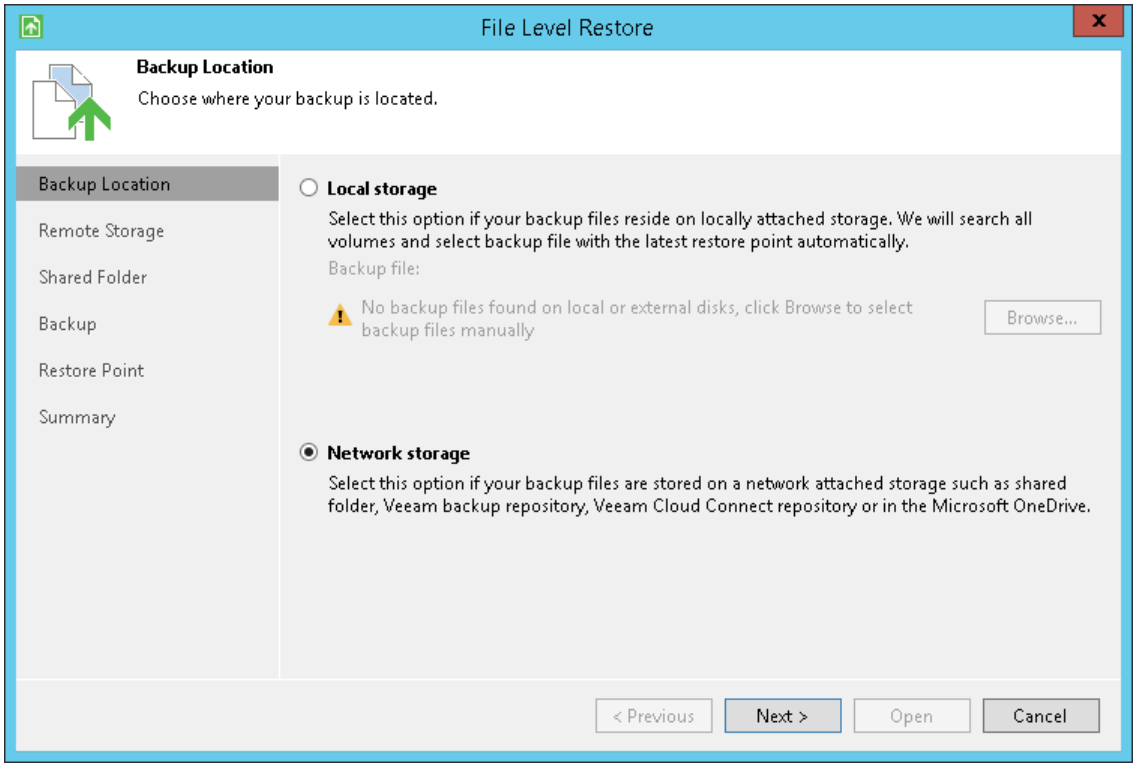


Step 2. Specify Backup File Location

At the **Backup Location** step of the wizard, specify where the backup file that you plan to use for restore resides.

If you configured only one backup job, Veeam Agent automatically locates the latest backup of this backup job on the target storage, and you pass immediately to the Restore Point step of the wizard. If there are several backup jobs configured or Veeam Agent fails to locate the backup for some reason, specify where the backup file resides:

- **Local storage** — select this option if the backup file resides on the computer drive, external drive or removable storage device that is currently connected to your computer. Click **Browse** and select a backup metadata file (VBM).
- **Network storage** — select this option if the backup file resides in a network shared folder, in a Microsoft OneDrive cloud storage, on a backup repository managed by a Veeam backup server or on a cloud repository exposed to you by a Veeam Cloud Connect service provider. In this case, the Veeam Recovery Media wizard will include additional steps for specifying the backup file location settings.



Step 3. Select Remote Storage Type

The **Remote Storage** step of the wizard is available if you have chosen to restore data from a backup file that resides in a remote location — in a network shared folder, on a backup repository or a cloud repository.

