Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication.

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified mailbox back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft Exchange populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

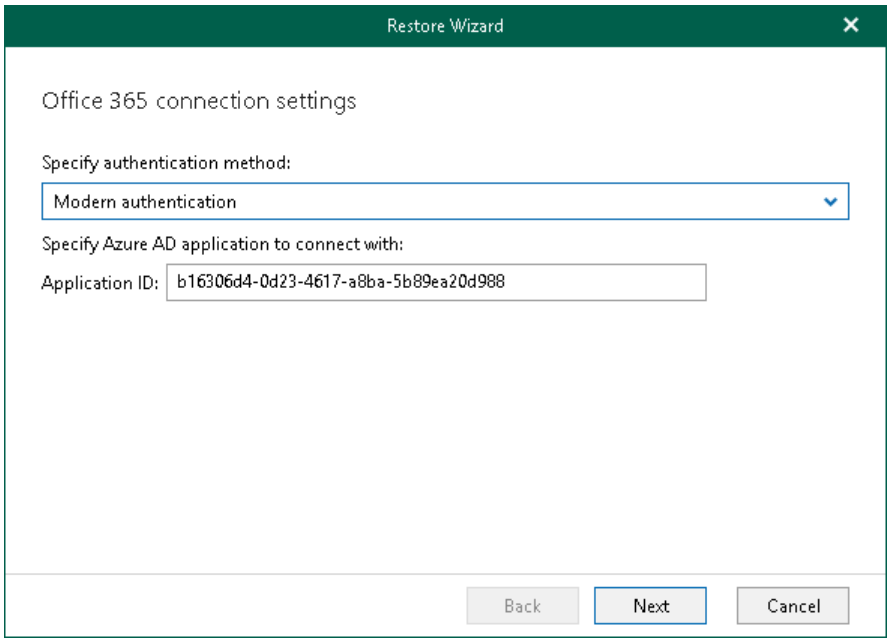

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the Exchange organization and click **Restore**.

   Make sure that the account that you use has both the *Exchange Admin* and the *Global Admin* roles.

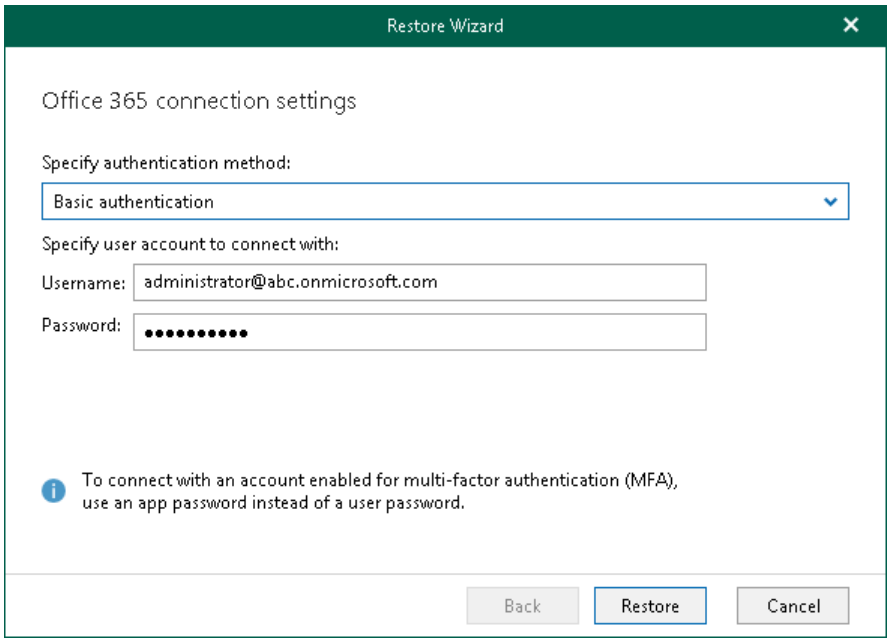

Step 3. Log In to Office 365

[This step is available only if you have selected the *Modern Authentication* option at the previous step.]

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has both the *ApplicationImpersonation* role and the *Exchange admin* role. For more information, see this Microsoft article.

4. Return to the **Restore Wizard** window and click **Restore**.

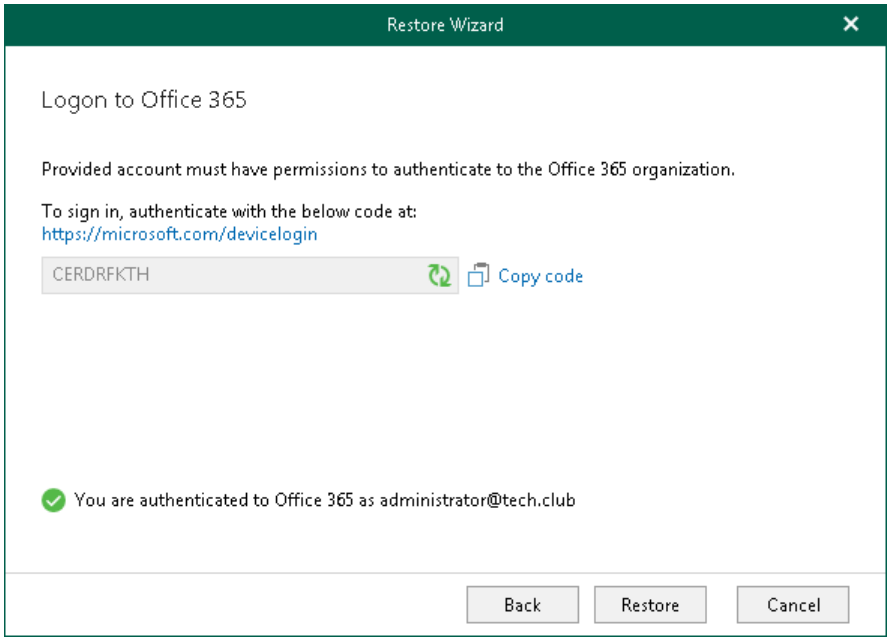

Restoring to On-Premises Microsoft Servers

To perform *1-Click Restore* to on-premises Microsoft servers, do the following:

1. Launch the Restore wizard.
2. Specify a user account.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

1. In the navigation pane, select a mailbox, folder or item that you want to restore.
2. On the **Mailbox/Folder/Items** tab, select **Restore Mailbox/Restore Folder/Restore Item** > **Restore to <mailbox_name>** or right-click a mailbox/folder/item and select **Restore to <mailbox_name>**.

Step 2. Specify User Account

At this step of the wizard, specify an account that you want to use to connect to the Exchange server.

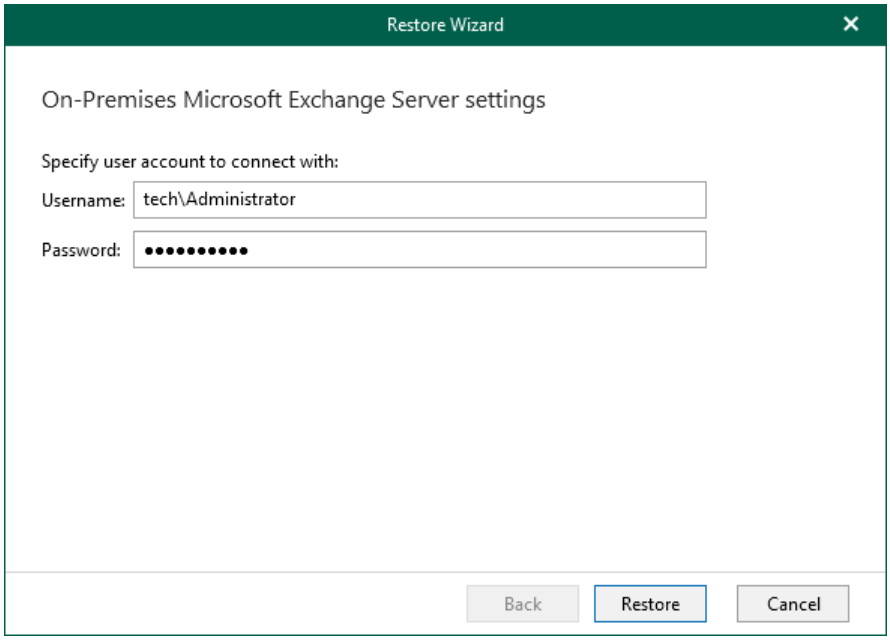

Data Export

Continue with this section to learn more about exporting Microsoft Exchange data.

> **TIP**
>
> Before exporting data, make sure to read Considerations and Limitations.

Exporting to Custom Location

To export objects to the custom location, do the following:

1. In the navigation pane, select a folder or in the preview pane, select an item.
2. On the **Folder/Items** tab, select **Export Folder/Item > Export to .pst file** or right-click a folder/item and select **Export to .pst file**.

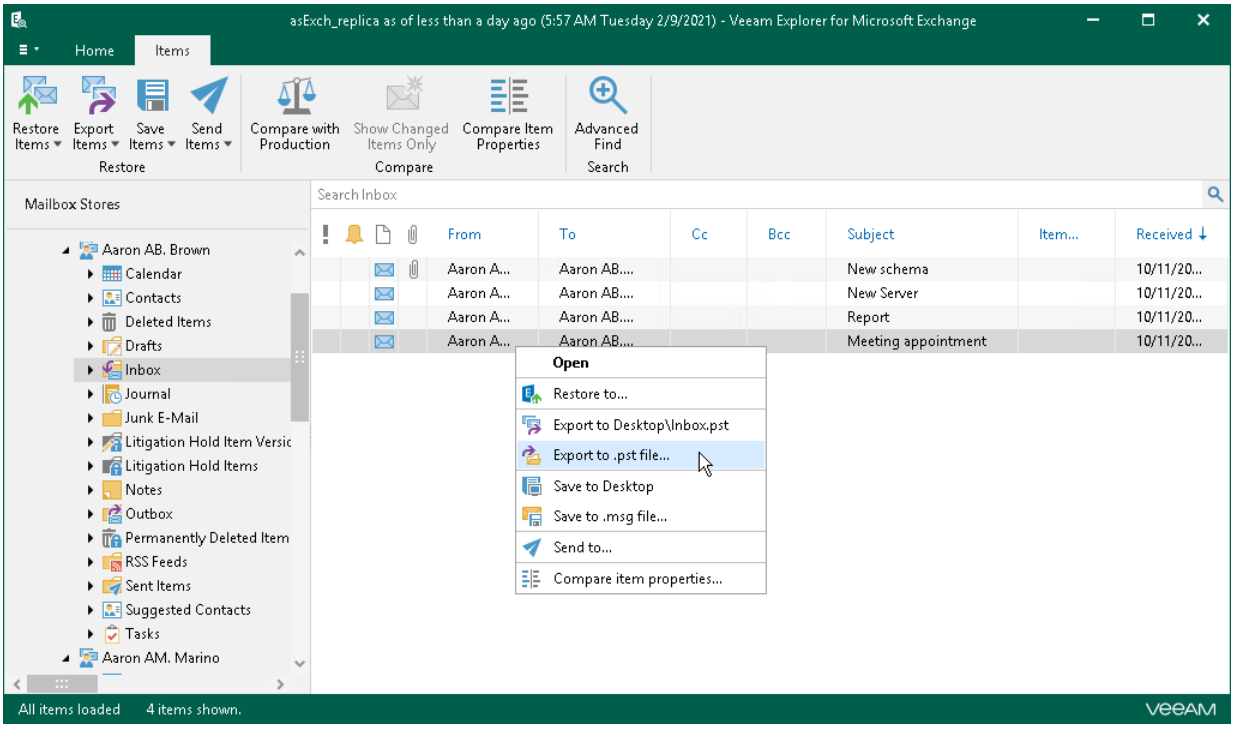

3. In the **Export Folder**/**Export Mail Item** dialog, do the following:
   - In the **Export location** field, specify the name and location for the new *.pst* file.
   - Select the **Apply filter** check box and enter keywords to export only those items that match specified criteria. To provide multiple keywords, use the *OR* logical operator as the delimiter.

     Veeam Explorer for Microsoft Exchange will check for the specified keywords in all fields such as **From**, **To**, **Subject** and **Body**.

     The **Apply filter** option is available only when exporting folders and mailboxes, not items.
   - Select the **Limit PST size to** check box and specify the size of which to export *.pst* files.

     With this option, Veeam Explorer for Microsoft Exchange will create multiple *.pst* files as of the specified size.

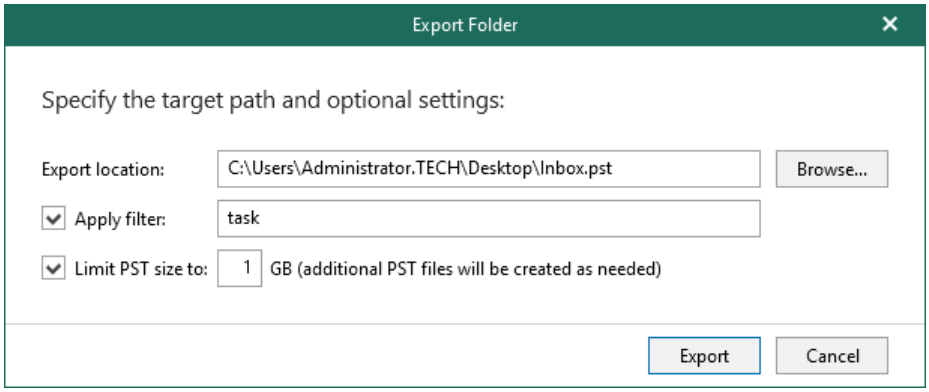

Using 1-Click Export

Veeam Explorer for Microsoft Exchange allows you to quickly export mailbox store data as *.pst* files.

> **NOTE**
>
> Export operations require Microsoft Outlook to be installed on a computer running Veeam Explorer for Microsoft Exchange.

Depending on the object type you want to export, do the following:

- To export a folder, select it in the navigation pane and click **Export Folder** > **Export to <default_location>\<original_folder_name>.pst** on the **Folder** tab.

  Alternatively, in the navigation pane, right-click a folder and select **Export to <default_location>\<original_folder_name>.pst**.

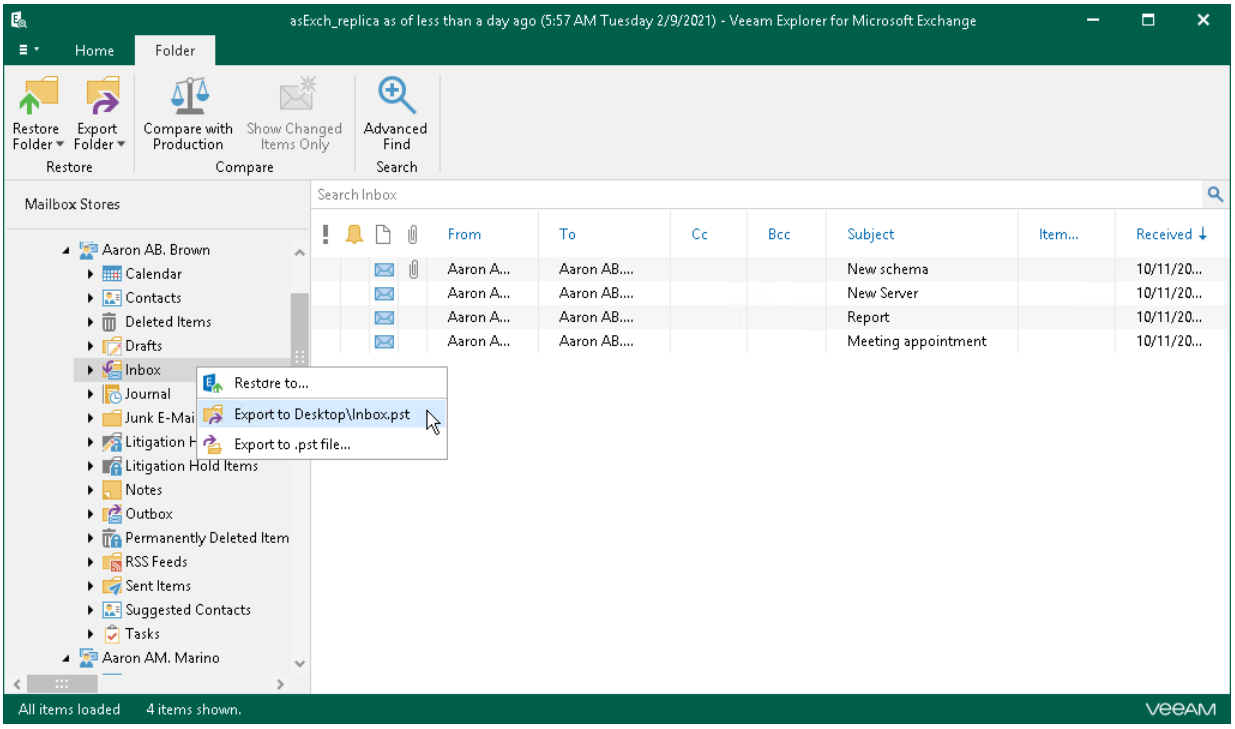

- To export an item, select it in the preview pane and click **Export Items** > **Export to <default_location>\<original_folder_name>.pst** on the **Items** tab.

  Alternatively, in the preview pane, right-click an item and select **Export to <default_location>\<original_folder_name>.pst**.

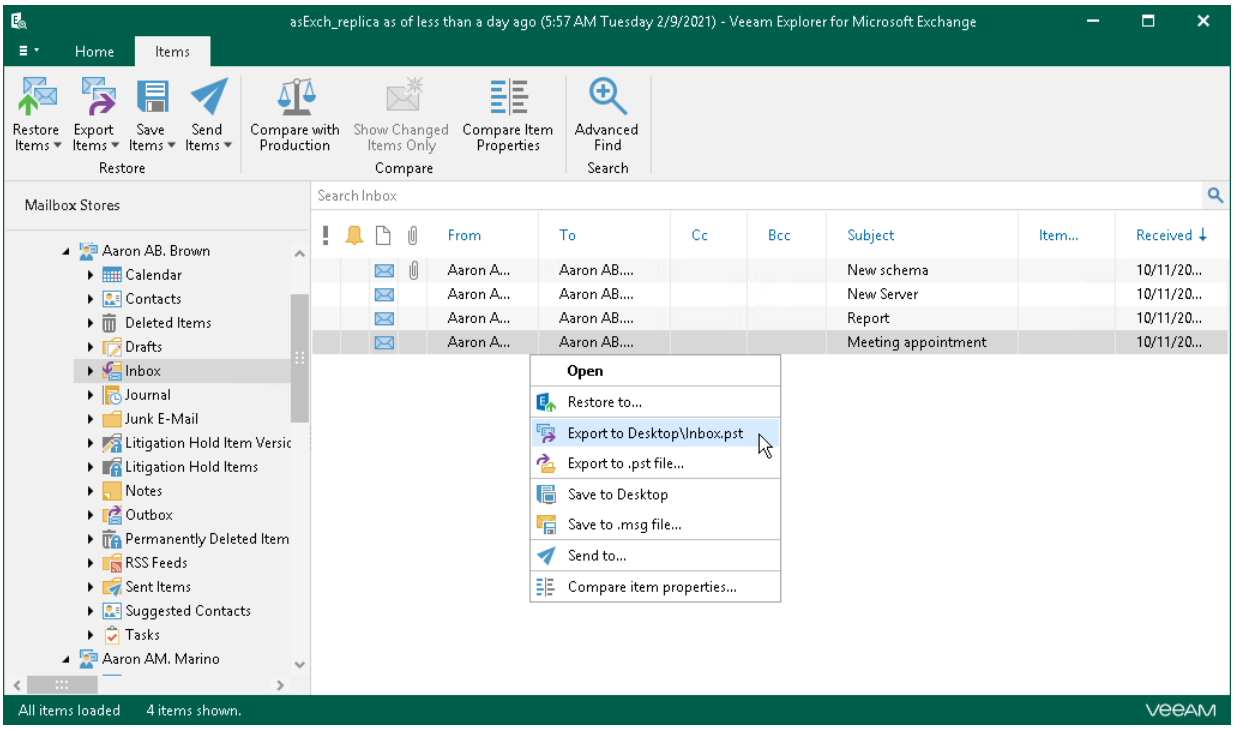

- To export a mailbox, select it in the navigation pane and click **Export Mailbox** > **Export to <default_location>\<original_mailbox_name>.pst** on the **Mailbox** tab.

  Alternatively, in the navigation pane, right-click a mailbox and select **Export to <default_location>\<original_mailbox_name>.pst**.

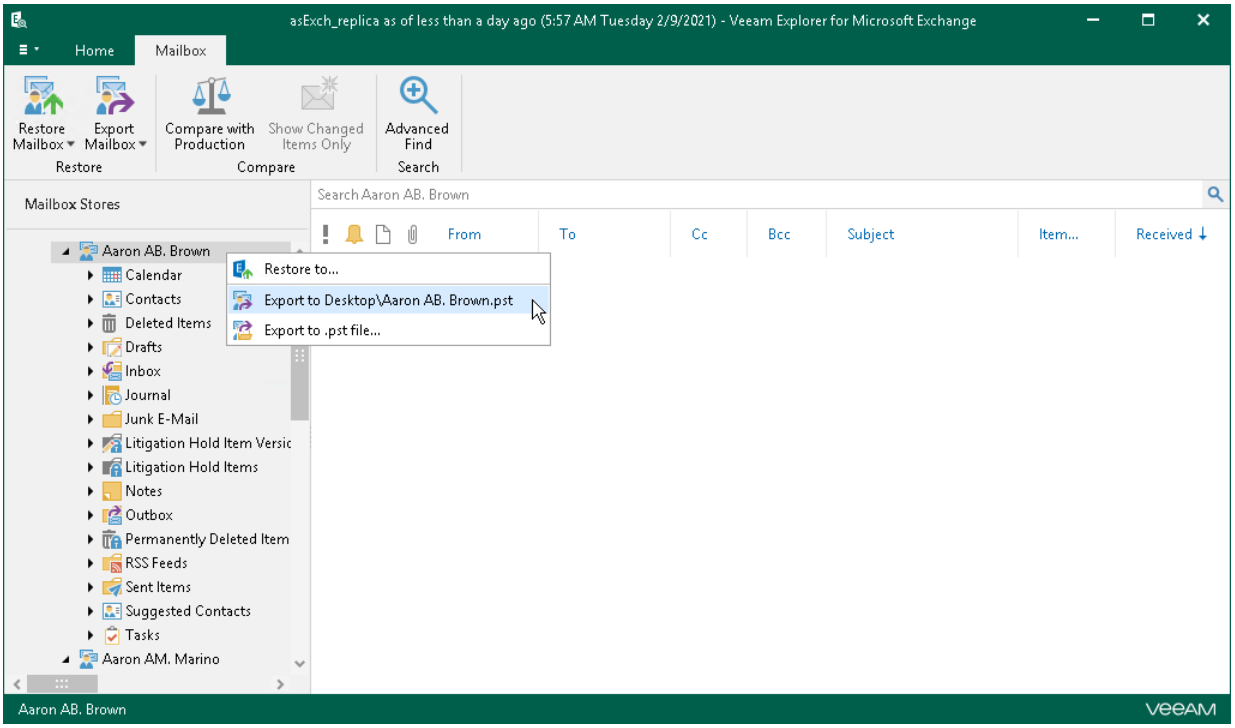

- To export a mailbox store or all stores added to the scope, select a store or select the **All Stores** node and click **Export Stores** > **Export to <default_location>\<original_store_name>.pst** on the **Home** tab.

  Alternatively, in the navigation pane, right-click a mailbox store and select **Export to <default_location>\<original_store_name>.pst**.

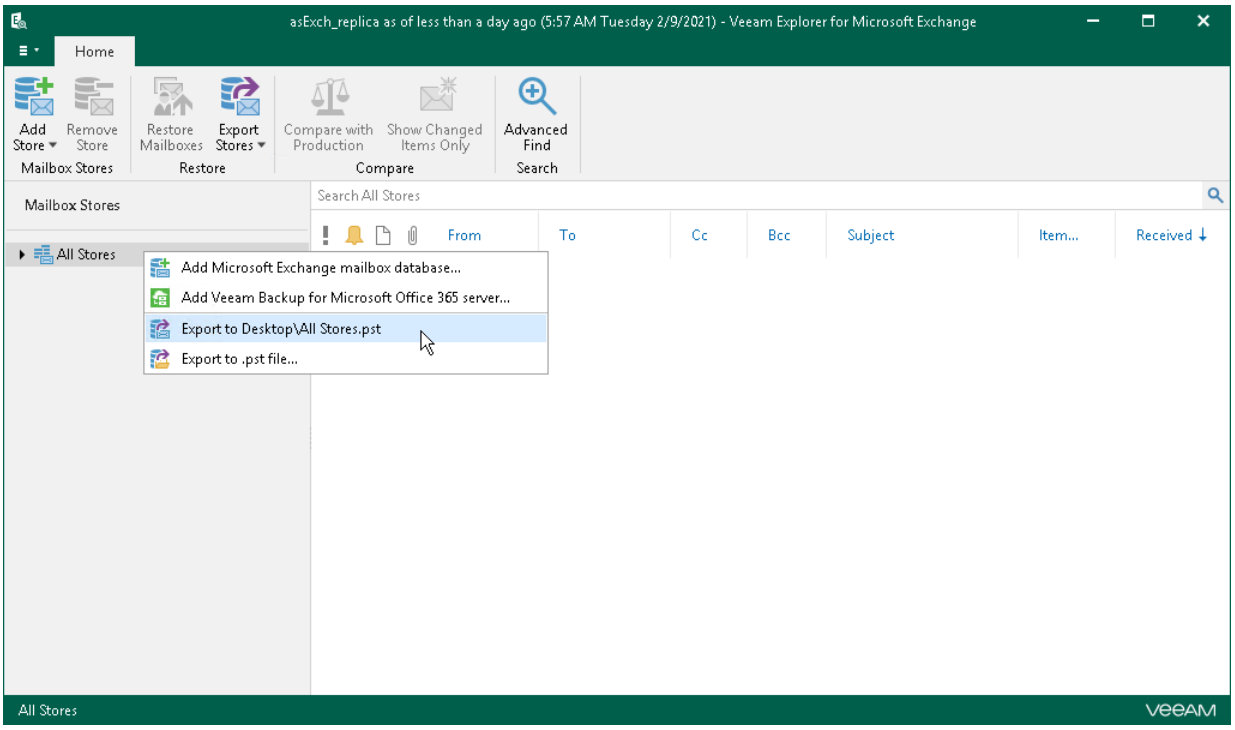

Receiving Export Reports

Veeam Explorer for Microsoft Exchange can deliver reports that show summary information about export operations.

Consider the following:

- If you want export reports to be delivered, make sure to configure an SMTP server in Veeam Explorer for Microsoft Exchange and/or Veeam Backup & Replication notification settings.

  For more information, see the Configuring SMTP Settings section of this guide and the Configuring Global Email Notification Settings section of the Veeam Backup & Replication User Guide.

- A server with Veeam Explorer for Microsoft Exchange and Veeam Backup & Replication must be in a domain.

- Export reports are delivered automatically upon each export session.

  You can use Export to Custom Location or 1-Click Export.

- Recipients of the report are chosen as follows:

  - The recipient is always the current user account under which Veeam Explorer for Microsoft Exchange is running when using only Veeam Explorer for Microsoft Exchange notification settings.
  - The recipient is always the user account provided in Veeam Backup & Replication notification settings when using only Veeam Backup & Replication settings.
  - If you have configured both Veeam Explorer for Microsoft Exchange and Veeam Backup & Replication notification settings, the report will be sent to both recipients.

The following table lists fields that are shown in the report.

Report field	Description
Performed by	Veeam Explorer for Microsoft Exchange
Version	Shows the Veeam Explorer for Microsoft Exchange build number.
Percent complete	Shows the number of items included in the *.pst* file.
Started by	Shows the user account generated the report.
Start time, End time	Shows the creation time of a report.
Root	Shows the object path.
Path	Shows the path to the database file.
Export type	Shows the export type.

Report field	Description
Stores	Shows a datastore name from which items have been exported.
Mailboxes to search	Shows a mailbox name from which items have been exported.
Mailboxes searched successfully, Mailboxes not searched successfully	Shows results for mailboxes to search.
Size	Shows the total size of exported items.
Items	Shows the number of exported items.
Errors	Shows errors that occurred during export.

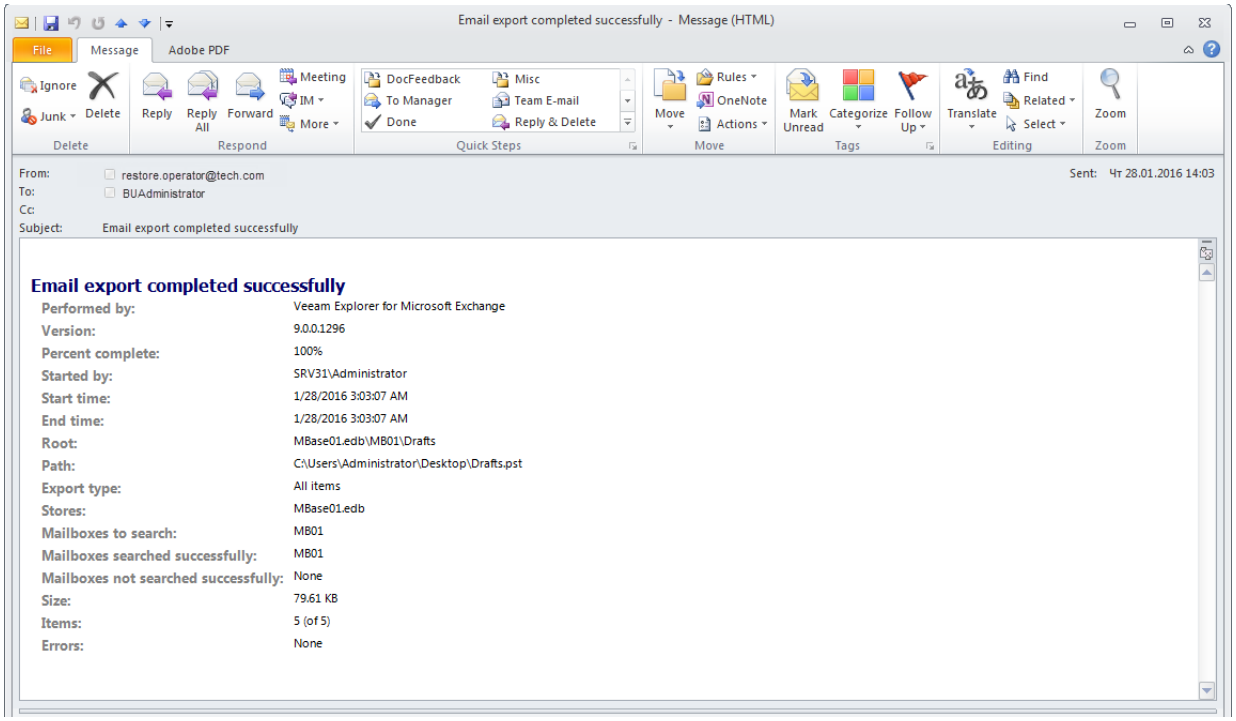

Data Compare

This section explains how to compare data in a backup file with that of the production state.

1. In the navigation pane, select an object and click **Compare with Production** on the toolbar.
2. Proceed to either Compare to Office 365 Organizations or Compare to On-Premises Microsoft Servers.
3. After the authorization process is complete, click **Show Changed Items Only** on the toolbar to view only those items that have been changed.

To view the details on what exactly has changed since the last backup, right-click an item and select **Compare item properties**. The discrepancies will be shown in the **Compare Message Properties** window. If both objects are equal, nothing will be displayed.

To show unchanged objects, select **Show unchanged properties** in the top-right corner. To show system properties, select **Show system properties**.

Double-click **Body** to see the body message.

To compare a single item, right-click an item and select **Compare item properties**.

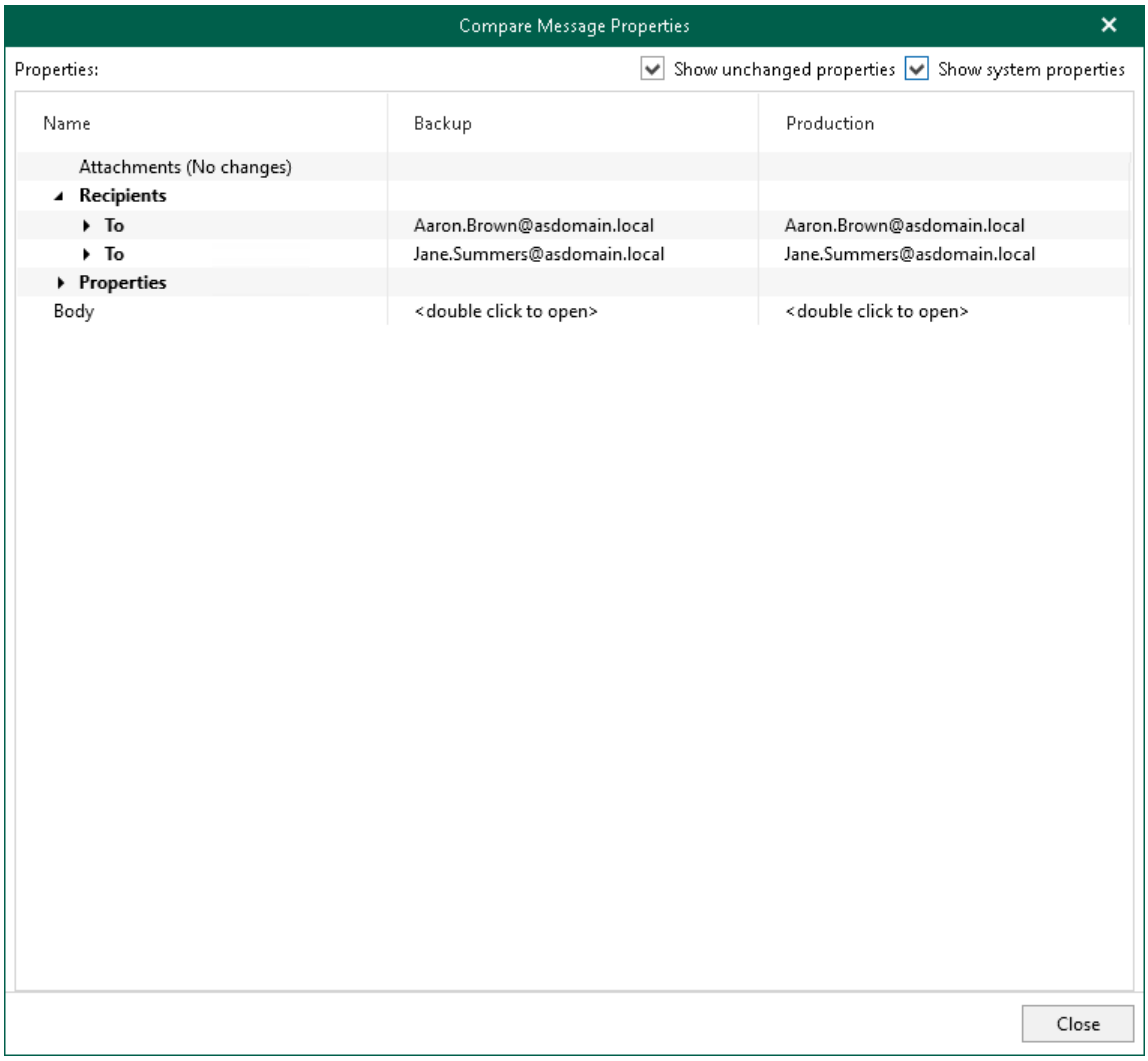

Compare to Office 365 Organizations

This operation is available only for backups created in Veeam Backup for Microsoft Office 365.

To compare data with your production Office 365 organizations, do the following:

1. Select Authentication Method.
2. Log In to Office 365.

Step 1. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication.

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified mailbox back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft Exchange populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

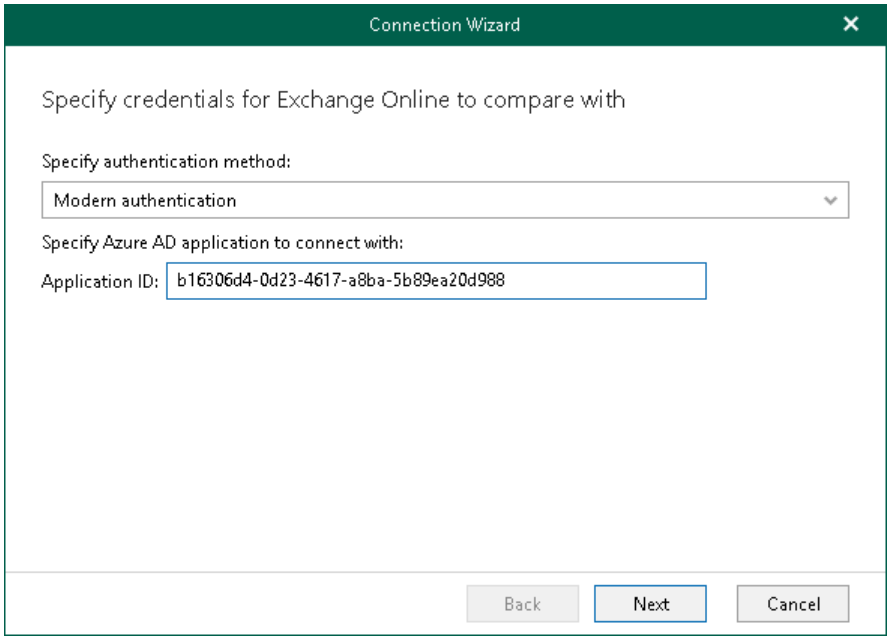

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the Exchange organization and click **Connect**.

   Make that the account that you use has both the *Exchange Admin* and the *Global Admin* roles.

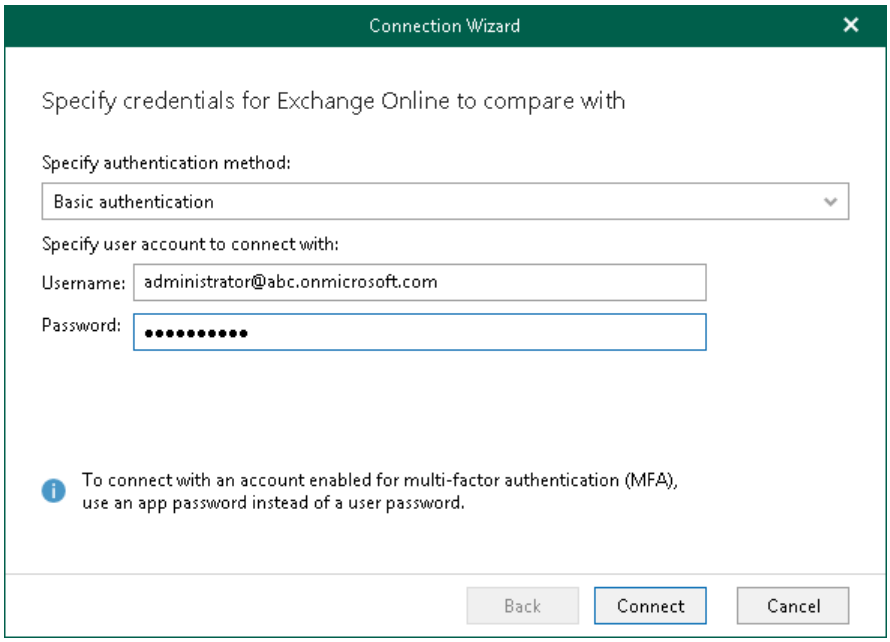

Step 2. Log In to Office 365

[This step is available only if you have selected the *Modern Authentication* option at the previous step.]

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has both the *ApplicationImpersonation* role and the *Exchange admin* role. For more information, see this Microsoft article.

4. Return to the **Connection Wizard** window and click **Connect**.

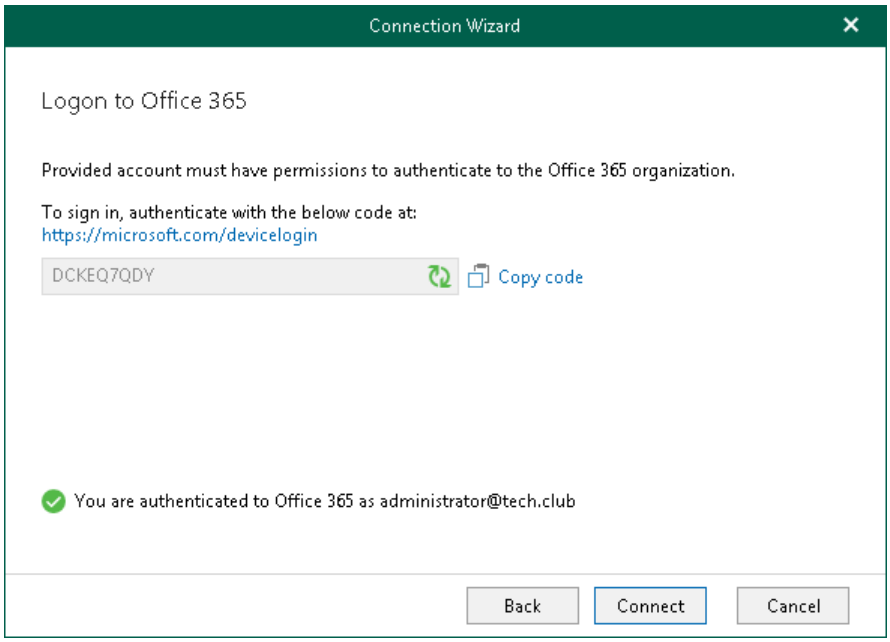

Compare to On-Premises Microsoft Servers

To compare data with your production on-premises Microsoft Exchange, do the following:

1. Specify a domain and credentials.
2. Specify a target mailbox server.

Step 1. Specify Domain and Credentials

At this step of the wizard, enter a domain name where your CAS server resides and provide credentials.

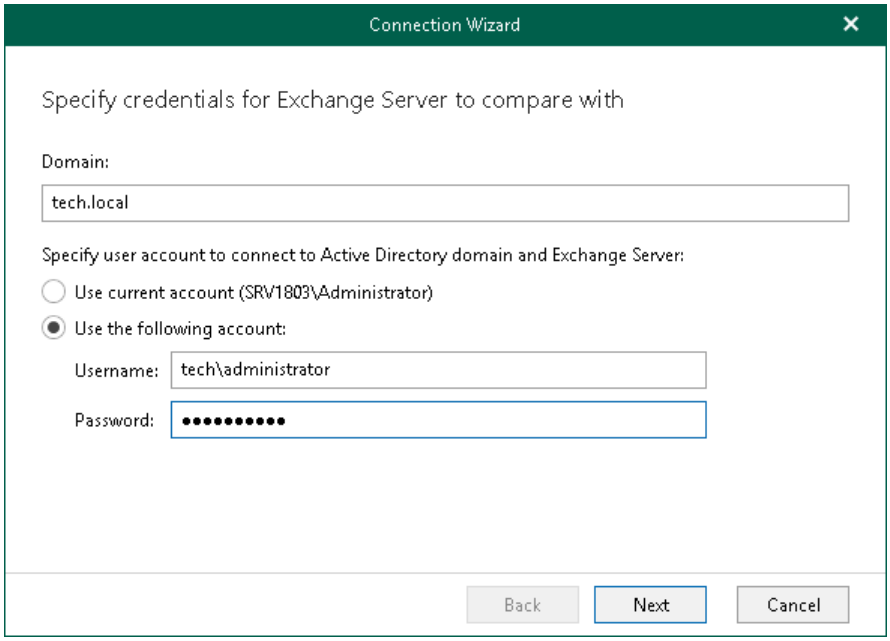

Step 2. Specify Target Mailbox Server

At this step of the wizard, specify a target server to which you want to restore mailboxes and click **Connect**.

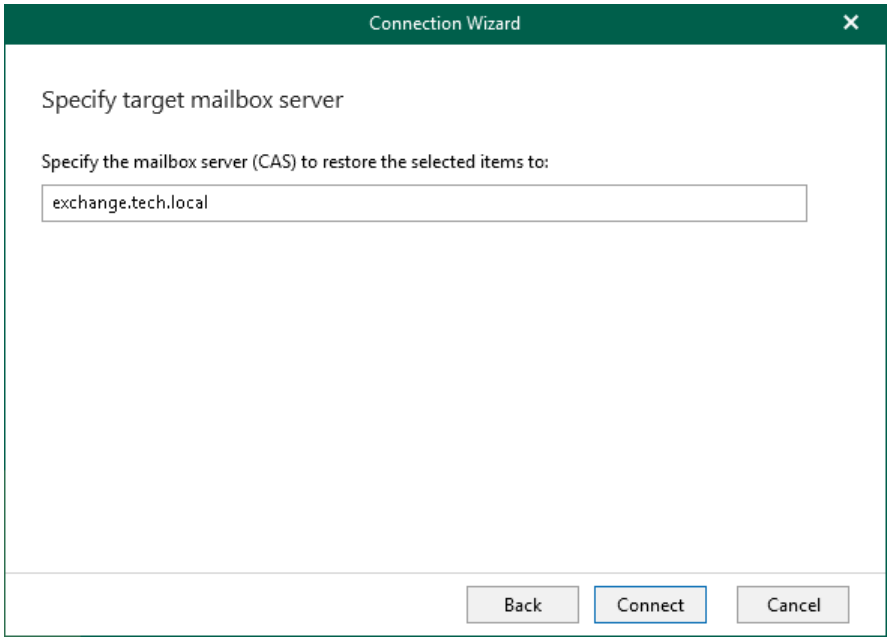

Saving Microsoft Exchange Items

Veeam Explorer for Microsoft Exchange allows you to save any item as *Microsoft Exchange Mail Document* (*.msg*) files.

Saving to Custom Location

To save items to a custom location, do the following:

1. In the navigation pane, select an item.
2. On the **Items** tab, select **Save Items** > **Save to .msg file** or right-click a folder/item and select **Save to .msg file**.

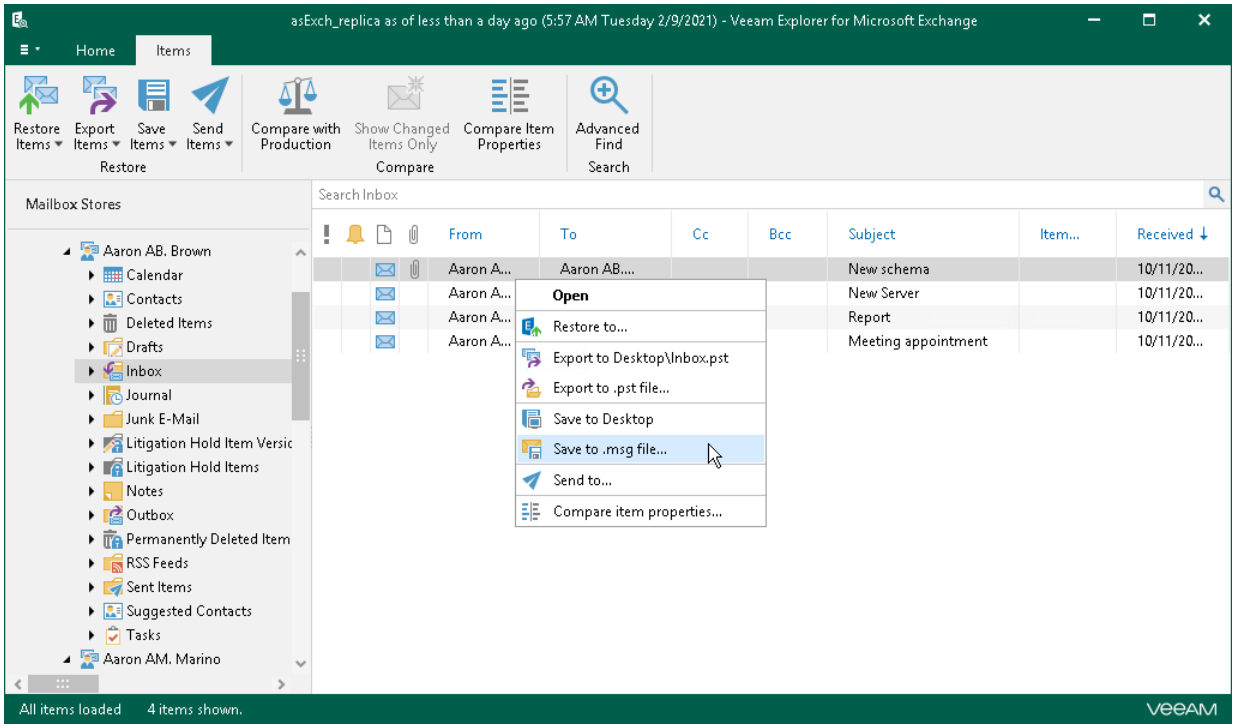

1-Click Save

To save your Exchange data, do the following:

1. In the navigation pane select a folder containing required items.
2. In the preview pane, select items.

   Use the **Search** field to find the necessary items.
3. On the **Items** tab, select **Save Items > Save to <folder_name>** or right-click a folder/item and select **Save to <folder_name>**.

   A **<folder_name>** name depends on the latest location that was used when saving items.

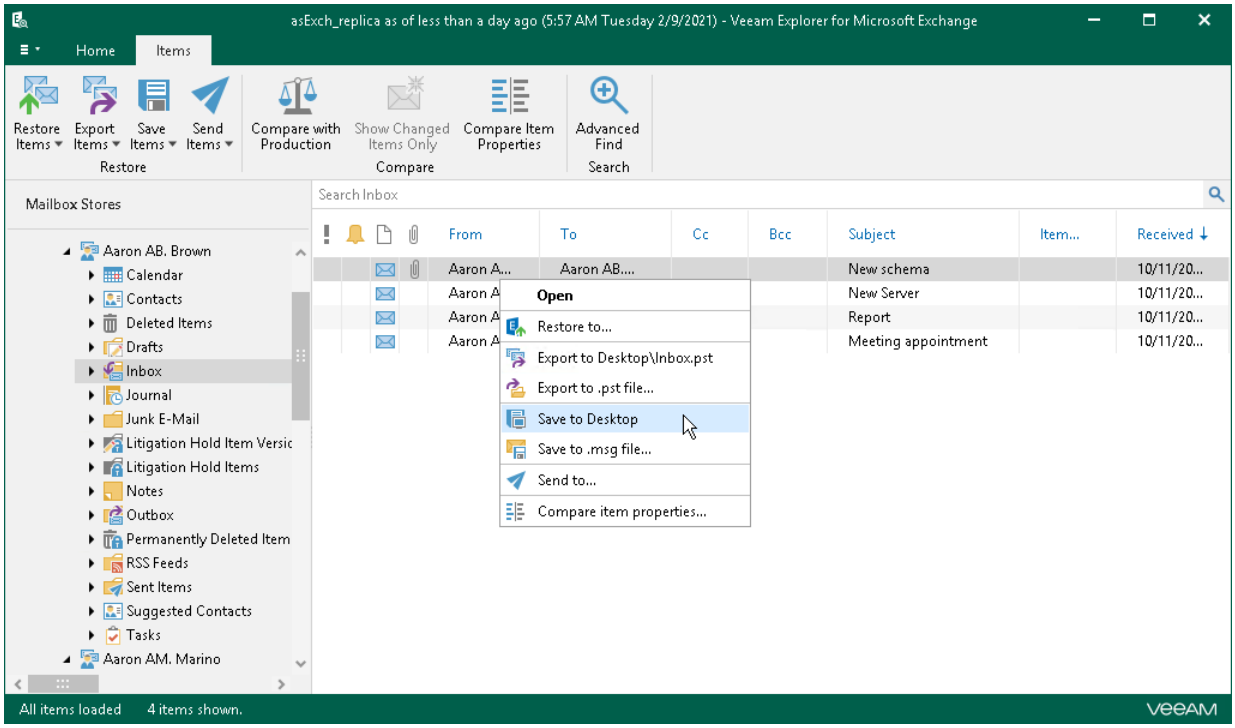

Sending Microsoft Exchange Items

Veeam Explorer for Microsoft Exchange allows you to send Exchange items via email.

> **TIP**
>
> Before sending documents, make sure to configure SMTP settings, as described in Configuring SMTP Settings. The amount of data you can send at a time depends on your SMTP server configuration.

To send items, do the following:

1. In the navigation pane select a folder.
2. In the preview pane, select items to send.

   Use the **Search** field to find the necessary items.
3. On the **Items** tab, select **Send Items** > **Send to** or right-click a folder/item and select **Send to**.

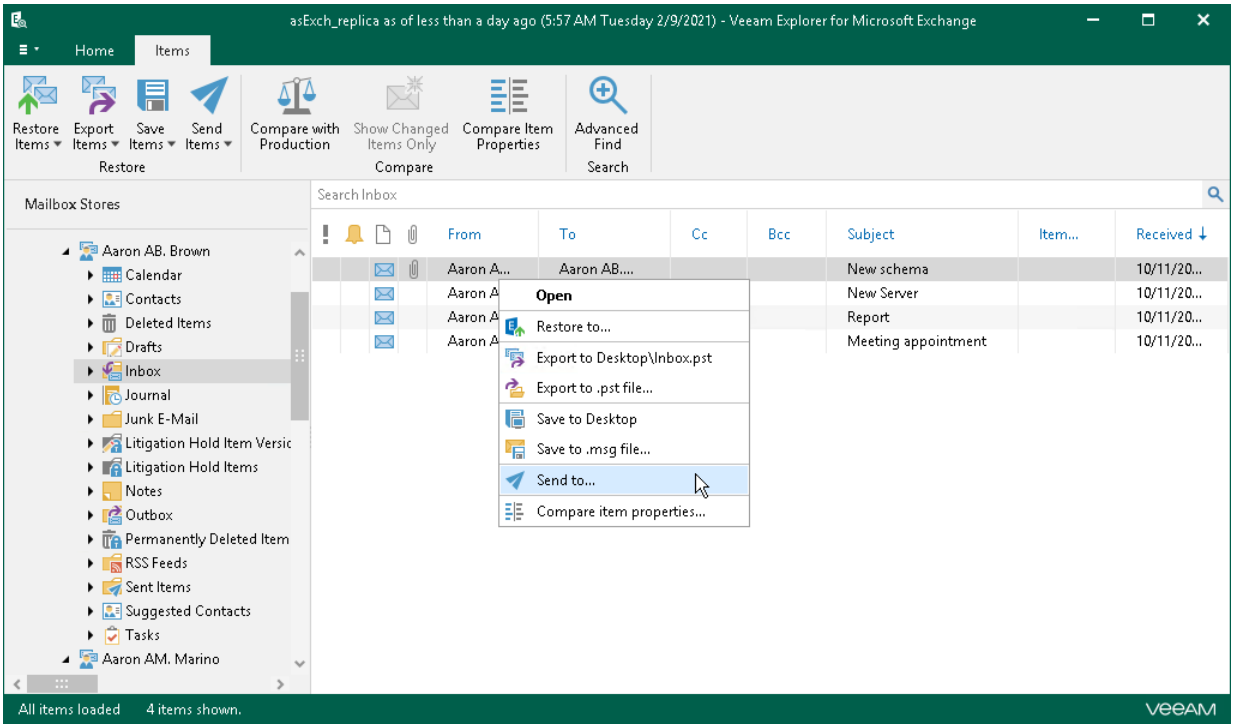

4. Provide a recipient address.

   The **From** field is filled automatically based on the address you have provided when configuring SMTP settings. To edit the message body, click **More Details**.

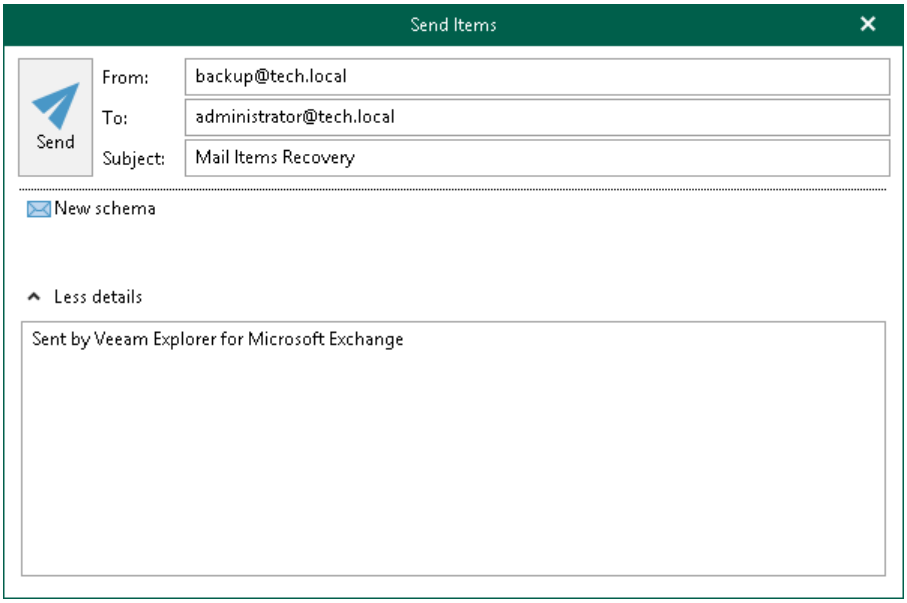

Veeam Explorer for Microsoft SharePoint

Veeam Explorer for Microsoft SharePoint allows you to restore and export Microsoft SharePoint objects from backups created in Veeam Backup & Replication and Veeam Backup for Microsoft Office 365.

Planning and Preparation

Continue with this section to learn how to configure your environment before start using Veeam Explorer for Microsoft SharePoint:

System Requirements

This section lists system requirements for Veeam Explorer for Microsoft SharePoint.

Component	Requirement
Microsoft SharePoint	For more information about supported versions of Microsoft SharePoint, see: • The Supported Applications subsection of the Veeam Backup & Replication User Guide. • The System Requirements section of the Veeam Backup for Microsoft Office 365 User Guide.

NOTE

Consider the following:

- To restore SharePoint content database items from a server running Microsoft Windows ReFS, the Veeam backup server or a management console must be installed on the Microsoft Windows Server 2012 or later.
- To restore from a server running Microsoft Windows ReFS 3.x, Veeam backup server or a management console must be installed on the Microsoft Windows Server 2016.

Used Ports

The following tables list network ports that must be opened to manage inbound/outbound traffic.

General

From	To	Protocol	Port	Notes
Veeam Explorer for Microsoft SharePoint	**SharePoint Web Application**	As recommended by Microsoft	As recommended by Microsoft	For more information on recommended port numbers and protocols for SharePoint web application, see this Microsoft article To discover ports currently used by your SharePoint web application, follow the steps described in this Microsoft article.
Veeam Backup Server / Standalone Veeam Backup Console / Mount server associated with the backup repository	**Staging Server**	TCP, UDP	135, 445	Ports used to deploy the runtime coordination process on a target guest OS.
		TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic RPC range used by the runtime coordination process that is deployed on a target guest OS.[1] For more information, see this Microsoft article.
		TCP	6160	Port used to communicate with the installer service.
Veeam Backup Server / Standalone Veeam Backup Console	**Staging Server**	TCP	1433, 1434 and other	Ports used to communicate with Microsoft SQL servers hosting content databases. Exact port numbers depend on the configuration of a Microsoft SQL server. For more information, see this Microsoft article.
		UDP	1434	Port used by the Microsoft SQL Server Browser service. For more information, see this Microsoft article.

[1] If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports: during setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *RPC function call failed* error, you need to configure

dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article.

> **NOTE**
>
> To restore database items or lists to a server that is running in a DMZ, the SharePoint web application ports will be used.

iSCSI Traffic

From	To	Protocol	Port	Notes
Staging Server	**Mount server associated with the backup repository**	TCP	3260 to 3270	Ports used by iSCSI initiator to connect to the iSCSI target.

Required Permissions

The following table lists required permissions for user accounts to back up and restore Microsoft SharePoint data.

Operation	Required Roles and Permissions
Backup	For more information, see: • The Required Permissions section of the Veeam Backup & Replication User Guide. • The Required Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.
Restore to on-premises Microsoft SharePoint from backups created in Veeam Backup & Replication and Veeam Backup for Microsoft Office 365	To restore data to on-premises Microsoft SharePoint, make sure to configure user accounts as follows: • The account must be granted *Full Control* to connect to the target SharePoint server. • The account must be assigned either the *Site Administrator* or *System Account* role to restore user permissions. • If permissions of items being restored are inherited from the parent one, the account must be granted *Full Control*. • If permissions of items being restored are not inherited from the parent one and items being restored replace the existing ones, the account must be granted *Contribute* and *Full Control*.

Operation	Required Roles and Permissions
Restore to Microsoft Office 365 from backups created in Veeam Backup for Microsoft Office 365	To restore data to SharePoint Online, make sure to configure user accounts as follows: **Restore Using Basic Authentication Method** • The account used to log in to Microsoft Office 365 must have the *Global Administrator* or *SharePoint Administrator* role assigned. • For restore of personal SharePoint sites, make sure to select the **Allow users to run custom script on personal sites** option in the SharePoint admin center. For more information, see this Microsoft article. • During restore, Veeam Backup for Microsoft Office 365 automatically assigns the *Site Collection Administrator* role to the user account. **Restore Using Modern App-Only Authentication Method** • The account used to log in to Microsoft Office 365 must have the *Global Administrator* or *SharePoint Administrator* role assigned. • For restore of personal SharePoint sites, make sure to select the **Allow users to run custom script on personal sites** option in the SharePoint admin center. For more information, see this Microsoft article. • During restore, Veeam Backup for Microsoft Office 365 automatically assigns the *Site Collection Administrator* role to the user account. • Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see Required Azure AD Application Settings. • If you restore data with Azure AD applications using a certificate, make sure that your Azure AD application is granted the required permissions. For more information, see the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

Consider the following:

- The current account can only be used to access a *local* staging server. To connect to a *remote* server, use appropriate authentication credentials to access that server.
- The account requires the *sysadmin* fixed server role on a staging Microsoft SQL server.
- For ADFS as an authentication provider:
  - When using *Windows Authentication*, you can use both you current account or provide another account.
  - When using *Forms Authentication*, the current account cannot be used.

Required Backup Job Settings

When you create a backup job, make sure to enable the **application-aware image processing** option, as described in the Specify Guest Processing Settings section of the Veeam Backup & Replication User Guide.

Staging SQL Server

To perform Microsoft SharePoint items recovery, Veeam Explorer requires a Microsoft SQL server to be used as a staging system.

The server you are going to use as a staging system must conform to the following:

- A staging system must have the same or later version of Microsoft SQL Server that hosts restored Microsoft SharePoint databases.
- Microsoft SQL Server that is included in Microsoft SQL Server Failover Cluster cannot be used as a staging system.
- Nodes participating in AlwaysOn Availability Groups are supported. However, using Availability Group Listeners as staging servers is not recommended.

> **NOTE**
>
> You can use Microsoft SQL Server 2012/2016 Express editions which are shipped with the Veeam Backup & Replication distribution package. Consider that due to Express edition limitations, the maximum database size that can be attached is 10 GB. For more information, see this Microsoft article.

Remote BLOB Stores Support

A BLOB store must either be included in the SharePoint backup created by Veeam Backup & Replication or Veeam Agent for Microsoft Windows (for automated discovery) or stored on a local machine running Veeam Explorer (for manual discovery).

> **IMPORTANT**
>
> RBS FILESTREAM is the only provider that is supported in the current version.

Make sure the staging SQL Server configuration meets the following requirements:

- **FILESTREAM** is enabled on a database server, as described in the following articles:
  - For SQL Server 2017 for Windows, see this Microsoft article.
  - For SQL Server 2016, see this Microsoft article.
  - For SQL Server 2014, see this Microsoft article.
  - For SQL Server 2012, see this Microsoft article.
  - For SQL Server 2008 R2, see this Microsoft article.
- **RBS Client Library** is installed on the database server.

  For Microsoft SQL Server 2014 and later, the Remote Blob Store setup is included in the installation media. For other versions, you can use the corresponding Microsoft SQL Server Remote Blob Store installation package. For more information on how to install the package, see this Microsoft article.

Considerations and Limitations

This section covers considerations and known limitations of Veeam Explorer for Microsoft SharePoint.

- General
- Status Recovery Limitations
- Restore of Documents, Libraries and Lists
- Restore of List Items
- Restore of List Items with Link Attachments
- Restore of Surveys
- Restore of Sites
- Export and Import
- Data Type Limitations
- Workflow-Related Considerations

General

- When Veeam Explorer for Microsoft SharePoint is installed on a server on which both Veeam Backup & Replication and Veeam Backup for Microsoft Office 365 are installed, the notification settings will be inherited from the Veeam Backup & Replication Global Notification settings.
- For VMs that run Microsoft SQL Server backed up by Veeam Agent for Microsoft Windows and Microsoft SharePoint backed up by Veeam Backup & Replication, to perform SharePoint item recovery, you must start Veeam Explorer for Microsoft SharePoint from the Veeam Backup Browser. To do this, start File Level Recovery for the backup of the required SQL Server VM, and then launch Veeam Explorer for Microsoft SharePoint from the toolbar in the Veeam Backup Browser.

  If you start item-level recovery directly for the backup of the SharePoint VM, mapping of a SharePoint site to a content database in the Restore wizard will fail, and Veeam Backup & Replication will display the following error: *Unable to find SQL Server VM hosting the content database for the selected SharePoint site*.

  Backup of the same VM by both Veeam Backup & Replication and Veeam Agent for Microsoft Windows may be required, for example, if Microsoft SQL Server running on the VM operates as part of a failover cluster.

Status Recovery Limitations

Consider the following when planning recovery of documents and list items:

- If a document/item was in *Check Out* state when the backup was created, item's last version will not be restored to the target SharePoint and will be available for viewing only. Previous versions (if any) will be restored.
- If the *Declare this item as a record* action was originally applied to list item, the corresponding status will not be preserved. Instead, the restored item status will be set in accordance with the target list/library content approval workflow.
- Original *On Hold* status will not be restored.

Restore of Documents, Libraries and Lists

- Versioning settings of SharePoint lists are not preserved during restore.
- Restoring Generic List and Pages Library may fail with the "*No content type 'XXX' found in web YYY*" error.
- *Modified By* field of restored documents is updated with the account performing restore.
- Some *Rating Settings* of *Discussion* lists values are not restored.
- Restore via PSDirect, VIX or Sphere API is not supported.
- Veeam Explorer for Microsoft SharePoint does not display hidden lists. Thus, such lists cannot be restored.

Restore of List Items

Consider that when restoring a list item, Veeam Explorer works in the following way:

- Deletes an existing item.
- Creates item's latest version anew using data from the backup.
- Checks whether it is declared as a record.
- If the check is a success, the process finishes.
- If not, the created version is deleted and item versions are restored sequentially.
- Restore via PSDirect, VIX or Sphere API is not supported.

This logic leads to several peculiarities of list item restore, as described below.

If a list or list items column is used as a lookup column in the dependent list, consider that restoring an item from the source list causes corresponding items deletion in the dependent list.

If a lookup column in the dependent list has the enforced relationship behavior set to **Cascade Delete**, then restoring an item from the source list may cause item deletion in the dependent list due to Microsoft implementation. For more information, see this Microsoft article.

To prevent this issue, it is recommended to turn off enforced relationship behavior. As a work-around, you can roll-back the SharePoint database using Veeam Explorer for Microsoft SQL Server (as described in the corresponding section of this guide) or roll-back the whole SharePoint server VM to the desired state using any Veeam Backup & Replication recovery option.

The following limitations should also be considered:

- If a lookup column in the dependent list has the enforced relationship behavior set to **RestrictDelete**, then item restore with Veeam Explorer will fail.
- If an .ASPX page references an item using ItemID, this reference may fail to restore (as the item will be created anew with a different ItemID).
- If a list item cannot be deleted (for example, site's "Welcome" page), consider that restore process will recover all versions of the item sequentially without deletions, adding them to *Version History*.
- *Restored Issue* list items are assigned new Issue ID.
- Restore of *Time Card* list is not supported.

Restore of List Items with Links (Attachments)

Consider the following when planning for the recovery of list items with links (attachments):

- If the retention policy for target list/document library was configured to **Declare record** automatically, only the last version of the item will be restored to target list/library. Target retention policy settings will be applied to restored item (**Declare record**). However, links (attachments) will not be restored.
- Alternatively (with different retention policy settings), all versions of the original item will be restored to target list/library; item links (attachments) will be restored only if such item does not exist on target SharePoint.
- If *Declare this item as a record* action was originally applied to list item, such item will not be restored.
- Restore via PSDirect, VIX or Sphere API is not supported.

Restore of Surveys

Consider the following limitations when planning for the recovery of surveys, survey questions and responses:

- Survey item(s) can be restored to a new survey, created automatically by Veeam Explorer for Microsoft SharePoint in the specified destination instead of the previously deleted survey. However, if a new survey is created by user from scratch (not replacing a deleted one) – items cannot be recovered to such a survey.
- A survey can be restored to an existing target survey only if that target survey includes at least one item (question) same as survey questions stored in the content database.
- If a survey question was not answered completely in the source survey, after restore the response status in the target survey will be set to *Completed*, anyway.
- When restoring a single response to a survey, target response item with the same number will be deleted and restored item will be placed in the target survey after the last numbered response.
  For example, if the target survey has responses #1-15 and you try to restore a response that used to be #6 on source – then target response #6 will be deleted and restored response will be assigned #16.
- Restore via PSDirect, VIX or Sphere API is not supported.

Restore of Sites

- If you plan to restore SharePoint site pages, consider that Veeam Explorer does not support recovery of items which are not stored in the SharePoint content database (in particular, pages, page references and items based on default templates). Such items cannot be restored (neither by **Restore** nor by **Save** menu option), except for data from Wiki Content (text and images) which is stored in the database. Thus, site pages containing only text and/or images can be restored and displayed properly.

> **NOTE**
>
> Consider the following:
>
> - In case you attempt to restore such items, the following error message will be logged: "*Item <item> is skipped: restoration of items based on SharePoint default template is not supported*."
> - In case you attempt to save such items, the following error message will be logged: "*Unable to save document <item>. Document content is not available*."

- Web features restore is not supported for SharePoint 2010 sites.

Export and Import

- Importing exported Picture Library may result in IDs changed for some items.
- Importing exported Project Tasks list does not preserve column order.
- Importing exported SharePoint list does not preserve Validation Settings.

Data Type Limitations

Consider the column (field) data type when planning for the recovery of your SharePoint libraries/lists:

- If source column (field) data type was set to *Lookup*, but the referenced list/library was deleted, such columns (fields) will not be restored even if you recover that referenced list. The reason is that if referenced list is deleted, the reference (link) to that list is no longer valid.
- If source column (field) data type was set to *Managed Metadata*, such columns (fields) will not be restored.

Workflow-Related Considerations

You may need to restore the item(s) originally belonging to a list with no content approval required, to another list. If you try to restore such item(s) to a list that requires content approval, item version and status will be modified in the following way:

1. If a target list is configured to include major versions only, then all versions of restored item will become major (despite the original versioning); item status will be set to *Pending*.
2. If a target list is configured to include both major and minor versions, then all versions of restored item will become minor (despite the original versioning); item status will be set as follows:
    - If the last version (original) was major, status will be set to *Pending*.
    - If the last version was minor, status will be set to *Draft*.

Also, consider the following when planning for the recovery of list items (with or without content approval originally required):

1. If the retention policy for target list/document library is configured to **Declare record** automatically, only the last version of the item will be restored to target list/library. Target retention policy settings will be applied to restored item (**Declare record**).
   Besides, if **Require content approval for submitted items** was enabled for the original list, then after recovery item status will be set to *Pending*.
2. Alternatively (with different retention policy settings), all versions of the original item will be restored to target list/library.
   Besides, if **Require content approval for submitted items** was enabled for the original list, then after recovery item status in the content approval workflow will be also restored, except for the states listed (see "Status Recovery Limitations" above).

Launching Application and Exploring Backups

To open Veeam Explorer for Microsoft SharePoint and load backups, you can use any of the following methods:

- The **Restore application item** option to load backups created in Veeam Backup & Replication.

  For more information, see the Application Items Restore section of the Veeam Backup & Replication User Guide.

- The **Explore** option to load backups created in Veeam Backup for Microsoft Office 365.

  For more information, see the Data Restore section of the Veeam Backup for Microsoft Office 365 User Guide.

  When loading backups located in object storage repositories, Veeam Explorer for Microsoft SharePoint will notify you about the location of these backups. You can select the **Don't show this message again** check box to unsubscribe from such notifications. To enable notifications, select the **Show pop-up notifications on startup** check box, as described in Configuring Advanced Settings.

- Go to **Start**, click **Veeam Explorer for Microsoft SharePoint** and manually open Microsoft SharePoint databases, as described in the Standalone Databases Management section.

If you open Veeam Explorer for Microsoft SharePoint from the **Start** menu on a machine that runs the Veeam Backup & Replication console and Veeam Explorers only, specify the following:

- The domain name or IP address of a server to which you want to connect
- The port number
- User credentials

  The account must be a member of the *Local Administrator* group on a target server. To use the account under which Veeam Explorer for Microsoft SharePoint is running, select **Use Windows session authentication**.

To save the connection shortcut to the desktop, click **Save shortcut** in the bottom-left corner.

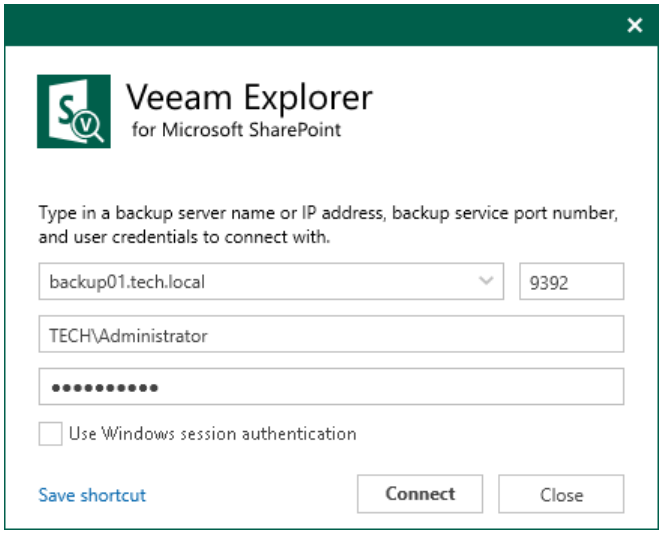

Understanding User Interface

Veeam Explorer for Microsoft SharePoint provides you with the convenient user interface that allows you to perform required operations in a user-friendly manner.

Main Menu

The main menu comprises the following features:

- **General Options**. Allows you to configure program options.
- **Help**.
  - **Online help**. Opens the online web help page.
  - **About**. Shows information about the product.
- **Exit**. Closes the program.

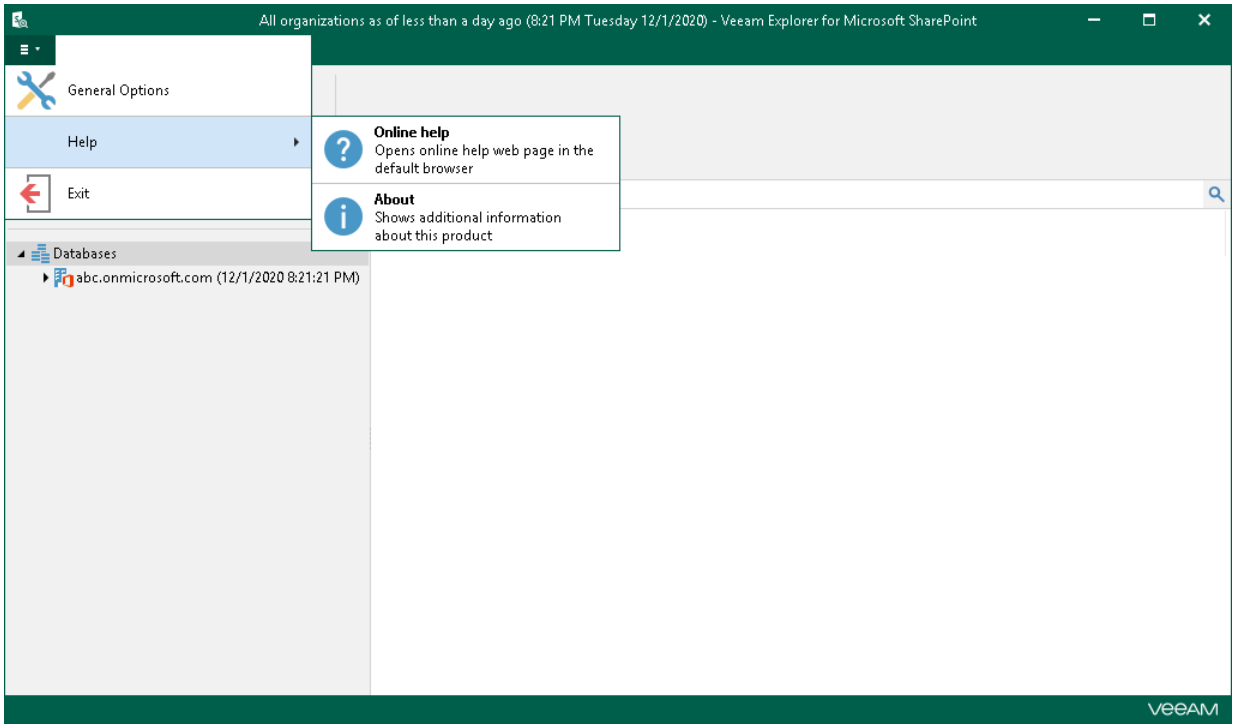

Main Application Window

The main application window can be divided into three categories:

- The ribbon menu, which contains general program commands organized into logical groups.
- The navigation pane, which allows you to browse through the hierarchy of your backup files.
- The preview pane, which shows you the details about objects you have selected in the navigation area.

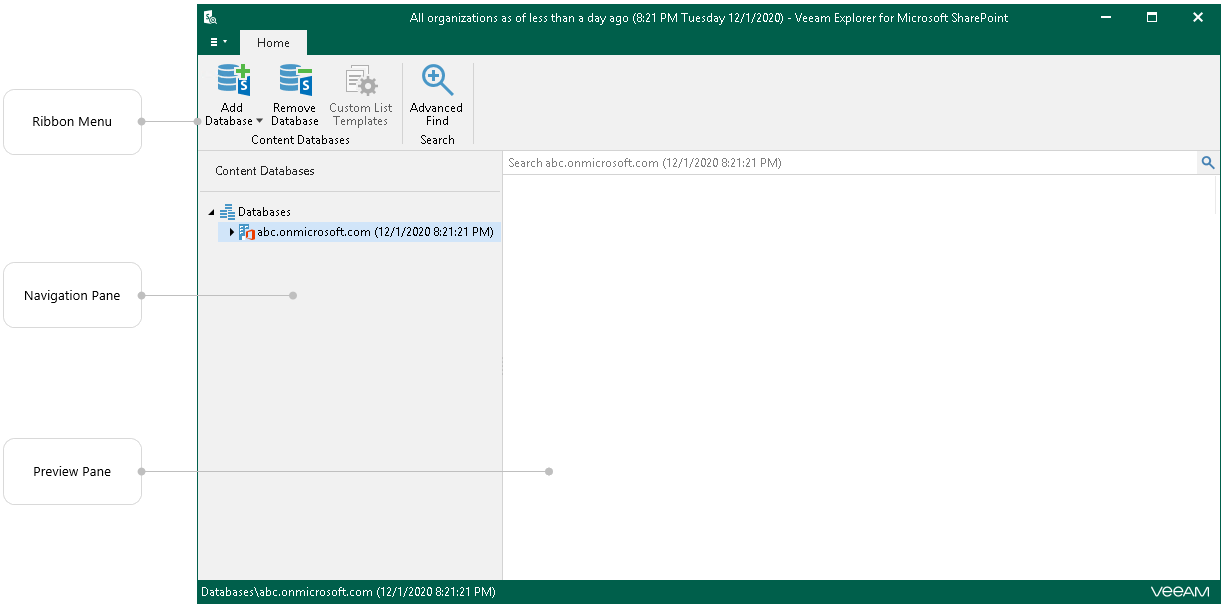

Browsing, Searching and Viewing Items

Continue with this section to learn about:

- Browsing backup content
- Viewing objects properties and open files
- Searching for objects in a backup file
- Using the advance search capabilities

Browsing

To view the contents of a backup file, you use the navigation pane which shows you the database structure containing your site items such as libraries and subsites.

After you select an object in the navigation pane, you can see its content in the preview pane.

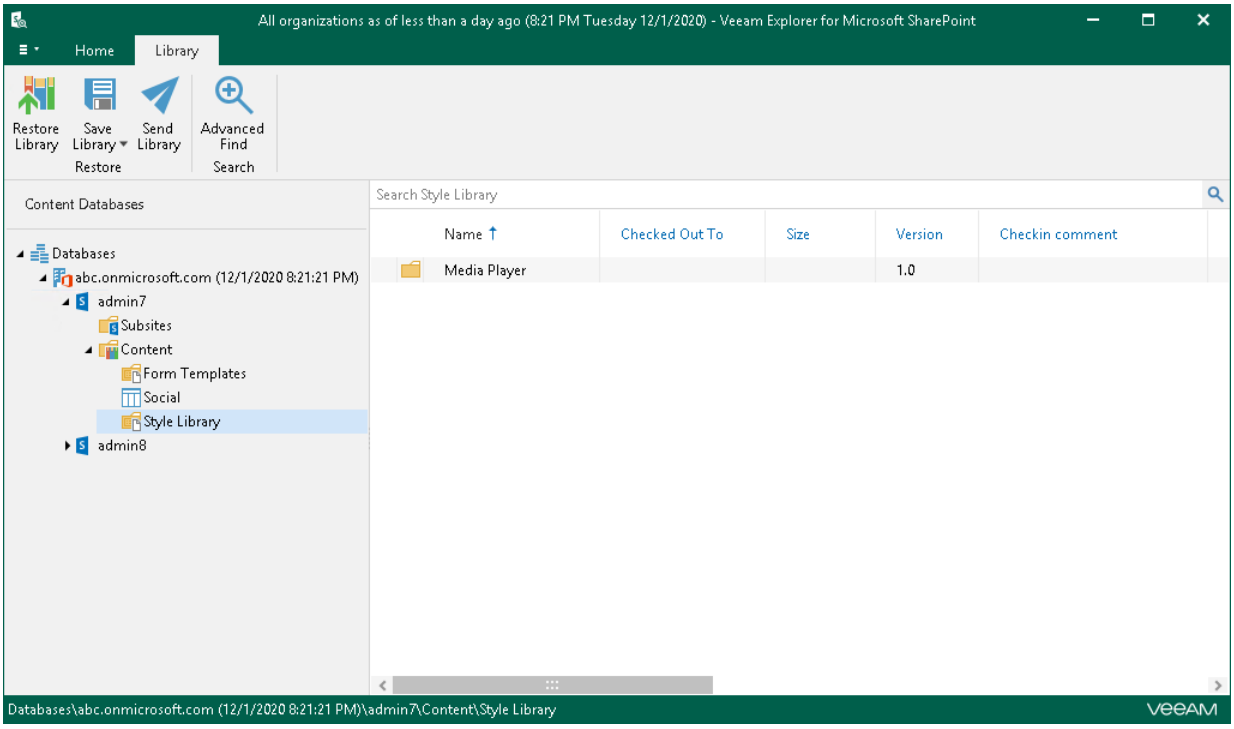

Viewing Properties and Opening Files

To view object properties, right-click an object in the preview pane and select **View Properties**.

To open a document using an associated application, right-click a document in the preview pane and select **Open**.

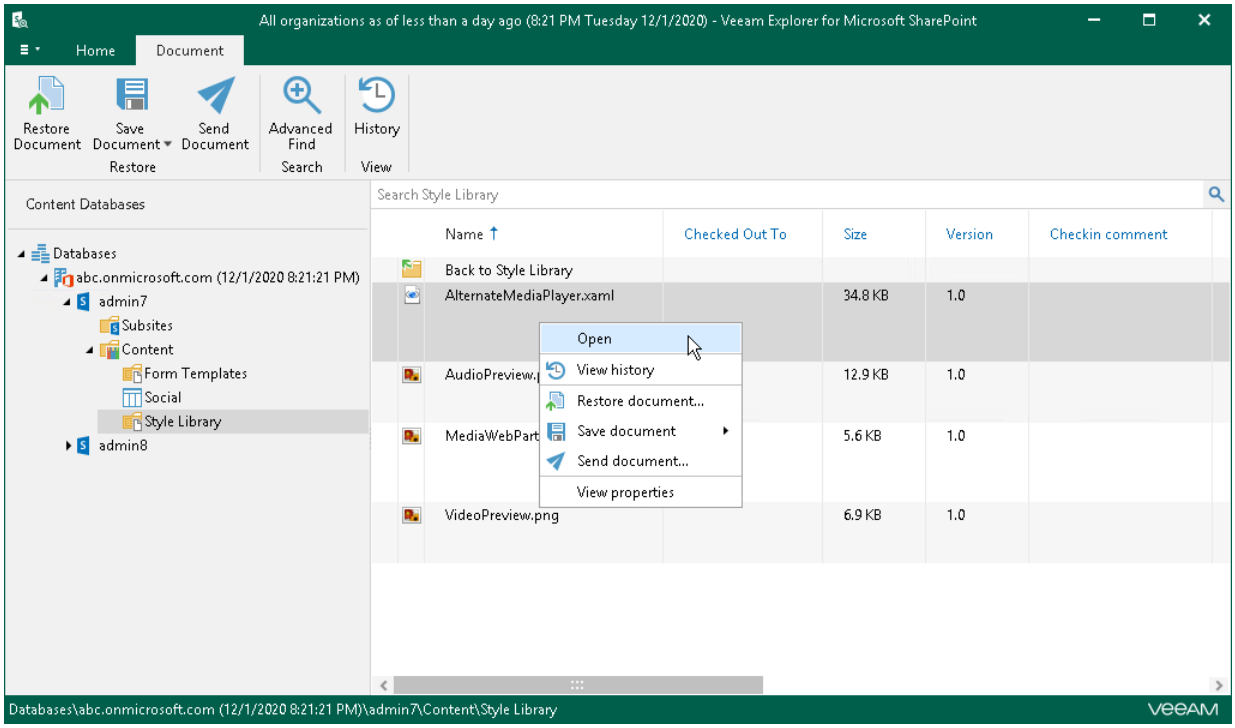

Searching

The search mechanism allows you to find items matching specified search criteria.

To search for required items, do the following:

1. In the navigation pane. select an object in which you want to find your data.
2. Enter a search query to the search field at the top of the preview pane.

> **NOTE**
>
> To find the exact phrase, use double quotes. For example, *"media player"*.

You can narrow your search results by specifying various search criteria using the *criteria:value* format.

*F*or example, to find all items that require approval in the list of decisions, you can use the following search query: *status:pending approval*.

You can also use logical upper-cased operators such as *AND*, *OR* and *NOT* along with wildcard characters such as *** and *?*. The search criteria are similar to those used for searching in Microsoft SharePoint. For more information, see this Microsoft article.

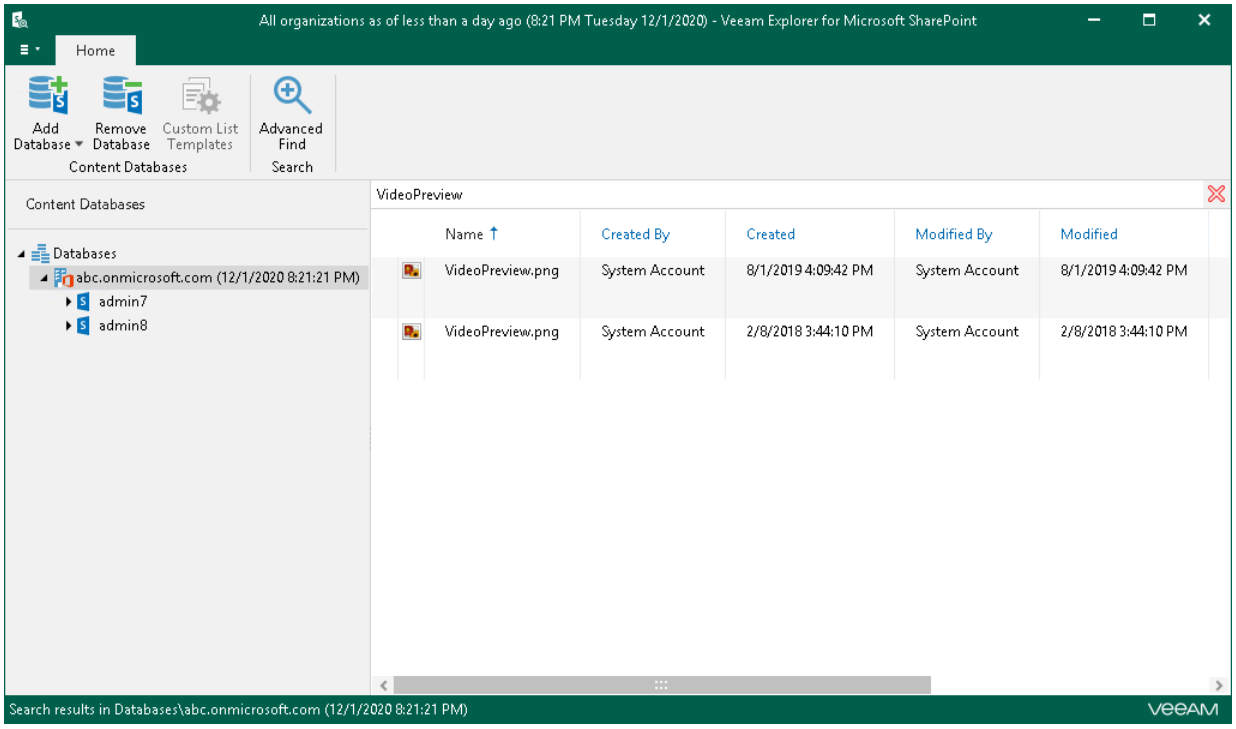

Using Advanced Find Capabilities

The **Advanced Find** mechanism allows you to define your search criteria more precisely.

For example, to find an object that starts with the word *Media*, do the following:

1. In the preview pane, select a content node and click **Advanced Find**.
2. In the **Define search criteria** section, select **Category** > **Document fields**.
3. In the **Field** list, select **File Name**.
4. In the **Condition** list, select **Starts With**.
5. In the **Value** field, specify a file name.
6. Click **Start**.

To remove a filter, click the cross mark next to it. To remove all configured filters, click **Reset**.

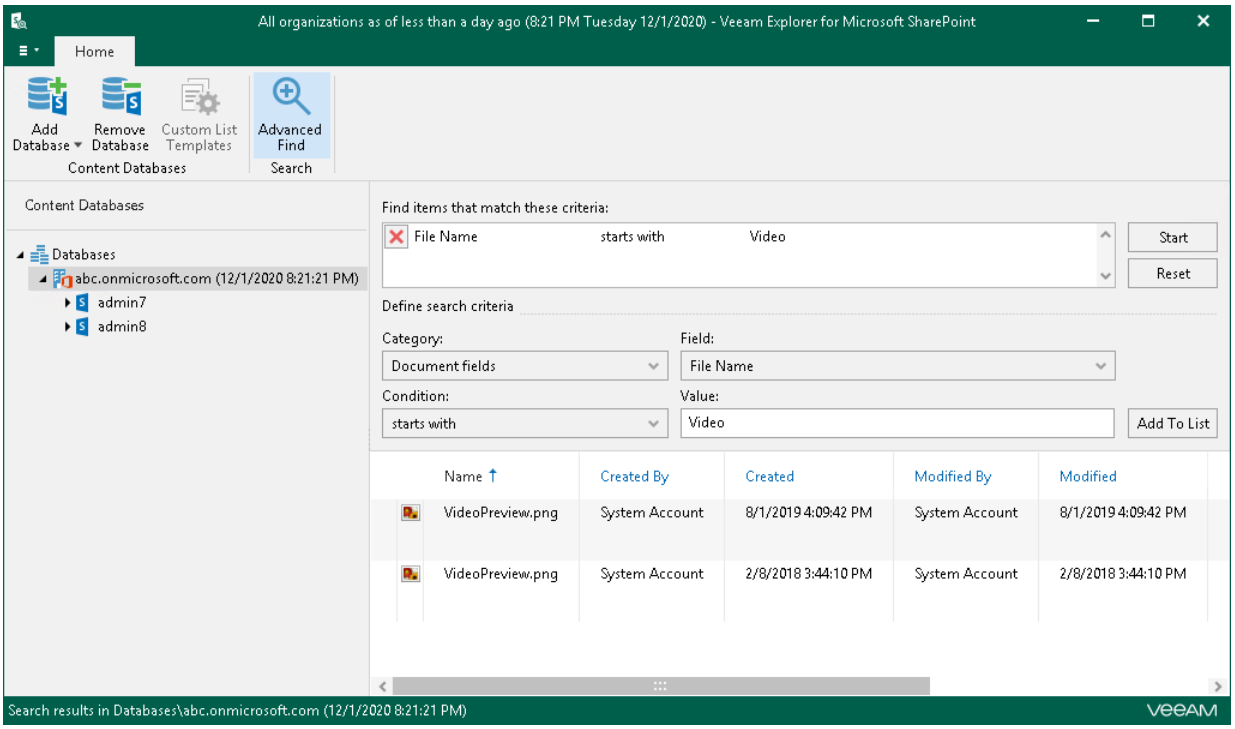

General Application Settings

Continue with this section to learn more about configuring required application settings and components.

Configuring Staging SQL Server

Veeam Explorer for Microsoft SharePoint requires a staging Microsoft SQL Server for the following operations:

- To browse VM backups that contain Microsoft SharePoint data and perform recovery of Microsoft SharePoint items.
- To add standalone Microsoft SharePoint databases to Veeam Explorer, as described in Adding Microsoft SharePoint Databases.

A machine used as a staging SQL Server for browsing and restore operations can reside in the same domain as the machine hosting Veeam Explorer, as well as in a trusted domain or untrusted domain.

For the add standalone database operation, consider the following:

- A machine used as a staging SQL Server can reside in the same domain as the machine hosting Veeam Explorer or in a trusted domain. You cannot use a staging server that belongs to an untrusted domain to add a standalone database to Veeam Explorer.
- Both *Windows authentication* and *SQL Server authentication* methods are supported for the staging SQL Server. If you want to use *Windows authentication*, complete the following steps to configure delegation settings:
    a. In **Active Directory Users and Computers**, select the necessary staging SQL server.
    b. Open server properties and select the **Delegation** tab.
    c. Select **Trust this computer for delegation to specified services only** and **Use any authentication protocol** options for the *cifs* service on a computer with Veeam Explorer.
    d. Restart the staging SQL Server VM.
    e. Select a user account to connect to the staging SQL server. Also, select its properties on the **Account** tab and make sure the **Account is sensitive and cannot be delegated** check box is cleared.

To configure a staging server, do the following:

1. Go to the main menu and click **General Options**.
2. Go to the **SQL Server Settings** tab.
3. Specify a SQL server you want to use as a staging server.

   Click **Browse** to select a server, as described in Browsing For Servers.
4. Specify the account to access the selected SQL server.
5. Specify the account to access a SQL server instance.

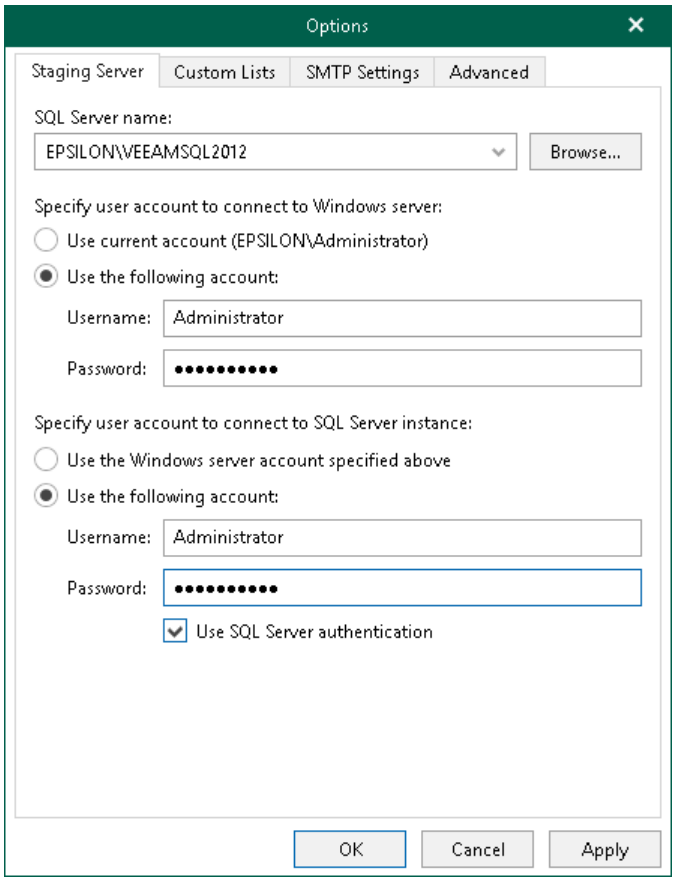

Browsing for Servers

To browse for a server, do either of the following:

- On the **Local Servers** tab, select a local SQL server that is located on a machine with Veeam Explorer.
- On the **Network Server** tab, select a SQL server over the network.

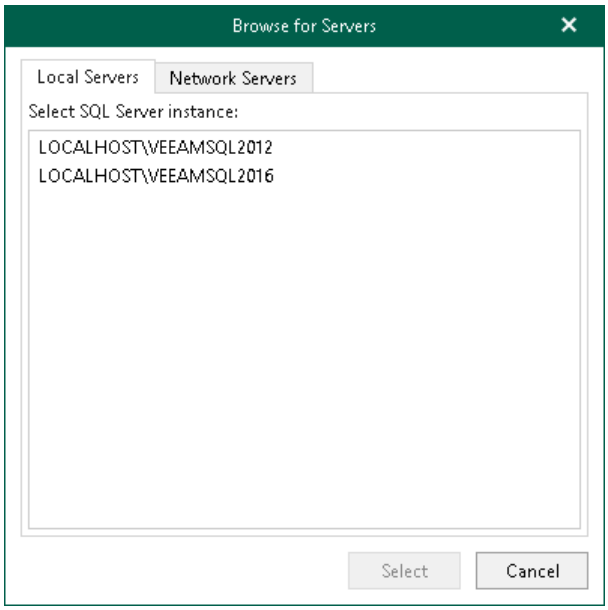

Configuring Custom Lists

Veeam Explorer for Microsoft SharePoint supports obtaining custom list templates from backups created in Veeam Backup & Replication.

The actual information about templates is stored directly in a backup file and obtained automatically so that you can view the custom list data without having to perform any additional operations.

> **NOTE**
>
> Configuring custom lists is not required when working with backups created in Veeam Backup for Microsoft Office 365.

To see templates with their corresponding names and features, go to the main menu, select **General Options** and click the link under the **Status** column which also comprises the total number of available templates.

You can select **Automatically update custom list templates from backups** to automatically update existing templates. Also, you can select **Warn me if attached database has new custom list templates** to be notified if the database contains new templates other than those you already have.

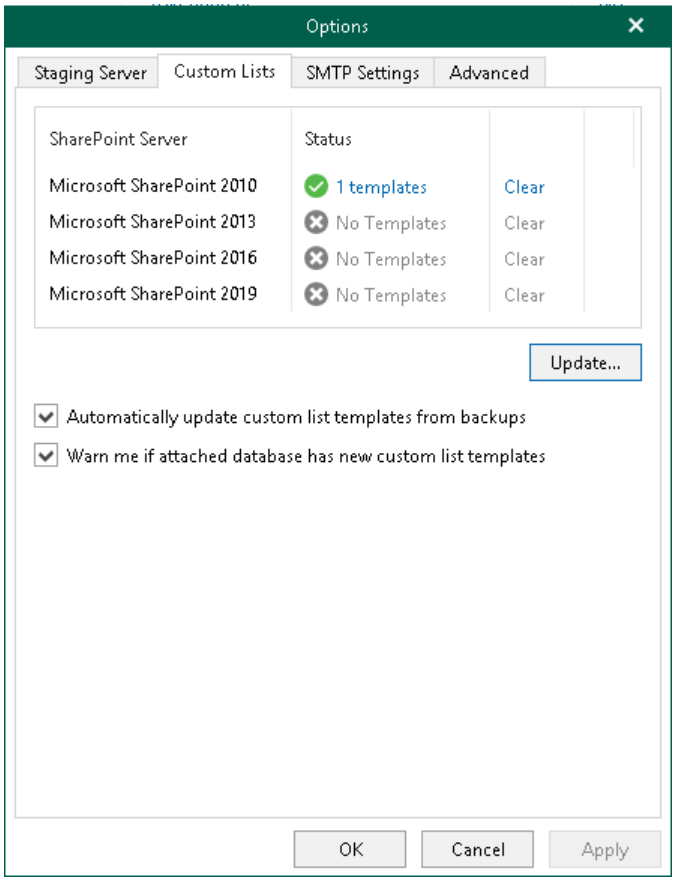

See the figure below showing you what custom list templates are available.

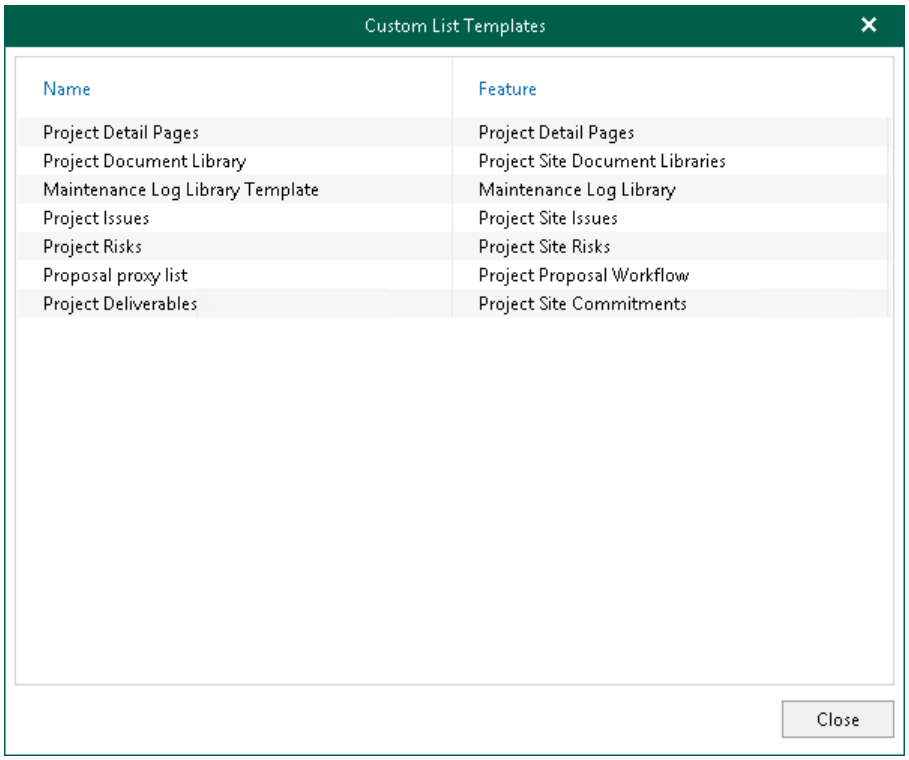

To remove templates, click **Clear**.

Importing Templates Manually

To import existing templates from the Microsoft SharePoint server manually, do the following:

1. Go to the main menu, select **General Options** > **Custom Lists** and click **Update**.
2. Specify the path to Microsoft SharePoint templates, select the language and click **OK**. The language set in the drop-down list depends on the installed language packages on a SharePoint server.

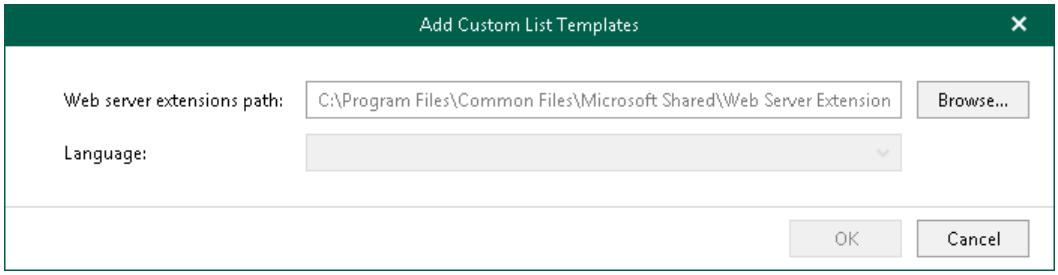

Importing Templates Using Ribbon Menu

> **NOTE**
>
> Available only for Microsoft SQL Server databases and requires a staging SQL server. For more information on configuring a staging server, see Configuring SQL Server Settings.

To import existing templates from the Microsoft SharePoint server using the ribbon menu, do the following:

1. In the navigation pane, select a SQL database.
2. Click **Custom List Templates** on the ribbon menu or right-click a SQL database and select **Custom list templates settings**.

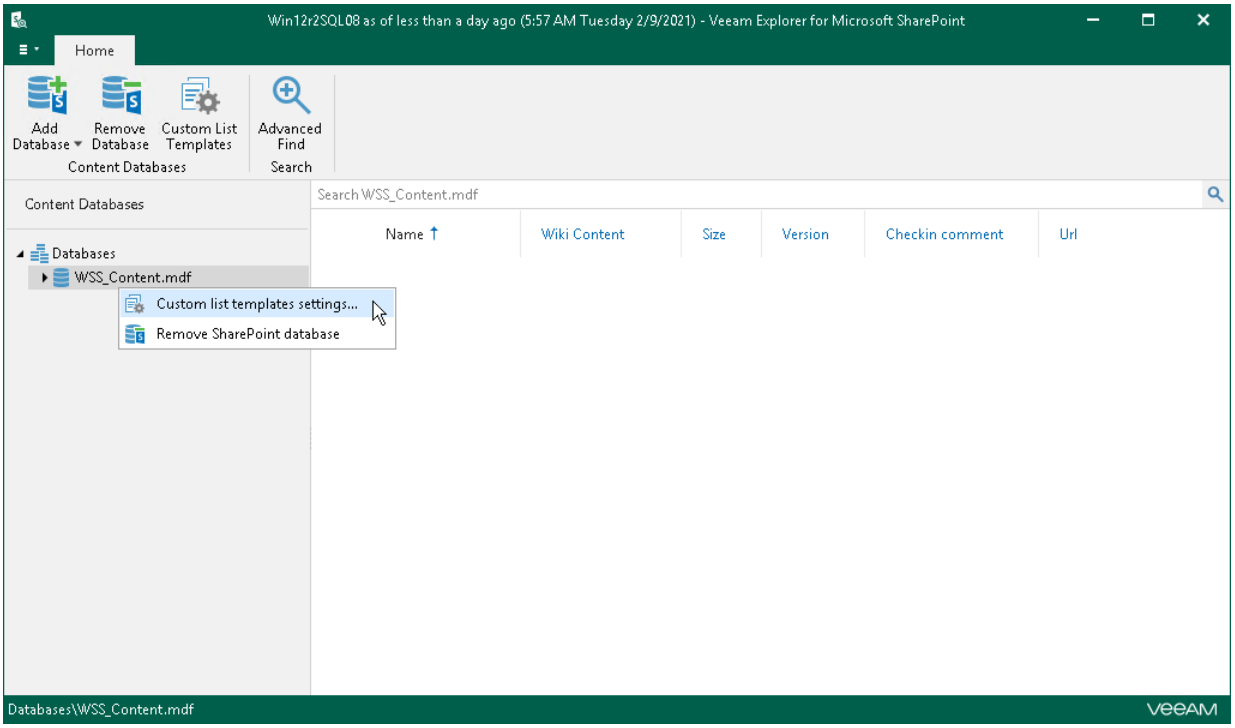

3. Choose how you want your templates to be applied:
   a. Select **Use global custom list settings** to apply global list templates to the select database.
   b. Select **Use the following custom list settings** to apply the templates to the selected database only.

      When using the latter option, specify the path to the Microsoft SharePoint templates, select the language and click **OK**. The language set in the drop-down list depends on the installed language packages on a SharePoint server.

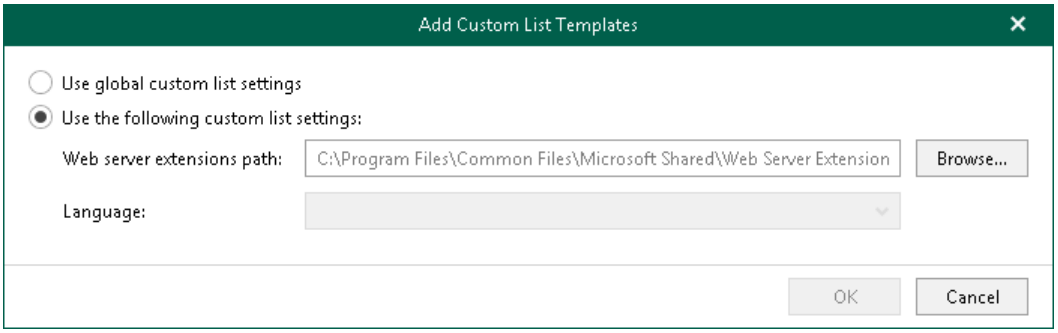

Configuring SMTP Settings

To send Microsoft SharePoint items as attachments, you must configure SMTP server settings.

To configure SMTP settings, do the following:

1. Go to the main menu and click **General Options**.
2. On the **SMTP Settings** tab, select the **Configure SMTP settings** check box and specify the following:
   - DNS name or IP address of the mail server.
   - SMTP communication port.
   - The sender email address. This address will appear in the **From** field when sending SharePoint items.
   - Select **Use authentication** check box if your SMTP server requires SMTP authentication for outgoing mail, and provide valid credentials.
   - Select **Enable SSL security** check box to enable SSL data encryption.
3. Click **Send** to send a test email message.
4. Click **Apply**.

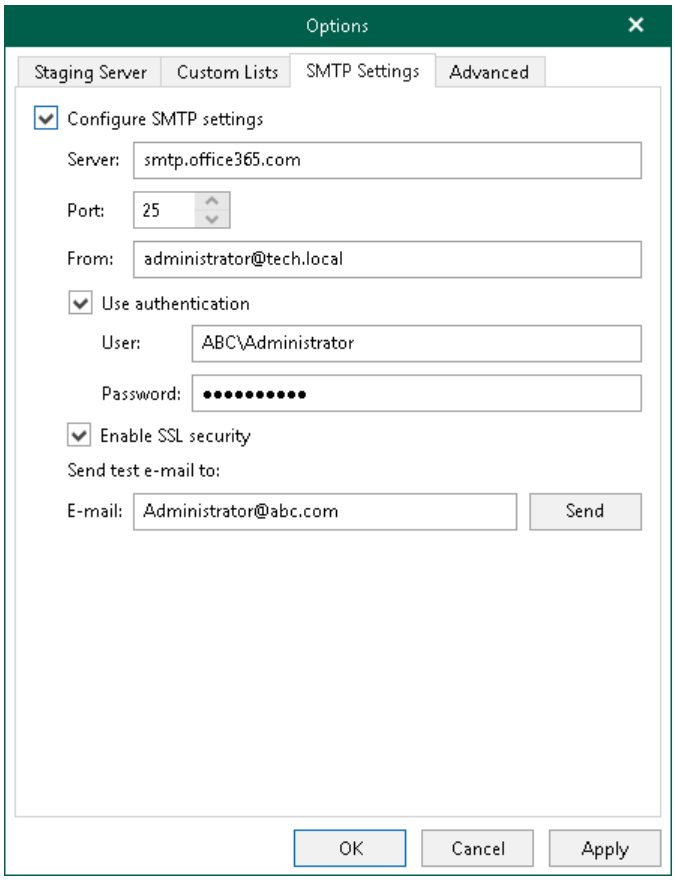

Configuring Advanced Settings

This section explains how to configure advanced settings of the application.

To configure advanced settings, do the following:

1. In the main menu, click **General Options**.
2. Open the **Advanced** tab.
3. If you want to enable extended logging mode to collect logs that contain more details on specific operations, select the **Enable extended logging** check box.

   After enabling extended logging mode, you can go back to the application and perform actions for which you want to collect additional information. Then you can collect logs.
4. By default, Veeam Explorer for Microsoft SharePoint displays notifications when loading backups from object storage repositories. If you want to disable these notifications, clear the **Show pop-up notifications on startup** check box.

   Applicable to backups created in Veeam Backup for Microsoft Office 365.

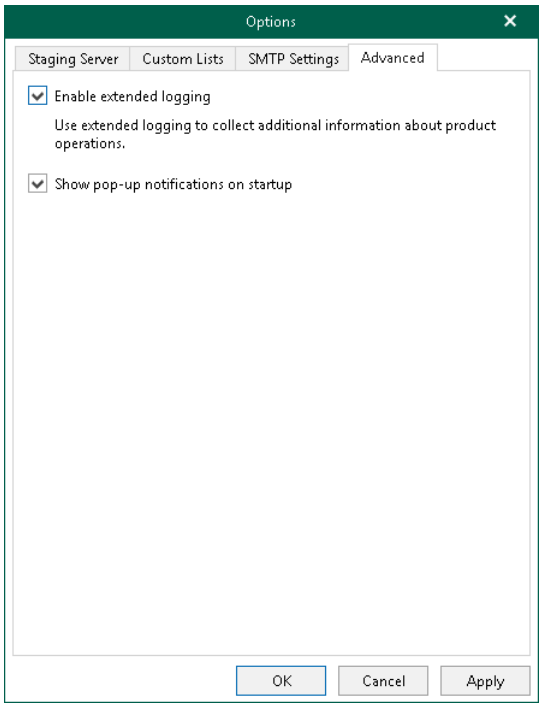

Standalone Databases Management

Continue with this section to learn more about adding and removing Microsoft SharePoint databases.

Adding Microsoft SharePoint Databases

You can add Microsoft SharePoint databases to the application scope manually. For example, you may want to do this to recover items from a Microsoft SharePoint database exported earlier using Veeam Explorer for Microsoft SQL Server. For more information, see Export as MDF.

When you add a standalone Microsoft SharePoint database to the application scope, Veeam Explorer for Microsoft SharePoint attaches the database to the staging SQL server, creating a temporary Microsoft SharePoint content database from which you can recover your items.

To manually add Microsoft SharePoint databases to the application scope, do the following:

1. Click **Add Database** > **Microsoft SharePoint Databases** on the toolbar or right-click the root **Database** node and select **Add database**.
2. Specify the location of the Microsoft SharePoint primary content database file (`.mdf`). The secondary database and the transaction log file (`.ldf`) will be added as well.

   To add remote BLOB stores, click **Add** next to the **Remote BLOB Stores** section.
3. Click **OK**.

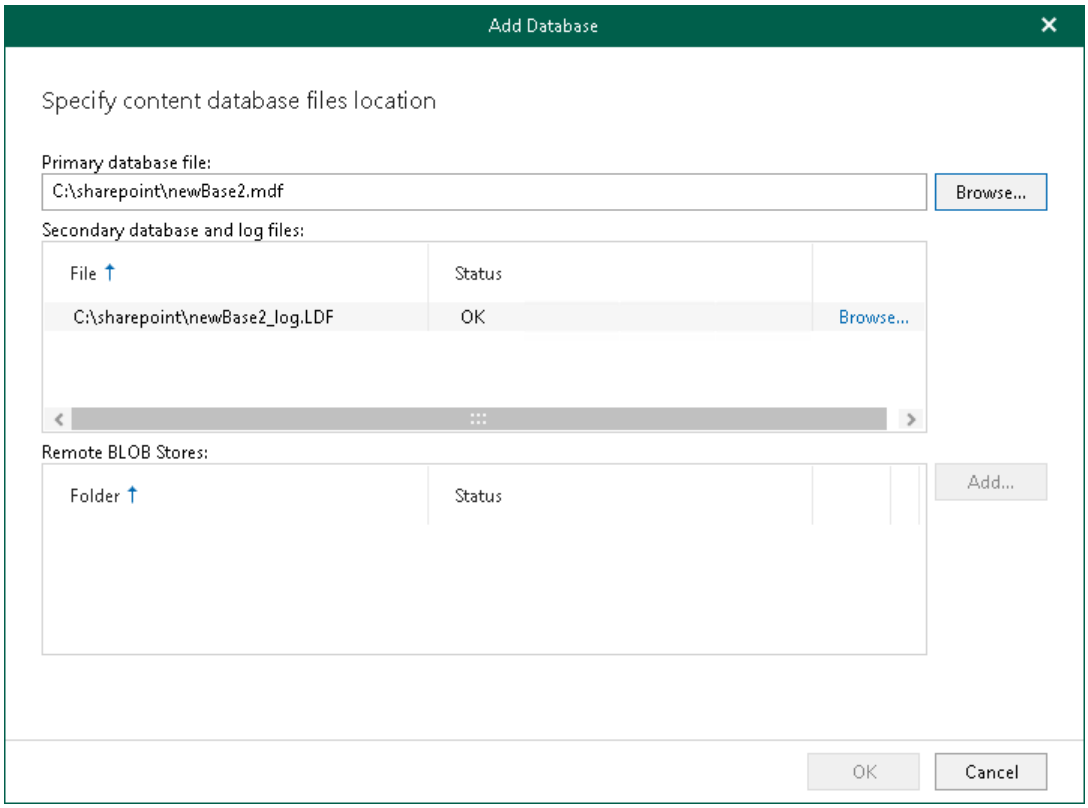

Adding Veeam Backup for Microsoft Office 365 Databases

To manually add databases that store Microsoft Office 365 organization data, do the following:

1. Click **Add Database** > **Veeam Backup for Microsoft Office 365 Databases** or right-click the root **Database** node and select **Veeam Backup for Microsoft Office 365 databases**.
2. Specify the database file location and log directory.
3. Click **Open**.

> **NOTE**
> Make sure you have disabled the **Veeam Backup Proxy for Microsoft Office 365** service when adding local databases. You can stop this service by using the `services.msc` console. If you try to add a database having this service still in progress, you will receive an error message and will not be able to access the database due to database lock.

Adding Veeam Backup for Microsoft Office 365 Server

You can use the built-in Veeam Explorer abilities to connect to another Veeam Backup for Microsoft Office 365 server and add its databases to the Veeam Explorer for Microsoft SharePoint scope.

To connect to another Veeam Backup for Microsoft Office 365 server remotely, do the following:

1. Click **Add Database** > **Veeam Backup for Microsoft Office 365 Server** on the ribbon menu or use the corresponding context menu command.
2. Specify connection settings under which to connect to the Veeam Backup for Microsoft Office 365 server and click **Connect**.

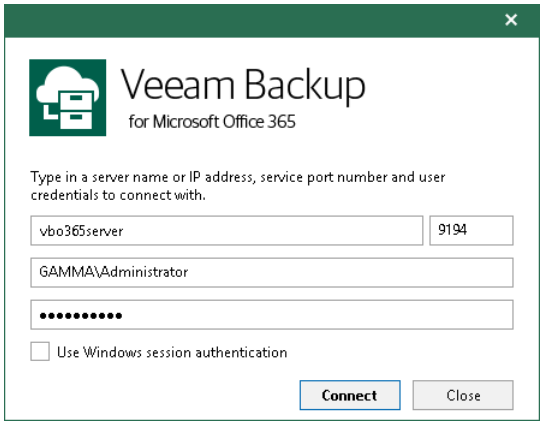

Adding Veeam Backup for Microsoft Office 365 Service Provider

You can use Veeam Explorer for Microsoft SharePoint to connect to a service provider server to add its backups to the application scope and perform a restore directly from such backups.

Consider the following:

- Both Veeam Explorer for Microsoft SharePoint and Veeam Backup & Replication must be installed on the same server.
- At least one service provider must be added to Veeam Backup & Replication, as described in the Connecting to Service Providers section of the Veeam Cloud Connect Guide.

To connect to a Veeam Backup for Microsoft Office 365 service provider, do the following:

1. From the **Start** menu, launch Veeam Explorer for Microsoft SharePoint.
2. On the ribbon, click **Add Database** > **Veeam Backup for Microsoft Office 365 service provider**.

   Alternatively, in the preview pane, right-click **Databases** and select **Veeam Backup for Microsoft Office 365 service provider**.
3. In the **Select service provider and authentication method** section, select from the drop-down lists a service provider to which you want to connect and an authentication method: *Modern authentication* or *Basic authentication*.
4. [For *Modern authentication*] Select a region to which an Office 365 organization of the selected service provider belongs.
5. [For *Modern authentication*] Copy the device code and use it to sign in to Microsoft Azure Active Directory.

   For *Basic authentication*, you must enter a user name and password to connect to the service provider server.
6. If you want to select a point-in-time state as of which you want to load Office 365 backups, click **Specify point in time** and select the necessary option:
   - **Use the latest available state**. To load data as of the latest backup state.
   - **Use the following point in time**. To load data as of the selected point in time.

     To select a point in time, use the calendar control.

7. Select the **Show items that have been deleted by user** check box to see all the items that were removed.
8. Select the **Show all versions of items that have been modified by user** check box if you want to load modified versions.
9. Click **Connect**.

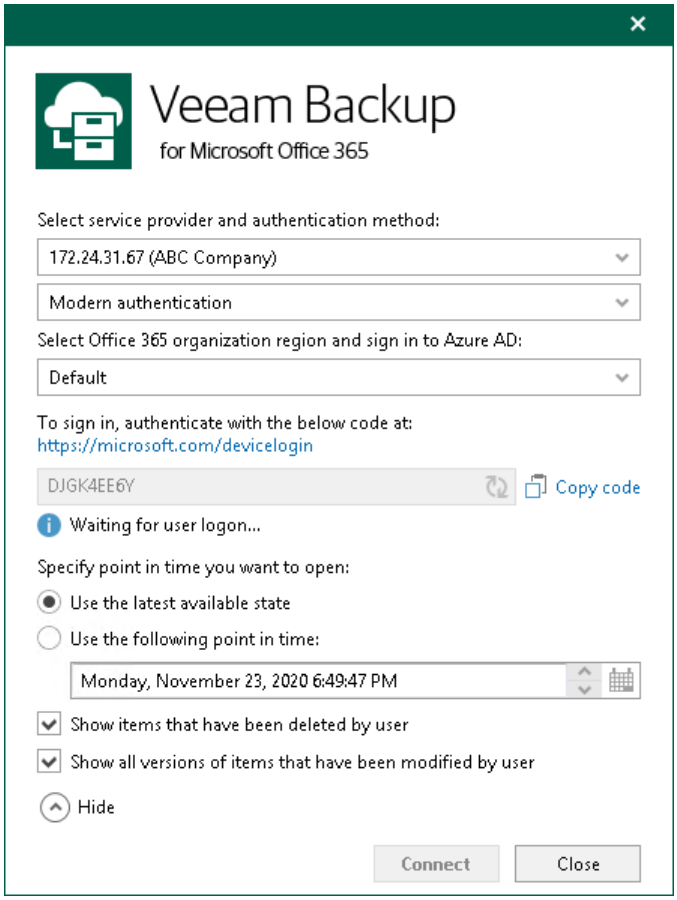

Removing Standalone Databases

Veeam Explorer for Microsoft SharePoint allows you to remove Microsoft SharePoint databases from the application scope when you no longer need it.

To remove a database from the application scope, right-click a database in the navigation pane and select **Remove SharePoint database**.

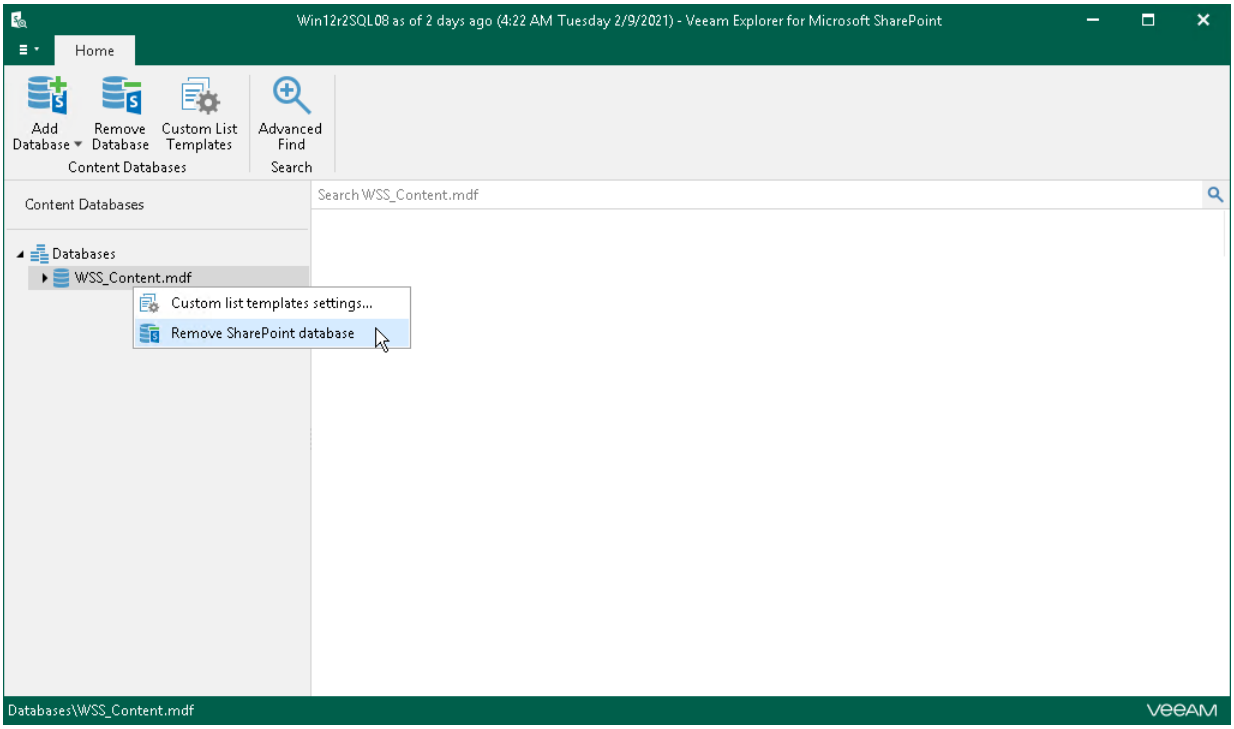

Data Restore

This section explains how to restore Microsoft SharePoint data.

Consider the following:

- To use an internet proxy server to restore backups created in Veeam Backup for Microsoft Office 365, make sure to provide appropriate proxy server address and the port number. For that, go to the **Control Panel** > **Internet Options Connections** tab, click **LAN Settings**, select the **Use a proxy server for your LAN** check box and specify a proxy server you want to use. Credentials for such a proxy (if needed) will be taken from the **Control Panel** > **Credential Manager** > **Windows Credentials** console.

  Using an internet proxy server is only possible in Veeam Explorer for Microsoft Exchange that comes as part of the Veeam Backup for Microsoft Office 365 distribution package.

- Make sure to read Unsupported SharePoint Lists before restoring Microsoft SharePoint lists.

Restoring Microsoft SharePoint Sites

To restore Microsoft SharePoint sites, do the following:

1. Launch the Restore wizard.
2. Specify credentials.
3. Log in to Office 365 (for data restore from backups created in Veeam Backup for Microsoft Office 365).
4. Specify the target site.
5. Specify restore options.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

1. In the navigation pane, select a site.
2. On the **Site** tab, select **Restore Site** or right-click an object and select **Restore site**.

> **NOTE**
>
> Sites can only be restored to the existing site collection. Creating a new collection is not supported.

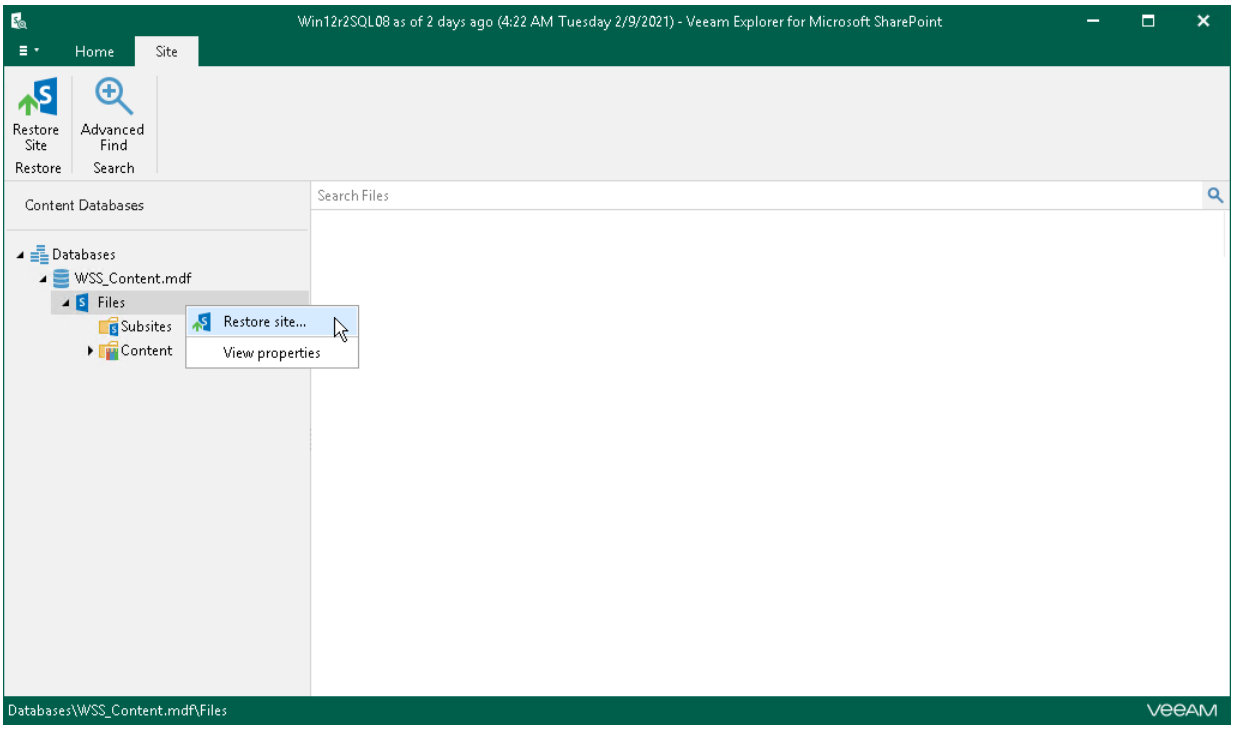

Step 2. Specify Credentials

At this step of the wizard, specify target server credentials. The credentials dialog differs depending on the organization type you are restoring from.

Credentials for Office 365 Organizations with Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified object back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft SharePoint populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

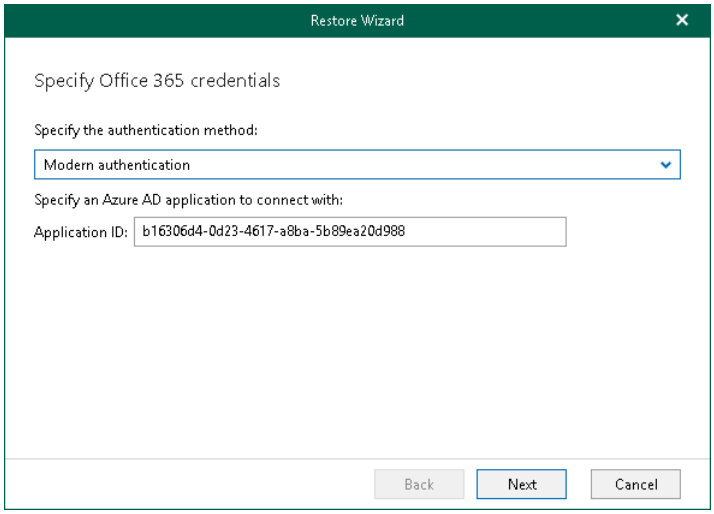

Credentials for Office 365 Organizations with Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the SharePoint organization.

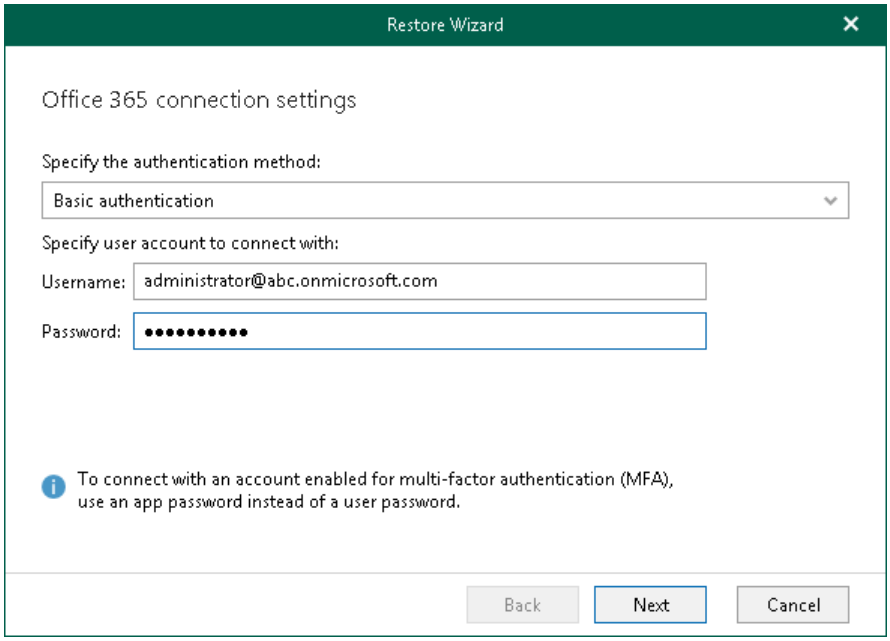

Credentials for On-Premises SharePoint Organizations

To authenticate on an on-premises Microsoft SharePoint server, provide a user name and password.

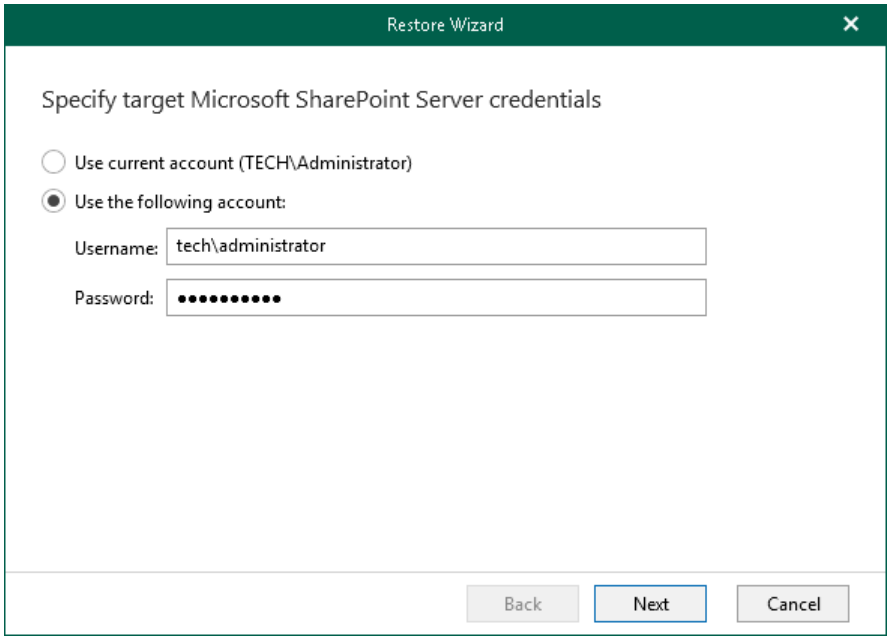

Restoring from Microsoft SQL Database

When restoring from a Microsoft SQL database, provide the following:

- The site URL to which you want to recover your site.
- Authentication credentials.

NOTE

When recovering Microsoft SharePoint sites from a Microsoft SQL database, you will be taken directly to the Specify Restore Options step.

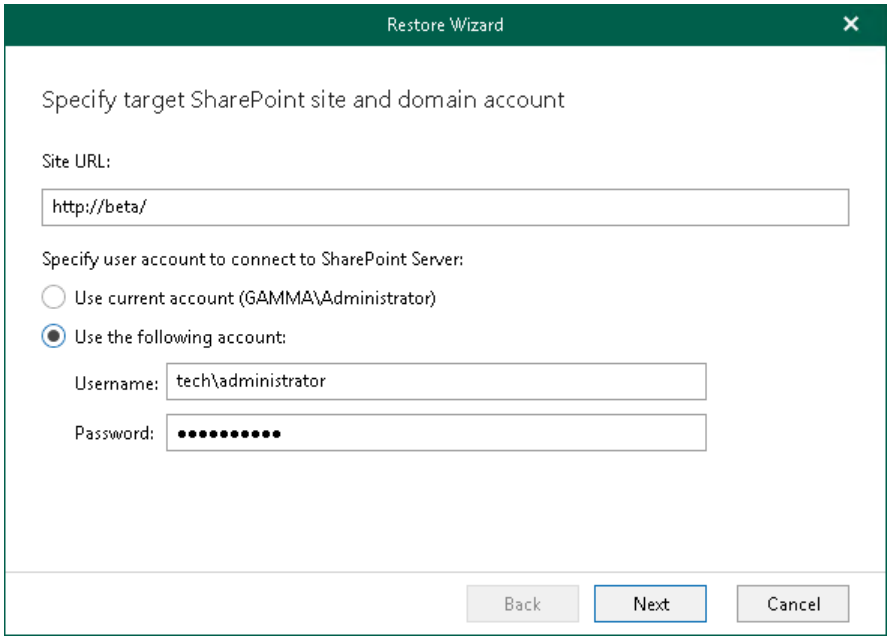

Step 3. Log In to Office 365

This step is available only if you restore data from a backup created in Veeam Backup for Microsoft Office 365 and have selected the *Modern Authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

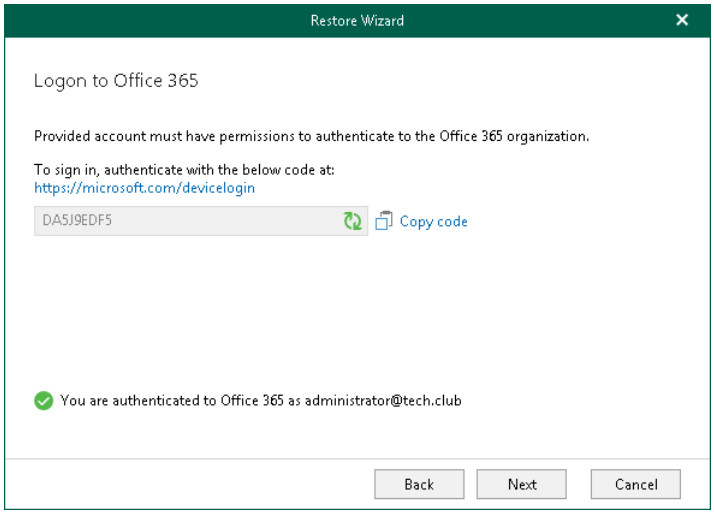

Step 4. Specify Target Site

At this step of the wizard. specify whether to restore a site back to the original location or choose site alias you want to use.

> **NOTE**
>
> Consider the following:
>
> - This step is available only when restoring from backups created in Veeam Backup for Microsoft Office 365.
> - This step is unavailable when restoring *Site Collections*.

You can select the following options:

- **Restore to the original site**. To restore data back to the original site.
- **Restore to the following site alias**. To restore data to a custom site using alias. If the provided site does not exist, it will be created automatically.

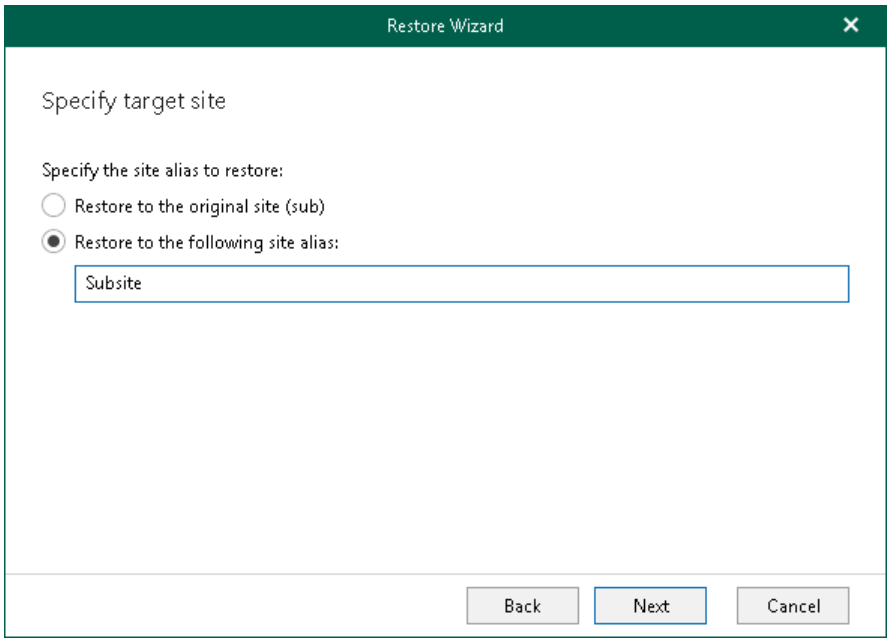

Step 5. Specify Restore Options

At this step of the wizard, specify restore options and click **Restore**.

The restore options are as follows:

- **Changed items**. Allows you to recover data that has been modified in your production environment.
- **Missing items**. Allows you to recover items that have been deleted in your production environment.
- **Restore permissions**. Allows you to recover permissions. If not selected, permissions will be inherited from the parent site.
- **Restore list views.** Allows you to recover your list views.
- **Restore subsites**. Allows you to recover subsites. If not selected, sites that are being restored will contain only document libraries and lists.
- **History restore options**. Allows you to select a version:
  - **Overwrite**. To overwrite data in the production environment by recovering only the latest version of the document from the backup.
  - **Merge**. To merge an existing version of the document with that of a backup version.

  If not selected, all the versions in the production environment will be replaced with the corresponding data from the backup file.

> **NOTE**
>
> The **Send a notification by e-mail to the users with permissions to the file** check box is available only when restoring data from backups created in Veeam Backup for Microsoft Office 365 for Microsoft Online organizations.

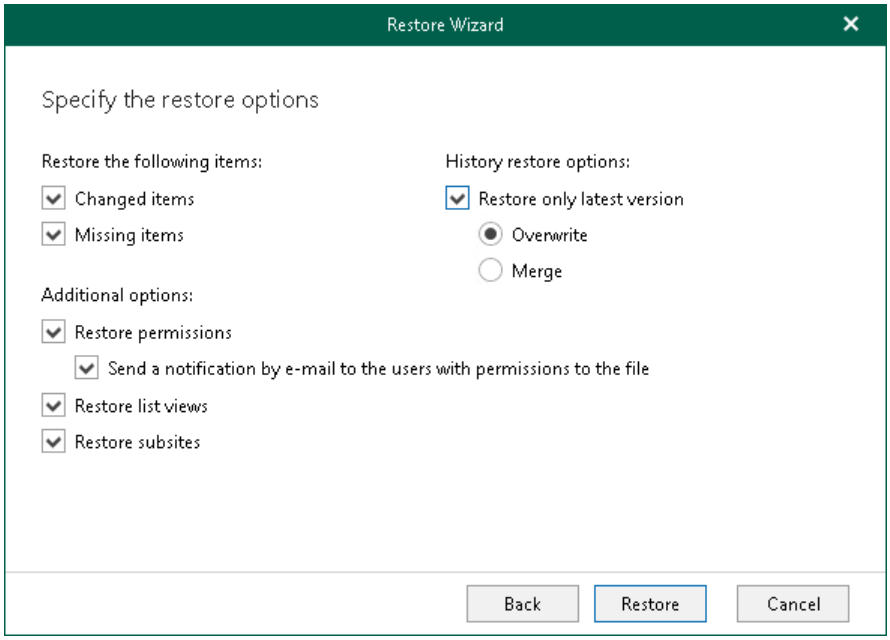

Restoring Microsoft SharePoint Document Libraries and Lists

To restore Microsoft SharePoint document libraries and lists, do the following:

1. Launch the Restore wizard.
2. Specify credentials.
3. Log in to Office 365 (for data restore from backups created in Veeam Backup for Microsoft Office 365).
4. Specify the target list.
5. Specify restore options.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

1. In the navigation pane, select a library or list.
2. On the **Library** tab, select **Restore Library/Restore List** or right-click an object and select **Restore Library/Restore List**.

> **NOTE**
>
> When restoring data from backups created with Veeam Backup for Microsoft Office 365, document libraries and lists can only be restored back to the source site from which they were backed up.

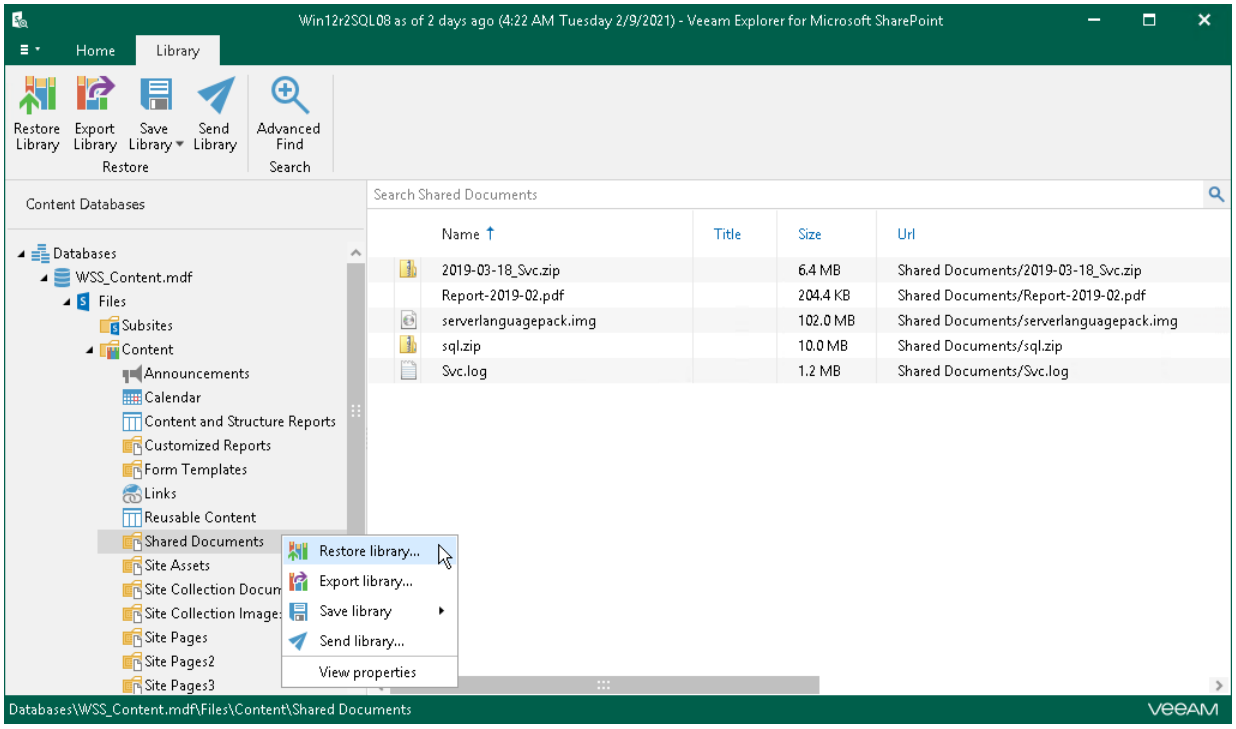

Step 2. Specify Credentials

At this step of the wizard, specify target server credentials. The credentials dialog differs depending on the organization type you are restoring from.

Credentials for Office 365 Organizations with Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified object back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft SharePoint populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

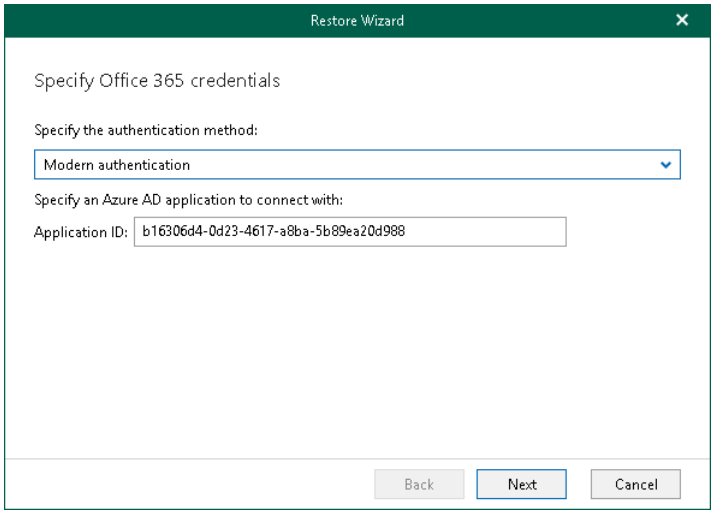

Credentials for Office 365 Organizations with Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the SharePoint organization.

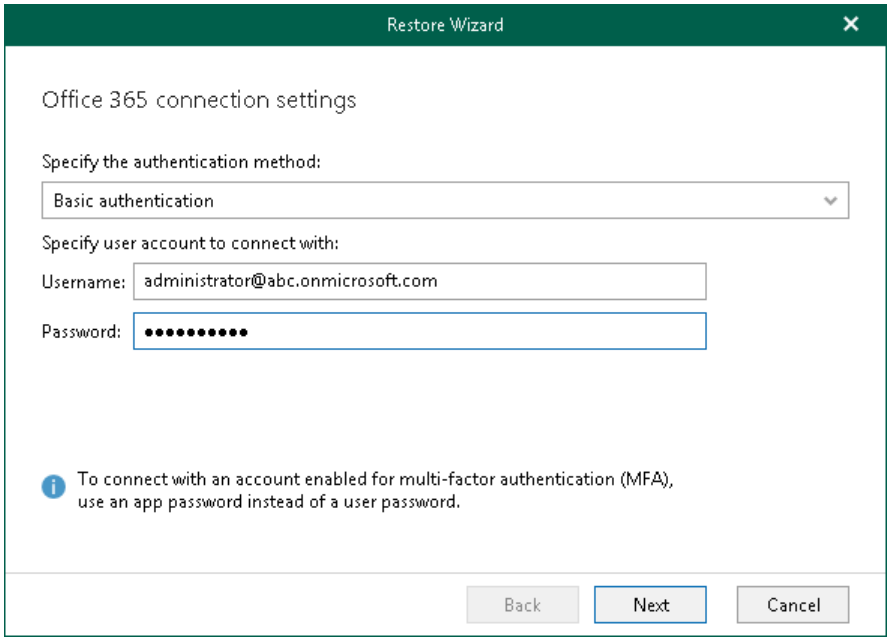

Credentials for On-Premises SharePoint Organizations

To authenticate on an on-premises Microsoft SharePoint server, provide a user name and password.

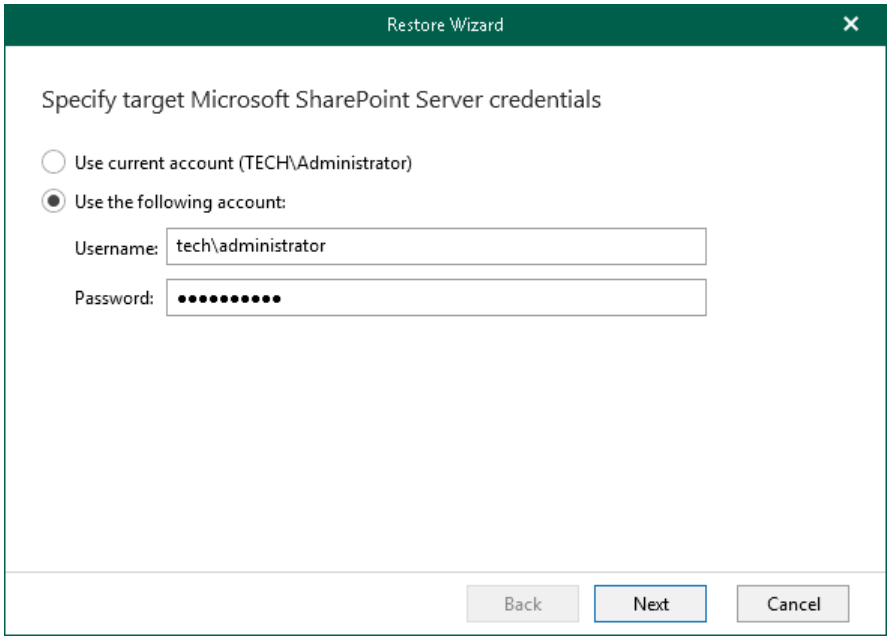

Restoring from Microsoft SQL Database

When restoring from a Microsoft SQL database, provide the following:

- The site URL to which you want to recover your site.
- Authentication credentials.

NOTE

When recovering Microsoft SharePoint sites from a Microsoft SQL database, you will be taken directly to the Specify Restore Options step.

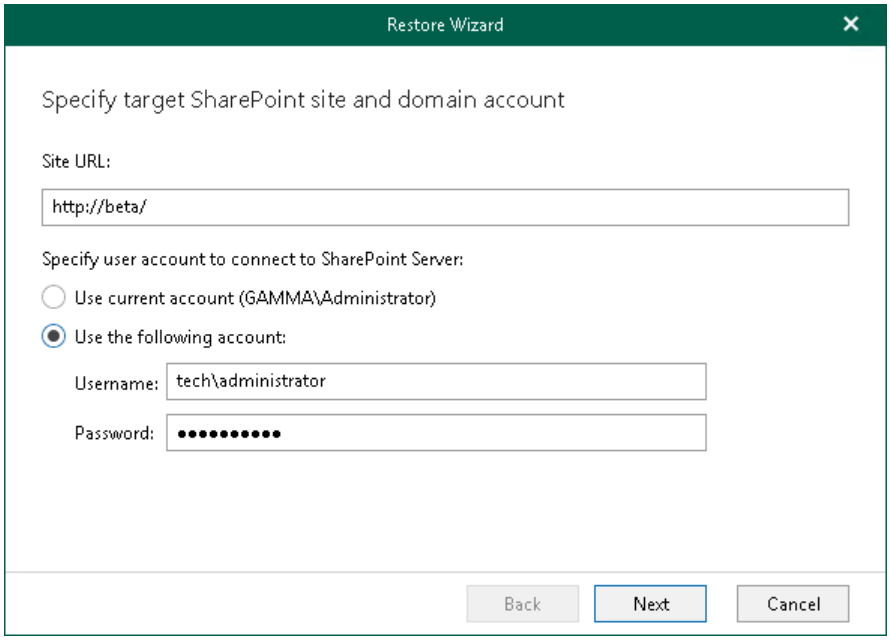

Step 3. Log In to Office 365

This step is available only if you restore data from a backup created in Veeam Backup for Microsoft Office 365 and have selected the *Modern Authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

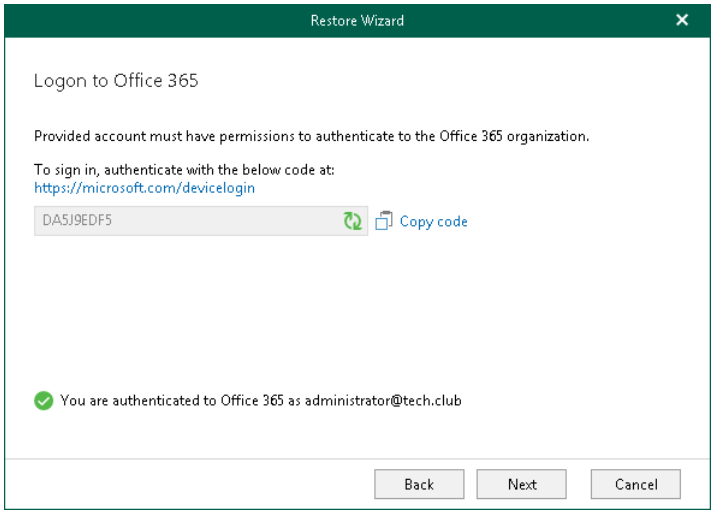

Step 4. Specify Target List

At this step of the wizard, select whether to restore data to the original list or specify a custom one.

You can select the following options:

- **Restore to the original list**. To restore data back to the original location.
- **Restore to the following list**. To restore data to a custom list.

  If the specified list does not exist, it will be created automatically.

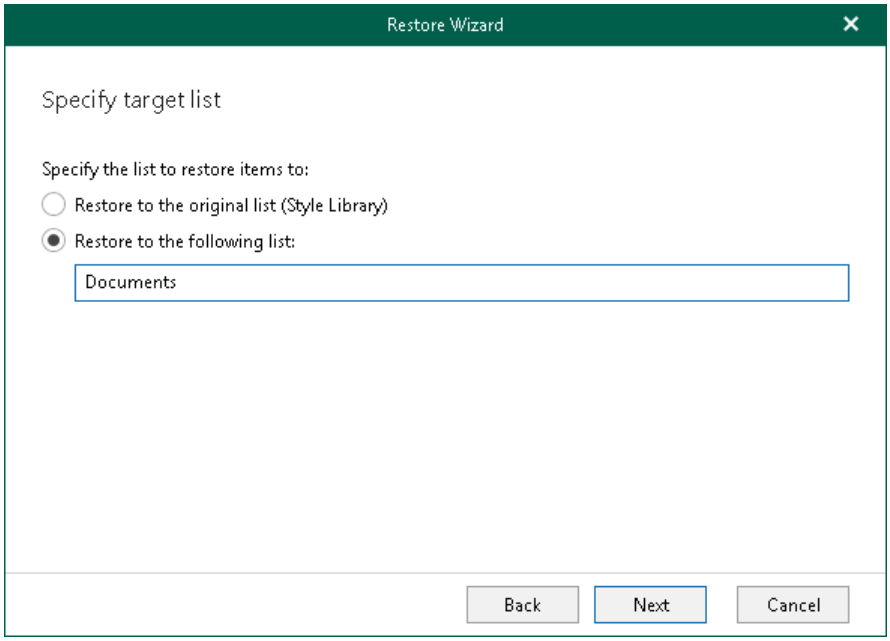

Step 5. Specify Restore Options

At this step of the wizard, specify restore options and click **Restore**.

The restore options are as follows:

- **Changed items**. Allows you to recover data that has been modified in your production environment.
- **Missing items**. Allows you to recover items that have been deleted in your production environment.
- **Restore permissions**. Allows you to recover permissions. If not selected, the permissions for the recovered document library or list will be set as follows:
  - If the library (or list) does not exist on target, it will be created inheriting permissions from the parent object.
  - If the library (or list) already exists on target, permissions will be preserved.
- **Restore list views**. Allows you to recover your list views.
- **History restore options**. Allows you to select a version:
  - **Overwrite**. To overwrite data in the production environment by recovering only the latest version of the document from the backup.
  - **Merge**. To merge an existing version of the document with that of a backup version.

  If not selected, all the versions in the production environment will be replaced with the corresponding data from the backup file.

> **NOTE**
>
> The **Send a notification by e-mail to the users with permissions to the file** check box is available only when restoring data from backups created in Veeam Backup for Microsoft Office 365 for Microsoft Online organizations.

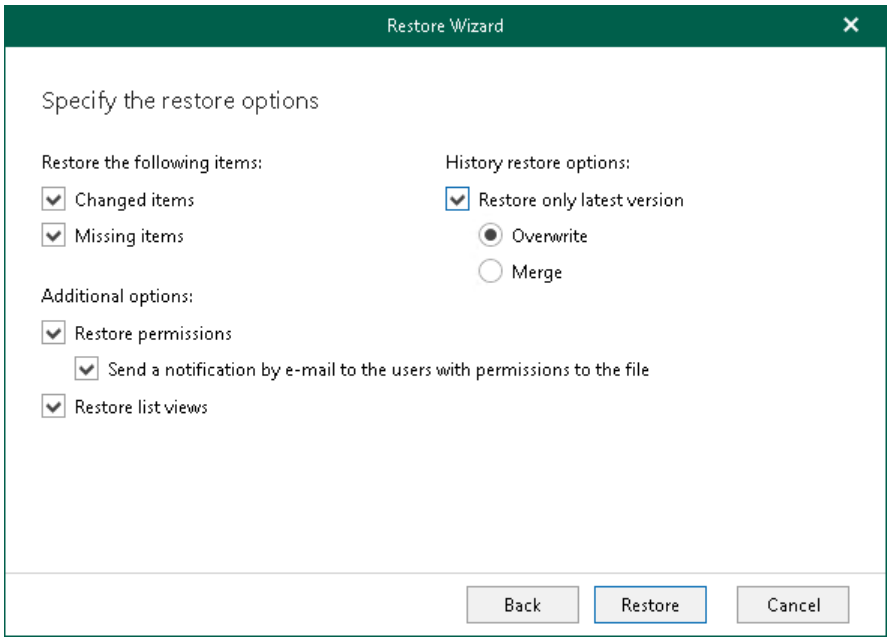

Restoring Microsoft SharePoint Documents and List Items

To restore Microsoft SharePoint documents or list items, do the following:

1. Launch the Restore wizard.
2. Specify credentials.
3. Log in to Office 365 (for data restore from backups created in Veeam Backup for Microsoft Office 365).
4. Specify target location.
5. Specify restore options.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

1. In the navigation pane, select a document or list item.
2. On the **Document** tab, select **Restore Document/Restore Item** or right-click an object and select **Restore Document/Restore Item**.

> **NOTE**
>
> When restoring backups created with Veeam Backup for Microsoft Office 365, document libraries and lists can only be restored back to the source site from which they were backed up.

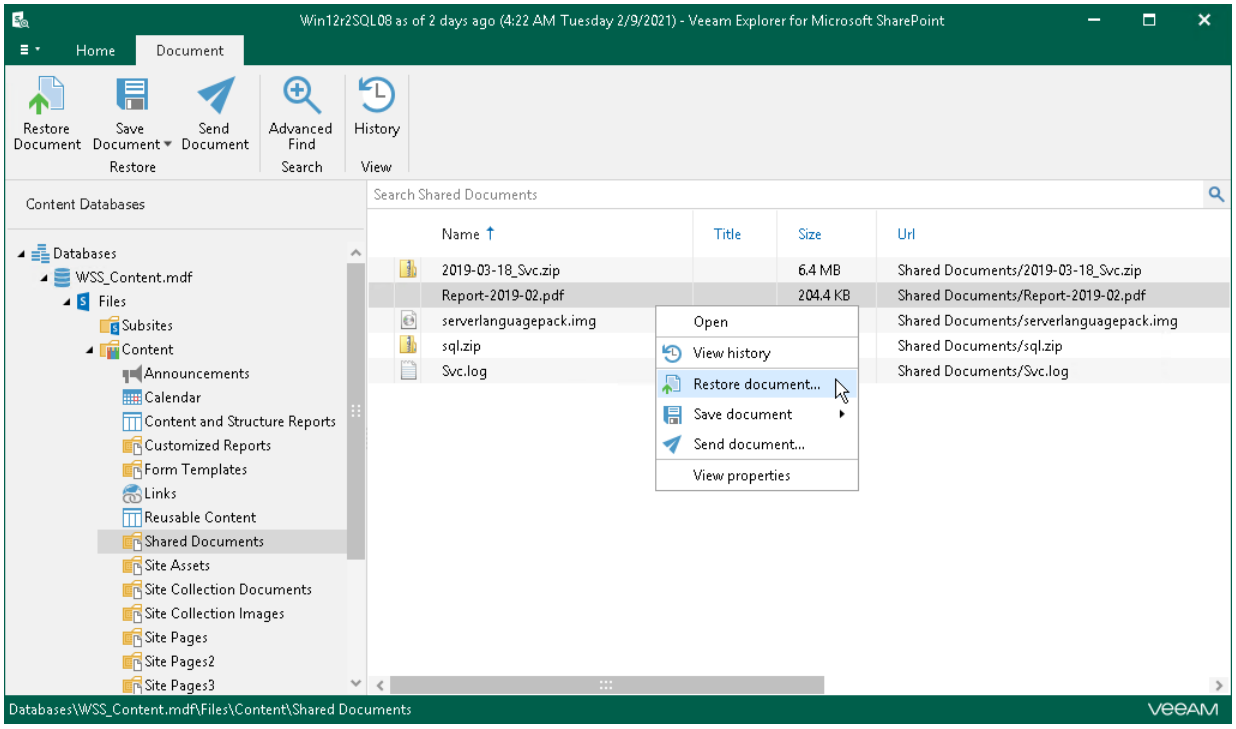

Step 2. Specify Credentials

At this step of the wizard, specify target server credentials. The credentials dialog differs depending on the organization type you are restoring from.

Credentials for Office 365 Organizations with Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified object back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft SharePoint populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

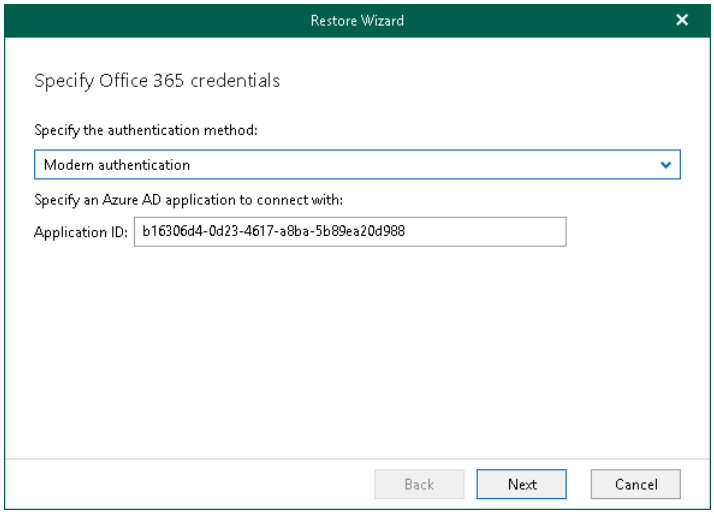

Credentials for Office 365 Organizations with Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the SharePoint organization.

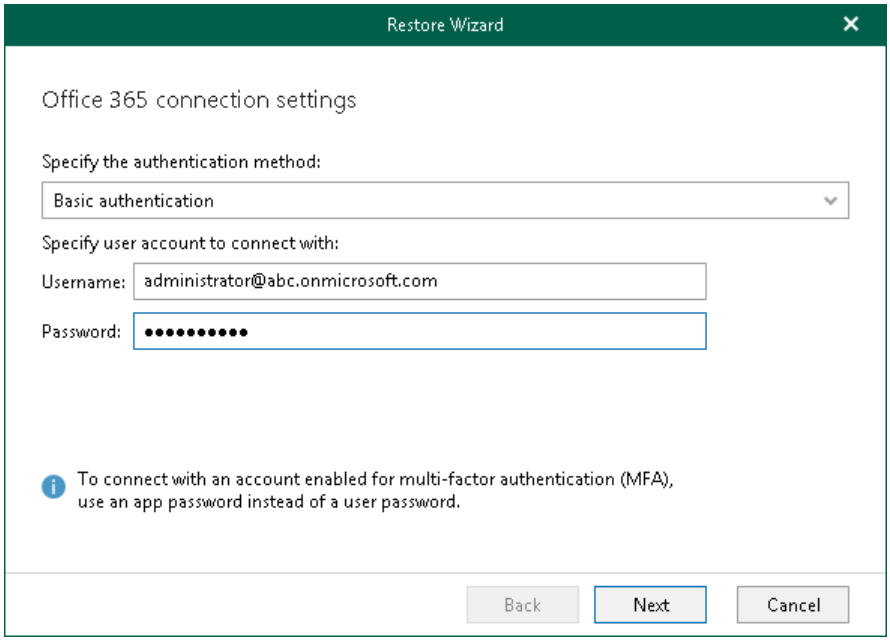

Credentials for On-Premises SharePoint Organizations

To authenticate on an on-premises Microsoft SharePoint server, provide a user name and password.

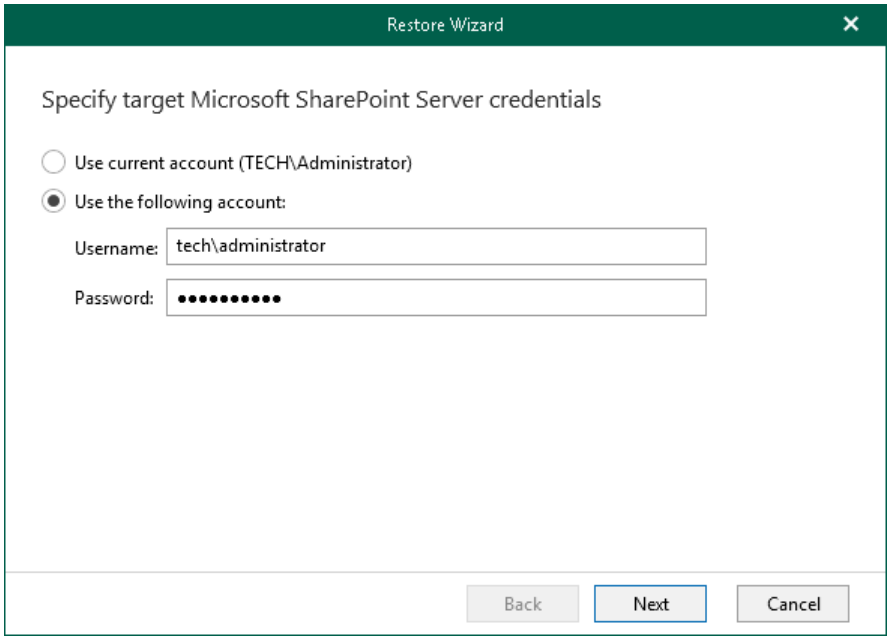

Restoring from Microsoft SQL Database

When restoring from a Microsoft SQL database, provide the following:

- The site URL to which you want to recover your site.
- Authentication credentials.

NOTE

When recovering Microsoft SharePoint sites from a Microsoft SQL database, you will be taken directly to the Specify Restore Options step.

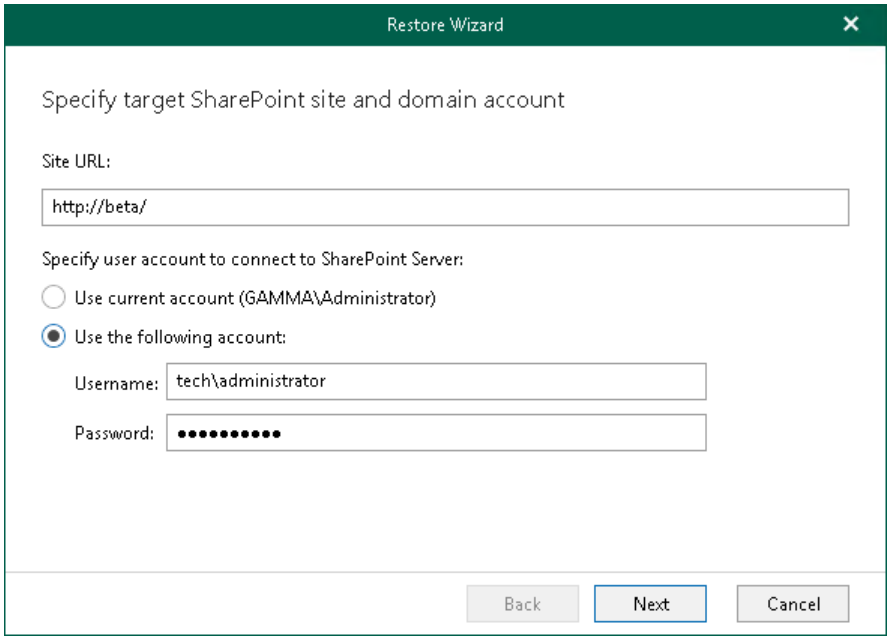

Step 3. Log In to Office 365

This step is available only if you restore data from a backup created in Veeam Backup for Microsoft Office 365 and have selected the *Modern Authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

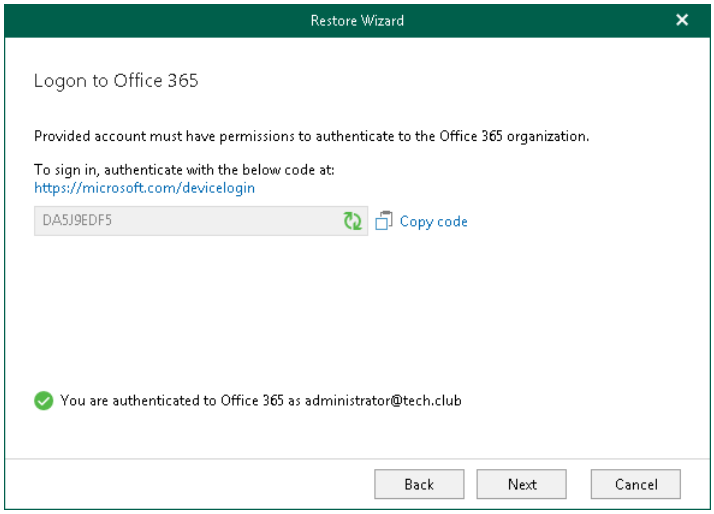

Step 4. Specify Target Location

At this step of the wizard, select whether to restore data to the original list or to a custom one.

You can select the following options:

- **Restore to the original list**. Select this option to restore items to the original list.
- **Restore to the following list**. Select this option to restore items to a custom list.

  Make sure to specify a list that exists in the original SharePoint site.

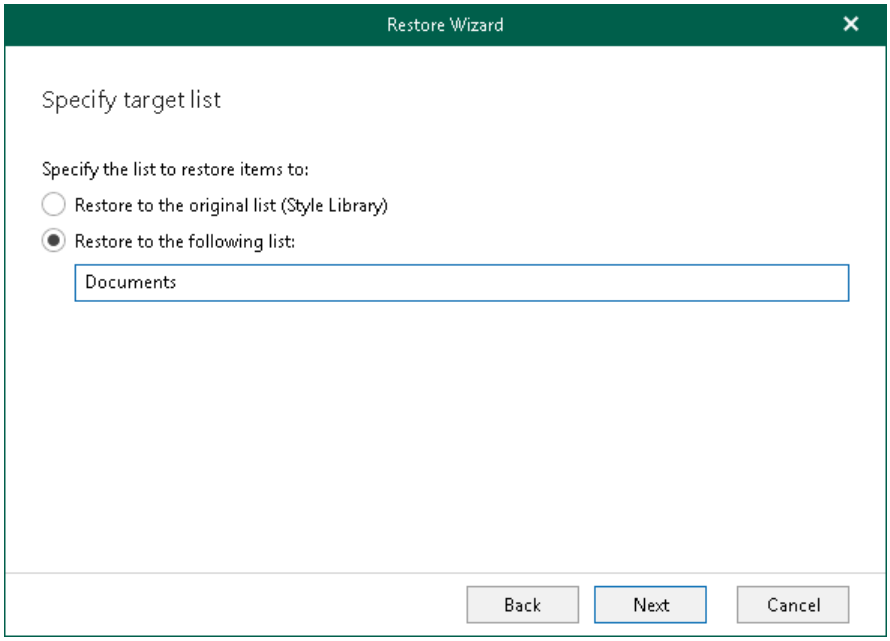

Step 5. Specify Restore Options

At this step of the wizard, specify restore options and click **Restore**.

The restore options are as follows:

- **Changed items**. Allows you to recover data that has been modified in your production environment.
- **Missing items**. Allows you to recover items that have been deleted in your production environment.
- **Restore permissions**. Allows you to recover permissions. If not selected, permissions for the recovered document library or list will be set as follows:
  - If the library (or list) does not exist on target, it will be created inheriting permissions from the parent object.
  - If the library (or list) already exists on target, permissions will be preserved.
- **History restore options**. Allows you to select a version:
  - **Overwrite**. To overwrite data in the production environment by recovering only the latest version of the document from the backup.
  - **Merge**. To merge an existing version of the document with that of a backup version.

  If not selected, all the versions in the production environment will be replaced with the corresponding data from the backup file.

> **NOTE**
>
> The **Send a notification by e-mail to the users with permissions to the file** check box is available only when restoring data from backups created in Veeam Backup for Microsoft Office 365 for Microsoft Online organizations.

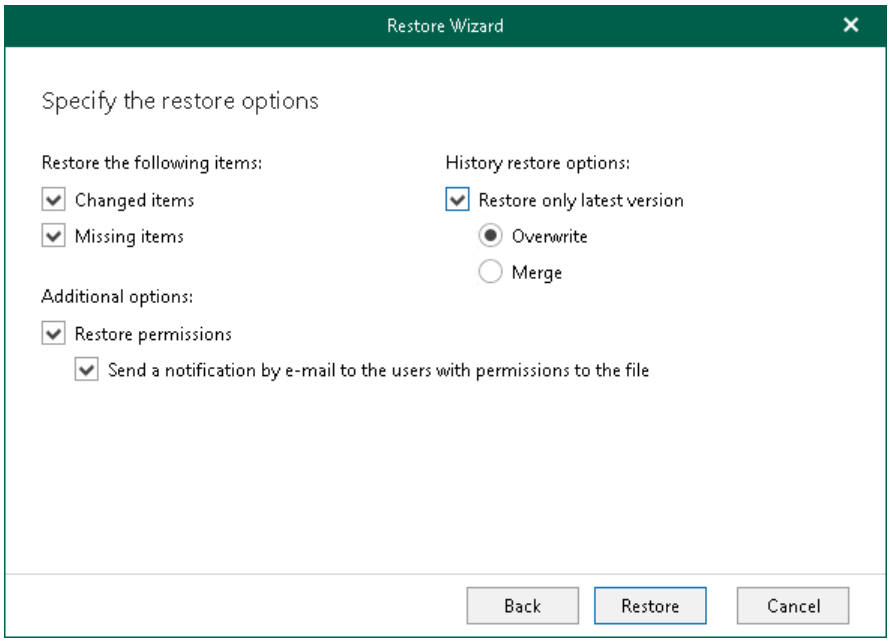

Data Export

This section explains how to export and import Microsoft SharePoint data.

Consider the following:

- Export is unavailable for backups created in Veeam Backup for Microsoft Office 365. The operation is available only for backups created in Veeam Backup & Replication.
- Make sure to read Unsupported SharePoint Lists before exporting Microsoft SharePoint lists.

Exporting Microsoft SharePoint Data

Veeam Explorer for Microsoft SharePoint allows you to export SharePoint document libraries and lists.

The exported content will be saved as a set of XML files and can be imported to another SharePoint database using PowerShell cmdlets, as described in Importing Microsoft SharePoint Data.

To export SharePoint data, do the following:

1. In the navigation pane, select a library or list.
2. On the **Library** tab, select **Export Library/Export List** or right-click an object and select **Export Library/Export List**.
3. Specify the output directory and click **OK**.

> **NOTE**
>
> Export of items is not supported in the current version.

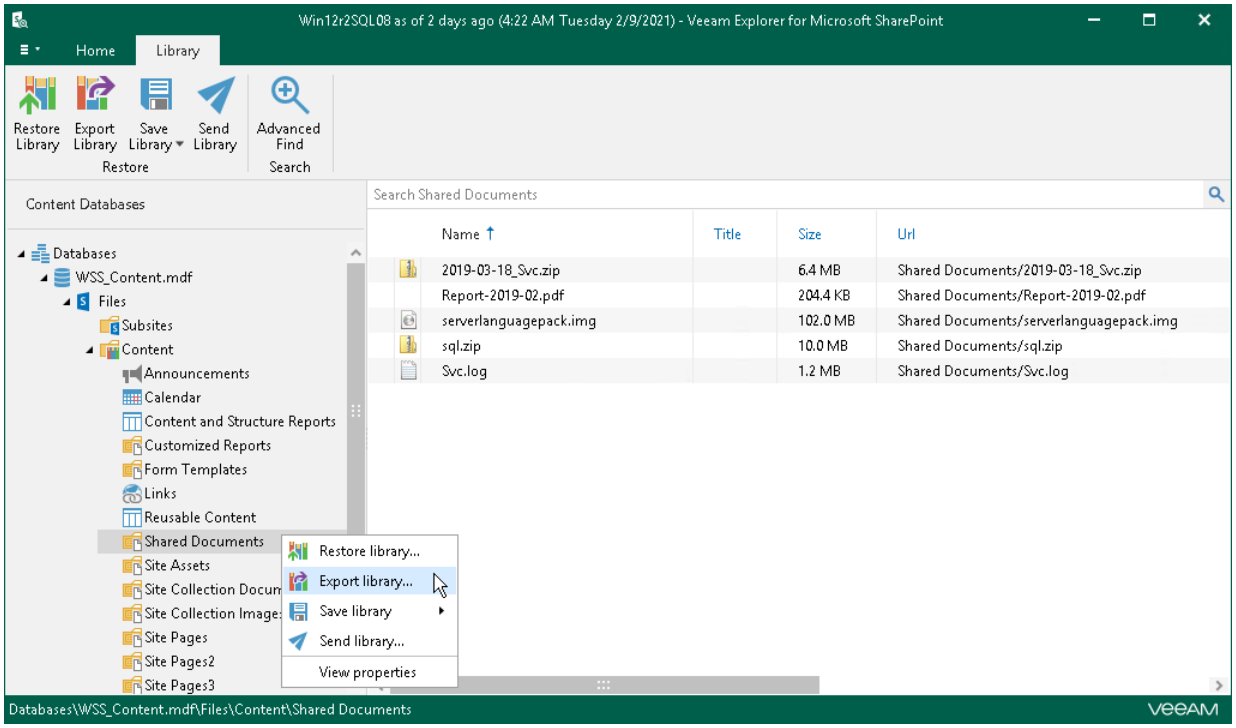

Importing Microsoft SharePoint Data

To import document library or list, use either of the following PowerShell cmdlets:

- For PowerShell snap-in, use the following command.

```
Add-PsSnapin Microsoft.SharePoint.PowerShell
Import-SPWeb -Identity "http://<web_server_name>/sites/<destination_site>"
-Path "C:\<export_folder>" -NoFileCompression –IncludeUserSecurity
```

- For SharePoint Management Shell, use the following command.

```
Import-SPWeb -Identity "http://<web_server_name>/sites/<destination_site>"
-Path "C:\<export_folder>" -NoFileCompression –IncludeUserSecurity
```

where:

- `<web_server_name>` — destination web server.
- `<destination_site>` — destination web site.
- `<export_folder>` — source folder containing exported library/list content.

To get extended help on the `Import-SPWeb` command, use the following command.

```
Get-Help Import-SPWeb -full
```

Saving Microsoft SharePoint Documents and Libraries

Veeam Explorer for Microsoft SharePoint allows you to save your libraries and library documents to a specified location.

To save a Microsoft SharePoint library or library documents, do the following:

1. In the navigation pane, select a library.
2. On the **Library** tab, select **Save Library** > **Save as files** or **Save Library** > **Save as ZIP** or right-click an object and select **Save Library** > **Save as files** or **Save Library** > **Save as ZIP**.

   To save documents from a library, select a document in the preview pane and click **Save Document** > **Save as files** or **Save Document** > **Save as ZIP**.

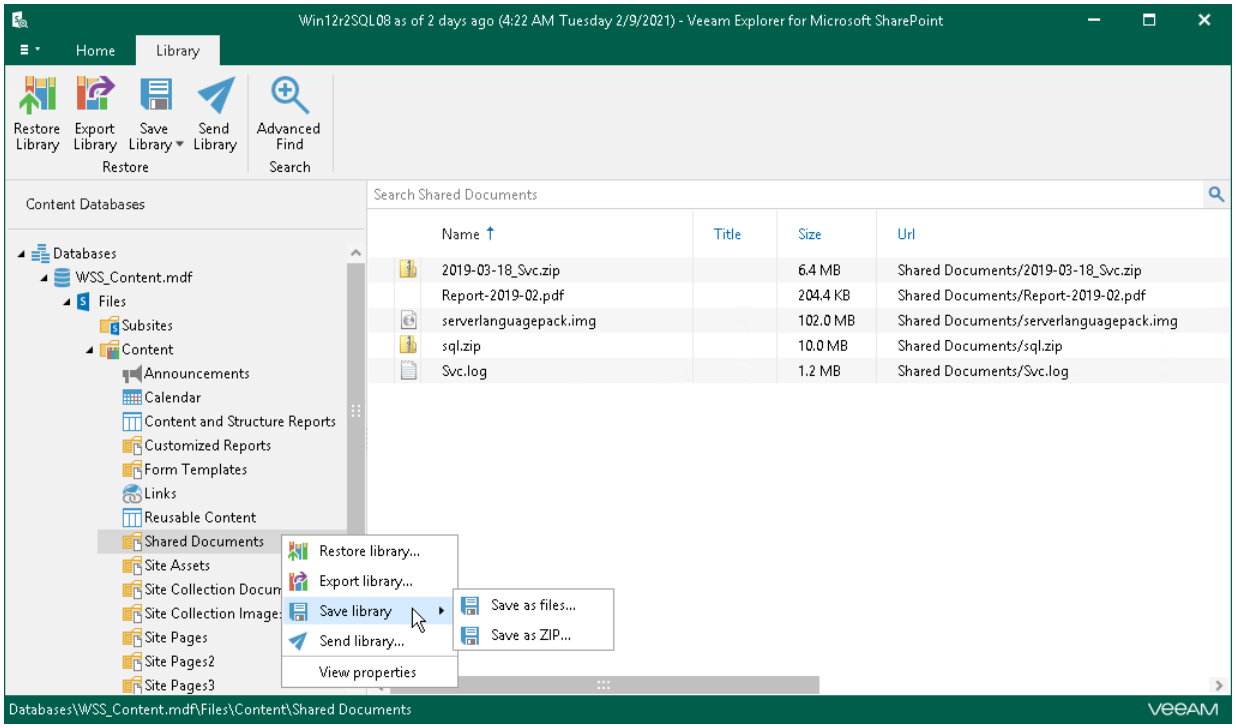

Sending Microsoft SharePoint Documents and Libraries

Veeam Explorer for Microsoft SharePoint allows you to send libraries or library documents to the specified recipients via email.

> **TIP**
>
> Before sending documents, make sure to configure SMTP settings, as described in Configuring SMTP Settings. The amount of data you can send at a time depends on your SMTP server configuration.

To send Microsoft SharePoint library or library documents, do the following:

1. In the navigation pane, select a library to send.
2. On the **Library** tab, select **Send Library** or right-click an object and select **Send Library**.

   To send documents from a library, select a document in the preview pane and click **Send Document**.

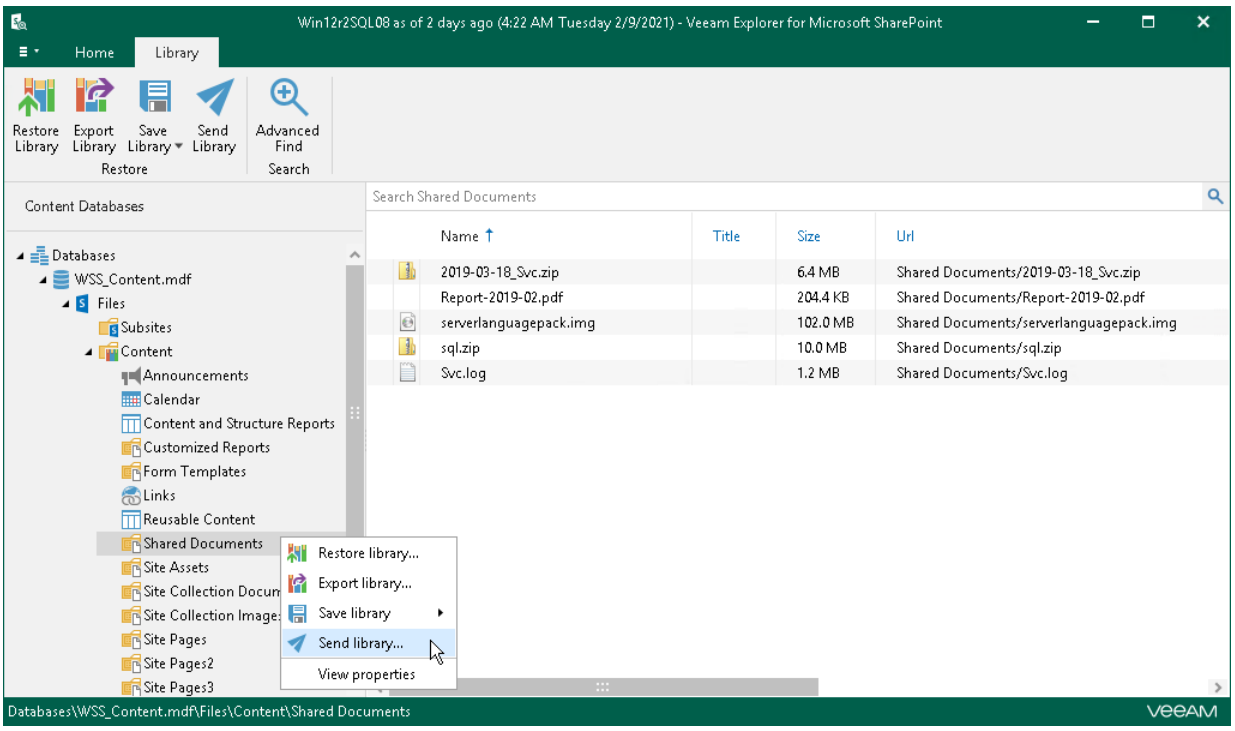

3. Provide a recipient address.

   The **From** field is filled automatically based on the address you have provided when configuring SMTP settings. To edit the message body, click **More details**.

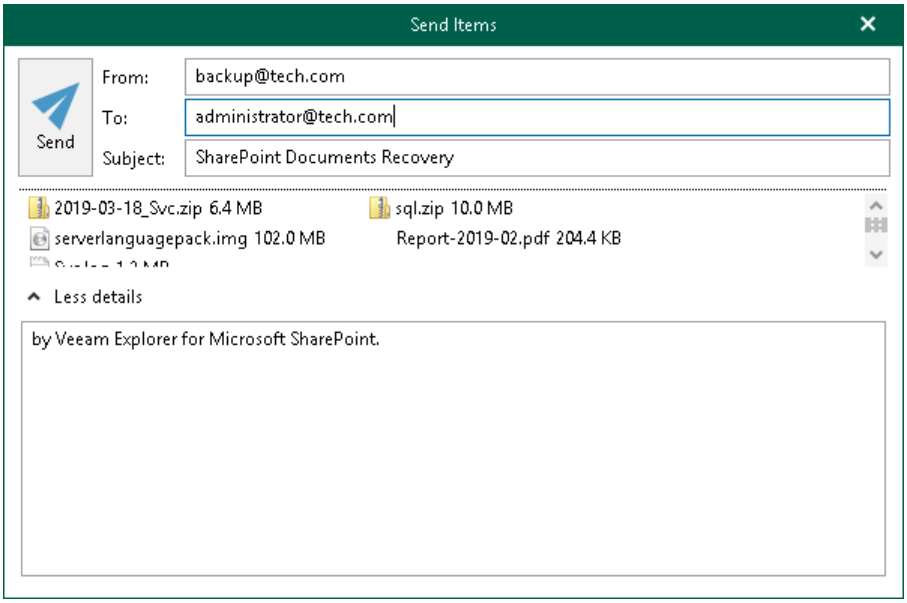

Appendix A: Unsupported SharePoint Lists

The following table contains unsupported Microsoft SharePoint lists.

ID	List Name	Version
101	SearchConfigListTemplate	2013/2016/2019
110	datasrcs	2010/2013/2016/2019
111	webtemp	2010/2013/2016/2019
112	users	2010/2013/2016/2019
113	wplib	2010/2013/2016/2019
114	listtemp	2010/2013/2016/2019
116	mplib	2010
117	nocodewf	2010/2013/2016/2019
118	workflowProcess	2010/2013/2016/2019
121	solutions	2010/2013/2016/2019
122	nocodepublicwf	2010/2013/2016/2019
123	themes	2010/2013/2016/2019
124	designs	2013/2016/2019
125	appdata	2013/2016/2019
126	appfiles	2019
160	accessrequests	2013/2016/2019
175	MaintenanceLogs	2013/2016/2019

ID	List Name	Version
200	meetings	2010/2013/2016/2019
212	wkspglib	2010/2013/2016/2019
331	ComponentManifests	2019
334	ClientSideAssets	2019
335	SiteCollectionAppCatalogs	2019
336	WebApiPermissionRequests	2019
401	fcgroups	2010/2013/2016/2019
401	MonitoredAppsList	2013/2016/2019
420	timecard	2013/2016/2019
425	whatsnew	2010/2013/2016/2019
530	ContentFollowingList	2013/2016/2019
544	MicroBlogList	2013/2016/2019
550	SocialDataStoreList	2013/2016/2019
563	AnnouncementTilesList	2016/2019
600	extlist	2010/2013/2016/2019
880	members	2013/2016/2019
925	abusereports	2013/2016/2019
1221	healthreports	2010/2013/2016/2019
1305	EDiscoverySources	2016/2019

ID	List Name	Version
1306	EDiscoverySourceInstances	2013/2016/2019
1309	DlpEDiscoveryQueries	2016/2019
1310	PreservationListTemplate	2013/2016/2019
2000	PNSubscribers	2013/2016/2019
2001	ExternalSubscriptions	2013/2016/2019
3100	AccSrvAddApp	2013/2016/2019
3300	SharingLinks	2016/2019
3400	HashtagStore	2019
3401	BlockedHashtags	2019
3412	HubSites	2019
3415	WebTemplateExtensions	2019
3500	IRCollection	2019
3501	IRReference	2019
3502	IRReferenceCollection	2019
4501	wfsvc	2013/2016/2019
10000	Academic Library	N/A
10001	eduentity"	2013/2016/2019
10007	EduWorkItem	2013/2016/2019
10008	EduQuiz	2013/2016/2019

ID	List Name	Version
10051	eduentity	2013/2016/2019
10060	eduusersetting	2013/2016/2019
10061	eduexternalsyncsetting	2013/2016/2019
10099	AcquisitionHistoryList	2013/2016/2019
10101	edudocument	2013/2016/2019
10102	IWConvertedForms	2010/2013/2016/2019
10401	eduannouncement	2013/2016/2019
10631	educalendar	2013/2016/2019

Veeam Explorer for Microsoft OneDrive for Business

Veeam Explorer for Microsoft OneDrive for Business allows you to restore Microsoft OneDrive data from backups created in Veeam Backup for Microsoft Office 365.

Planning and Preparation

Veeam Explorer for Microsoft OneDrive for Business comes as part of the Veeam Explorer for Microsoft SharePoint installation package and has the same system requirements.

Consider that when Veeam Explorer for Microsoft SharePoint is installed on a server on which both Veeam Backup & Replication and Veeam Backup for Microsoft Office 365 are installed, the notification settings will be inherited from the Veeam Backup & Replication Global Notification settings.

Launching Application and Exploring Backups

To open Veeam Explorer for Microsoft OneDrive for Business and load backups, you can use one of the following methods:

- Use the **Explore** option to load backups created in Veeam Backup for Microsoft Office 365.

  For more information, see the Data Restore section of the Veeam Backup for Microsoft Office 365 User Guide.

  When loading backups located in object storage repositories, Veeam Explorer for Microsoft Teams will notify you about the location of these backups. You can select the **Don't show this message again** check box to unsubscribe from such notifications. To disable these notifications, clear the **Show pop-up notifications on startup** check box, as described in Configuring Advanced Settings.

- Go to **Start**, click **Veeam Explorer for Microsoft OneDrive for Business** and manually open Microsoft OneDrive databases, as described in the Standalone Databases Management section.

If you open Veeam Explorer for Microsoft OneDrive for Business from the **Start** menu on a machine that runs the Veeam Backup & Replication console and Veeam Explorers only, specify the following:

- The domain name or IP address of a server to which you want to connect
- The port number
- User credentials

  The account must be a member of the *Local Administrator* group on the target server. To use the account under which Veeam Explorer for Microsoft OneDrive for Business is running, select **Use Windows session authentication**.

To save the connection shortcut to the desktop, click the **Save shortcut** link in the bottom-left corner.

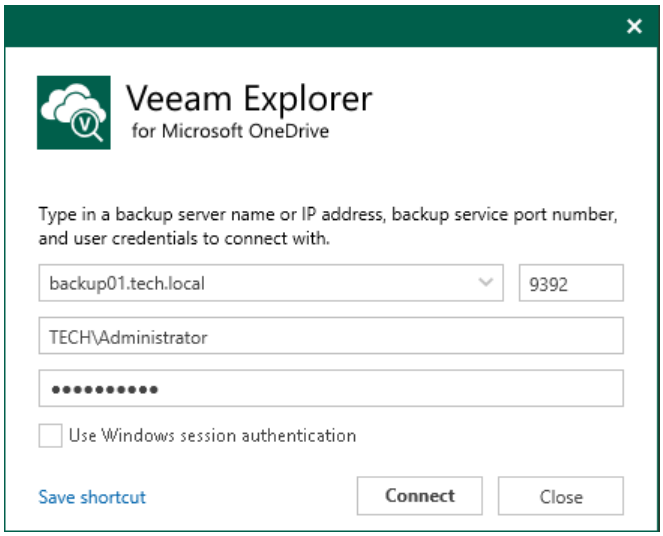

Understanding User Interface

Veeam Explorer for Microsoft OneDrive for Business provides you with the convenient user interface that allows you to perform required operations in a user-friendly manner.