Specify where the backup file resides:

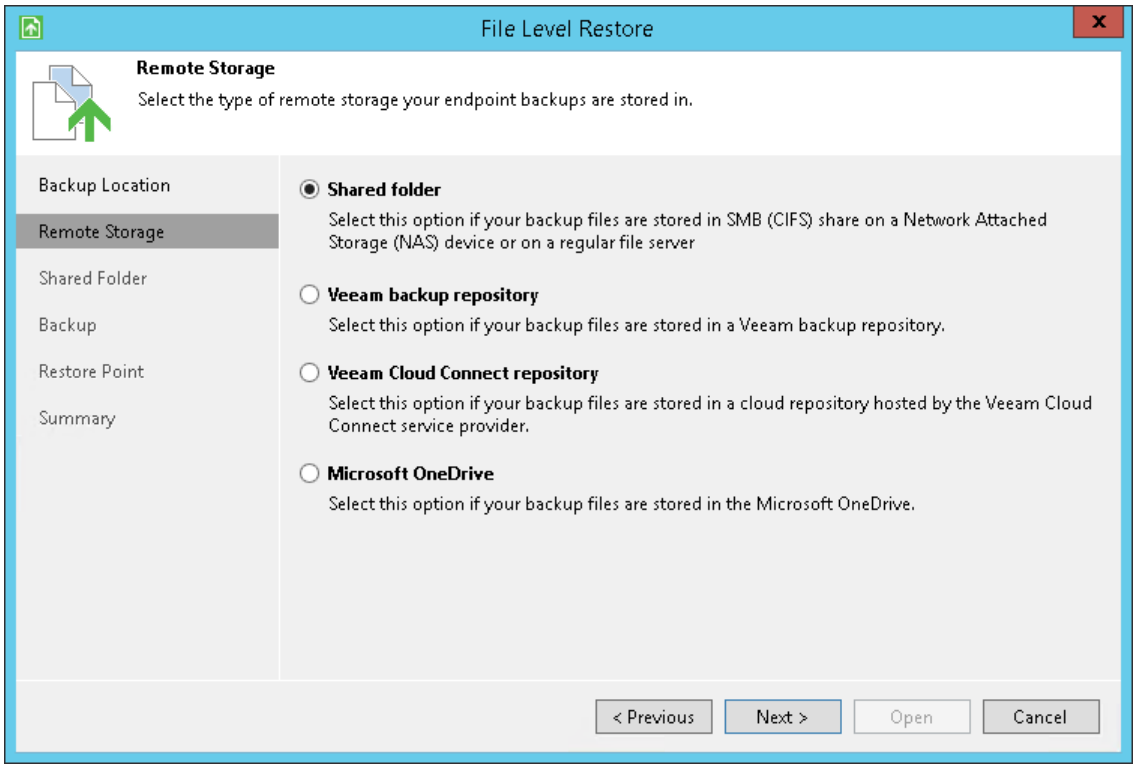
- **Shared folder** — select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the Shared Folder step of the wizard.
- **Veeam backup repository** — select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the Backup Server step of the wizard.
- **Veeam Cloud Connect repository** — select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the Service Provider step of the wizard.
- **Microsoft OneDrive** — select this option if the backup file resides in the Microsoft OneDrive cloud storage. With this option selected, you will pass to the Microsoft OneDrive step of the wizard.

NOTE

Consider the following:

- The Microsoft OneDrive option is not available if the Veeam Agent computer runs a Microsoft Windows Server OS.
- The **Veeam Cloud Connect repository** option is also available at the **Network Storage** step of the wizard. However, for the Free product edition, you cannot store backups on a Veeam Cloud Connect repository. As a result, the restore process for the backup stored on a cloud repository is not described in this guide.