Main Menu

The main menu comprises the following features:

- **General Options**. Allows you to configure program options. See Performing Initial Configuration Settings.
- **Help**.
  - **Online help**. Opens the online web help page.
  - **About**. Shows current product information.
- **Exit**. Closes the program.

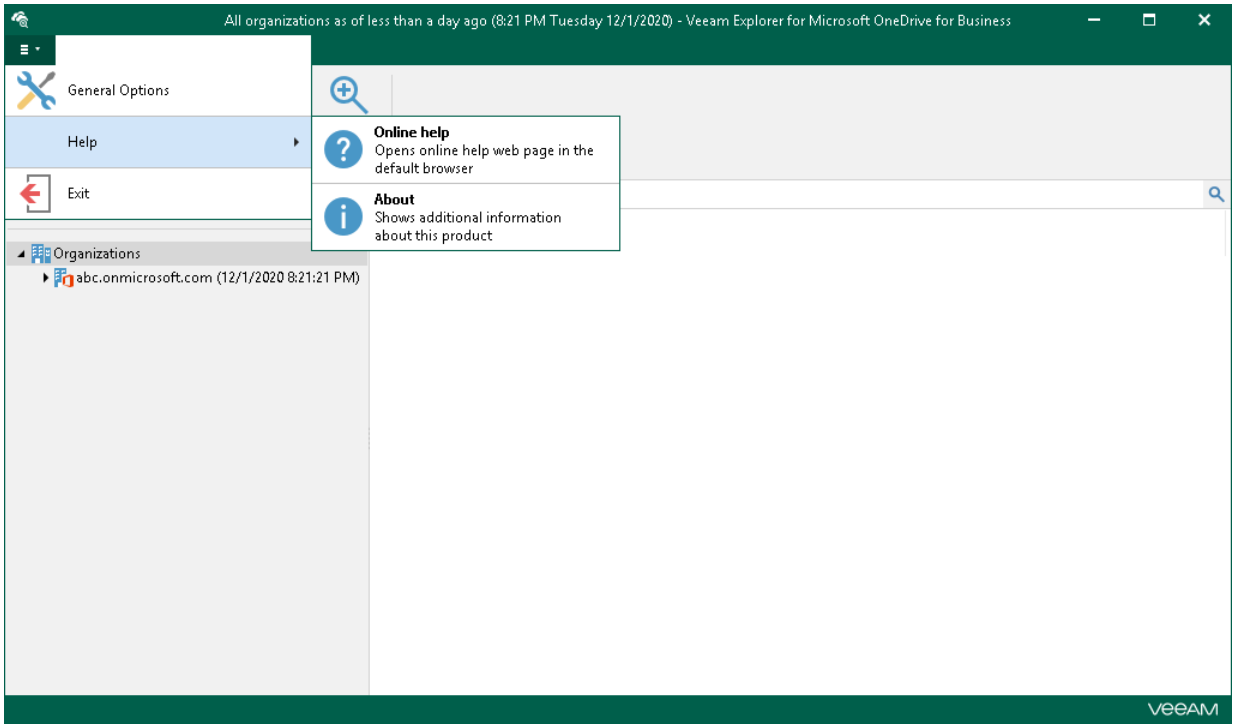

Main Application Window

The main application window might be divided into three categories:

1. The ribbon menu, which contains general program commands organized into logical groups.
2. The navigation area, which allows you to browse through the hierarchy of your backup files.
3. The preview pane, which shows you the details about objects you have selected in the navigation area.

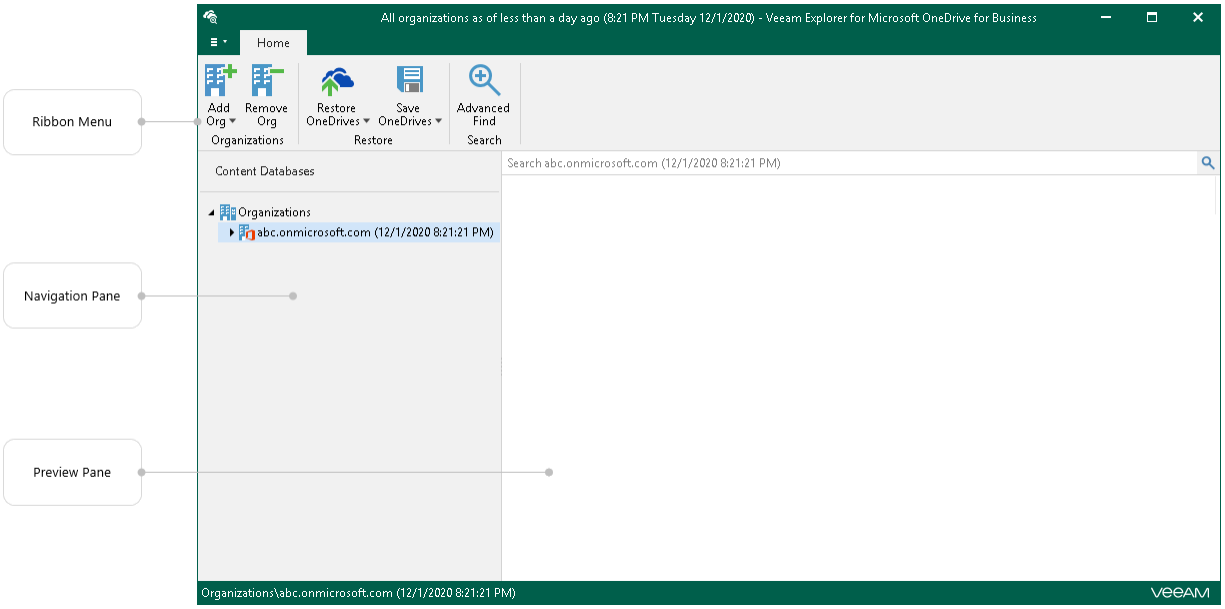

Browsing, Searching and Viewing Items

Continue with this section to learn more about:

- Browsing backup content
- Viewing objects properties and open files
- Searching for objects in a backup file
- Using the advance search capabilities

Browsing

To view the contents of a backup file, you use the navigation pane which shows you the database structure containing your OneDrive documents.

After you select an object in the navigation pane, you can see its content in the preview pane.

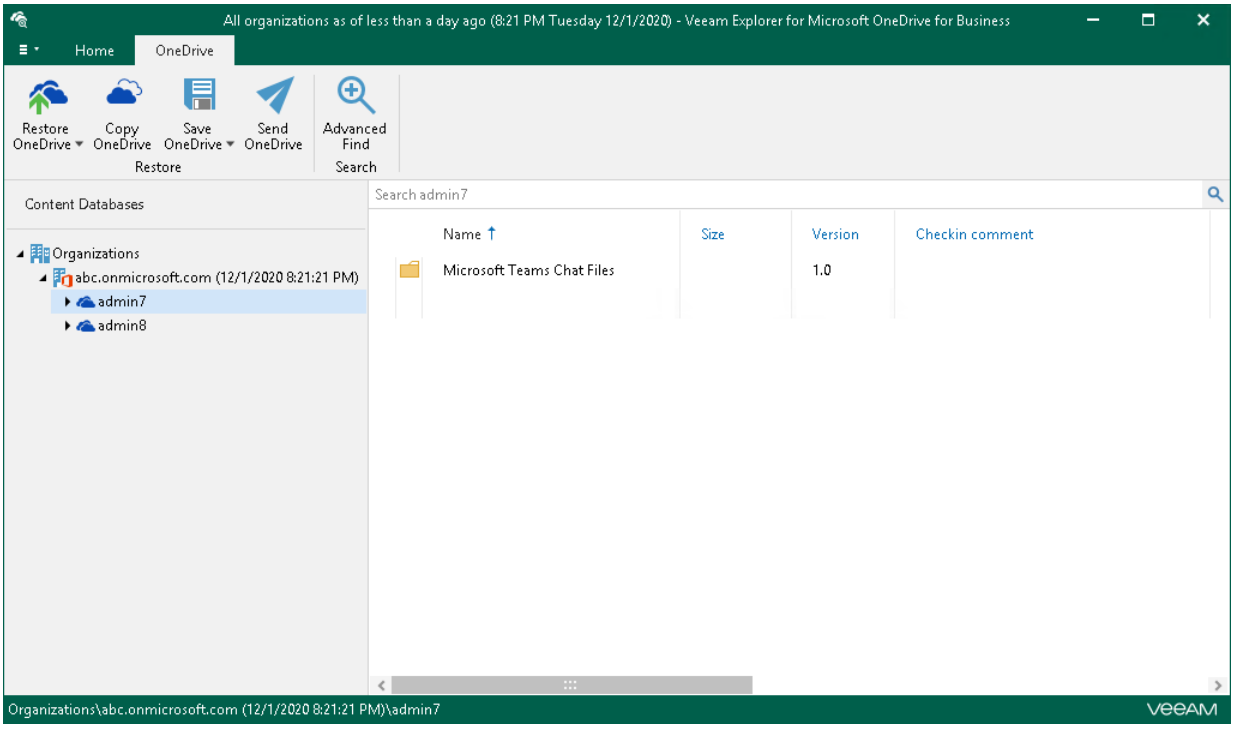

Viewing Properties and Opening Files

To view object properties, right-click an object in the preview pane and select **View Properties**.

To open a document using an associated application, right-click a document in the preview pane and select **Open**.

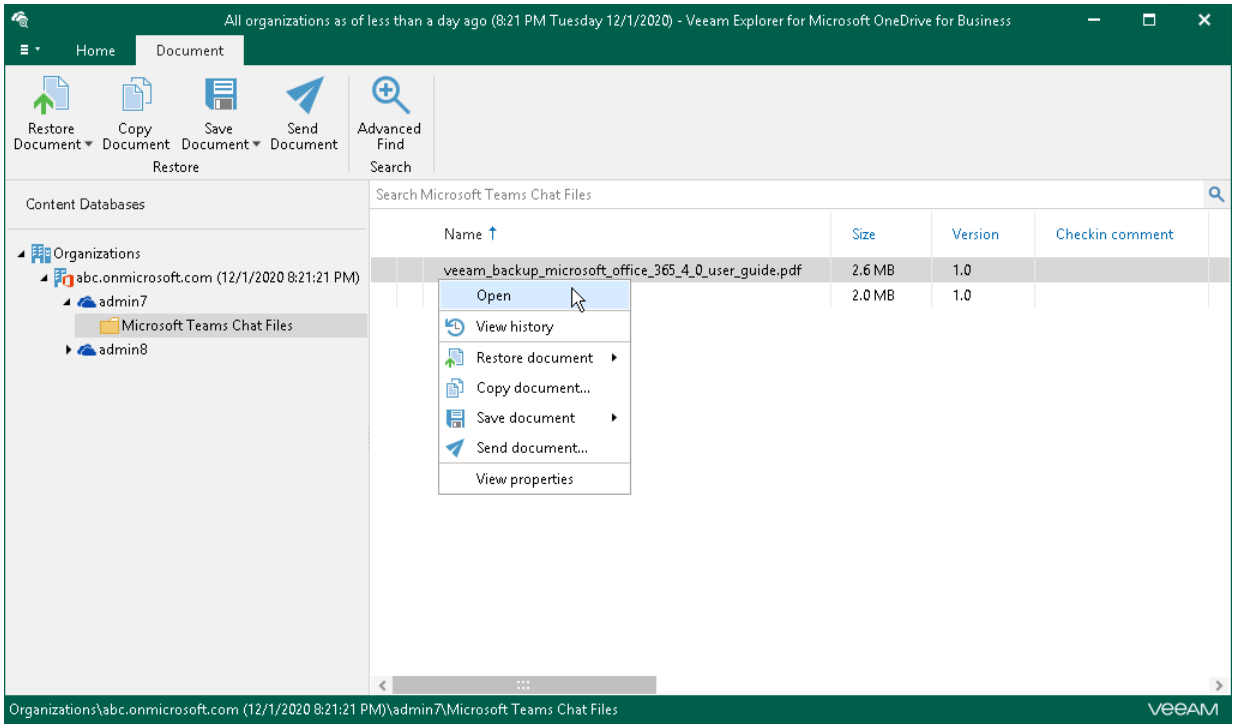

Searching

The search mechanism allows you to find items matching specified search criteria.

To search for required items, do the following:

1. In the navigation pane. select an object in which you want to find your data.
2. Enter a search query to the search field at the top of the preview pane.

> **NOTE**
>
> To find the exact phrase, use double quotes. For example, *"Attachments"*.

You can narrow your search results by specifying various search criteria using the *criteria:value* format. You can also use logical upper-cased operators such as *AND*, *OR* and *NOT* along with wildcard characters such as *** and *?*.

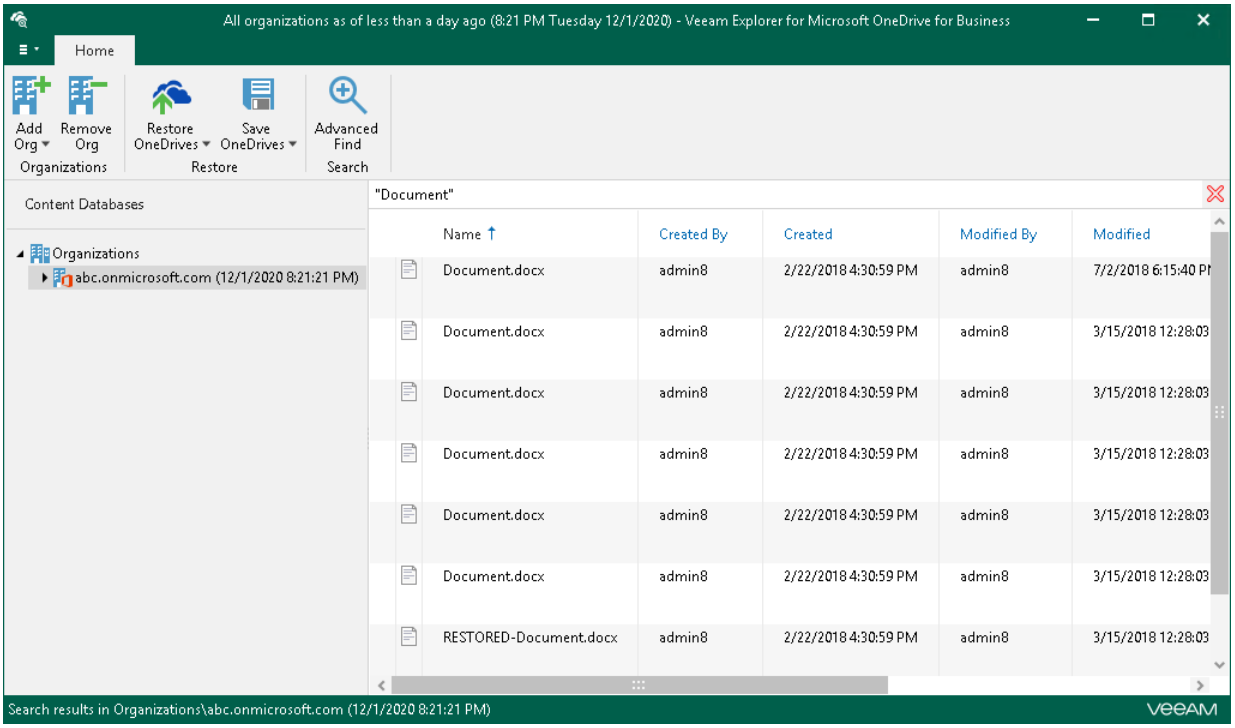

Using Advanced Find Capabilities

The **Advanced Find** mechanism allows you to define your search criteria more precisely.

For example, to find an object that starts with the word *Attachme*, do the following:

1. In the preview pane, select a content node and click **Advanced Find**.
2. In the **Define search criteria** section, select **Category** > **Document fields**.
3. In the **Field** list, select **File Name**.
4. In the **Condition** list, select **Starts With**.
5. In the **Value** field, specify a file name.
6. Click **Start**.

To remove a filter, click the cross mark next to it. To remove all configured filters, click **Reset**.

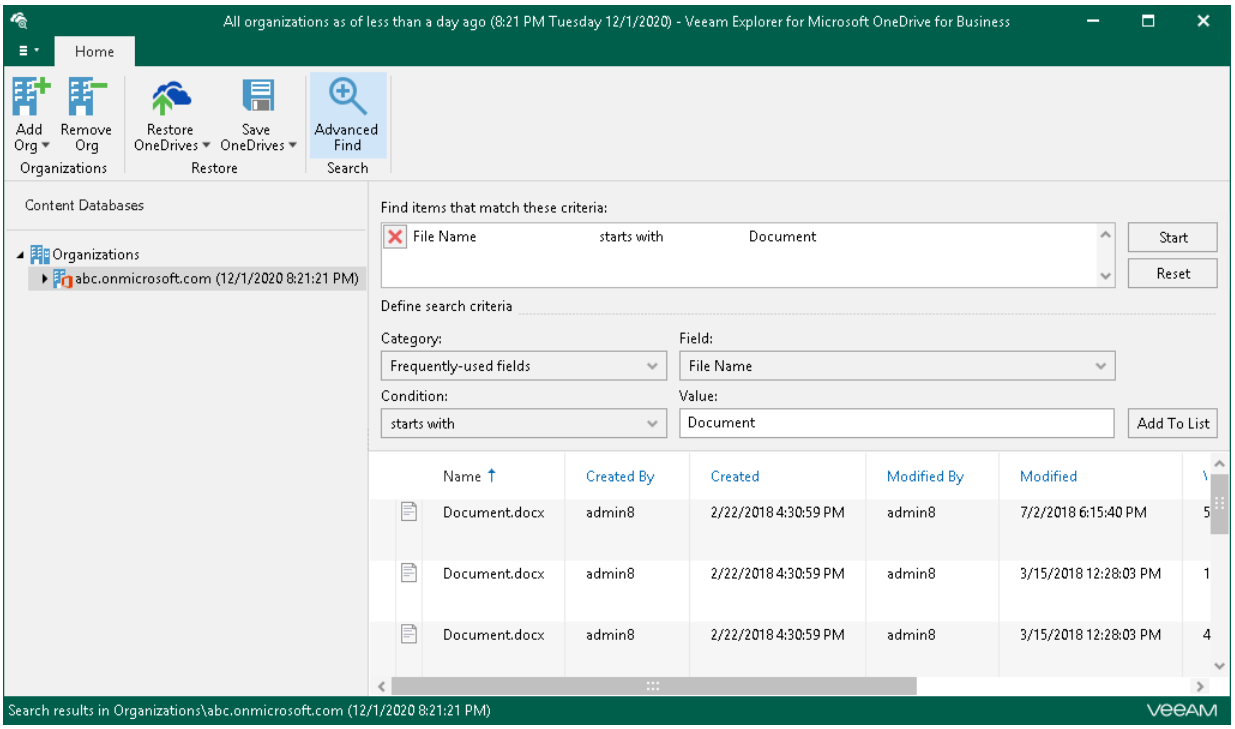

General Application Settings

Continue with this section to learn more about configuring required application settings and components.

Configuring SMTP Settings

To send Microsoft OneDrive items as attachments, you must configure SMTP server settings.

To configure SMTP settings, do the following:

1. Go to the main menu and click **General Options**.
2. On the **SMTP Settings** tab, select the **Configure SMTP settings** check box and specify the following:
   - DNS name or IP address of the mail server.
   - SMTP communication port.
   - The sender email address. This address will appear in the **From** field when sending OneDrive items. See Sending Microsoft OneDrive Documents.
   - Select **Use authentication** check box if your SMTP server requires SMTP authentication for outgoing mail, and provide valid credentials.
   - Select **Enable SSL security** check box to enable SSL data encryption.
3. Click **Send** to send a test email message.
4. Click **Apply**.

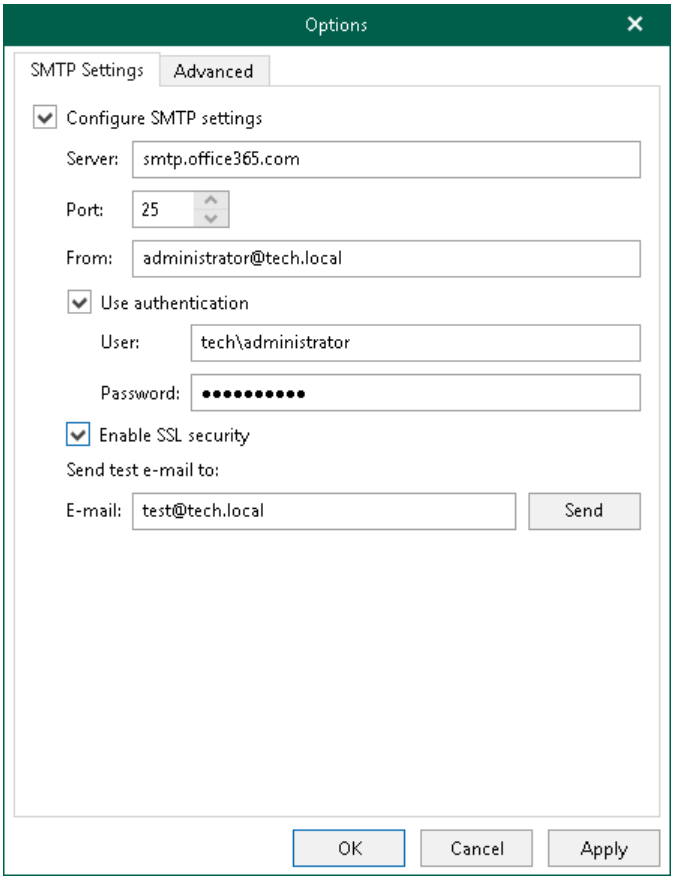

Configuring Advanced Settings

This section explains how to configure advanced settings of the application.

To configure advanced settings, do the following:

1. In the main menu, click **General Options**.
2. Open the **Advanced** tab.
3. If you want to enable extended logging mode to collect logs that contain more details on specific operations, select the **Enable extended logging** check box.

   After enabling extended logging mode, you can go back to the application and perform actions for which you want to collect additional information. Then you can collect logs.
4. By default, Veeam Explorer for Microsoft OneDrive for Business displays notifications when loading backups from object storage repositories. If you want to disable these notifications, clear the **Show pop-up notifications on startup** check box.

   Applicable to backups created in Veeam Backup for Microsoft Office 365.

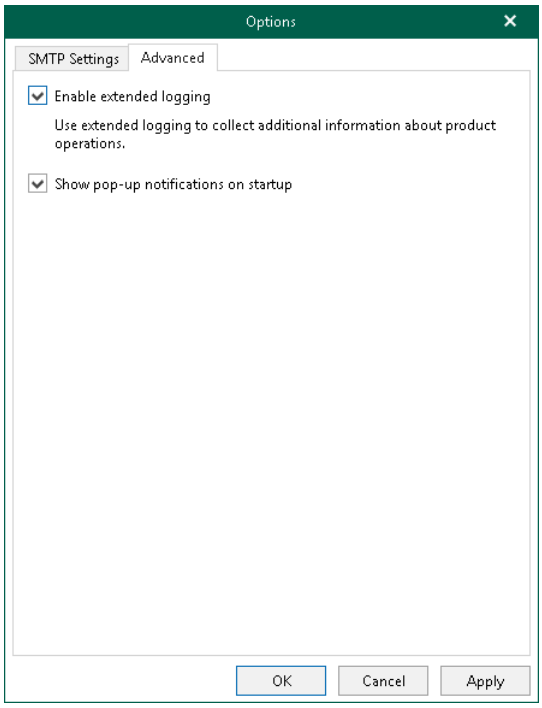

Standalone Databases Management

Continue with this section to learn more about adding and removing OneDrive for Business databases.

Adding Veeam Backup for Microsoft Office 365 Databases

To manually add databases that store Microsoft Office 365 organization data, do the following:

1. Click **Add Org** > **Veeam Backup for Microsoft Office 365 Databases** on the ribbon menu or use the corresponding context menu command.
2. Specify the database file location and log directory.
3. Click **Open**.

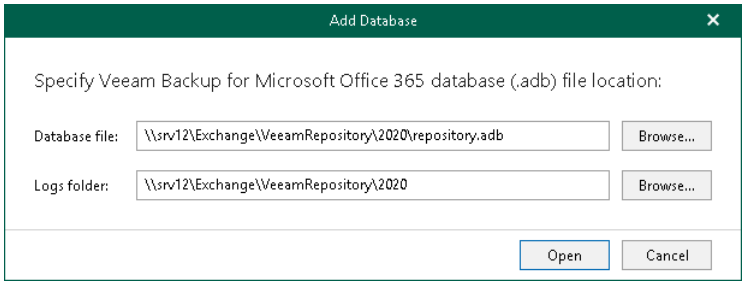

Adding Veeam Backup for Microsoft Office 365 Server

You can use the built-in Veeam Explorer abilities to connect to another Veeam Backup for Microsoft Office 365 server and add its databases to the Veeam Explorer for Microsoft OneDrive for Business scope.

To connect to another Veeam Backup for Microsoft Office 365 server, do the following:

1. Click **Add Org** > **Veeam Backup for Microsoft Office 365 Server** on the toolbar or use the corresponding context menu command.
2. Specify connection settings under which to connect to the Veeam Backup for Microsoft Office 365 server and click **Connect**.

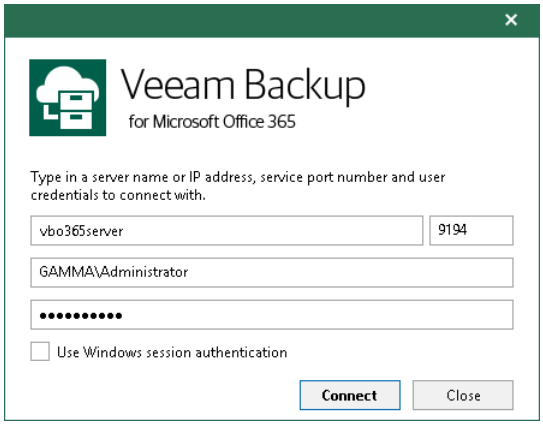

Adding Veeam Backup for Microsoft Office 365 Service Provider

You can use Veeam Explorer for Microsoft OneDrive for Business to connect to a service provider server to add its backups to the application scope and perform a restore directly from such backups.

Consider the following:

- Both Veeam Explorer for Microsoft OneDrive for Business and Veeam Backup & Replication must be installed on the same server.
- At least one service provider must be added to Veeam Backup & Replication, as described in the Connecting to Service Providers section of the Veeam Cloud Connect Guide.

To connect to a Veeam Backup for Microsoft Office 365 service provider, do the following:

1. From the **Start** menu, launch Veeam Explorer for Microsoft OneDrive for Business.
2. On the ribbon, click **Add Org** > **Veeam Backup for Microsoft Office 365 service provider**.

   Alternatively, in the preview pane, right-click **Organizations** and select **Veeam Backup for Microsoft Office 365 service provider**.
3. In the **Select service provider and authentication method** section, select from the drop-down lists a service provider to which you want to connect and an authentication method: *Modern authentication* or *Basic authentication*.
4. [For *Modern authentication*] Select a region to which an Office 365 organization of the selected service provider belongs.
5. [For *Modern authentication*] Copy the device code and use it to sign in to Microsoft Azure Active Directory.

   For *Basic authentication*, you must enter a user name and password to connect to the service provider server.
6. If you want to select a point-in-time state as of which you want to load Office 365 backups, click **Specify point in time** and select the necessary option:
   - **Use the latest available state**. To load data as of the latest backup state.
   - **Use the following point in time**. To load data as of the selected point in time.

     To select a point in time, use the calendar control.

7. Select the **Show items that have been deleted by user** check box to see all the items that were removed.
8. Select the **Show all versions of items that have been modified by user** check box if you want to load modified versions.
9. Click **Connect**.

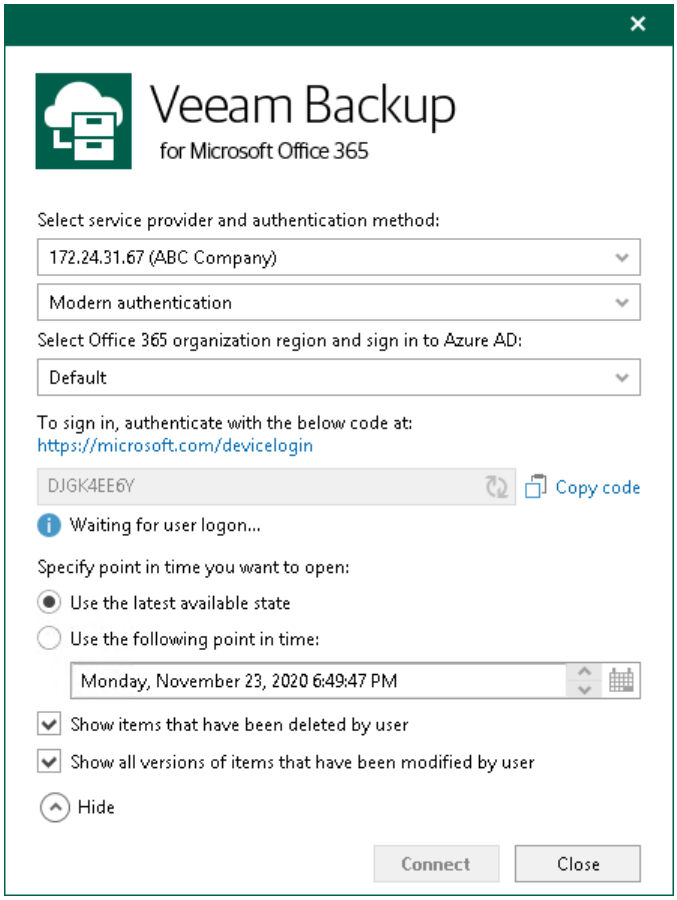

Removing Standalone Databases

Veeam Explorer for Microsoft OneDrive for Business allows you to remove an organization from the application scope when you no longer need it.

To remove an organization from the application scope, right-click an organization in the navigation pane and select **Remove organization**.

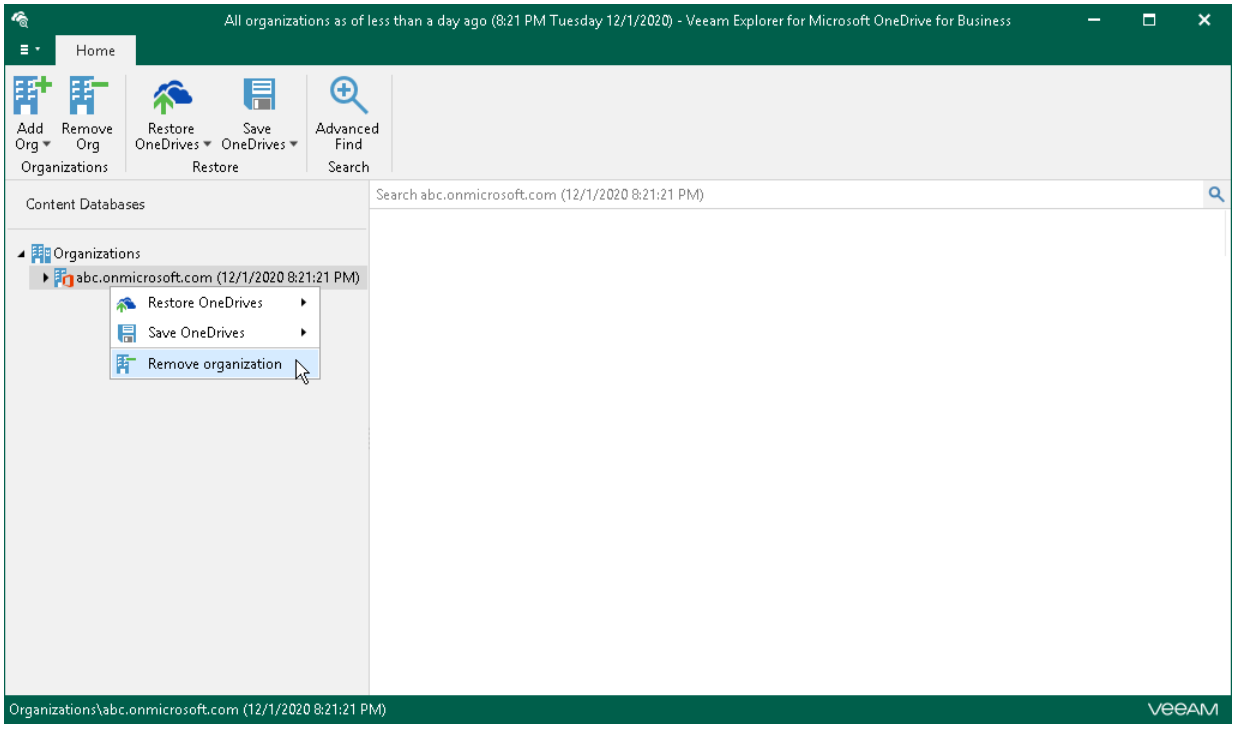

Restoring Microsoft OneDrive Data

Veeam Explorer for Microsoft OneDrive for Business allows you to restore entire OneDrive or certain documents and/or folders of the selected OneDrive.

Restoring Single OneDrive

To restore data of a single OneDrive, do the following:

1. Launch the Restore wizard.
2. Select an authentication method.
3. Log in to Office 365.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

1. Select an object that you want to restore:
   - To restore OneDrive, select the necessary OneDrive in the inventory pane.
   - To restore a folder, select the necessary folder in the inventory pane.
   - To restore a document, select a folder that contains the document in the inventory pane and select the document in the preview pane.
2. On the **OneDrive/Folder/Document** tab, click **Restore OneDrive/Restore Folder/Restore Document and select one of the following options:**
   - **Overwrite** — to completely overwrite existing OneDrive data.
   - **Keep** — to preserve existing data and recover items with the *RESTORED* prefix (`RESTORED-<file_name>.ext`).

Alternatively, right-click an object that you want to restore and select **Restore OneDrive/Restore folder/Restore document** > **Overwrite** or **Restore OneDrive/Restore folder/Restore document** > **Keep**.

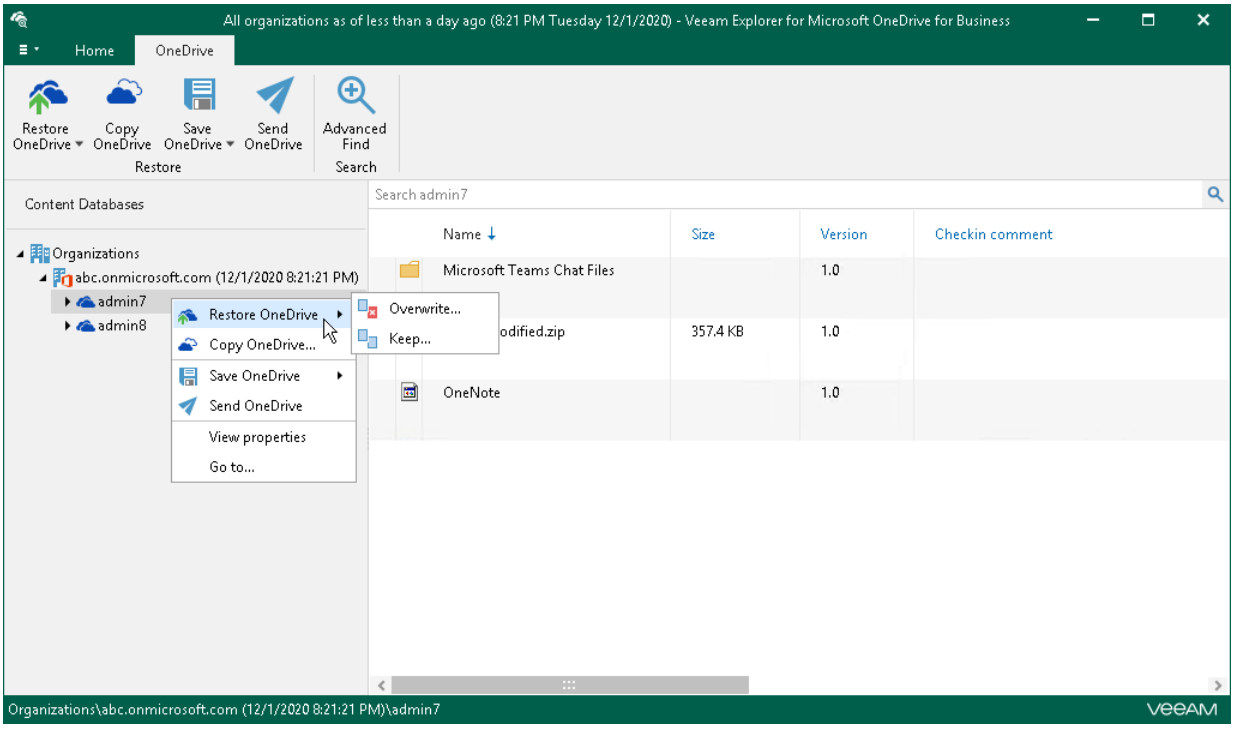

Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication.

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified object back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft OneDrive for Business populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

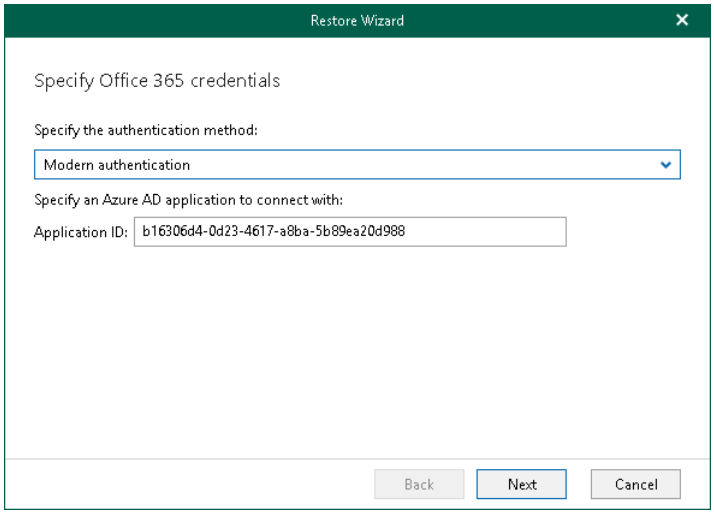

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the SharePoint organization.

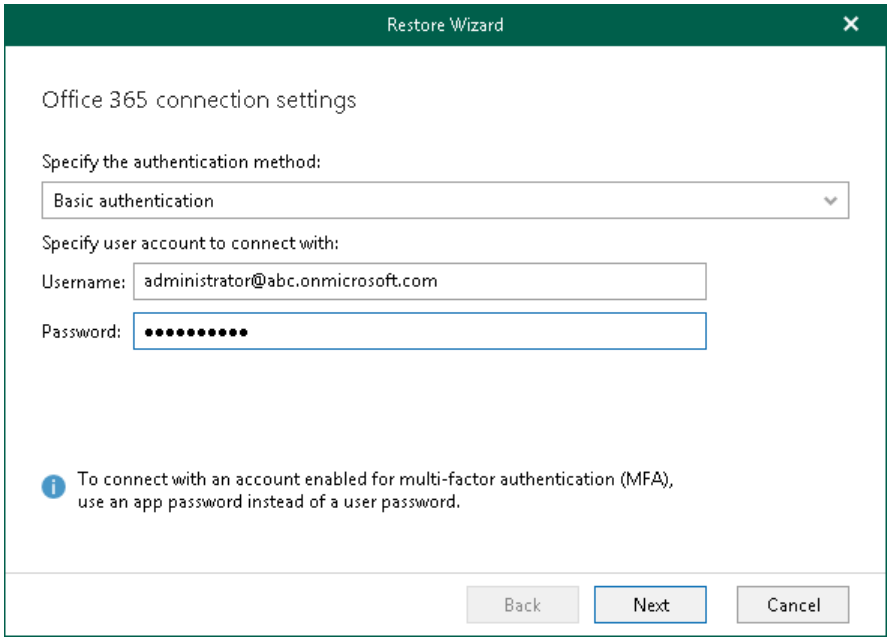

On-Premises Organizations

To authenticate on an on-premises Microsoft SharePoint server, provide a user name and password.

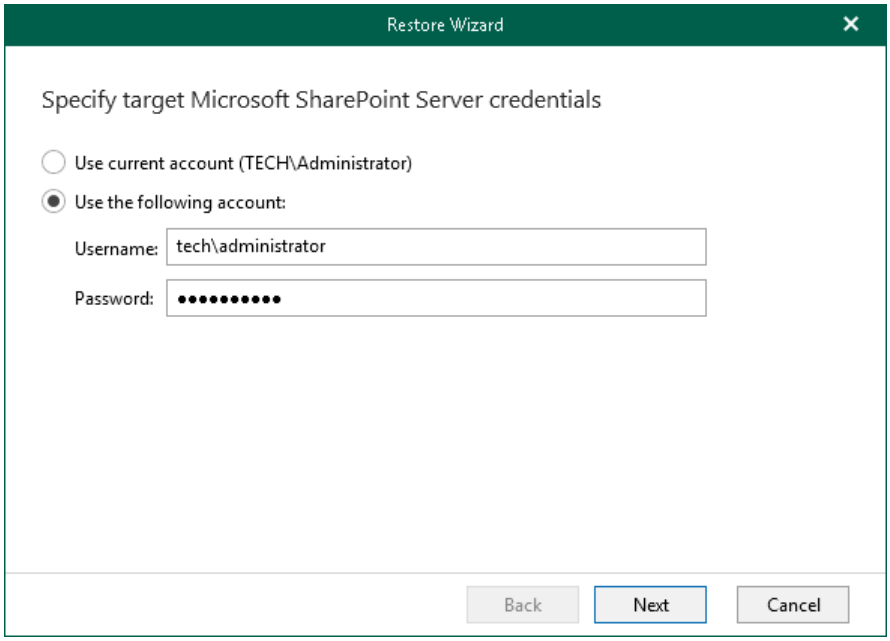

Step 3. Log In to Office 365

This step is available only if you have selected the *Modern authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For Veeam Explorer for Microsoft OneDrive for Business, required permissions are the same as for Veeam Explorer for Microsoft SharePoint. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Restore**.

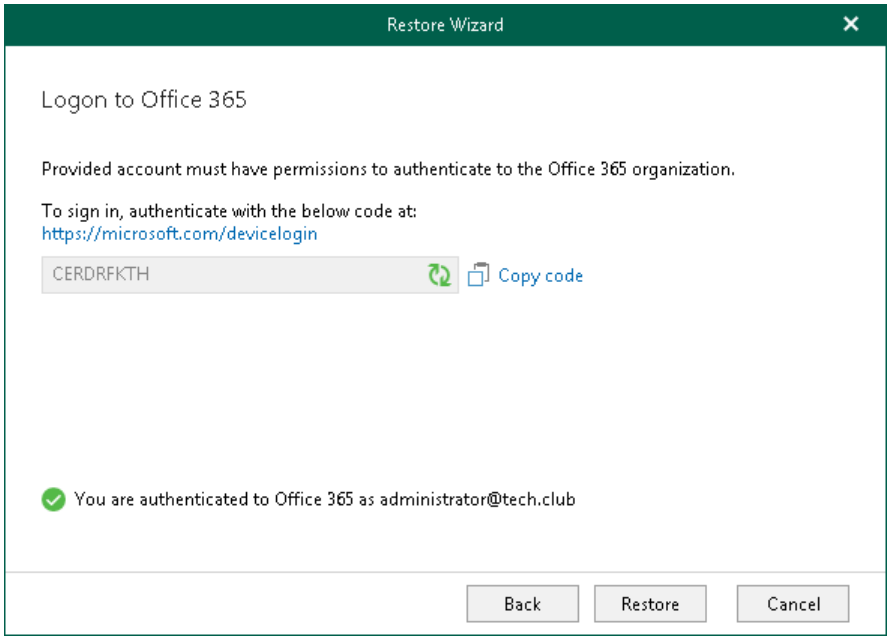

Restoring Multiple OneDrives

To restore multiple OneDrives, do the following:

1. Launch the Restore wizard.
2. Select an authentication method.
3. Log in to Office 365.
4. Select OneDrives.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

1. In the navigation pane, select an organization with OneDrives.
2. On the **Home** tab, click **Restore OneDrives** and select either of the following:
   - **Overwrite** — to completely overwrite existing OneDrives.
   - **Keep** — to preserve existing data and recover items with the *RESTORED* prefix (`RESTORED-<file_name>.ext`).

Alternatively, in the navigation pane, right-click an organization with OneDrives and select **Restore OneDrives** > **Overwrite** or **Restore OneDrives** > **Keep**.

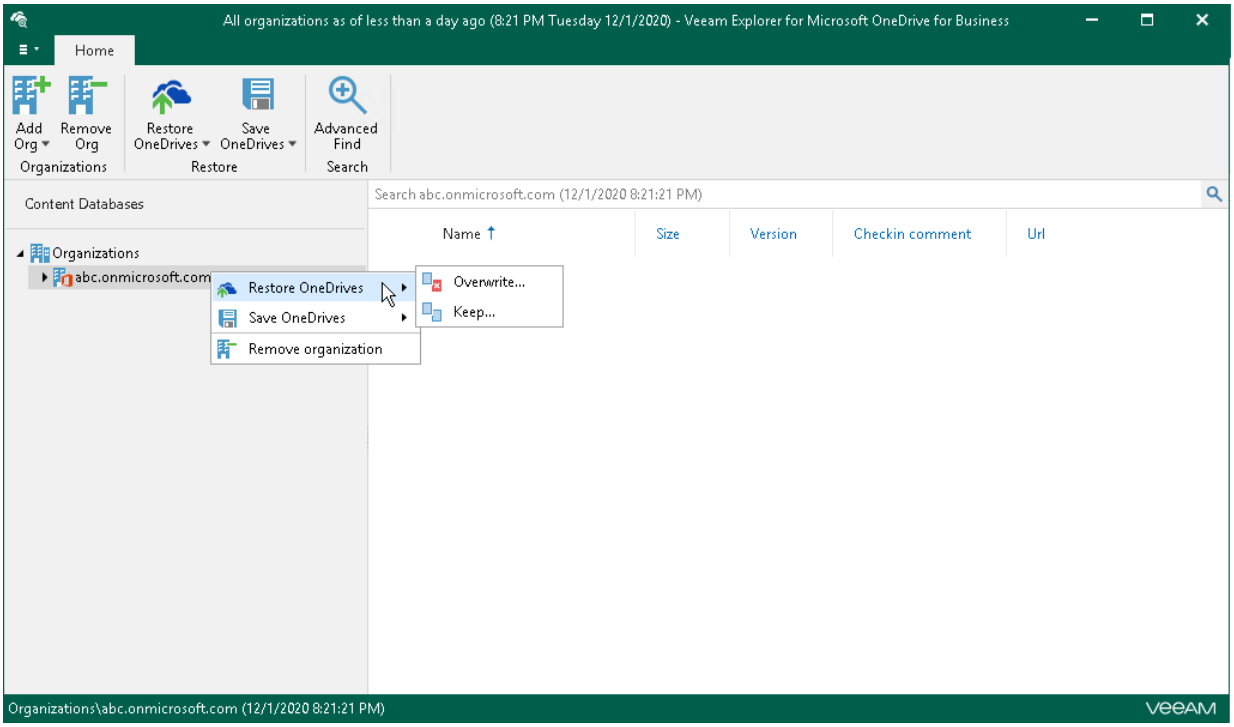

Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication.

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified object back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft OneDrive for Business populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

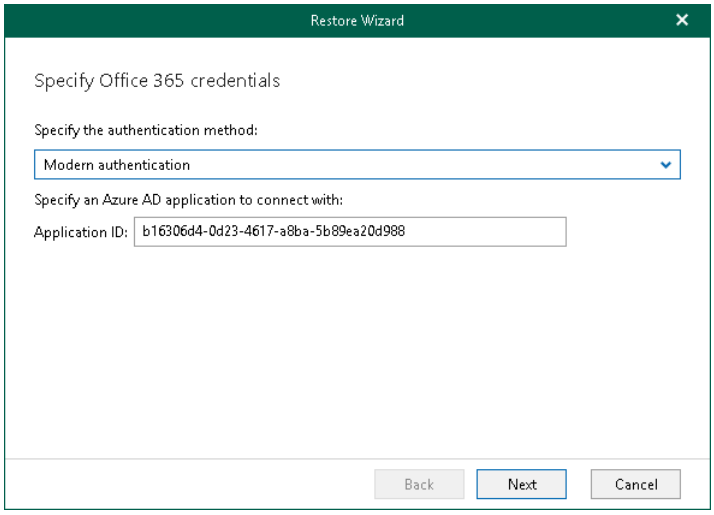

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the SharePoint organization.

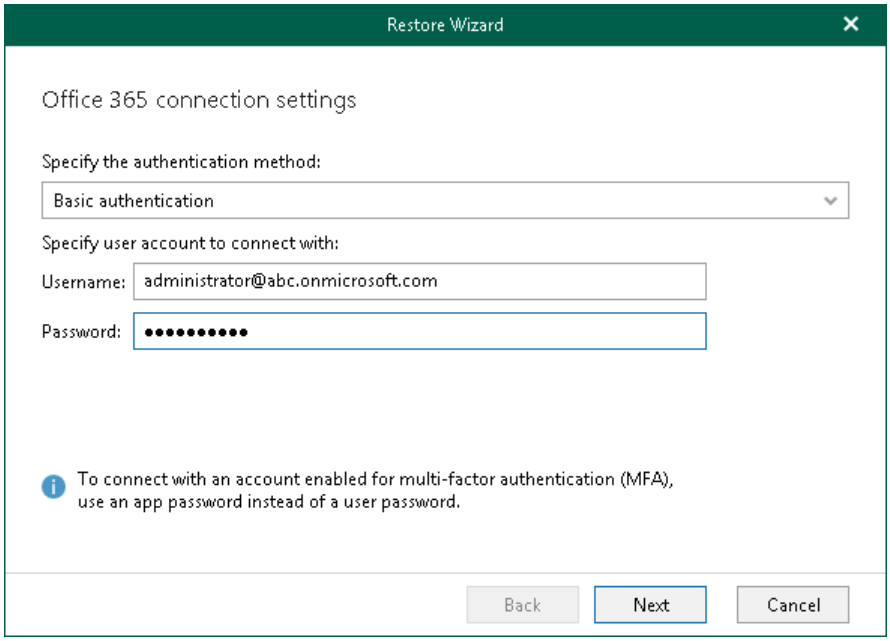

On-Premises Organizations

To authenticate on an on-premises Microsoft SharePoint server, provide a user name and password.

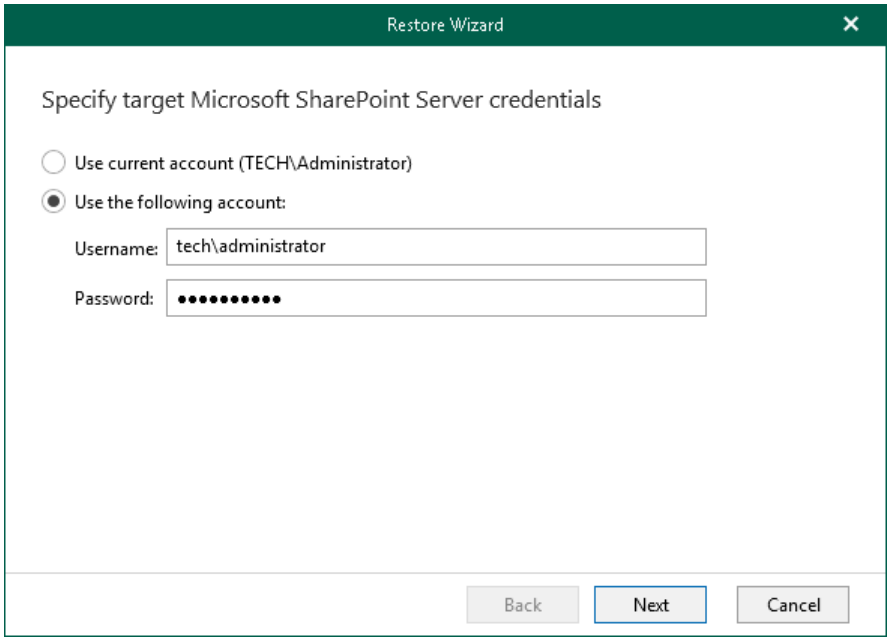

Step 3. Log In to Office 365

This step is available only if you have selected the *Modern authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For Veeam Explorer for Microsoft OneDrive for Business, required permissions are the same as for Veeam Explorer for Microsoft SharePoint. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

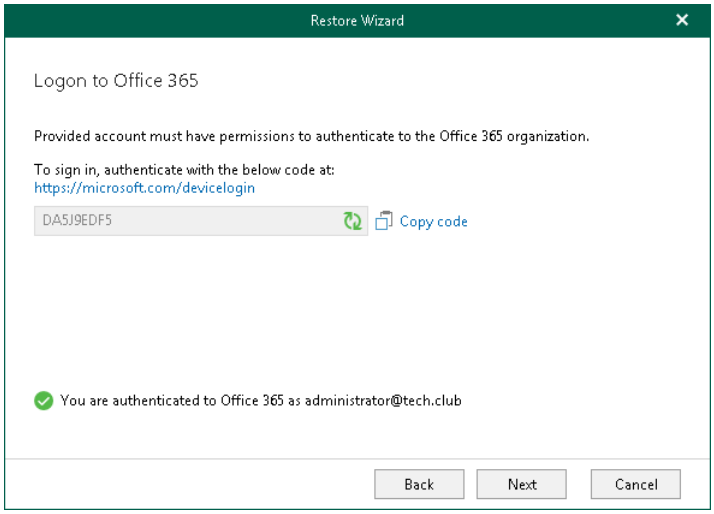

Step 4. Select OneDrives

At this step of the wizard, select OneDrives that you want to restore and click **Restore**.

If you want to check whether user accounts whose OneDrives you want to restore exist, click **Resolve selected accounts**.

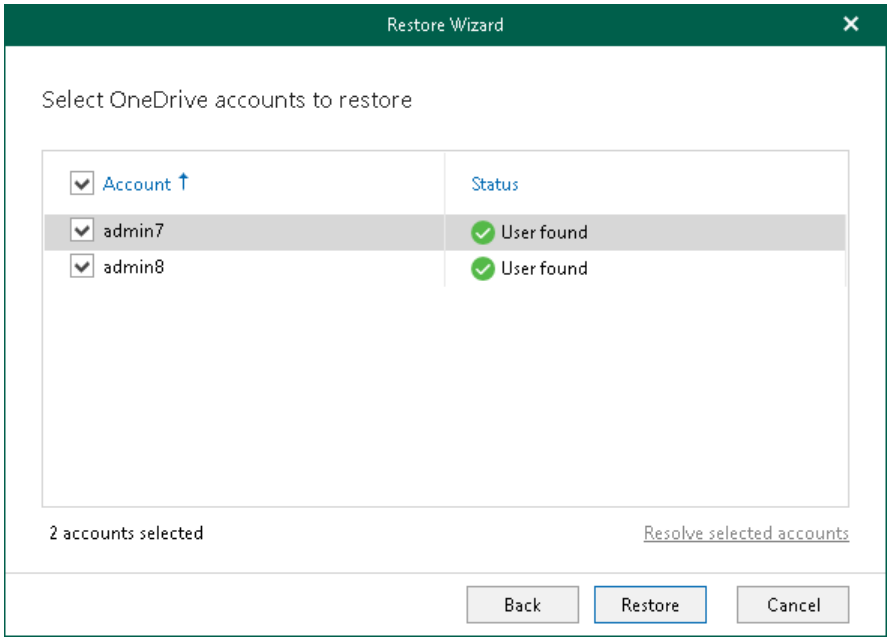

Copying Microsoft OneDrive Data

To copy OneDrive data to the same or different user, do the following:

1. Launch the Restore wizard.
2. Select an authentication method.
3. Log in to Office 365.
4. Specify the target user.
5. Specify the target folder.
6. Specify restore options.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

1. Select an object.
2. On the **Folder** tab, select **Copy OneDrive/Copy Folder/Copy Document** or right-click an object and select **Copy OneDrive/Copy Folder/Copy Document**.

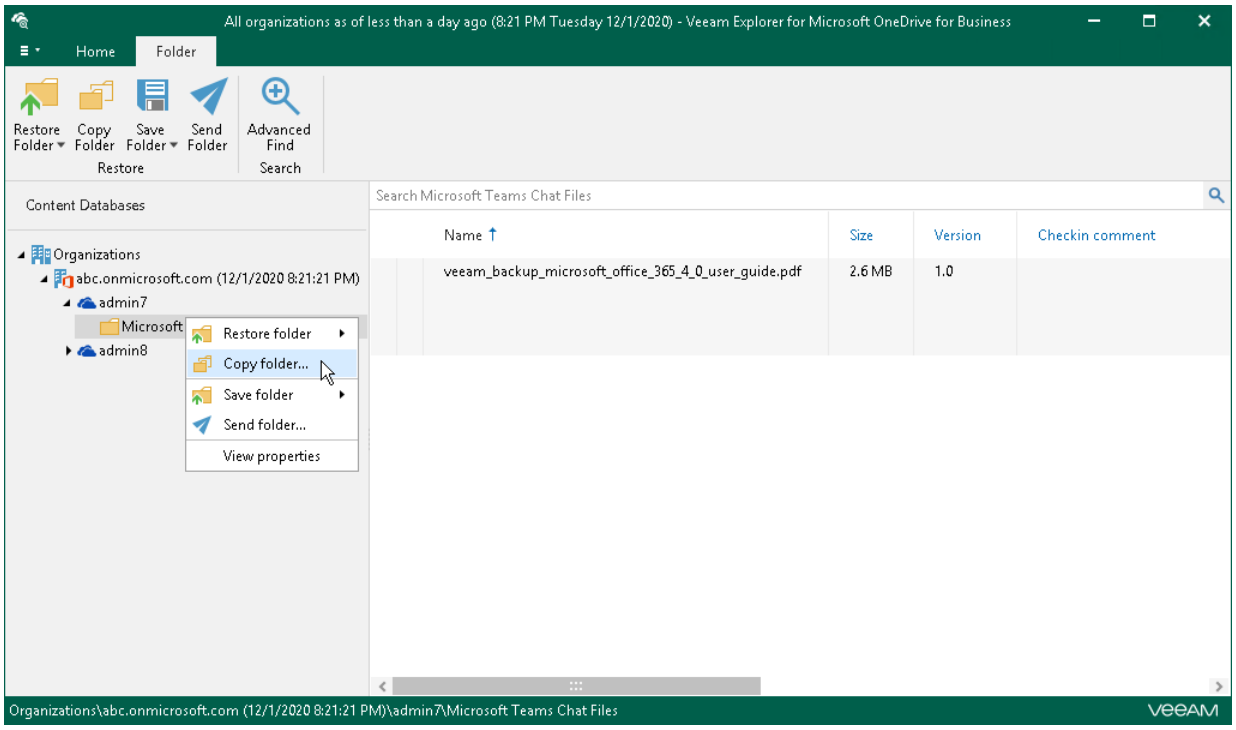

Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication.

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for a restore. Such an application is used to restore the specified object back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for a restore.

   By default, Veeam Explorer for Microsoft OneDrive for Business populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

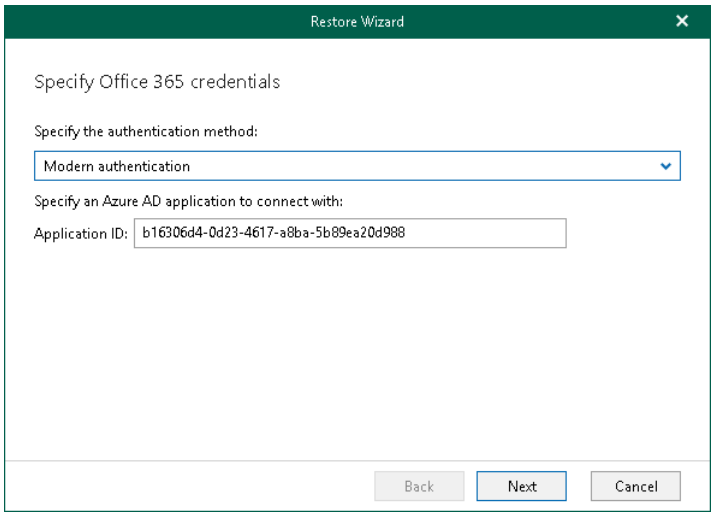

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify authentication method** drop-down list, select *Basic authentication*.
2. Enter credentials to connect to the SharePoint organization.

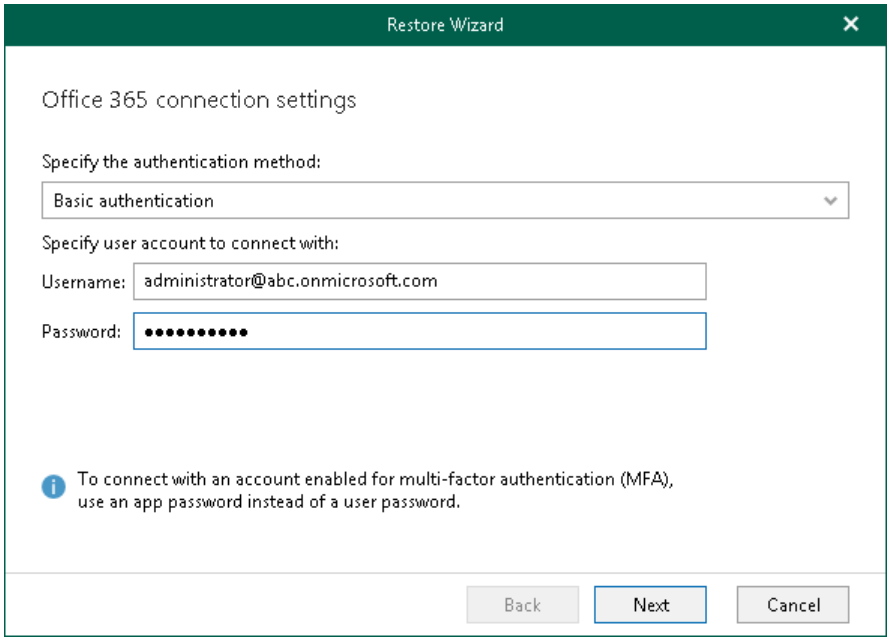

On-Premises Organizations

To authenticate on an on-premises Microsoft SharePoint server, provide a user name and password.

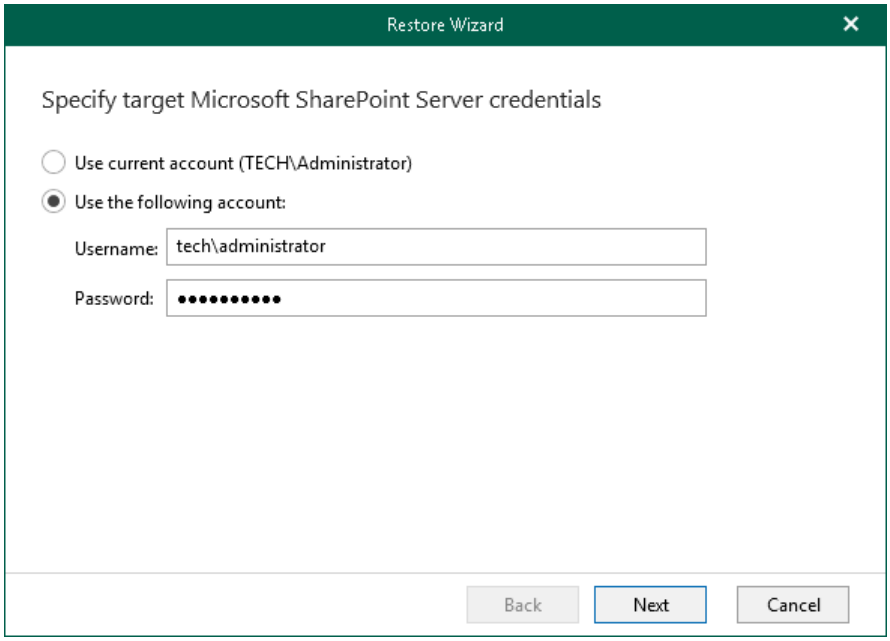

Step 3. Log In to Office 365

This step is available only if you have selected the *Modern authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.
2. Click the link to the Microsoft portal.
3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has the *SharePoint Administrator* role.
4. Return to the **Restore Wizard** window and click **Next**.

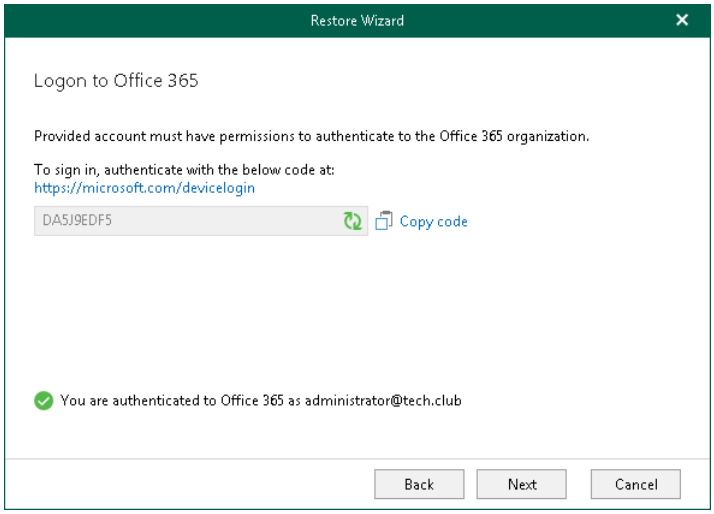

Step 4. Specify Target User

At this step of the wizard, specify the destination user to which you want to copy your data.

You can select either of the following options:

- **Restore to original user**. To restore data back to the original user.
- **Restore to the following user**. To restore data to a custom user.

  To select a user, click **Browse**.

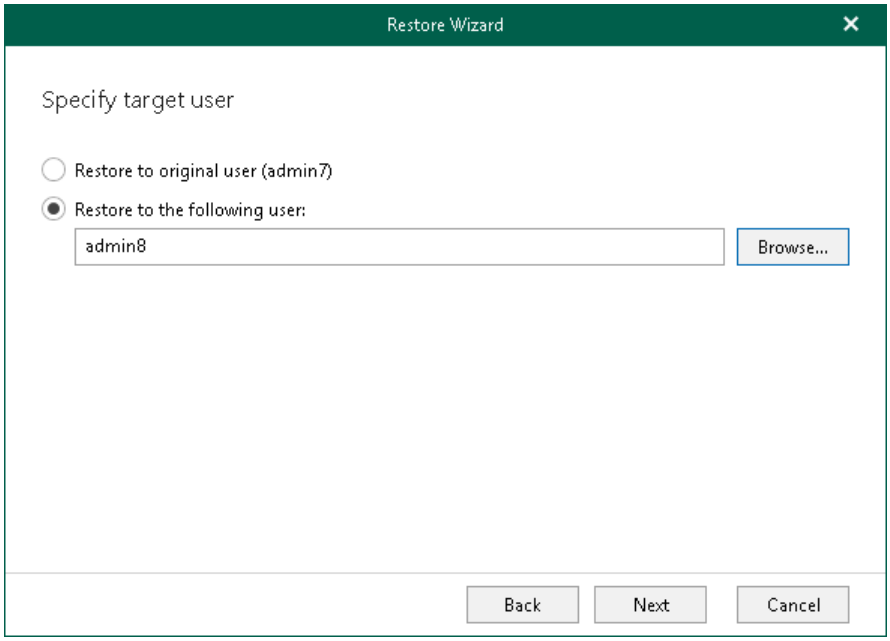

Step 5. Specify Target Folder

At this step of the wizard, specify the target directory to which you want to copy your data.

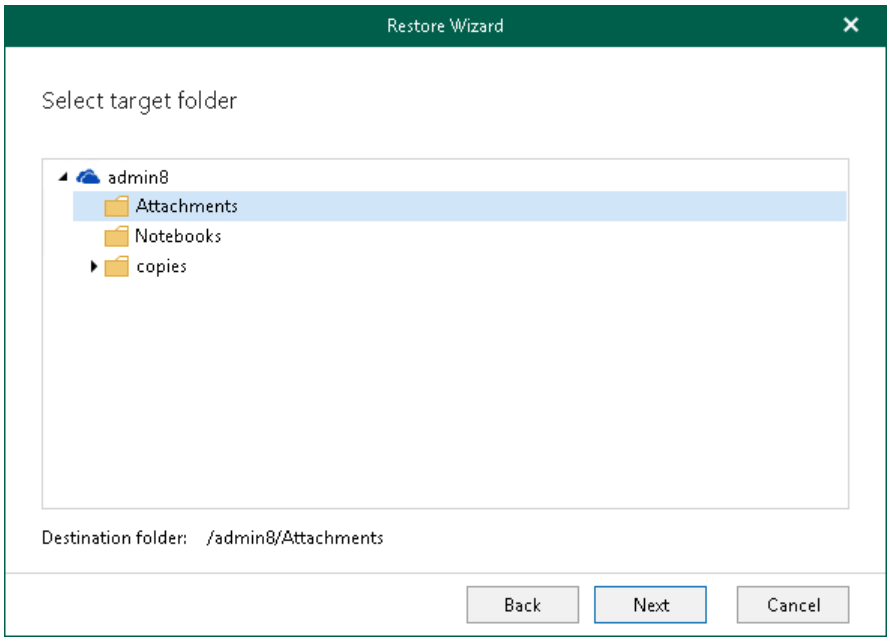

Step 6. Specify Restore Options

At this step of the wizard, specify restore options and click **Restore**.

The restore options are as follows:

- **Changed items**. Allows you to restore data that has been modified in your production environment.
- **Missed items**. Allows you to restore missed items.
- **Restore shared access**. Allows you to restore shared access.
- **History restore options**. Allows you to select a version:
  - **Overwrite**. To overwrite data in the production environment by recovering only the latest version of the document from the backup.
  - **Merge**. To merge an existing version of the document with that of a backup version.

  If not selected, all the versions in the production environment will be replaced with the corresponding data from the backup file.

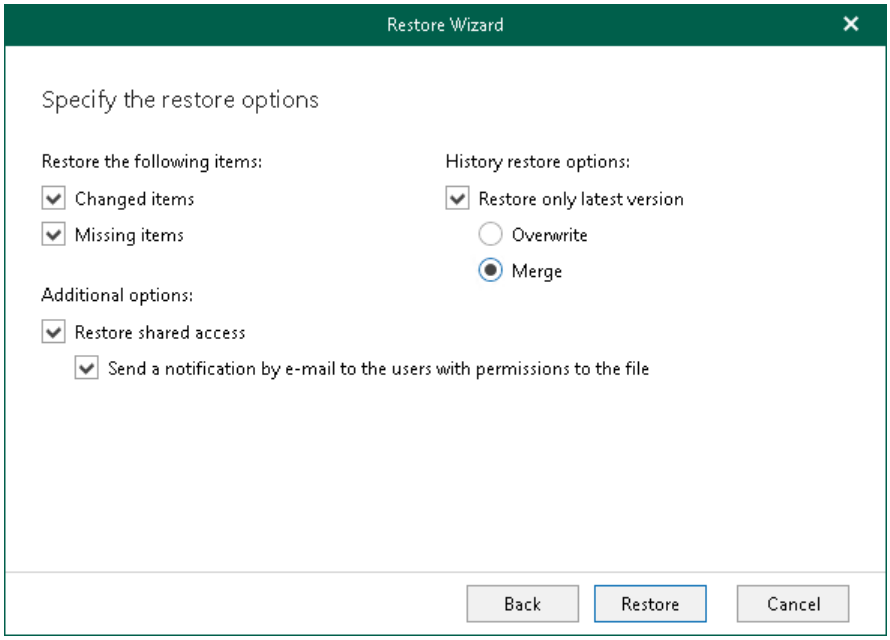

Saving Microsoft OneDrives

Veeam Explorer for Microsoft OneDrive for Business allows you to save OneDrive content to the specified location.

To save data, do the following:

1. Select OneDrive.
2. On the **Home** tab, select **Save OneDrives** > **Save files/Save OneDrive** > **Save files** or right-click an object and select **Save OneDrives** > **Save files/Save OneDrive** > **Save files**.

   Select **Save as ZIP** to save OneDrive documents and/or folders as ZIP.

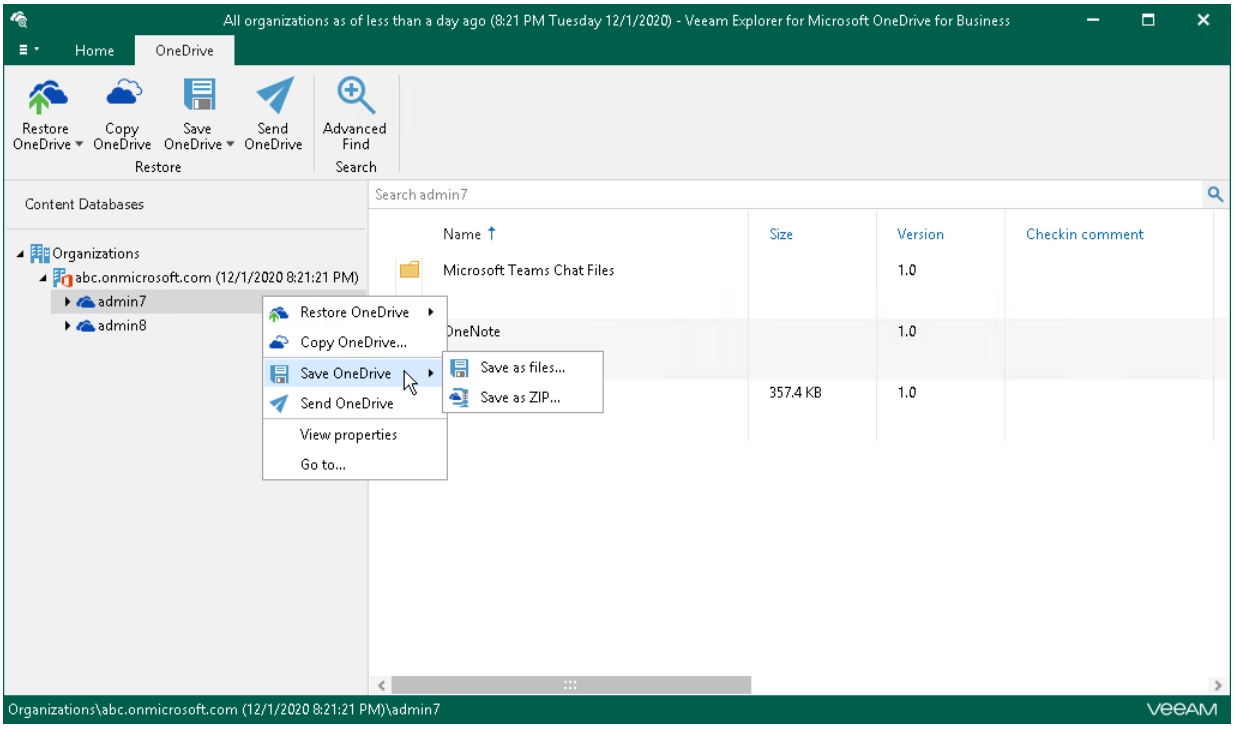

3. Select OneDrive accounts to save.

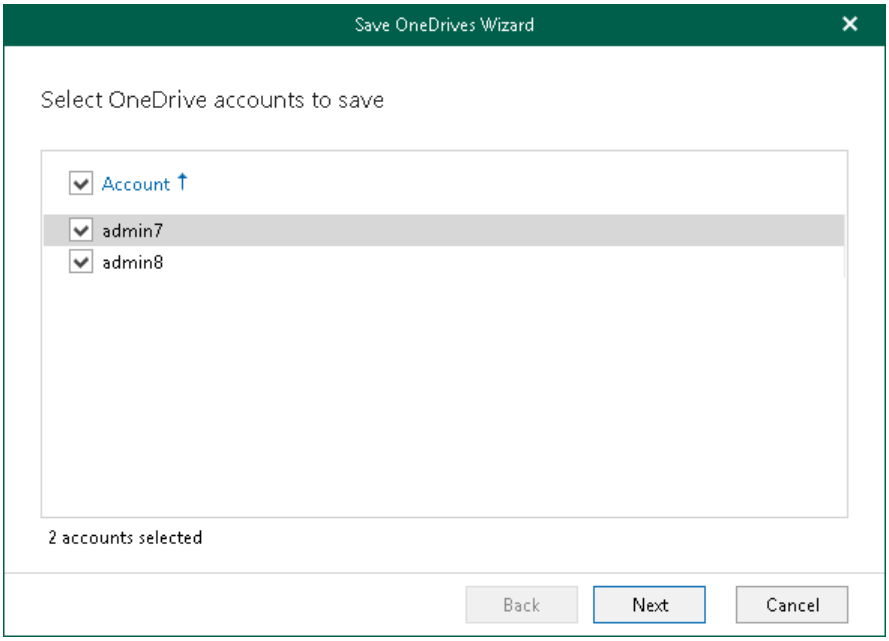

4. Specify a directory to save OneDrive data and click **Finish**.

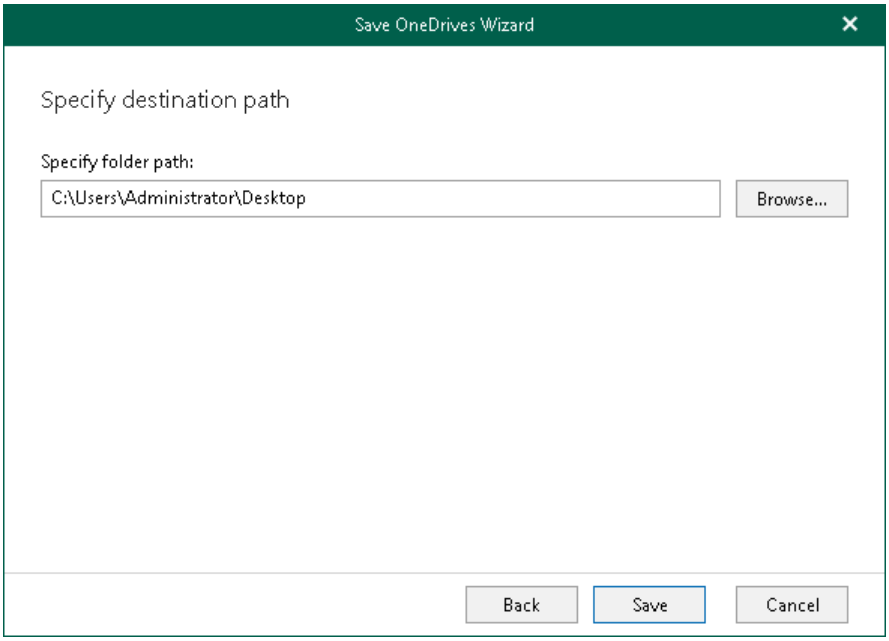

Saving Microsoft OneDrive Documents and Folders

Veeam Explorer for Microsoft OneDrive for Business allows you to save your documents and folders located on users OneDrives to a specified location.

To save a document or folder, do the following:

1. Select a folder or document.
2. On the **Home** tab, select **Save Document** > **Save files/Save Folder** > **Save files** or right-click an object and select **Save Document** > **Save files/Save Folder** > **Save files**.

   Select **Save as ZIP** to save OneDrive documents and/or folders as ZIP.
3. Specify the destination folder and click **Select Folder**.

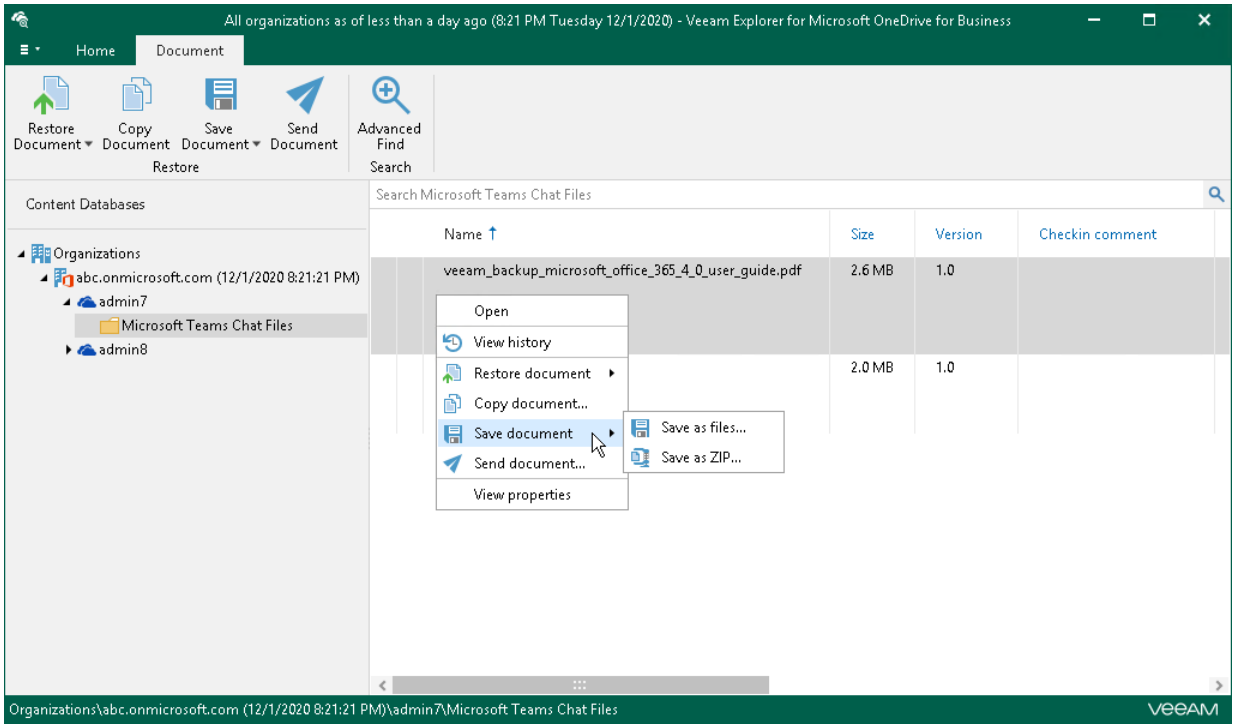

Sending Microsoft OneDrive Documents

Veeam Explorer for Microsoft OneDrive for Business allows you to send OneDrive documents to specified recipients via email.

> **TIP**
>
> Before sending documents, make sure to configure SMTP settings, as described in Configuring SMTP Settings. The amount of data you can send at a time depends on your SMTP server configuration.

To send OneDrive documents via email, do the following:

1. Select OneDrive or a document.
2. On the **Document** tab, select **Send OneDrive/Send Document** or right-click an object and select **Send OneDrive/Send Document**.

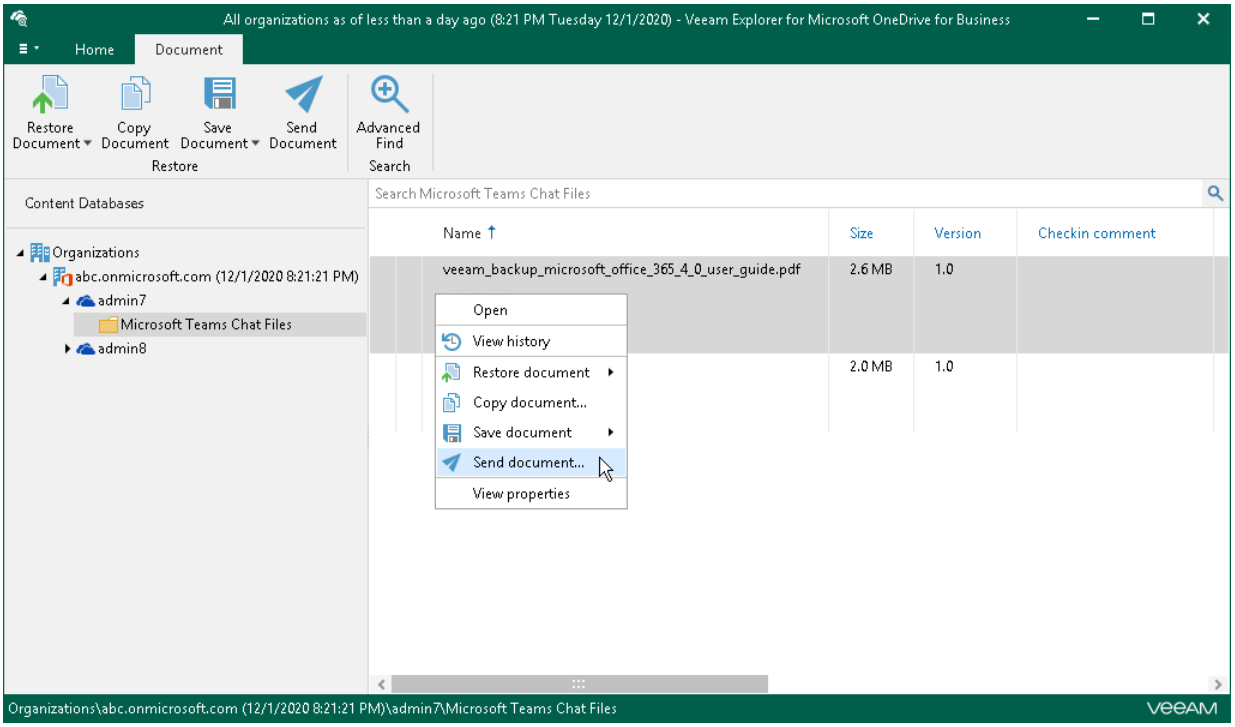

3. Provide a recipient address.

   The **From** field is filled automatically based on the address you have provided when configuring SMTP settings. To edit the message body, click **More Details**.

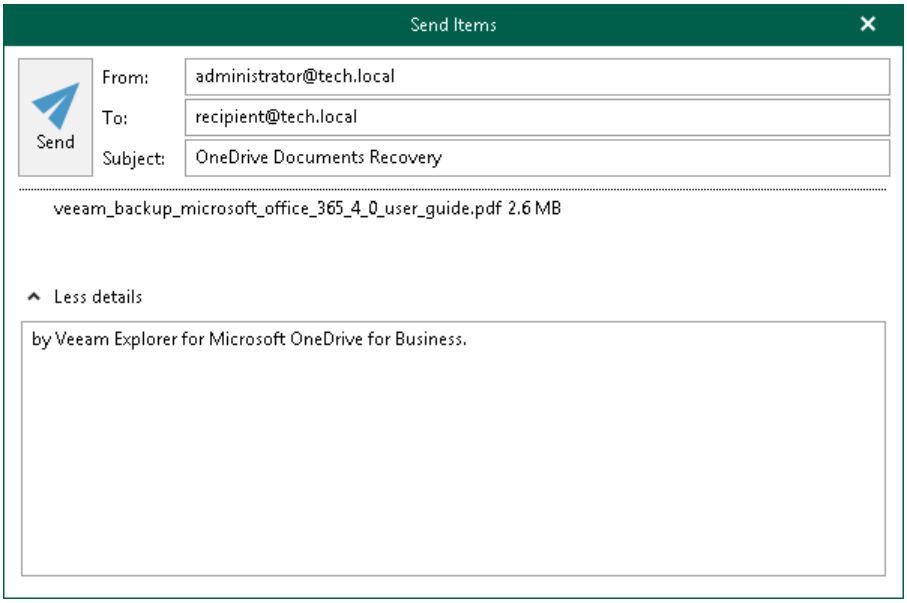

Veeam Explorer for Microsoft Teams

Veeam Explorer for Microsoft Teams allows you to restore Microsoft Teams data from backups created in Veeam Backup for Microsoft Office 365.

You can also save data from backups to a local computer and send data from backups via email.

Planning and Preparation

Continue with this section to learn how to configure your environment before start using Veeam Explorer for Microsoft Teams:

System Requirements

This section lists system requirements for Veeam Explorer for Microsoft Teams.

Component	Requirement
Microsoft Teams	For information about supported versions of Microsoft Office 365 components required for backup of Microsoft Teams data, see the System Requirements section of the Veeam Backup for Microsoft Office 365 User Guide.

Used Ports

For restore of Microsoft Teams data, consider opening the same ports that are required for restore of Microsoft Exchange and Microsoft SharePoint data. For more information, see Used Ports: Veeam Explorer for Microsoft Exchange and Used Ports: Veeam Explorer for Microsoft SharePoint.

Required Permissions

The following table lists required permissions for data restore.

Operation	Required Roles and Permissions
Restore to Microsoft Office 365	To restore Microsoft Teams data, make sure to configure user accounts as follows: **Restore Using Basic Authentication Method** • The user account must have a Microsoft Office 365 license that permits access to Microsoft Teams API. The minimum sufficient license is Microsoft Teams Exploratory experience. • To restore Microsoft Teams data, the user account must have the *Teams Administrator* role. **Restore Using Modern App-Only Authentication Method** • The user account must have a Microsoft Office 365 license that permits access to Microsoft Teams API. The minimum sufficient license is Microsoft Teams Exploratory experience. • To restore Microsoft Teams data, the user account must have the *Teams Administrator* role. • Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide. • If you restore data with Azure AD applications using a certificate, make sure that your Azure AD application is granted the required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

Considerations and Limitations

This section covers considerations and known limitations of Veeam Explorer for Microsoft Teams.

Data Restore

Consider the following when planning to restore Microsoft Teams data:

- Use of modern authentication with legacy protocols allowed is not supported for data restore with Veeam Explorer for Microsoft Teams.
- You can restore Microsoft Teams data to the original organization only.
- Veeam Explorer for Microsoft Teams does not change roles for team owners during restore. For example, you create a backup of your organization, and then change role for a team member from *Member* to *Owner*. In this case, if you restore this team member from the backup, Veeam Explorer for Microsoft Teams will not set their role to *Member*.
- Restore of OneNote notebooks from backups of Microsoft Teams data for organizations with enabled security defaults using multi-factor authentication is not supported.
- When restoring a channel tab, Veeam Explorer for Microsoft Teams does not preserve relation between the link to a file published on the tab and the file itself. You will need to link the tab to the file manually after restore. This limitation does not apply to the scenario where you restore an entire team.
- Veeam Explorer for Microsoft Teams does not restore posts to their original location in the team channel. Instead, Veeam Explorer for Microsoft Teams exports posts to a file of the HTML format, creates a separate tab in the original channel and attaches the HTML file to this tab.

Launching Application and Exploring Backups

To open Veeam Explorer for Microsoft Teams and load backups, do either of the following:

- Use the **Explore** option to load backups created in Veeam Backup for Microsoft Office 365.

  For more information, see the Data Restore section of the Veeam Backup for Microsoft Office 365 User Guide.

  When loading backups located in object storage repositories, Veeam Explorer for Microsoft Teams will notify you about the location of these backups. You can select the **Don't show this message again** check box to unsubscribe from such notifications. To disable these notifications, clear the **Show pop-up notifications on startup** check box, as described in Configuring Advanced Settings.

- Go to **Start**, click **Veeam Explorer for Microsoft Teams** and connect to the Veeam Backup for Microsoft Office 365 server or manually open databases that contain Microsoft Office 365 data, as described in the Standalone Databases Management section.

If you open Veeam Explorer for Microsoft Teams from the **Start** menu on a machine that runs the Veeam Backup & Replication console and Veeam Explorers only, specify the following:

- The domain name or IP address of a server to which you want to connect
- The port number
- User credentials

  The account must be a member of the *Local Administrator* group on the target server. To use the account under which Veeam Explorer for Microsoft OneDrive for Business is running, select **Use Windows session authentication**.

To save the connection shortcut to the desktop, click the **Save shortcut** link in the bottom-left corner.

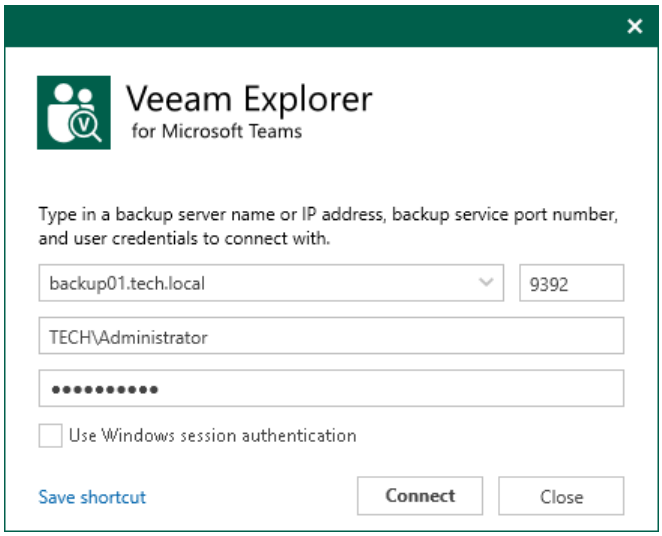

Understanding User Interface

Veeam Explorer for Microsoft Teams provides you with the convenient user interface that allows you to perform required operations in a user-friendly manner.

Main Menu

The main menu comprises the following features:

- **General Options**. Allows you to configure program options. For more information, see General Application Settings.
- **Help**.
  - **Online help**. Opens the online web help page.
  - **About**. Shows current product information.
- **Exit**. Closes the program.

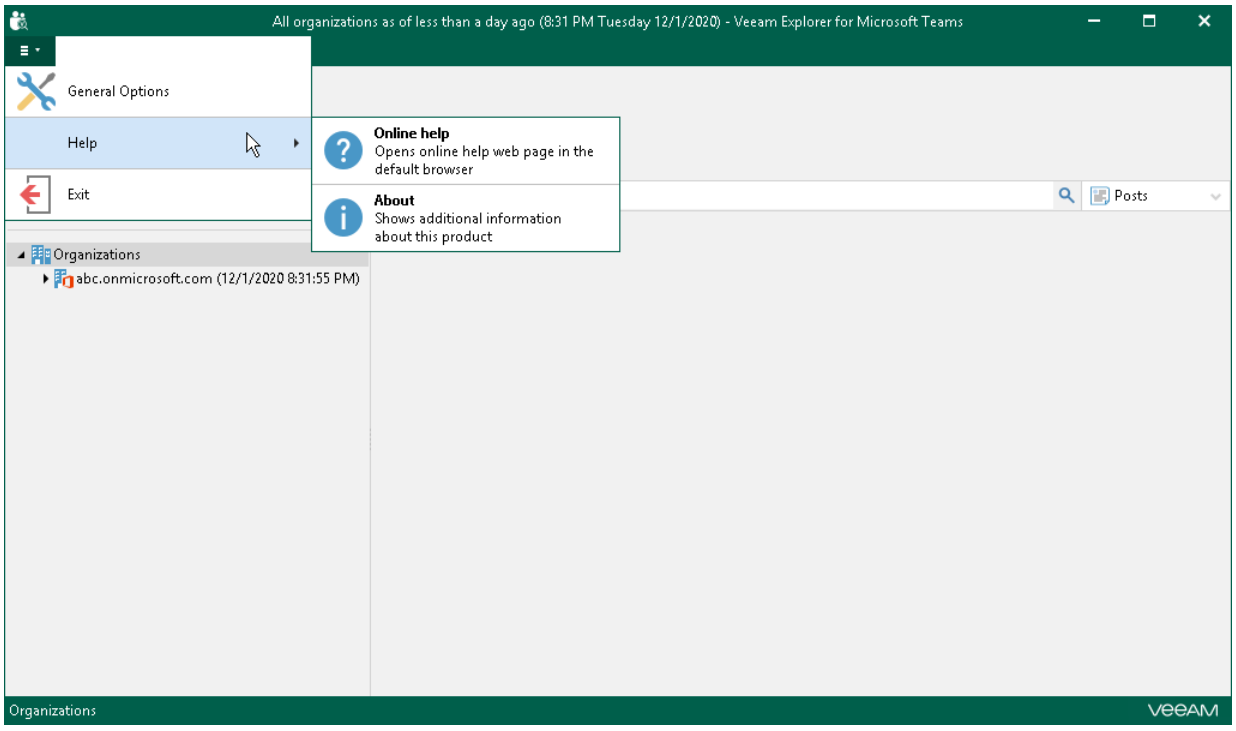

Main Application Window

The main application window might be divided into three categories:

1. The ribbon menu, which contains general program commands organized into logical groups.
2. The navigation pane, which allows you to browse through the hierarchy of your backup files.
3. The preview pane, which shows you the details about objects you have selected in the navigation pane.

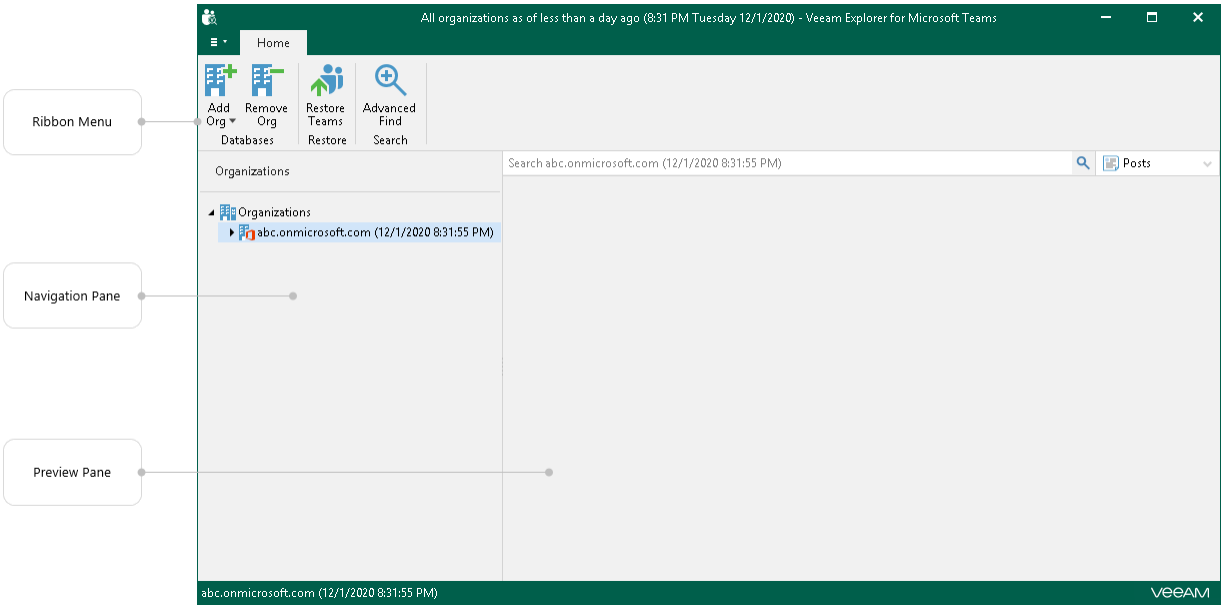

Browsing, Searching and Viewing Items

Continue with this section to learn about the following operations with the backed-up data:

- Browsing backup content
- Viewing team properties
- Opening posts and files
- Searching for objects in a backup file
- Using the advanced search capabilities

Browsing Backups

To view the contents of a backup file, you use the navigation pane which displays the database structure containing your Microsoft Teams objects: teams, channels and channel tabs. Veeam Explorer for Microsoft Teams groups channel tabs in the following nodes of the navigation pane:

- **Posts**. Contains the *Posts* tab of the channel.
- **Files**. Contains the *Files* tab of the channel.
- **Other tabs**. Contains other tabs of the channel that link to sites, documents, applications, and so on.