Specify settings to connect to Microsoft OneDrive:



Step 4. Specify Remote Storage Settings

Specify settings for the remote storage that contains a backup file from which you plan to restore data:

- Shared folder settings — if you have selected the **Shared folder** option at the Network Storage step of the wizard.
- Veeam backup repository settings — if you have selected the **Veeam backup repository** option at the Network Storage step of the wizard.
- Veeam Cloud Connect repository settings — if you have selected the **Veeam Cloud Connect repository** option at the Network Storage step of the wizard.
- Microsoft OneDrive settings — if you have selected the **Microsoft OneDrive** option at the Network Storage step of the wizard.

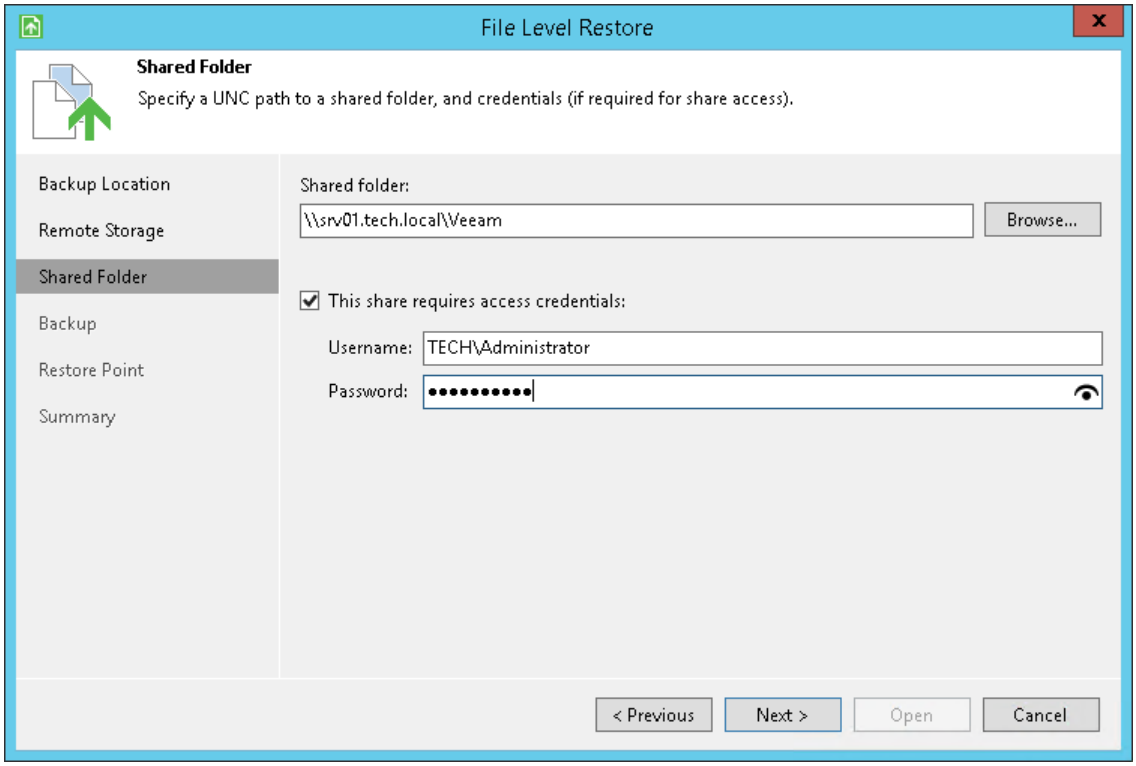
Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have chosen to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. In the **Shared folder** field, type in a UNC name of the network shared folder with the backup file. Keep in mind that the UNC name always starts with two back slashes (\\).
2. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the down-level logon name format. For example, *DOMAIN\UserName* or *HOSTNAME\UserName*.

   If you do not select the **This share requires access credentials** check box, Veeam Agent for Microsoft Windows will connect to the shared folder using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed.
3. To view the entered password, click and hold the eye icon on the right of the **Password** field.