After you select an object in the navigation pane, you can see its content in the preview pane.

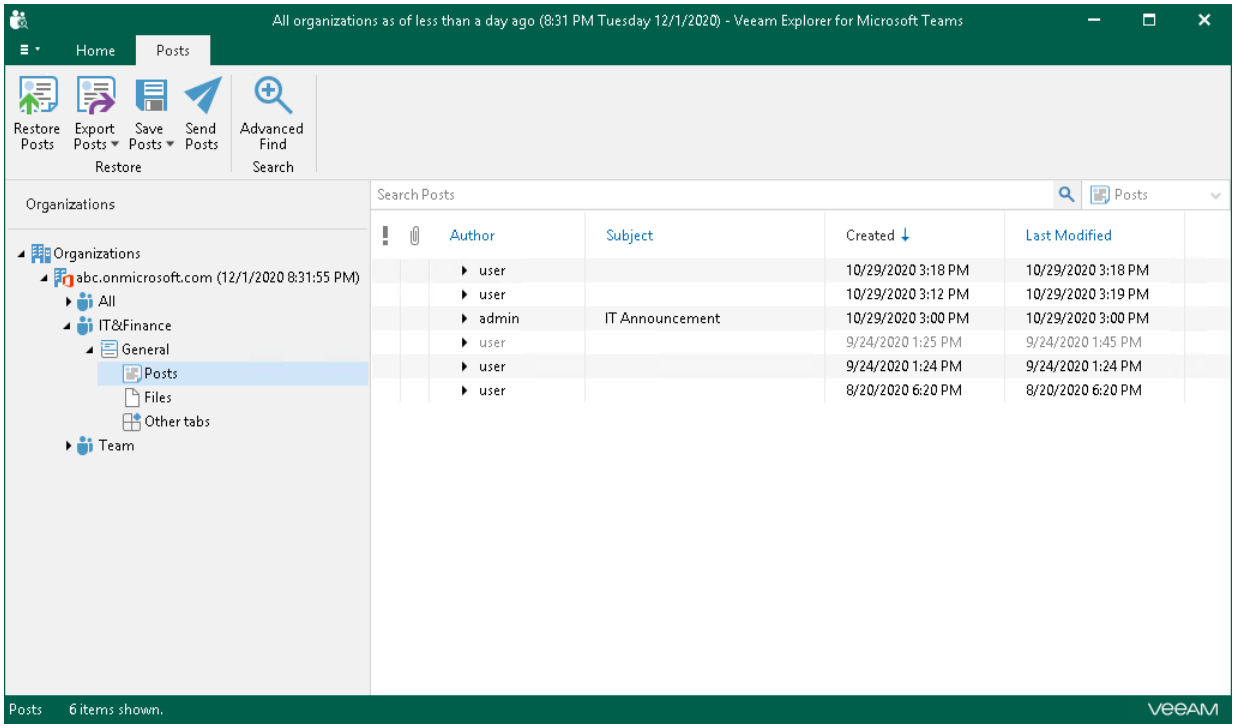

Viewing Team Properties

You can view properties of a backed-up team. This may be useful, for example, in case you want to decide whether to restore team properties during restore of a team.

To view team properties, in the navigation pane, right-click a necessary team and select **View properties**.

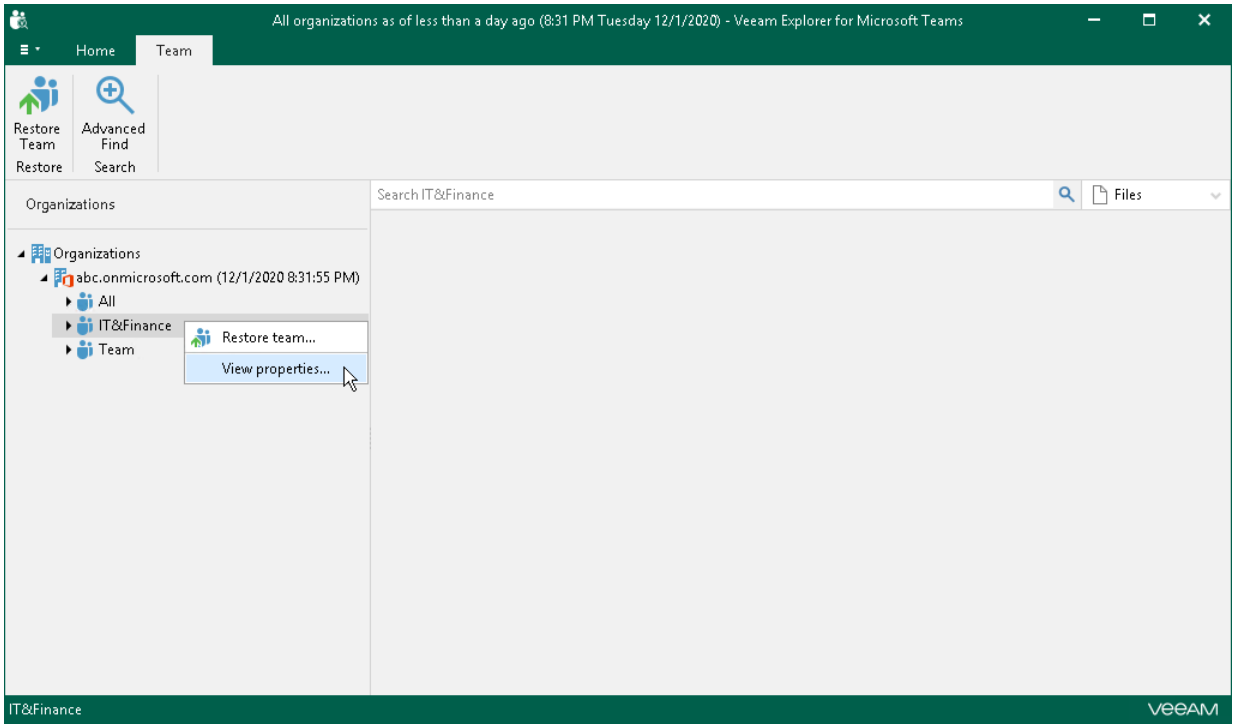

Opening Posts and Files

You can use Veeam Explorer for Microsoft Teams to open backed-up posts and files.

To open a post or file, do the following:

1. In the navigation pane, browse to the channel whose post or file you want to open and click the necessary node: **Posts** or **Files**.
2. In the preview pane, right-click a post or file and select **Open**.

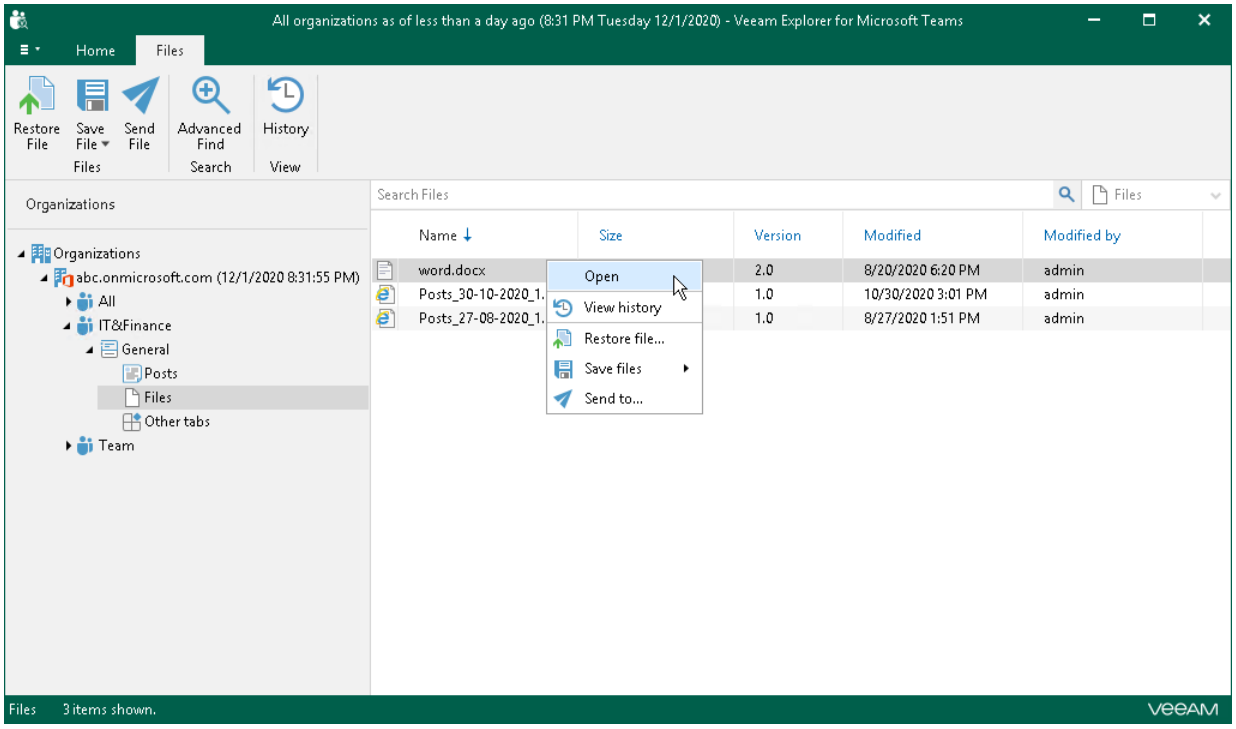

Searching for Objects

The search mechanism allows you to find items (posts, files and tabs) matching specified search criteria.

To search for required items, do the following:

1. In the navigation pane, select a node that contains an item you want to find: **Posts**, **Files** or **Other tabs**.

   Alternatively, in the navigation pane, select a team or channel whose items you want to find. Then, from the drop-down list next to the search field, select what type of items you want to find: *Files*, *Posts* or *Other tabs*.
2. Enter a search query to the search field.

> **NOTE**
>
> To find the exact phrase, use double quotes. For example, "*Document*".

You can narrow your search results by specifying various search criteria using the *criteria:value* format. You can also use logical upper-cased operators such as *AND*, *OR* and *NOT* along with wildcard characters such as *** and *?*.

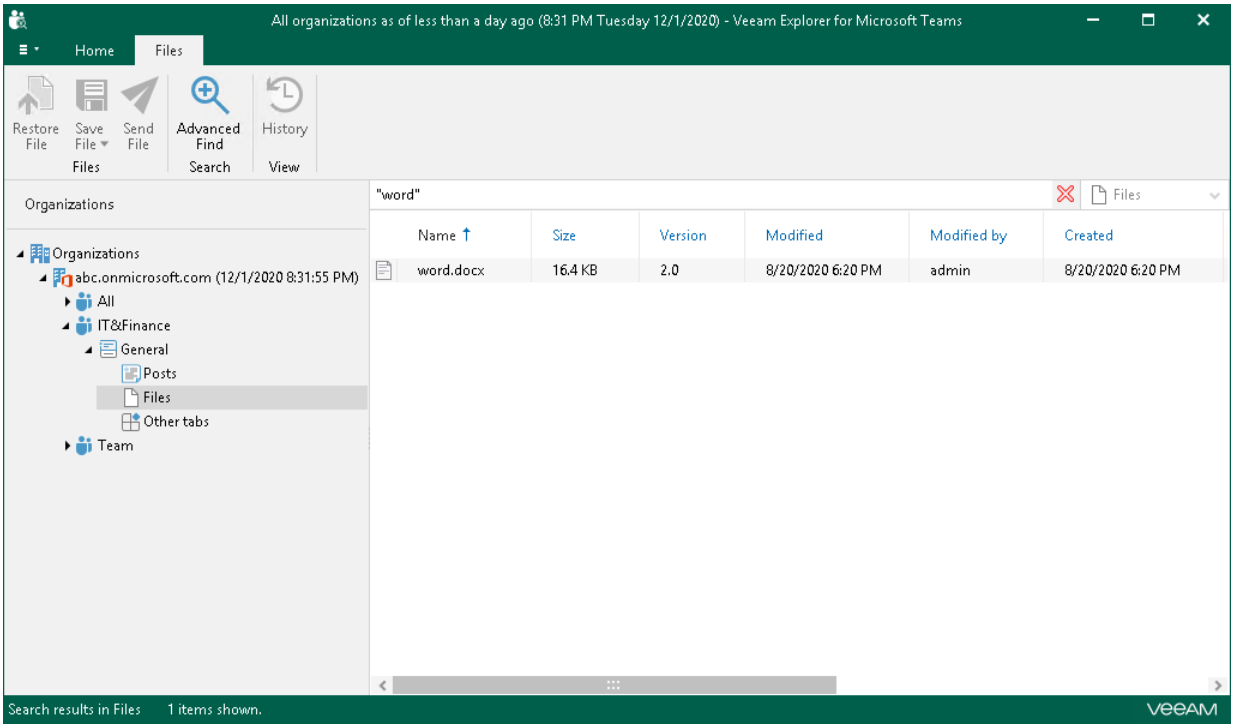

Using Advanced Find Capabilities

The **Advanced Find** mechanism allows you to define your search criteria more precisely.

> **NOTE**
>
> The **Advanced Find** mechanism is supported for posts and files only. The mechanism is not supported for tabs.

For example, to find a file whose name starts with *"word"*, do the following:

1. In the preview pane, select a node that contains objects you want to find and click **Advanced Find** on the ribbon.
2. In the **Define search criteria** section, select **Category** > **File fields**.
3. In the **Field** list, select **Name**.
4. In the **Condition** list, select **starts with**.
5. In the **Value** field, specify a file name.
6. Click **Add To List**, then click **Start**.

To remove a filter, click the cross mark next to it. To remove all configured filters, click **Reset**.

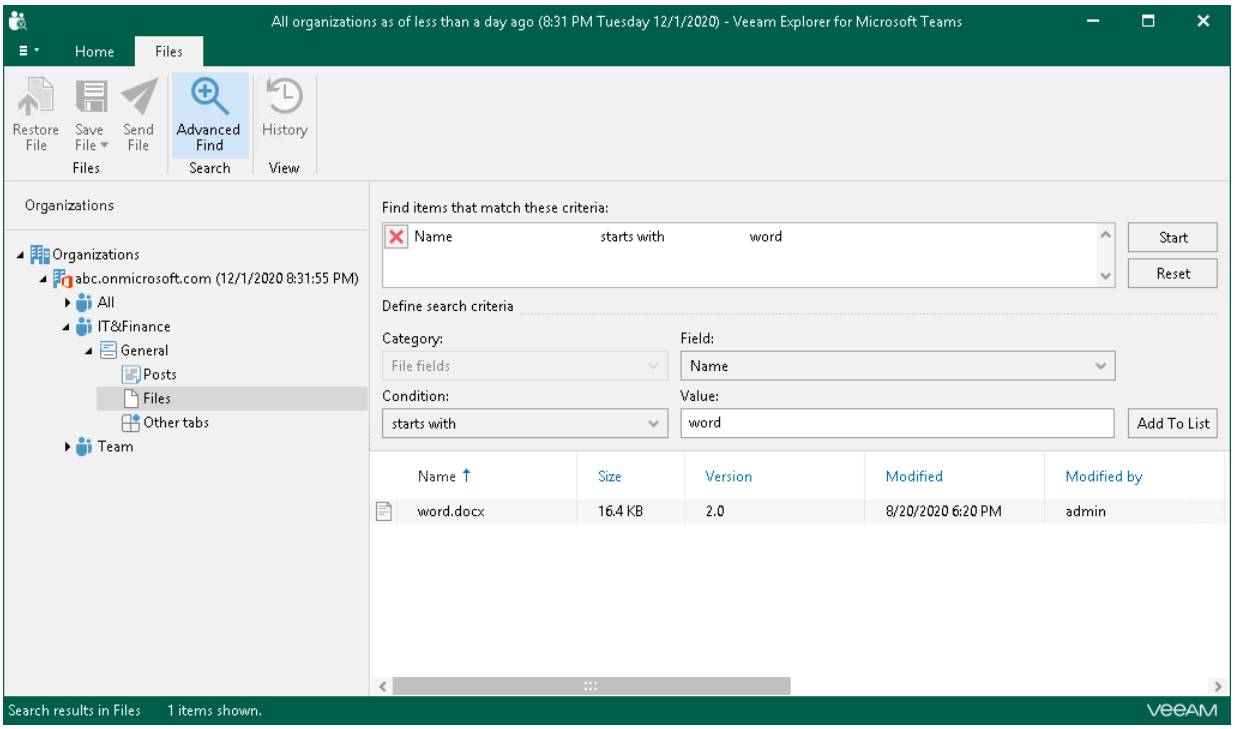

General Application Settings

Continue with this section to learn more about configuring general application settings and components.

Configuring SMTP Settings

To send Microsoft Teams items as attachments, you must configure SMTP server settings.

To configure SMTP settings, do the following:

1. In the main menu, click **General Options**.
2. Open the **SMTP Settings** tab.
3. Select the **Configure SMTP settings** check box and do the following:
   a. In the **Server** field, specify a DNS name or IP address of the SMTP server.
   b. In the **Port** field, specify a port number of the specified SMTP server.

      The default port number is *25*. If you want to use SSL data encryption, specify port *587*.
   c. In the **From** field, specify the sender email address. This address will appear in the **From** field when sending Microsoft Teams posts or files. For more information, see Sending Posts and Files.
   d. If the specified SMTP server requires an authentication for outgoing mail, select the **Use authentication** check box and provide valid credentials.
   e. Select the **Enable SSL security** check box to enable SSL data encryption.
4. Click **Send** to send a test email message.
5. Click **Apply**.

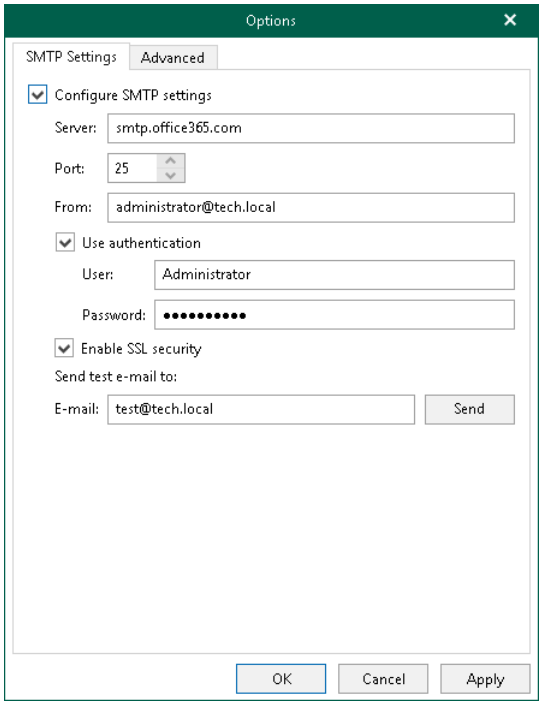

Configuring Advanced Settings

Veeam Explorer for Microsoft Teams allows you to configure advanced settings such as extended logging mode and pop-up notifications.

To configure advanced settings, do the following:

1. In the main menu, click **General Options**.
2. Open the **Advanced** tab.
3. If you want to enable extended logging mode to collect logs that contain additional information on specific operations, select the **Enable extended logging** check box.

   After enabling extended logging mode, you can go back to the application and perform actions for which you want to collect additional information. Then you can collect logs.
4. By default, Veeam Explorer for Microsoft Teams displays notifications when loading backups from object storage repositories. If you want to disable these notifications, clear the **Show pop-up notifications on startup** check box.

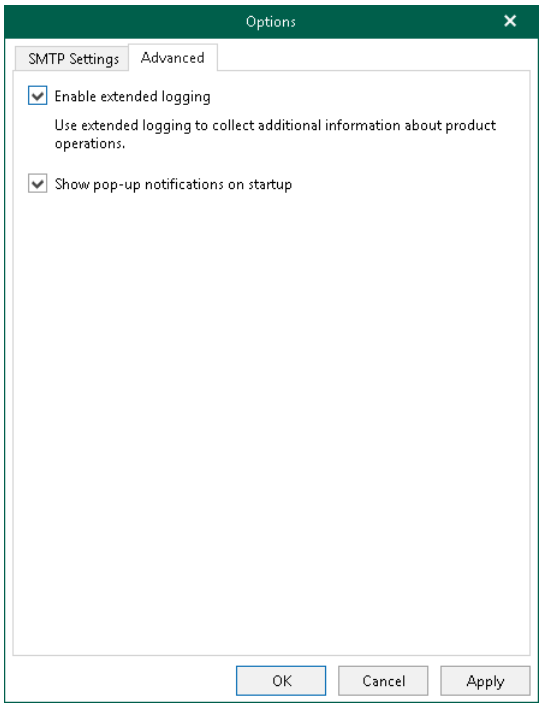

Standalone Databases

This section explains how to add standalone databases to the Veeam Explorer for Microsoft Teams scope and how to connect to a Veeam Backup for Microsoft Office 365 server and service provider server to view backups located on these servers.

Adding Veeam Backup for Microsoft Office 365 Databases

To manually add databases that store Microsoft Office 365 organization data, do the following:

1. On the ribbon, click **Add Org** > **Veeam Backup for Microsoft Office 365 database**.

   Alternatively, in the preview pane, right-click **Organizations** and select **Veeam Backup for Microsoft Office 365 database**.

2. Specify the database file location and log directory.

3. Click **Open**.

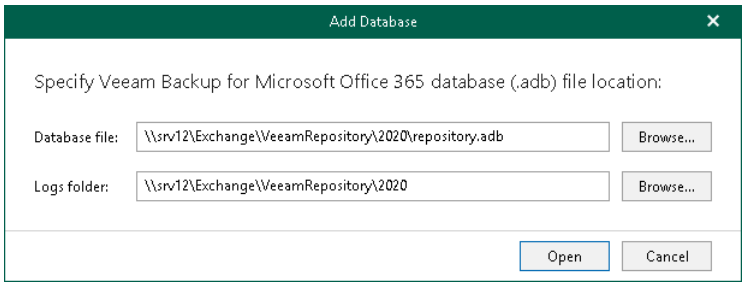

Adding Veeam Backup for Microsoft Office 365 Server

You can use Veeam Explorer for Microsoft Teams to connect to another Veeam Backup for Microsoft Office 365 server and add its databases to the Veeam Explorer scope.

To connect to another Veeam Backup for Microsoft Office 365 server, do the following:

1. On the ribbon, click **Add Org** > **Veeam Backup for Microsoft Office 365 server**.

   Alternatively, in the preview pane, right-click **Organizations** and select **Veeam Backup for Microsoft Office 365 server**.

2. Specify settings to connect to the Veeam Backup for Microsoft Office 365 server and click **Connect**.

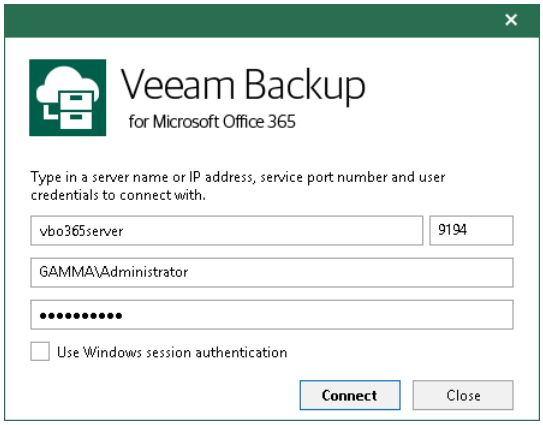

Adding Veeam Backup for Microsoft Office 365 Service Provider

You can use Veeam Explorer for Microsoft Teams to connect to a service provider server to add its backups to the application scope and perform a restore directly from such backups.

Consider the following:

- Both Veeam Explorer for Microsoft Teams and Veeam Backup & Replication must be installed on the same server.
- At least one service provider must be added to Veeam Backup & Replication, as described in the Connecting to Service Providers section of the Veeam Cloud Connect Guide.

To connect to a Veeam Backup for Microsoft Office 365 service provider, do the following:

1. On the ribbon, click **Add Org** > **Veeam Backup for Microsoft Office 365 service provider**.

   Alternatively, in the preview pane, right-click **Organizations** and select **Veeam Backup for Microsoft Office 365 service provider**.

2. From the **Select service provider and authentication method** drop-down list, select a service provider to which you want to connect and choose an authentication method that can be either *Modern authentication* or *Basic authentication*.

4. [For *Modern authentication*] Select a region to which an Office 365 organization of the selected service provider belongs.

5. [For *Modern authentication*] Copy the device code and use it to sign in to Microsoft Azure Active Directory.

   For *Basic authentication*, you must enter a user name and password to connect to the service provider server.

6. If you want to select a point-in-time state as of which you want to load Office 365 backups, click **Specify point in time** and select the necessary option:

   - **Use the latest available state**. To load data as of the latest backup state.
   - **Use the following point in time**. To load data as of the selected point in time.

     To select a point in time, use the calendar control.

7. Select the **Show items that have been deleted by user** check box to see all the items that were removed.
8. Select the **Show all versions of items that have been modified by user** check box if you want to load modified versions.
9. Click **Connect**.

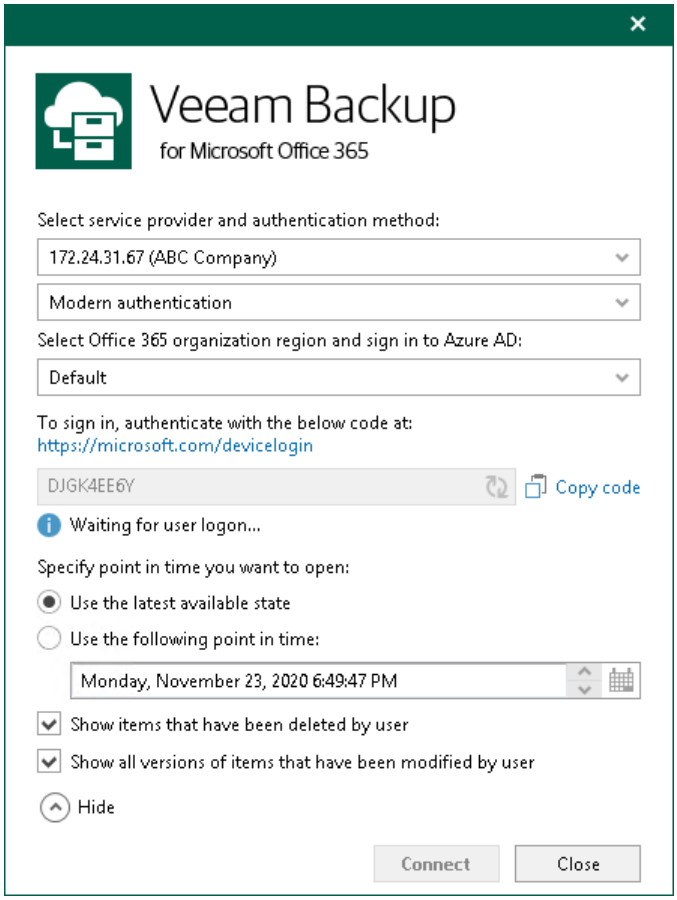

Removing Standalone Databases

Veeam Explorer for Microsoft Teams allows you to remove an organization from the application scope when you no longer need it.

To remove an organization from the application scope, right-click an organization in the navigation pane and select **Remove organization**.

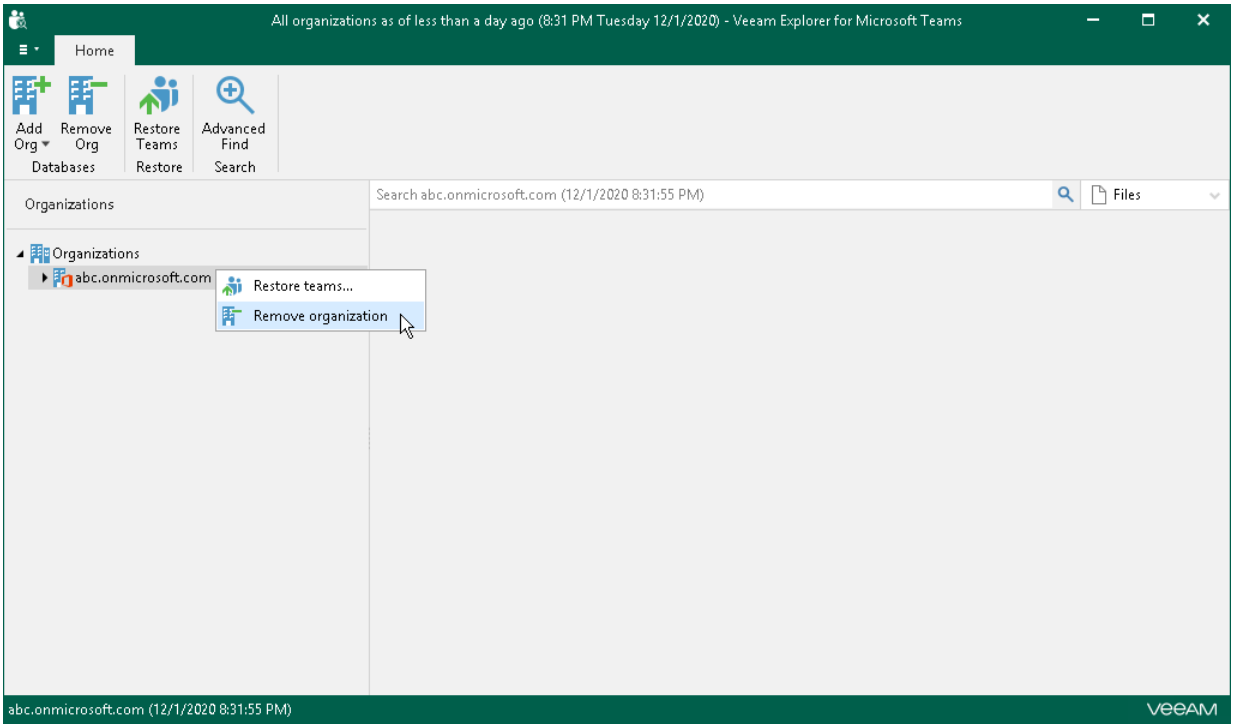

Data Restore

This section explains how to restore Microsoft Teams data to Microsoft Office 365 organizations.

Restoring Teams

You can use Veeam Explorer for Microsoft Teams to restore all teams of a Microsoft Office 365 organization or a specific team of an organization.

> **TIP**
>
> Before restoring team data, make sure to read Considerations and Limitations.

To restore teams, do the following:

1. Launch the Restore wizard.
2. Select an authentication method.
3. Log in to Office 365.
4. Select teams to restore.
5. Specify restore options.
6. Select team members to restore.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

- If you want to restore multiple teams of a Microsoft Office 365 organization, do either of the following:
  - In the navigation pane, select an organization whose teams you want to restore and click **Restore Teams** on the ribbon.
  - In the navigation pane, right-click an organization whose teams you want to restore and select **Restore teams**.

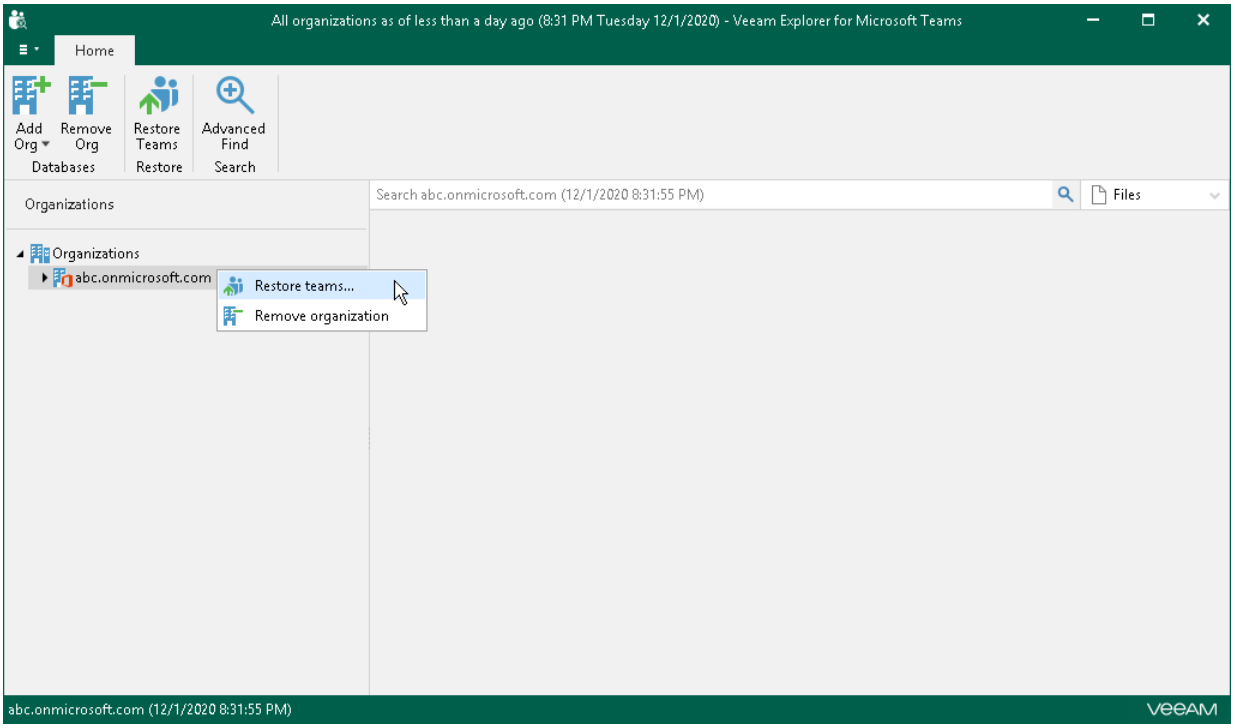

- If you want to restore a specific team, do either of the following:
  - In the navigation pane, select a team that you want to restore and click **Restore Team** on the ribbon.
  - In the navigation pane, right-click a team that you want to restore and select **Restore team**.

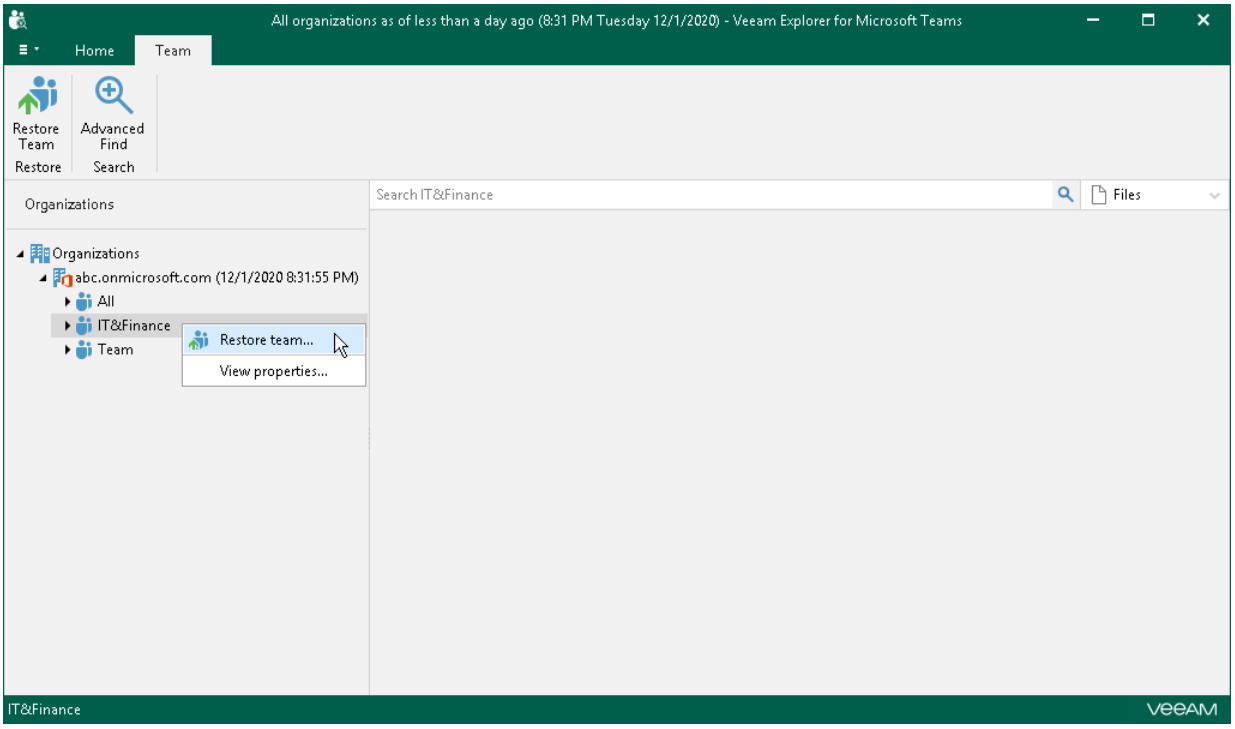

Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication and specify authentication settings.

- Modern authentication
- Basic authentication

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for data restore. Such an application is used to restore the specified team back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for data restore.

   By default, Veeam Explorer for Microsoft Teams populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

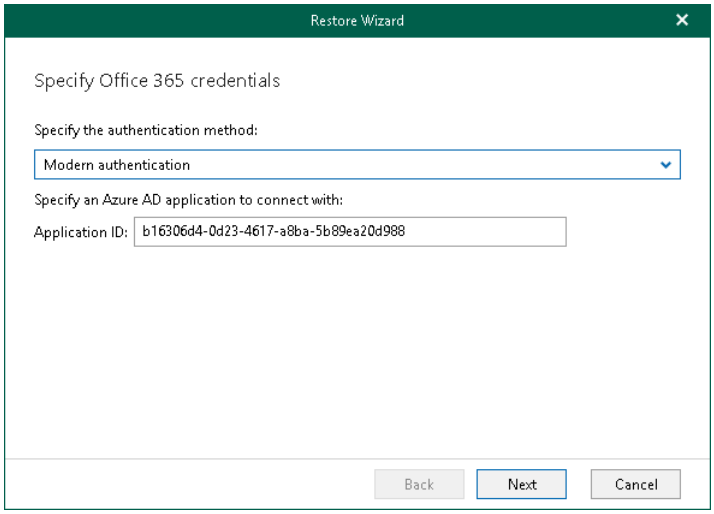

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. In the **Username** and **Password** fields, enter credentials to connect to the Microsoft Office 365 organization.

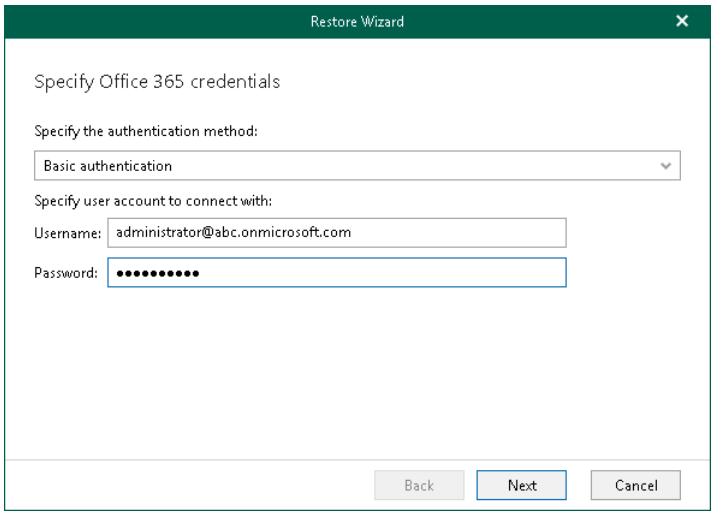

Step 3. Log In to Office 365

This step is available if you selected the *Modern authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

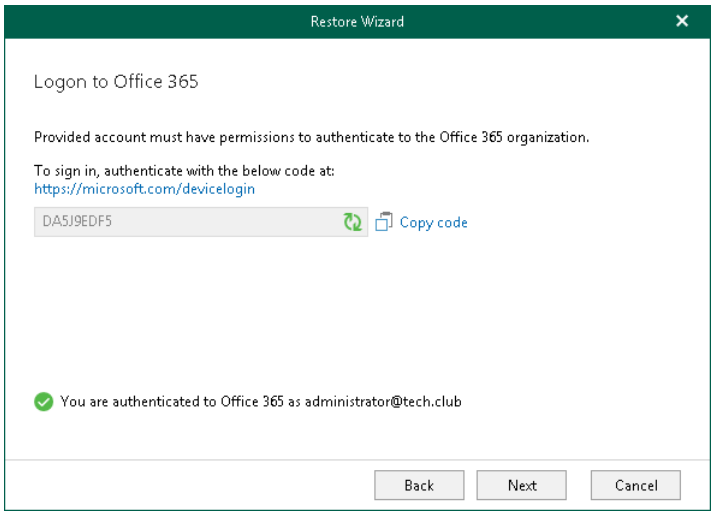

Step 4. Select Teams

This step is available if you launched the **Restore** wizard to restore multiple teams.

At this step of the wizard, select check boxes next to the teams that you want to restore.

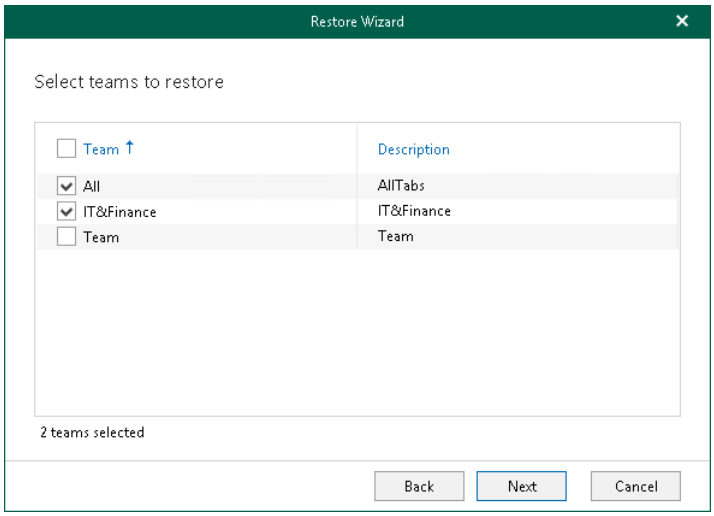

Step 5. Specify Restore Options

At this step of the wizard, specify data restore options:

1. Select the **Changed items** check box if you want to restore items that have changed since the time when the backup was created. When you select this option, Veeam Explorer for Microsoft Teams overwrites existing items in the original team.
2. Select the **Missing items** check box if you want to restore items that are missing in the original team. For example, some of the items were removed and you want to restore them from the backup.
3. Select the **Restore team settings** check box if you want to replace the current team settings with team settings from the backup.

   You cannot select this check box if you clear both the **Changed items** and **Missing items** check boxes.

   For information on how to view team settings in the backup, see Viewing Team Properties.
4. If you want to restore information about team members and their roles from the backup, make sure that the **Restore membership and their permissions** check box is selected.
   - If you launched the **Restore** wizard to restore a specific team, you will be able to review team members at the next step of the wizard. Click **Next**.
   - If you launched the **Restore** wizard to restore multiple teams, you cannot review team members. Click **Restore** to start the restore process.

> **TIP**
>
> If you clear the **Restore membership and their permissions** check box, you can start the restore operation immediately. To do this, click **Restore**.

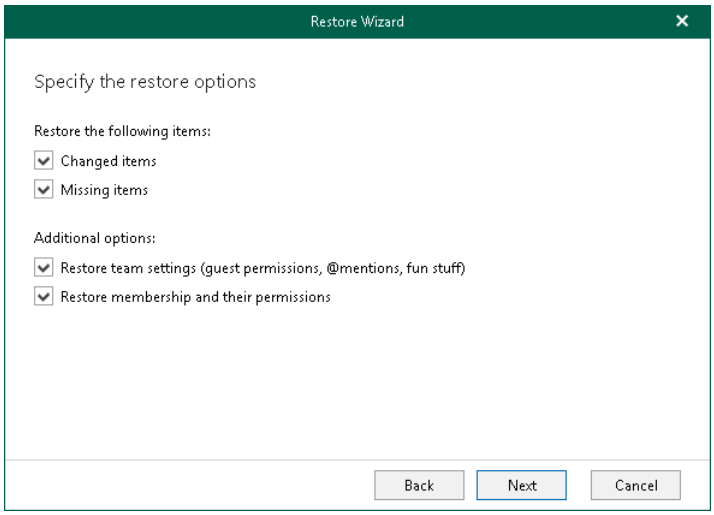

Step 6. Review Team Members

This step is available if you launched the **Restore** wizard to restore a specific team and selected the **Restore membership and their permissions** check box at the previous step of the wizard.

At this step of the wizard, review the list of users whose membership and roles in the team will be restored and click **Restore**.

To quickly find the necessary team members, use the search field at the top of the wizard window.

> **NOTE**
>
> Veeam Explorer for Microsoft Teams does not change roles for team owners during restore. For example, you create a backup of your organization, and then change role for a team member from *Member* to *Owner*. In this case, if you choose to restore this team member from the backup, Veeam Explorer for Microsoft Teams will not set their role to *Member*.
>
> In the opposite case, if a team member role was changed from *Owner* to *Member* after the backup was created, and you choose to restore team membership from the backup, Veeam Explorer for Microsoft Teams will elevate the role to *Owner* during restore.

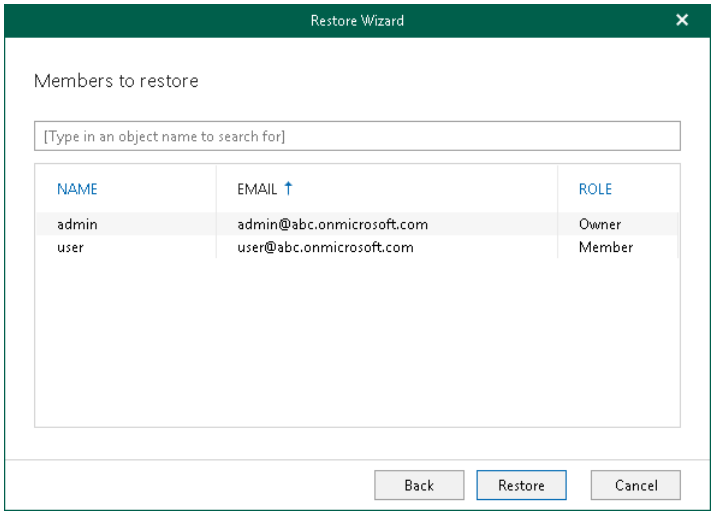

Restoring Channels and Tabs

You can use Veeam Explorer for Microsoft Teams to restore individual channels of a team. You can restore an entire channel or specific tabs of a channel.

To restore channels and tabs, do the following:

1. Launch the Restore wizard.
2. Select an authentication method.
3. Log in to Office 365.
4. Specify restore options.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

- If you want to restore a channel, do either of the following:
  - In the navigation pane, select a channel that you want to restore and click **Restore Channel** on the ribbon.
  - In the navigation pane, right-click a channel that you want to restore and select **Restore channel**.

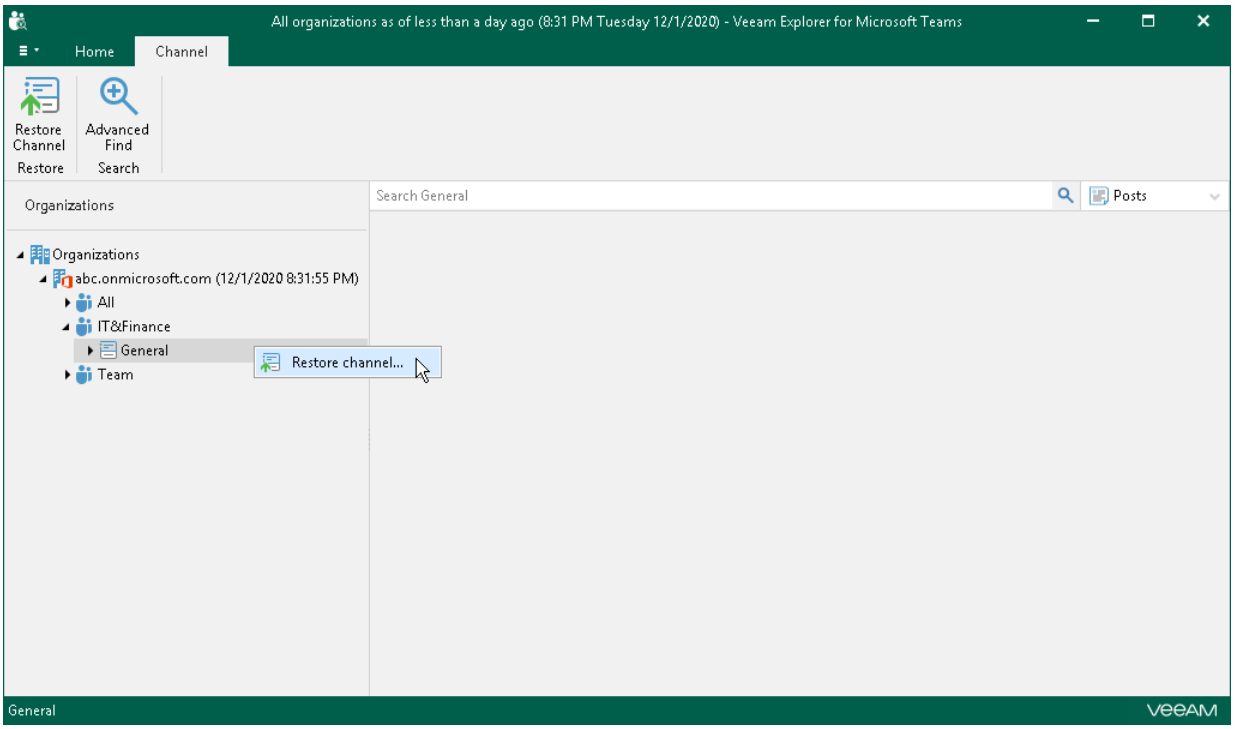

- If you want to restore all tabs of a channel, do either of the following:
  - In the navigation pane, expand a channel whose tabs you want to restore, select **Other tabs** and click **Restore Tabs** on the ribbon.
  - In the navigation pane, expand a channel whose tabs you want to restore, right-click **Other tabs** and select **Restore tabs**.

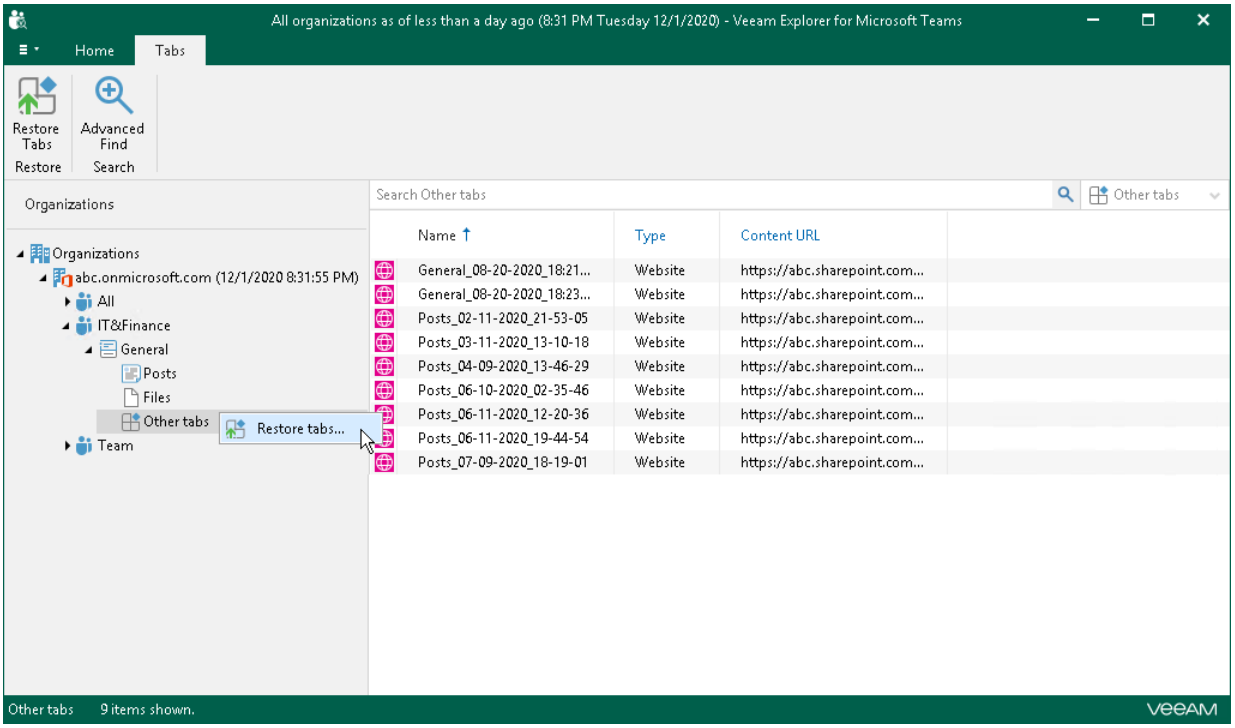

- If you want to restore individual tabs of a channel, do the following:
  a. In the navigation pane, expand a channel whose tabs you want to restore and select **Other tabs**.
  b. In the preview pane, select a necessary tab and click **Restore Tabs** on the ribbon or right-click a tab and select **Restore tabs**.

TIP

To select more than one tab in the preview pane, press and hold the **[CTRL]** key and select the necessary tabs.

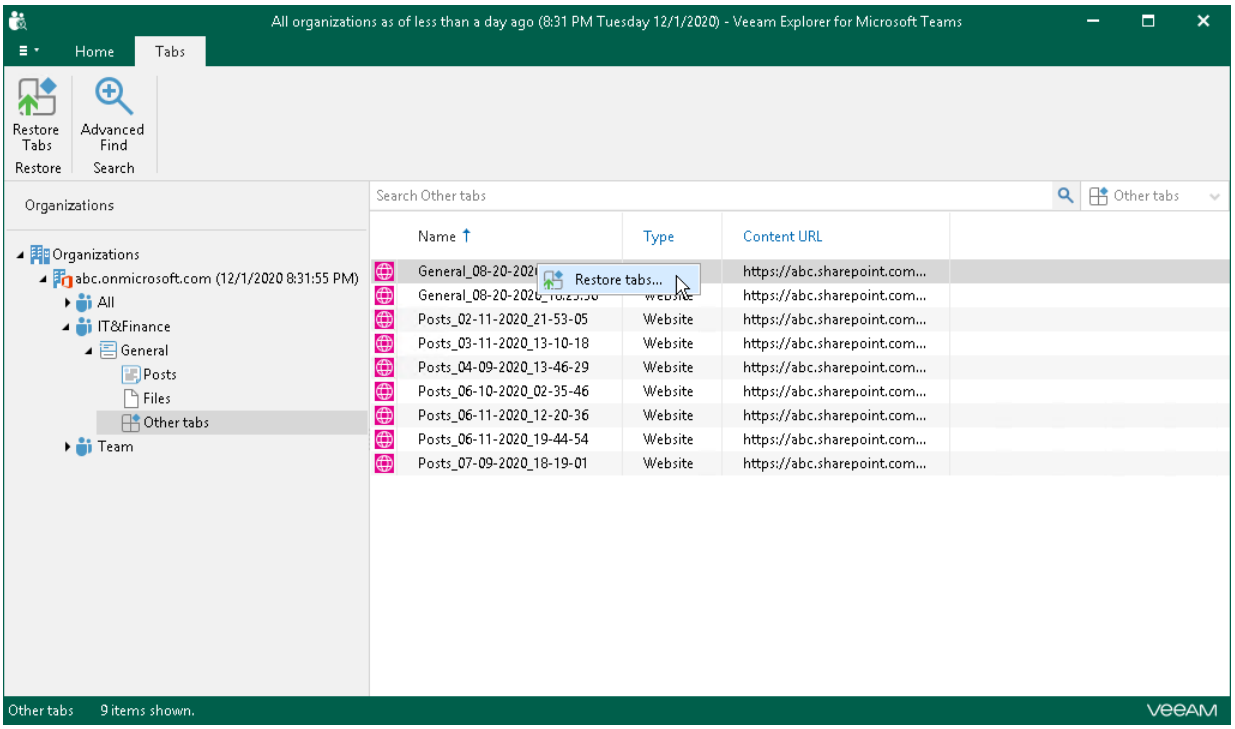

Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication and specify authentication settings.

- Modern authentication
- Basic authentication

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for data restore. Such an application is used to restore the specified team back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for data restore.

   By default, Veeam Explorer for Microsoft Teams populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

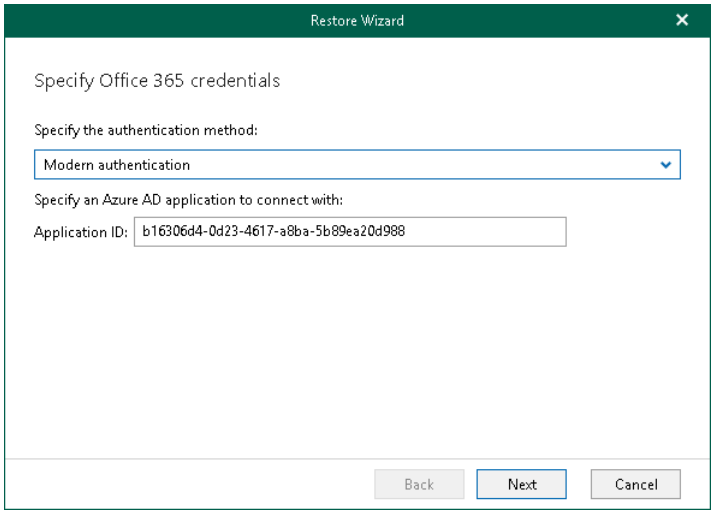

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. In the **Username** and **Password** fields, enter credentials to connect to the Microsoft Office 365 organization.

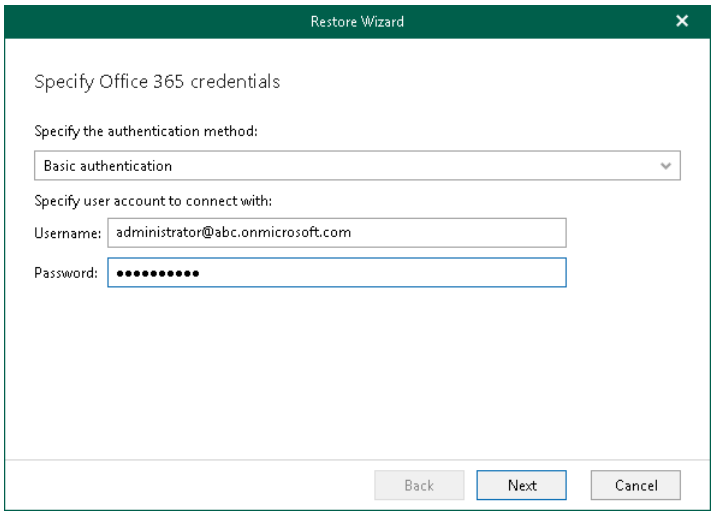

Step 3. Log In to Office 365

This step is available if you selected the *Modern authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

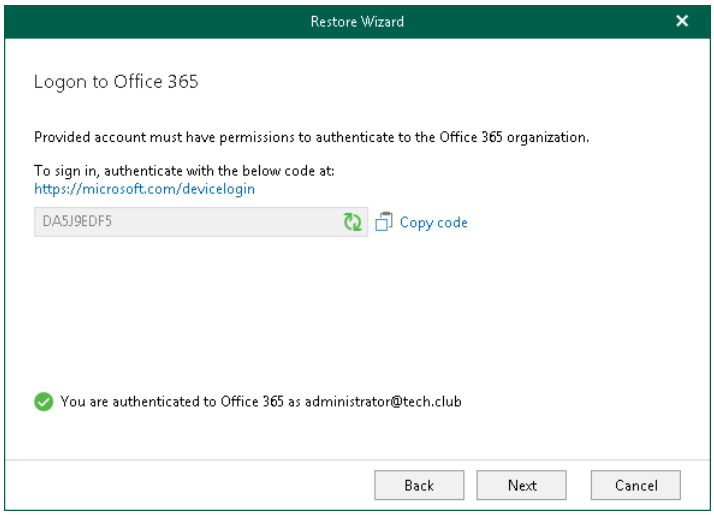

Step 4. Specify Restore Options

At this step of the wizard, specify data restore options and start the restore process:

1. Select the **Changed items** check box if you want to restore items that have changed since the time when the backup was created. When you select this option, Veeam Explorer for Microsoft Teams overwrites existing items in your target location.
2. Select the **Missing items** check box if you want to restore items that are missing in the target location. For example, some of the items were removed and you want to restore them from the backup.
3. Click **Restore**.

> **NOTE**
>
> During restore of a tab, Veeam Explorer for Microsoft Teams restores the name of the tab and a link to an object published on the tab. The object itself, for example, a SharePoint site or document, is not restored from the backup in this scenario. To restore such objects along with tabs, consider restoring an entire team. For more information, see Restoring Teams.

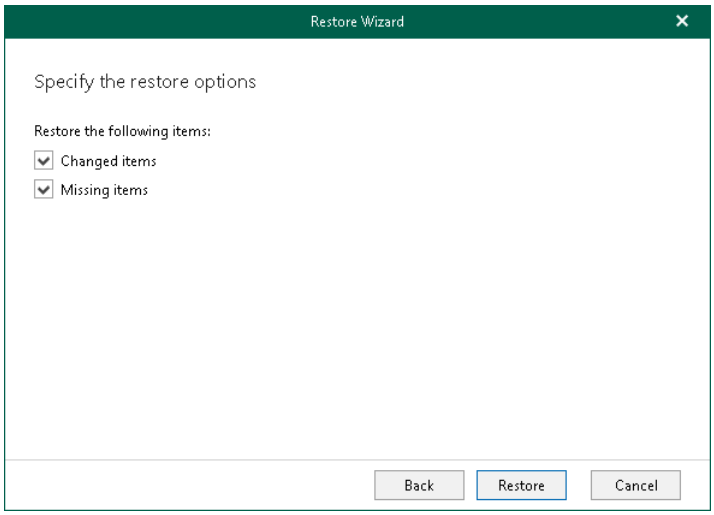

Restoring Posts

You can use Veeam Explorer for Microsoft Teams to restore posts of a team channel. When you restore posts, you can choose to restore all posts of a channel or restore posts published within a specified time period.

> **NOTE**
>
> Veeam Backup for Microsoft Office 365 does not restore posts to their original location in the team channel. Instead, Veeam Backup for Microsoft Office 365 creates a new tab in the original channel and restores posts to this tab.
>
> If you want to restore posts to another channel, use the export operation. For more information, see Data Export.

To restore posts, do the following:

1. Launch the Restore wizard.
2. Select an authentication method.
3. Log in to Office 365.
4. Specify restore options.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do either of the following:

- In the navigation pane, expand a channel whose posts you want to restore, select **Posts** and click **Restore Posts** on the ribbon.
- In the navigation pane, expand a channel whose posts you want to restore, right-click **Posts** and select **Restore posts.**

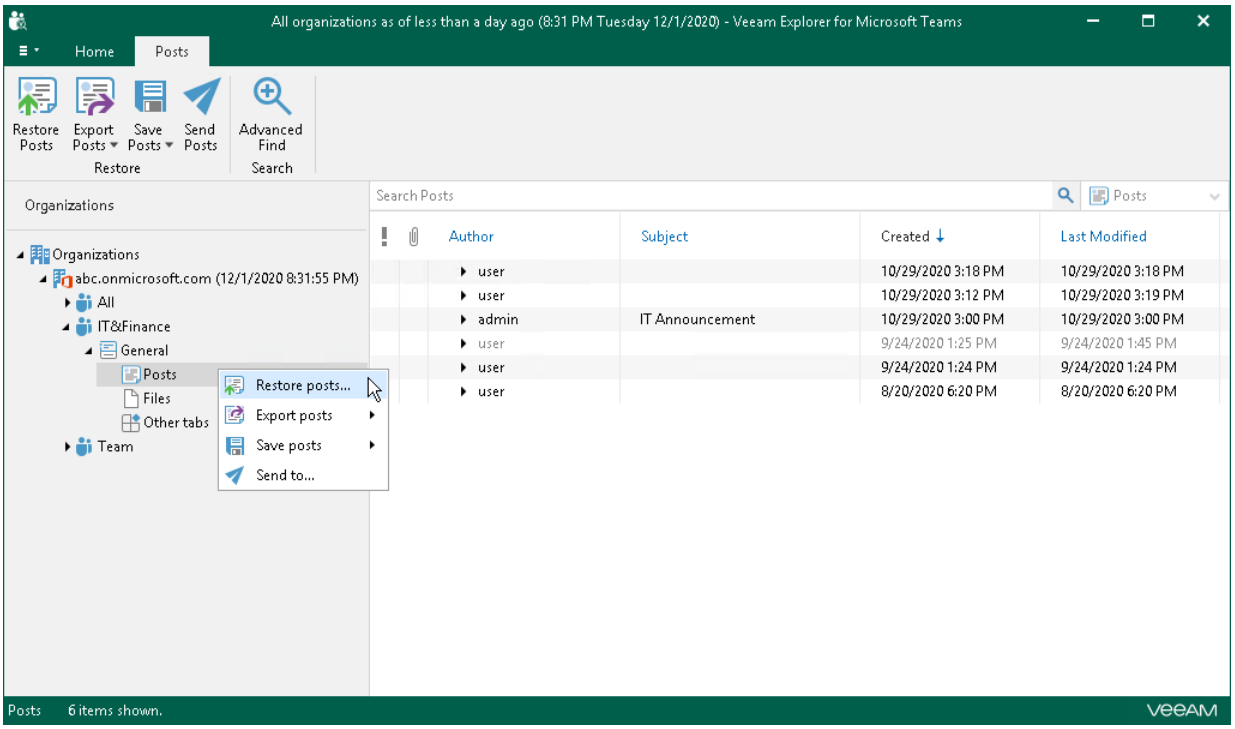

Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication and specify authentication settings.

- Modern authentication
- Basic authentication

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for data restore. Such an application is used to restore the specified team back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for data restore.

   By default, Veeam Explorer for Microsoft Teams populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

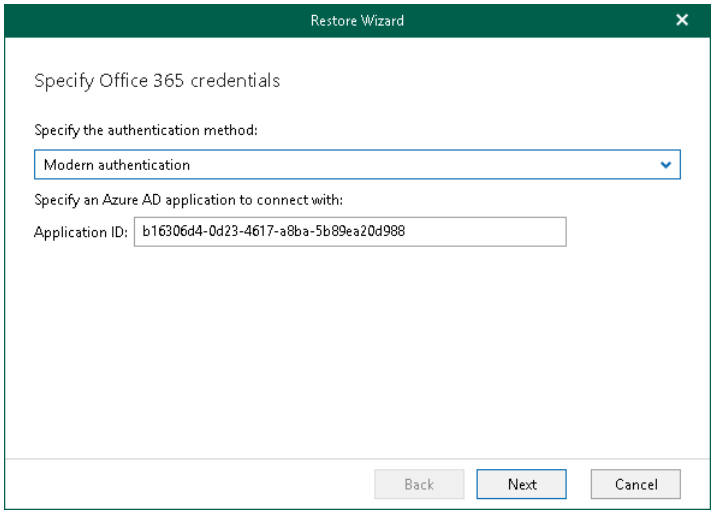

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. In the **Username** and **Password** fields, enter credentials to connect to the Microsoft Office 365 organization.

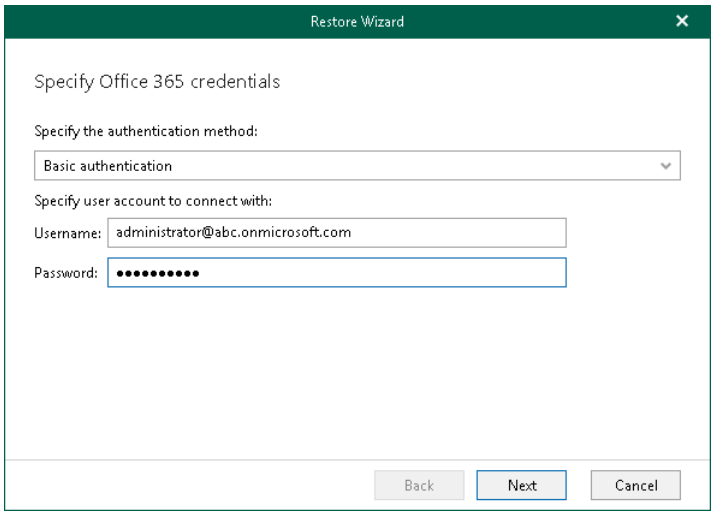

Step 3. Log In to Office 365

This step is available if you selected the *Modern authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

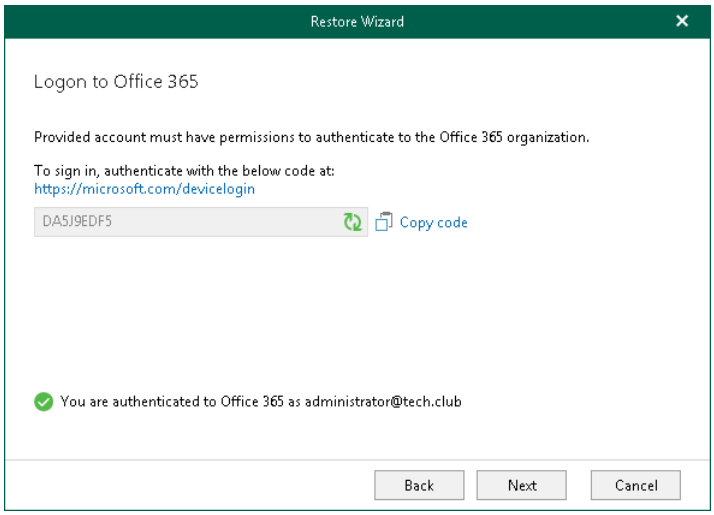

Step 4. Specify Restore Options

At this step of the wizard, select restore options and click **Restore**.

You can select one of the following options:

- **Restore all posts**

  Select this option if you want to restore all posts published in the channel.

- **Restore posts for the specified time period**

  Select this check box if you want to restore posts published in the channel during the specified time period.

  If you select this option, in the **From** and **To** fields, specify the dates that define the start and end of the desired time period. Veeam Explorer for Microsoft Teams will restore posts whose last modification date belongs to the specified period.

When you restore posts, Veeam Backup for Microsoft Office 365 will create a separate tab in the original team channel and restore posts to this tab.

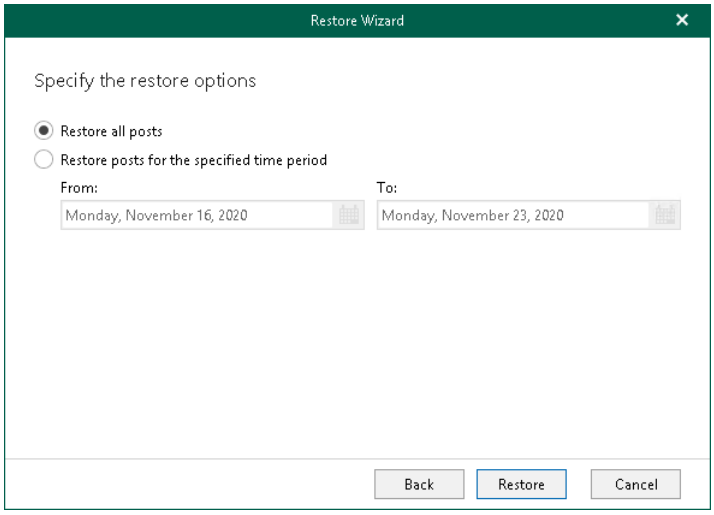

Restoring Files

You can use Veeam Explorer for Microsoft Teams to restore files of a team channel.

To restore files, do the following:

1. Launch the Restore wizard.
2. Select an authentication method.
3. Log in to Office 365.
4. Specify restore options.

Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do the following:

- If you want to restore all files of a channel, do either of the following:
  - In the navigation pane, expand a channel whose files you want to restore and click **Restore File** on the ribbon.
  - In the navigation pane, right-click a channel whose files you want to restore and select **Restore files**.

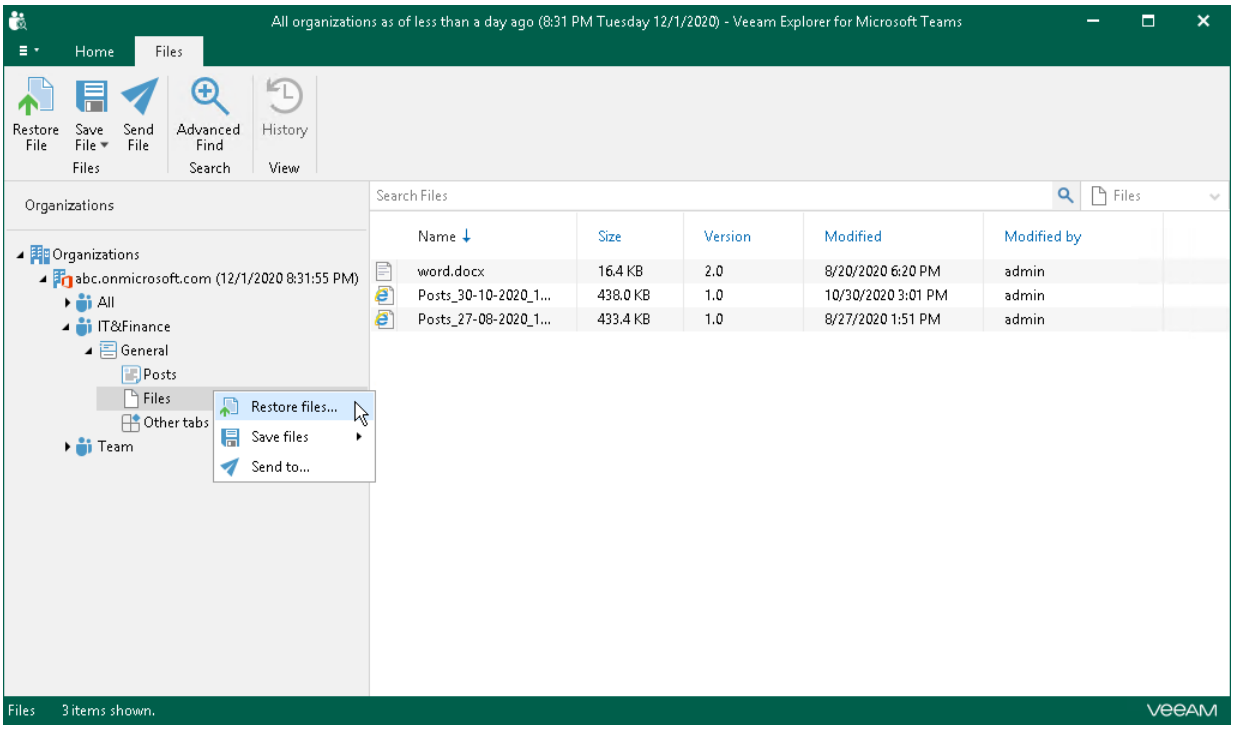

- If you want to restore individual files, do the following:
  a. In the navigation pane, expand a channel whose files you want to restore and select **Files**.
  b. In the preview pane, select a necessary file and click **Restore File** on the ribbon or right-click a file and select **Restore file**.

TIP

To select more than one file in the preview pane, press and hold the **[CTRL]** key and select the necessary files.

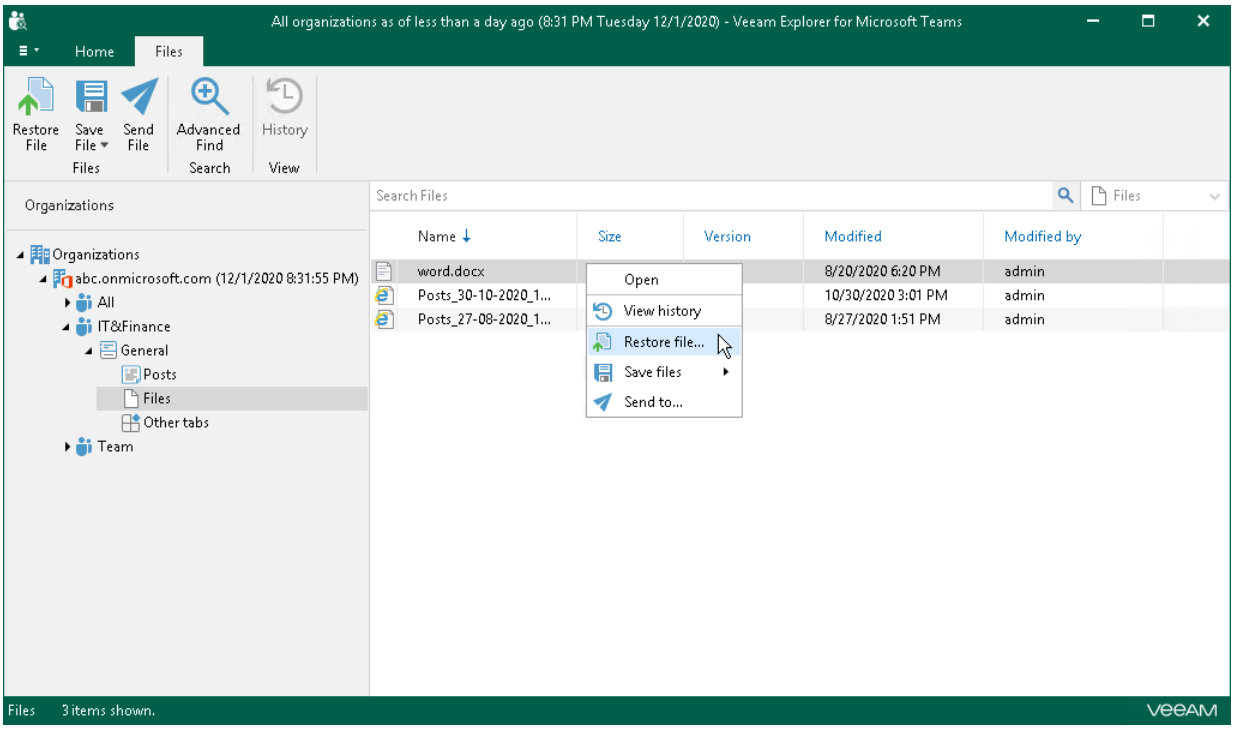

Step 2. Select Authentication Method

At this step of the wizard, choose whether you want to use modern or basic authentication and specify authentication settings.

- Modern authentication
- Basic authentication

Modern Authentication

To use modern authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Modern authentication*.

   This will allow Veeam Backup for Microsoft Office 365 to use an Azure AD application for data restore. Such an application is used to restore the specified team back to Office 365 organizations with enabled multi-factor authentication (MFA) and enabled or disabled security defaults.

   For more information, see the Microsoft Office 365 Organizations section of the Veeam Backup for Microsoft Office 365 User Guide.

2. In the **Application ID** field, enter an identification number of the Azure AD application that you want to use for data restore.

   By default, Veeam Explorer for Microsoft Teams populates this field with the identification number of the application that was used during a backup session. If you want to use another application, make sure to grant this application required permissions, as described in the Azure AD Application Permissions section of the Veeam Backup for Microsoft Office 365 User Guide.

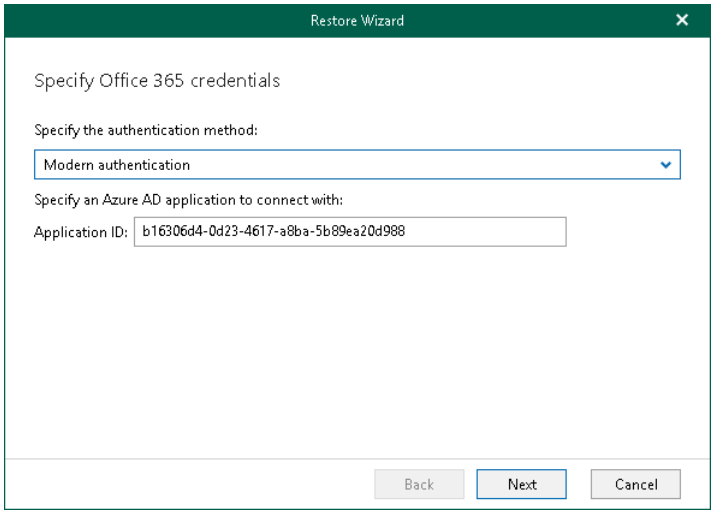

Basic Authentication

To use basic authentication, do the following:

1. From the **Specify the authentication method** drop-down list, select *Basic authentication*.
2. In the **Username** and **Password** fields, enter credentials to connect to the Microsoft Office 365 organization.

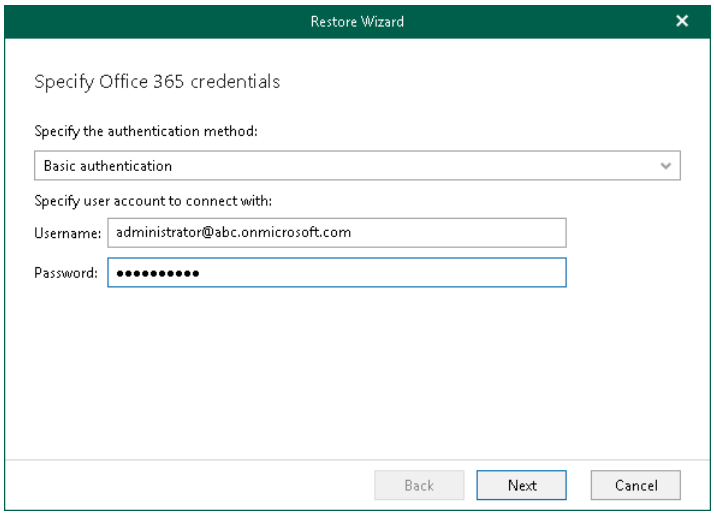

Step 3. Log In to Office 365

This step is available if you selected the *Modern authentication* option at the previous step of the wizard.

At this step of the wizard, sign in to your Microsoft Office 365 organization.

To sign in to the Microsoft Office 365 organization, do the following:

1. Click **Copy code** to copy an authentication code.

   Mind that a code is valid for 15 minutes. You can click **Refresh** to request a new code from Microsoft.

2. Click the link to the Microsoft portal.

3. On the **Microsoft Azure device logon** page, paste the code that you have copied and sign in to Azure.

   Make sure to log in with the user account that has required permissions assigned. For more information, see Required Permissions.

   Make sure that the required settings are specified for the Azure AD application used for restore. For more information, see the Required Azure AD Application Settings section of the Veeam Backup for Microsoft Office 365 User Guide.

4. Return to the **Restore Wizard** window and click **Next**.

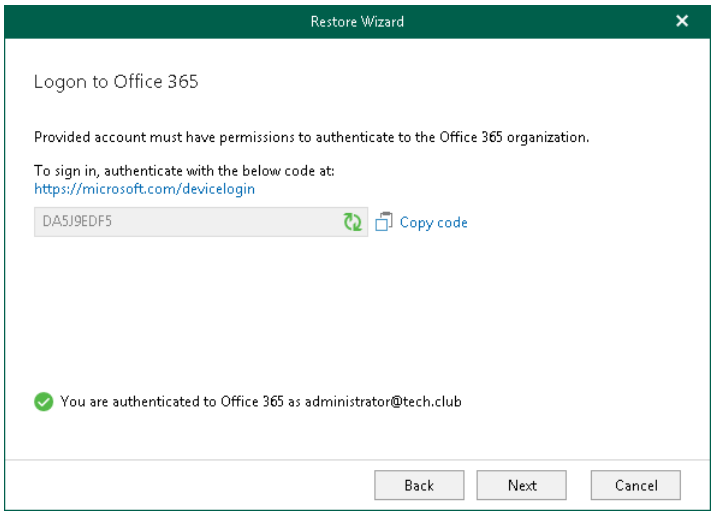

Step 4. Specify Restore Options

At this step of the wizard, specify data restore options and start the restore process:

1. Select the **Changed items** check box if you want to restore items that have changed since the time when the backup was created.
2. Select the **Missing items** check box if you want to restore items that are missing in the target location. For example, some of the items were removed and you want to restore them from the backup.
3. By default, Veeam Explorer for Microsoft Teams restores all versions of backed-up files from the backup. If you want to restore only the latest version of a file, select the **Restore only the latest version** check box. If you select this option, Veeam Explorer for Microsoft Teams will merge original files in the production environment with the latest version of the files in the backup.
4. Click **Restore**.

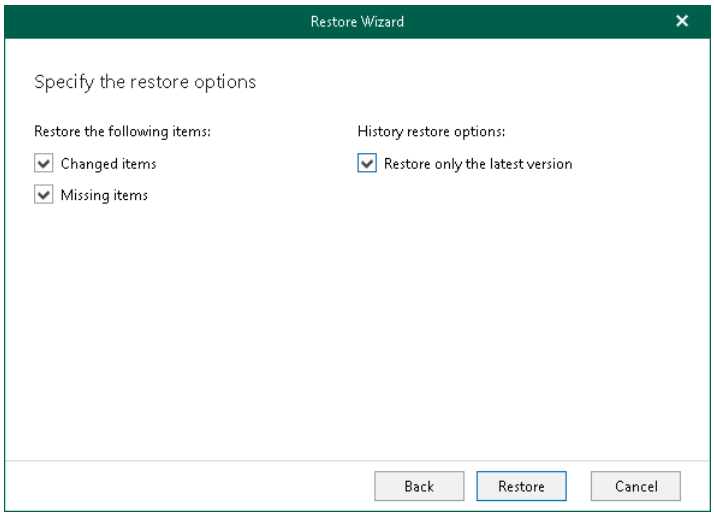

Data Export

You can use Veeam Explorer for Microsoft Teams to export posts published in a team channel to a file of the HTML format. You can choose to export all posts of a channel, export posts created during a specified time period or export individual posts.

The export operation can be useful, for example, if you want to restore posts of a team channel to another channel. To do this, you can export posts to an HTML file, publish this file in Microsoft SharePoint and publish a link to this file in a custom tab of the necessary channel.

> **TIP**
>
> Before exporting data, make sure to read Considerations and Limitations.

Exporting Multiple Posts

To export posts of a channel, do the following:

1. In the navigation pane, expand a channel whose posts you want to export.
2. In the navigation pane, select **Posts**, click **Export Posts** on the ribbon and select how you want to export posts:
   - **Export to Desktop**. Select this option if you want to perform the 1-Click Export operation. Veeam Explorer for Microsoft Teams will immediately export all posts of the channel to the *Desktop* folder.
   - **Export to**. Select this option if you want to specify what posts to export and where to export posts.

   Alternatively, in the navigation pane, right-click **Posts** and select **Export posts** > **Export to Desktop** or **Export posts** > **Export to**.

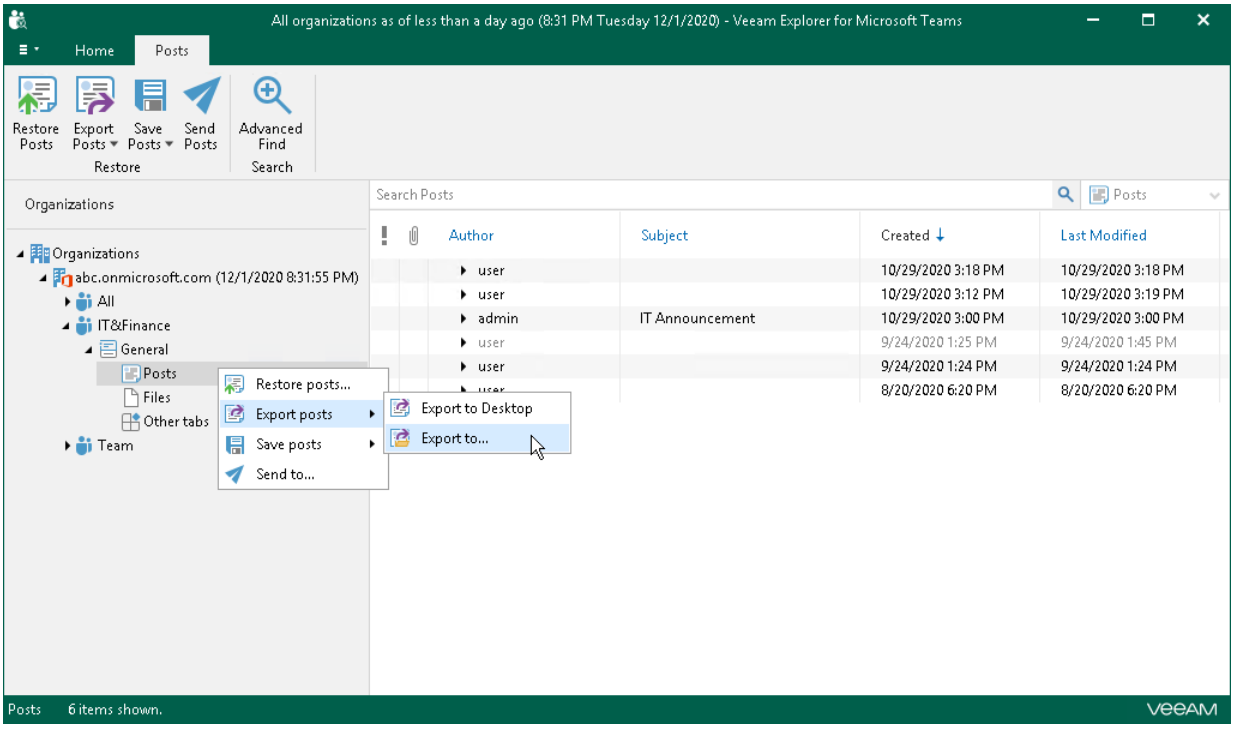

3. If you selected the **Export to** option at the step 2, in the **Export** wizard, do the following:
   a. Specify what posts you want to export. You can select from the following options:
      - **Export all posts**. If you select this option, Veeam Explorer for Microsoft Teams will export all posts published in the channel to an HTML file.
      - **Export posts for the specified time period**. If you select this option, you can define a time period when posts that you want to export were created. To do this, specify the start date and end date of the period in the **From** and **To** fields. Veeam Explorer for Microsoft Teams will export posts whose last modification date belongs to the specified period.

b. In the **Save as** field, click **Browse** and browse to a directory where you want to save an HTML file with exported posts. By default, Veeam Explorer for Microsoft Teams exports posts to a file with the following name: *<Channel_name>_<date_and_time>.html*. For example: *General_2020_08_31_04_47_34.html*.

c. Select the **Open file after export** check box if you want to open the resulting HTML file right after the export operation is completed.

d. Click **Export**.

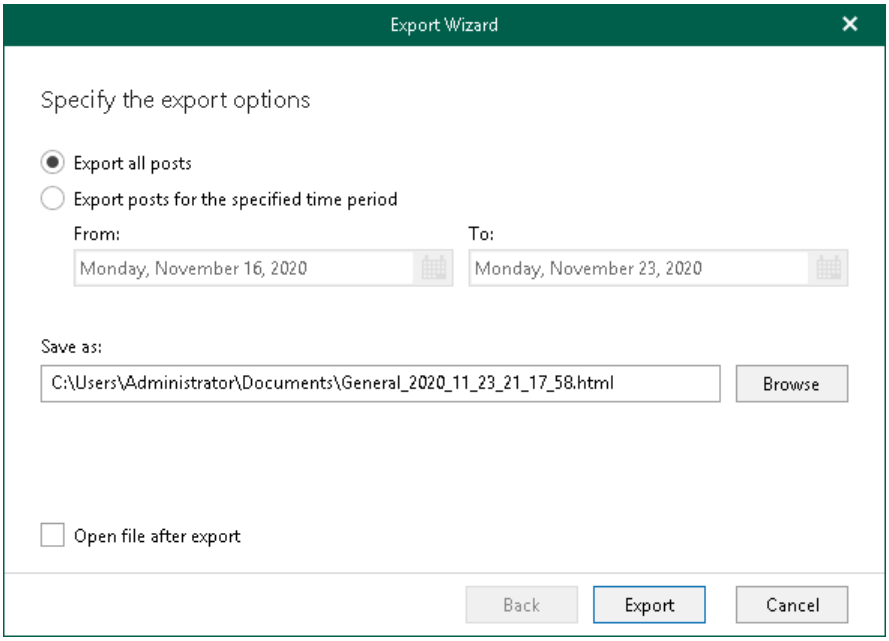

Exporting Individual Posts

To export a specific post of a channel, do the following:

1. In the navigation pane, expand a channel whose post you want to export and click **Posts**.
2. In the preview pane, select a necessary post, click **Export Posts** on the ribbon and select how you want to export the post:
   - **Export to Desktop**. Select this option if you want to perform the 1-Click Export operation. Veeam Explorer for Microsoft Teams will immediately export the post to the *Desktop* folder.
   - **Export to**. Select this option if you want to export the post to a custom location.

   Alternatively, in the preview pane, right-click a necessary post and select **Export to Desktop** or **Export to**.

> **TIP**
>
> To select more than one post in the preview pane, press and hold the **[CTRL]** key and select the necessary posts.

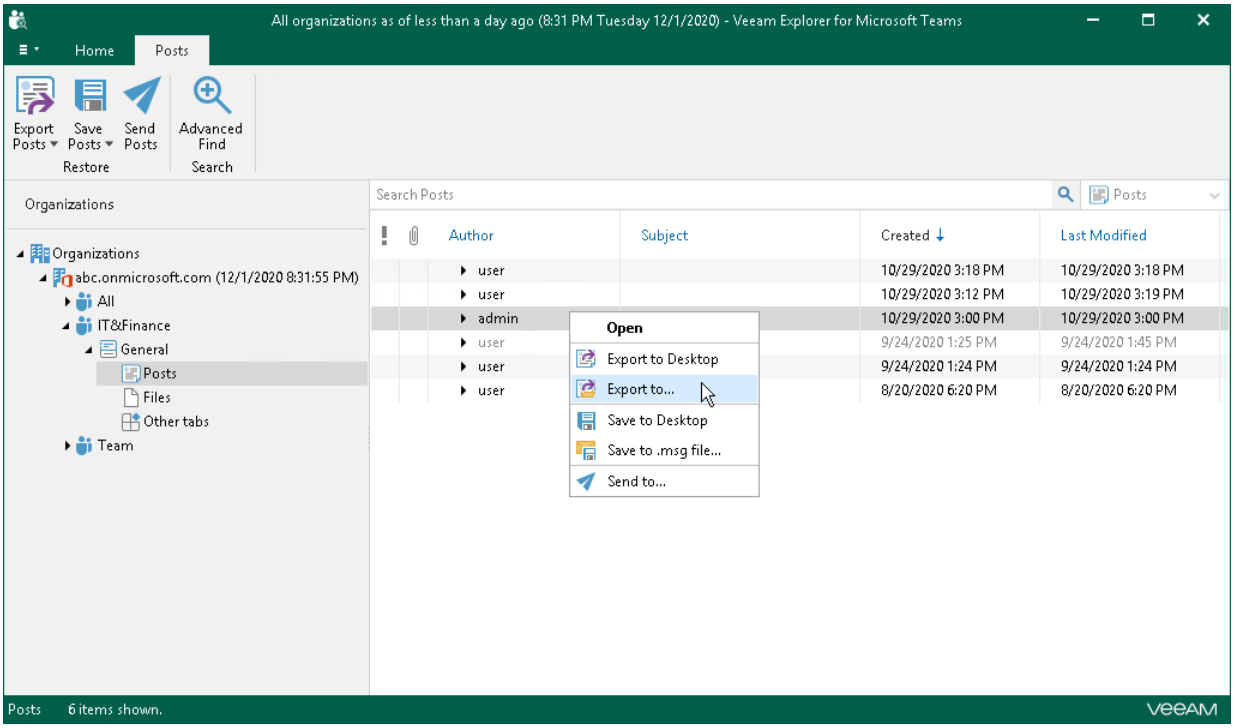

3. If you selected the **Export to** option at the step 2, in the **Export** wizard, do the following:

   a. In the **Save as** field, click **Browse** and browse to a directory where you want to save an HTML file with exported posts. By default, Veeam Explorer for Microsoft Teams exports posts to a file with the following name: *<Channel_name>_<date_and_time>.html*. For example: *General_2020_08_31_04_47_34.html*.

   b. Select the **Open file after export** check box if you want to open the resulting HTML file right after the export operation is completed.

   c. Click **Export**.

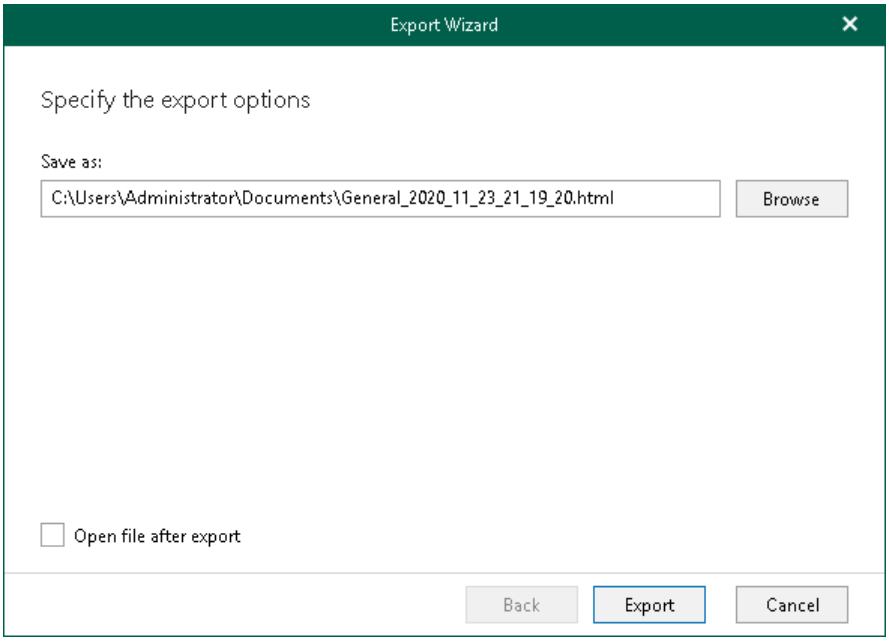

Saving Posts

You can use Veeam Explorer for Microsoft Teams to save posts published in a team channel as *Microsoft Exchange Mail Document* (*.msg*) files. You can save all posts of a channel or specific posts. Veeam Explorer for Microsoft Teams saves each post to a separate *.msg* file in the specified location.

- Saving All Posts of a Channel
- Saving Individual Posts

Saving All Posts of a Channel

To save posts published in a team channel, do the following:

1. In the navigation pane, expand a channel whose posts you want to save.
2. In the navigation pane, select **Posts**, click **Save Posts** on the ribbon and select where you want to save posts:
   - **Save to <folder_name>**. Select this option if you want to perform the 1-Click Save operation. Veeam Explorer for Microsoft Teams will save posts to the latest location that was used when saving posts. If you have not saved posts yet, Veeam Explorer for Microsoft Teams save posts to the *Desktop* folder.
   - **Save to .msg file**. Select this option if you want to save posts to a custom location. To specify a location, in the displayed window, browse to the necessary folder and click **Select Folder**.

   Alternatively, in the navigation pane, right-click **Posts** and select **Save posts** > **Save to <folder_name>** or **Save posts** > **Save to .msg file**.

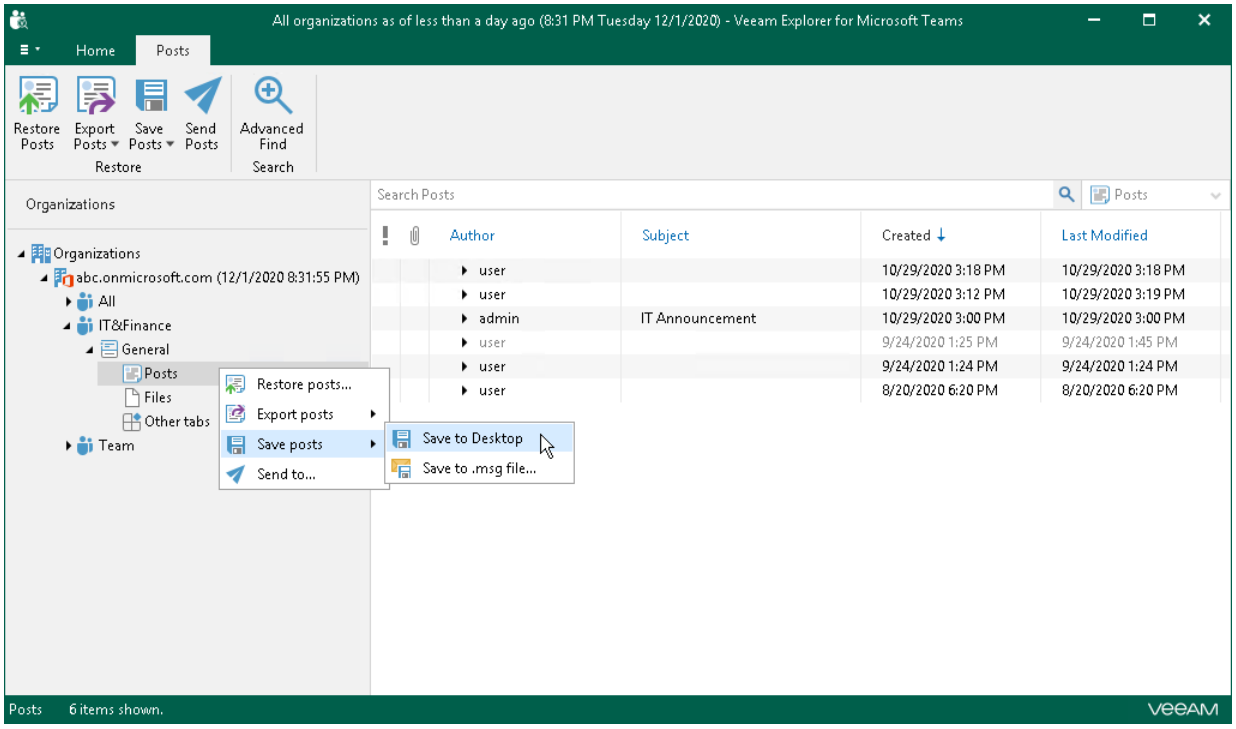

Saving Individual Posts

To save a specific post published in a team channel, do the following:

1. In the navigation pane, expand a channel whose post you want to save and select **Posts**.
2. In the preview pane, select a necessary post, click **Save Posts** on the ribbon and select where you want to save posts:
   - **Save to <folder_name>**. Select this option if you want to perform the 1-Click Save operation. Veeam Explorer for Microsoft Teams will save posts to the latest location that was used when saving posts. If you have not saved posts yet, Veeam Explorer for Microsoft Teams save posts to the *Desktop* folder.
   - **Save to .msg file**. Select this option if you want to save posts to a custom location. To specify a location, in the displayed window, browse to the necessary folder and click **Select Folder**.

   Alternatively, in the preview pane, right-click a necessary post and select **Save to <folder_name>** or **Save to .msg file**.

> **TIP**
>
> To select more than one post in the preview pane, press and hold the **[CTRL]** key and select the necessary posts.

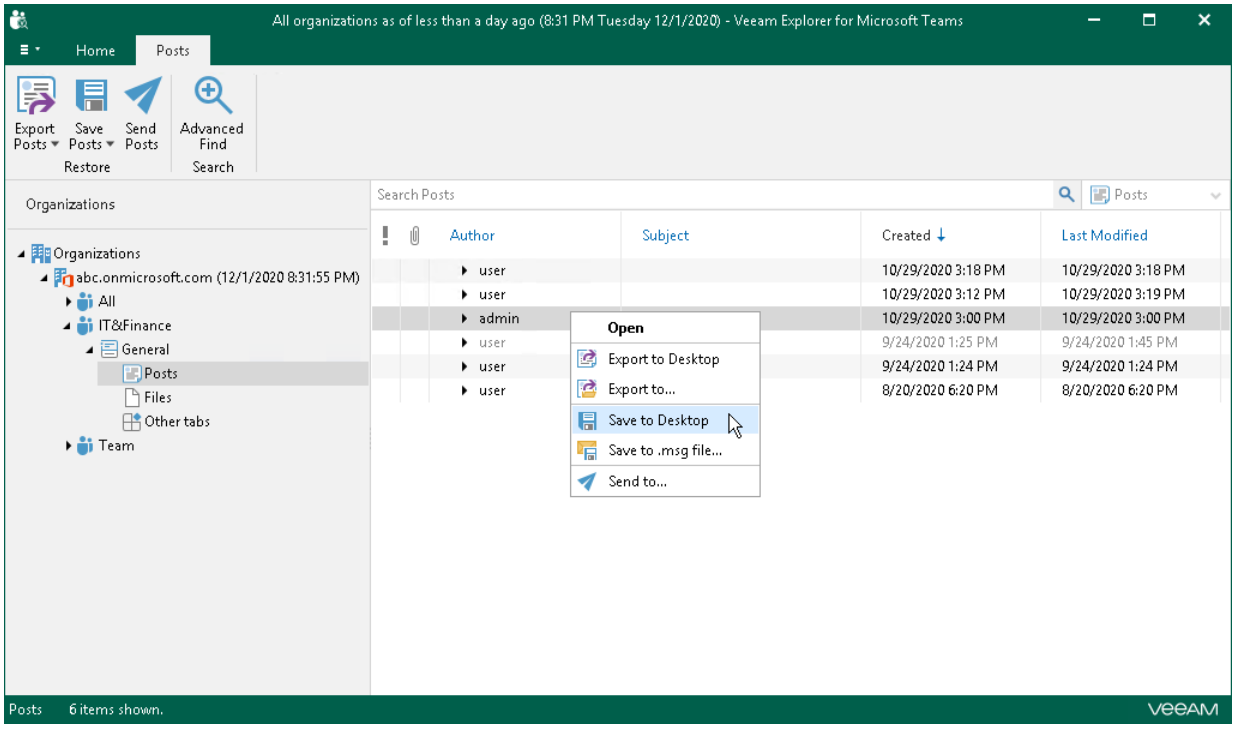

Saving Files

You can use Veeam Explorer for Microsoft Teams to save files published in a team channel to a specified location. You can save all files of a channel or specific posts.

- Saving All Files of a Channel
- Saving Individual Files

Saving All Files of a Channel

To save files published in a team channel, do the following:

1. In the navigation pane, expand a channel whose posts you want to save.
2. In the navigation pane, select **Files**, click **Save File** on the ribbon and select how you want to save files:
   - **Save as files**. Select this option if you want to save each file separately to a specified location. To specify a location, in the displayed window, browse to the necessary folder and click **Select Folder**.
   - **Save as ZIP**. Select this option if you want to add saved files to a ZIP archive and save the resulting archive to a specified location. To specify a location, in the displayed window, browse to the necessary folder and click **Select Folder**. Veeam Explorer for Microsoft Teams will save files to an archive with the *archive.zip* name.

   Alternatively, in the navigation pane, right-click **Files** and select **Save files** > **Save as files** or **Save files** > **Save as ZIP**.

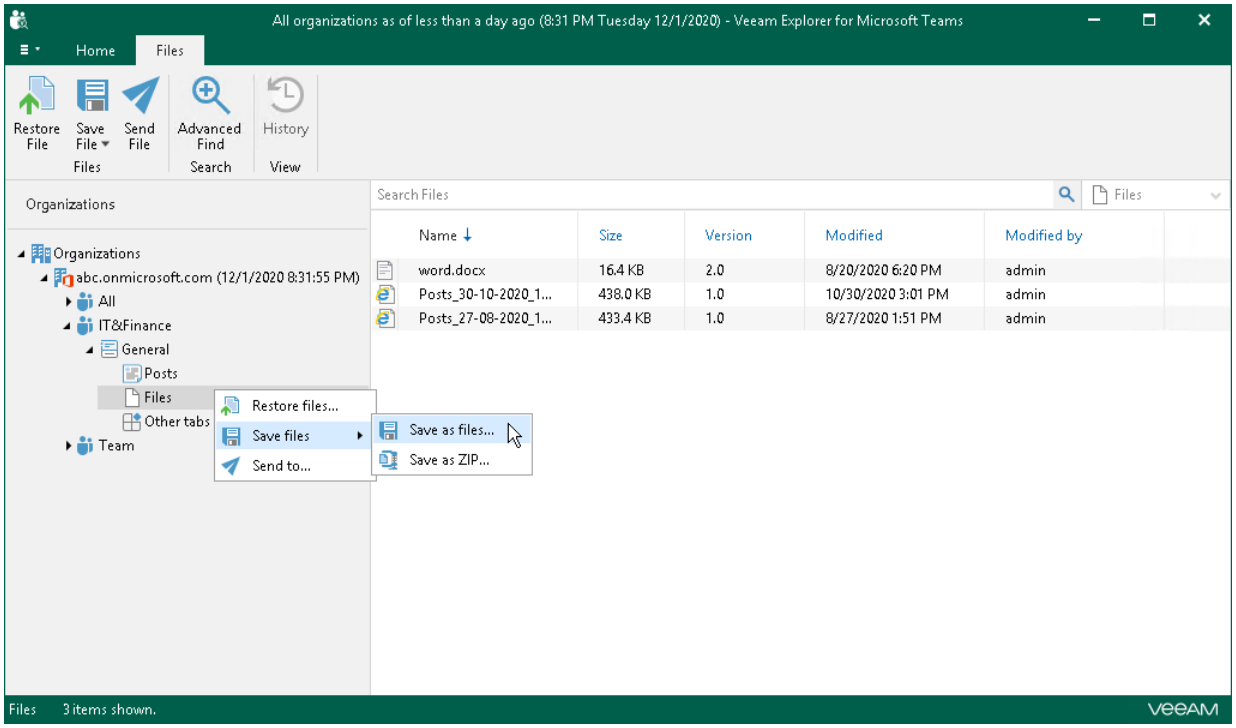

Saving Individual Files

To save a specific file published in a team channel, do the following:

1. In the navigation pane, expand a channel whose post you want to save and select **Files**.
2. In the preview pane, select a necessary file, click **Save File** on the ribbon and select how you want to save files:
   - **Save as files**. Select this option if you want to save each file separately to a specified location. To specify a location, in the displayed window, browse to the necessary folder and click **Select Folder**.
   - **Save as ZIP**. Select this option if you want to add saved files to a ZIP archive and save the resulting archive to a specified location. To specify a location, in the displayed window, browse to the necessary folder and click **Select Folder**. Veeam Explorer for Microsoft Teams will save files to an archive with the *archive.zip* name.

   Alternatively, in the preview pane, right-click a necessary file and select **Save file** > **Save as files** or **Save file** > **Save as ZIP**.

> **TIP**
>
> To select more than one file in the preview pane, press and hold the **[CTRL]** key and select the necessary files.

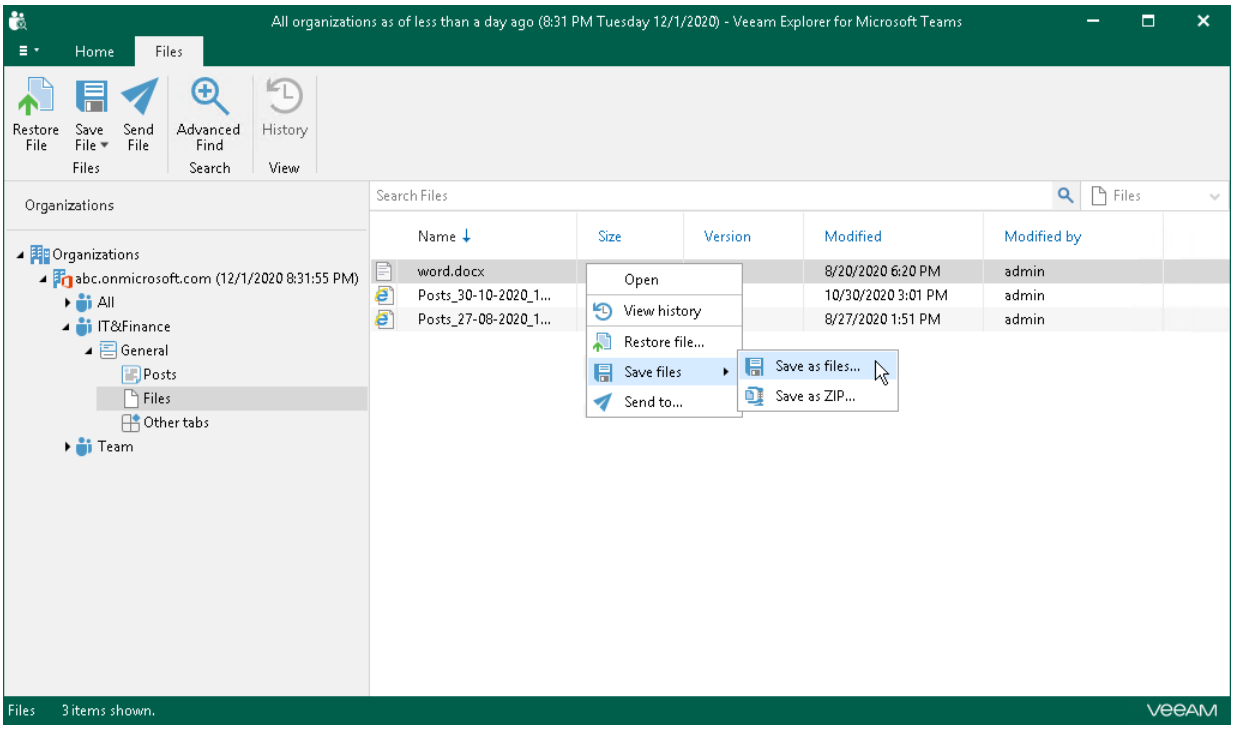

Sending Posts and Files

You can use Veeam Explorer for Microsoft Teams to send posts and files published in a team channel to specified recipients by email. You can send all posts or files of a channel or send individual posts or files.

- Sending All Posts or Files of a Channel
- Sending Individual Posts or Files

> **TIP**
>
> Before sending documents, make sure to configure SMTP settings, as described in SMTP Settings. The amount of data you can send at a time depends on your SMTP server configuration.

Sending All Posts or Files of a Channel

To send posts or files of a team channel by email, do the following:

1. In the navigation pane, expand a channel whose posts or files you want to send.
2. In the navigation pane, select **Posts** and click **Send Posts** on the ribbon or select **Files** and click **Send File** on the ribbon.

   Alternatively, in the navigation pane, right-click **Posts** and select **Send to** or right-click **Files** and select **Send to**.

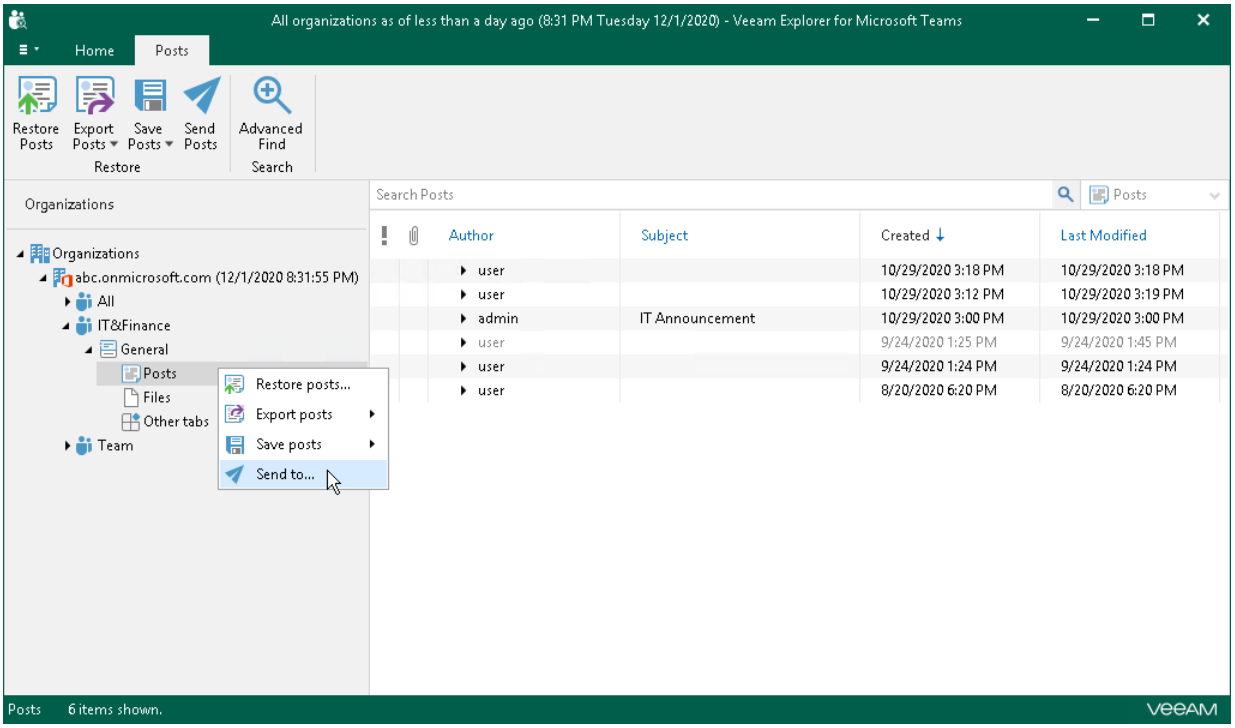

3. In case the size of the message with files exceeds the size allowed by your SMTP server configuration, Veeam Explorer for Microsoft Teams will prompt whether you want to send files. In the notification window, click **Yes**.
4. In the **Send Items** window, provide a recipient address.

   The **From** field is filled automatically based on the address you have provided when configuring SMTP settings.

You can also edit the message subject and body. To edit the message body, click **More details**.

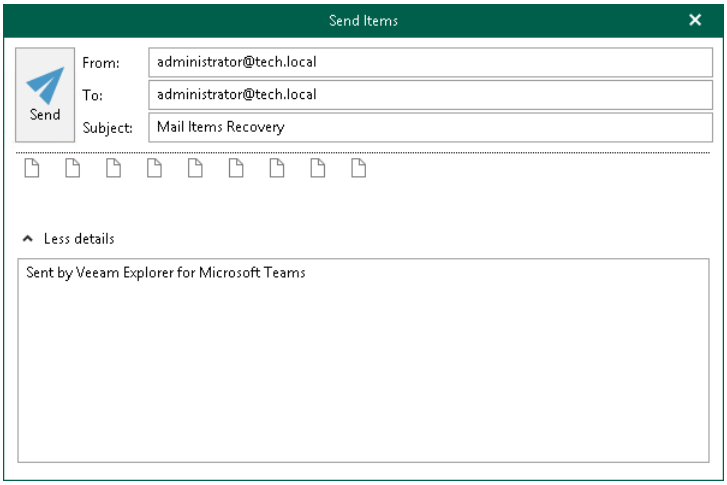

Sending Individual Posts or Files

To send a specific post or file of a team channel by email, do the following:

1. In the navigation pane, expand a channel whose post or file you want to send and select **Posts** or **Files**.
2. In the preview pane, select a necessary post or file and click **Send Posts** or **Send File** on the ribbon.

   Alternatively, in the preview pane, right-click a necessary post or file and select **Send to**.

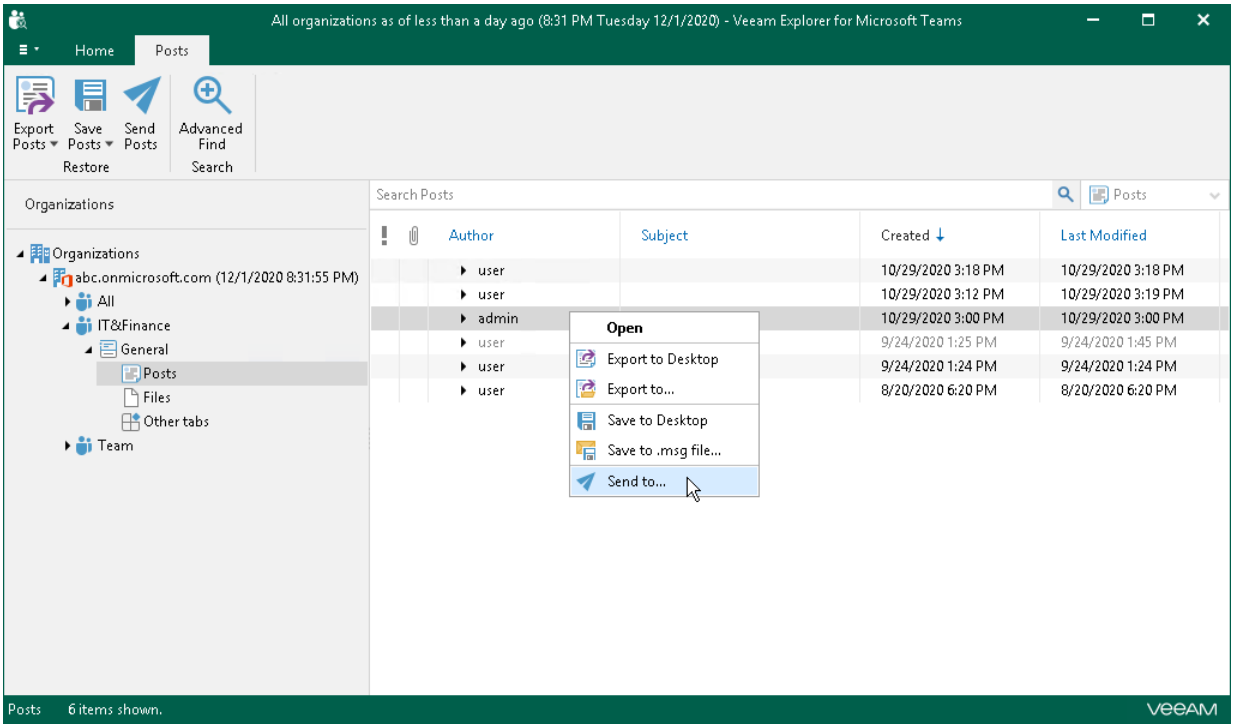

3. In the **Send Items** window, provide a recipient address.

   The **From** field is filled automatically based on the address you have provided when configuring SMTP settings.

   You can also edit the message subject and body. To edit the message body, click **More details**.

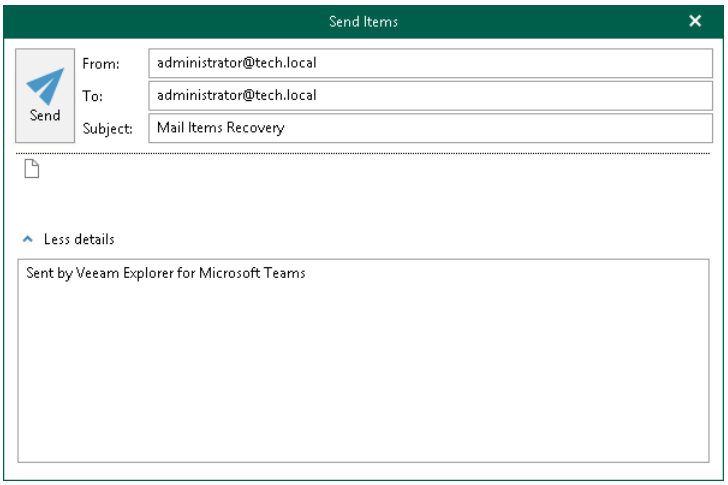

Veeam Plug-ins for Enterprise Applications

Version 11

User Guide

March, 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 7
ABOUT THIS DOCUMENT 8
VEEAM PLUG-INS OVERVIEW 9
VEEAM PLUG-IN FOR SAP HANA 10
How Veeam Plug-in for SAP HANA Works 11
Planning and Preparation 12
System Requirements 13
Required Permissions 14
Used Ports 15
Licensing 18
Environment Planning 20
Veeam Backup Repositories 26
Access and Encryption Settings on Repositories 30
Deployment and Configuration 32
Installing Plug-in for SAP HANA 33
Configuring Plug-in for SAP HANA 35
Automating Configuration of Plug-in for SAP HANA 40
Upgrading Plug-in for SAP HANA 41
Importing Backup Files 43
Upgrading Backup Files 45
Uninstalling Plug-in for SAP HANA 47
Database Protection 48
Database Backup (HDBSQL Scripts) 49
Database Backup (SAP HANA Studio) 51
Database Backup (SAP HANA Cockpit) 54
Backup Job in Veeam Backup & Replication Console 56
Database Recovery 58
Restoring Databases (HDBSQL Commands) 59
Restoring Databases (SAP HANA Studio) 60
Restoring SYSTEMDB (SAP HANA Cockpit) 67
Recovering Tenant Databases with SAP HANA Cockpit 72
Recovering Databases to Other Servers (System Copy) 76
Restore from Backup Copy 78
Retention of SAP HANA Backups 79
Deleting Backups Using SAP HANA Tools 80
Deleting Backups Manually 81

Configuring Force Deletion of Backups 82
Removing Backups from Configuration 83
Backup Copy for SAP HANA Backups 84
Creating Backup Copy Job 85
Converting Backup Copy to Backup 96
Logs and Support 97
VEEAM PLUG-IN FOR ORACLE RMAN 98
How Veeam Plug-in for Oracle RMAN Works 99
Planning and Preparation 101
System Requirements 102
Required Permissions 104
Used Ports 106
Licensing 109
Oracle Environment Planning 111
Veeam Environment Planning 115
Veeam Backup Repositories 117
Access and Encryption Settings on Repositories 120
Deployment and Configuration 122
Installing Veeam Plug-in for Oracle RMAN 123
Configuring Veeam Plug-in for Oracle RMAN 129
Upgrading Veeam Plug-in for Oracle RMAN 138
Importing/Exporting Plug-in Settings 141
Importing Backup Files 142
Upgrading Backup Files 145
Uninstalling Veeam Plug-in for Oracle RMAN 147
Database Protection 148
Oracle RMAN Full Backup 149
Oracle RMAN Channel Allocation 150
Backup Job in Veeam Backup & Replication 151
Database Recovery 154
Restore to Original Server 155
Restore to Another Server 156
Restore from Backup Copy 158
Restore with Veeam Explorer for Oracle 160
Restore of Control File from Autobackup 161
Retention of RMAN Backups and Archived Logs 163
Configuring Retention Policy 164
Configuring Archived Redo Log Retention 165
Configuring Force Deletion of Backups 166

Deleting Backups Manually 167
Removing Backups from Configuration 168
Backup Copy for Oracle RMAN Backups 169
Creating Backup Copy Job 170
Converting Backup Copy to Backup 181
Logs and Support 183
VEEAM PLUG-IN FOR SAP ON ORACLE 184
How Veeam Plug-in for SAP on Oracle Works 185
Planning and Preparation 187
System Requirements 188
Required Permissions 189
Used Ports 190
Licensing 193
Environment Planning 195
Veeam Backup Repositories 197
Access and Encryption Settings on Repositories 200
Deployment and Configuration 202
Installing Veeam Plug-in for SAP on Oracle 203
Configuring Plug-in for SAP on Oracle 205
Configuring Parallelism for Redo Logs 209
Importing Backup Files 210
Uninstalling Plug-in for SAP on Oracle 212
Database Protection 213
Limitations and Considerations 214
SAP on Oracle Backup Using Backint 215
SAP on Oracle Backup Using RMAN_UTIL 218
Backup Job in Veeam Backup & Replication Console 221
Database Recovery 223
Oracle Databases Restore 224
Redo Logs Restore 225
Database Restore to Another Server (System Copy) 226
Restore from Backup Copy 228
Retention of SAP on Oracle Backups 229
Configuring Retention Policy for Backups 230
Deleting Backups Manually 231
Removing Backups from Configuration 232
Backup Copy for SAP on Oracle Backups 233
Creating Backup Copy Job 234
Converting Backup Copy to Backup 245

Logs and Support 246

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

The document describes how to deploy, configure and use Veeam Plug-in for SAP HANA, Veeam Plug-in for Oracle RMAN and Veeam Plug-in for SAP on Oracle.

Intended Audience

This document is intended for DBA's, backup administrators and other IT specialists who use Veeam to back up and restore SAP HANA and Oracle databases.

Veeam Plug-ins Overview

Veeam Plug-ins for Enterprise Applications extend the functionality of Veeam Backup & Replication and allow you to create transactionally-consistent backups of SAP HANA and Oracle databases.

- Veeam Plug-in for SAP HANA — an SAP-certified backup and recovery solution that allows you to back up and restore SAP HANA databases.
- Veeam Plug-in for Oracle RMAN — an Oracle-certified backup and recovery solution that allows you to back up and restore Oracle databases.
- Veeam Plug-in for SAP on Oracle — an SAP-certified backup and recovery solution that allows you to back up and restore Oracle databases to which an SAP application is connected.

IMPORTANT

Veeam Plug-ins store database and log backups in repositories added to the Veeam Backup & Replication infrastructure. Thus, to use Veeam Plug-ins, you must have a Veeam Backup & Replication server deployed in your infrastructure. For instructions on how to deploy Veeam Backup & Replication, see the Deployment section of the Veeam Backup & Replication User Guide.

Veeam Plug-in for SAP HANA

Veeam Plug-in for SAP HANA is an SAP-certified backup tool that integrates with SAP backint and allows you to store transactionally-consistent SAP HANA database backups and logs in repositories connected to Veeam Backup & Replication.

> **NOTE**
>
> If you want to protect the SAP HANA server itself, you can use the image-level and file-level backup functionality of Veeam Backup & Replication or Veeam Agent for Linux. Note that image- and file-level backups of SAP HANA servers do not guarantee transaction-consistency of database backups.

How Veeam Plug-in for SAP HANA Works

Veeam Plug-in acts as an agent between an SAP HANA server and Veeam backup repositories. The plug-in integrated with databases through the SAP HANA Backint component. Backint for SAP HANA is an API that enables Veem Plug-in to directly connect to the SAP HANA database and send the database backup files to Veeam repositories.

Veeam Plug-in compresses, deduplicates database backups and transfers them to a backup repository connected to the Veeam Backup & Replication infrastructure. After you install and configure Veeam Plug-in on the SAP HANA server, you can perform all backup and restore operations with HDBSQL scripts and with native SAP HANA tools, such as SAP HANA Studio and SAP HANA Cockpit.

When Veeam Plug-in is configured, SAP Backint performs a database backup in the following way:

1. When you start a database backup, the SAP HANA Backint starts Veeam Plug-in services on the SAP HANA server.
2. Veeam Plug-in connects to the Veeam Backup & Replication server and creates a backup job (if it has not been created before). In the Veeam Backup & Replication console, Veeam backup administrators can use the backup job to monitor SAP HANA backups.
3. Veeam Plug-in starts Veeam Data Mover services on the SAP HANA server and on a backup repository. According to a specified number of parallel backint channels, Veeam Data Movers create channels to transfer backup data.
4. Veeam Data Movers transport backup data to the backup repository.

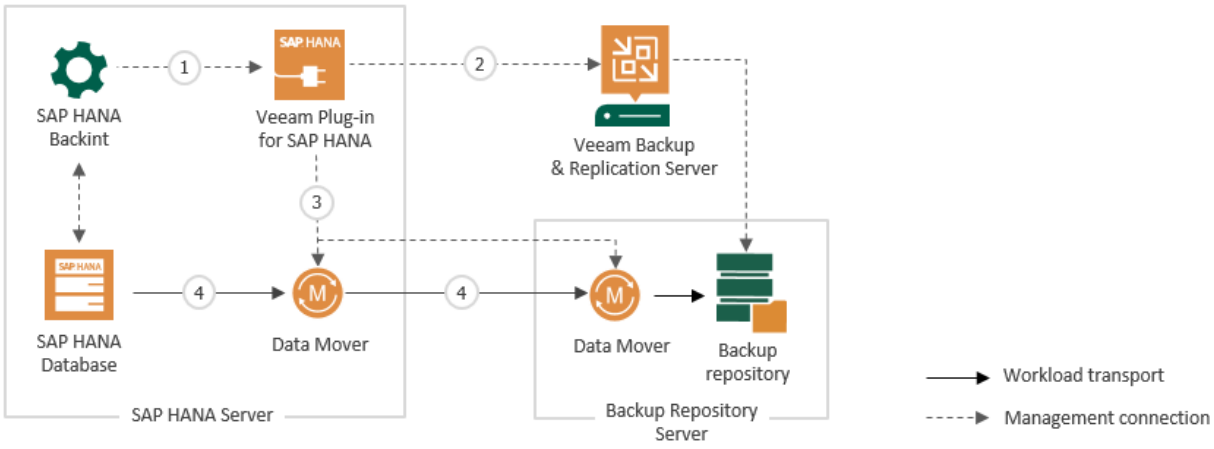

Planning and Preparation

Before you start to use Veeam Plug-in for SAP HANA, read the environment planning recommendations and make sure that your environment meets system requirements.

- System Requirements
- Required Permissions
- Used Ports
- Licensing
- Environment Planning
- Veeam Backup Repositories
- Access and Encryption Settings on Repositories

System Requirements

Before you start using Veeam Plug-in for SAP HANA, make sure the following requirements are met.

Supported OSes

Veeam Plug-in for SAP HANA is supported for the following OSes:

- **SLES for SAP Applications 15** (x86_64): GA, SP1, SP2, SP3.
- **SLES for SAP Applications 12** (x86_64): GA, SP1, SP2, SP3, SP4, SP5.
- **RHEL for SAP Solutions 8** (x86_64): 8.0, 8.1, 8.2, 8.4, 8.4.
- **RHEL for SAP Solutions 7** (x86_64): 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 7.9

Supported SAP HANA Versions

Veeam Plug-in for SAP HANA supports the following versions of SAP HANA:

- **SAP HANA 2.0**: SPS 02, SPS 03, SPS 04, SPS 05 (only with Backint version 1.0), SPS06. Express Edition is not supported.
- **SAP HANA 1.0**: SPS12 and newer.

> **NOTE**
>
> To check whether an OS version is compatible with the SAP HANA version you want to use, see the SAP HANA Administration Guide.

Veeam Backup & Replication

Mind the following compatibility of Veeam Backup & Replication and Veeam Plug-in versions:

- **Veeam Plug-in for SAP HANA 11** supports integration with Veeam Backup & Replication version 11, 11a Cumulative Patch P20211217.
- **Veeam Plug-in for SAP HANA 10.0.1.4854 (10a Cumulative Patch 20201202)** supports integration with Veeam Backup & Replication version 10, 11.
- **Veeam Plug-in for SAP HANA 10 (earlier than 10.0.1.4854)** supports integration only with Veeam Backup & Replication version 10.
- **Veeam Plug-in for SAP HANA 9.5 Update 4** supports integration with Veeam Backup & Replication version 9.5 Update 4, 9.5 Update 4a, and 9.5 Update 4b, 10.

Note that if you want to use the latest functionality, you must upgrade both Veeam Backup & Replication and Veeam Plug-in to the latest version.

Network

Veeam Plug-in should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Plug-in cannot work with the Veeam Backup & Replication server that is located behind the NAT gateway.

Required Permissions

User Rights on the SAP HANA Server

The account used for installing and updating Veeam Plug-in must have root privileges.

Veeam Backup Server User

- The account specified in the Veeam Plug-in configuration settings must be able to authenticate against the Veeam Backup & Replication server. For details, see Configuring Veeam Plug-in for SAP HANA.
- The account specified in the Veeam Plug-in configuration settings must be granted access rights on the Veeam backup repository where you want to store backups.

  To learn how to grant permissions on Veeam repositories, see Granting Permissions on Repositories.
- You can work with backups created by Veeam Plug-in only with the account used for creating the backups. If you want to use another account, see required permissions in Configuring Veeam Plug-in for SAP HANA.

Used Ports

To enable proper work of Veeam Plug-ins, make sure that the following ports are open.

SAP HANA Server

The following table describes network ports that must be opened to ensure proper communication of the SAP HANA server and backup infrastructure components.

From	To	Protocol	Port	Notes
SAP HANA server where Veeam Plug-in is installed	Veeam Backup & Replication server	TCP	10006	Default port used for communication with the Veeam Backup & Replication server. Note that data between Veeam Plug-ins and backup repositories is transferred directly, bypassing the Veeam Backup & Replication server.
	Backup repository server or gateway server*	TCP	2500 to 3300**	Default range of ports used as data transmission channels. For every TCP connection that a backup process uses, one port from this range is assigned.

* For NFS share, SMB share repositories, and Dell EMC Data Domain, HPE StoreOnce deduplication storage appliances, Veeam Backup & Replication uses an auxiliary backup infrastructure component — gateway server. For details, see the Gateway Server section of the Veeam Backup & Replication User Guide.

** This range of ports applies to newly added backup infrastructure components. If you upgrade to Veeam Backup & Replication 10.0 from earlier versions of the product, the range of ports from 2500 to 5000 applies to the already added components.

Backup Repositories and Gateway Servers

On backup infrastructure components, Veeam Backup & Replication automatically creates firewall rules for the required ports. These rules allow communication between the components. Depending on the type of backup repositories that you use for Veeam Plug-in backups, the following ports must be open to allow communication between backup infrastructure components.

From	To	Protocol	Port	Notes
Veeam Backup & Replication server	Backup repository server or gateway server*	TCP	2500 to 3300**	Default range of ports used as data transmission channels. For every TCP connection that a backup process uses, one port from this range is assigned.
Direct Attached Storage				

From	To	Protocol	Port	Notes
Veeam Backup & Replication server	Linux server used as a backup repository or gateway server	TCP	22	Port used as a control channel from the Veeam Plug-in server to the target Linux host.
	Microsoft Windows server used as a backup repository or gateway server	TCP UDP	135, 137 to 139, 445	Ports used as a management channel from the Veeam Plug-in server to the Repository/Gateway server. Also, the ports are used to deploy Veeam components.
		TCP	6160, 6162	Default ports used by the Veeam Installer Service and Veeam Data Mover Service
Network Attached Storage				
Gateway server (specified in the SMB share repository settings)	SMB server	TCP	445	Default port used by the SMB transport protocol.
		TCP UDP	135, 137 to 139	SMB/Netbios name resolution for the SMB protocol (needed in some cases). For details, see the Used Ports section of the Veeam Backup & Replication User Guide.
Gateway server (specified in the NFS share repository settings)	NFS server	TCP UDP	111, 2049	Standard NFS ports used as a transmission channel from the gateway server to the target NFS share.
Dell EMC Data Domain				
Veeam Backup & Replication server or **Gateway server**	Dell EMC Data Domain For more information, see Dell EMC Documents.	TCP	111	Port used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned.
		TCP	2049	Main port used by NFS. To change the port, you can use the `'nfs set server-port'` command. Note that the command requires SE mode.

From	To	Protocol	Port	Notes
		TCP	2052	Main port used by NFS MOUNTD. To change the port, you can use the `'nfs set mountd-port'` command. Note that the command requires SE mode.
HPE StoreOnce				
Veeam Backup & Replication server or **Gateway server**	HPE StoreOnce	TCP	9387	Default command port used for communication with HPE StoreOnce.
			9388	Default data port used for communication with HPE StoreOnce.
ExaGrid				
Veeam Backup & Replication server	ExaGrid	TCP	22	Default command port used for communication with ExaGrid.
Quantum DXi				
Veeam Backup & Replication server	Quantum DXi	TCP	22	Default command port used for communication with Quantum DXi.

For detailed list of ports used by Veeam Backup & Replication server and backup repositories, see the Used Ports section of the Veeam Backup & Replication User Guide.

Licensing

To use the Veeam Plug-in functionality, you must have a valid Veeam Backup & Replication license. Licenses are installed and managed on the Veeam Backup & Replication server that is connected to the Veeam Plug-in server. If the license is not valid or out of resources, Veeam Plug-in backup jobs fail.

This guide provides information only on specifics of Veeam licenses for Veeam Plug-ins. For terminology and general information about Veeam Licensing, see Veeam Licensing Policy.

In this section:

- Licensed Objects
- Supported License Types and Packages
- Obtaining and Managing Licenses

Licensed Objects

If you are using any instance-based (Veeam Universal Licensing) license on your Veeam Backup & Replication, you don't need to install any additional license keys.

A machine where SAP HANA is deployed is assumed protected if it has been processed by a Veeam Plug-in backup job in the last 31 days. When you back up SAP HANA databases on one host, one License Unit is consumed from the Veeam Backup & Replication license. A machine protected by both Veeam Plug-in and Veeam Backup & Replication will consume a License Unit only once. For example, you have an SAP HANA server that you back up using Veeam Plug-in. You can also back up this server using image-level backup functionality of Veeam Backup & Replication. In this case, only one License Unit will be consumed.

NOTE

If you are using a legacy perpetual per-socket license, a license is required for each hypervisor CPU socket occupied by protected SAP HANA servers.

A socket is consumed from the license only if the hypervisor where protected servers reside is added to the Veeam Backup & Replication infrastructure. If the hypervisor is not added to the Veeam Backup & Replication infrastructure, an instance unit will be consumed from the license. To learn how to add a hypervisor to the Veeam Backup & Replication infrastructure, see the Virtualization Servers and Hosts section of the Veeam Backup & Replication User Guide.

IMPORTANT

If you have an SAP HANA Scale-Out Cluster, each node will consume one License Unit. The License Units are consumed for all cluster nodes, even if Veeam Plug-in is installed only on one of the nodes.

Supported License Types and Packages

You can use Veeam Plug-ins with the following license types and packages. Note that this guide contains information on specifics of Veeam license packages only for Veeam Plug-ins. For the full list of license packages, see Pricing and Packaging.

- **For Veeam Universal Licensing**:

  You can use Veeam Plug-ins with all license packages (*Veeam Backup Essentials*, *Veeam Backup & Replication*, *Veeam Availability Suite*).

Note that if you use the *Rental* license type, functionality of Veeam Plug-ins is supported only for the *Enterprise Plus* edition of Veeam Backup & Replication.

- **For Perpetual Socket license**:

  Functionality of Veeam Plug-ins is supported only for the *Enterprise Plus* edition of Veeam Backup & Replication.

Obtaining and Managing Licenses

To learn how to install a license and monitor licensed objects, see the Licensing section in the Veeam Backup & Replication User Guide.

Environment Planning

Integration of SAP HANA and Veeam Plug-in requires additional environment planning. When you deploy the plug-in, keep in mind the following requirements and limitations.

Compression

Veeam Plug-in uses built-in compression functionality of Veeam Backup & Replication. If you want to disable the compression, do the following:

1. Open the `/opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml` file with a text editor.
2. In the `veeam_config.xml` file, find the `<AgentParams />` line and add the following parameter:

```
<AgentParams compression="NoCompression" />
```

Scheduling

Veeam Plug-in forwards the backups created by SAP HANA integrated backup application to a Veeam backup repository. You can schedule backup operations with all SAP HANA relevant scheduling options like SAP HANA Cockpit (HANA Cockpit 2.0 SPS 06 or later version), SAP DB13 (NW 7.02 SP17 or later version) or external schedulers like cron, UC4, TWS and others.

For instructions on how to configure external schedulers, see the Veeam Plug-in for SAP HANA Best Practices.

> **NOTE**
>
> For SAP Management Software, make sure SAP HANA 2.0 systems are configured in the Multiple-Container mode. Otherwise, backups will fail with the following error: `[110091] Invalid path selection for data backup using backint`. For details, see the SAP HANA Multitenant Database Containers section of the SAP HANA Master Guide.

Veeam User Management

Veeam Plug-in for SAP HANA uses the Windows authentication methods of the Veeam Backup & Replication server to establish a connection to this server and to the target backup repository. It is recommended to create one specific user for each Veeam Plug-in server or for each scale-out cluster.

If this user will be later changed manually, the new user must have at least the *Veeam Backup Operator* and *Veeam Restore Operator* rights within the Veeam Backup & Replication user management. For instructions on how to assign Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication Guide.

SAP HANA Backup Channels and Veeam Repository Task Slots

By default, SAP HANA uses one channel per data backup operation. You can configure SAP HANA to use additional channels. When multiple channels are used, SAP HANA distributes the data equally across available channels.

To control the number of parallel channels used for each SAP HANA Backint instance, you can edit the `parallel_data_backup_backint_channels` parameter in the SAP HANA `global.ini` file. For instructions, see the Multistreaming Data Backups with Third-Party Backup Tools section of the SAP HANA Administration Guide

> **NOTE**
>
> The number of multistreaming channels applies to all data backup services larger than 128GB. Data backup services smaller than 128GB use only one channel.

Basically, the more channels used in parallel, the faster is the data flow between SAP HANA and the source Veeam Transport Agent. However, the more channels used in parallel, the more resources are used on the SAP HANA server, network, Veeam backup repository, backup source and target disk systems. You should find the right mix between performance and resource allocation for your specific business need.

The following hardware resources are recommended based on tests on Skylake processors:

- **SAP HANA server**: 1 CPU core and 200 MB of RAM per currently used channel.
- **Backup repository server**: 1 CPU core and 1 GB of RAM per 5 currently used channels.

  These resources are recommended only if you use a dedicated backup repository for Veeam Plug-in backups. If you use the same backup repository for Veeam Plug-in backups and VM backups created by Veeam Backup & Replication or Veeam Agents, consider adding the mentioned above hardware resources based on usual load on your backup repository. For details on hardware requirements for a backup repository, see the System Requirements section of the Veeam Backup & Replication User Guide.

  We recommend to contact your Veeam system engineer to optimize the channel settings and resource allocation. Also, mind the following:

  - It is not recommended to use more than 64 channels in parallel as the overhead will reduce individual channel performance. Set the `max_recovery_backint_channels` setting in `global.ini` to 64 or below depending on available hardware resources.
  - It is recommended to use a separate backup repository for Veeam Plug-in backups.
  - If you want to improve backup performance, the SAP HANA buffer must be increased for additional used channels. For details, consult with your SAP HANA database administrator.
  - SAP HANA can back up individual databases and tenants in parallel. To optimize resources, you can back up databases sequentially.
  - If there are not enough available repository task slots, SAP HANA waits till repository task slots become available.
  - During restore, the order of repository task slots is ignored, and channels are used as requested by SAP HANA.
- **Veeam Backup & Replication server**: during manual metadata operations such as import of backup files, the Veeam Backup & Replication server needs additional 15 GB of RAM per 1 million files located in the same backup job folder.

You can use the following examples as a reference:

- **Example 1: Backing up all databases in parallel**

  In this example, there is a system with 2 tenant databases, each database has 4 services. The databases are backed up in parallel. The SAP HANA channel setting is 6. The following maximum repository task slots and SAP channels are used:

  - Up to 4 task slots/channels are used by SYSTEMDB and its 4 services (all below 128 GB)

- Up to 6 task slots/channels are used for the index service of the tenant database 1 (the database is bigger than 128 GB)
- Up to 3 task slots/channels are used for the rest of the 3 remaining services of the tenant database 1 (all below 128 GB)
- Up to 6 task slots/channels are used for the index service of the tenant database 2 (the database is bigger than 128 GB)
- Up to 3 task slots/channels are used for the rest of the 3 remaining services of the tenant database 2 (all below 128 GB)
- If the log backups are below 128GB, you must reserve at least 3 channels for the log backup of SYSTEMDB, tenant database 1, and tenant database 2. These log backups are started automatically on their own schedule or when the maximum file size of the log file is reached.

In total, for backup processes of all databases started in parallel you need up to 27 available task slots.

- **Example 2: Backup of all databases sequentially**

  In this example, there is system with 2 tenant databases, each database has 4 services. The databases are backed up sequentially. The SAP HANA channel setting is 6. The following maximum repository task slots and SAP channels are used:

  - Up to 6 task slots/channels are used for the index service of a tenant database (the database is bigger than 128 GB).
  - Up to 3 task slots/channels are used for the rest of the 3 remaining services of the same tenant database (all below 128 GB).
  - If the log backups are below 128GB, you must reserve at least 3 channels for the log backup of SYSTEMDB, tenant database 1, and tenant database 2. These log backups are started automatically on their own schedule or when the maximum file size of the log file is reached. Assuming that the log file backups are below 128 GB and do not use additional channels.

  In total, for backup processes of sequential started database backups, 12 task slots must be available.

SAP HANA Encryption

Veeam Plug-in supports SAP HANA integrated encryption. The encryption processes are performed on the SAP HANA side. Veeam Plug-in is not involved in encryption processing.

Plan the protection of the encryption environment carefully. In case the encryption keys are lost, Veeam Plug-in can only provide an access to the encrypted backup file. You will have to decrypt data in SAP HANA. For details, see the Managing Data Encryption section of the SAP HANA Administration Guide.

SAP HANA Catalog Backup with Backint

To back up the SAP HANA catalog using Backint, change the settings of the **catalog_backup_using_backint** parameter in the **backup** section of the `global.ini`.

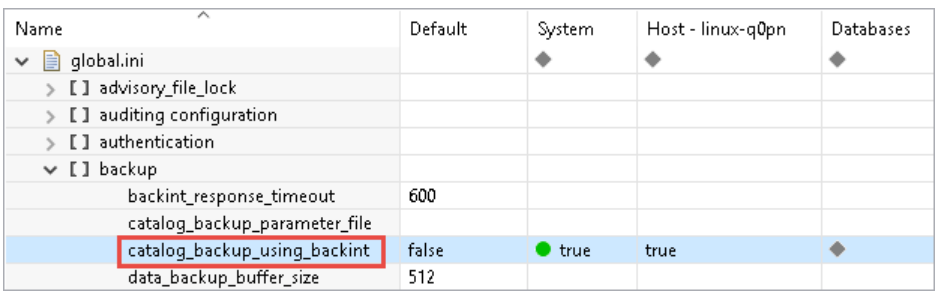

SAP HANA Backint Parameter File

Veeam Plug-in does not use the Backint parameter file. Leave these fields empty when asked for.

SAP HANA Scale-Out Cluster

Veeam Plug-in supports SAP HANA scale-out clusters with the following limitations:

- Due to design of SAP HANA databases, the same Veeam Plug-in configuration must be set on all scale-out cluster members, including stand-by nodes.
  - On all cluster nodes, Veeam Plug-in must be configured to transfer backups to the same repository.
  - Each cluster node must use the same credentials to connect to Veeam servers.
- All backup tasks across the SAP HANA scale-out cluster are performed in parallel.

SAP HANA System Replication Failover

SAP HANA does not allow you to back up from replicas. You can back up these databases only after a failover. To prepare the replication target system for backups after the failover, you can configure Veeam Plug-in as usual for a new scale-up or scale-out system. The Veeam backup job object will be created at first backup run and reflect the hostname of each system.

> **IMPORTANT**
>
> You must perform full database backup at least once after each failover or failback, so that SAP HANA starts to create automatic log backups.

After the failover, if you want to restore backups created before the failover, you must configure the plug-in to be able to access the backup files from the original source system:

1. Go to `/opt/veeam/VeeamPluginforSAPHANA` and run the Veeam Plug-in configuration tool with the following parameter.

```
VM2ADM:/opt/veeam/VeeamPluginforSAPHANA> SapBackintConfigTool --set-restor
e-server
```

2. Select the original source server.

```
Select source SAP HANA plug-in server to be used for system copy restore:
1. SAP-VM1
2. SAP-VM02
Enter server number: 1
```

3. Specify the backup repository where the required source server backup is stored.

```
Available backup repositories:
1. serv10_repo
Enter repository number: 1
```

4. Perform system copy restore. For instructions, see Recovering Databases to Other Servers.

5. Later, if you want to restore from the new backup chain created from the system replication server, you must run the command again and select the system replication server as a source for restore.

Hosting Environments

By default, Veeam Plug-in uses a hostname to create the Veeam Backup & Replication job object and a folder where the backups will be stored. If server names match, you can set the following entry in the Veeam configuration XML file (`/opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml`) to be able to distinguish servers:

```
<PluginParameters customServerName="hostname.domain.tld" />
```

Example:

If your servers in multiple environments have the name *sap1* and the domains for the 2 environments are `customer1.local` and `customer2.local` you have to set the following entries:

```
<PluginParameters customServerName="sap1.customer1.local" />
<PluginParameters customServerName="sap1.customer2.local" />
```

Additional Files to Back Up

SAP HANA INI Files

SAP HANA does not back up the SAP configuration stored in INI files. Contact your SAP HANA database administrator to discuss the backup of the following files:

- `/usr/sap/<SID>/SYS/global/hdb/custom/config`
- `/usr/sap/<SID>/<INSTANCE>/<FQDN>`
- `/usr/sap/<SID>/SYS/global/hdb/custom/config`

Also, to backup SAP HANA configuration files, you can use file or image-level backup options of Veeam Backup & Replication or Veeam Agent for Linux.

SAP HANA Server for Disaster Recovery

You can use Veeam Backup & Replication or Veeam Agent to create an image-level backup of the SAP HANA server. Note that to create transaction-consistent backups, you must use pre-freeze and post-thaw scripts.

Veeam Plug-in Configuration File

To back up the configuration file of Veeam Plug-in, back up the following file:
`/opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml`

Backup Files Format

Veeam Plug-in stores backup files in the following formats:

- A .VAB file stores compressed and deduplicated copy of a SAP HANA database. Veeam Plug-in creates .VAB files for all types of backups.

- A .VASM file stores metadata that contain information about the backup. A .VASM file is created for each .VAB file. .VASM files are used by Veeam Backup & Replication to get data about Veeam Plug-in backups.
- A .VACM file stores metadata of a backup job object.

Veeam Plug-in backup file names match the backup file ID's (EBID) created by SAP HANA.

Veeam Backup Repositories

Veeam Plug-ins store backup files in repositories added to the Veeam Backup & Replication infrastructure. In this section, you can find the list of supported backup repositories and limitations for Veeam Plug-in backups.

Supported Backup Repositories

Veeam Plug-in for SAP HANA supports integration with the following types of repositories added to the Veeam Backup & Replication infrastructure:

- Windows Server
- Linux Server
- CIFS (SMB) Share
- Dell EMC Data Domain
- HPE StoreOnce. If you plan to use HPE StoreOnce as a backup repository for Veeam Plug-in backups, the total number of stored files (data and metadata) must not exceed 3,000,000 per Catalyst store. If necessary, multiple Catalyst stores may be created on the same StoreOnce system.
- Quantum DXi
- NFS File Share
- ExaGrid. If you plan to use an ExaGrid appliance as a backup repository for Veeam Plug-in backups, mind the following:
  - Make sure the repository is configured as described in the ExaGrid section of the Veeam Backup & Replication User Guide.
  - In the `global.ini` settings of SAP HANA, you must set the `max_recovery_backint_channels` parameter value to a number lower than the number of repository task slots. ExaGrid recommends setting it to *1*, and adjust gradually if needed.

You can also use scale-out backup repositories that contain repositories supported by Veeam Backup & Replication.

Backup Repositories Limitations

- For Veeam Plug-in backups, the warning which indicates that free space on a storage device has reached a specified threshold is configured in the **veeam_config.xml** file of Veeam Plug-in. The warning settings in the Veeam Backup & Replication console does not affect this setting.

  To configure the warning settings, add the following parameter in the `/opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml` file.

  ```
  <PluginParameters repositoryFreeSpacePercentWarning="10" />
  ```

- Due to specific design of SAP HANA backups, Veeam Plug-in does not use fast cloning. Backups transferred to repositories that use ReFS or XFS as a file system are processed the same way as with NTFS repositories.

- The plug-in configuration wizard will not show repositories where the **Encrypt backups stored in this repository option** is enabled. To learn how to disable the encryption option, see Access and Encryption Settings on Repositories.
- Make sure Veeam backup repositories have enough free space to store database backups and transaction log backups. If required, you can use a scale-out backup repository.
- If you use Veeam Plug-in for SAP HANA version 9.5 Update 4, Veeam Backup & Replication backup copy jobs does not process and copy the backups created by the plug-in.
- Veeam extract utility cannot extract backup files created by Veeam Plug-in.
- For security reasons, it is recommended to use separate repositories for different users and grant access to backup repositories only for required users.

Scale-Out Backup Repositories

If you want to store Veeam Plug-in backups on scale-out backup repositories, mind the following:

- For Veeam Plug-in backups and backup copies, the *Performance* policy of a scale-out repository functions differently:
  a. Veeam Backup & Replication checks if there are extents without warning on free space insufficiency. If all extents have the warning, Veeam Backup & Replication uses an extent with the largest amount of free space that has a free task slot.
  b. If there are extents without the warning, Veeam Backup & Replication checks if there are incremental extents with free task slots. If there are no incremental extents with free task slots, Veeam Backup & Replication uses a full extent with the least amount of used task slots.
  c. If there are incremental extents with free task slots, Veeam Backup & Replication sends backup files to an incremental extent with the least amount of used task slots. If the amount of used tasks is the same, an extent with the largest amount of free space.
- If a scale-out repository is configured in the **Data locality** policy, repository extents will be selected according to the amount of free space for each SAP HANA Backint connection. If there are two extents with one slot on each extent, the backup will be launched in two streams (one on each extent).
- If you want to add a backup repository as an extent to a scale-out backup repository and Veeam Plug-in backups are present on this backup repository, you must do the following:
  a. In the Veeam Backup & Replication console, select Veeam Plug-in backup files that reside in this backup repository and remove them from configuration. For details, see Removing backups from configuration. Note that this action does not delete the backups from the repository.
  b. In the Veeam Backup & Replication console, delete the Veeam Plug-in backup job. For details, see Deleting Jobs.
  c. Add the repository as an extent to the scale-out repository. For details, see Extending Scale-Out Repositories.
  d. Rescan the scale-out repository. For details, see Rescanning Scale-Out Repositories.

> **NOTE**
>
> Names of backup files and paths to backup files must contain only allowed characters:
>
> - Alphanumeric characters: `a-zA-Z0-9`
> - Special characters: `_-.+=@^`
> - Names of backup files and paths to backup files must not contain spaces.

e. On the Veeam Plug-in server, set the scale-out repository as the target for backups using the following command:

```
SapBackintConfigTool --set-repository
```

f. Map the imported backups using the following command:

```
SapBackintConfigTool --map-backup
```

Capacity Tier

You can configure Veeam Backup & Replication to transfer Veeam Plug-in backup files to a capacity tier. Both policies (Move policy, Copy policy) are supported for Veeam Plug-in backups with the following limitations:

- For Veeam Plug-in backup files, capacity tier does not verify whether data that is being moved is unique and has not been offloaded earlier. Thus, it is highly recommended to check the pricing plans of your cloud storage provider to avoid additional costs for offloading and downloading backup data.
- Capacity tier does not track dependencies of full and incremental Veeam Plug-in backup files. Thus, mind the following:
  - [For the Move policy] When backup files are transferred to the capacity tier, Veeam Backup & Replication takes into account only the creation time of backup files. Make sure that the operational restore window is not longer than the whole backup chain cycle period. Otherwise, you may encounter the scenario when full backup files are transferred to the capacity tier and their increment backup files still remain in the performance tier.
  - The capacity tier immutability expiration date does not have the additional block generation period. The immutability expiration date is based only on the number of days specified in settings of the object storage backup repository.
- If a scale-out repository is down, you cannot restore from the Veeam Plug-in backup files stored on the capacity tier. In this case, you can only import the backup files manually and then perform the data recovery operations.
- If you use a capacity tier that has been created in Veeam Backup & Replication version 10, you cannot transfer Veeam Plug-in backup files to a capacity tier. However, if you want to transfer them manually, do the following:
  - If the backup files are created by Veeam Plug-in version 10, upgrade the metadata of backup files as described in Upgrading Metadata Files to New Format.
  - Run the Set-VBRScaleOutBackupRepository PowerShell command with the `-EnablePluginBackupOffload` parameter to offload backup files to the capacity tier.

- If you want to restore from backups stored on the capacity extent, at least one performance extent must be available or you must switch the `catalog_backup_using_backint` parameter to the *False* state in the `global.ini` file. Otherwise, at the end of the restore process, SAP Backint will not be able to back up the catalog and restore will fail.

Access and Encryption Settings on Repositories

While you configure Veeam Plug-in, you specify an account that must be used to connect to the Veeam Backup & Replication server. To be able to store backups on a backup repository, the specified account must have access permissions on the target backup repository.

To grant access permissions, do the following:

1. In Veeam Backup & Replication, open the **Backup Infrastructure** view.
2. In the inventory pane, click the **Backup Repositories** node or the **Scale-out Repositories** node.
3. In the working area, select the necessary backup repository, click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**.

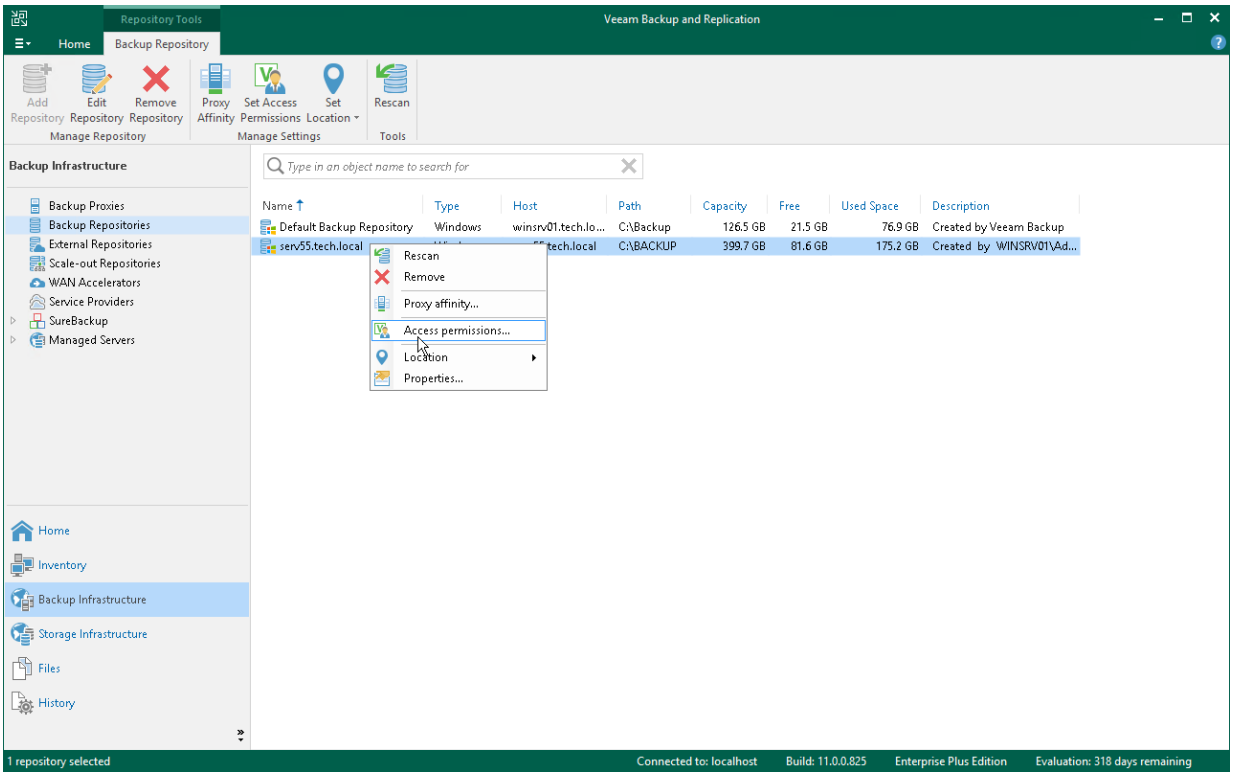

4. In the **Access Permissions** window, specify to whom you want to grant access permissions on this backup repository:
   - *Allow to everyone* – select this option if you want to grant repository access to any user. This option is equal to granting access rights to the *Everyone* group in Microsoft Windows (anonymous users are excluded). For security reasons, the option is not recommended for production environments.
   - *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 repository. Click **Add** to add the necessary users and groups to the list.

5. Veeam Plug-ins cannot send backups or backup copies to a repository where encryption is enabled. Thus, unselect the **Encrypt backups stored in this repository** check box.

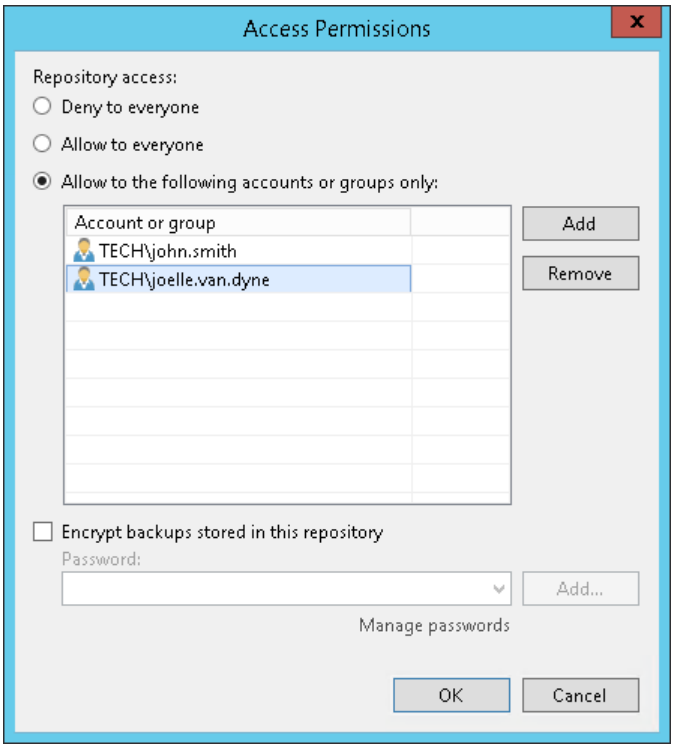

Deployment and Configuration

To deploy Veeam Plug-in, you must install the plug-in on a SAP HANA server and configure plug-in integration settings. In this section:

- Installing Veeam Plug-in for SAP HANA
- Configuring Veeam Plug-in for SAP HANA
- Automatic Configuration of Veeam Plug-in for SAP HANA
- Upgrading Plug-in for SAP HANA
- Importing Backups
- Upgrading Backup Files

This guide gives instructions on how to deploy Veeam Plug-in assuming that you have already deployed a Veeam Backup & Replication server and configured a backup repository. If you need instructions on how to deploy Veeam Backup & Replication, see the Veeam Backup & Replication User Guide for your platform.

Installing Plug-in for SAP HANA

Veeam Plug-in for SAP HANA is an additional component of Veeam Backup & Replication, and the installation package of the plug-in is included in the Veeam Backup & Replication installation ISO file.

You can install the plug-in using the `.RPM` package or extract the plug-in files from the `.TAR.GZ` archive. Depending on the type of package suitable for your OS, perform steps in one of the following guides:

- Installing Plug-in from .RPM Package
- Unpacking Plug-in from .TAR.GZ Archive

> **IMPORTANT**
>
> Mind the following:
>
> - Veeam Plug-in for SAP HANA must be installed on the SAP HANA server.
> - The `/opt/veeam` directory must be writable.
> - To install the plug-in, use the `sudo` command or a user with root privileges.
> - If you want to install Veeam Plug-in on an SAP HANA scale-out cluster, repeat the described installation process on all cluster nodes.

Installing Veeam Plug-in from .RPM Package

To install Veeam Plug-in, do the following:

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.X.X.XXXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk image from the Veeam Backup & Replication: Download page.

2. Open the mounted disk image and go to the `/Plugins/SAP HANA/x64` directory.
3. Upload the `VeeamPluginforSAPHANA-11.X.X.XXXX.x86_64.rpm` file to the SAP HANA server.
4. To install Veeam Plug-in, run the following command:

   ```
   rpm -i VeeamPluginforSAPHANA-11.X.X.XXXX.x86_64.rpm
   ```

Unpacking Veeam Plug-in from .TAR.GZ Archive

To extract plug-in files from the `.TAR.GZ` archive, perform the following:

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.X.X.XXXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk image from the Veeam Backup & Replication: Download page.

2. Open the mounted disk image and go to the `/Plugins/SAP HANA/x64` directory.
3. Upload the `VeeamPluginforSAPHANA.tar.gz` file to the SAP HANA server.

4. Create the `/opt/veeam` directory.

```
mkdir /opt/veeam
```

3. In the terminal, open the folder that contains the `VeeamPluginforSAPHANA.TAR.GZ` archive.
4. Unpack the plug-in files from the archive to the `/opt/veeam` directory.

```
tar -xzvf VeeamPluginforSAPHANA.tar.gz -C /opt/veeam
```

Configuring Plug-in for SAP HANA

When you configure Veeam Plug-in settings, you set up integration settings between a SAP HANA server, Veeam Backup & Replication server and backup repositories where backup files will be stored. Veeam Plug-in uses the **SapBackintConfigTool** wizard to configure the integration settings. The wizard configures the SAP HANA Backint settings and creates the `/opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml` file.

NOTE

- The configuration of Veeam Plug-in must be performed by a user with database administrator rights on all SAP HANA instances of the server.
- The SAP HANA High Level Isolation mode is not supported.

See the following instructions:

- Veeam Plug-in Configuration
- Configuration of Veeam Plug-in on Multiple SAP HANA Instances
- Verifying Configuration of Veeam Plug-in for SAP HANA
- Configuration Tool Commands

Veeam Plug-in Configuration

To configure Veeam Plug-in, do the following:

1. Log in with operating system user (*<sid>adm* or a user with similar rights) and run the following command to launch the Veeam Plug-in configuration tool. You do not need root privileges if you have configured group access as described in the Required Permissions section.

   ```
   SapBackintConfigTool --wizard
   ```

   If you have extracted files form the .TAR.GZ archive, go to the `/opt/veeam/VeeamPluginforSAPHANA` folder and run the following command:

   ```
   ./SapBackintConfigTool --wizard
   ```

2. Specify the DNS name or IP address of your Veeam Backup & Replication server.

   ```
   Enter backup server name or IP address: serv02.tech.local
   ```

3. Specify the port which will be used to communicate with the backup server. Default port: *10006*.

   ```
   Enter backup server port number: 10006
   ```

4. Specify credentials to authenticate against the Veeam Backup & Replication server.

```
Enter username: serv02\administrator
Enter password for serv02\administrator:
```

IMPORTANT

You can work with backups created by Veeam Plug-in only with the account used for creating the backups. If you want to use another account, assign the *Veeam Backup Administrator* role or the *Veeam Backup Operator* and *Veeam Restore Operator* roles to the account.

To learn how to assign Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication Guide.

5. Select the backup repository where you want to store backups. In the terminal dialog, enter the number of the repository from the list of available repositories.

```
Available backup repositories:
1. serv10_repo
2. serv07_repo
Enter repository number: 1
Configuration result:
SID SH2 has been configured
```

IMPORTANT

- The used account must have access to Veeam backup repositories that you plan to use.
- Encryption must be disabled on the repository.

Otherwise, the repositories will not be listed as available. To learn how to configure access and encryption settings on repositories, see Access and Encryption Settings on Repositories.

If you start the wizard for the first time on an SAP HANA scale-out cluster, the wizard asks you for a cluster name. The cluster name will be used by Veeam Backup & Replication to identify the backup job for the cluster. Further runs of the wizard within the SAP HANA scale-out cluster will not ask for this entry again.

NOTE

[For SAP HANA 1.0] If the wizard finishes with an error that required `hdbbackint` symlink cannotbe created, see this Veeam KB.

Configuration of Veeam Plug-in on Multiple SAP HANA Instances

Configuration of Veeam Plug-in includes configuration or creation of the SAP HANA Backint symlinks on all SAP HANA instances. To be able to do this for multiple SAP HANA instances at the same time, the configuration must be performed by a user with root privileges. Alternatively, you can use an account from the *sapsys* user group to configure the plug-in and set the symlink for SAP HANA instances where the account has access rights. You can repeat the wizard under another account to configure additional SAP HANA instances.

Alternatively, you can configure a Linux security group. To do that, you must add all Veeam Plug-in admins to this security group and set the following rights:

```
chown root:<youradmingroup> /opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml
chmod 664 /opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml
```

To learn about required permissions for backup and restore operations within SAP HANA, see the Authorization for Backup and Recovery section of the SAP HANA Administration Guide.

Verifying Configuration of Veeam Plug-in for SAP HANA

When you finish the plug-in configuration wizard, the plug-in creates a soft link in the `/hana/shared/<SID>/global/hdb/opt` directory.

To verify that the Backint Agent is configured correctly, do the following:

1. Connect to the database using SAP HANA Studio or Cockpit.
2. Go to `Backup/Configuration`.
3. In the **Backint Agent** field, make sure that the specified path leads to `/opt/veeam/VeeamPluginforSAPHANA/hdbbackint`.

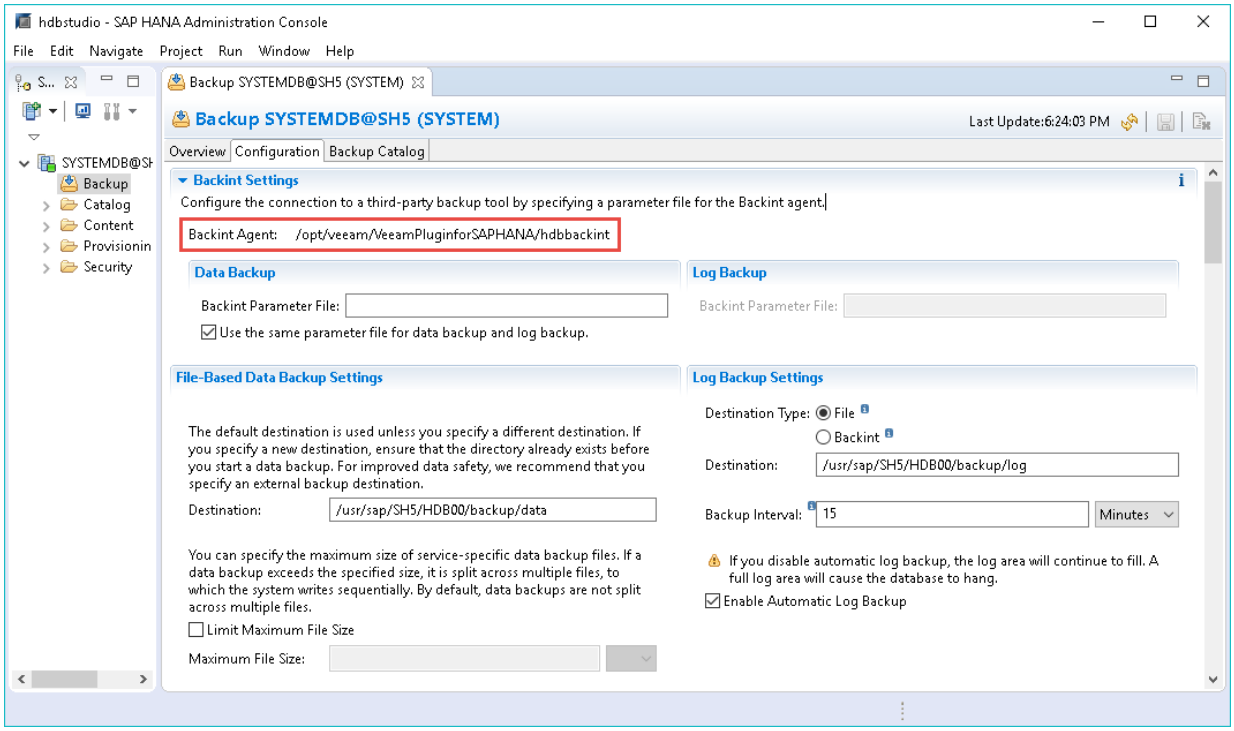

Configuration Tool Commands

Apart from running a configuration wizard, you can use the **SapBackintConfigTool** tool to change a specific parameter in the `veeam_config.xml` file or enable/disable Veeam Plug-in features.

See the list of available commands for **SapBackintConfigTool**:

Command	Description
--help	Shows the list of tool parameters.
--show-config	Shows configuration parameters.
--wizard	Starts the wizard to configure the plug-in settings. This wizard edits the `veeam_config.xml` file or creates a new one if the configuration file was removed from the `/opt/veeam/VeeamPluginforSAPHANA` directory.
--set-credentials <"serv\username"> <password>	Specifies credentials to log in to the Veeam Backup & Replication server.
--set-host <hostname>	Specifies the IP address or hostname of the Veeam Backup & Replication server.
--set-port <port_number>	Specifies a port number that will be used to communicate with the Veeam Backup & Replication server.
--set-repository	Launches a wizard to select a backup repository. A backup repository is selected from repositories which are available in the connected Veeam Backup & Replication instance.
--set-restore-server	[for System Copy] Specifies the backup that will be copied.
--map-backup	Maps the imported backups.
--set-force-delete	Deletes backup files after specified days.
--configure-restore-from-copy	Enables restore from backup copy. Note that if you enable restore from backup copy, you cannot back up databases with Veeam Plug-in. To revert changes, you must disable restore from backup copy. Note that when you launch the command, the wizard will ask you to reconfigure the catalog backup from backint to disk.
--promote-backup-copy-to-primary	Maps the imported backup copy to a regular Veeam Plug-in backup chain.

Example:

To specify credentials that will be used to log in to the Veeam Backup & Replication server, use the plug-in configuration tool with the following command.

```
SapBackintConfigTool --set-credentials "serv02\Administrator" "password"
```

Automating Configuration of Plug-in for SAP HANA

To automate the configuration of Veeam Plug-in for SAP HANA, do the following:

1. Copy the `veeam_config.xml` file to other servers where you want to configure the plug-in.
2. The password stored in the configuration file is encrypted with a machine key. Thus, on each machine, after the `veeam_config.xml` file was copied, you must reset the password of the account used to log in to the Veeam Backup & Replication server. To reset the password, use the following command. Note that the operation requires *root* privileges.

```
SapBackintConfigTool --set-credentials <"serv\username"> <password>
```

Upgrading Plug-in for SAP HANA

Periodically, Veeam releases a new version of Veeam Backup & Replication that contains new features and bug fixes. The release package also contains a new version of Veeam Plug-ins.

If you want to upgrade Veeam Plug-in, note that Veeam Backup & Replication must be the same or newer that the version of Veeam Plug-in. If you want to use the latest functionality, you must upgrade both Veeam Backup & Replication and to the latest version. After the upgrade, you don't need to to re-run the Veeam Plug-in configuration wizard, the plug-in configuration files will be preserved.

IMPORTANT

Mind the following:

- Version of Veeam Backup & Replication must be the same or newer that the version of Veeam Plug-in. First, you must upgrade Veeam Backup & Replication, then you can upgrade Veeam Plug-ins. For instructions on how to upgrade Veeam Backup & Replication, see the Upgrading to Veeam Backup & Replication 11 section of the Veeam Backup & Replication User Guide.
- Operations in the terminal of the Linux machine require root privileges.
- If you want to upgrade Veeam Plug-in on an SAP HANA scale-out cluster, repeat the described upgrade process on all cluster nodes.

Before You Begin

Veeam Plug-in installation files are included in the installation disk image of Veeam Backup & Replication. You must upload the installation file to the SAP HANA server. To do this, perform the following steps.

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.X.X.XXXX.iso`).
2. Open the mounted disk image and go to the `Plugins\SAP HANA\x64` directory.
3. Select the the Veeam Plug-in installation file and upload it to the SAP HANA server.

For instructions on how to upgrade Veeam Plug-in for SAP HANA, see the following guides:

- Upgrading Plug-in on Linux (RPM)
- Upgrading Plug-in on Linux (TAR.GZ)

Upgrading Plug-in on Linux (.RPM)

To upgrade Veeam Plug-in for SAP HANA from the `.RPM` package, perform the following:

1. Upload the new `VeeamPluginforSAPHANA-11.X.X.XXXX.x86_64.rpm` package to the SAP HANA server.
2. Run the following command. Note that the operation requires *root* privileges.

```
rpm -U VeeamPluginforSAPHANA-11.X.X.XXXX.x86_64.rpm
```

> **TIP**
>
> To find out which version of Veeam Plug-in is installed on your server, you can use the following command: `rpm -qa | grep VeeamPlugin*`

Upgrading Plug-in on Linux (.TAR.GZ)

To upgrade Veeam Plug-in for SAP HANA on a Linux machine from the `.TAR.GZ` archive, do the following:

1. Upload the `VeeamPluginforSAPHANA.tar.gz` file to the SAP HANA server.
2. In the terminal, open the folder that contains the `VeeamPluginforSAPHANA.TAR.GZ` archive.
3. Unpack the plug-in files from the archive to the `/opt/veeam` directory. Old Veeam Plug-in files will be replaced by new files.

```
tar -xzvf VeeamPluginforSAPHANA.tar.gz -C /opt/veeam
```

Importing Backup Files

If the Veeam Backup & Replication server has failed and you have restored it in a new location, you can copy the backup files to a new repository and re-map the Veeam Plug-in backup files.

Limitations and Prerequisites

Mind the following limitations:

- If backup files are not imported according to instructions given in this section, Veeam Plug-in backup and restore operations may fail.
- The repository from which you plan to import backups must be added to the Veeam Backup & Replication infrastructure. Otherwise you will not be able to access backup files.
- [For backups of scale-out clusters and servers with the `customServerName` option] To avoid mapping failure, the cluster name must be the same as the name used before importing backups.
- If you are importing backup files from a scale-out backup repository, the names of backup files and paths to backup files must contain only allowed characters:
  - Alphanumeric characters: `a-zA-Z0-9`
  - Special characters: `_-.+=@^`
  - Names of backup files and paths to backup files must not contain spaces.

How to Import Veeam Plug-in Backup Files

To import Veeam Plug-in backup files, do the following:

1. Copy the backup file folder to a backup repository or add a new backup repository with this folder as a subfolder.

   > **TIP**
   >
   > Starting from version 11, each Veeam Plug-in backup file (.vab) has its own metadata file (.vasm). Make sure that you import backup files and all related metadata files. Also, you must import the backup job metadata file (.vacm) which is stored in the same folder.

2. Log in to the Veeam Backup & Replication console.
3. Open the **Backup Infrastructure** view.
4. In the inventory pane of the **Backup Infrastructure** view, select the **Backup Repositories** node.

5. In the working area, select the required backup repository and click **Rescan** on the ribbon. Alternatively, you can right-click the backup repository and select **Rescan**.

   During the rescan operation, Veeam Backup & Replication gathers information about backups that are currently available in the backup repository and updates the list of backups in the configuration database. After the rescan operation, backups that were not in this configuration database will be shown on the **Home** view in the **Backups > Disk (Imported)** node.

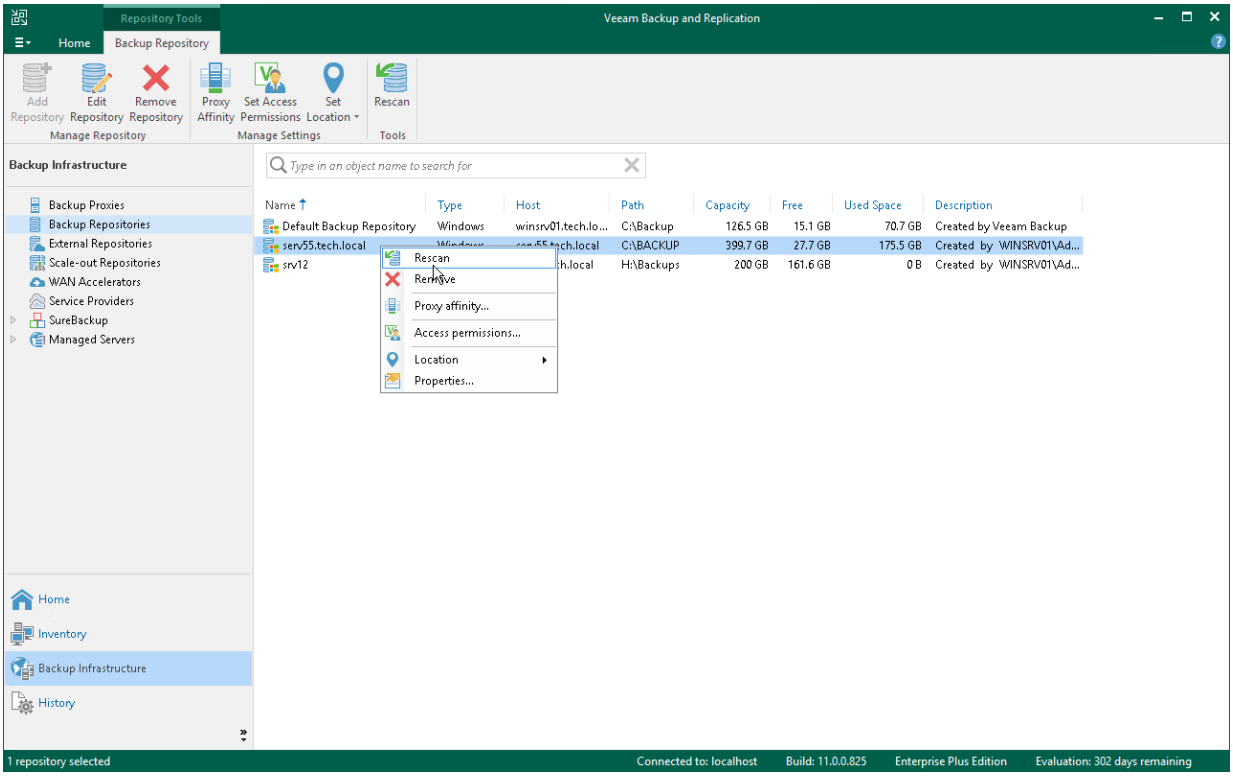

6. On the SAP HANA server, set the new repository as a target in the Veeam Plug-in settings:

```
sudo SapBackintConfigTool --set-repository
Available backup repositories:
1. serv55.tech.local
2. serv07_repo
Enter repository number: 1
Configuration result:
SID SH2 has been configured
```

7. Start the Veeam Plug-in configuration wizard with the following parameter:

```
sudo SapBackintConfigTool --map-backup
```

Upgrading Backup Files

Since version 11, Veeam Plug-in uses a new format of backup files: instead of one metadata file for all backup files there are separate metadata files (.vasm) for each database backup file (.vab). The new metadata format allows to optimize the productivity of backup and restore operations.

For Veeam Plug-in 11, the backup files upgrade is not obligatory. However, in the next version, backup files created by Veeam Plug-in version 10 will not be supported.

> **IMPORTANT**
>
> If you do not upgrade backup files, you will get the following warning in the job session logs: *Backup metadata is not up to date. Please upgrade the backup*. If you want to disable the warning, see instructions in this Veeam KB.

Prerequisites

Before upgrading backup files, make sure the following requirements are met.

- Make sure that you have upgraded Veeam Plug-in on the source server. If the plug-in is not upgraded to version 11 and you upgrade the backup files, then all next backup job runs will fail.
- Make sure that you have disabled the backup job whose backup files you want upgrade. You must also disable the backup copy jobs that use these backup files as a source.
- If the backup files reside on the scale-out backup repository, all repository extents must be available. Also, the extents must not be in the seal or maintenance mode.
- If you want to upgrade backup files created by a backup copy job, you must meet the same requirements as for the backup job files.
- During the process of the metadata upgrade, you cannot run the target backup job and you cannot restore from the backup files.

  The upgrade process duration depends on the number of backup files in the backup set, type of the backup repository and workload level on the file system.

  For example, there are backup files of the SAP HANA server that contains 10 instances and is backed up every 15 minutes with the retention policy set for 2 weeks. The upgrade of backup files can have the following duration on not overloaded file systems:

  - Microsoft Windows: 30 minutes
  - Linux: from 30 minutes to 3 hours
  - SMB/NFS: 1.5 hours
  - Data Domain Boost/Quantum DXi/ExaGrid/CIFS (SMB)/NFS file share: 3-4 hours
  - HPE StoreOnce: up to 10 hours (due to specifics of this repository type for processing large number of files)

Upgrading Backup Files in Veeam Backup & Replication Console

To upgrade backup files in the Veeam Backup & Replication console, do the following:

1. Open the **Home** view.
2. In the inventory pane, expand the **Backups** view and select **Disk**.
3. In the working area, right-click the job or the restore point and select **Upgrade**.

   Alternatively, you can select the job or the restore point and click **Upgrade** on the ribbon.

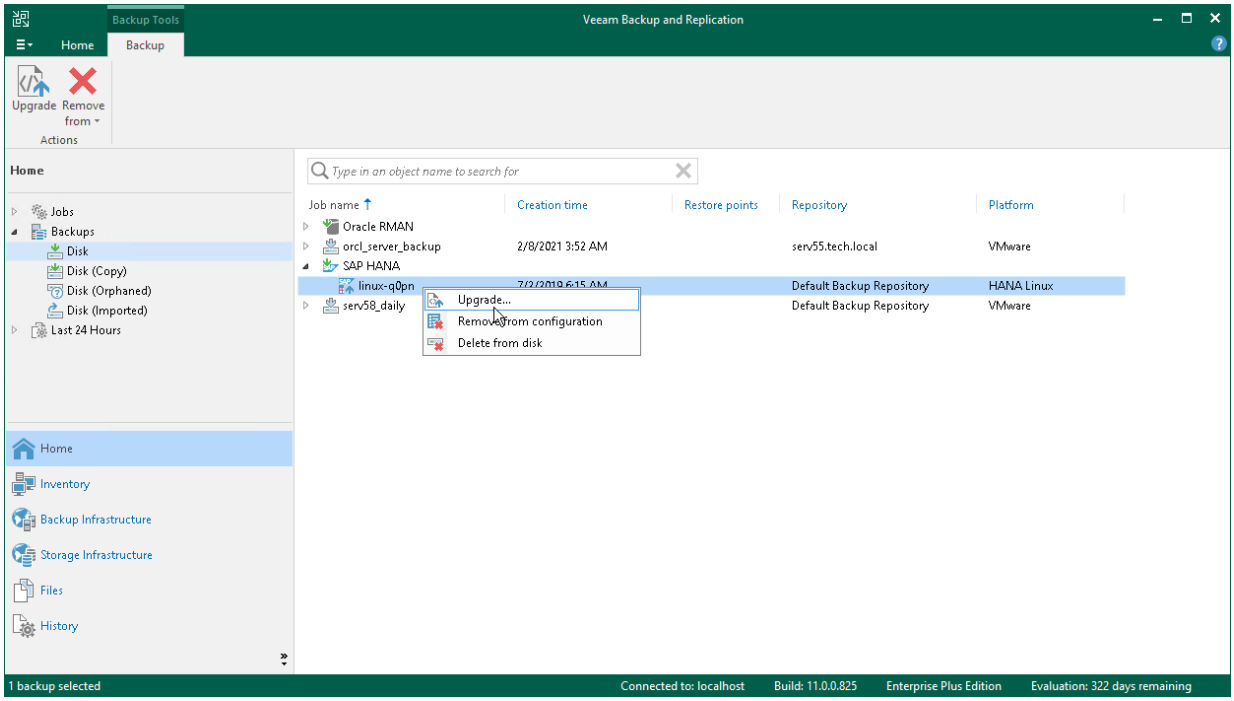

Uninstalling Plug-in for SAP HANA

To uninstall Veeam Plug-in for SAP HANA on a Linux machine, go to the directory with the Veeam Plug-in package and run the following command. Note that the operation requires *root* privileges.

```
rpm -e VeeamPluginforSAPHANA
```

Database Protection

After you configure Veeam Plug-in, you can back up databases with SAP HANA backup tools. Veeam Plug-in will automatically transfer data to the Veeam backup repository and store this data in Veeam proprietary format. The backup process itself is performed by SAP HANA Backint.

Keep in mind that examples in this section are provided only for demonstrating purposes. For details on full backup functionality of SAP HANA tools, see the SAP HANA Backup section of the SAP HANA Administration Guide.

> **IMPORTANT**
>
> Veeam Plug-in transfers backup data to the Veeam backup repository only when you perform the backup using SAP Backint.

To back up SAP HANA databases, you can use SQL commands or SAP HANA administration tools. For examples, see the following guides:

- Backing Up Databases Using SQL Commands
- Backing Up Databases with SAP HANA Studio
- Backing Up Databases with SAP HANA Cockpit

Database Backup (HDBSQL Scripts)

After you configure Veeam Plug-in settings, you can use HDBSQL to back up and restore SAP HANA databases. For details on the HDBSQL backup, see the BACKUP DATA Statement section of the SAP HANA SQL and System Views Reference.

Prerequisites

Before the backup process, you can use the `hdbuserstore` tool to set secure storage of SAP HANA connection details.

To configure `hdbuserstore`, you must log in to SAP HANA HDBSQL as the operating system administrator (*<sid>adm*) and run the following commands. For details, see the Secure User Store section of the SAP HANA Security Guide.

```
sh4adm@linux-q0pn:/usr/sap/SH4/HDB01> hdbuserstore SET <key> hostname:30013@SID
<username> <password>
sh4adm@linux-q0pn:/usr/sap/SH4/HDB01> hdbsql -U <key>
```

Backing Up SAP HANA Databases Using Backint

To back up the database with Backint, use one of the following commands depending on which type of backup you want to perform:

- Full backup of a tenant database.

  ```
  backup data for <TENANT_DATABASE_NAME> using backint ('backup_name_prefix'
  );
  ```

- Differential backup of a tenant database.

  ```
  backup data differential for <TENANT_DATABASE_NAME> using backint ('backup
  _name_prefix');
  ```

- Incremental backup of a tenant database.

  ```
  backup data incremental for <TENANT_DATABASE_NAME> using backint ('backup_
  name_prefix');
  ```

- Full backup of a tenant database with the ASYNCHRONOUS option. The ASYNCHRONOUS option can be helpful if you monitor SAP HANA backups on another host and just want to run the backup command from a script. The option runs the backup job in the background and closes the current script session.

  ```
  backup data for <TENANT_DATABASE_NAME> using backint ('backup_name_prefix'
  ) ASYNCHRONOUS;
  ```

- Full backup of SYSTEMDB.

```
backup data using backint ('backup_name_prefix');
```

- Differential backup of SYSTEMDB.

```
backup data differential using backint ('backup_name_prefix');
```

- Incremental backup of SYSTEMDB.

```
backup data incremental using backint ('backup_name_prefix');
```

Database Backup (SAP HANA Studio)

After you configure Veeam Plug-in settings, you can back up your databases using SAP HANA Studio. Veeam Plug-in will automatically transform backup files to a Veeam backup repository.

The example provided below is for demonstration purposes only. For details on the full backup functionality of SAP HANA Studio, see the Creating Data Backups and Delta Backups section of the SAP HANA Administration Guide.

To perform Backint backup with SAP HANA Studio, do the following:

1. In SAP HANA Studio, connect to the database as a user with DATABASE ADMIN privileges.
2. In the **Systems** view, right-click the database.
3. Select **Backup and Recovery** and then select **Back Up System Database** or **Back Up Tenant Database**.

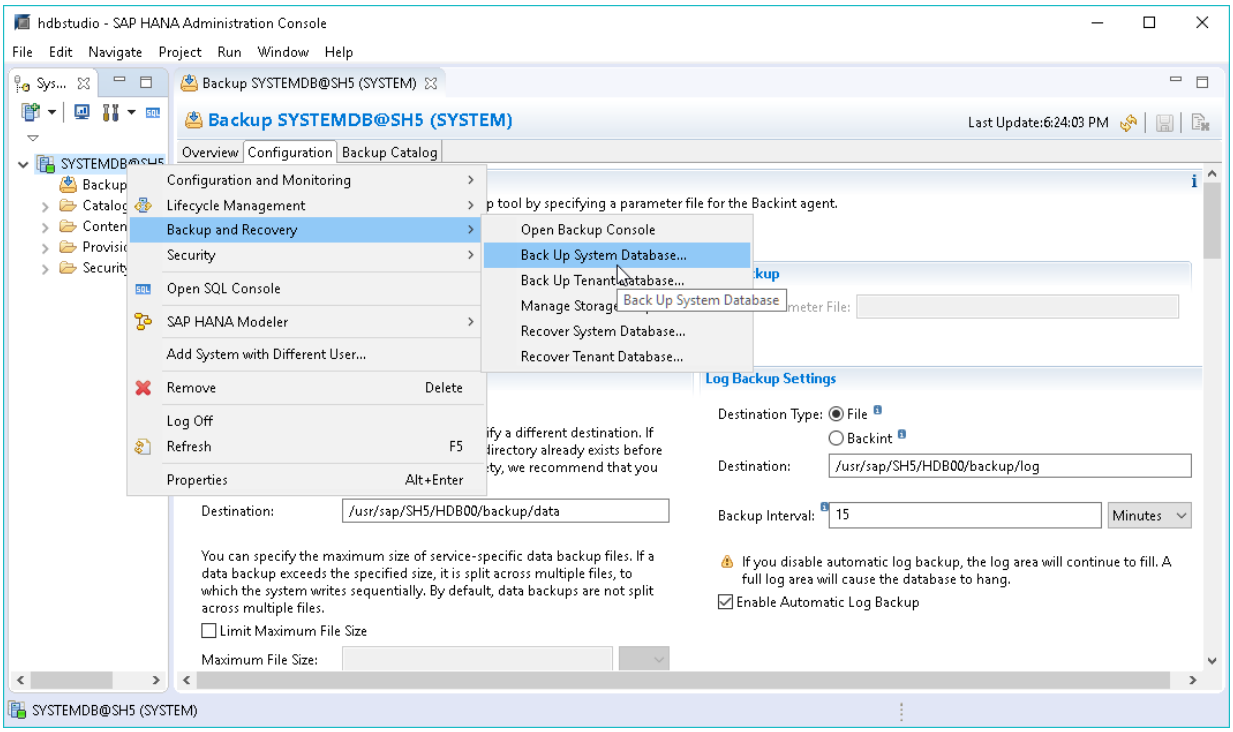

4. In the backup wizard, specify backup settings:
   a. Select the required backup type:
      - **Complete Data Backup**: backup of all data structures required to recover the database.
      - **Incremental Data Backup**: backup of data changed since the last full data backup or the last delta backup.
      - **Differential Data Backup**: backup of data changed since the last full data backup.
   b. In the **Destination Type** list, select *Backint*. With this option selected, Veeam Plug-in will transfer the backup file to Veeam backup repository.
   c. Change the default backup prefix, if needed.

d. Click **Next**.