Backup Server Settings

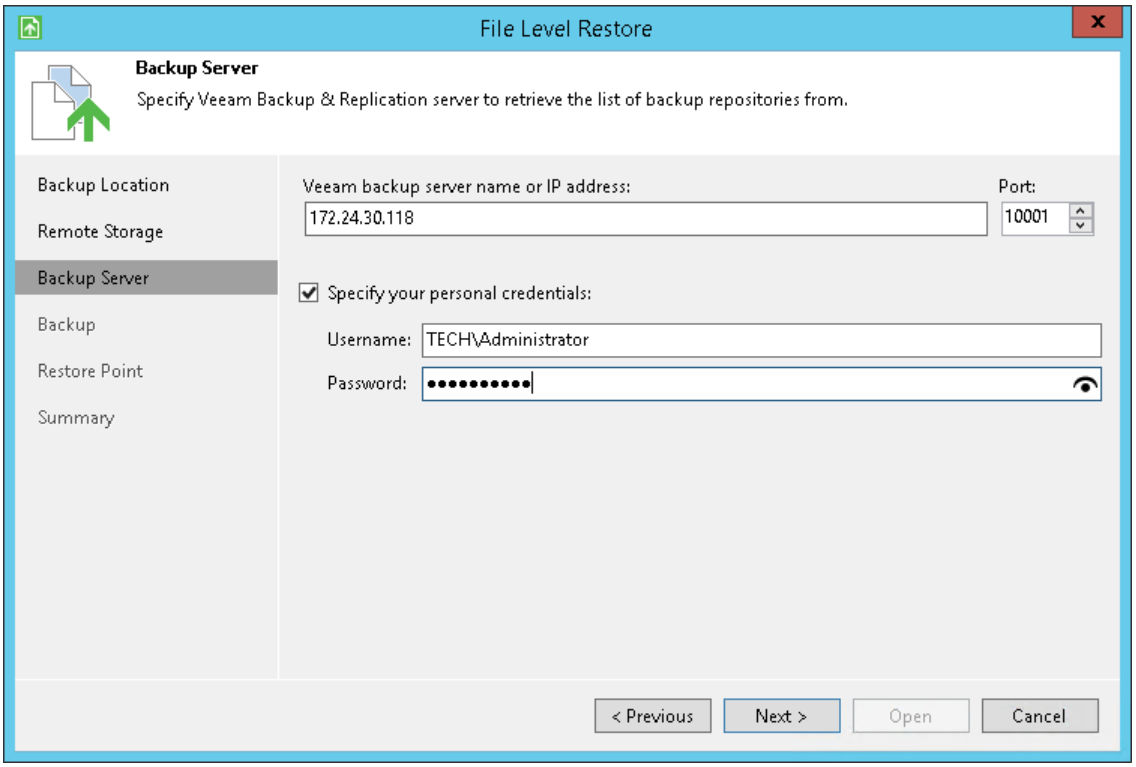
The **Backup Server** step of the wizard is available if you have chosen to restore data from a backup file located on a backup repository.

Specify settings for the Veeam backup server that manages the backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup repository. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. Select the **Specify your personal credentials** check box. In the **Username** and **Password** fields, enter a user name and password of the account that has access to this backup repository. Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see Setting Up User Permissions on Backup Repositories.

   If you do not select the **Specify your personal credentials** check box, Veeam Agent for Microsoft Windows will connect to the backup repository using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can add the computer account (*DOMAIN\COMPUTERNAME$*) to an AD group and grant access rights on the backup repository to this group.

   Setting access permissions on the backup repository to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). If you have set such permissions on the backup repository, you can omit specifying credentials. However, we recommend this scenario for demo environments only.



IMPORTANT

If you specify a DNS name of the Veeam backup server, make sure that the Veeam backup server name is resolved into IPv4 address on the machine where Veeam Agent for Microsoft Windows is installed. The Veeam Backup Service in Veeam Backup & Replication listens on IPv4 addresses only. If the Veeam backup server name is resolved into IPv6 address, Veeam Agent for Microsoft Windows will fail to connect to the Veeam backup server.

Service Provider Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. Specify service provider settings.
2. Verify the TLS certificate and specify user account settings.

Specifying Service Provider Settings

The **Service provider** step of the wizard is available if you have chosen to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.