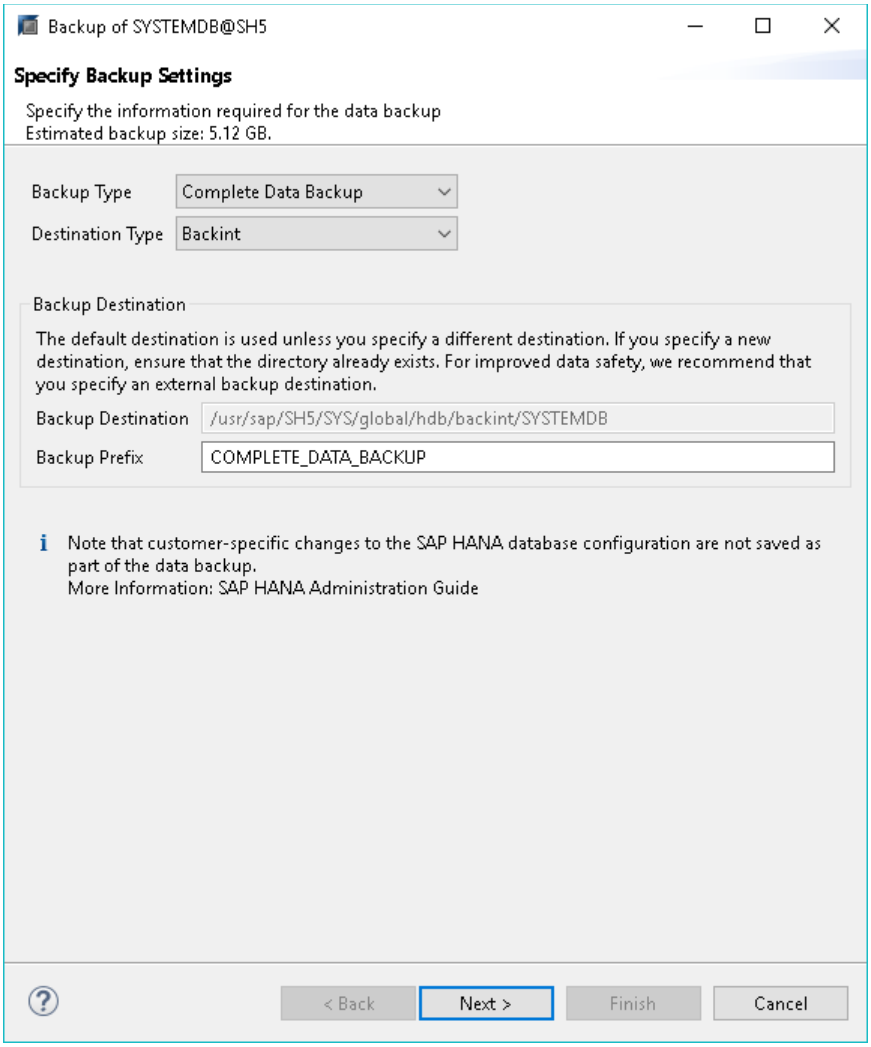
Backup of SYSTEMDB@SH5
Specify Backup Settings
Specify the information required for the data backup
Estimated backup size: 5.12 GB.
Backup Type
Complete Data Backup
Destination Type
Backint
Backup Destination
The default destination is used unless you specify a different destination. If you specify a new destination, ensure that the directory already exists. For improved data safety, we recommend that you specify an external backup destination.
Backup Destination
/usr/sap/SH5/SYS/global/hdb/backint/SYSTEMDB
Backup Prefix
COMPLETE_DATA_BACKUP
Note that customer-specific changes to the SAP HANA database configuration are not saved as part of the data backup.
More Information: SAP HANA Administration Guide
< Back
Next >
Finish
Cancel

5. In the **Review Backup Settings** step of the wizard, click **Finish** to start the backup process.

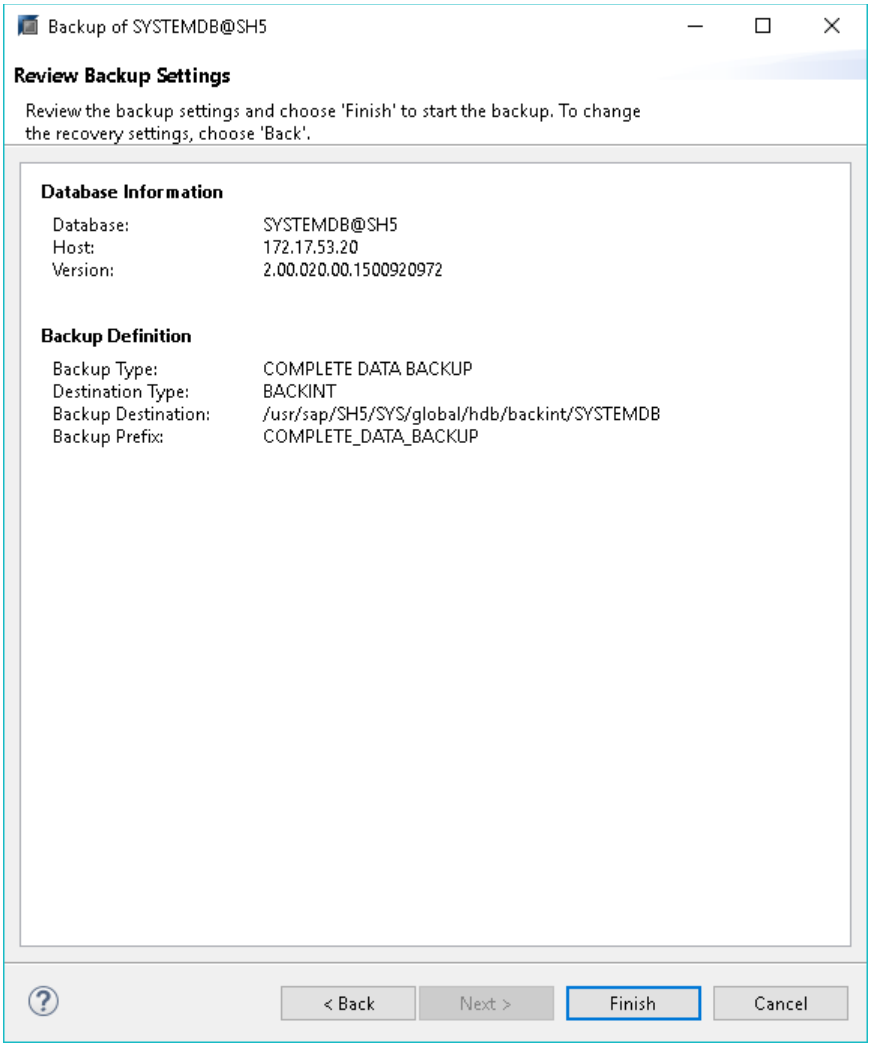

After you launch the backup process, SAP HANA Studio will back up the database, and Veeam Plug-in will forward backup files to the backup repository that is specified in the Veeam Plug-in settings.

Database Backup (SAP HANA Cockpit)

After you configure Veeam Plug-in settings, you can back up your databases with SAP HANA Cockpit 2.0. You can perform complete, incremental, and differential backups of SYSTEMDB and tenant databases. Veeam Plug-in will automatically transform backup files to Veeam backup repository. Keep in mind that you must select the Backint option as a destination target.

The example provided below is for demonstrating purposes only. For details on the full backup functionality of SAP HANA Cockpit, see the Create Data Backups and Delta Backups of the SAP HANA Administration Guide.

To perform Backint backup with SAP HANA Cockpit, do the following:

1. In the **System Overview** page, go to **Database Administration** and select **Manage Database Backups**.

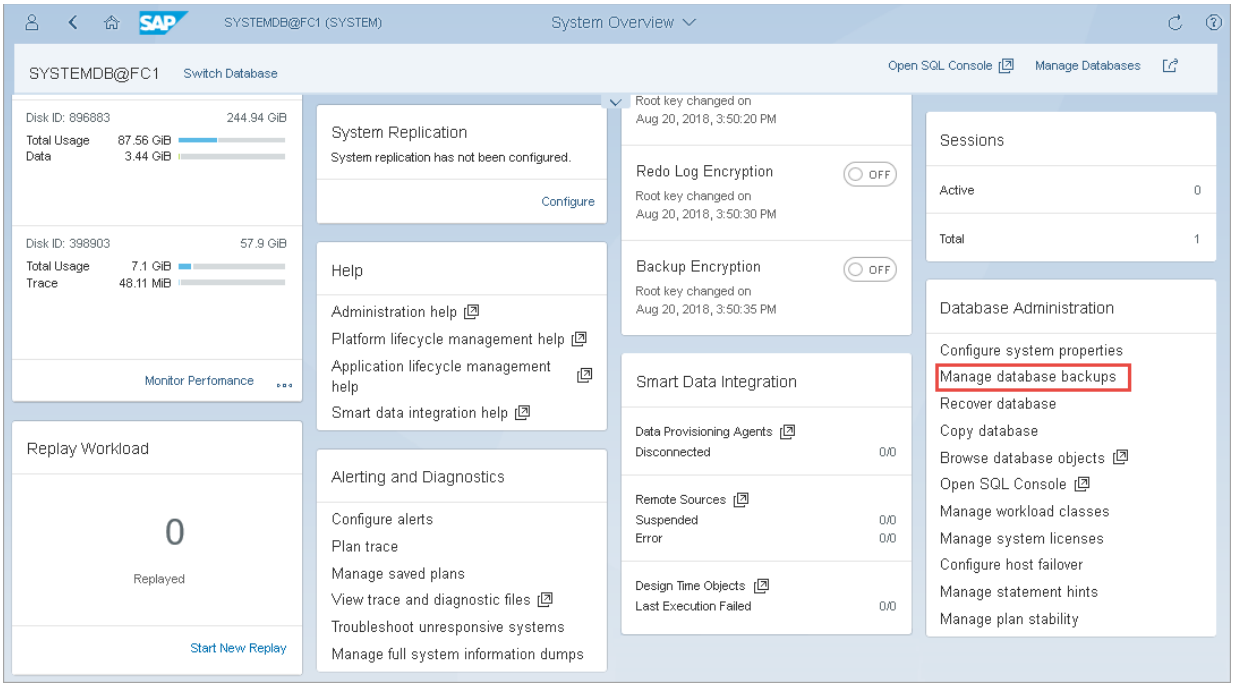

2. At the **Backup Catalog** section, click **Create Backup**.

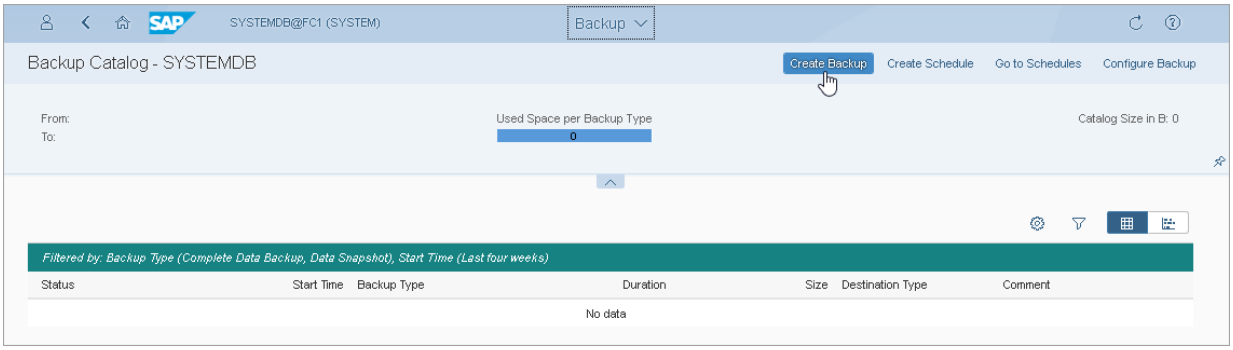

3. Specify backup settings:

   a. Select the required backup type:

   - **Complete Data Backup**: backup of all data structures required to recover the database.
   - **Incremental Data Backup**: backup of data changed since the last full data backup or the last delta backup (incremental or differential).
   - **Differential Data Backup**: backup of data changed since the last full data backup.

b. In the **Destination Type** setting, select *Backint*. With this option selected, Veeam Plug-in will transfer the backup to Veeam backup repository.

c. Change the default prefix for the backup file, if needed.

d. To start the backup, click **Back Up**.

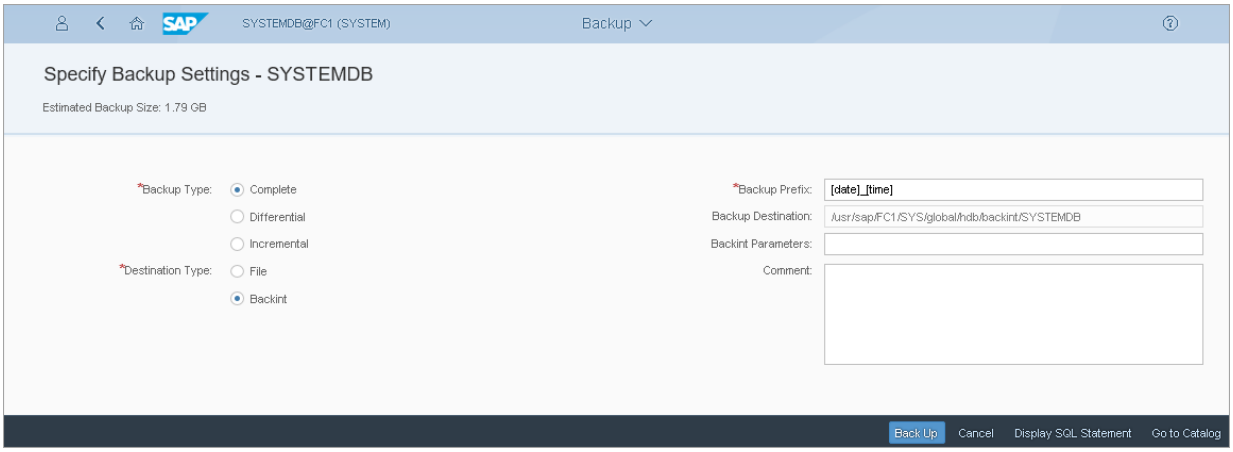

Backup Job in Veeam Backup & Replication Console

After you start a backup process with SAP HANA Backint, Veeam Backup & Replication creates a backup job. You can use this job to view the statistics on the backup process, generate backup job reports or you can also disable the backup job. You cannot launch or edit SAP HANA backup jobs in the Veeam Backup & Replication console. You can manage backup operations only on the SAP HANA side using SAP HANA Studio, SAP HANA Cockpit or HDBSQL.

Mind the following regarding the naming of SAP HANA backup jobs:

- For a standalone SAP HANA server (scale-up), the backup job name is generated automatically based on names of the SAP HANA server and selected repository.
- For a scale-out SAP HANA cluster: When you run the Veeam Plug-in configuration wizard for the first time in one of the SAP HANA cluster nodes, the wizard asks for the cluster name. The cluster name will be used in the backup job name along with the repository name.

> **NOTE**
>
> Due to specifics of the SAP HANA backup process, the progress bar of a running SAP HANA backup job is not available.

To view details of a backup job process, do the following.

1. Open the Veeam Backup & Replication console.
2. In the **Home** view, expand the **Jobs** node and click **Backup**.
3. In the list of jobs, select the SAP HANA backup job to see details of the current backup process or the last backup job session.

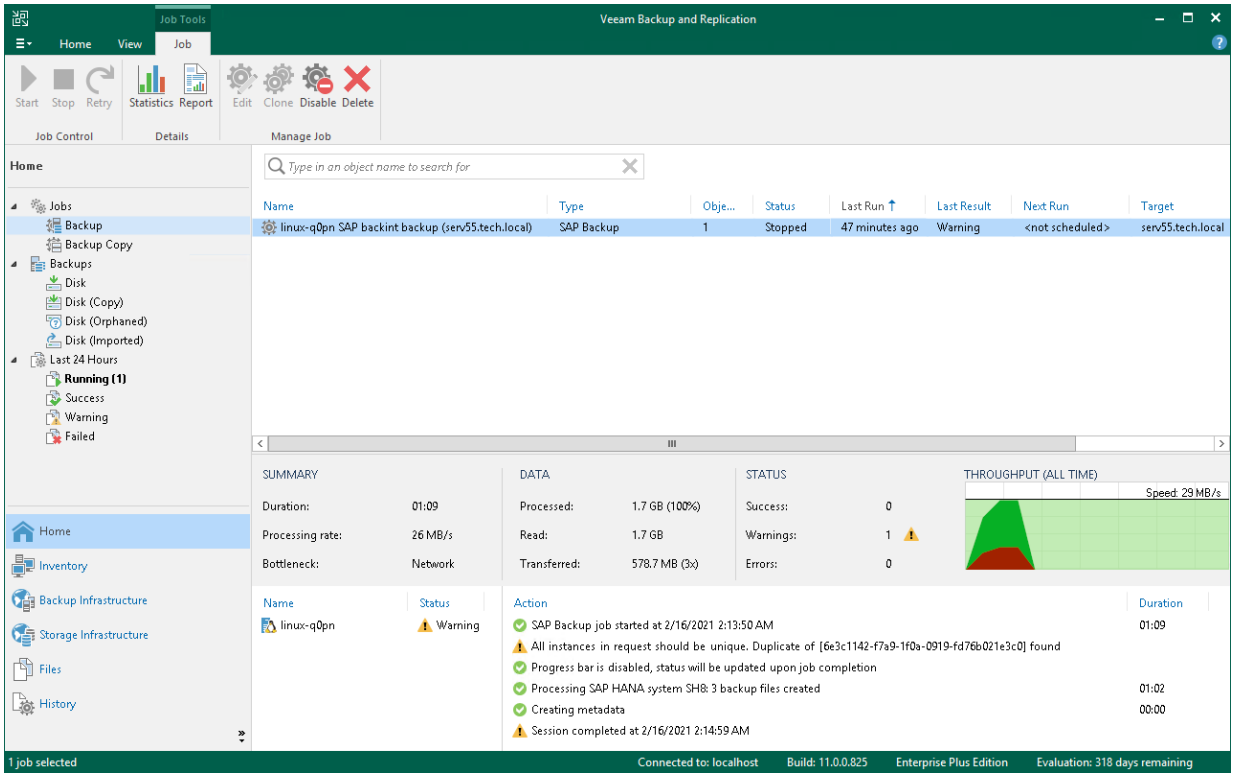

Generating Backup Job Reports

Veeam Backup & Replication can generate reports with details about an SAP HANA backup job session performance. The session report contains the following session statistics: session duration details, details of the session performance, amount of read, processed and transferred data, backup size, compression ratio, list of warnings and errors (if any).

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

Disabling Backup Job

You can disable SAP HANA backup jobs in the Veeam Backup & Replication console. If you disable the job, you will not be able to run SAP Backint backup commands on the SAP HANA server.

To disable a backup job:

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

Database Recovery

With the configured Veeam Plug-in you can restore SAP HANA databases from the backups that reside on the Veeam backup repository. All restore operations are performed on the SAP HANA side. To restore databases, you can use SAP HANA Cockpit, SAP HANA Studio, or HDBSQL.

Keep in mind that examples provided in this section are for demonstrating purposes only. To see the full restore functionality of SAP HANA tools, see the SAP HANA Recovery section of SAP HANA Administration Guide.

To learn how to recover SAP HANA databases from backups stored on Veeam repositories, see:

- Recovering Databases Using SQL Commands
- Recovering Databases with SAP HANA Studio
- Recovering SYSTEMDB with SAP HANA Cockpit
- Recovering Tenant Databases with SAP HANA Cockpit
- Recovering Databases to Other Servers

Restoring Databases (HDBSQL Commands)

You can use HDBSQL to restore SAP HANA databases from backups stored on Veeam backup repositories. For details on the HDBSQL restore, see the RECOVER DATABASE Statement section of the SAP HANA SQL and System Views Reference.

To recover SAP HANA databases from backups stored on Veeam backup repositories, do the following:

1. Log in to SAP HANA HDBSQL as the HDB administrator. Use HDBUSERSTORE to securely store connection details on a client machine. For details, see the Secure User Store section of the SAP HANA Security Guide.

```
sh4adm@linux-q0pn:/usr/sap/SH4/HDB01> hdbuserstore SET <key> hostname:3001
3@SID <username> <password>
sh4adm@linux-q0pn:/usr/sap/SH4/HDB01> hdbsql -U <key>
```

2. Recover a tenant database to the latest state using Backint. As the timestamp, specify the current data and time or future date and time.

```
alter stop database <DATABASE_NAME>;
recover database for <DATABASE_NAME> until timestamp '2020-01-01 12:00:00'
using catalog backint;
```

Restoring Databases (SAP HANA Studio)

You can restore SAP HANA databases from the Veeam Plug-in backups using SAP HANA Studio.

The example below is provided for demonstration purposes only. For details on the full restore functionality of SAP HANA tools, see the Recovering an SAP Database section of the SAP HANA Administration Guide.

To perform a Backint recovery from Veeam Plug-in backups, do the following:

1. Log in to SYSTEMDB using as a user with DATABASE ADMIN privileges.
2. Right-click the SYSTEMDB database.
3. Click **Backup and Recovery** and select **Recover System Database** or **Recover Tenant Database**.

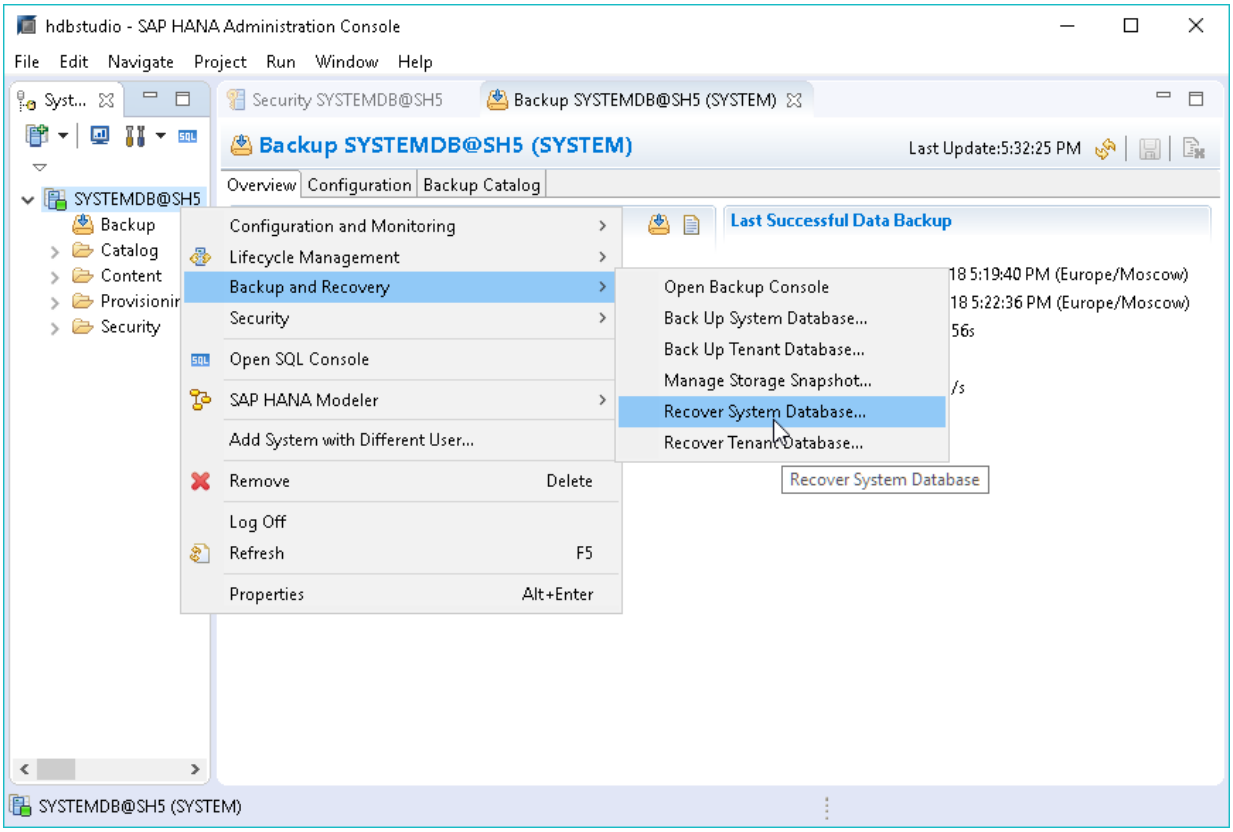

4. Enter the operating system user credentials.

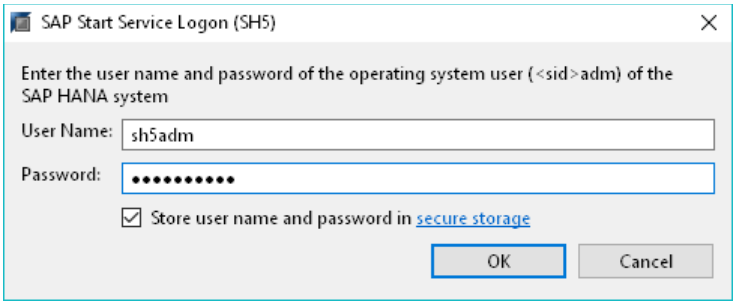

5. Recovery process requires the database to be shut down. In the pop-up window, click **OK** to confirm the database shutdown.

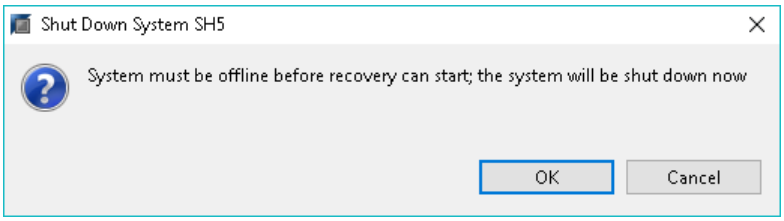

6. At the **Specify Recovery Type** step of the recovery wizard, select the required restore point or the option to restore the database to the most recent state.

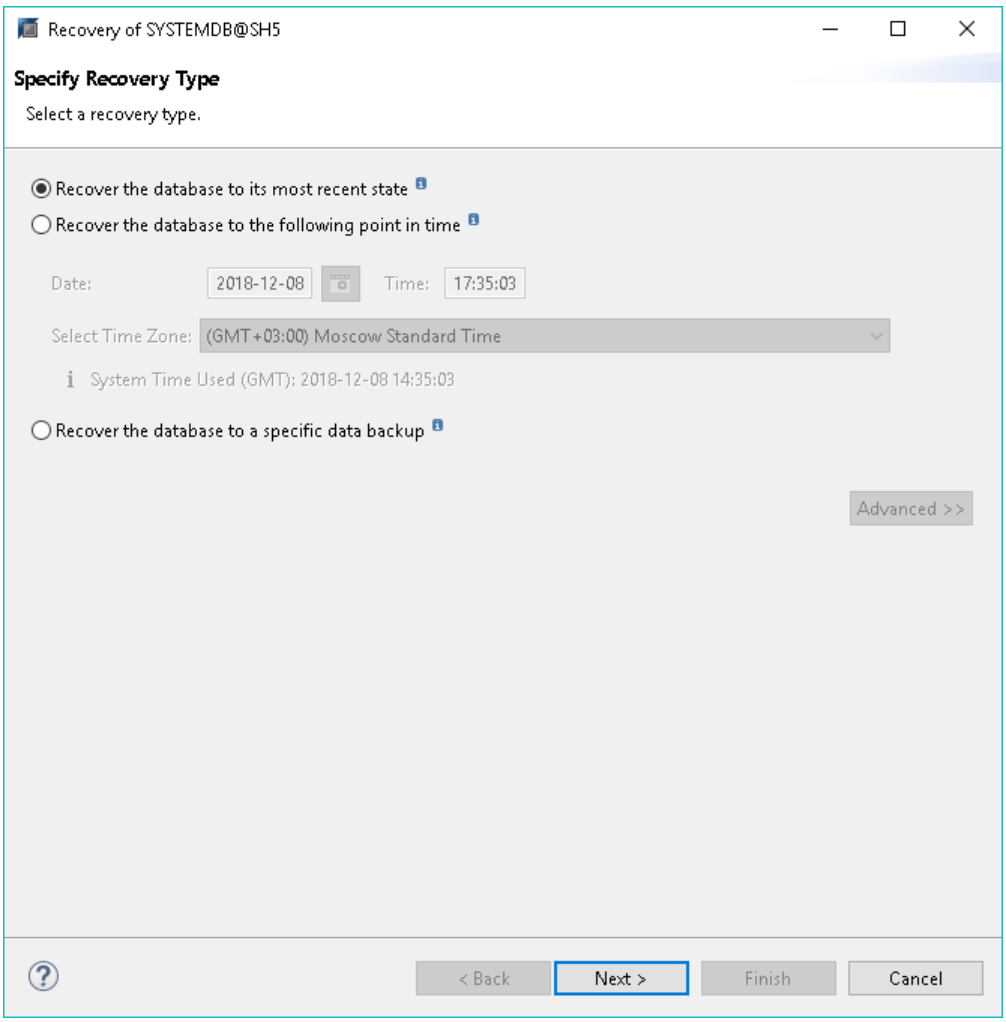

7. At the **Locate Backup Catalog** step of the wizard, select one of the following, depending on where your backup catalogs reside:
   - **Search for the backup catalog in the file system only**.

- Search for the backup catalog in Backint only.

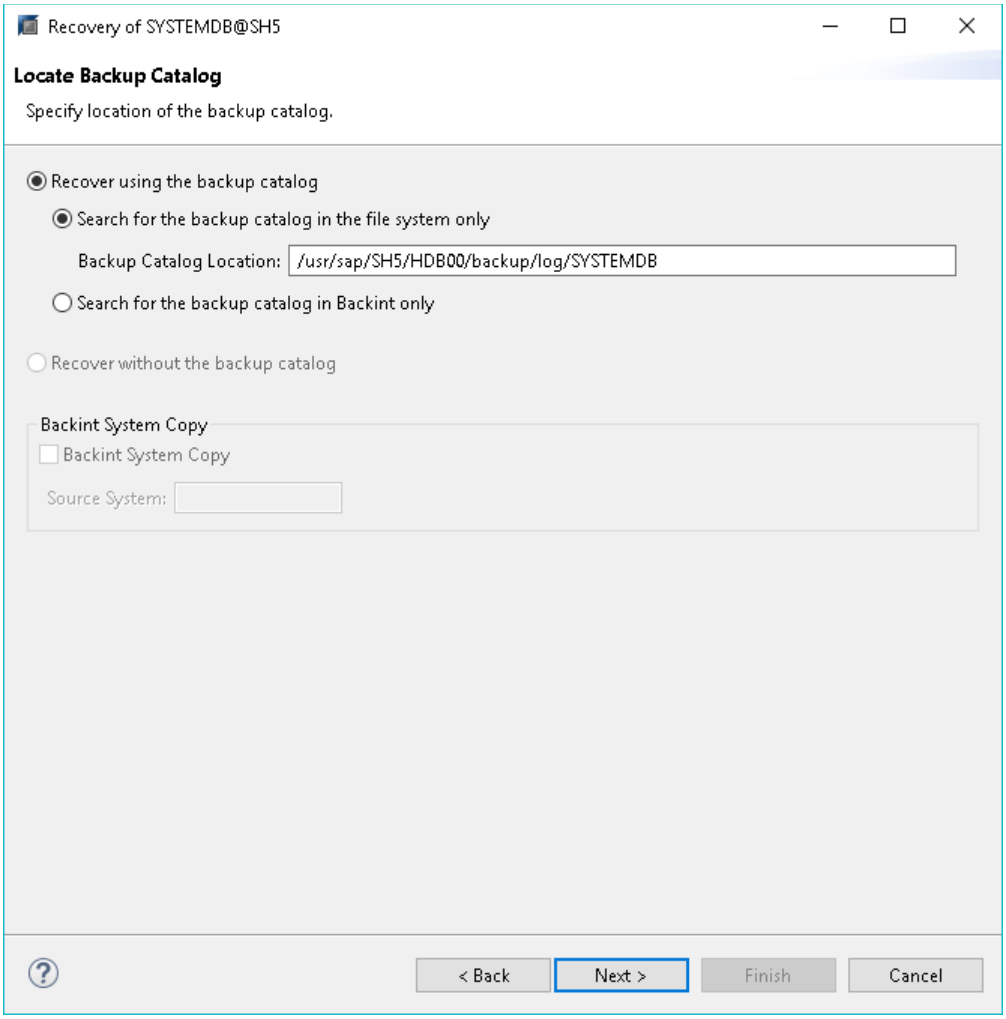

8. Select the required backup to restore.

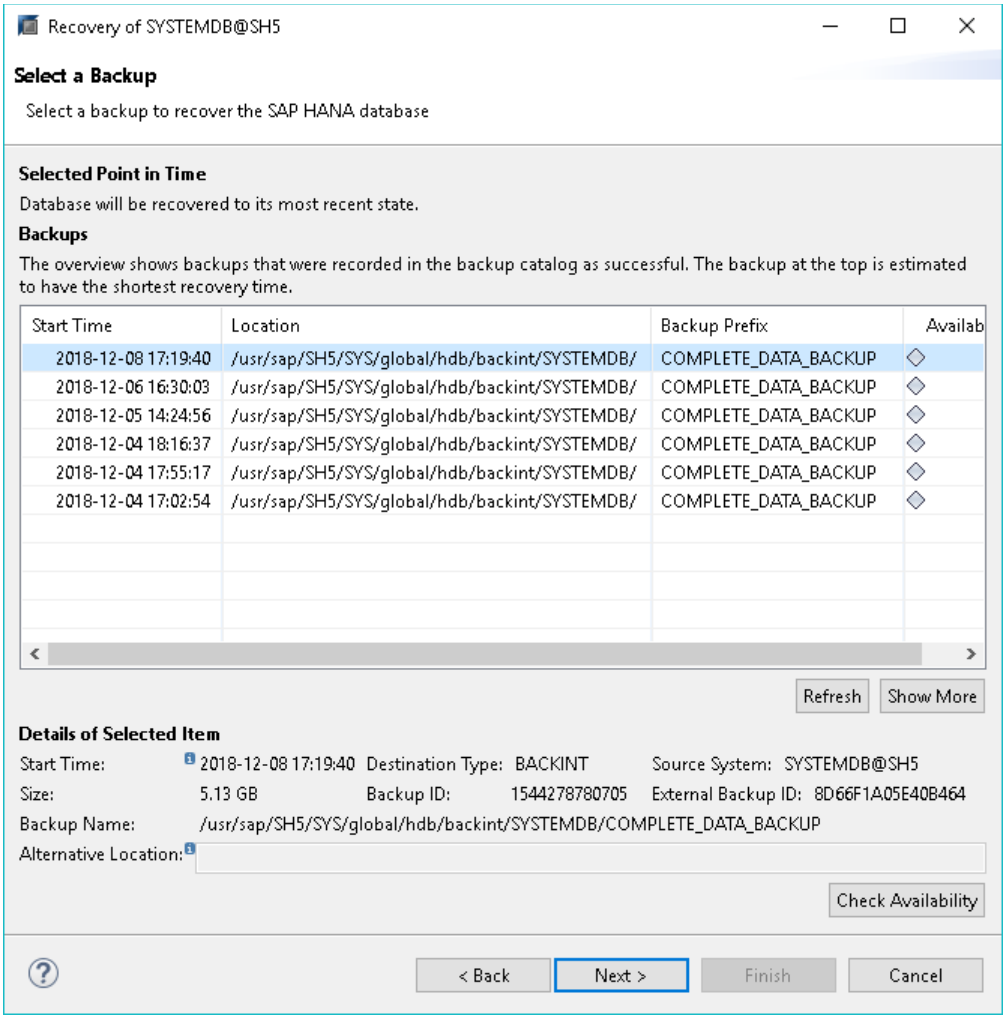

9. At the **Locate Log Backups** step, click **Next**.

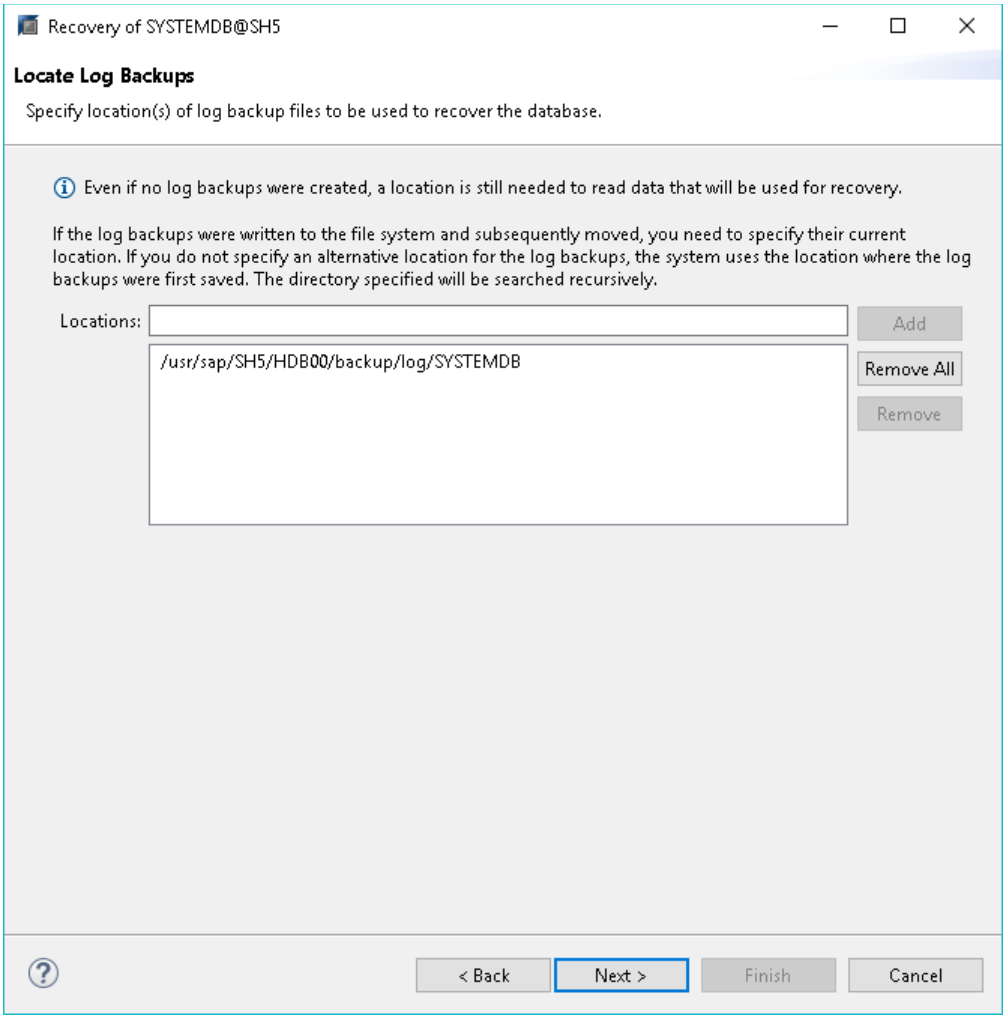

10. At the **Other Settings** step:

    a. Switch on the availability check for the Backint backups.

    b. If you are recovering the database to a database with a new SID or landscape ID, select the **Install New License Key** check box and specify the path to the license file.

c. Click **Next**.

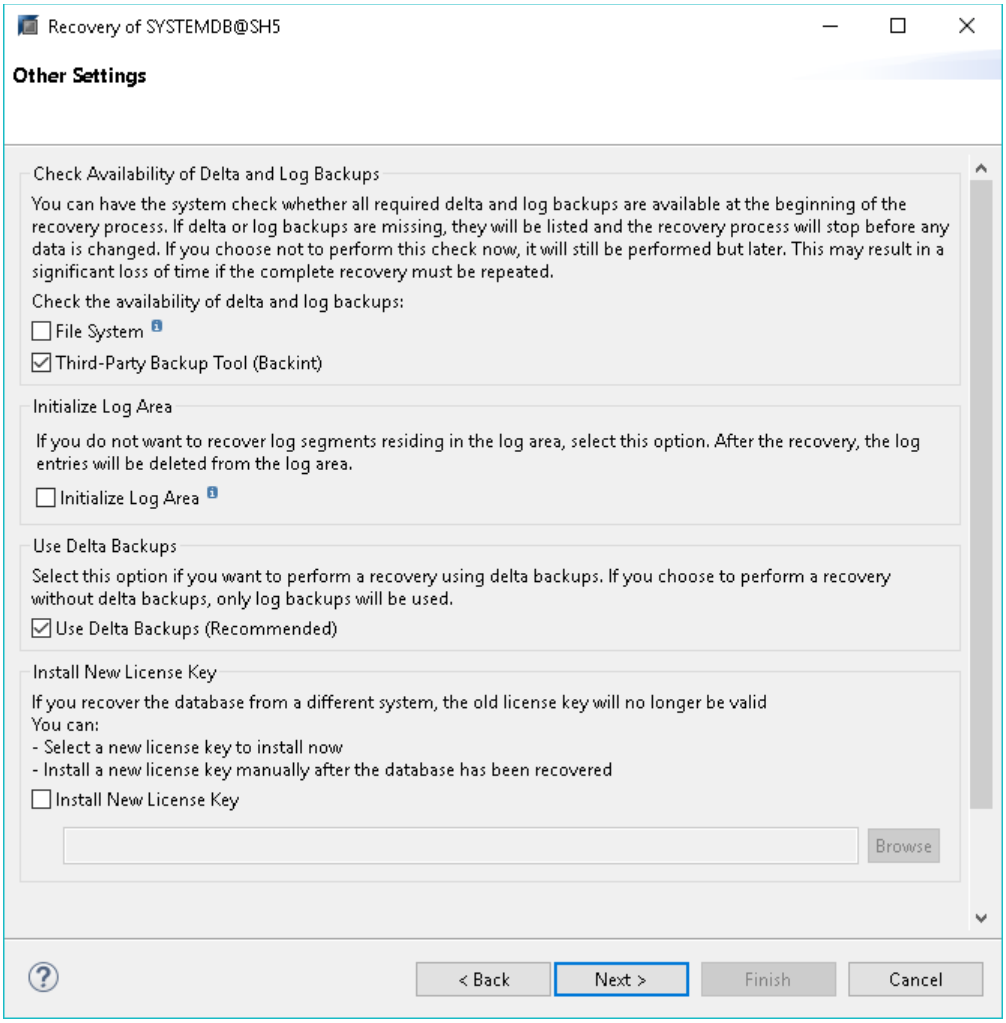
Recovery of SYSTEMDB@SH5
Other Settings
Check Availability of Delta and Log Backups
You can have the system check whether all required delta and log backups are available at the beginning of the recovery process. If delta or log backups are missing, they will be listed and the recovery process will stop before any data is changed. If you choose not to perform this check now, it will still be performed but later. This may result in a significant loss of time if the complete recovery must be repeated.
Check the availability of delta and log backups:
File System
Third-Party Backup Tool (Backint)
Initialize Log Area
If you do not want to recover log segments residing in the log area, select this option. After the recovery, the log entries will be deleted from the log area.
Initialize Log Area
Use Delta Backups
Select this option if you want to perform a recovery using delta backups. If you choose to perform a recovery without delta backups, only log backups will be used.
Use Delta Backups (Recommended)
Install New License Key
If you recover the database from a different system, the old license key will no longer be valid
You can:
- Select a new license key to install now
- Install a new license key manually after the database has been recovered
Install New License Key
Browse
< Back
Next >
Finish
Cancel

11. At the **Review Recovery Settings** step, click **Finish**.

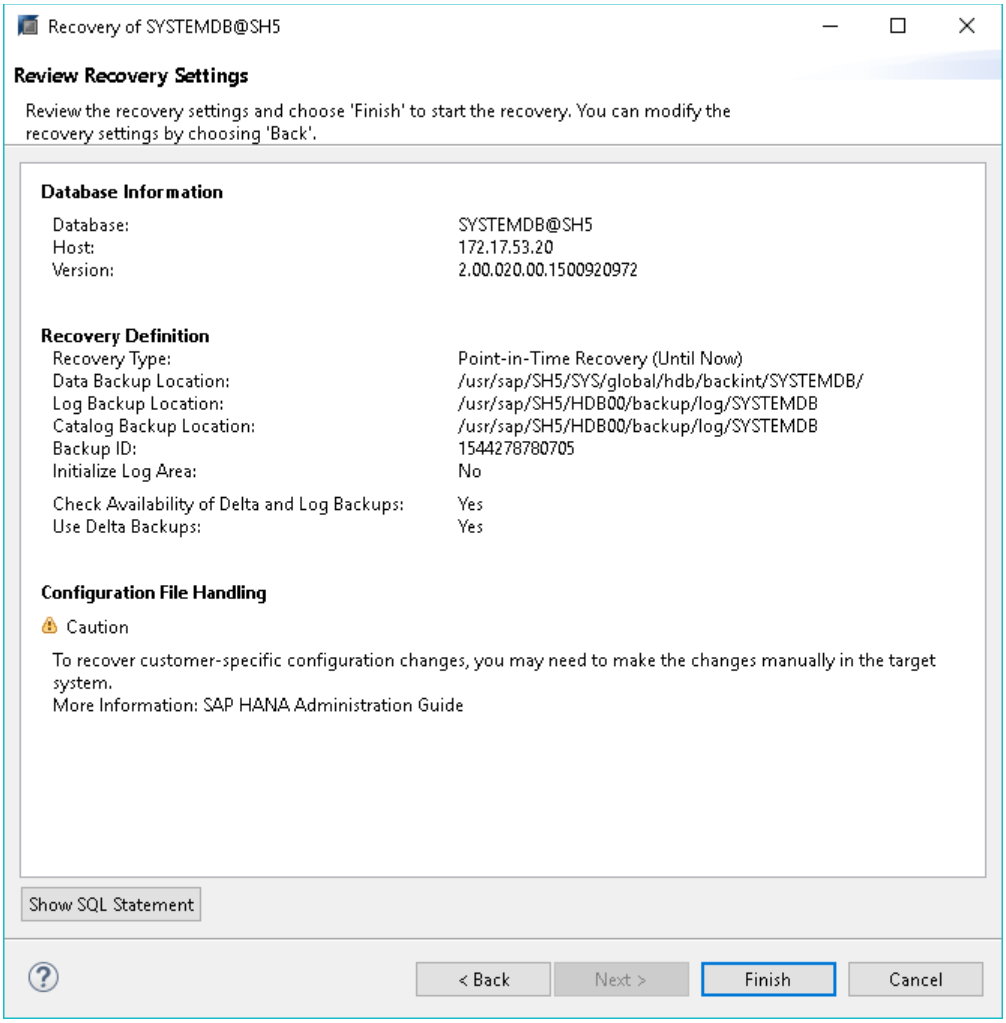

Restoring SYSTEMDB (SAP HANA Cockpit)

You can restore SAP HANA SYSTEMDB databases from the Veeam Plug-in backups using SAP HANA Cockpit.

The example below is provided for demonstration purposes only. For details on the full restore functionality of SAP HANA Cockpit, see the Recovering an SAP HANA Database section of the SAP HANA Administration Guide.

Before You Begin

Before you start the recovery, shut down the database that you want to recover:

1. In the SAP HANA Cockpit console, locate the database that you want to recover.
2. In the **Overall Database Status** block, click **Stop System**.

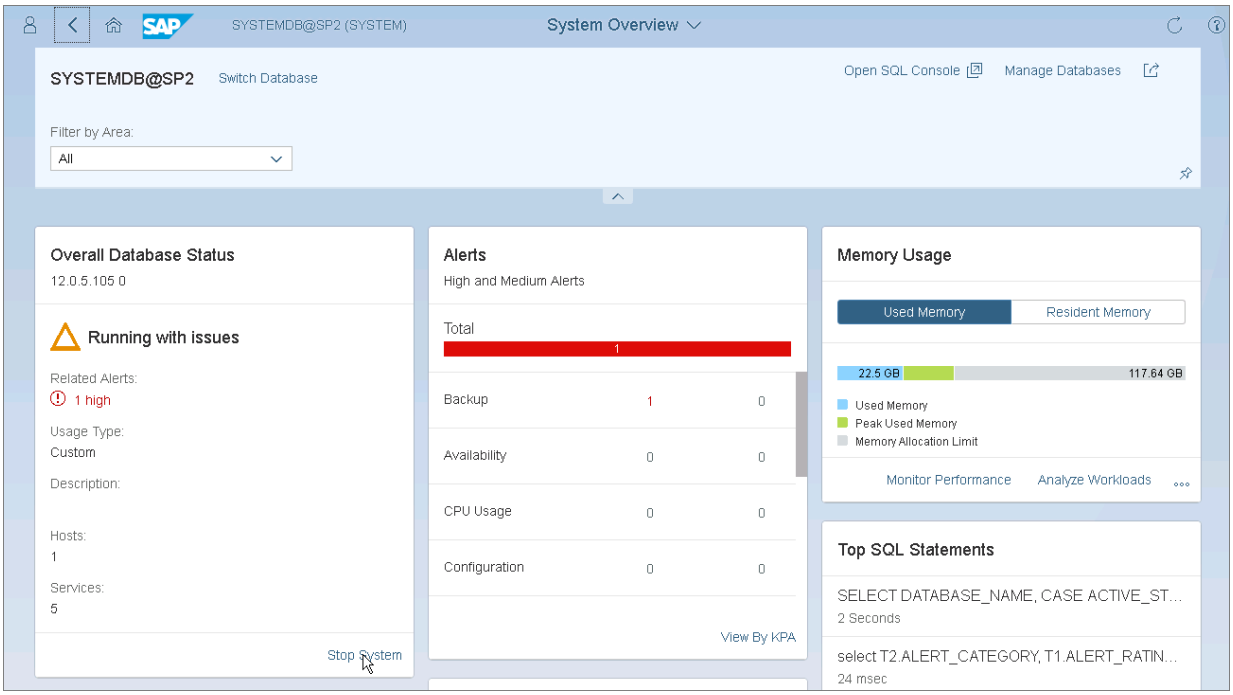

3. At the **Manage Services** section, click **Stop System** and select the **Softly** option to shut down the database after SAP HANA finishes running statements.

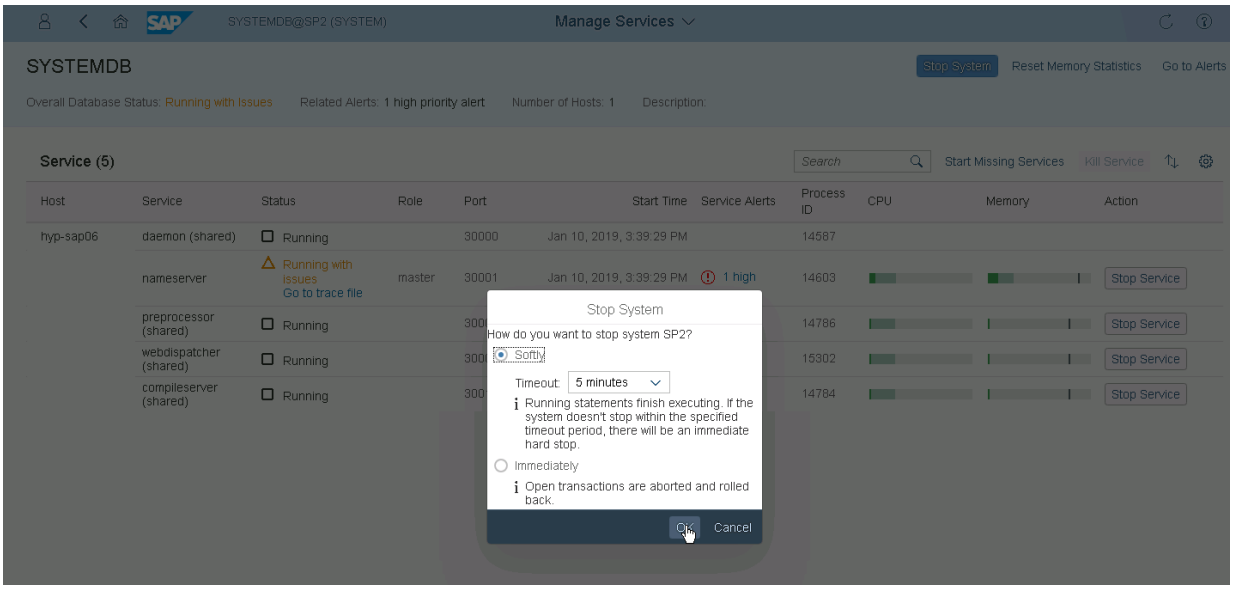

Performing Recovery

To perform a Backint recovery of SYSTEMDB from a Veeam Plug-in backup, do the following:

1. In the **System Overview** block, go to the **Database Administration** section and click **Recover database**.

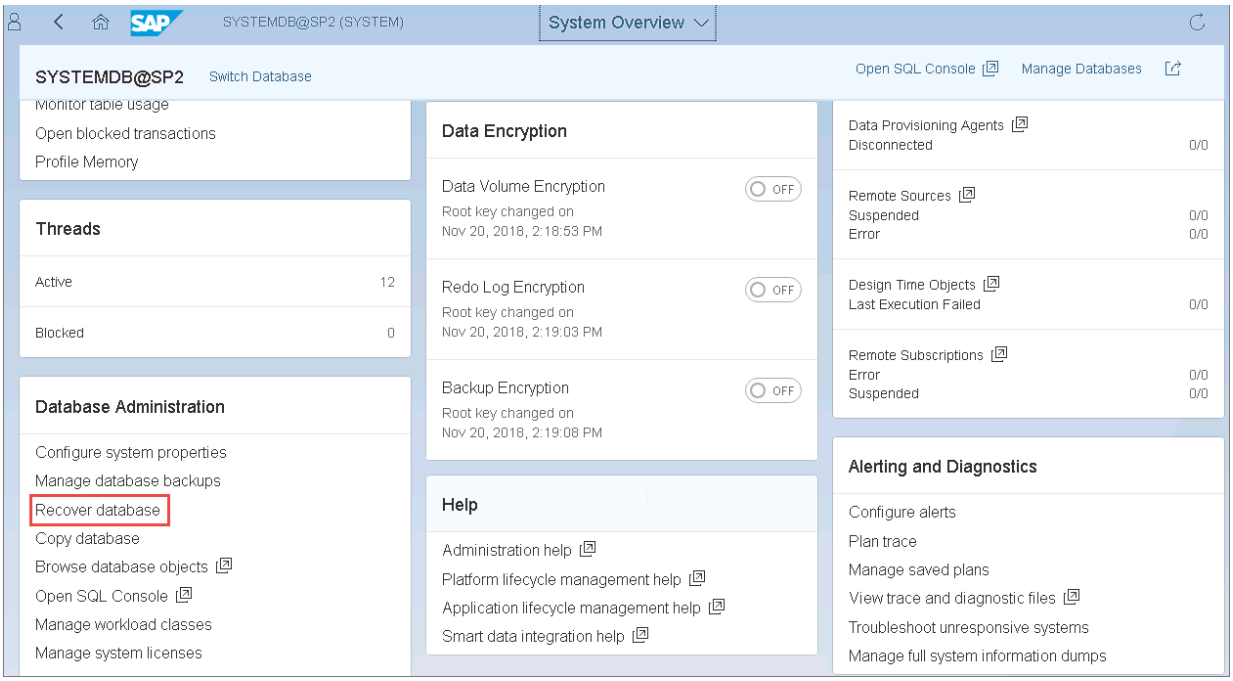

2. At the **Recovery Target** step, select the required restore point or the option to restore the database to the most recent state. Then, click **Step 2**.

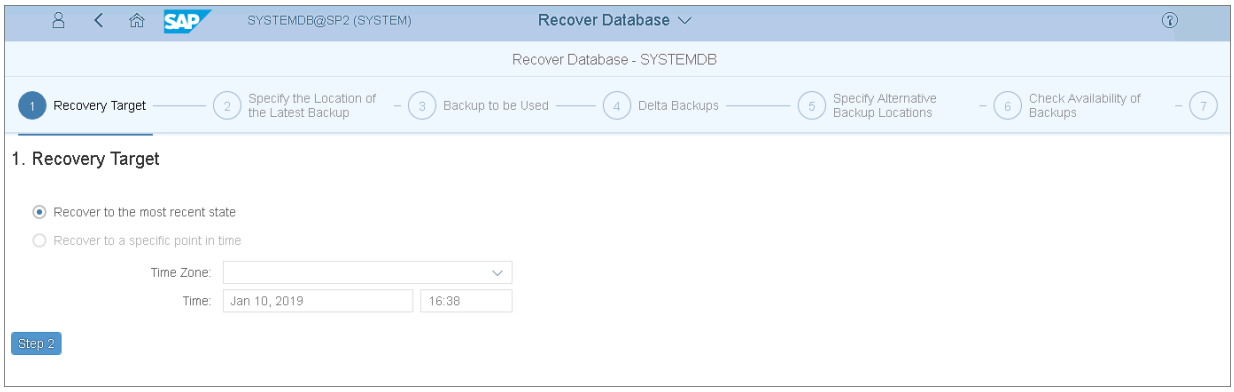

3. Specify the location of the latest backup catalog and click **Step 3**.

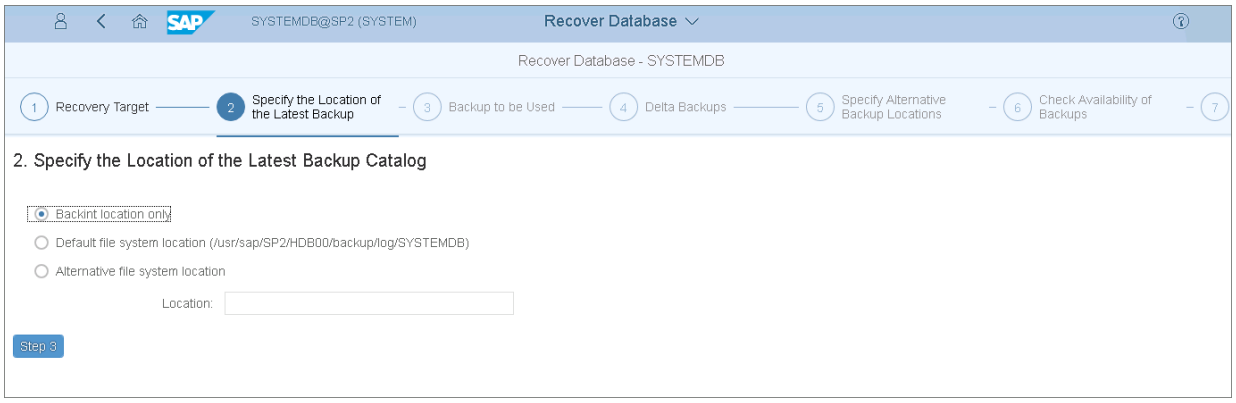

4. At the **Backup to be Used** step, select the backup and click **Step 4**.

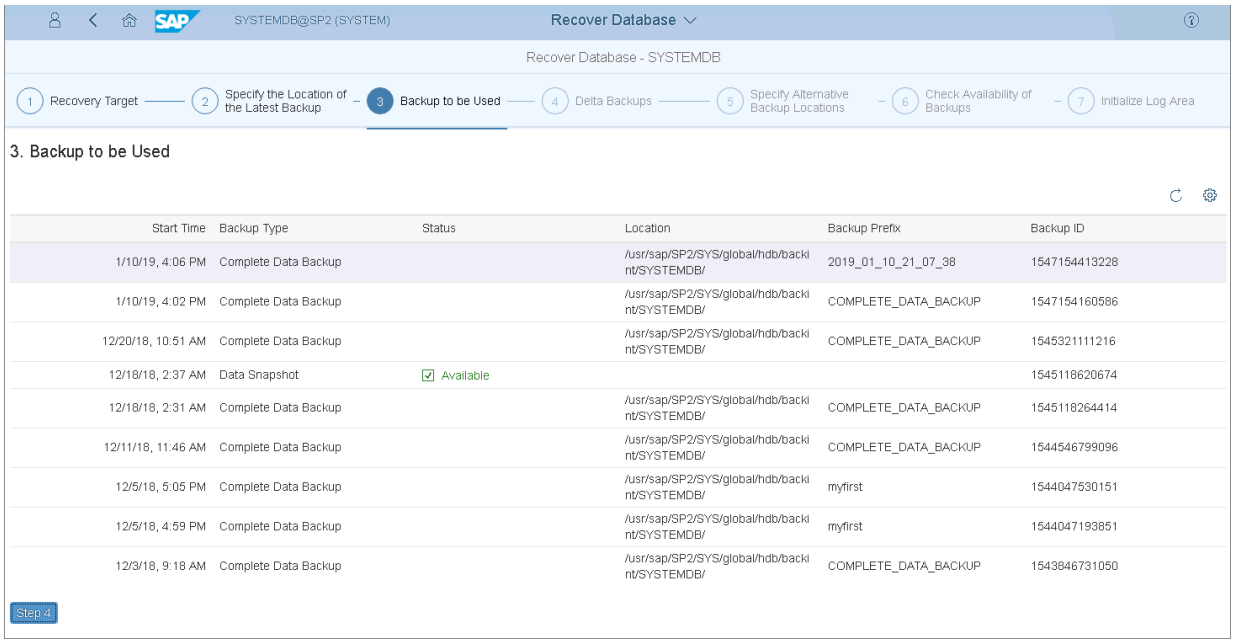

5. At the **Delta Backups** step, select **Yes** to use delta backups.

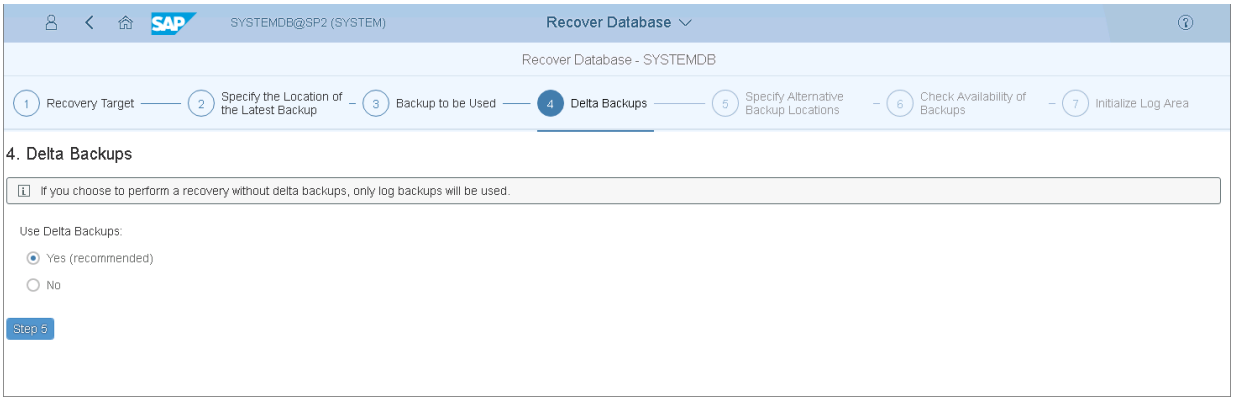

6. At the **Specify Alternative Backup Locations** step, if you want to use backups that are not in the backup catalog, specify their locations. You can also change the location for log backups.

   If you you leave the fields empty, SAP HANA will use the locations specified in the backup catalog.

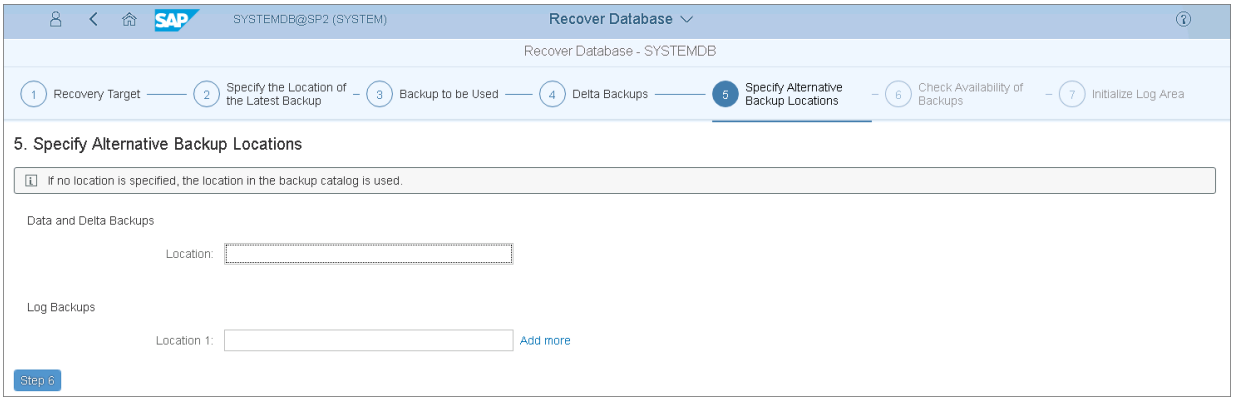

7. At the **Check Availability of Backups** step, select **Yes** or **No** options, to check if the backups are available. Note that at this step SAP HANA does not check the integrity of the backup content on the block level.

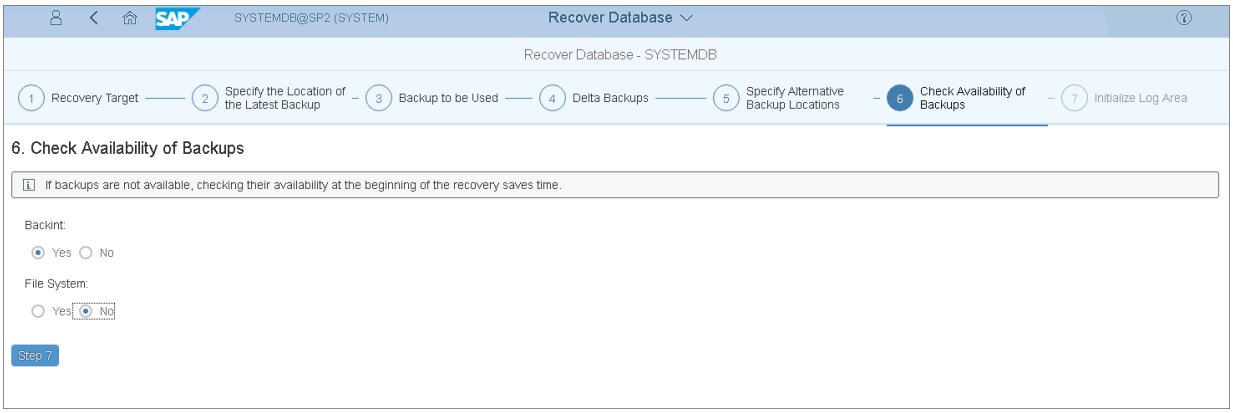

8. At the **Initialize Log Area** step, select **No** to initialize the log area and click **Review**. You must initialize the log are only if the log area is unavailable or if you are recovering the database to a different system.

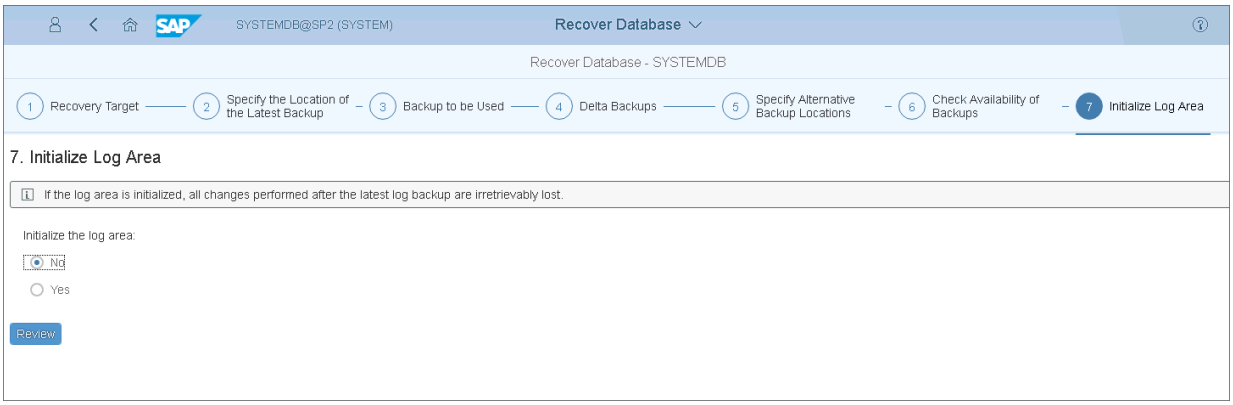

9. Review the recovery options and click **Start Recovery**.

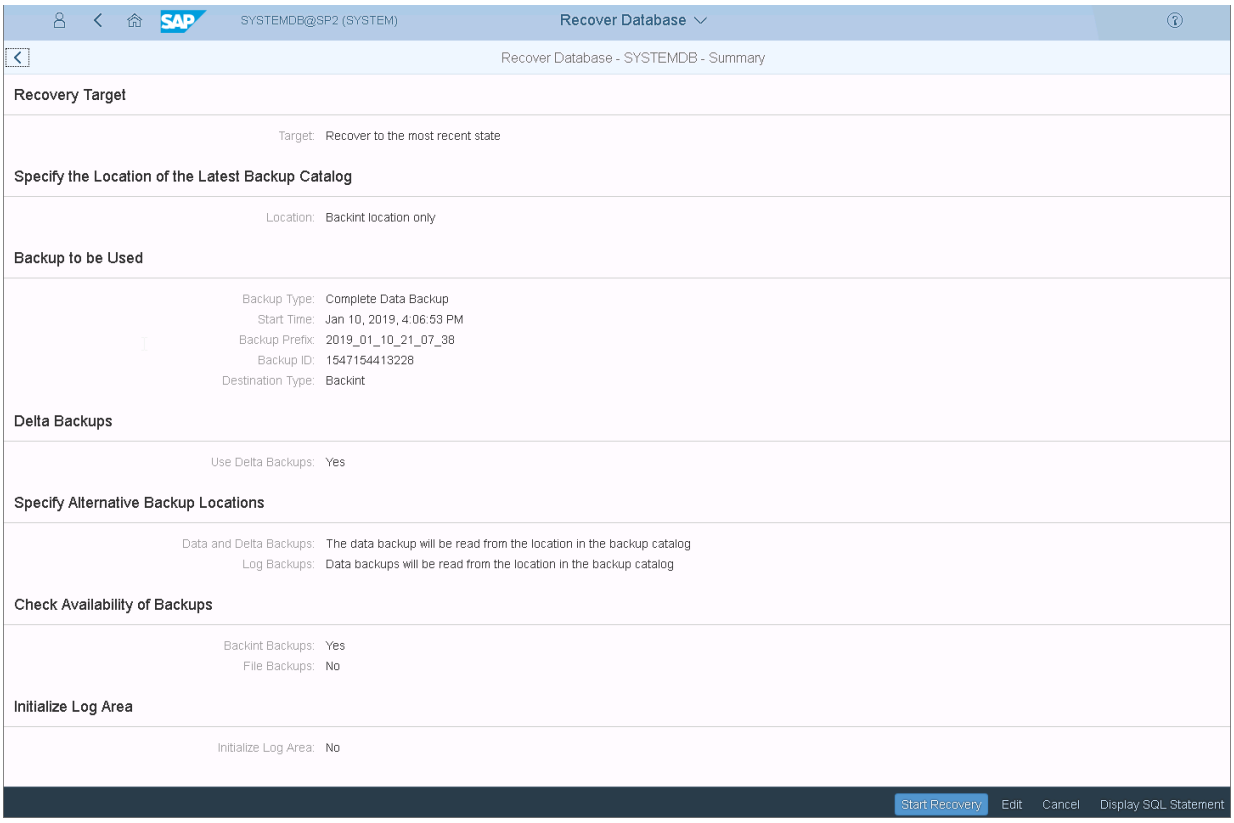

Recovering Tenant Databases with SAP HANA Cockpit

You can restore SAP HANA tenant databases from the Veeam Plug-in backups using SAP HANA Cockpit.

The example below is provided for demonstration purposes only. For details on the full restore functionality of SAP HANA tools, see the Recovering an SAP HANA Database of the SAP HANA Administration Guide.

To perform a Backint recovery of an SAP HANA tenant database from a Veeam Plug-in backup, do the following:

1. In the **System Overview** page of the required system, click **Manage Databases**.

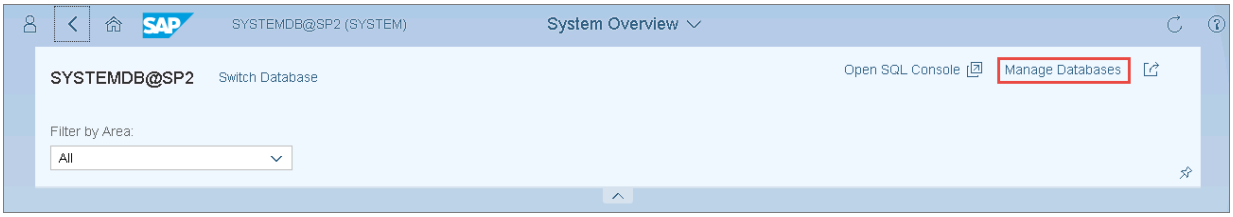

2. In the **Manage Databases** page, expand the toolbar options and select **Recover Tenant**.

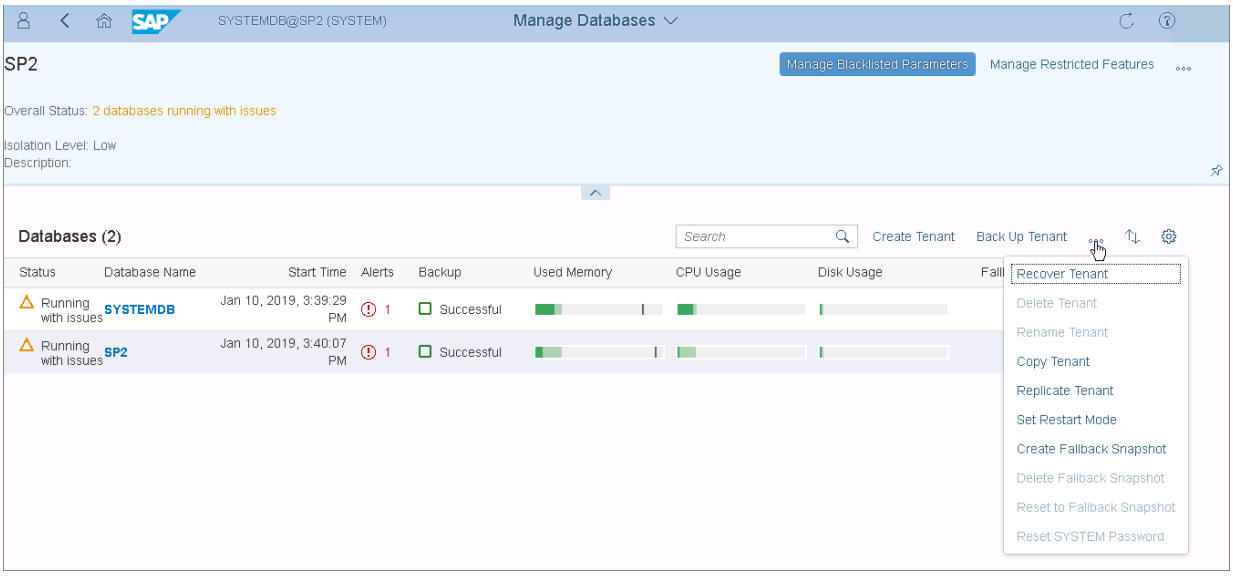

3. After you launch the recovery wizard, SAP HANA will issue the warning that the database must be stopped for recovery. Click **Stop Tenant** in the warning window.

4. At the **Recovery Target** step of the wizard, select the required restore point and click **Step 2**.

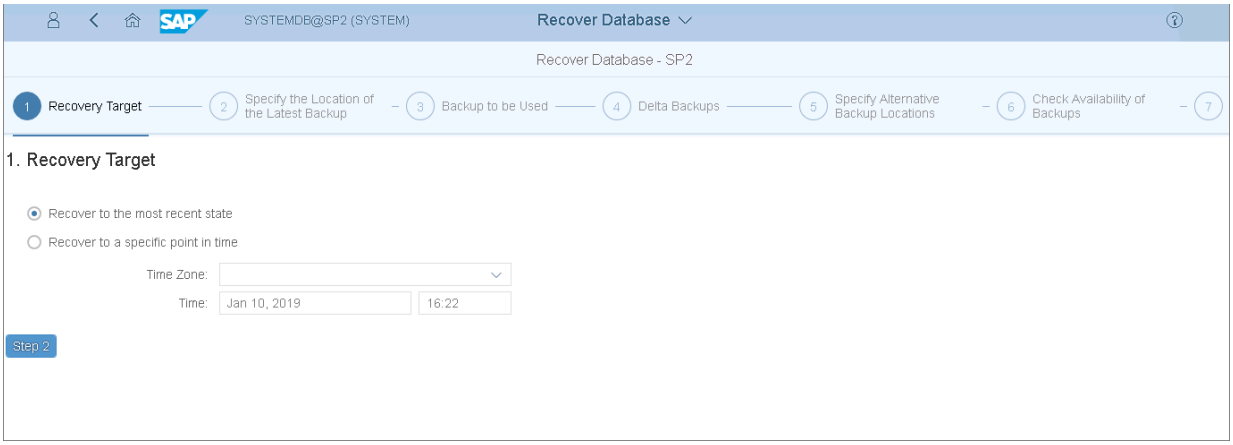

5. Specify the location of the latest backup catalog and click **Step 3**.

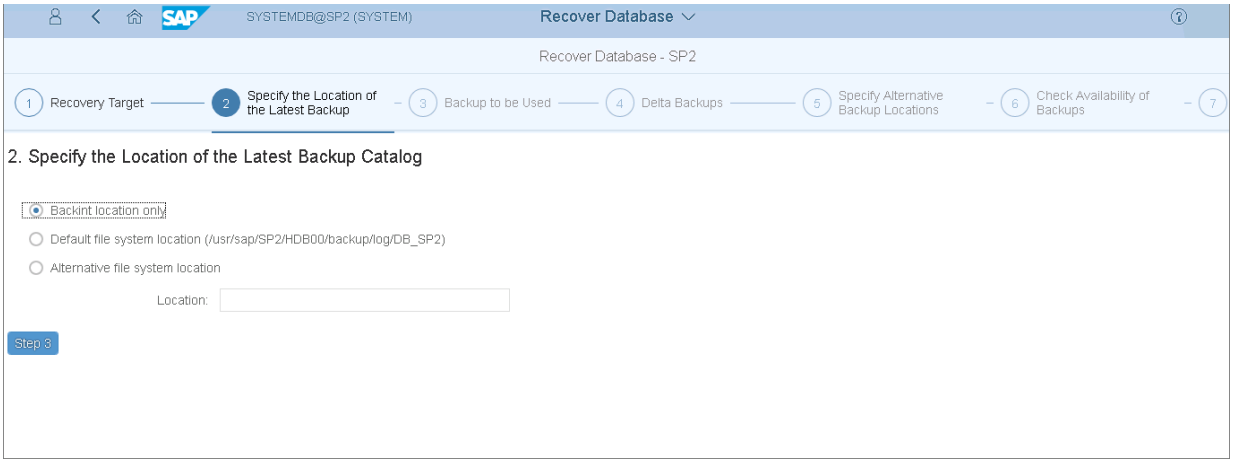

6. At the **Backup to be Used** step, select the backup and click **Step 4**.

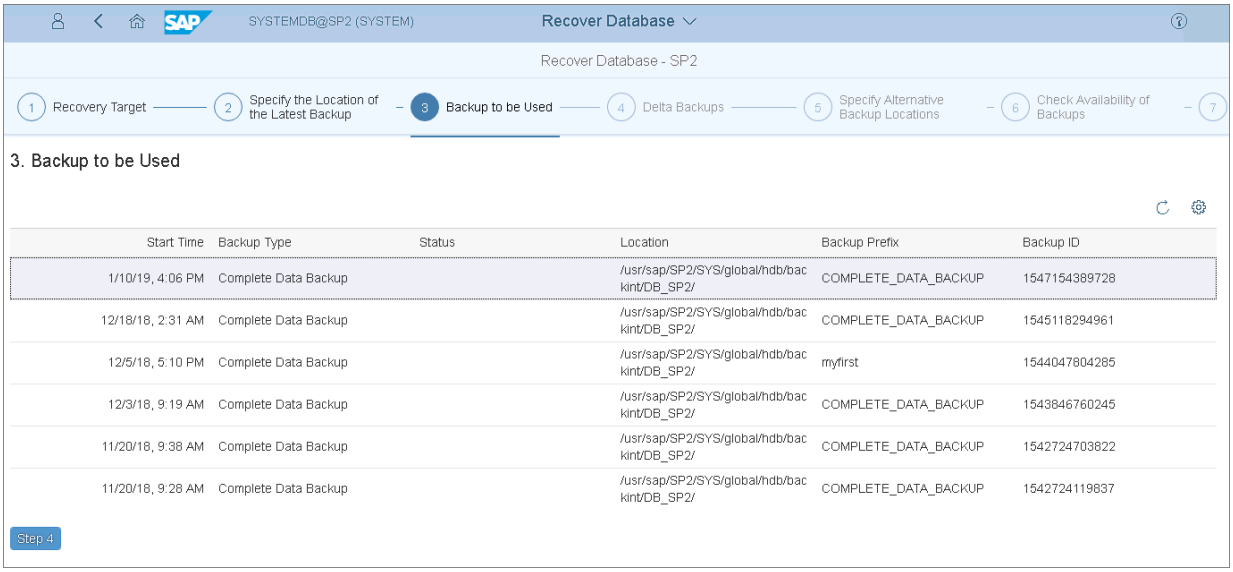

7. At the **Delta Backups** step, select **Yes** to use delta backups.

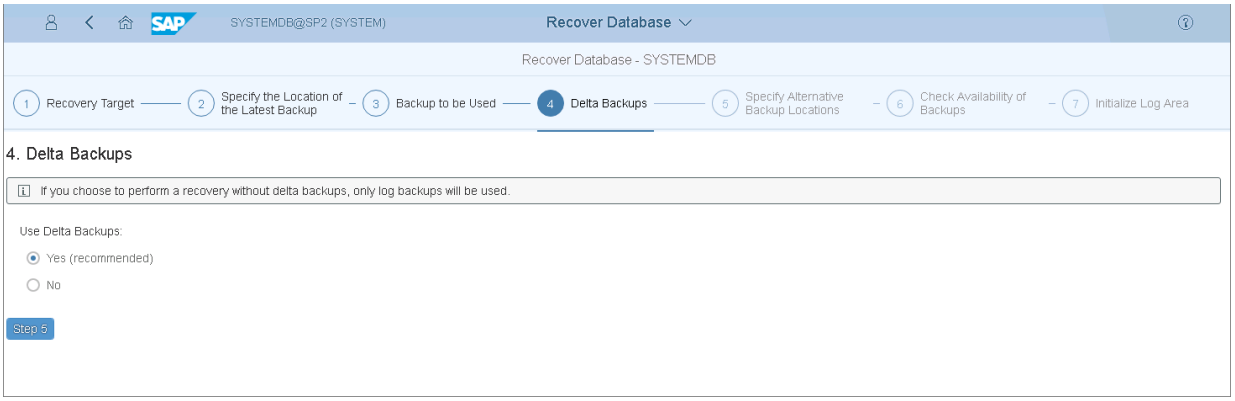

8. At the **Specify Alternative Backup Locations** step, if you want to use backups that are not included in the backup catalog, specify their locations. You can also change the location for log backups.

   If you leave the fields empty, SAP HANA will use the locations specified in the backup catalog.

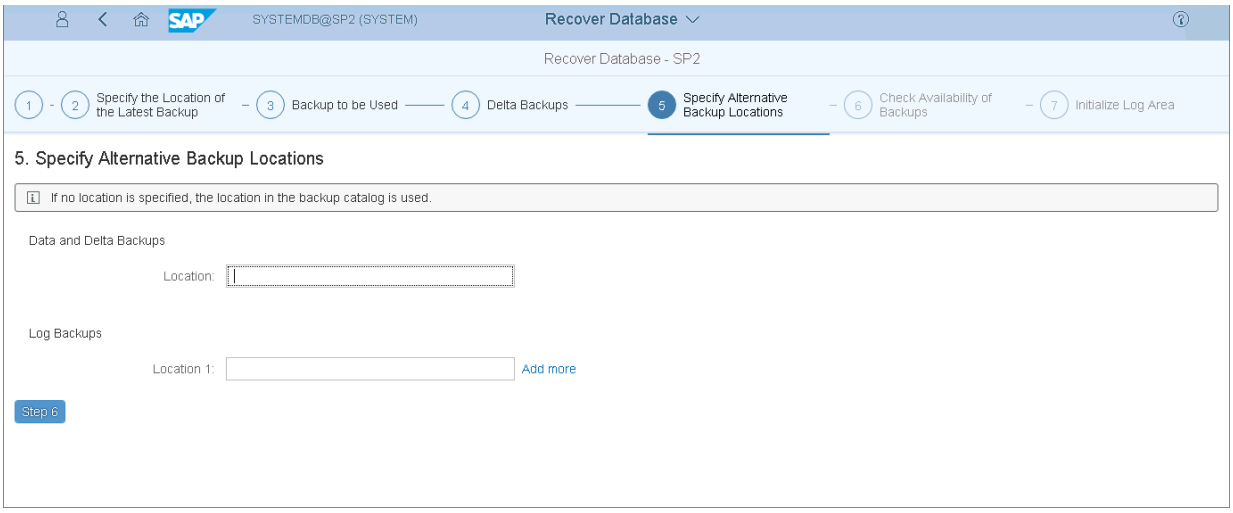

9. Select **Yes** or **No**, to check if the backups are available. Note that at this stage SAP HANA does not check the integrity of the backup content on the block level.

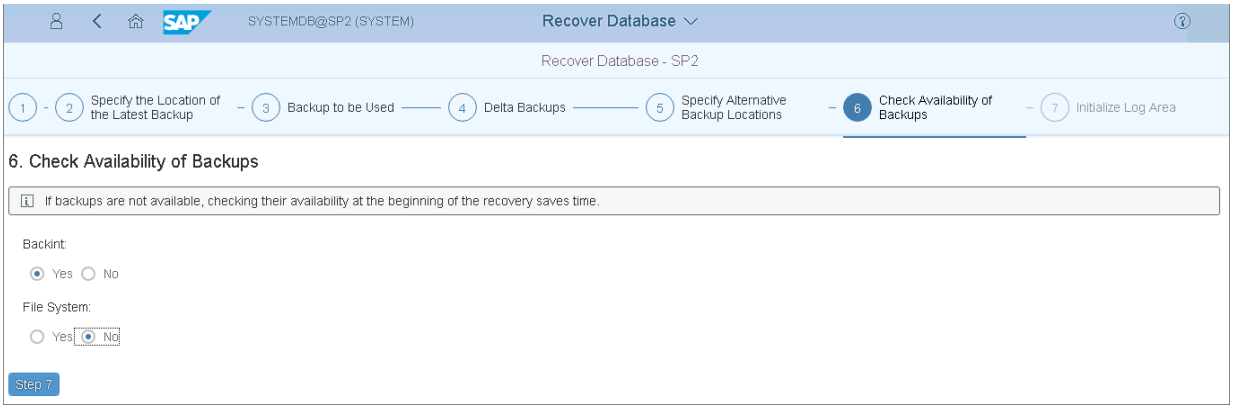

10. Select **No** to initialize the log area and click **Review**. You must initialize the log area only if the log area is unavailable or if you are recovering the database to a different system.

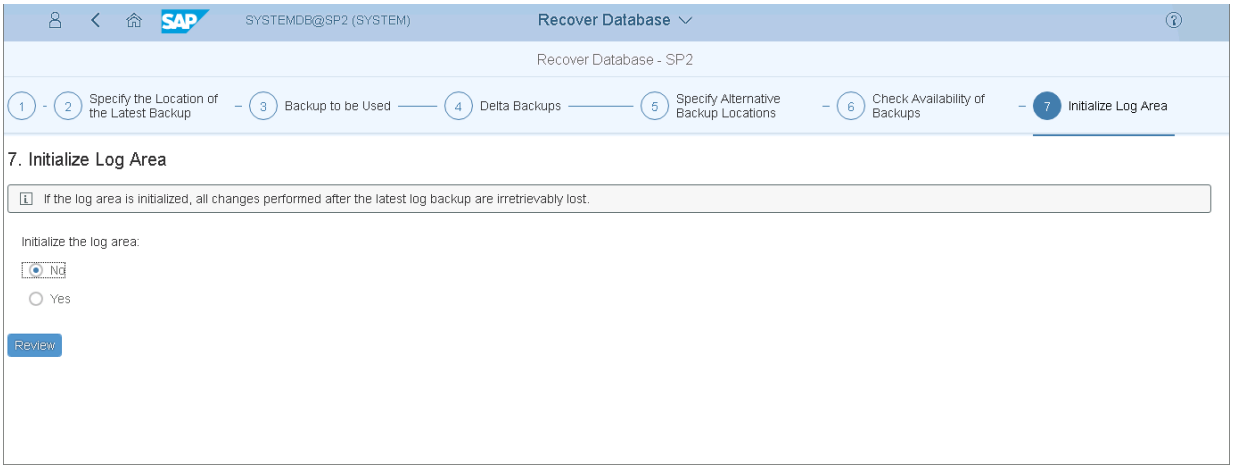

11. Review the recovery options and click **Start Recovery**.

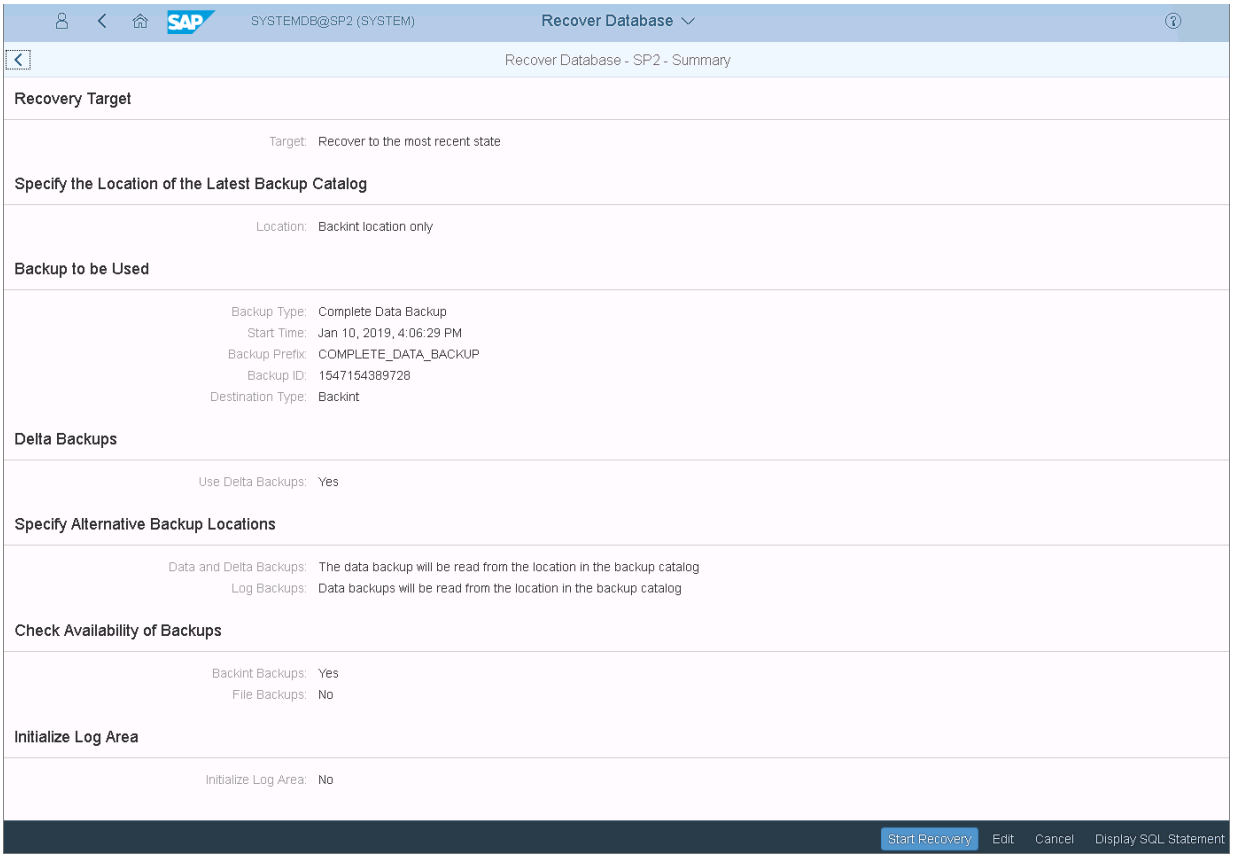

Recovering Databases to Other Servers (System Copy)

You can restore SAP HANA databases from Veeam Plug-in backups to another server. To restore databases to another server, you must reconfigure settings of Veeam Plug-in as shown below.

For security reasons, you can restore databases to another server only in the following condition. The account you use to connect to Veeam Backup & Replication server must be the same account that performed the backup of the source system. If you want to use another account, you can assign the **Veeam Backup Administrator** or **Veeam Restore Operator** roles to the required account. For details on assigning Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication User Guide.

Procedure

To restore databases to another server, you must reconfigure settings of Veeam Plug-in as shown below:

1. Go to `/opt/veeam/VeeamPluginforSAPHANA` and run the following command to select the source server whose backups you want to use during restore.

```
VM2ADM:/opt/veeam/VeeamPluginforSAPHANA> SapBackintConfigTool --set-restor
e-server
Select source SAP HANA plug-in server to be used for system copy restore:
1. SAP-VM1
2. SAP-VM02
Enter server number: 1
```

2. Specify a backup repository where the required backup files are stored.

```
Available backup repositories:
1. serv10_repo
Enter repository number: 1
```

NOTE

- The account used to connect to Veeam Backup & Replication server must have access permissions on the required repository. Otherwise the repository will not be displayed in the list of available repositories. To learn how to configure access permissions on repositories, see Setting Up User Permissions on Backup Repositories.
- The wizard does not import existing backups from the repository. To perform a System Copy restore from the imported backup, you must map the backup. For details, see Importing Backups.

3. Perform the SAP HANA System Copy based restore following this SAP KB article.
4. After the restore, you must revert back the restore-server option of the Veeam Plug-in configuration wizard. Otherwise, you will not be able to restore data from the actual server backup file. If you perform only restore to other server, leave this setting enabled. It will not affect the backups of the actual system.

```
VM2ADM:/opt/veeam/VeeamPluginforSAPHANA> SapBackintConfigTool --set-restor
e-server
Select source SAP HANA plug-in server to be used for system copy restore:
1. SAP-VM1
2. SAP-VM02
Enter server number: 2
Available backup repositories:
1. serv10_repo
Enter repository number: 1
```

NOTE

If you are performing a system copy from another database to a database that was previously backed up by Veeam Plug-in, mind the following.
If the catalog_backup_using_backint parameter is enabled, after performing a system copy, SAP HANA automatically starts a new log chain and sends it to the backint along with a new catalog backup. This new catalog backup overwrites the previous catalog, making it impossible to access pre-restore backups for this database.

If you are planning to restore the older state of this database, you can disable the catalog_backup_using_backint parameter before performing the system copy.

Also, to have access to pre-restore backups, you can store a copy of the old backup catalog outside the default directory and specify this catalog during the restore.

Restore from Backup Copy

You can restore from backups and backup copies. To restore from backup copies, you must enable the **restore from backup copy** option in the Veeam Plug-in wizard.

> **IMPORTANT**
>
> If the **restore from backup copy** option is enabled, you cannot back up databases using Veeam Plug-in, and you cannot restore from backups created by primary Veeam Plug-in backup jobs. You can restore only from backup copy files until you disable the **restore from backup copy** option.

- Enabling Restore from Backup Copy
- Disabling Restore from Backup Copy

Enabling Restore from Backup Copy

To be able to restore from backup copies, do the following:

1. In the machine where Veeam Plug-in is installed, open the terminal and run the following command:

   ```
   SapBackintConfigTool --configure-restore-from-copy
   ```

2. Select the number of the backup copy job you want to use:

   ```
   Select secondary job for failover:
   0. Disable
   1. Plug-ins backup copy job\linuxq01 SAP HANA backup <serv10_repo>
   Select secondary job for failover:1
   ```

   > **IMPORTANT**
   >
   > The account used to connect to the Veeam Backup & Replication server must have access permissions on the required repository.

Disabling Restore from Backup Copy

To be able to back up with Veeam Plug-in and restore from backups, disable the restore from backup copies (set the parameter back to **0**):

```
SapBackintConfigTool --configure-restore-from-copy
Select secondary job for failover:
0. Disable
1. Plug-ins backup copy job\linuxq01 SAP HANA backup <serv10_repo>
Select secondary job for failover:0
```

Retention of SAP HANA Backups

In the main scenario, when using Veeam Plug-in for SAP HANA, you must configure the retention policy using native SAP HANA tools. For details the SAP HANA housekeeping options:

- Deleting Backups Using SAP HANA Tools

Also, you can manually delete backups from a backup repository using the Veeam Backup & Replication console and enable the force deletion functionality of Veeam Plug-in. For details, see:

- Deleting Backups Manually Using Veeam Backup & Replication Console
- Configuring Force Deletion of Backups

Deleting Backups Using SAP HANA Tools

To configure retention policies for SAP HANA backups, you can use the SAP HANA housekeeping options:

- Manual Deletion of Backups in SAP HANA Studio
- Configuring Retention Policy in SAP HANA Cockpit
- Deletion of Catalog and Backups Using Scripts

> **IMPORTANT**
>
> If you delete backups from a backup catalog using scripts or SAP HANA Studio and don't select the option to delete backup physically from the backup location, backups will remain in the backup repository. In this case, we recommend to enable the options for physical deletion of backups in used SAP HANA retention tools or you must enable the force deletion feature of Veeam Plug-in for SAP HANA. Otherwise, you will run out of space on the backup repository.

Manual Deletion of Backups in SAP HANA Studio

For details, see the Housekeeping for Backup Catalog and Backup Storage section of the SAP HANA Administration Guide.

To physically delete the backups, you must select the **Catalog and Backup Location** option. Note that if you have physical backups in both the file system and a Veeam backup repository, you can choose to delete data backups in only one location.

Configuring Retention Policy in SAP HANA Cockpit

For details, see the Retention Policy section of the SAP HANA Administration with SAP HANA Cockpit Guide. Note that the retention policy functionality is supported only in SAP HANA 2.0 SPS03 and later versions.

When you configure a retention policy in SAP HANA Cockpit, make sure that you have selected the **Also delete physically from Backint** check box in the **Options for Backup Deletion** section. Otherwise the backups will not be deleted from the repository.

Deletion of Catalog and Backups Using Scripts

Deletion of catalog and backups using scripts. For details see the BACKUP CATALOG DELETE Statement section of the SAP HANA SQL and Views Reference.

To physically delete backups from the backup repository, you must include the **WITH BACKINT and WITH FILE** options in the script.

Deleting Backups Manually

In the main scenario, when using Veeam Plug-in for SAP HANA, you must configure the retention policy using native SAP HANA tools. For details on the SAP HANA housekeeping options, see Deleting Backups Using SAP HANA Tools.

If you have lost the backup catalog, you can delete the backups manually from Veeam backup repositories using the Veeam Backup & Replication console.

> **NOTE**
>
> If you remove backups from a backup repository manually, the backup catalog will not be updated.

To remove a backup from a backup repository, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the **Inventory** pane, select **Backups**.
3. In the working area, right-click the backup job object name and select **Delete from disk**.

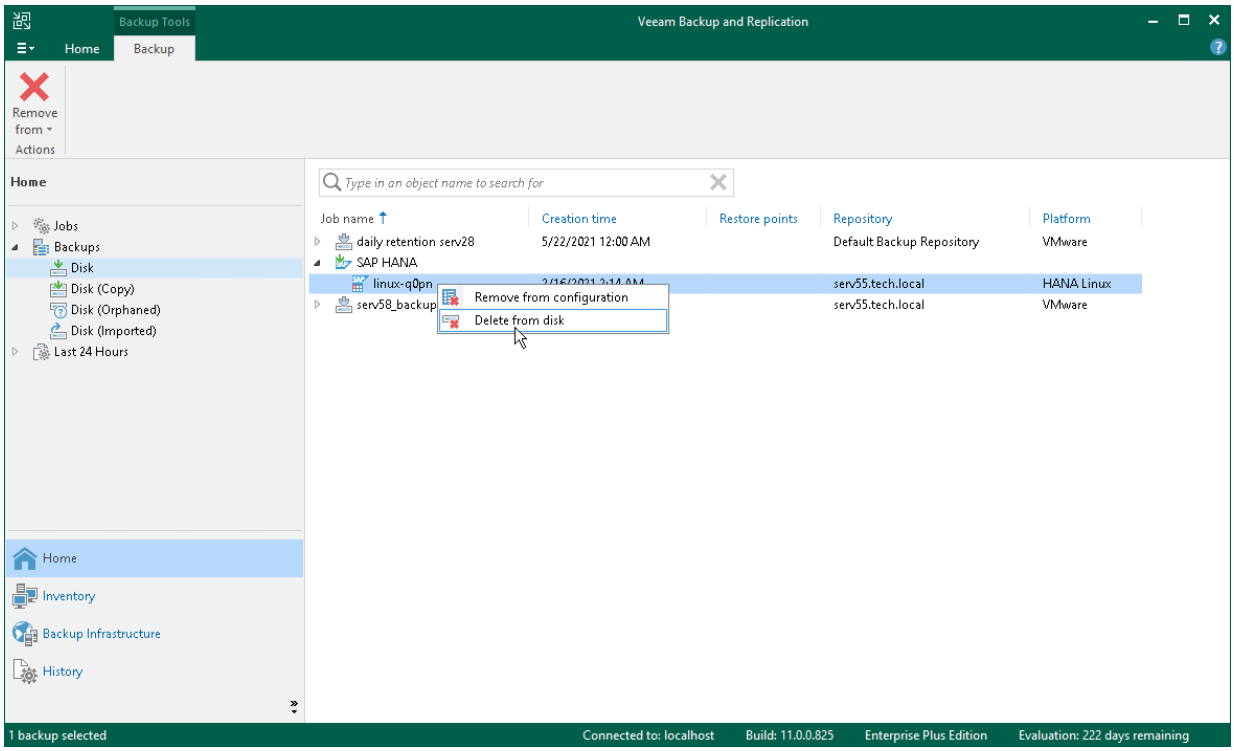

Configuring Force Deletion of Backups

In the main scenario, when using Veeam Plug-in for SAP HANA, you must configure the retention policy using native SAP HANA tools. For details, see Retention of SAP HANA Backups.

Veeam Plug-in for SAP HANA has a functionality that automatically force deletes backup files which are older than specified number of days. For example, you can use it if a backup repository contains backup files that are no longer in the backup catalog.

To enable force deletion of backup files, do the following:

1. On the SAP HANA server, run the following command.

```
SapBackintConfigTool --set-force-delete
```

2. Enter the number of days after which Veeam Plug-in will force delete backup files on all configured Veeam backup repositories.

```
Garbage collector automatically deletes backup files older than the specif
ied number of days.
Make sure the number of days value exceeds your retention policy.
To disable this functionality, set the number of days to 0.
Enter the number of days to delete backups after, between 7 and 999 [0]:
```

By default, the force delete functionality is disabled (set to *0*).

IMPORTANT

- A value for the **number of days** setting must be at least 1 backup generation period longer than the retention period for your SAP HANA backups. Otherwise, Veeam Plug-in will delete earliest backups created within the retention period.
- If a backup repository contains backups older than the specified retention period, Veeam Plug-in removes old backup files only after the next run of the Backint backup.

Removing Backups from Configuration

If you want to remove records about backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation.

When you remove a backup from the configuration, backup files (VAB, VBM) remain on the backup repository. You can import backup files later and restore data from them.

To remove a backup from the configuration:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, select the backup and click **Remove from** > **Configuration** on the ribbon. You can also right-click the backup and select **Remove from configuration**.

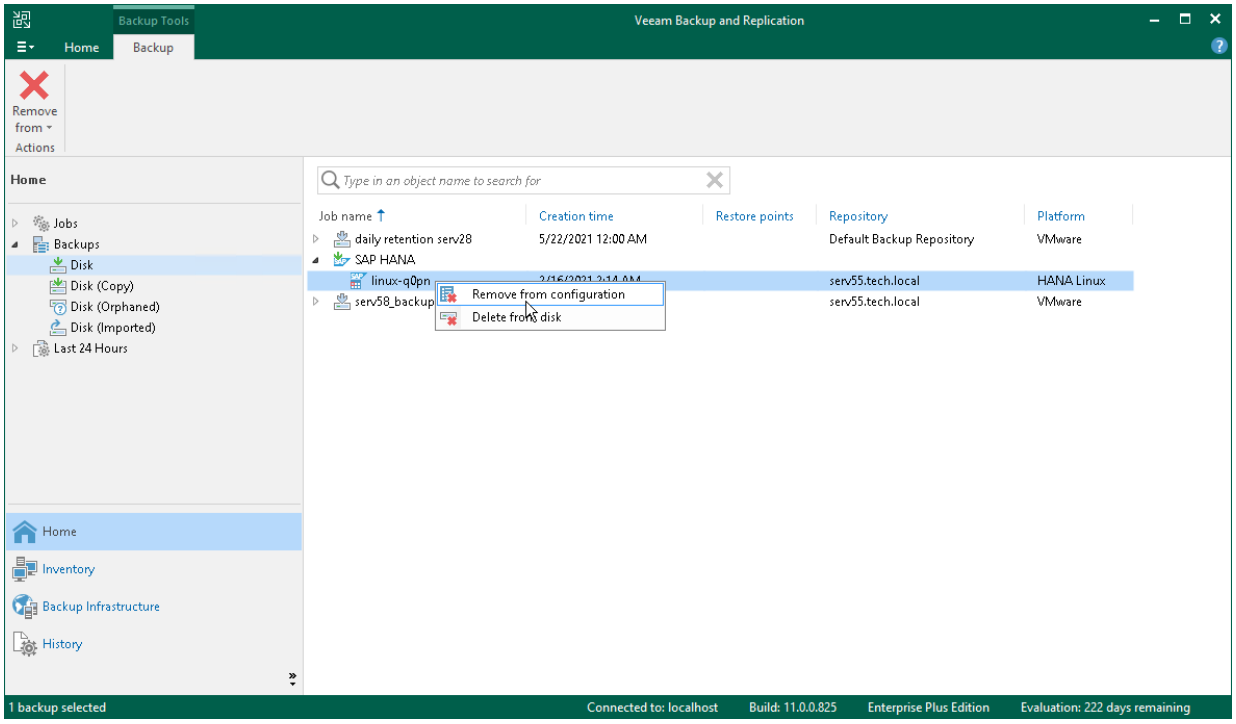

Backup Copy for SAP HANA Backups

Having just one backup does not provide the necessary level of safety. The primary backup may get destroyed together with production data, and you will have no backups from which you can restore data.

To build a successful data protection and disaster recovery plan, it is recommended that you follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups.
- 2: You must use at least two different types of media to store the copies of your data, for example, local disk and cloud.
- 1: You must keep at least one backup offsite, for example, in the cloud or in a remote site.

Thus, you must have at least two backups and they must be in different locations. If a disaster takes out your production data and local backup, you can still recover from your offsite backup.

To help you adopt the 3-2-1 rule, Veeam Backup & Replication offers the backup copy functionality that allows you to create several instances of the same backup in different locations, whether onsite or offsite. Backup copies have the same format as those created by backup jobs and you can recover your data from them when you need it.

Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes. Backup copy is a job-driven process. When enabled, the backup copy job for Veeam Plug-in backups runs continuously.

For details on how it works, see the Backup Copy section of the Veeam Backup & Replication User Guide.

In This Section

- Creating Backup Copy Job
- Converting Backup Copy to Backup

Creating Backup Copy Job

To copy backups to a secondary location, you must configure a backup copy job. The backup copy job defines how, where and when to copy backups. One job can be used to process backups of one or more machines.

You can configure a job and start it immediately or save the job to start it later.

Before creating a job, check prerequisites. Then use the **New Backup Copy Job** wizard to configure a backup copy job.

1. Launch Backup Copy Job wizard.
2. Specify a job name and description.
3. Selects backups to process.
4. Define backup copy target.
5. Specify advanced settings.
6. Define backup copy schedule.
7. Finish working with the wizard.

Before You Begin

Before you create a backup copy job, check the prerequisites and limitations:

- If you create a backup copy job for a backup of Veeam Plug-in version 9.5 Update 4, the backups will not be copied. Veeam Plug-in will issue a warning that you must upgrade Veeam Plug-in to version 10. After you upgrade the plug-in and create a new restore point, backup copy job will start to function and copy backup files created by both plug-in versions.
- Backup infrastructure components that will take part in the backup copy process must be added to the backup infrastructure and properly configured. These include source and target backup repositories between which backups must be copied.
- The target backup repository must have enough free space to store copied backups. To receive alerts about low space on the backup repository, configure global notification settings. For more information, see Specifying Other Notification Settings.
- For Veeam Plug-in backup copy jobs, you cannot select a Veeam Cloud Connect repository as a backup copy target.
- If you have upgraded the backup files, make sure that you have upgraded Veeam Plug-in on the source server. If the plug-in is not upgraded to version 11 and you convert backup copy files to backup files, then the next backup job runs will fail.

Step 1. Launch Backup Copy Job Wizard

To create a backup copy job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Click the **Backup Copy** tab and select **Application backup**.

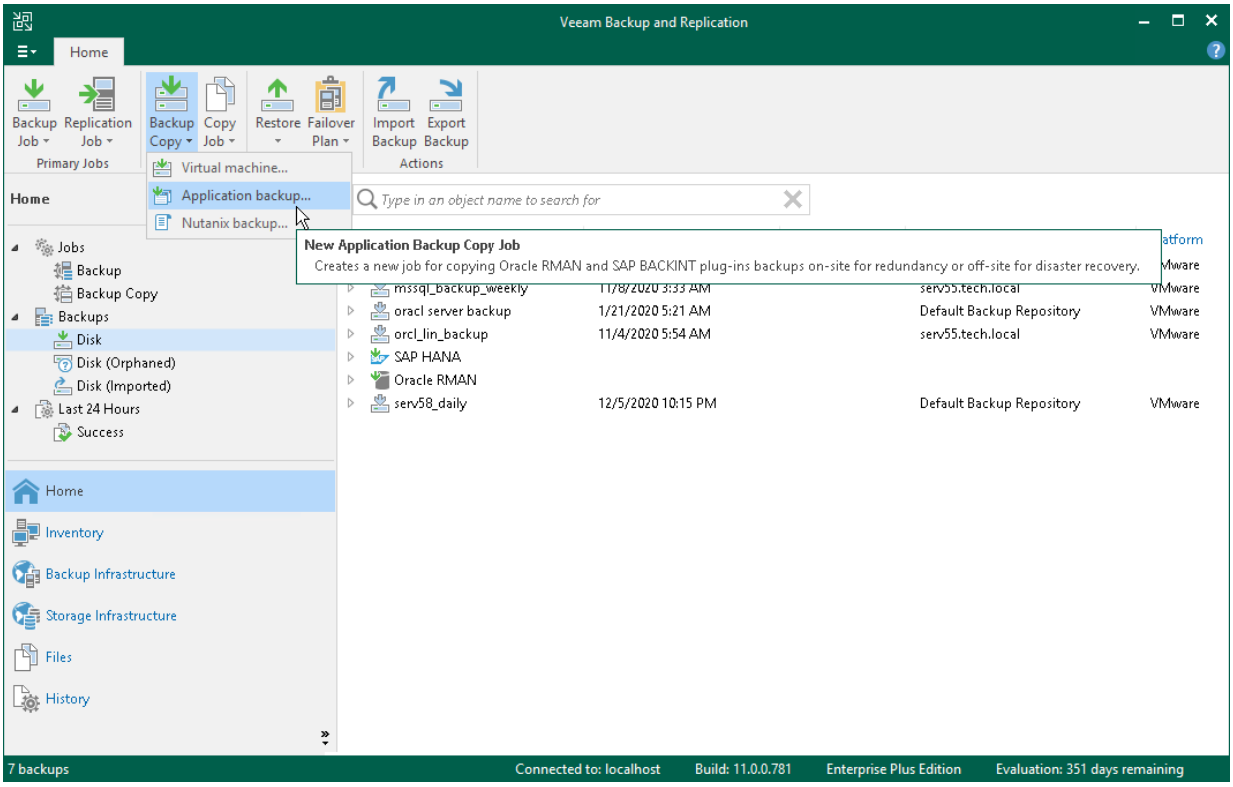

Step 2. Specify Job Name and Description

At the **Job** step of the wizard, specify a name and description for the backup copy job.

1. In the **Name** field, enter a name for the job.
2. In the **Description** field, enter a description for the job. The default description contains information about the user who created the job, date and time when the job was created.

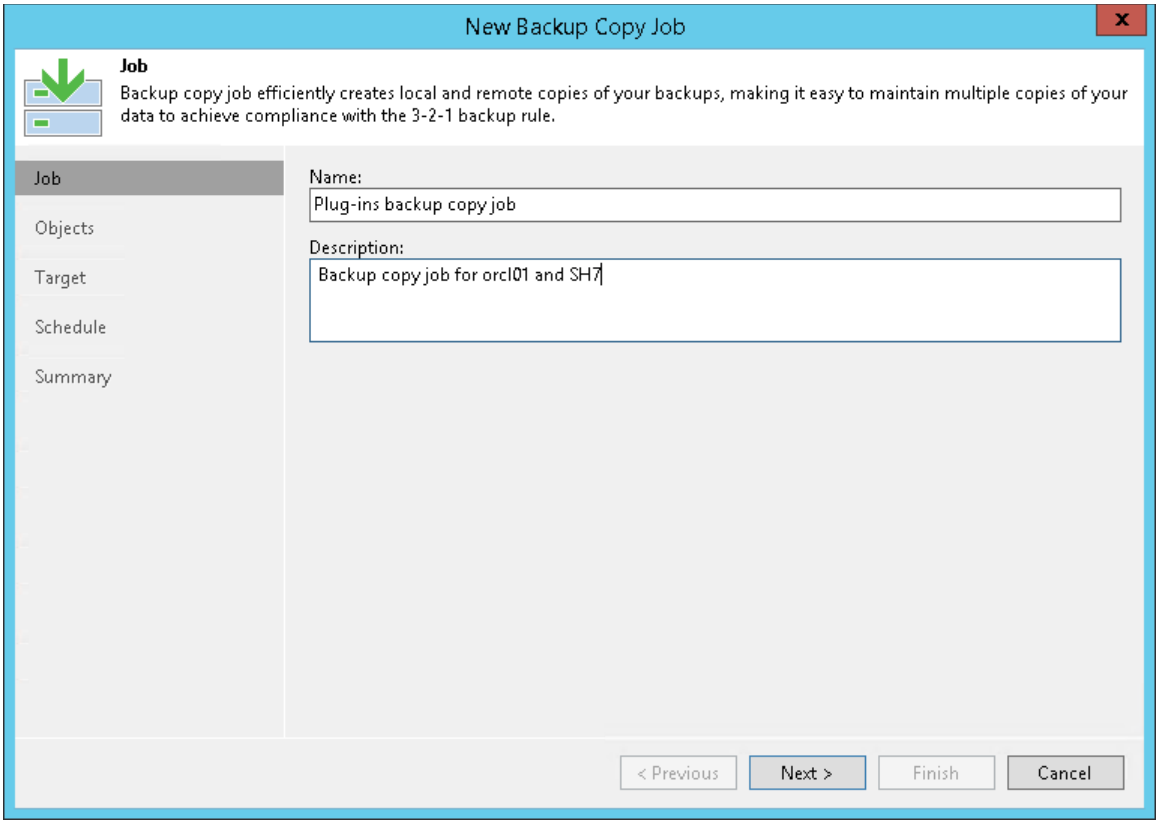

Step 3. Selects Backups to Process

At the **Object** step of the wizard, select machines whose backups you want to copy to the target repository.

1. Click the **Add** button and select from which entity you want to process the machines.
   - **From jobs**: You can select Veeam Plug-in backup jobs. When a backup copy job runs, Veeam Backup & Replication will search for backup files created by selected jobs.
   - **From repositories**: You can select repositories where Veeam Plug-in backups are stored. When a backup copy job runs, Veeam Backup & Replication will search for backup files created by Veeam Plug-in in selected repositories.
2. Use the **Remove** button if you want to remove selected jobs or repositories from processing.
3. If you have added jobs from a repository and want to exclude from processing some of the backup jobs on the selected repository, click **Exclusions** and select the jobs that you want to exclude.

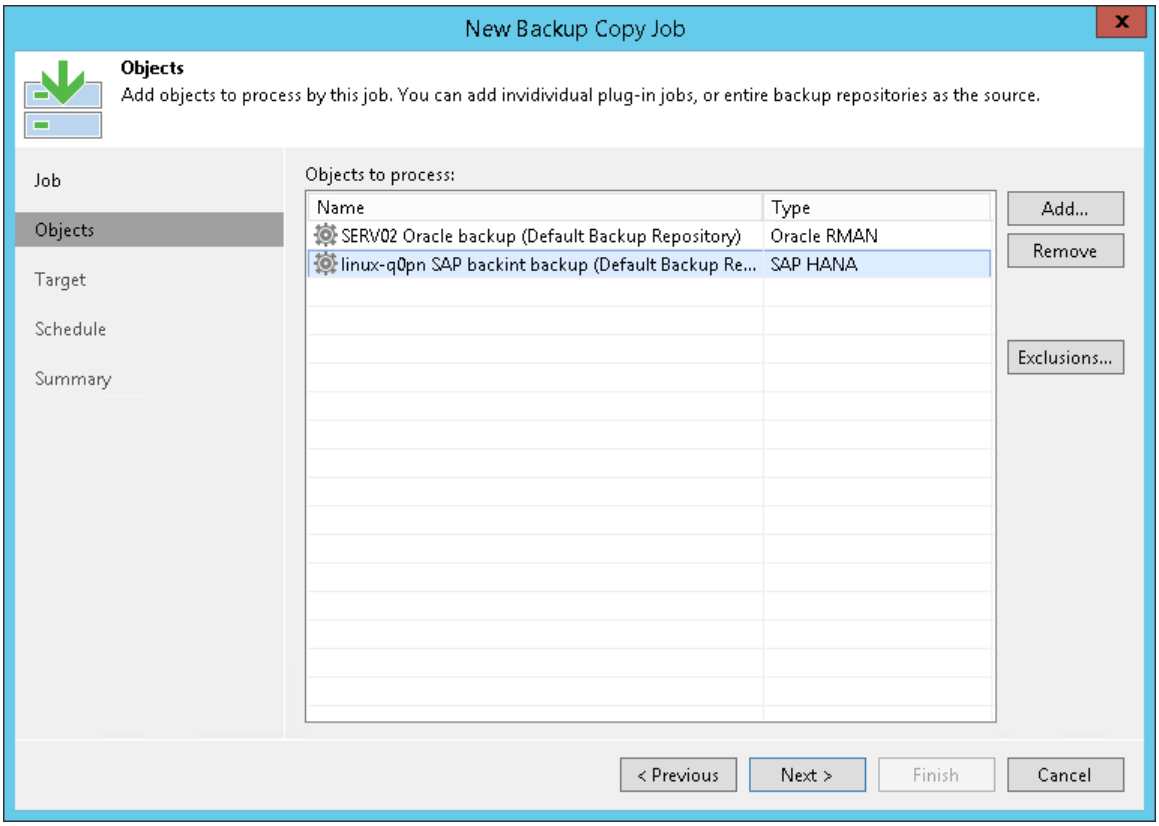

Step 4. Define Backup Copy Target

At the **Target** step of the wizard, configure the target repository settings.

1. From the **Backup repository** list, select a backup repository in the target site where copied backups must be stored. When you select a target backup repository, Veeam Backup & Replication automatically checks how much free space is available on it. Make sure that you have enough free space to store copied backups.

   > **IMPORTANT**
   >
   > For Veeam Plug-in backup copy jobs, you cannot select a Veeam Cloud Connect repository as a backup copy target.

2. If the target repository contains a Veeam Plug-in backup that was excluded from the backup copy job, and if you don't want to transfer duplicate data, you can use the mapping feature.

   After you configure mapping, if some of backup files (.VAB) of the source backup are missing in the target backup copy, these files are uploaded to the target backup copy.

   > **NOTE**
   >
   > Veeam Plug-in backup copy jobs does not use WAN accelerators.

   To map a backup copy job to the backup:

   a. Click the **Map backup** link.

   b. Point the backup copy job to the backup in the target backup repository. Backups in the target backup repository can be easily identified by backup job names. To facilitate search, you can use the search field at the bottom of the window.

   > **IMPORTANT**
   >
   > - Used account must have access to Veeam backup repositories that you plan to use.
   > - Encryption must be disabled on the repository.
   >
   > Otherwise, the repositories will not be listed as available. To learn how to configure access permissions and encryption settings on repositories, see Access and Encryption Settings on Repositories.

3. You can specify the number of days after which the backup copy will be deleted from the repository. Note that the countdown starts from the moment when source backup has been created.

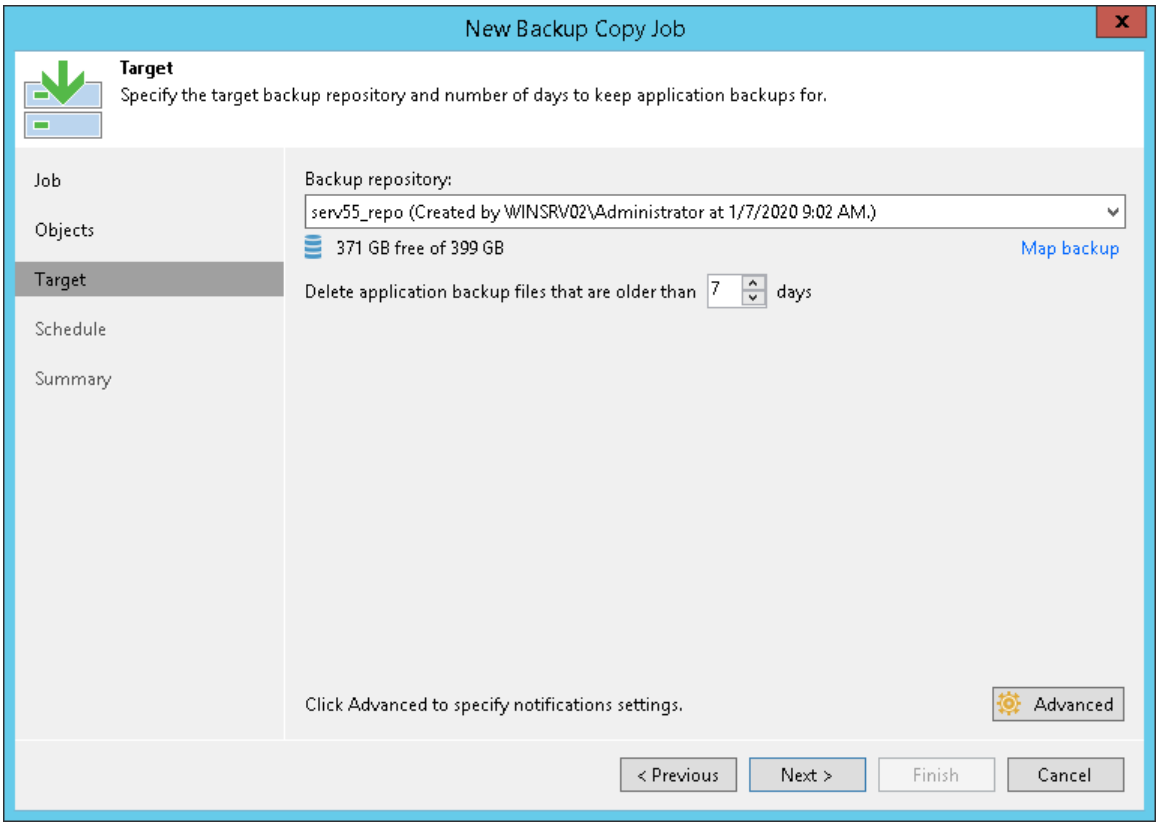

Step 5. Specify Advanced Settings

At the **Target** step of the wizard, click **Advanced** to configure compression, RPO and notifications settings.

- Compression and Deduplication
- RPO
- Notifications

Compression and Deduplication

At the **Storage** tab, define compression and deduplication settings.

By default, Veeam Backup & Replication performs deduplication before storing copied data on the target backup repository. Deduplication provides a smaller size of the resulting backup file but may reduce the job performance.

1. You can disable data deduplication. To do this, clear the **Enable inline data deduplication** check box.
2. From the Compression level list, choose a compression level to be used: **Auto**, **None**, **Dedupe-friendly**, **Optimal**, **High** or **Extreme**. The recommended level of compression for backup copy jobs is **Auto**. In this case, Veeam Backup & Replication uses compression settings of the copied backup files. For more information, see Compression and Deduplication.

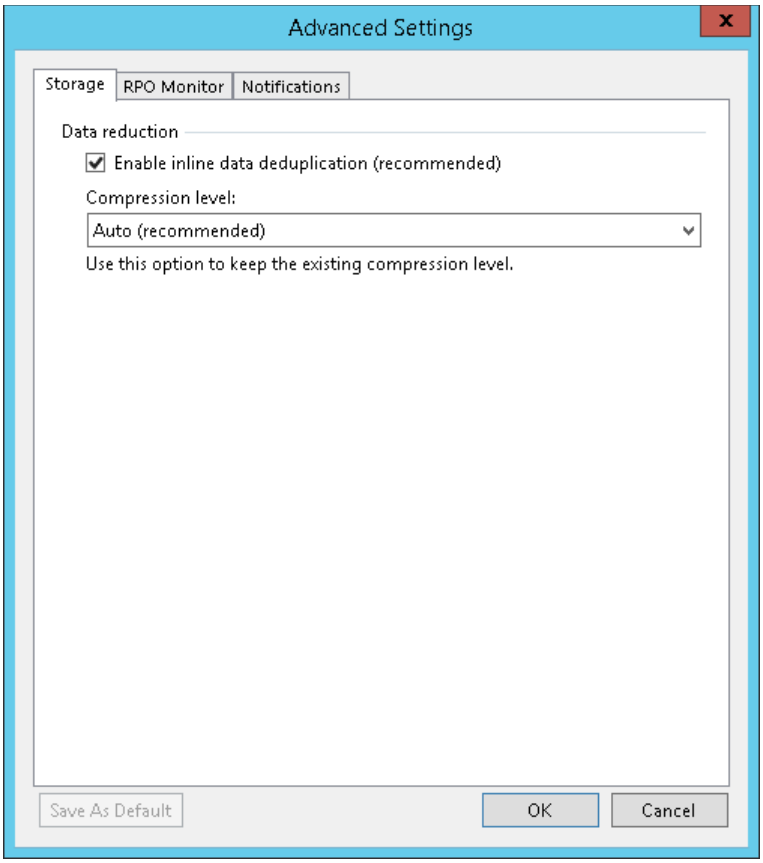

RPO Monitor

At the **RPO Monitor** tab, specify RPO warning settings.

Enable the **Warn me if backup is not copied within** check box and specify the time period in **minutes**, **hours**, or **days**.

If the backup copy is not created within the specified time period, the backup copy job will finish with the *Warning* status. The countdown starts from the moment when the required backup is finished and ready to be copied.

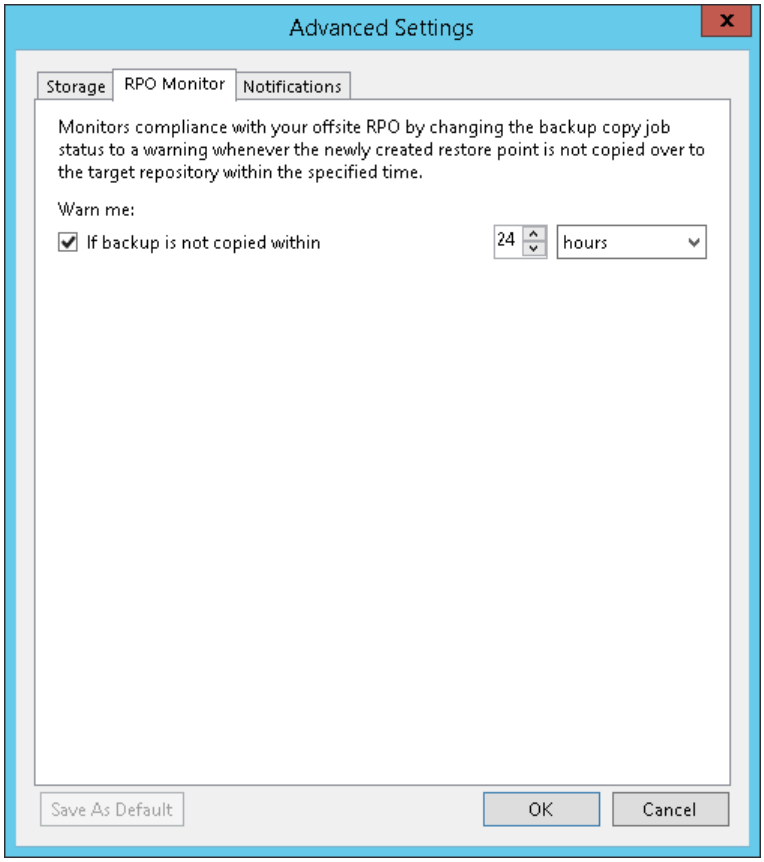

Notifications

At the **Notifications** tab, to specify notification settings for the backup copy job:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully. 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 by email in case of job failure or success. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon.
5. Veeam Backup & Replication sends a consolidated email notification once for the specified backup copy interval. Even if the synchronization process is started several times within the interval, for example, due to job retries, only one email notification will be sent.
6. 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.

7. At the **Send** at field, specify the time when you want to receive notifications. Note that you will receive a notification on a status of the job status once a day.

8. 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%*, *%VmCount%* (number of machines in the job) and *%Issues%* (number of machines in the job that have been processed with the **Warning** or **Failed** status).

     ii. Select the **Notify on success**, **Notify on warning** and **Notify on error** check boxes to receive email notification if data processing within the backup copy interval completes successfully, fails or completes with a warning.

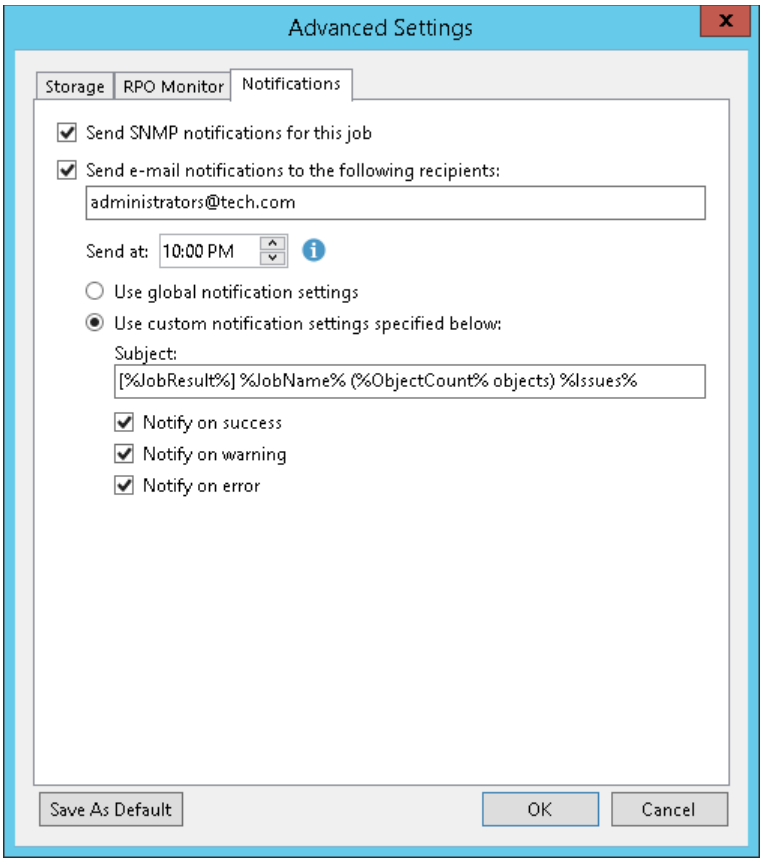

Step 6. Define Backup Copy Schedule

At the **Schedule** step of the wizard, define a time span in which the backup copy job must not transport data between source and target backup repositories. For more information, see Backup Copy Window.

To define a backup window for the backup copy job:

1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enable** and **Disable** options to mark the selected area as allowed or prohibited for the backup copy job.

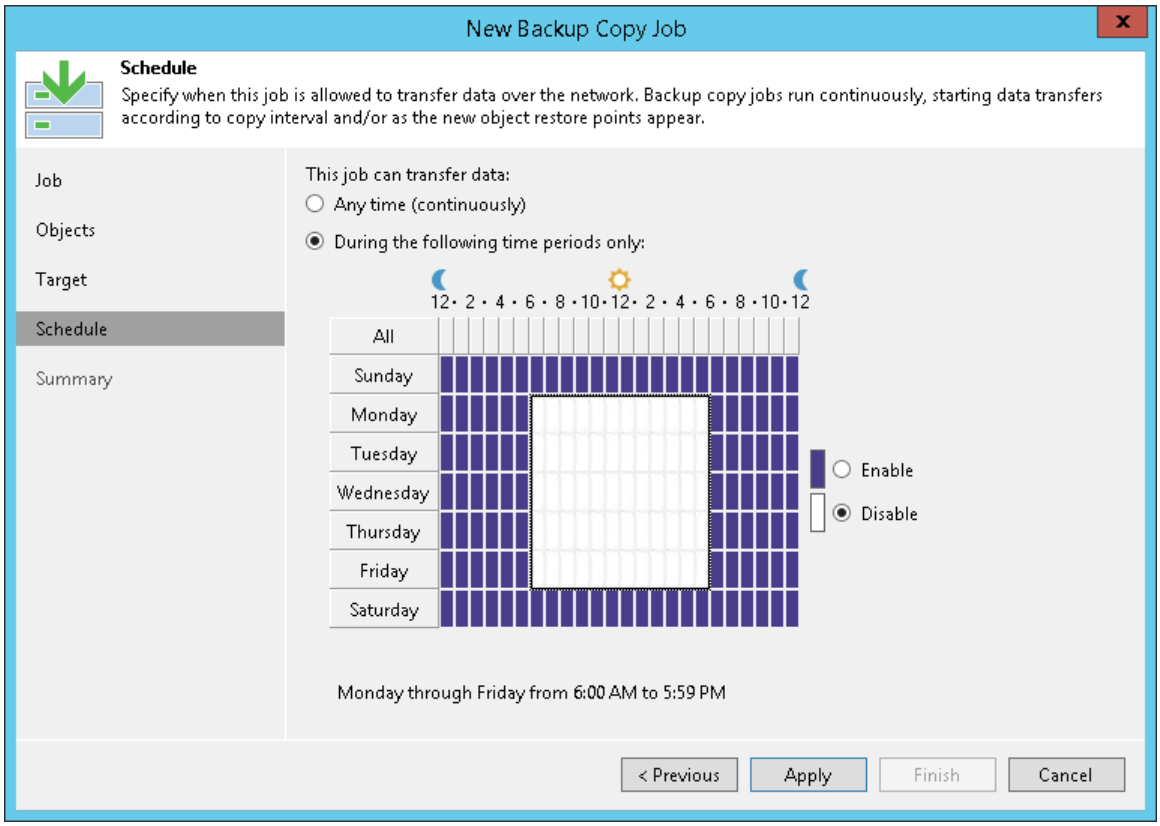

Step 7. Finish Working with the Wizard

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

1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

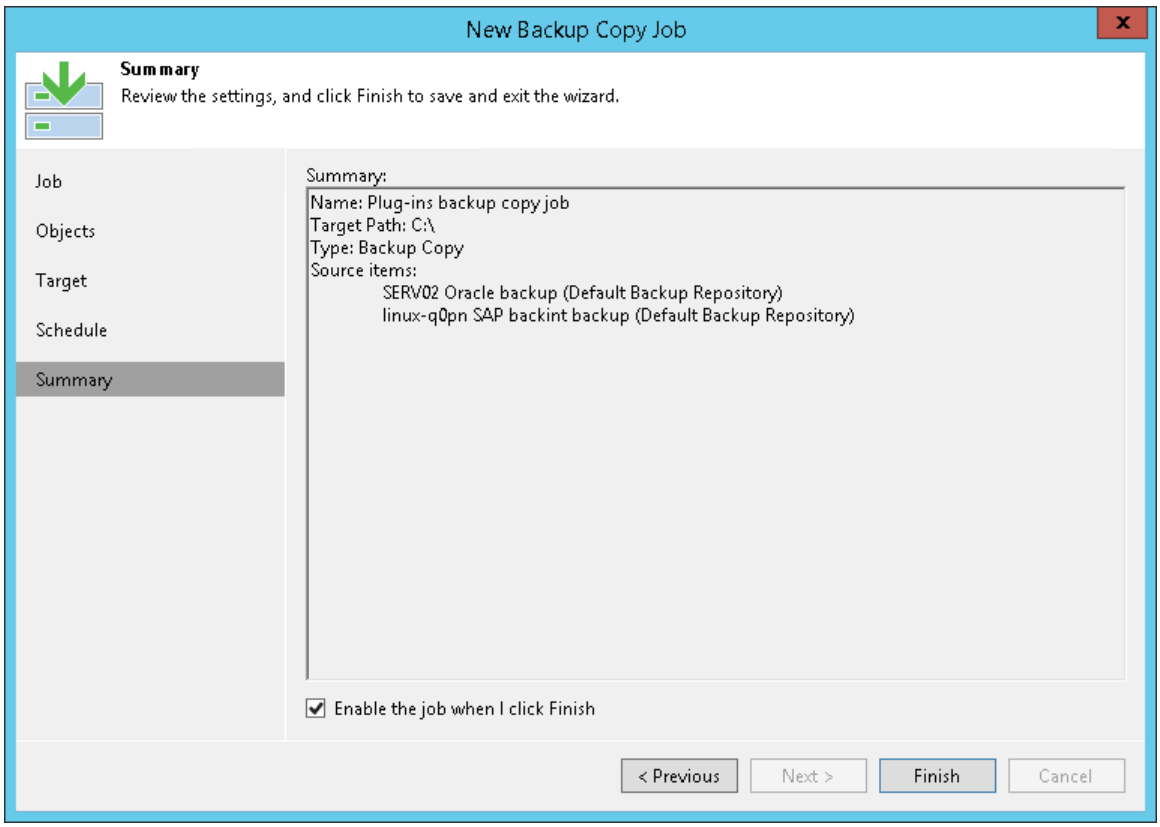

Converting Backup Copy to Backup

If you have imported Veeam Plug-in backup copies from another server, you can convert them into regular backup files. When you convert a backup copy to a backup, Veeam Plug-in creates a backup job with the converted backup. You can use this backup job to continue a backup chain and use the converted backup as a restore point.

You can convert unbind Veeam Plug-in backups into regular Veeam Plug-in backup files in the following cases:

- If you have deleted a backup copy job which created the backup copy.
- If you have excluded a backup job from a backup copy job that used multiple backup jobs as a source.
- If you imported a Veeam Plug-in backup copy from another host.

> **NOTE**
>
> If you want to restore from a backup copy, you don't need to convert the backup copy to backup. For details, see Restore from Backup Copy.

Converting Backup Copy to Backup for SAP HANA

To convert a backup copy to a primary backup, use the **--promote-backup-copy-to-primary** parameter as shown below:

```
SapBackintConfigTool --promote-backup-copy-to-primary
Backup copies available for promotion to primary backup:
1. Backup Copy Job 1\saprhel01-localdomain SAP backint backup (Default Backup R
epository)
Select backup: 1
Promotion of backup copy to a primary backup will reconfigure the plug-in to us
e a different repository. Continue? (y/N): y
```

> **IMPORTANT**
>
> [For backups of scale-out clusters and servers with the `customServerName` option] To avoid failure of conversion of backup copies, the cluster name must be the same as the name used in the backup copy.

Logs and Support

If you have any questions or issues with Veeam Plug-in for SAP HANA or Veeam Backup & Replication, you can search for a resolution on Veeam Community Forums or submit a support case in the Veeam Customer Support Portal.

When you submit a support case, we recommend you attach necessary logs related to Veeam Plug-in operations.

For instructions on how to collect logs, see this Veeam KB.

Veeam Plug-in for Oracle RMAN

Veeam Backup & Replication offers two options to protect Oracle databases:

- Veeam Plug-in for Oracle RMAN: for transactionally-consistent RMAN-based backups of Oracle databases.
- Veeam Backup & Replication or Veeam Agents: for image-level backups of Oracle servers.

You can use both or one of the options depending on your environment specifics and approach to handle Oracle databases.

Veeam Plug-in for Oracle RMAN

Veeam Plug-in uses the backup and restore functionality of Oracle Recovery Manager (RMAN) and transfers backups to Veeam backup repositories.

Use Veeam Plug-in to back up Oracle databases in the following cases:

- If you want the Oracle database administrator to fully control the backup and recovery processes.
- If you want to use existing Oracle RMAN scripts or external schedulers.
- If you use Oracle RAC.
- If you use ASM disks on a physical server.

Veeam Backup & Replication or Veeam Agents

Veeam Backup & Replication (or Veeam Agent) performs image-level/file-level backup and restore of Oracle servers. Use Veeam Backup & Replication or Veeam Agents to back up Oracle servers in the following cases:

- If you do not have Oracle database administrators.
- If you want to control backup and restore processes on the Veeam Backup & Replication side.

If you want to use Veeam Backup & Replication to protect Oracle servers, see the Creating Backup Jobs section of the Veeam Backup & Replication User Guide.

If you want to use Veeam Agents to protect Oracle servers, see one of the following guides: Veeam Agent for Linux, Veeam Agent for Microsoft Windows, Veeam Agent for IBM AIX.

How Veeam Plug-in for Oracle RMAN Works

Veeam Plug-in functions as an agent between Oracle RMAN and Veeam backup repository.

By default, RMAN sends backups to a native RMAN location on disk (`DEFAULT DEVICE TYPE TO DISK`). When you configure the Veeam Plug-in, the default device type is changed to `SBT_TAPE`, which gives control over backup media management to Veeam Plug-in. Thus, after you deploy Veeam Plug-in on an Oracle server, you can perform all backup and restore operations in the Oracle RMAN console. Veeam Plug-in compresses, deduplicates database backups and transfers them to a backup repository connected to Veeam Backup & Replication.

When use Oracle RMAN integrated with Veeam Plug-in, the database backup is performed in the following way:

1. After you launch a database backup process in the Oracle RMAN console, RMAN launches Veeam Plug-in services.
2. Veeam Plug-in connects to the Veeam Backup & Replication server and creates a backup job (if it hasn't been created earlier).
3. Veeam Plug-in starts Veeam Data Movers on the Oracle server and on the Veeam backup repository. Depending on the configured limit of RMAN channels, there will be multiple connections started in parallel.
4. Veeam Data Movers transport the backup data to the backup repository.

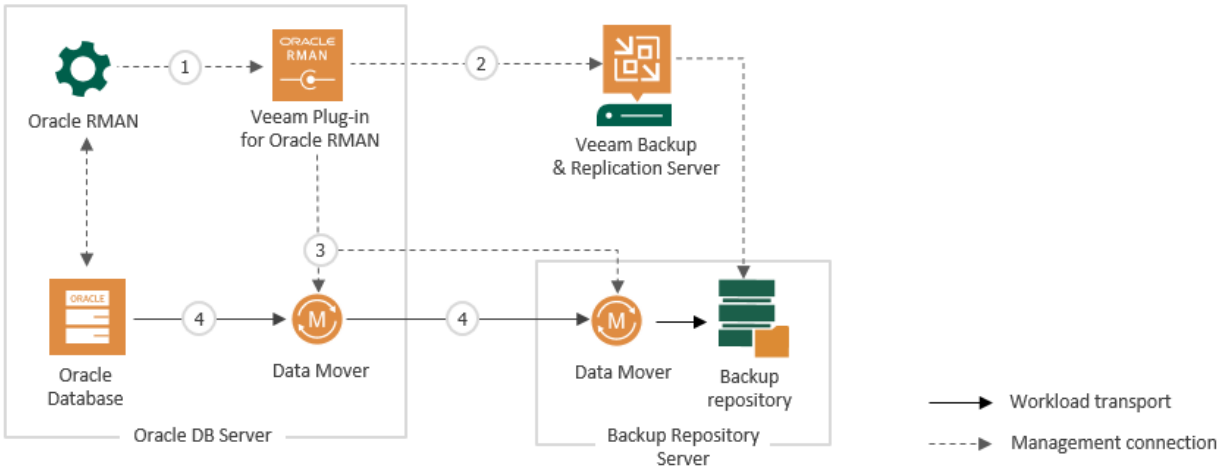

Multiple Repositories Deployment

Veeam Plug-in allows you to add up to 4 backup repositories. The backup process can be run in multiple channels. For each channel Veeam Plug-in creates a separate agent process.

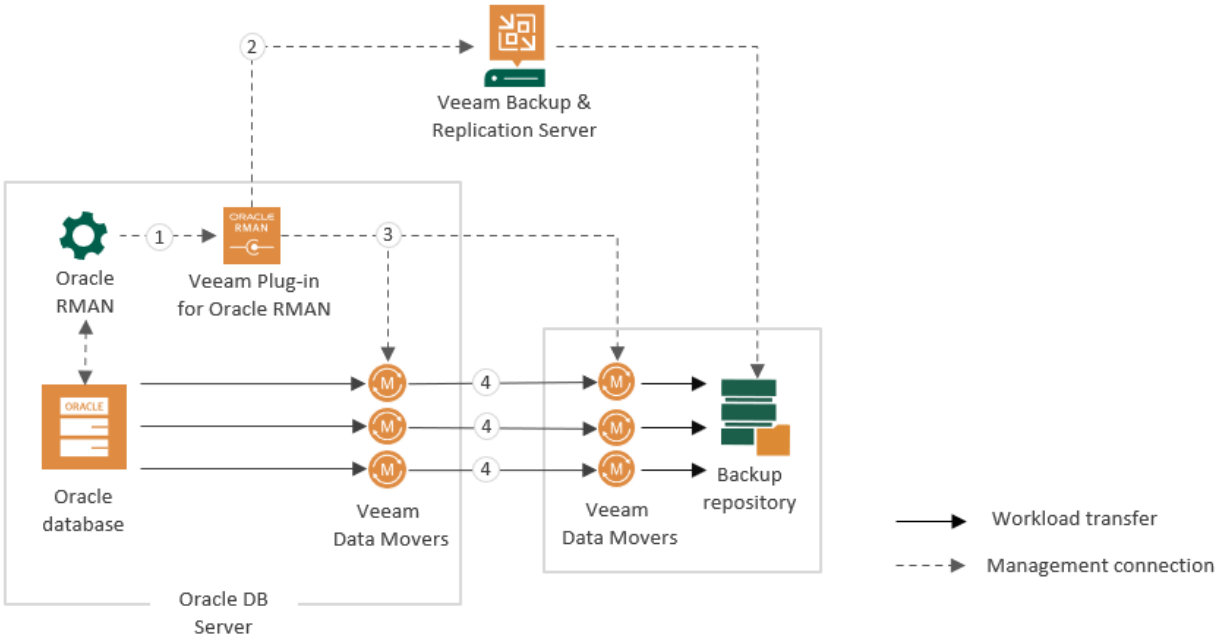

Planning and Preparation

Before you start to use Veeam Plug-in for Oracle RMAN, read the environment planning recommendations and make sure that your environment meets system requirements.

- System Requirements
- Required Permissions
- Used Ports
- Licensing
- Oracle Environment Planning
- Veeam Environment Planning
- Veeam Backup Repositories
- Access and Encryption Settings on Repositories

System Requirements

Before you start using Veeam Plug-in for Oracle RMAN, make sure the following requirements are met.

Supported OSes

Veeam Plug-in for Oracle RMAN is supported for the following OSes:

- **Microsoft Windows:**
  - Microsoft Windows Server 2008 R2
  - Microsoft Windows Server 2012/2012 R2
  - Microsoft Windows Server 2016
  - Microsoft Windows Server 2019

> **NOTE**
>
> The Veeam Plug-in for Oracle RMAN installation wizard also installs Microsoft .NET Framework 4.6 if it does not detect this component on the machine during the product installation.

- **Linux:**
  - SUSE Linux Enterprise Server 11, 12, 15 (x86 and x86_64)
  - Red Hat Enterprise Linux 6.4–8.x (x86 and x86_64)
  - Oracle Linux 6.4–8.x (x86 and x86_64)
  - CentOS 6.4–8.x (x86 and x86_64): For non-production environments, as it is not officially supported by Oracle for their databases.
- **Unix:**
  - Oracle Solaris 10, 11 (x86_64, SPARC)
  - AIX 6.1, 7.1, 7.2

Oracle Database

Veeam Plug-in for Oracle RMAN supports Oracle Database 11gR2, 12c, 18c, 19c, 21c: Standard and Enterprise Edition (Express Edition is not supported).

Supported Oracle RMAN features

Veeam Plug-in for Oracle RMAN will be registered as an SBT_TAPE device. All Oracle RMAN functionality that is supported with the SBT_TAPE device type will work (including RAC, ASM and Container DBS).

Veeam Backup & Replication

Mind the following compatibility of Veeam Backup & Replication and Veeam Plug-in versions:

- **Veeam Plug-in for Oracle RMAN 11.0.101.1264** supports integration only with Veeam Backup & Replication version 11a Cumulative Patch P20211217.

- **Veeam Plug-in for Oracle RMAN 11.0.100.1261 (11a Cumulative Patch P20211123)** supports integration with Veeam Backup & Replication version 11, 11a.
- **Veeam Plug-in for Oracle RMAN 11** supports integration only with Veeam Backup & Replication version 11.
- **Veeam Plug-in for Oracle RMAN 10.0.1.4854 (10a Cumulative Patch 20201202)** supports integration with Veeam Backup & Replication version 10, 11.
- **Veeam Plug-in for Oracle RMAN 10 (earlier than 10.0.1.4854)** supports integration only with Veeam Backup & Replication version 10.
- **Veeam Plug-in for Oracle RMAN 9.5 Update 4** supports integration with Veeam Backup & Replication version 9.5 Update 4, 9.5 Update 4a, and 9.5 Update 4b, 10.

Note that if you want to use the latest functionality, you must upgrade both Veeam Backup & Replication and Veeam Plug-in to the latest version.

Network

Veeam Plug-in should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Plug-in cannot work with the Veeam Backup & Replication server that is located behind the NAT gateway.

Required Permissions

Mind the required permissions for the following user accounts:

- OS User that Configures Veeam Plug-in
- User that Launches Backup/Restore in RMAN
- Veeam Backup Server User

OS User That Configures Veeam Plug-in

The account used for configuring Veeam Plug-in must have the following permissions.

- **For Linux and Unix machines**:

  To configure Veeam Plug-in on a Linux or Unix machine, use an account which is a member of the OSDBA (typically called as "dba") group and has SYSDBA privileges.

- **For Microsoft Windows machines**:

  To configure Veeam Plug-in on a Microsoft Windows machine, use an account which is a member of the ORA_DBA group and has SYSDBA privileges.

User That Launches Backup/Restore in RMAN

The account used for starting Oracle RMAN backup and restore processes Veeam Plug-in must have the following permissions.

- **For Linux and Unix**:

  To launch RMAN backup or restore, you can use any user account that has required set of privileges for backup operations on the Oracle side. Starting from Oracle Database 12c, Oracle recommends to use the SYSBACKUP role. For details, see the Configuring Privilege and Role Authorization section of the Oracle Database Security Guide.

  During the backup process, Veeam Plug-in connects to the database to get DB properties. Thus, Linux/Unix user that started the RMAN client must be a member of the OSDBA (typically called as "dba") group and has SYSDBA privileges.

  > **IMPORTANT**
  >
  > If you use the CONNECT string with the channel allocation command in the Oracle RAC environment, the plug-in manager process will be started by the owner of the Oracle listener, not by the user that started the RMAN client. Thus, if the listener is owned by a cluster service user (`grid`) that is not a member of the OSDBA group and doesn't have SYSDBA privileges, the plug-in manager will not be able to collect database properties and the backup will fail. As a workaround, you can add DBA privileges to the `grid` user.

- **For Microsoft Windows**:

  To launch RMAN backup or restore, you can use any user account that has required set of privileges for backup operations on the Oracle side. Starting from Oracle Database 12c, Oracle recommends to use the SYSBACKUP role. For details, see the Configuring Privilege and Role Authorization section of the Oracle Database Security Guide.

During the backup process, Veeam Plug-in connects to the database to get DB properties. Thus, the Oracle Home user must be a member of the ORA_DBA group and the OS authentication must be enabled for this user.

Veeam Backup Server User

The account which is used to authenticate against Veeam Backup & Replication must have access permissions on required Veeam repository servers. To learn how to configure permissions on repositories, see Granting Access to Repositories.

The Veeam Plug-in for Oracle RMAN uses Windows authentication methods of the Veeam Backup & Replication server to establish a connection to this server and to the backup target. It is recommended to create one user for each Veeam Plug-in server or RAC.

To work with backups created by Veeam Plug-in, you can use only the account used for creating the backup. If you want to use another account, assign the *Veeam Backup Administrator* role or *Veeam Backup Operator* and *Veeam Restore Operator* roles to the account. To learn how to assign Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication Guide.

Used Ports

To enable proper work of Veeam Plug-ins, make sure that the following ports are open.

Oracle Database Server

The following table describes network ports that must be opened to ensure proper communication of the Oracle server and backup infrastructure components.

From	To	Protocol	Port	Notes
Oracle server where Veeam Plug-in is installed	Veeam Backup & Replication server	TCP	10006	Default port used for communication with the Veeam Backup & Replication server. Note that data between Veeam Plug-ins and backup repositories is transferred directly, bypassing the Veeam Backup & Replication server.
	Backup repository server or gateway server*	TCP	2500 to 3300**	Default range of ports used as data transmission channels. For every TCP connection that a backup process uses, one port from this range is assigned.
	Oracle server (localhost)	TCP	6791+; 2500 to 3300**	Local connections between Veeam Plug-in and source Data Movers.

* For NFS share, SMB share repositories, and Dell EMC Data Domain, HPE StoreOnce deduplication storage appliances, Veeam Backup & Replication uses an auxiliary backup infrastructure component — gateway server. For details, see the Gateway Server section of the Veeam Backup & Replication User Guide.

** This range of ports applies to newly added backup infrastructure components. If you upgrade to Veeam Backup & Replication 10.0 from earlier versions of the product, the range of ports from 2500 to 5000 applies to the already added components.

Backup Repositories and Gateway Servers

Depending on the type of backup repositories that you use for Veeam Plug-in backups, the following ports must be open to allow communication between backup infrastructure components.

From	To	Protocol	Port	Notes
Veeam Backup & Replication server	Backup repository server or gateway server*	TCP	2500 to 3300**	Default range of ports used as data transmission channels. For every TCP connection that a backup process uses, one port from this range is assigned.

From	To	Protocol	Port	Notes
Direct Attached Storage				
Veeam Backup & Replication server	Linux server used as a backup repository or gateway server	TCP	22	Port used as a control channel from the Veeam Plug-in server to the target Linux host.
	Microsoft Windows server used as a backup repository or gateway server	TCP UDP	135, 137 to 139, 445	Ports used as a management channel from the Veeam Plug-in server to the Repository/Gateway server. Also, the ports are used to deploy Veeam components.
		TCP	6160, 6162	Default ports used by the Veeam Installer Service and Veeam Data Mover Service
Network Attached Storage				
Gateway server (specified in the SMB share repository settings)	SMB server	TCP	445	Default port used by SMB transport protocol.
		TCP UDP	135, 137 to 139	SMB/Netbios name resolution for SMB protocol (needed in some cases). For details, see the Used Ports section of the Veeam Backup & Replication User Guide.
Gateway server (specified in the NFS share repository settings)	NFS server	TCP UDP	111, 2049	Standard NFS ports used as a transmission channel from the gateway server to the target NFS share.
Dell EMC Data Domain				
Veeam Backup & Replication server	Dell EMC Data Domain	TCP	111	Port used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned.

From	To	Protocol	Port	Notes
or **Gateway server**	For more information, see Dell EMC Documents.	TCP	2049	Main port used by NFS. To change the port, you can use the 'nfs set server-port' command. Note that the command requires SE mode.
		TCP	2052	Main port used by NFS MOUNTD. To change the port, you can use the 'nfs set mountd-port' command. Note that the command requires SE mode.
HPE StoreOnce				
Veeam Backup & Replication server or **Gateway server**	HPE StoreOnce	TCP	9387	Default command port used for communication with HPE StoreOnce.
			9388	Default data port used for communication with HPE StoreOnce.
ExaGrid				
Veeam Backup & Replication server	ExaGrid	TCP	22	Default command port used for communication with ExaGrid.
Quantum DXi				
Veeam Backup & Replication server	Quantum DXi	TCP	22	Default command port used for communication with Quantum DXi.

* For NFS share, SMB share repositories, and Dell EMC Data Domain, HPE StoreOnce deduplication storage appliances, Veeam Backup & Replication uses an auxiliary backup infrastructure component – gateway server. For details, see the Gateway Server section of the Veeam Backup & Replication User Guide.

** This range of ports applies to newly added backup infrastructure components. If you upgrade to Veeam Backup & Replication 10.0 from earlier versions of the product, the range of ports from 2500 to 5000 applies to the already added components.

For detailed list of ports used by Veeam Backup & Replication server and backup repositories, see the Used Ports section of the Veeam Backup & Replication User Guide.

Licensing

To use the Veeam Plug-in functionality, you must have a valid Veeam Backup & Replication license. Licenses are installed and managed on the Veeam Backup & Replication server that is connected to the Veeam Plug-in server. If the license is not valid or out of resources, Veeam Plug-in backup jobs fail.

This guide provides information only on specifics of Veeam licenses for Veeam Plug-ins. For terminology and general information about Veeam Licensing, see Veeam Licensing Policy.

In this section:

- Licensed Objects
- Supported License Types and Packages
- Obtaining and Managing Licenses

Licensed Objects

An Oracle server is assumed protected if it has been processed by a Veeam Plug-in backup job in the last 31 days.

If you are using any instance-based (Veeam Universal Licensing) license on your Veeam Backup & Replication, you don't need to install any additional licenses. A protected Oracle server consumes one instance unit from the license. Oracle servers processed by backup copy jobs are not regarded as protected VMs, these types of jobs provide an additional protection level for VMs that are already protected with Veeam Plug-in backup jobs.

A machine protected by both Veeam Plug-in and Veeam Backup & Replication will consume a license only once. For example, you have an Oracle server that you back up using Veeam Plug-in. You also back up this server using image-level backup functionality of Veeam Backup & Replication. In this case, only one license will be consumed.

NOTE

[For Perpetual per-socket licenses] If you are using a legacy perpetual per-socket license, a license is required for each hypervisor CPU socket occupied by protected Oracle servers.

A socket is consumed from the license only if the hypervisor where protected servers reside is added to the Veeam Backup & Replication infrastructure. If the hypervisor is not added to the Veeam Backup & Replication infrastructure, an instance unit will be consumed from the license. To learn how to add a hypervisor to the Veeam Backup & Replication infrastructure, see the Virtualization Servers and Hosts section of the Veeam Backup & Replication User Guide.

IMPORTANT

[For Oracle RAC] The license is required for all cluster nodes, even if Veeam Plug-in is installed only on one of the nodes.

Supported License Types

You can use Veeam Plug-ins with the following license types and packages. Note that this guide contains information on specifics of Veeam license packages only for Veeam Plug-ins. For the full list of license packages, see Pricing and Packaging.

- **For Veeam Universal Licensing**:

  You can use Veeam Plug-ins with all license packages (*Veeam Backup Essentials*, *Veeam Backup & Replication*, *Veeam Availability Suite*).

  Note that if you use the *Rental* license type, functionality of Veeam Plug-ins is supported only for the *Enterprise Plus* edition of Veeam Backup & Replication.

- **For Perpetual Socket license**:

  Functionality of Veeam Plug-ins is supported only for the *Enterprise Plus* edition of Veeam Backup & Replication.

Obtaining and Managing Licenses

To learn how to install a license and monitor licensed objects, see the Licensing section in the Veeam Backup & Replication User Guide.

Oracle Environment Planning

Before you deploy Veeam Plug-in, mind the following requirements and limitations.

Deployment

[For Linux OS and Unix] To install Veeam Plug-in, the `/opt/veeam` directory must be writable.

Oracle Temp Tablespace

Veeam Plug-in runs SQL queries on the Oracle database to collect statistical information about the RMAN job processes. As with any other SQL queries, Oracle can decide based on availability of hardware resources to use Oracle Temp Tablespace for these queries. Make sure that you configure the Temp Tablespace resources to avoid shortage of temporary tablespace.

Scheduling

You can schedule backup processes with all Oracle RMAN relevant scheduling options like Cron, Windows Task Scheduler, UC4 and TWS.

If you use Veeam Backup & Replication or Veeam Agents to create image-level backups of the Oracle server, you can schedule the backup job and run Oracle RMAN backup scripts along with the backup job. For detailed instructions on how to add Oracle RMAN scripts to a backup job, see the one of the following guides:

- For Veeam Backup & Replication: Pre-Freeze and Post-Thaw Scripts section of the Veeam Backup & Replication User Guide.
- For Veeam Agent for Windows: Pre-Freeze and Post-Thaw Scripts section of the Veeam Agent for Windows User Guide.
- For Veeam Agent for Linux: Pre-Freeze and Post-Thaw Scripts section of the Veeam Agent for Linux User Guide.

> **NOTE**
>
> If you want to use Oracle RMAN scripts within backup jobs of Veeam Backup & Replication or Veeam Agents, mind the following:
>
> - A backup job does not control the workflow of Oracle RMAN scripts. The backup job invokes the script and gets its exit status when the script is finished. Backup job logs show whether the script executed successfully or failed. The script is considered to be executed successfully if "0" is returned. In order to see that the script failed, configure the script to return an exit status different than "0" in case of any errors.
> - The default timeout for a custom script in a backup job is 10 minutes. If it takes longer than 10 minutes to run the script, you can open a support ticket to increase the timeout.

Oracle RAC

Mind the following for Oracle RAC:

- It is recommended to install Veeam Plug-in on each RAC server that is responsible for the backup operations. If the plug-in is not installed on all nodes, the backup process may fail when RMAN selects another node.

- Veeam Plug-in supports parallel execution of all operations supported by Oracle RMAN: backup, restore, crosscheck, remove. This applies to execution of these commands on one or multiple databases residing on one or multiple RAC nodes.
- If you use Veeam Explorer for Oracle to restore a RAC database with different settings, it will not be restored as a cluster database. It will be restored as a standalone database.

Oracle Backup Encryption

The Oracle Secure Backup SBT library is the only interface that supports RMAN encrypted backups. Veeam Plug-ins does not support encrypted backups of Oracle databases. If the backup encryption is enabled, Veeam Plug-in backup jobs fail with the following error: `ORA-19919: encrypted backups to tertiary storage require Oracle Secure Backup`. To avoid the error, you must disable the backup encryption on the Oracle side.

Backup File Naming

During backup and restores, Veeam Plug-in uses the `repositoryID` parameter from the plug-in configuration file (`veeam_config.xml`). In the `veeam_config.xml` file, the repository ID is stored in the following format: `repositoryID="xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx"`.

To name backup files, Veeam Plug-in uses the following format by default:

```
CONFIGURE CHANNEL DEVICE TYPE sbt PARMS 'SBT_LIBRARY=/opt/veeam/VeeamPluginforO
racleRMAN/libOracleRMANPlugin.so' FORMAT 'xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx/
RMAN_%I_%d_%T_%U.vab';
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE SBT_TAPE TO '%F_RMAN_AU
TOBACKUP.vab';
```

- If you select multiple Veeam backup repositories and disable RMAN copy processing:

  Backup files will be read by default from the first selected repository. Same naming rules apply as in the single repository scenario.
- If you select multiple Veeam backup repositories and enable RMAN copy processing, you must use the repository ID with the "/" sign as a prefix for the backup file names (see the example above). This allows RMAN to directly access the requested backup file on one of the copy extents.

NOTE

[For Linux, Unix, Windows] In a backup file name, you cannot use symbols reserved by Microsoft Windows: "<", ">", ":", "/", "\", "|", "?", "*". To learn more about file naming conventions, see Microsoft Documentation.

Backup of Control File and SPFILE

If you set CONFIGURE CONTROLFILE AUTOBACKUP to ON, then RMAN automatically creates a control file and an SPFILE backup after you run the BACKUP command. For details, see the Oracle Documentation.

Controlfile Autobackup File Naming

If you perform the restore with different name and settings using Veeam Explorers for Oracle, you must enable the autobackup of the control file. If you use the Controlfile Autobackup option, the Veeam Plug-in configuration wizard creates the following RMAN configuration entry:

```
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE SBT_TAPE TO '%F_RMAN_AU
TOBACKUP.vab';
```

If you enable the control file autobackup after configuring Veeam Plug-in, you can start the configuration wizard again. The control file backup naming option will be set to default.

Parallel Processing

Veeam Plug-in supports parallel backup processing for up to 4 backup repositories or scale-out backup repositories. In the plug-in configuration wizard, if you select more than one repository, the parallelism functionality will be enabled automatically.

Note that your Oracle Enterprise Edition must be able to use RMAN parallel processing.

> **TIP**
>
> If you want to recover the Veeam Plug-in job folder to a specific point in time state, you can use Storage Replication (plus import) or the File Backup to Tape job of Veeam Backup & Replication.

Additional Files to Back up

Veeam Plug-in creates backups of databases and logs. Apart from these files,

Oracle Home

- It is recommended to back up the Oracle home folder in addition to RMAN backups. You can back it up with Veeam Backup & Replication or Veeam Agents.
- If the Oracle home folder is on a shared disk, you can use the file-level backup functionality of Veeam Backup & Replication or Veeam Agent for Linux. Alternatively, you can copy the Oracle home folder to a non-shared disk before the backup.

Oracle Recovery Catalog

You can back up the Oracle Recovery Catalog with Veeam Plug-in on the Recovery Catalog server according to the Oracle procedures. For details, see the Managing a Recovery Catalog section of the Database Backup and Recovery User's Guide.

Veeam Plug-in for RMAN Configuration File

You can back up the Veeam Plug-in configuration file. The file is located in the following directory:

- [Linux or Unix]: `/opt/veeam/VeeamPluginforOracleRMAN/veeam_config.xml`
- [Windows]: `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\veeam_config.xml`

NOTE

You can also create an image-level backup of the Oracle server using the image-level backup functionality of Veeam Backup & Replication or Veeam Agents.

Multiple Oracle Versions on One Server

Multiple Oracle Database versions deployed on the same machine are supported only for Linux or Unix based operating systems.

Disabling Veeam Explorer Processing

You can disable Veeam Explorer based restore for specific Oracle servers. To disable the restore on the Veeam Explorer for Oracle side, do the following:

- [Linux or Unix] On the Oracle server, log in as a user with the Oracle Administrator rights and create an empty file in the following directory: `/etc/veeam/disablerestore`
- [Windows] On the Oracle server, create an empty file in the following directory: `%ProgramData%\Veeam\disablerestore`

Oracle Data Guard

Since version 11, Veeam Plug-in for Oracle RMAN has an official support of Oracle Data Guard.

Veeam Environment Planning

Before you deploy Veeam Plug-in, keep in mind the following requirements and limitations.

- RMAN Channels and Resource Consumption
- Veeam Backup Job Name
- Hosting Environments

RMAN Channels and Resource Consumption

Any parallel channel started by RMAN will use one Veeam backup repository task slot. By design, Oracle Standard Edition can work with one channel. Oracle Enterprise Edition has the option to use multiple channels and you can configure them in the Veeam Plug-in configuration wizard or at the ALLOCATE CHANNEL definition in RMAN scripts. It is recommended to carefully plan repository task slots, so that Oracle RMAN can work with multiple channels in parallel when configured.

The following hardware resources are recommended based on tests on Skylake processors:

- **Oracle server**: 1 CPU core and 200 MB of RAM per currently used channel. Note that resource consumption on the Oracle server depends on hardware and Oracle settings.
- **Backup repository server**: 1 CPU core and 1 GB of RAM per 5 currently used channels.

  These resources are recommended only if you use a dedicated backup repository for Veeam Plug-in backups. If you use the same backup repository for Veeam Plug-in backups and VM backups created by Veeam Backup & Replication or Veeam Agents, consider adding the mentioned above hardware resources based on usual load on your backup repository. For details on hardware requirements for a backup repository, see the System Requirements section of the Veeam Backup & Replication User Guide.

  We recommend to contact your Veeam system engineer to optimize the channel settings and resource allocation. Also, mind the following:

  - It is recommended to use a separate backup repository for Veeam Plug-in backups.
  - The control file does not use a repository task slot and will be processed even if there are no free task slots.
- **Veeam Backup & Replication server**: during manual metadata operations such as import of backup files, the Veeam Backup & Replication server needs additional 15 GB of RAM per 1 million files located in the same backup job folder.

Veeam Backup Job Name

- On the Veeam Backup & Replication server, the backup job name will be created automatically based on the server or cluster name and selected repository.
- For environments that use Oracle RMAN copy processing, one job per repository is created.

Backup Files

- A .VAB file stores a compressed copy of an Oracle database. Veeam Plug-in creates .VAB files for both full and incremental backups.

- A .VASM file stores metadata that contain information about the backup. A .VASM file is created for each .VAB file. The .VASM files are used by Veeam Backup & Replication to get data about Veeam Plug-in backups.
- A .VACM file stores metadata of a backup job object.

Hosting Environments

By default, Veeam Plug-in uses the Oracle server hostname to create a Veeam Backup & Replication job object and backup folder. To be able to distinguish individual servers, it is recommended to set the following entry within the Veeam configuration XML file: `<PluginParameters customServerName="hostname.domain.tld" />`

- [Linux or Unix]: `/opt/veeam/VeeamPluginforOracleRMAN/veeam_config.xml`
- [Windows]: `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\veeam_config.xml`

If your servers that have the same hostname in multiple environments, you must add the following entries in the plug-in configuration file:

```
<PluginParameters useFQDNInServerName="true" />
```

IMPORTANT

For security reasons, it is recommended to use separate repositories for different customers and limit the Veeam Repository Authentication to the specific customer.

Veeam Backup Repositories

Veeam Plug-ins store backup files in repositories added to the Veeam Backup & Replication infrastructure. In this section, you can find the list of supported backup repositories and limitations for Veeam Plug-in backups.

Supported Backup Repositories

Veeam Plug-in for Oracle RMAN supports integration with the following types of repositories added to the Veeam Backup & Replication infrastructure:

- Windows Server
- Linux Server
- CIFS (SMB) Share
- Dell EMC Data Domain Boost
- Quantum DXi
- ExaGrid
- HPE StoreOnce. If you plan to use HPE StoreOnce as a backup repository for Veeam Plug-in backups, the total number of stored files (data and metadata) must not exceed 3,000,000 per Catalyst store. If necessary, multiple Catalyst stores may be created on the same StoreOnce system.
- NFS File Share

You can also use scale-out backup repositories that contain supported repository types.

Backup Repositories Limitations

- For Veeam Plug-in backups, the warning which indicates that free space on a storage device has reached a specified threshold is configured in the `veeam_config.xml` file of Veeam Plug-in. The warning settings in the Veeam Backup & Replication console does not affect this setting.

  To configure the warning settings, add the following parameter in the `veeam_config.xml` file.

  ```
  <PluginParameters repositoryFreeSpacePercentWarning="10" />
  ```

- The plug-in configuration wizard will not show repositories where the **Encrypt backups stored in this repository** option is enabled. To learn how to disable the encryption option, see Access and Encryption Settings on Repositories.

  If you want to use the same backup target with the repository-based encryption and Veeam Plug-ins, create a second repository in the subfolder for Veeam Plug-in backups.

- If you use Veeam Plug-in version 9.5 Update 4, Veeam backup copy jobs are not able to process and copy the backups created by the plug-in.

- Veeam extract utility cannot extract Veeam Plug-in backup files. By design of Oracle RMAN, these files cannot be imported "as files" to RMAN as they contain additional metadata bound to the used SBT device.

Veeam Scale-Out Backup Repositories

If you want to store Veeam Plug-in backups in scale-out backup repositories, mind the following:

- If you want to add a backup repository as an extent to a scale-out backup repository and Veeam Plug-in backups are present on this backup repository, you must do the following:

  a. In the Veeam Backup & Replication console, select Veeam Plug-in backup files that reside in this backup repository and remove them from configuration. For details, see Removing backups from configuration. Note that this action does not delete the backups from the repository.

  b. In the Veeam Backup & Replication console, delete the Veeam Plug-in backup job. For details, see Deleting Jobs.

  c. Add the repository as an extent to the scale-out repository. For details, see Extending Scale-Out Repositories.

  d. Rescan the scale-out repository. For details, see Rescanning Scale-Out Repositories.

  > **NOTE**
  >
  > Names of backup files and paths to backup files must contain only allowed characters:
  >
  > - Alphanumeric characters: `a-zA-Z0-9`
  > - Special characters: `_-.+=@^`
  > - Names of backup files and paths to backup files must not contain spaces.

  e. On the Veeam Plug-in server, set the scale-out repository as the target for backups using the following command:

  ```
  SapBackintConfigTool --set-repository
  ```

  f. Map the imported backups using the following command:

  ```
  SapBackintConfigTool --map-backup
  ```

- For Veeam Plug-in backups and backup copies, the *Performance* policy of a scale-out repository functions differently:

  a. Veeam Backup & Replication checks if there are extents without warning on free space insufficiency. If all extents have the warning, Veeam Backup & Replication uses an extent with the largest amount of free space that has a free task slot.

  b. If there are extents without the warning, Veeam Backup & Replication checks if there are incremental extents with free task slots. If there are no incremental extents with free task slots, Veeam Backup & Replication uses a full extent with the least amount of used task slots.

  c. If there are incremental extents with free task slots, Veeam Backup & Replication sends backup files to an incremental extent with the least amount of used task slots. If the amount of used tasks is the same, an extent with the largest amount of free space.

  To learn more about file placement policies of scale-out repositories, see Backup File Placement section of the Veeam Backup & Replication guide.

- If a scale-out repository is configured in the **Data locality** policy, repository extents will be selected according to the amount of free space for each Oracle RMAN connection. If there are two extents with one slot on each extent, the backup will be launched on two streams (one on each extent).

Capacity Tier

You can configure Veeam Backup & Replication to transfer Veeam Plug-in backup files to a capacity tier. Both policies (Move policy, Copy policy) are supported for Veeam Plug-in backups with the following limitations:

- For Veeam Plug-in backup files, capacity tier does not verify whether data that is being moved is unique and has not been offloaded earlier. Thus, it is highly recommended to check the pricing plans of your cloud storage provider to avoid additional costs for offloading and downloading backup data.
- Capacity tier does not track dependencies of full and incremental Veeam Plug-in backup files. Thus, mind the following:
    - [For the Move policy] When backup files are transferred to the capacity tier, Veeam Backup & Replication takes into account only the creation time of backup files. Make sure that the operational restore window is either longer than the whole backup chain cycle period or exceeds that period. Otherwise, you may encounter the scenario when full backup files are transferred to the capacity tier and their increment backup files still remain in the performance tier.
    - The capacity tier immutability expiration date does not have the additional block generation period. The immutability expiration date is based only on the number of days specified in settings of the object storage backup repository.
- If a scale-out repository is down, you cannot restore from the Veeam Plug-in backup files stored on the capacity tier. In this case, you can only import the backup files to Veeam Backup & Replication manually and then perform data recovery operations.
- If you use a capacity tier that has been created in Veeam Backup & Replication version 10, you cannot transfer Veeam Plug-in backup files to a capacity tier. However, if you want to transfer them manually, do the following:
    - If the backup files are created by Veeam Plug-in version 10, upgrade the metadata of backup files as described in Upgrading Metadata Files to New Format.
    - Run the Set-VBRScaleOutBackupRepository PowerShell command with the `-EnablePluginBackupOffload` parameter to offload backup files to the capacity tier.

Access and Encryption Settings on Repositories

While you configure Veeam Plug-in, you specify an account that must be used to connect to the Veeam Backup & Replication server. To be able to store backups on a backup repository, the specified account must have access permissions on the target backup repository.

To grant access permissions, do the following:

1. In Veeam Backup & Replication, open the **Backup Infrastructure** view.
2. In the inventory pane, click the **Backup Repositories** node or the **Scale-out Repositories** node.
3. In the working area, select the necessary backup repository, click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**.

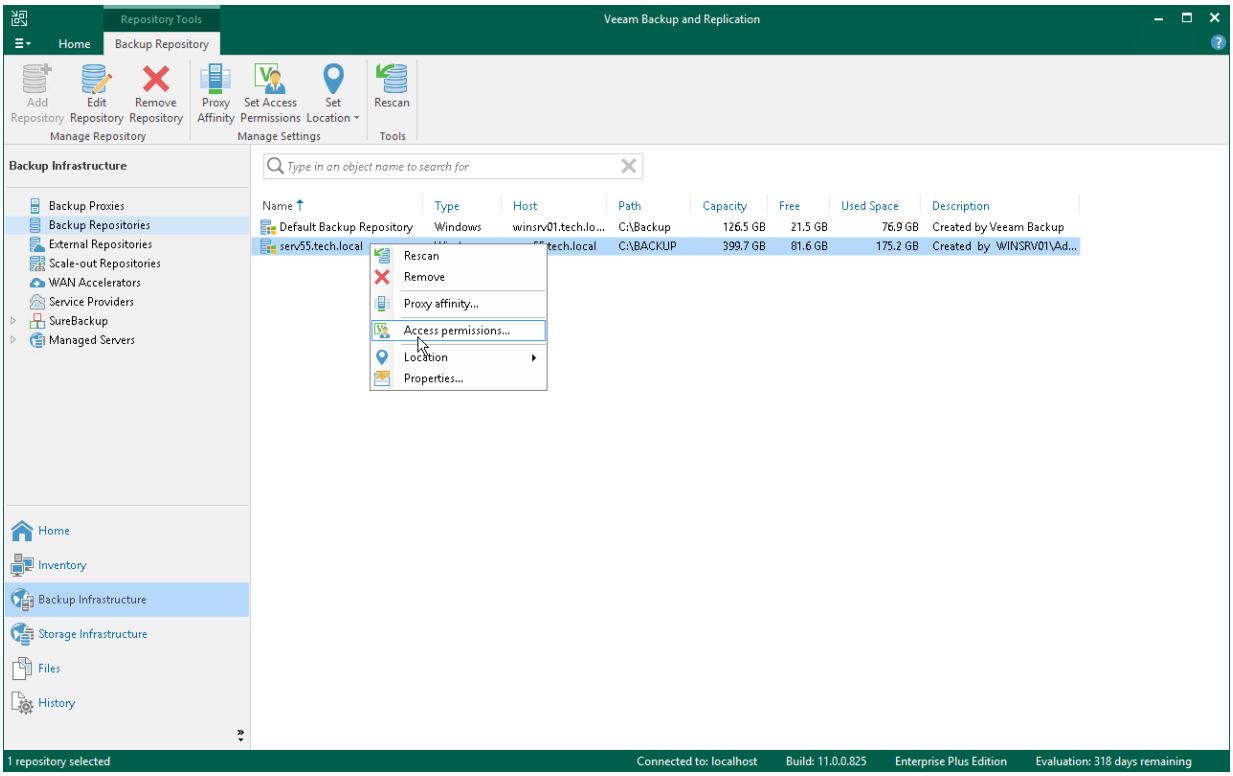

4. In the **Access Permissions** window, specify to whom you want to grant access permissions on this backup repository:
   - *Allow to everyone* – select this option if you want to grant repository access to any user. This option is equal to granting access rights to the *Everyone* group in Microsoft Windows (anonymous users are excluded). For security reasons, the option is not recommended for production environments.
   - *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 repository. Click **Add** to add the necessary users and groups to the list.

5. Veeam Plug-ins cannot send backups or backup copies to a repository where encryption is enabled. Thus, unselect the **Encrypt backups stored in this repository** check box.

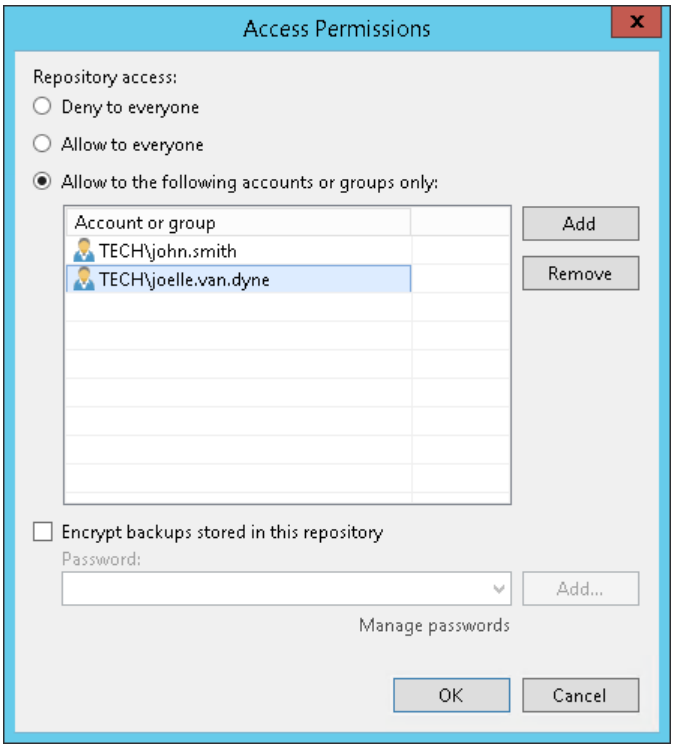

Deployment and Configuration

Veeam Plug-in for Oracle RMAN is a feature of Veeam Backup & Replication. This guide gives instructions on how to deploy Veeam Plug-in assuming that you already have deployed Veeam Backup & Replication and configured a backup repository. To learn how to deploy Veeam Backup & Replication, see the Veeam Backup & Replication User Guide for your platform.

To be able to use Veeam Plug-in for Oracle RMAN, you must install the plug-in on the Oracle server and configure the plug-in settings:

- Installing Veeam Plug-in for Oracle RMAN
- Configuring Veeam Plug-in for Oracle RMAN

See Also

- Upgrading Veeam Plug-in for Oracle RMAN
- Importing/Exporting Plug-in Settings
- Importing Backups
- Upgrading Backup Files
- Uninstalling Veeam Plug-in for Oracle RMAN

Installing Veeam Plug-in for Oracle RMAN

See one of the following guides depending on which OS is installed on the target machine.

- Installing Veeam Plug-in on Linux machines
- Installing Veeam Plug-in on Solaris machines
- Installing Veeam Plug-in on Windows machines
- Installing Veeam Plug-in on IBM AIX machine

Installing Plug-in on Linux

Veeam Plug-in for Oracle RMAN is an additional component of Veeam Backup & Replication, and the installation package of the plug-in is included in the Veeam Backup & Replication installation ISO file.

You can install the plug-in using the `.RPM` package or extract the plug-in files from the `.TAR.GZ` archive. Depending on the type of package suitable for your OS, perform steps in one of the following guides:

- Installing Plug-in from .RPM Package
- Unpacking Plug-in from .TAR.GZ Archive

IMPORTANT

Mind the following:

- Veeam Plug-in for Oracle RMAN must be installed on the Oracle Database server.
- The `/opt/veeam` directory must be writable.
- To install the plug-in, use the `sudo` command or a user with root privileges.

Installing Plug-in from .RPM Package

To install Veeam Plug-in on a Linux machine, perform the following steps.

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.X.X.XXXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk image from the Veeam Backup & Replication: Download page.

2. Open the mounted disk image and go to the `Plugins\Oracle RMAN\Linux` directory.
3. Upload the `VeeamPluginforOracleRMAN-11.X.X.XXXX.x86_64.rpm` package to the Oracle server. If you need the 32-bit version, choose the `i386` package.
4. To install Veeam Plug-in, run the following command:

```
rpm -i VeeamPluginforOracleRMAN-11.X.X.XXXX.x86_64.rpm
```

Unpacking Plug-in from .TAR.GZ Archive

To extract plug-in files from the `.TAR.GZ` archive, perform the following:

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.X.X.XXXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk image from the Veeam Backup & Replication: Download page.

2. Open the mounted disk image and go to the `Plugins\Oracle RMAN\Linux` directory.
3. Upload the `VeeamPluginforOracleRMAN.tar.gz` file to the Oracle server.
4. Create the `/opt/veeam` directory.

   ```
   mkdir /opt/veeam
   ```

5. Unpack the plug-in files from the archive to the `/opt/veeam` directory.

   ```
   tar -xzvf VeeamPluginforOracleRMAN.tar.gz -C /opt/veeam
   ```

Installing Plug-in on Microsoft Windows

You can install Veeam Plug-in for Oracle RMAN on Windows machines using a wizard or in an unattended mode. For instructions, see:

- Installing Veeam Plug-in on Windows machines
- Installing Veeam Plug-in in an unattended mode

> **NOTE**
>
> When you launch the installation file, it also installs Microsoft .NET Framework 4.6 if it does not detect this component on the machine during the product installation. In some cases, installation of .NET Framework requires a reboot of the machine. This can happen, for example, if you have an earlier version of .NET Framework installed on the machine and during the installation process it is used by third-party software.

Installing Plug-in on Windows Machine

Veeam Plug-in for Oracle RMAN is an additional component of Veeam Backup & Replication, and the installation package of the plug-in is included in the Veeam Backup & Replication installation ISO file.

To install Veeam Plug-in for Oracle RMAN on a Windows machine, do the following:

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.0.0.XXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk at: https://www.veeam.com/backup-replication-vcp-download.html.

2. In the installation disk folder go to `Plugins\Oracle RMAN\Windows`.

3. To launch the installation wizard, run the `VeeamPluginforOracleRMAN.exe` file.
4. At the welcome screen of the installation wizard, click **Next**.

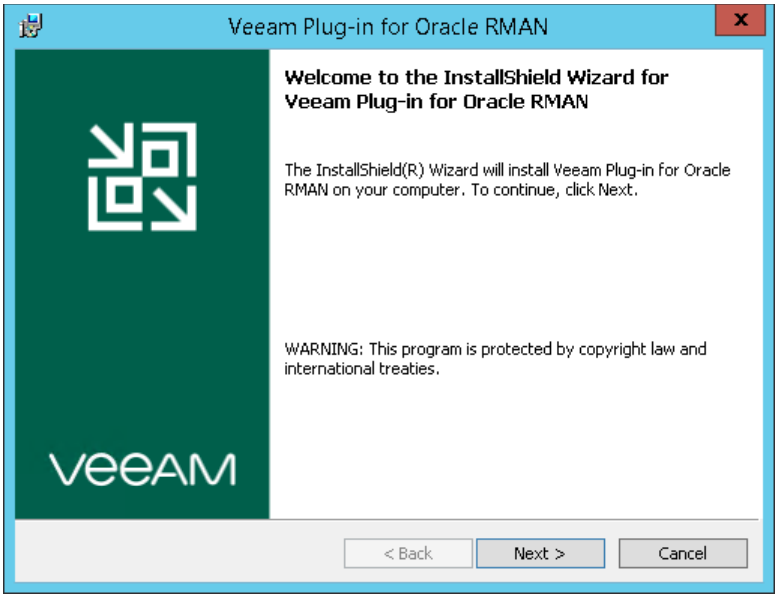

5. At the **License Agreement** step of the wizard, accept the terms of license agreements and click **Next**.

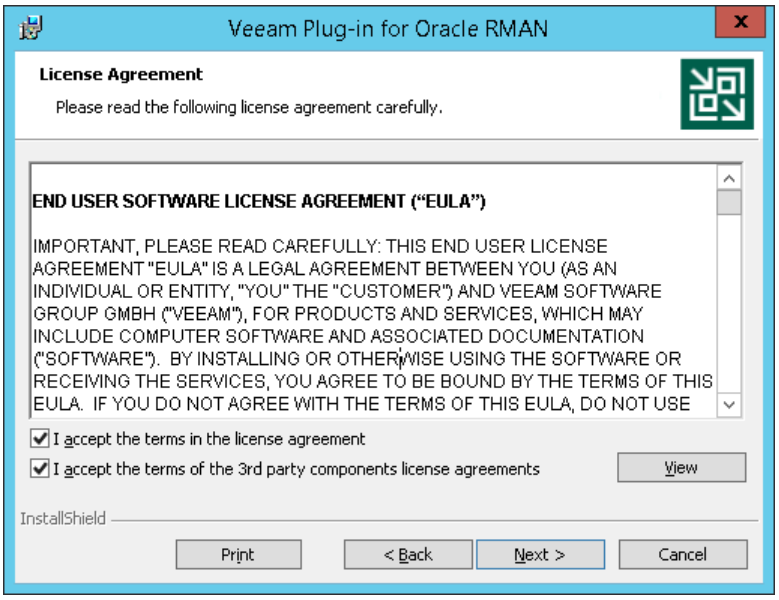

6. At the **Custom Setup** step of the wizard, specify the installation path for Veeam Plug-in and click **Next**.

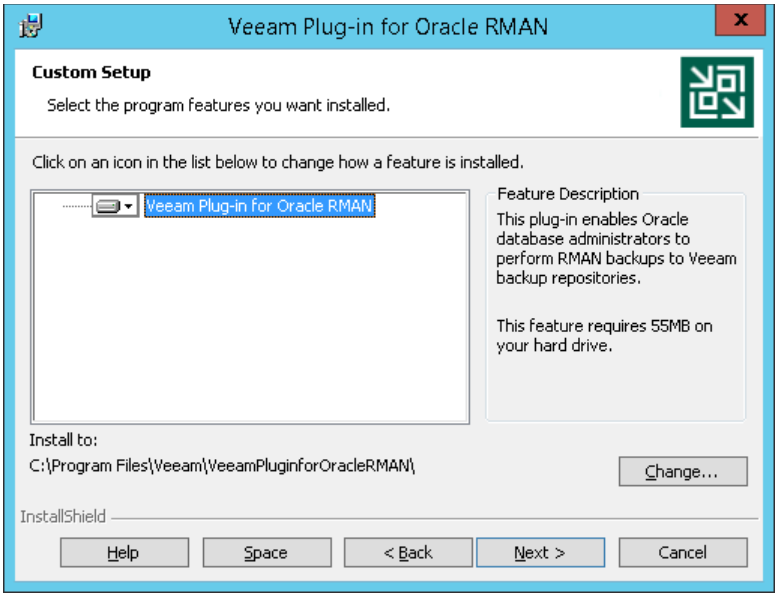

7. At the **Ready to Install the Program** step of the wizard, click **Install**.

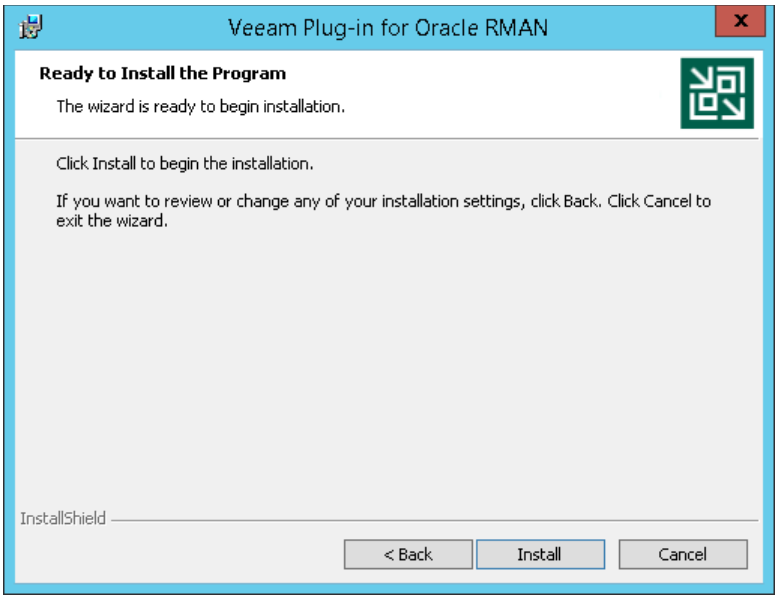

Installing the Plug-in in Unattended Mode

You can install Veeam Plug-in for Oracle RMAN on a Windows machine in the unattended mode using the command line. Go to folder where the `VeeamPluginforOracleRMAN.exe` file resides and run the following command:

```
<path_to_exe>\VeeamPluginforOracleRMAN.exe /silent /accepteula /acceptthirdpartylicenses
```

where <path_to_exe> is a path to the Veeam Plug-in for Oracle RMAN installation file.

/silent	Enables the silent mode.
/accepteula	Accepts EULA terms.
/acceptthirdpartylicenses	Accepts terms of third-party licenses.

Veeam Plug-in for Oracle RMAN uses the following codes to report about the installation results:

- 1000 — Veeam Plug-in for Oracle RMAN has been successfully installed.
- 1001 — prerequisite components required for Veeam Plug-in for Oracle RMAN have been installed on the machine. Veeam Plug-in for Oracle RMAN has not been installed. The machine needs to be rebooted.
- 1002 — Veeam Plug-in for Oracle RMAN installation has failed.
- 1101 — Veeam Plug-in for Oracle RMAN has been installed. The machine needs to be rebooted.

Installing Plug-in on Oracle Solaris

Veeam Plug-in for Oracle RMAN is an additional component of Veeam Backup & Replication, and the installation package of the plug-in is included in the Veeam Backup & Replication installation ISO file.

Veeam Plug-in for Oracle RMAN must be installed on a machine where the target Oracle Database is deployed.

> **NOTE**
>
> To install Veeam Plug-in, the /opt/veeam directory must be writable.

To install Veeam Plug-in for Oracle RMAN on a Solaris machine, do the following:

1. Mount the Veeam Backup & Replication installation disk (VeeamBackup&Replication_11.0.0.XXXX.iso).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk at: https://www.veeam.com/backup-replication-vcp-download.html.

2. In the installation disk folder, go to Plugins\Oracle RMAN\Solaris. Select your system: i386 or SPARC.

3. Copy the Veeam Plug-in installation package (VeeamPluginforOracleRMAN-11.X.X.XXXX-1.SPARC.pkg) to Oracle Solaris server.

4. Install the plug-in from package with root privileges. Make sure the root user has privileges to add the PKG file.

   ```
   pkgadd -d /tmp/VeeamPluginforOracleRMAN-11.X.X.XXXX.SPARC.pkg
   ```

5. Once Veeam Plug-in is installed, you can configure the plug-in settings. For details, see Configuring Plug-in on Linux and Unix Machines.

Installing Plug-in on IBM AIX

Veeam Plug-in for Oracle RMAN is an additional component of Veeam Backup & Replication, and the installation package of the plug-in is included in the Veeam Backup & Replication installation ISO file.

Veeam Plug-in for Oracle RMAN must be installed on a machine where the target Oracle Database is deployed.

> **NOTE**
>
> To install Veeam Plug-in, the `/opt/veeam` directory must be writable.

To install Veeam Plug-in for Oracle RMAN on an IBM AIX machine, do the following:

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.0.0.XXXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk at: https://www.veeam.com/backup-replication-vcp-download.html.

2. In the installation disk folder, go to `Plugins\Oracle RMAN\AIX\ppc64`.

3. Copy the Veeam Plug-in installation package (`VeeamPluginforOracleRMAN-11.X.X.XXX-1.aix6.1.ppc.rpm`) to the AIX server where the target Oracle database is deployed.

4. Install the plug-in from package with root privileges. Make sure the root user has privileges to add the PKG file.

   ```
   rpm -i VeeamPluginforOracleRMAN-11.X.X.XXXX-1.aix6.1.ppc.rpm
   ```

5. Once Veeam Plug-in is installed, you can configure the plug-in settings. For details, see Configuring Plug-in on Linux and Unix Machines.

Configuring Veeam Plug-in for Oracle RMAN

By default, RMAN sends backups to a native RMAN location on disk (`DEFAULT DEVICE TYPE TO DISK`). When you configure Veeam Plug-in, the default device type is changed to `SBT_TAPE`, which gives control over backup media management to Veeam Plug-in. Thus, after you deploy Veeam Plug-in on an Oracle server, you can perform all backup and restore operations in the Oracle RMAN console. Veeam Plug-in compresses, deduplicates database backups and transfers them to a backup repository connected to Veeam Backup & Replication.

To use Veeam Plug-in you must configure the connection between the Oracle server, Veeam Backup & Replication server and backup repositories where backup files will be stored.

- Configuring Veeam Plug-in for Oracle RMAN on Linux or Unix Machines
- Configuring Veeam Plug-in for Oracle RMAN on Windows Machines

Configuring Plug-in on Linux or Unix

To configure Veeam Plug-in, you can use **OracleRMANConfigTool**. The tool configures Oracle RMAN integration settings and creates the `veeam_config.xml` file which is stored in the installation folder of the plug-in: `/opt/veeam/VeeamPluginforOracleRMAN`.

Note that the Veeam Plug-in configuration tool changes the settings of Oracle RMAN. All original settings of Oracle RMAN are saved in the `/opt/veeam/VeeamPluginforOracleRMAN/RMANParameters.xml` file.

To configure Veeam Plug-in for Oracle RMAN, do the following:

1. Log in to the Oracle server with an account which is a member of the DBA group.
2. Launch the configuration wizard:

   ```
   OracleRMANConfigTool --wizard
   ```

   If you have extracted files form the .TAR.GZ archive, go to the `/opt/veeam/VeeamPluginforOracleRMAN` folder and run the following command:

   ```
   ./OracleRMANConfigTool --wizard
   ```

3. Specify the DNS name or IP address of the Veeam Backup & Replication server.

   ```
   Enter backup server name or IP address: 172.24.164.68
   ```

4. Specify the port which will be used to communicate with the Veeam Backup & Replication server. Default port: *10006*.

   ```
   Enter backup server port number: 10006
   ```

5. Specify credentials to authenticate against the Veeam Backup & Replication server. The account must have the local Administrator permissions on the Veeam Backup & Replication server.

   Also, note that the account must have access permissions on the required backup repository. To learn how to configure access permissions and encryption settings on repositories, see Access and Encryption Settings on Repositories.

```
Enter username: serv17\administrator
Enter password for serv17\administrator:
```

   IMPORTANT

   You can work with backups created by Veeam Plug-in only with the account used for creating the backups. If you want to use another account, assign the *Veeam Backup Administrator* role or *Veeam Backup Operator* and *Veeam Restore Operator* roles to the account.

   To learn how to assign Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication Guide.

6. Select the backup repository where you want to store the backups. In the wizard dialog, enter the number of the repository from the list of available repositories. If you want to add several repositories, enter the required numbers separated by blank spaces.

   If you want to use the Oracle RMAN Backup Duplexing functionality, you can select up to four repositories. The copies of backups will be sent to all selected repositories. Note that Oracle Database Standard Edition does not allow using more than one RMAN channel. Thus, if you use Standard Edition, you can select only one repository.

```
Available repositories are:
1. serv10_repo
2. serv02_repo
Specify up to 4 Veeam repositories to use as target using whitespace as a
separator: 2
```

   IMPORTANT

   - The account must have access to backup repositories that you plan to use.
   - Encryption must be disabled on the repository.

   Otherwise, the repositories will not be listed as available. To learn how to configure access permissions and encryption settings on repositories, see Access and Encryption Settings on Repositories.

7. Specify the number of parallel data streams for each backup repository. Note that Oracle Database Standard Edition does not allow using more than one RMAN channel. Thus, if you use Standard Edition, you can select only one data stream.

```
Enter the number of data streams (From 1 to 254) to run in parallel for ea
ch repository (RMAN DEVICE PARALLELISM value).
Channel count per device: 4
```

8. If you want to enable Veeam compression of backup files, type **y**. For details, see the Data Compression and Deduplication section of Veeam Backup & Replication User Guide.

```
Do you want to use Veeam compression (Y/n):y
```

9. At the last step of the plug-in wizard, you can export configuration files (the Veeam Plug-in configuration file and RMAN configuration file). You can import these configuration files to other servers to apply the same settings.

```
Save configuration?
1. Apply configuration to the Oracle environment
2. Export configuration into a file for manual setup
3. Cancel without saving
Enter:1
*** Database instance ORCL is configured ***
```

NOTE

When you export the configuration files, Veeam Plug-in automatically enables Oracle's Controlfile Autobackup feature. This feature is required for restoring with different settings using Veeam Explorer for Oracle.

TIP

It is recommended to save the configuration files, so that you can use it as a reference. For example, if you are planning to manually allocate channels for backup and restore operations, you will need the repository UUID. The RMAN configuration file (`rman_config.txt`) contains an example for channel allocation definition for the target repository. You can use this statement in your backup/restore scripts.

Configuration Tool Commands

Apart from running a configuration wizard, you can use the **OracleRMANConfigTool** tool to change a specific parameter in the `veeam_config.xml` file or enable/disable Veeam Plug-in features.

See the list of available commands for **OracleRMANConfigTool**:

Command	Description
--help	Shows the list of parameters of the plug-in configuration tool.
--show-config	Shows configuration parameters.
--wizard	Starts the wizard to configure the plug-in settings. This wizard edits the `veeam_config.xml` file or creates a new one if the configuration file was removed from the `/opt/veeam/VeeamPluginforOracleRMAN` directory.

Command	Description
--set-credentials <"serv\username"> <"password">	Specifies credentials to connect to the Veeam backup server.
--set-host <hostname>	Specifies the host of the Veeam backup server.
--set-port <port_number>	Specifies the host to connect to the Veeam Backup & Replication server.
--set-repositories	Launches a wizard to select a backup repository. A backup repository is selected from repositories which are available in the connected Veeam Backup & Replication instance.
--set-parallelism <number_of_channels>	Configures RMAN parallelism settings.
--compression <y/n>	Enables/disables Veeam proprietary feature which compresses backup files.
--map-backup	Maps the imported backups.
--set-force-delete	Configures the auto-deletion of backup files after specified days.
--configure-restore-from-copy	Enables restore from backup copy. Note that if you enable restore from backup copy, you cannot back up databases with Veeam Plug-in. To revert changes, you must disable restore from backup copy.
--promote-backup-copy-to-primary	Maps the imported backup copy to a regular Veeam Plug-in backup chain.

Example:

To specify credentials that will be used to log in to the Veeam Backup & Replication server, use the plug-in configuration tool with the following command.

```
OracleRMANConfigTool --set-credentials "serv04\joelle" "password"
```

Configuring Plug-in on Windows

To configure backup and restore settings, use the **Veeam Plug-in for Oracle RMAN** configuration wizard. The wizard configures Oracle RMAN settings and creates the `veeam_config.xml` file which is stored in the `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN` folder.

Note that configuration wizard of Veeam Plug-in for Oracle RMAN changes the settings of Oracle RMAN. All original settings of Oracle RMAN are saved in the `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\RMANParameters.xml` file.

To configure Veeam Plug-in, do the following:

1. On the Oracle server, click launch the Veeam RMAN Configuration Wizard (`%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\Veeam.Backup.RMAN.Configuration.exe`).

2. At the **Backup Server** step of the wizard, specify the DNS name of the Veeam Backup & Replication server and credentials that will be used to connect to the server. The specified account must have the local Administrator permissions on the Veeam Backup & Replication server.

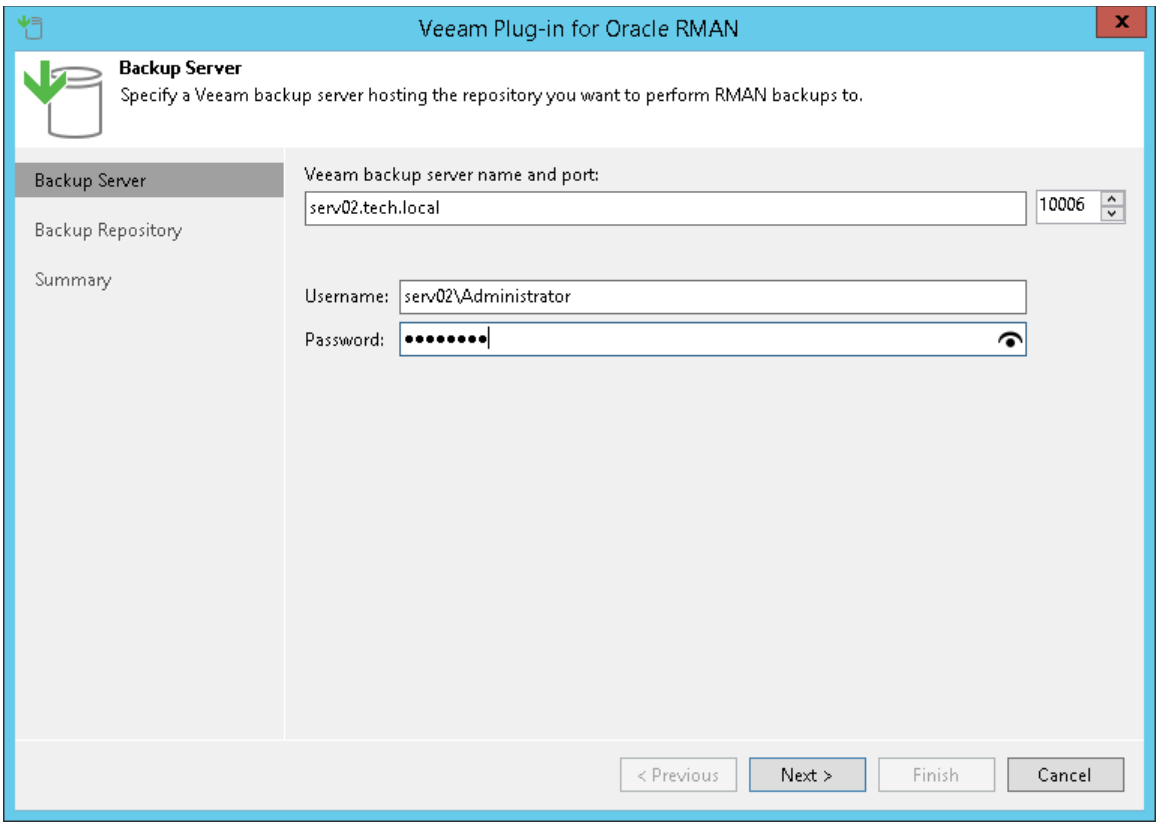

3. At the **Backup Repository** step of the wizard:

   a. Click **Add** and select the required repository. For Oracle Database Standard Edition, you can select only one repository. If you want to use the Oracle RMAN Backup Duplexing functionality, you can select up to four repositories. The copies of backups will be sent to all selected repositories.

      You must allow access to Veeam backup repositories that you plan to use. Also, the encryption on the repository must be disabled. To learn how to configure access and encryption on repositories, see Access and Encryption Settings on Repositories.

   b. If you want to use another repository, select the repository from the list and click **Remove**. Then, you can add another repository

   c. At the **Channels per repository** field, specify the number of allowed parallel data streams for each repository.

   d. If you want to enable Veeam Plug-in compression, select the **Enable backup compression by the plug-in** check box.

      Note that if you use Veeam Plug-in compression in combination with Oracle RMAN integrated compression (BACKUP AS COMPRESSED commands), it can slow down processing.

e. Click **Next**.

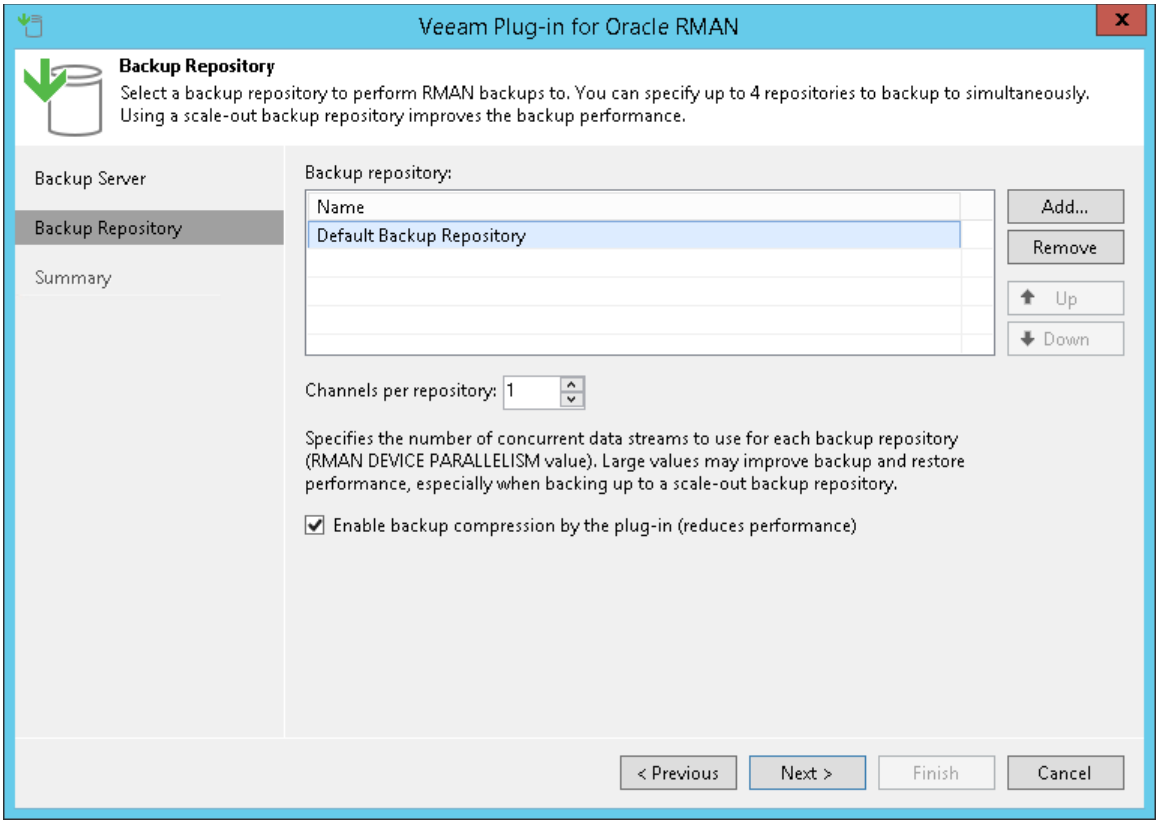

IMPORTANT

You can work with backups created by Veeam Plug-in only with the account used for creating the backups. If you want to use another account, assign the *Veeam Backup Administrator* role or *Veeam Backup Operator* and *Veeam Restore Operator* roles to the account.

To learn how to assign Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication Guide.

4. At the **Summary** step of the wizard, you can export the plug-in and Oracle RMAN configuration files. You can use the configuration files to apply the plug-in settings on other servers.

   a. To export the `veeam_config.xml` file click **Export** and select **Veeam Config**.

b. Click **Finish**.

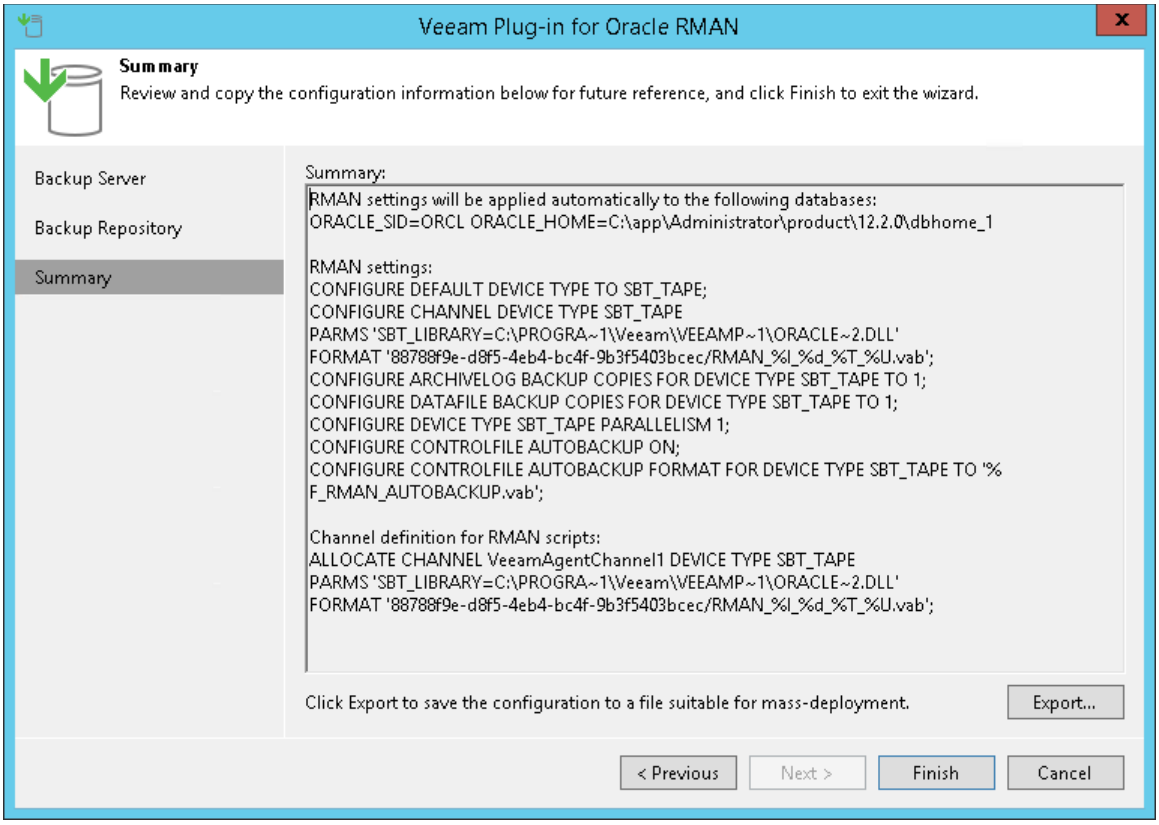

NOTE

When you export the configuration files, Veeam Plug-in automatically enables Oracle's Controlfile Autobackup feature. This feature is required for restoring with different settings using Veeam Explorer for Oracle.

TIP

It is recommended to save the configuration files, so that you can use it as a reference. For example, if you are planning to manually allocate channels for backup and restore operations, you will need the repository UUID. The RMAN configuration file (`rman_config.txt`) contains an example for channel allocation definition for the target repository. You can use this statement in your backup/restore scripts.

Configuration Tool Commands

Apart from running a configuration wizard, you can use the **OracleRMANConfigTool.exe** tool to change a specific parameter in the `veeam_config.xml` file or enable/disable Veeam Plug-in features.

To run a specific command, do the following:

1. On the Oracle server, go to `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN`.

2. Run the required **OracleRMANConfigTool.exe** command from the table below.

   For example, to specify credentials that will be used to log in to the Veeam Backup & Replication server, use the plug-in configuration tool with the following command:

```
OracleRMANConfigTool.exe --set-credentials "serv02\Joelle" "pas
sword"
```

Command	Description
--help	Shows the list of parameters of the plug-in configuration tool.
--show-config	Shows configuration parameters.
--wizard	Starts the wizard to configure the plug-in settings. This wizard edits the `veeam_config.xml` file or creates a new one if the configuration file was removed from the `/opt/veeam/VeeamPluginforOracleRMAN` directory.
--set-credentials <"serv\username"> <"password">	Specifies credentials to connect to the Veeam Backup & Replication server.
--set-host <hostname>	Specifies the host of the Veeam Backup & Replication server.
--set-port <port_number>	Specifies the host to connect to the Veeam Backup & Replication server.
--set-repositories	Launches a wizard to select a backup repository. A backup repository is selected from repositories which are available in the connected Veeam Backup & Replication instance.
--set-parallelism <number_of_channels>	Configures Oracle RMAN parallelism settings.
--compression <y/n>	Enables/disables Veeam proprietary feature which compresses backup files.
--map-backup	Maps the imported backups.
--set-force-delete	Deletes backup files after specified days.
--configure-restore-from-copy	Enables restore from backup copy. Note that if you enable restore from backup copy, you cannot back up databases with Veeam Plug-in. To revert changes, you must disable restore from backup copy.

	`OracleRMANConfigTool.exe --set-credentials "serv02\Joelle" "password"`
--promote-backup-copy-to-primary	Maps the imported backup copy to a regular Veeam Plug-in backup chain.

Upgrading Veeam Plug-in for Oracle RMAN

Periodically, Veeam releases a new version of Veeam Backup & Replication that contains new features and bug fixes. The release package also contains a new version of Veeam Plug-ins.

If you want to upgrade Veeam Plug-in, note that Veeam Backup & Replication must be the same or newer that the version of Veeam Plug-in. If you want to use the latest functionality, you must upgrade both Veeam Backup & Replication and to the latest version. After the upgrade, you don't need to to re-run the Veeam Plug-in configuration wizard, the plug-in configuration files will be preserved.

> **IMPORTANT**
>
> The Veeam Backup & Replication version must be the same or newer that the Veeam Plug-in version. First, you must upgrade Veeam Backup & Replication, then you can upgrade Veeam Plug-ins. For instructions on how to upgrade Veeam Backup & Replication, see the Upgrading to Veeam Backup & Replication 11 section of the Veeam Backup & Replication User Guide.

Before You Begin

Veeam Plug-in installation files are included in the installation disk image of Veeam Backup & Replication. You must upload the installation file to the Oracle Database server. To do this, perform the following steps.

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.X.X.XXXX.iso`).
2. Open the mounted disk image and go to the `Plugins\Oracle RMAN\Linux` directory.
3. Select the the Veeam Plug-in installation file suitable for your OS and upload it to the Oracle Database server.

For instructions on how to upgrade Veeam Plug-in for Oracle RMAN, see the following guides:

- Upgrading Plug-in on Microsoft Windows
- Upgrading Plug-in on Linux (RPM)
- Upgrading Plug-in on Linux (TAR.GZ)
- Upgrading Plug-in on Oracle Solaris
- Upgrading Plug-in on IBM AIX

Upgrading Plug-in on Windows

To upgrade Veeam Plug-in for Oracle RMAN on a Windows machine, upload the `VeeamPluginforOracleRMAN.exe` file of the new version to the Oracle Database server and install it. The old version will be replaced with the new version automatically.

You can upgrade Veeam Plug-in for Oracle RMAN to a newer version in the unattended mode using the same command that is used for unattended installation. For details, see Installing Veeam Plug-in on Windows Machine in Unattended Mode.

Upgrading Plug-in on Linux (.RPM)

To upgrade Veeam Plug-in for Oracle RMAN on a Linux machine, do the following:

1. Upload the `VeeamPluginforOracleRMAN-11.X.X.XXXX.x86_64.rpm` package to the Oracle server. If you need the 32-bit version, choose the `i386` package.
2. To upgrade Veeam Plug-in, run the following command. Note that the operation requires *root* privileges.

```
rpm -U VeeamPluginforOracleRMAN-11.X.X.XXXX.x86_64.rpm
```

> **TIP**
>
> To find out which version of Veeam Plug-in is installed on your server, you can use the following command: `rpm -qa | grep VeeamPlugin*`

Upgrading Plug-in on Linux (.TAR.GZ)

To upgrade Veeam Plug-in for Oracle RMAN on a Linux machine from the .TAR.GZ archive, do the following:

1. Upload the `VeeamPluginforOracleRMAN.tar.gz` file to the Oracle server.
2. Unpack the plug-in files from the archive to the `/opt/veeam` directory. Old Veeam Plug-in files will be replaced by new files.

```
tar -xzvf VeeamPluginforOracleRMAN.tar.gz -C /opt/veeam
```

Upgrading Plug-in on Oracle Solaris

To upgrade Veeam Plug-in for Oracle RMAN on an Oracle Solaris machine, do the following:

1. Upload the `VeeamPluginforOracleRMAN-11.X.X.XXXX.SPARC.pkg` package to the Oracle server. If you need the 32-bit version, choose the `i386` package.
2. Make sure the `pkgadd` administration file (`admin_file`) contains the following entry: `"instance=overwrite"`. For details, see the Avoiding User Interaction When Adding Packages section of the Oracle Solaris Administration Guide.
3. To upgrade Veeam Plug-in, run the following command:

```
pkgadd -a admin_file -d /tmp/VeeamPluginforOracleRMAN-11.X.X.XXXX.SPARC.pkg
```

Upgrading Plug-in on IBM AIX

To upgrade Veeam Plug-in for Oracle RMAN on an IBM AIX machine, do the following:

1. Upload the `VeeamPluginforOracleRMAN-11.X.X.XXX-1.aix6.1.ppc.rpm` package to the Oracle server.

2. To upgrade Veeam Plug-in, run the following command. Note that the operation requires *root* privileges.

```
rpm -U VeeamPluginforOracleRMAN-11.X.X.XXX-1.aix6.1.ppc.rpm
```

Importing/Exporting Plug-in Settings

You can export a Veeam Plug-in configuration file and apply the plug-in settings to other severs.

> **IMPORTANT**
>
> The password included in the configuration file is encrypted. Thus, after you import the configuration file, you must set the credentials manually in the Veeam Plug-in configuration wizard.

To export the configuration file to another Linux server, do the following:

[For Linux and Unix]:

a. On the server where Veeam Plug-in is installed, go to `/opt/veeam/VeeamPluginforOracleRMAN`.

b. Copy the `veeam_config.xml` file to the server where you want to configure the plug-in.

c. Install Veeam Plug-in on the new server and place the copied XML file in the `/opt/veeam/VeeamPluginforOracleRMAN` folder.

d. Set new credentials using the following command:

```
OracleRMANConfigTool --set-credentials "serv\username" "password"
```

[For Windows]:

a. On the server where Veeam Plug-in is installed, go to `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\`.

b. Copy the `veeam_config.xml` file to the server where you want to configure the plug-in.

c. Install Veeam Plug-in on the new server and place the copied XML file in the `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\` folder.

d. Set new credentials using the following command:

```
C:\Program Files\Veeam\VeeamPluginforOracleRMAN\OracleRMANConfigTool.exe --set-credentials "serv\username" "password"
```

Importing Backup Files

If the Veeam Backup & Replication server has failed and you have restored it in a new location, you can copy the backup files to a new repository and re-map the Veeam Plug-in backup files.

Limitations and Prerequisites

Mind the following limitations:

- If backup files are not imported according to instructions given in this section, Veeam Plug-in backup and restore operations may fail.
- The repository from which you plan to import backups must be added to the Veeam Backup & Replication infrastructure. Otherwise you will not be able to access backup files.
- If you are importing backup files from a scale-out backup repository, the names of backup files and paths to backup files must contain only allowed characters:
  - Alphanumeric characters: `a-zA-Z0-9`
  - Special characters: `_-.+=@^`
  - Names of backup files and paths to backup files must not contain spaces.

How to Import Veeam Plug-in Backup Files

To import Veeam Plug-in backup files, do the following:

1. Move the folder with the backup file to the required backup repository or create a new backup repository with this folder as a subfolder.

   > **TIP**
   >
   > Starting from version 11, each Veeam Plug-in backup file (.vab) has its own metadata file (.vasm). Make sure you import backup files and all related metadata files. Also, you must import the backup job metadata file (.vacm) which is stored in the same folder.

2. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
3. In the inventory pane of the **Backup Infrastructure** view, select the **Backup Repositories** node.

4. In the working area, select the required backup repository and click **Rescan** on the ribbon. Alternatively, you can right-click the backup repository and select **Rescan**.

   During the rescan operation, Veeam Backup & Replication gathers information about backups that are currently available in the backup repository and updates the list of backups in the configuration database. After the rescan operation, backups that were not in this configuration database will be shown on the **Home** view in the **Backups > Disk (Imported)** node.

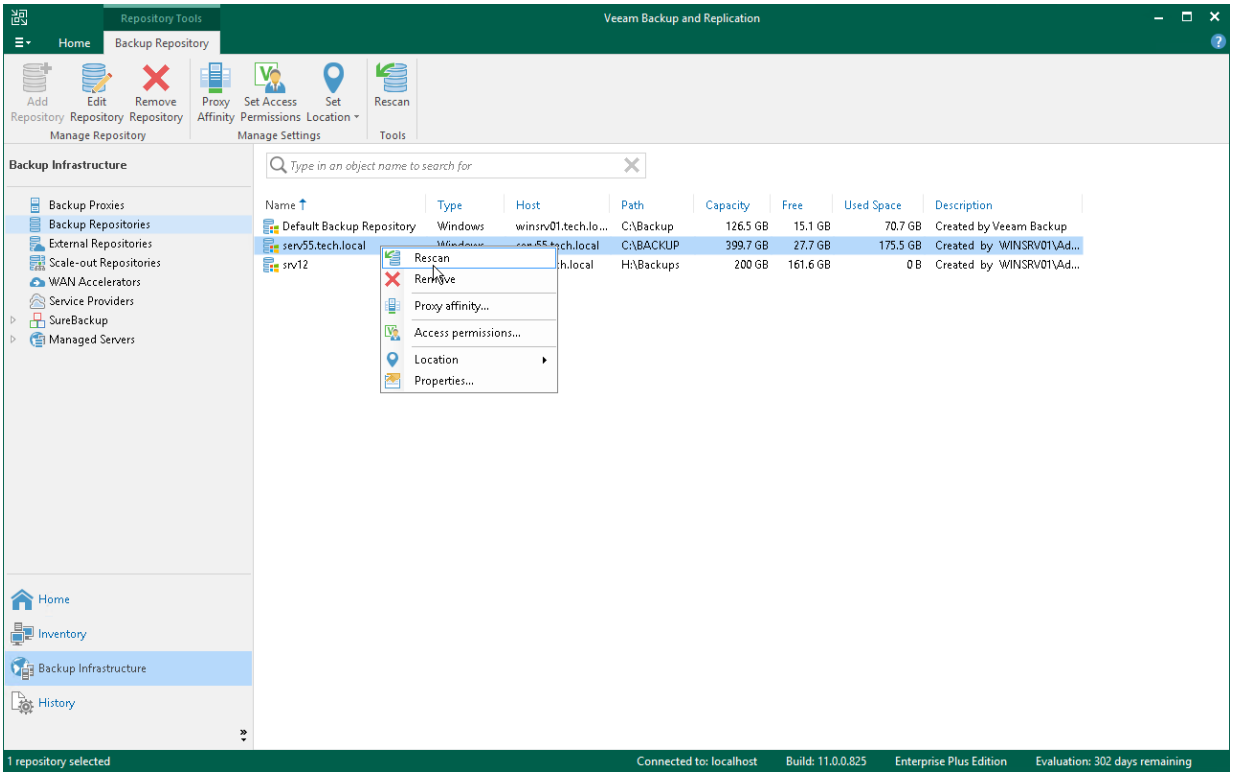

5. On the Oracle server, use the terminal to set the new repository as a target in the Veeam Plug-in settings:

   - **For Windows:**

     ```
     "C:\Program Files\Veeam\VeeamPluginforOracleRMAN\OracleRMANConfigTool"
     --set-repositories
     Available backup repositories:
     1. serv55.tech.local
     2. serv07_repo
     Enter repository number: 1
     ```

   - **For Linux and Unix:**

     ```
     OracleRMANConfigTool --set-repositories
     Available backup repositories:
     1. serv55.tech.local
     2. serv07_repo
     Enter repository number: 1
     ```

6. Start the Veeam Plug-in configuration wizard with the `--map-backup` parameter to re-map the plug-in backups:
   - **For Windows:**

     ```
     "C:\Program Files\Veeam\VeeamPluginforOracleRMAN\OracleRMANConfigTool"
     --map-backup
     ```

   - **For Linux and Unix:**

     ```
     OracleRMANConfigTool --map-backup
     ```

Upgrading Backup Files

Since version 11, Veeam Plug-in uses a new format of backup files: instead of one metadata file for all backup files there are separate metadata files (.vasm) for each database backup file (.vab). The new metadata format allows to optimize the productivity of backup and restore operations.

For Veeam Plug-in 11, the backup files upgrade is not obligatory. However, in the next version, backup files created by Veeam Plug-in version 10 will not be supported.

> **IMPORTANT**
>
> If you do not upgrade backup files, you will get the following warning in the job session logs: *Backup metadata is not up to date. Please upgrade the backup*. If you want to disable the warning, see instructions in this Veeam KB.

Prerequisites

Before upgrading backup files, make sure the following requirements are met.

- Make sure that you have upgraded Veeam Plug-in on the source server. If the plug-in is not upgraded to version 11 and you upgrade the backup files, then all next backup job runs will fail.
- Make sure that you have disabled the backup job whose backup files you want upgrade. You must also disable the backup copy jobs that use these backup files as a source.
- If the backup files reside on the scale-out backup repository, all repository extents must be available. Also, the extents must not be in the Seal or Maintenance mode.
- If you want to upgrade backup files created by a backup copy job, you must meet the same requirements as for the backup job files.
- During the process of the metadata upgrade, you cannot run the target backup job and you cannot restore from the backup files.

  The upgrade process duration depends on the number of backup files in the backup set, type of the backup repository and workload level on the file system.

  For example, there are backup files of an Oracle server that contains 10 instances and is backed up every 15 minutes with the retention policy set for 2 weeks. The upgrade of backup files can have the following duration on not overloaded file systems:

  - Microsoft Windows: 30 minutes
  - Linux: from 30 minutes to 3 hours
  - SMB/NFS: 1.5 hours
  - Data Domain Boost/Quantum DXi/ExaGrid/CIFS (SMB)/NFS file share: 3-4 hours
  - HPE StoreOnce: up to 10 hours (due to specifics of this repository type for processing large number of files)

Upgrading Backup Files in Veeam Backup & Replication Console

To upgrade backup files in the Veeam Backup & Replication console, do the following:

1. Open the **Home** view.
2. In the inventory pane, expand the **Backups** view and select **Disk**.
3. In the working area, right-click the job or the restore point and select **Upgrade**.

   Alternatively, you can select the job or the restore point and click **Upgrade** on the ribbon.

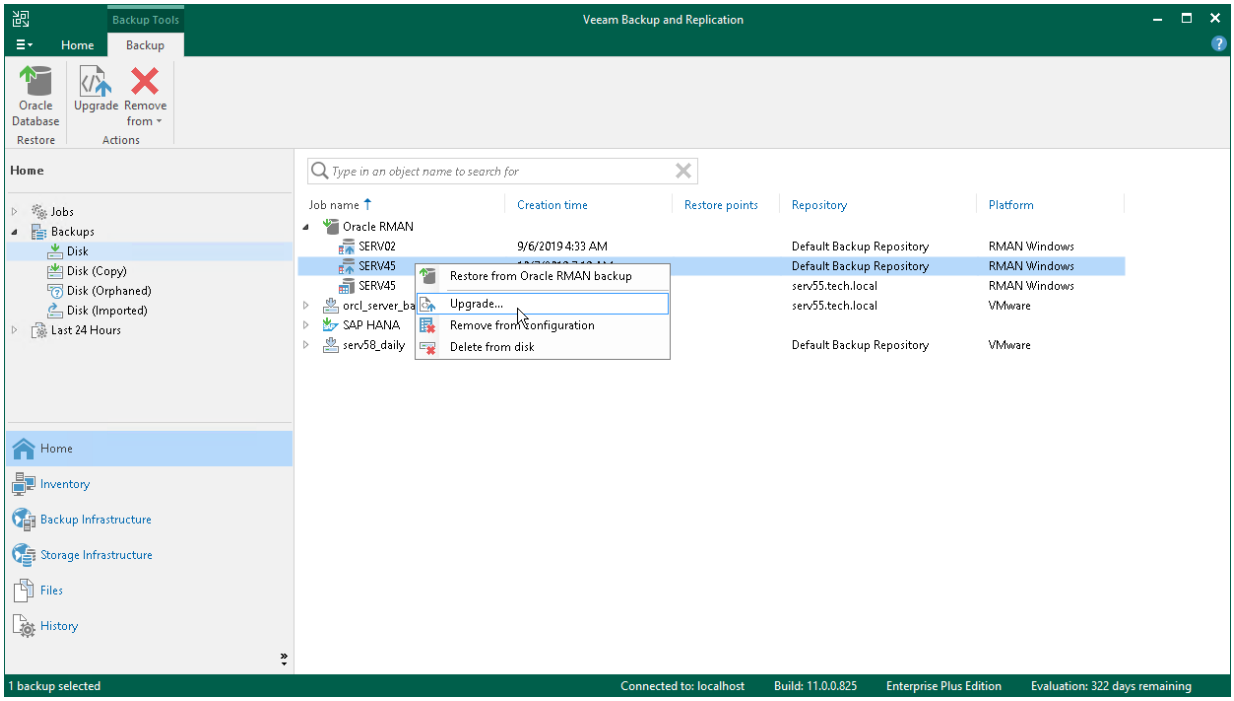

Uninstalling Veeam Plug-in for Oracle RMAN

You can uninstall the Veeam Plug-in and undo the configuration changes made by Veeam Plug-in.

When you configure Veeam Plug-in, original settings of Oracle RMAN are saved in the `/opt/veeam/VeeamPluginforOracleRMAN/RMANParameters.xml` file (or in the `C:\Program Files\Veeam\VeeamPluginforOracleRMAN\RMANParameters.xml` for Windows). If you uninstall Veeam Plug-in for Oracle RMAN, original settings of Oracle RMAN are restored from the `RMANParameters.xml` file.

Uninstalling Veeam Plug-in on Linux or IBM AIX Machines

To uninstall Veeam Plug-in and undo the configuration changes, run the following command. Note that the operation requires *root* privileges.

```
rpm -e VeeamPluginforOracleRMAN
```

Uninstalling Veeam Plug-in on Windows Machines

To uninstall Veeam Plug-in and undo the configuration changes, do the following:

1. Open the **Control Panel** and click **Programs and Features**.
2. In the list of programs, select **Veeam Plug-in for Oracle RMAN** and click **Uninstall**.

Uninstalling Veeam Plug-in on Solaris Machines

To uninstall Veeam Plug-in and undo the configuration changes, run the following command:

```
pkgrm VeeamPluginforOracleRMAN
```

Database Protection

After you configure Veeam Plug-in settings, you can use the Oracle RMAN functionality to back up databases. Veeam Plug-in will automatically transfer the backup files to the Veeam backup repository.

The examples given below are for demonstration purposes only. The backup process is performed on the Oracle RMAN side. Consider configuring required RMAN-specific parameters that may affect the backup process. For details on the backup functionality of Oracle RMAN, see the Backing Up the Database section of the Oracle's Database Backup and Recovery User's Guide.

NOTE

Mind the following:

- In the Veeam Plug-in configuration wizard, you can enable/disable Veeam Plug-in Data Compression and Deduplication. If you enable the Veeam Plug-in compression, do not use Oracle RMAN integrated compression as well. It can slow down the backup and restore processes.
- It is Oracle's best practice to add the `EXIT;` command at the bottom of the script to shut down the RMAN utility. Without the `EXIT;` command in the script, it is up to Oracle RMAN to decide when to close the backup session, which can lead to multiple unclosed RMAN backup sessions.

TIP

If you have configured the retention policy, run the `DELETE OBSOLETE` command after the database backup to delete obsolete backups from the repository.

In this Section

- Oracle RMAN Full Backup
- Oracle RMAN Channel Allocation
- Backup Job in Veeam Backup & Replication

Oracle RMAN Full Backup

After you configure Veeam Plug-in settings, you can use the Oracle RMAN functionality to back up databases. Veeam Plug-in will automatically transfer the backup files to the Veeam backup repository. You can create a consistent backup of Oracle databases in the ARCHIVELOG mode and in the NOARCHIVELOG mode. For details on the backup process in different modes, see the Choosing Between NOARCHIVELOG and ARCHIVELOG Mode section of the Oracle's Database Administrator's Guide.

> **NOTE**
>
> The examples given below are for demonstration purposes only. The backup process is performed on the Oracle RMAN side. Consider configuring required RMAN-specific parameters that may affect the backup process. For details on the backup functionality of Oracle RMAN, see the Backing Up the Database section of the Oracle Database Backup and Recovery User's Guide.

Consistent Backup of Oracle Database in ARCHIVELOG Mode

To create a consistent backup of an Oracle database and redo logs in the ARCHIVELOG mode, run the following script. In this example, Oracle RMAN will back up the entire database and available archived redo logs. The current online redo log will be archived to make sure all redo changes are transferred to the archived redo log chain. In the ARCHIVELOG mode, there will be no downtime as you do not have to shut down the database.

```
rman TARGET /
RUN {
BACKUP DATABASE PLUS ARCHIVELOG;
}
EXIT;
```

Consistent Backup of Oracle Database in NOARCHIVELOG Mode

To create a consistent backup of an Oracle database operating in the NOARCHIVELOG mode, start the Oracle RMAN console and run the following script. In this example, the database instance will be started after the backup process is complete. Note that the database will be unavailable during the backup.

```
rman TARGET /
RUN {
SHUTDOWN TRANSACTIONAL;
STARTUP MOUNT;
BACKUP DATABASE;
STARTUP;
}
EXIT;
```

Oracle RMAN Channel Allocation

If you want to manually allocate channels for backup operations, you must specify the Veeam backup repository UUID in the channel parameters. The `ALLOCATE CHANNEL` command must be issued within a RUN block. It allocates a channel only in the block where the command is issued. See the following example.

```
RUN {
     ALLOCATE CHANNEL ch1 DEVICE TYPE SBT_TAPE PARMS 'SBT_LIBRARY=/opt/veeam/Ve
eamPluginforOracleRMAN/libOracleRMANPlugin.so' FORMAT 'd8338780-1aec-4c36-b17c-
e1ea3ea2ca93/RMAN_%I_%d_%T_%U.vab';
     BACKUP DATABASE;
     RELEASE CHANNEL ch1;
}
EXIT;
```

Run the Oracle RMAN script with the following parameters:

1. Use the `ALLOCATE CHANNEL` command to manually allocate a channel or channels between RMAN and the database instance. Specify the following parameters:
   a. Specify the channel ID. For example: `ch1`.
   b. Specify the `SBT_TAPE` option for the `DEVICE TYPE` parameter.
   c. Specify `PARMS` to define other parameters for the `sbt_tape` channel.
   d. Specify which media library must be used for this `sbt_tape` channel. For Linux or Unix, set the path to the `libOracleRMANPlugin.so` file as the `SBT_LIBRARY`. For Windows, set the path to `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\OracleRMANPlugin.dll`.
   e. Use the Veeam backup repository UUID in the argument for the `FORMAT` parameter. You can find the required backup repository UUID in the `rman_config.txt` file saved during the Veeam Plug-in configuration process or in logs.

      [For Linux] If you have exported Veeam Plug-in configuration files, run the following command to see open the configuration file.

      ```
      cat /tmp/rman_config.txt
      ```

      Alternatively, you can find the repository UUID in logs:

      ```
      grep "received repos" /tmp/veeam_plugin_logs/oracle/OracleRMANConfigToo
      l.log | tail
      ```

      [For Windows] If you have exported Veeam Plug-in configuration files, find the channel allocation definition in the configuration file.

      Alternatively, you can find the repository UUID in logs. Go to the `%\ProgramData\Veeam\Backup\RmanPluginLogs\SERV_NAME` directory and search for "*received repos id*" in the `OracleRMANConfigToolLib.log` file.
2. Use the `BACKUP` command with required parameters to create a database backup.
3. [Optional] Use the `RELEASE CHANNEL` command. By default, RMAN automatically releases all normal channels when the `RUN` command terminates.

Backup Job in Veeam Backup & Replication

After you start a backup process in Oracle RMAN, Veeam Backup & Replication creates a backup job. You can use this job to view the statistics on the backup process, generate backup job reports or you can also disable the backup job.

You cannot start or edit Oracle RMAN backup jobs in the Veeam Backup & Replication console. You can manage backup operations only on the Oracle side using RMAN.

> **NOTE**
>
> The progress bar of a running Oracle database backup job is available only for backups of standalone Oracle databases. It is not available for Oracle RAC backups.

Viewing Backup Job Statistics

To view details of a backup job process, do the following:

1. Open the Veeam Backup & Replication console.
2. In the **Home** view, expand the **Jobs** node and click **Backup**.
3. In the list of jobs, select the Oracle RMAN backup job to see details of the current backup process or the last backup job session.

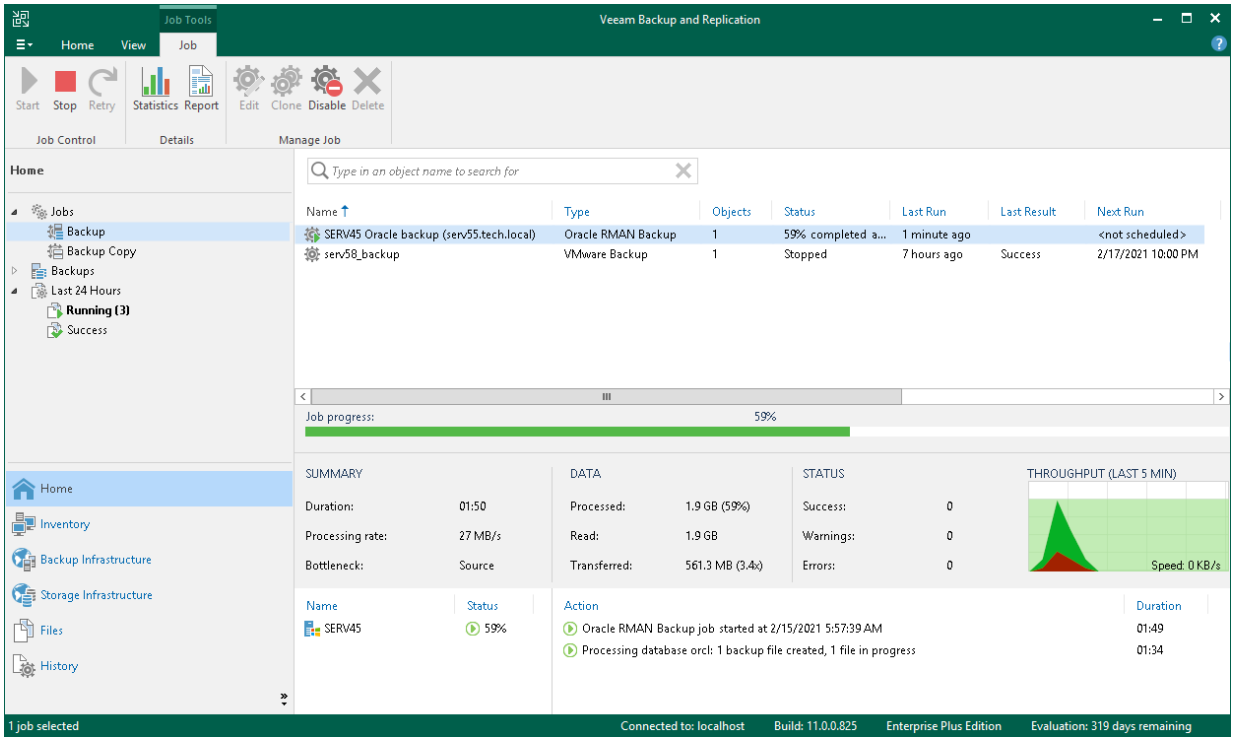

Generating Backup Job Reports

Veeam Backup & Replication can generate reports with details about an Oracle RMAN backup job session performance. The session report contains the following session statistics: session duration details, details of the session performance, amount of read, processed and transferred data, backup size, compression ratio, list of warnings and errors (if any).

To generate a report, do the following:

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

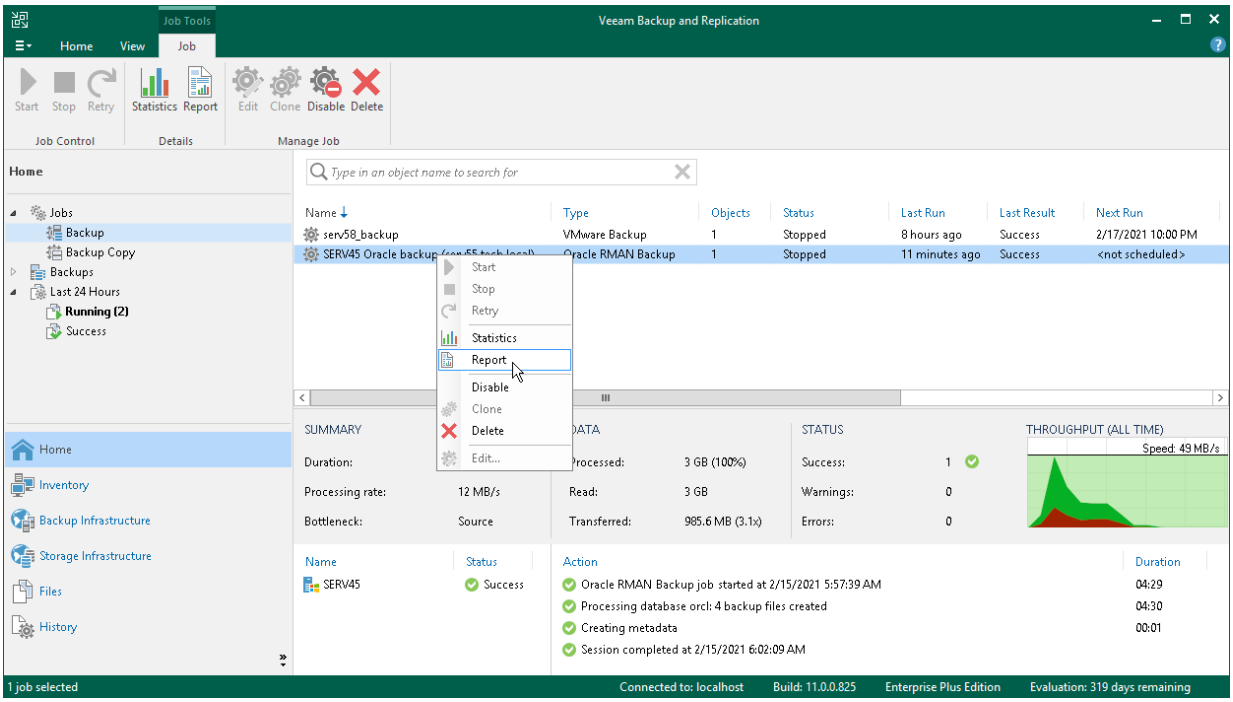

Disabling Backup Job

You can disable Oracle RMAN backup jobs in the Veeam Backup & Replication console. If you disable the job, you will not be able to run RMAN backup commands on the Oracle server.

To disable a backup job, do the following:

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

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

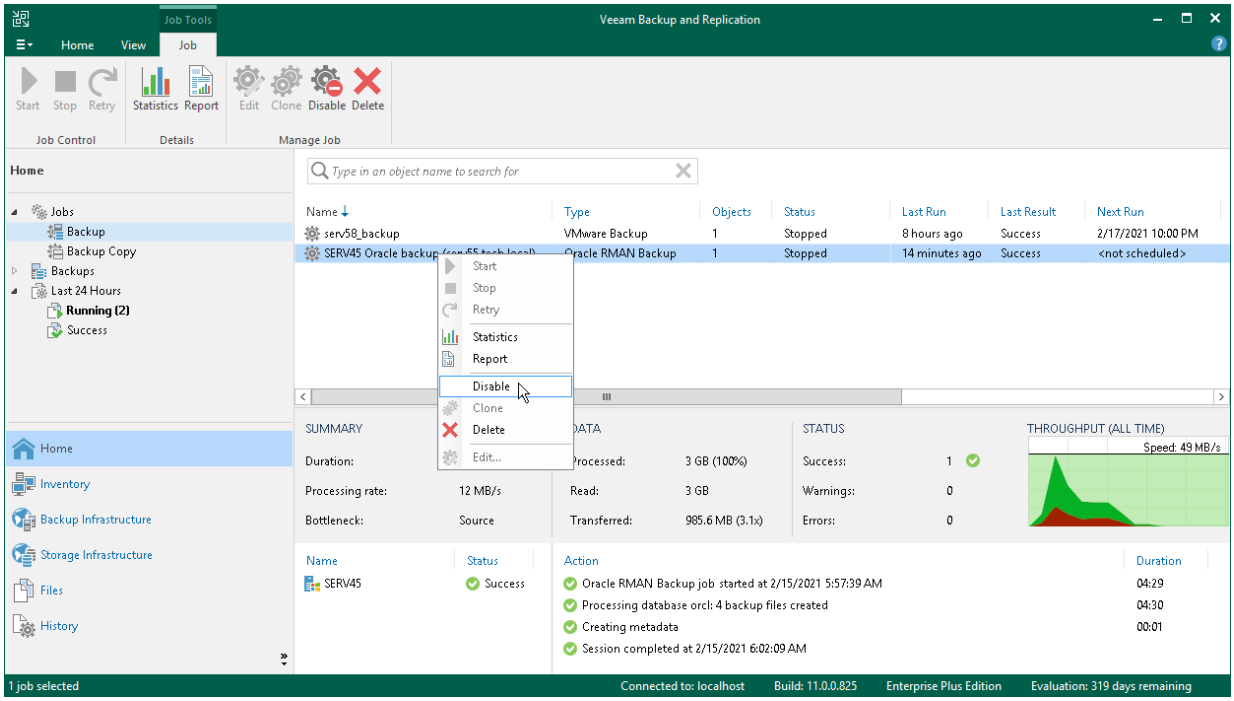

Database Recovery

With the configured Veeam Plug-in for Oracle RMAN, you can perform all kinds of database restore operations available in Oracle RMAN. You can also restore from database backups in the Veeam Backup & Replication console.

- Restore to Original Server
- Restore to Another Server
- Restore from Backup Copy
- Restore with Veeam Explorer for Oracle
- Restore of Control File from Autobackup

Restore to Original Server

Veeam Plug-in for Oracle RMAN allows you to restore databases using built-in Oracle RMAN functionality. When you launch a restore, RMAN restores the necessary database from the backup stored in the Veeam backup repository.

If you want to change the repository or channel settings, you must modify the Veeam Plug-in settings. For details, see Configuring Veeam Plug-in for Oracle RMAN.

To restore the Oracle database, you must connect to the database with RMAN and run the restore command. You may need to run additional commands depending on your database infrastructure. Consider configuring required RMAN-specific parameters that may affect the backup process. For details on all restore capabilities of Oracle RMAN, see the Performing Complete Database Recovery section of the Oracle's Database Backup and Recovery User's Guide.

```
rman TARGET /
RUN {
SHUTDOWN IMMEDIATE;
STARTUP MOUNT;
RESTORE DATABASE;
RECOVER DATABASE;
}
EXIT;
```

NOTE

If you use the SEND command on the target server to point to the source server, you can run operations like DUPLICATE. For details, see Restore to Another Server Using RMAN.

Restore to Another Server

If you want to restore Oracle databases from a Veeam Plug-in backup to another server, follow instructions in the Restoring a Database on a New Host section of Oracle's Database Backup and Recovery User's Guide and mind the specifics described in this section.

Restore from Backup to Another Server

After you allocate channels, you must use the SEND command with the `srcSrv=originalServerName` parameter, where `originalServerName` is the hostname of the protected server. In case of RAC, use the cluster SCAN name as `originalServerName`.

> **TIPS**
>
> Mind the following:
>
> - Veeam Plug-in backup job name contains the name of the original server. You can find out what to use as `originalServerName` if you look at the first part of the Veeam Plug-in backup job name.
> - The ALLOCATE CHANNEL and SEND commands must be issued only within a RUN block for a specific restore operation.

Example of a script for restoring Oracle database to another server using the SEND command:

```
rman TARGET /
RUN {
ALLOCATE CHANNEL a1 TYPE sbt_tape PARMS "SBT_LIBRARY=/opt/veeam/VeeamPluginforO
racleRMAN/libOracleRMANPlugin.so" SEND "srcSrv=server01";
RESTORE DATABASE;
RECOVER DATABASE;
}
EXIT;
```

Use the `ALLOCATE CHANNEL` command to manually allocate a channel or channels between RMAN and the database instance. Specify the following parameters:

1. Assign an ID for the channel. For example: `ch1`.
2. Specify the `SBT_TAPE` option for the `DEVICE TYPE` or `TYPE` parameter.
3. Specify `PARMS` to define other parameters for the `sbt_tape` channel.
4. Specify which media library must be used for this `sbt_tape` channel. For Linux or Unix, set the path to the `libOracleRMANPlugin.so` file as the `SBT_LIBRARY`. For Windows, set the path to `%PROGRAMFILES%\Veeam\VeeamPluginforOracleRMAN\OracleRMANPlugin.dll`.
5. Use the `SEND` command to specify the original server hostname (the `srcSrv` parameter). For example: `"srcSrv=server01"`.

Restore from Backup Copy to Another Server

If you want to restore a database from a backup copy to another server, do the following:

1. On the Oracle server, go to the `\Veeam\VeeamPluginforOracleRMAN` folder (`/opt/veeam/VeeamPluginforOracleRMAN/` for Linux OS).
2. Open the `veeam_config.xml` file with a text editor.
3. Change the `<PluginParameters />` line as follows:
   - For a standalone Oracle Server:

     ```
     <PluginParameters customServerName="original_server_hostname" />
     ```

   - For Oracle RAC:

     ```
     <PluginParameters customServerName="original_cluster_scan_name" />
     ```

     IMPORTANT

     Mind the following:

     - Veeam Plug-in for Oracle RMAN must be installed and configured on the target server.
     - Before you perform any other backup or restore job, revert back the changes in the `veeam_config.xml` file. The default configuration is: `<PluginParameters />`
4. Enable the **restore from backup copy** option, as described in Restore from Backup Copy.
5. Perform the restore.

```
rman TARGET /
RUN {
ALLOCATE CHANNEL a1 TYPE sbt_tape PARMS "SBT_LIBRARY=/opt/veeam/VeeamPlugi
nforOracleRMAN/libOracleRMANPlugin.so";
RESTORE DATABASE;
RECOVER DATABASE;
}
EXIT;
```

Restore from Backup Copy

You can restore from backups and backup copies. To restore from backup copies, you must enable the **restore from backup copy** option in the Veeam Plug-in wizard.

If the **restore from backup copy** option is enabled, you cannot back up databases using Veeam Plug-in, and you cannot restore from backups created by primary Veeam Plug-in backup jobs. You can restore only from backup copy files until you disable the **restore from backup copy** option.

> **IMPORTANT**
>
> Mind the following:
>
> - You can restore from Veeam Plug-in backups using Veeam Explorer for Oracle, even if the restore from backup copy option is enabled. For details, see Veeam Explorers User Guide.
> - You can restore only from backup copies that are located in repositories with primary backups. If you do not have access to repositories with primary backups, you can convert a backup copy to a primary backup. For details, see Converting Backup Copy to Backup.
> - If you want to restore from a backup copy to another server, see Restore from Backup Copy to Another Server.

Enabling Restore from Backup Copy

To be able to restore from backup copies, do the following.

- **For Linux or Unix OS:**
  a. Open the terminal and launch the following command:

```
OracleRMANConfigTool --configure-restore-from-copy
```

  b. Select the number of the required backup copy job:

```
Select secondary job for failover:
0. Disable
1. Plug-ins backup copy job\SERV02 Oracle backup <serv10_repo>
Select secondary job for failover:1
```

> **IMPORTANT**
>
> The account used to connect to the Veeam Backup & Replication server must have access permissions on the required repository.

- **For Windows OS:**
  a. Go to the Veeam Plug-in installation directory (by default, **C:\Program Files\Veeam\VeeamPluginforOracleRMAN**) and launch the following command:

```
OracleRMANConfigTool.exe --configure-restore-from-copy
```

b. Select the number of the required backup copy job:

```
Select secondary job for failover:
0. Disable
1. Plug-ins backup copy job\SERV02 Oracle backup <serv10_repo>
Select secondary job for failover:1
```

Disabling Restore from Backup Copy

To be able to back up with Veeam Plug-in and restore from backups, disable the restore from backup copies:

- **For Linux or Unix OS:**

```
OracleRMANConfigTool --configure-restore-from-copy
Select secondary job for failover:
0. Disable
1. Plug-ins backup copy job\SERV02 Oracle backup <serv10_repo>
Select secondary job for failover:0
```

- **For Windows OS:**

a. Go to the Veeam Plug-in installation directory (by default, `C:\Program Files\Veeam\VeeamPluginforOracleRMAN\`).

b. Run the following command and select *0* as a secondary job for failover:

```
OracleRMANConfigTool.exe --configure-restore-from-copy
Select secondary job for failover:
0. Disable
1. Plug-ins backup copy job\SERV02 Oracle backup <serv10_repo>
Select secondary job for failover:0
```

Restore with Veeam Explorer for Oracle

You can restore Oracle databases from Veeam Plug-in backups in the Veeam Backup & Replication console. To restore Oracle databases Veeam Backup & Replication uses Veeam Explorer for Oracle. For details, see the Restoring Oracle RMAN Backups section of the Veeam Explorers User Guide.

IMPORTANT

Mind the following:

- Veeam Explorer for Oracle does not support restore of encrypted Oracle databases.
- [For Solaris OS and IBM AIX] Veeam Explorer for Oracle does not support restore of Oracle databases deployed on Solaris OS and IBM AIX. You can restore Oracle databases on Solaris OS and IBM AIX only with RMAN. For details, see the Restore with Oracle RMAN section of the Veeam Plug-ins for Enterprise Applications User Guide.

TIP

Mind the following:

- To perform restore from Oracle databases you can also use Veeam Explorer's cmdlets. For details, see the Veeam Explorer for Oracle section of the Veeam Explorers Powershell Reference.
- For details on Veeam Explorer for Oracle, see the Veeam Explorer for Oracle section of the Veeam Explorers User Guide.

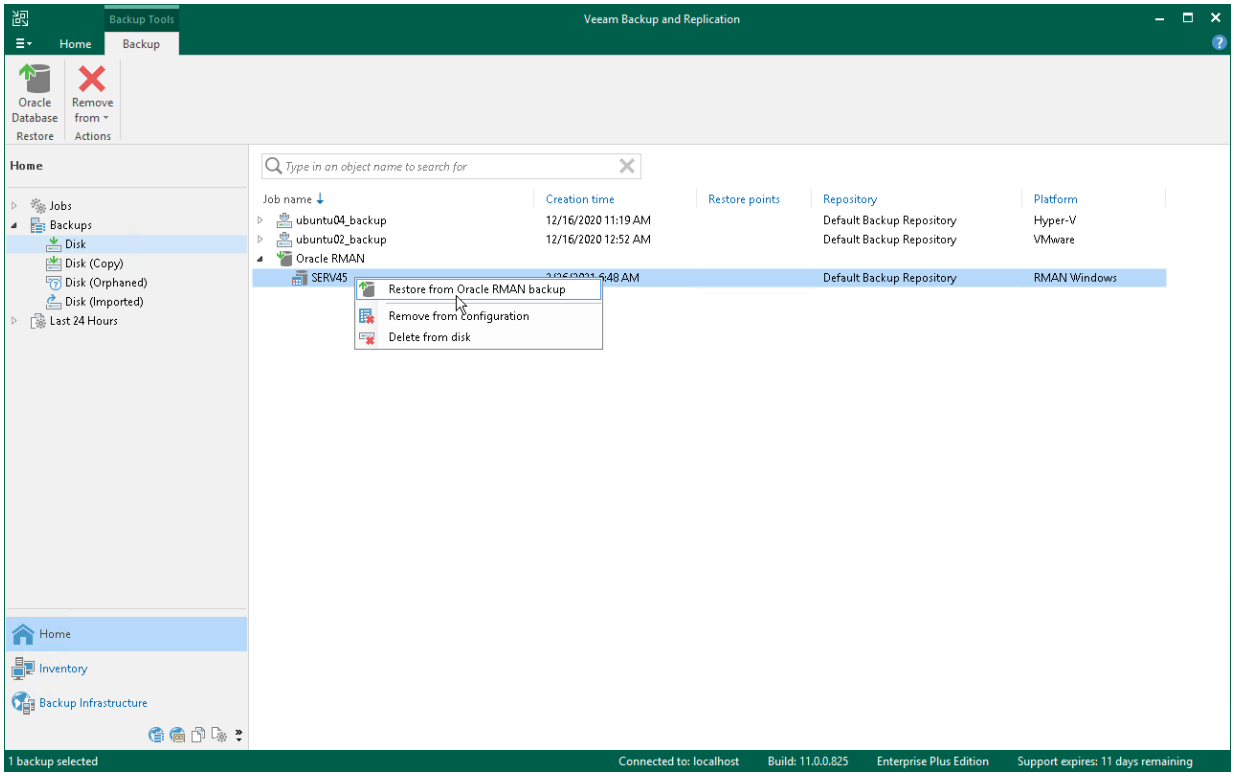

Restore of Control File from Autobackup

You may need to restore the Oracle database control file in the following cases:

- If you want to restore the database to a new location where the control file does not exist
- If the database control file is lost or corrupted

If you use Veeam Plug-in for Oracle RMAN and want to restore the Oracle database control file from autobackup, the autobackup format must be set to the SBT_TAPE device type.

To check if persistent configuration for the control file autobackup format is set to the SBT_TAPE device type, you can run the SHOW ALL or SHOW CONTROLFILE AUTOBACKUP FORMAT commands in the RMAN console. If the persistent configuration is set, you don't need to set the control file autobackup format before the restore command. If it is not set, you must run the SET CONTROLFILE AUTOBACKUP command before the restore process. See the following examples.

> **NOTE**
>
> To restore the control file from autobackup, the database must be in the NOMOUNT state.

Restoring Control File if Persistent Configuration Setting is NOT Set

If the persistent configuration for the control file autobackup format is NOT set to the SBT_TAPE device type, you must set the autobackup format before running the control file restore.

```
RUN {
ALLOCATE CHANNEL c0 DEVICE TYPE sbt PARMS 'SBT_LIBRARY=/opt/veeam/VeeamPluginfo
rOracleRMAN/libOracleRMANPlugin.so';
SET CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO '%F_RMAN_AUTOBA
CKUP.vab';
RESTORE controlfile FROM 'c-4097408439-20200410-00_RMAN_AUTOBACKUP.vab';
}
EXIT;
```

Restoring Control File if Persistent Configuration Setting is Set

If persistent configuration for the control file autobackup format is set to the SBT_TAPE device type, you must set the autobackup format before running the control file restore. To restore the control file, run the following script in the RMAN console:

```
RUN {
ALLOCATE CHANNEL c0 DEVICE TYPE sbt PARMS 'SBT_LIBRARY=/opt/veeam/VeeamPluginfo
rOracleRMAN/libOracleRMANPlugin.so';
RESTORE controlfile FROM 'c-4097408439-20200410-00_RMAN_AUTOBACKUP.vab';
}
EXIT;
```

See Also

- For details on restoring the control file, see the RMAN Restore: Restoring Lost Database Files from Backup section of the Database Backup and Recovery Basics guide.
- For details on the control file autobackup format, see the Configuring the Control File Autobackup Format section of the Database Backup and Recovery Basics.

Retention of RMAN Backups and Archived Logs

In the main scenario, when using Veeam Plug-in for Oracle RMAN, you must configure retention policies for RMAN backups and archived redo logs using native Oracle RMAN functionality:

- Configuring Retention Policy
- Configuring Archived Redo Log Retention

Also, you can manually delete backups from a backup repository using the Veeam Backup & Replication console and enable the force deletion functionality of Veeam Plug-in for Oracle RMAN. For details, see:

- Configuring Force Deletion of Backups
- Deleting Backups Manually
- Deleting Backups from Configuration

Configuring Retention Policy

If you want to edit a retention policy for Oracle RMAN backups, you must connect to the target database in RMAN console and configure one of the following retention policies:

- **Redundancy-Based Retention**: The `redundancy` parameter specifies how many full or level 0 backups of each data file and control file should RMAN keep in the repository. If the number of backups exceeds the specified value, RMAN considers the oldest backups as obsolete. By default, the redundancy parameter is set to 1. To configure a redundancy-based retention policy, run the following command:

```
CONFIGURE RETENTION POLICY TO REDUNDANCY 7;
```

- **Recovery Window-Based Retention**: The `recovery window` parameter specifies the number of days between the current time and the earliest point of recoverability. RMAN does not consider any full or level 0 incremental backups as obsolete if it falls within the recovery window. To configure a window-based retention policy, run the following command:

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 days;
```

Deleting Obsolete Backups

To delete obsolete backups, run the following command after each backup:

```
DELETE OBSOLETE;
```

Disabling Retention Policy

To disable retention policy for RMAN backups, run the following command:

```
CONFIGURE RETENTION POLICY TO NONE;
```

For details, see the Configuring the Backup Retention Policy section of the Oracle Database Backup and recovery User's Guide.

Configuring Archived Redo Log Retention

By default, archivelog deletion policy is disabled.

To configure archivelog deletion policy, run the following command. When the number of archived logs exceeds the specified number, RMAN deletes the oldest archived log.

```
CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 3 TIMES TO SBT;
```

To delete obsolete archivelogs, run the following command. If you do not run this command, the archived logs will be deleted from the fast recovery area (FRA) only.

```
DELETE ARCHIVELOG ALL;
```

To disable archivelog deletion policy, run the following command:

```
CONFIGURE ARCHIVELOG DELETION POLICY TO NONE;
```

For details, see the Configuring an Archived Redo Log Deletion Policy section of the Oracle Database Backup and Recovery User's Guide.

Configuring Force Deletion of Backups

In the main scenario, when using Veeam Plug-in for Oracle RMAN, you must configure the retention policy using native Oracle RMAN tools. For details, see Configuring Retention Policy.

Veeam Plug-in for Oracle RMAN has a functionality that automatically force deletes backup files which are older than specified number of days. For example, you can use it if a backup repository contains backup files that are no longer in the backup catalog.

To enable force deletion of backup files, do the following:

1. On the Oracle server, run the following command.
   - **For Linux and Unix:**

     ```
     OracleRMANConfigTool --set-force-delete
     ```

   - **For Microsoft Windows:**

     ```
     OracleRMANConfigTool.exe --set-force-delete
     ```

     By default, the `OracleRMANConfigTool.exe` file is located in `C:\Program Files\Veeam\VeeamPluginforOracleRMAN`.
2. Enter the number of days after which Veeam Plug-in will force delete backup files on all configured Veeam backup repositories.

   ```
   Garbage collector automatically deletes backup files older than the specif
   ied number of days.
   Make sure the number of days value exceeds your retention policy.
   To disable this functionality, set the number of days to 0.
   Enter the number of days to delete backups after, between 7 and 999 [0]:
   ```

   By default, the force delete functionality is disabled (set to *0*).

IMPORTANT

- A value for the **number of days** setting must be at least 1 backup generation period longer than the retention period for your Oracle Database backups. Otherwise, Veeam Plug-in will delete earliest backups created within the retention period.
- If a backup repository contains backups older than the specified retention period, Veeam Plug-in removes old backup files only after the next run of the RMAN backup.

Deleting Backups Manually

If you want to delete backups files, you can use the Oracle RMAN housekeeping functionality. For details, see the Deleting RMAN Backups and Archived Redo Logs section of the Oracle Database Backup and Recovery User's Guide.

If you have lost the recovery catalog, you can remove the backups manually from a Veeam backup repository.

> **IMPORTANT**
>
> If you remove backups from a Veeam backup repository manually, the metadata about these backups is NOT deleted from the recovery catalog. Thus, if you have a recovery catalog, it is not recommended to manually delete backup files. Otherwise, the recovery catalog will remain in the outdated state.

To remove a backup from Veeam backups repository, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the **Inventory** pane, select **Backups**.
3. In the working area, right-click the backup job object name and select **Delete from disk**.

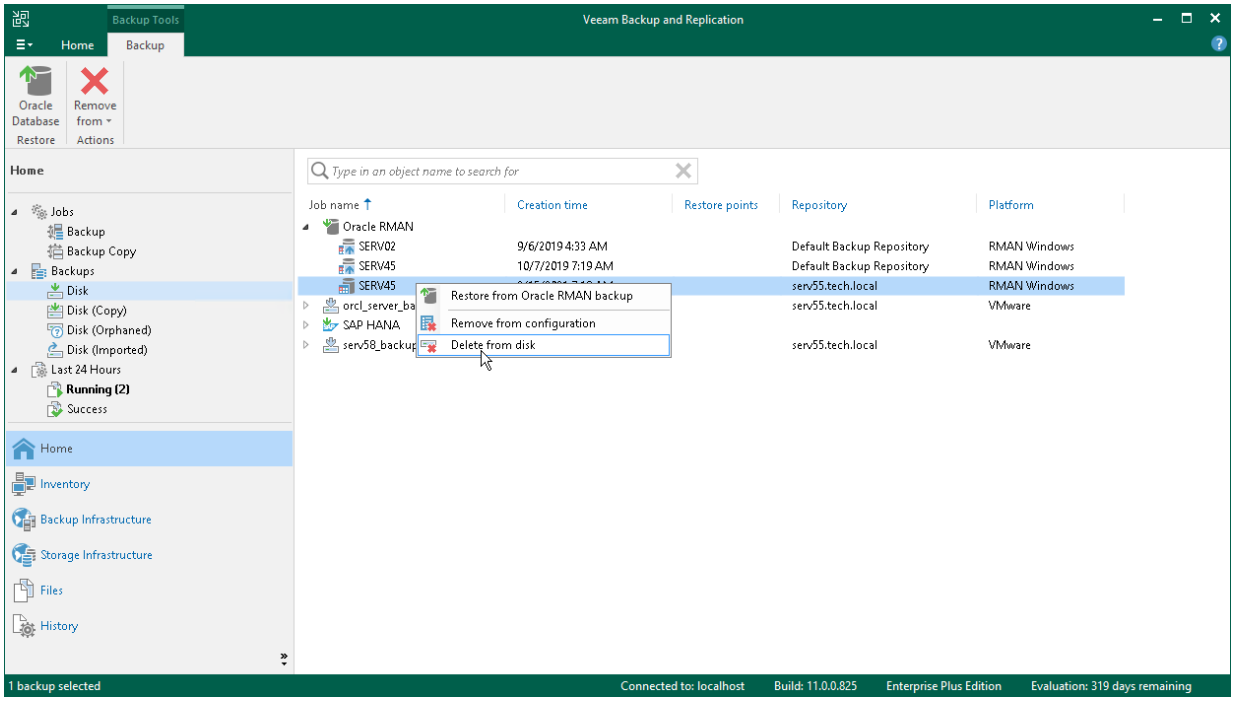

Removing Backups from Configuration

If you want to remove records about backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation.

When you remove a backup from the configuration, backup files (VAB, VBM) remain on the backup repository. You can import the backup later and restore data from it.

To remove a backup from the configuration:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, select the backup and click **Remove from** > **Configuration** on the ribbon. You can also right-click the backup and select **Remove from configuration**.

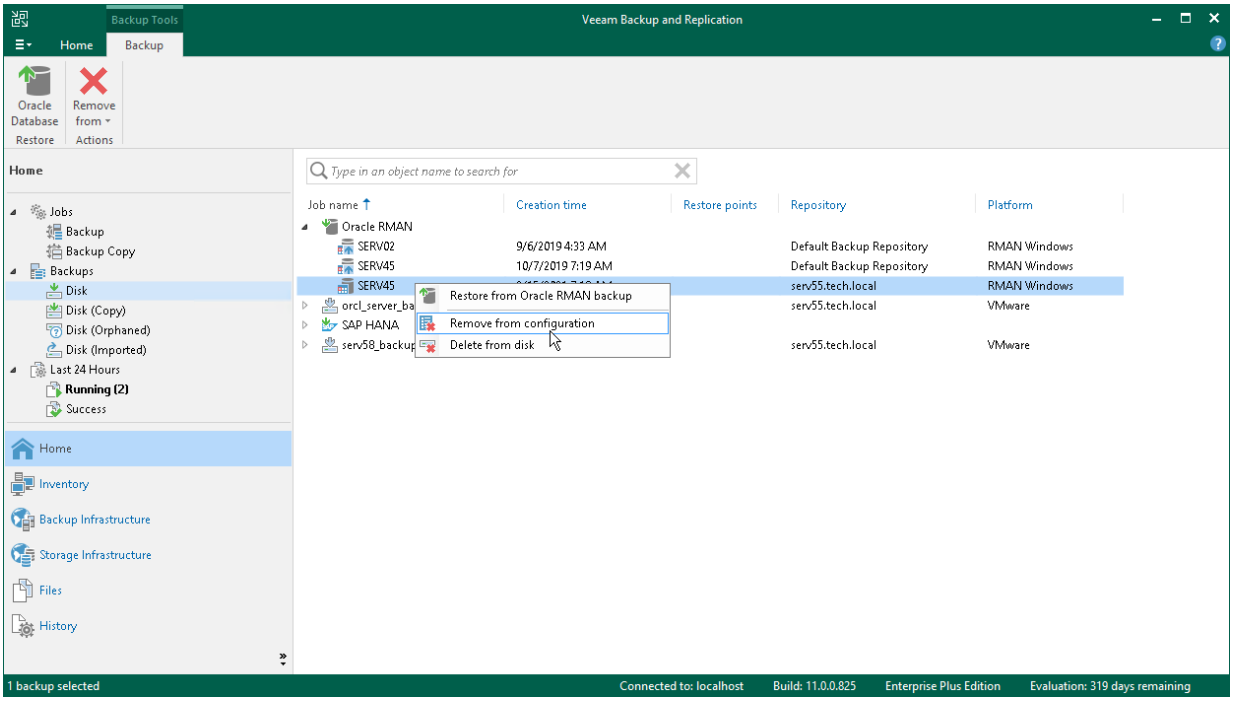

Backup Copy for Oracle RMAN Backups

Having just one backup does not provide the necessary level of safety. The primary backup may get destroyed together with production data, and you will have no backups from which you can restore data.

To build a successful data protection and disaster recovery plan, it is recommended that you follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups.
- 2: You must use at least two different types of media to store the copies of your data, for example, local disk and cloud.
- 1: You must keep at least one backup offsite, for example, in the cloud or in a remote site.

Thus, you must have at least two backups and they must be in different locations. If a disaster takes out your production data and local backup, you can still recover from your offsite backup.

To help you adopt the 3-2-1 rule, Veeam Backup & Replication offers the backup copy functionality that allows you to create several instances of the same backup in different locations, whether onsite or offsite. Backup copies have the same format as those created by backup jobs and you can recover your data from them when you need it.

Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes. Backup copy is a job-driven process. When enabled, the backup copy job for Veeam Plug-in backups runs continuously.

For details on how it works, see the Backup Copy section of the Veeam Backup & Replication User Guide.

In This Section

- Creating Backup Copy Job
- Converting Backup Copy to Backup

Creating Backup Copy Job

To copy backups to a secondary location, you must configure a backup copy job. The backup copy job defines how, where and when to copy backups. One job can be used to process backups of one or more machines.

You can configure a job and start it immediately or save the job to start it later.

Before creating a job, check prerequisites. Then use the **New Backup Copy Job** wizard to configure a backup copy job.

1. Launch Backup Copy Job wizard.
2. Specify a job name and description.
3. Selects backups to process.
4. Define backup copy target.
5. Specify advanced settings.
6. Define backup copy schedule.
7. Finish working with the wizard.

Before You Begin

Before you create a backup copy job, check the prerequisites and limitations:

- If you create a backup copy job for a backup of Veeam Plug-in version 9.5 Update 4, the backups will not be copied. Veeam Plug-in will issue a warning that you must upgrade Veeam Plug-in to version 10. After you upgrade the plug-in and create a new restore point, backup copy job will start to function and copy backup files created by both plug-in versions.
- Backup infrastructure components that will take part in the backup copy process must be added to the backup infrastructure and properly configured. These include source and target backup repositories between which backups must be copied.
- The target backup repository must have enough free space to store copied backups. To receive alerts about low space on the backup repository, configure global notification settings. For more information, see Specifying Other Notification Settings.
- For Veeam Plug-in backup copy jobs, you cannot select a Veeam Cloud Connect repository as a backup copy target.
- If you have upgraded the backup files, make sure that you have upgraded Veeam Plug-in on the source server. If the plug-in is not upgraded to version 11 and you convert backup copy files to backup files, then the next backup job runs will fail.

Step 1. Launch Backup Copy Job Wizard

To create a backup copy job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Click the **Backup Copy** tab and select **Application backup**.

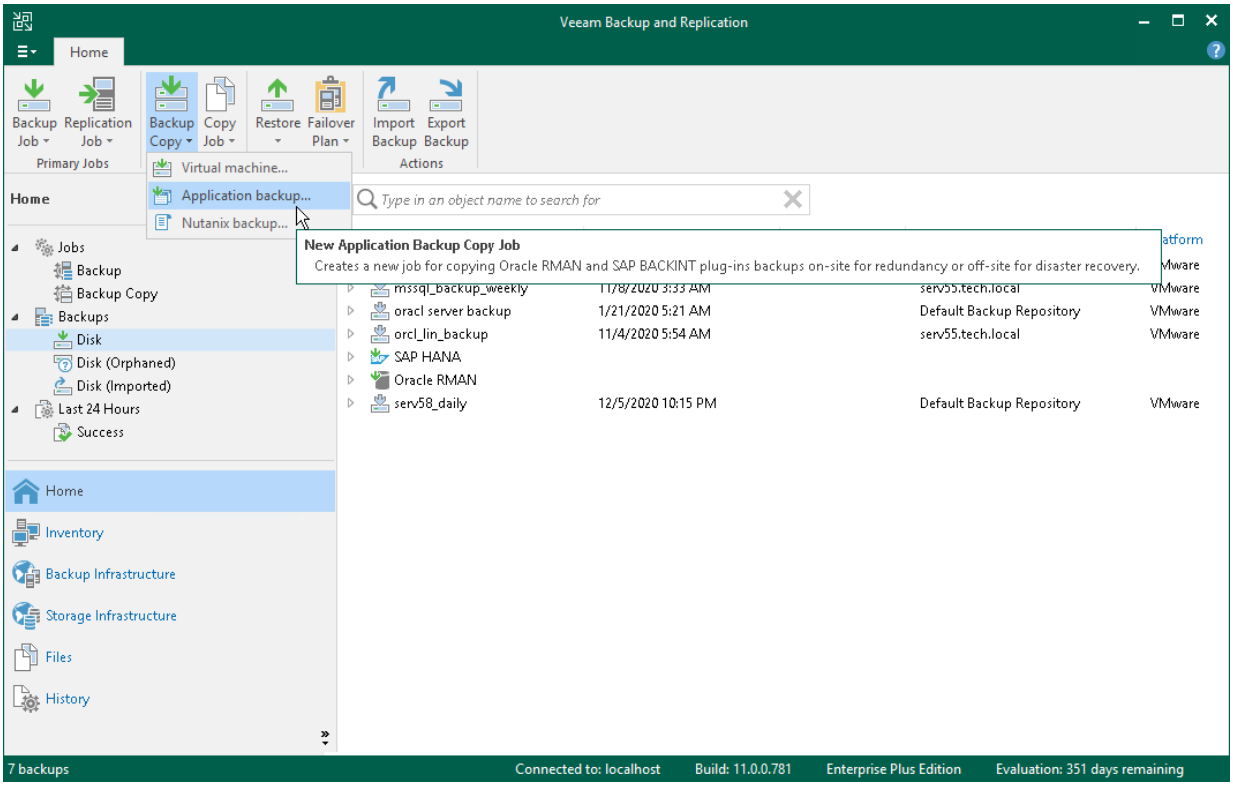

Step 2. Specify Job Name and Description

At the **Job** step of the wizard, specify a name and description for the backup copy job.

1. In the **Name** field, enter a name for the job.
2. In the **Description** field, enter a description for the job. The default description contains information about the user who created the job, date and time when the job was created.

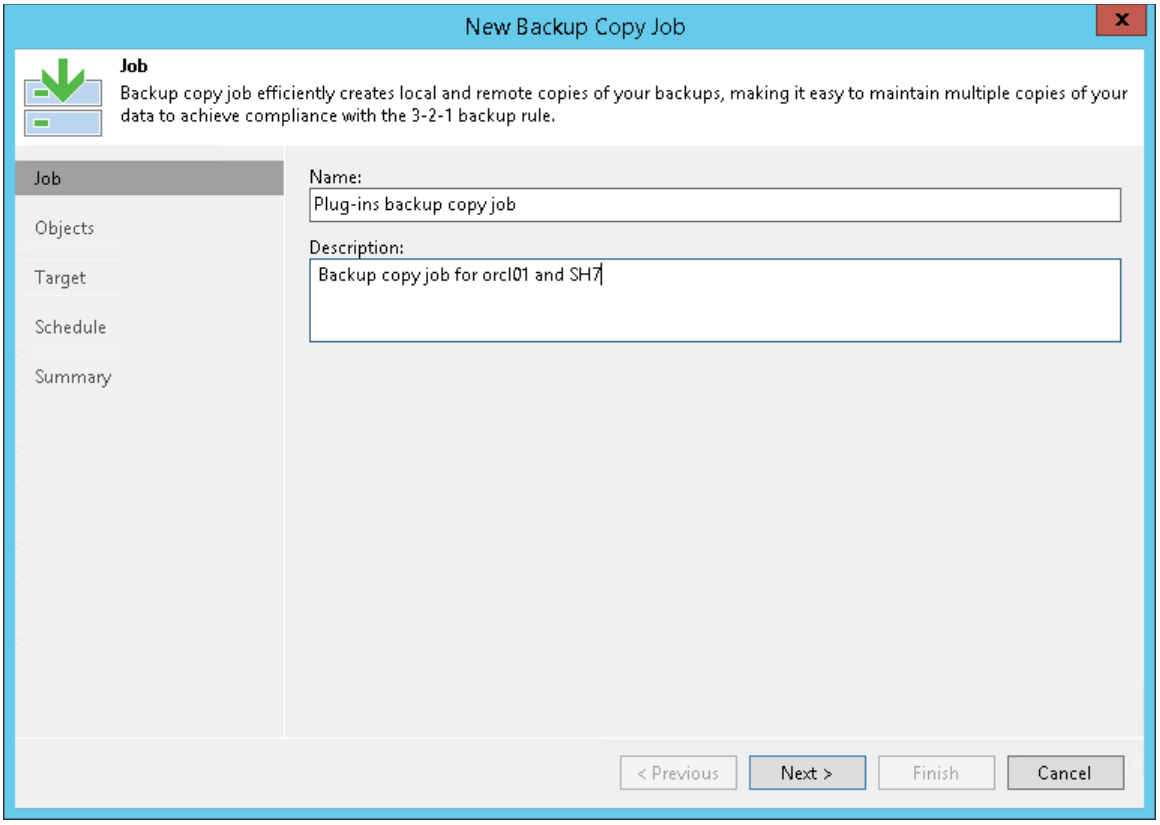

Step 3. Selects Backups to Process

At the **Object** step of the wizard, select machines whose backups you want to copy to the target repository.

1. Click the **Add** button and select from which entity you want to process the machines.
   - **From jobs**: You can select Veeam Plug-in backup jobs. When a backup copy job runs, Veeam Backup & Replication will search for backup files created by selected jobs.
   - **From repositories**: You can select repositories where Veeam Plug-in backups are stored. When a backup copy job runs, Veeam Backup & Replication will search for backup files created by Veeam Plug-in in selected repositories.
2. Use the **Remove** button if you want to remove selected jobs or repositories from processing.
3. If you have added jobs from a repository and want to exclude from processing some of the backup jobs on the selected repository, click **Exclusions** and select the jobs that you want to exclude.

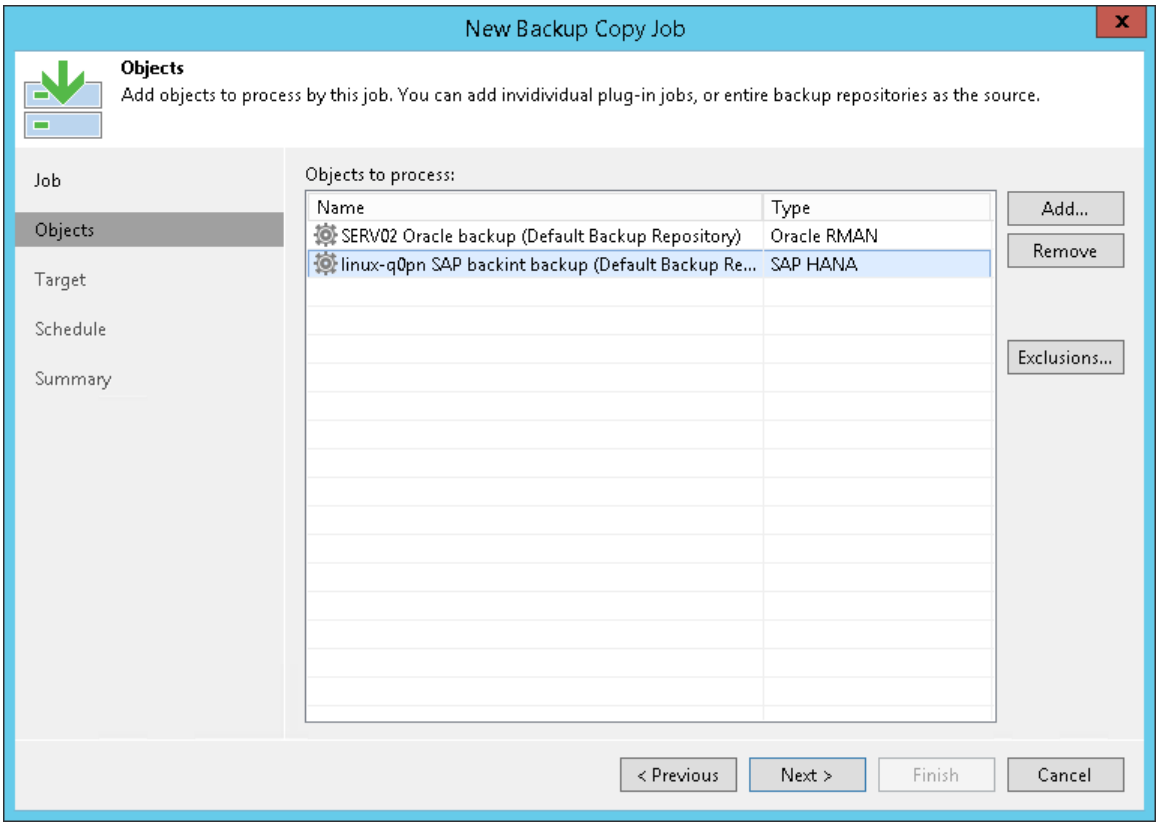

Step 4. Define Backup Copy Target

At the **Target** step of the wizard, configure the target repository settings.

1. From the **Backup repository** list, select a backup repository in the target site where copied backups must be stored. When you select a target backup repository, Veeam Backup & Replication automatically checks how much free space is available on it. Make sure that you have enough free space to store copied backups.

   > **IMPORTANT**
   >
   > For Veeam Plug-in backup copy jobs, you cannot select a Veeam Cloud Connect repository as a backup copy target.

2. If the target repository contains a Veeam Plug-in backup that was excluded from the backup copy job, and if you don't want to transfer duplicate data, you can use the mapping feature.

   After you configure mapping, if some of backup files (.VAB) of the source backup are missing in the target backup copy, these files are uploaded to the target backup copy.

   > **NOTE**
   >
   > Veeam Plug-in backup copy jobs does not use WAN accelerators.

   To map a backup copy job to the backup:

   a. Click the **Map backup** link.

   b. Point the backup copy job to the backup in the target backup repository. Backups in the target backup repository can be easily identified by backup job names. To facilitate search, you can use the search field at the bottom of the window.

   > **IMPORTANT**
   >
   > - Used account must have access to Veeam backup repositories that you plan to use.
   > - Encryption must be disabled on the repository.
   >
   > Otherwise, the repositories will not be listed as available. To learn how to configure access permissions and encryption settings on repositories, see Access and Encryption Settings on Repositories.

3. You can specify the number of days after which the backup copy will be deleted from the repository. Note that the countdown starts from the moment when source backup has been created.

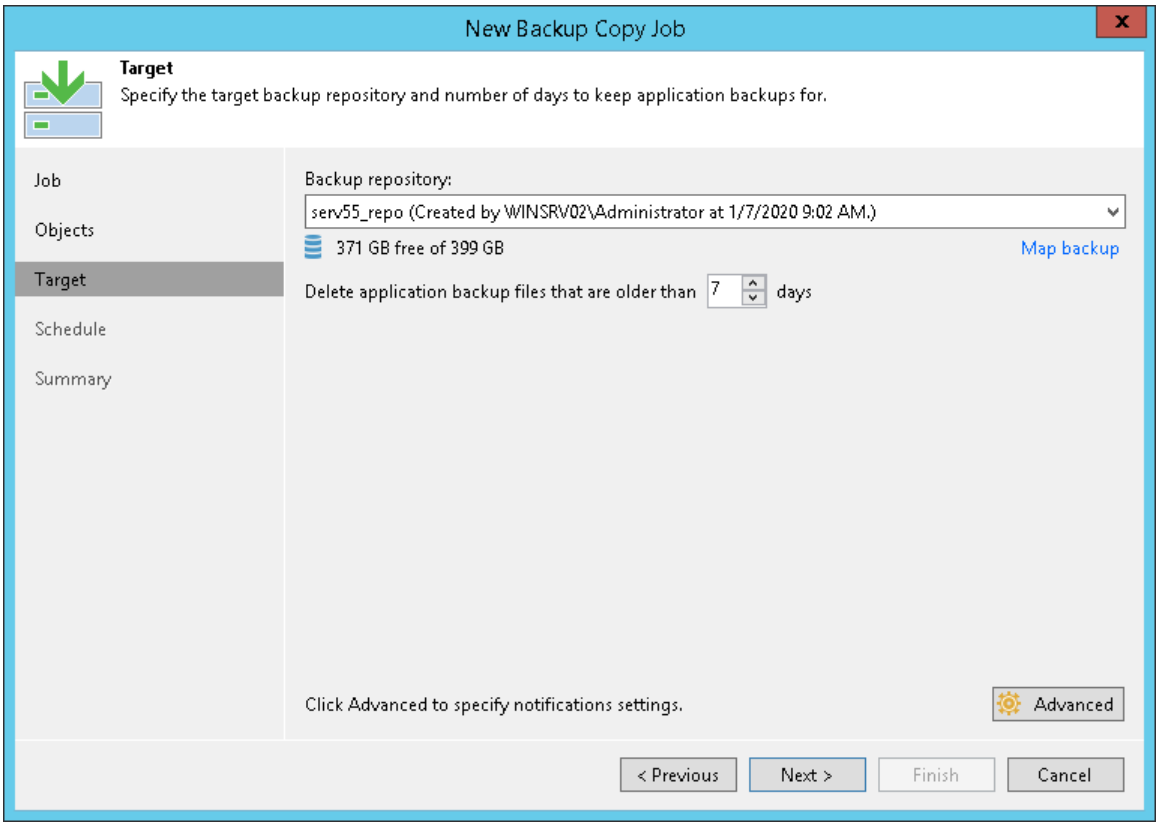

Step 5. Specify Advanced Settings

At the **Target** step of the wizard, click **Advanced** to configure compression, RPO and notifications settings.

- Compression and Deduplication
- RPO
- Notifications

Compression and Deduplication

At the **Storage** tab, define compression and deduplication settings.

By default, Veeam Backup & Replication performs deduplication before storing copied data on the target backup repository. Deduplication provides a smaller size of the resulting backup file but may reduce the job performance.

1. You can disable data deduplication. To do this, clear the **Enable inline data deduplication** check box.
2. From the Compression level list, choose a compression level to be used: **Auto**, **None**, **Dedupe-friendly**, **Optimal**, **High** or **Extreme**. The recommended level of compression for backup copy jobs is **Auto**. In this case, Veeam Backup & Replication uses compression settings of the copied backup files. For more information, see Compression and Deduplication.

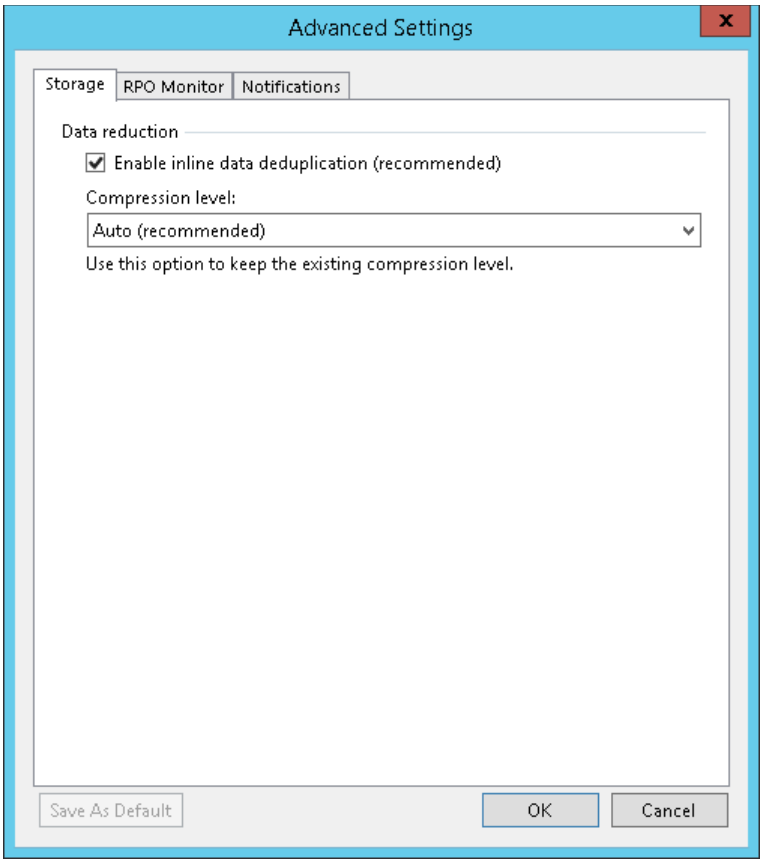

RPO Monitor

At the **RPO Monitor** tab, specify RPO warning settings.

Enable the **Warn me if backup is not copied within** check box and specify the time period in **minutes**, **hours**, or **days**.

If the backup copy is not created within the specified time period, the backup copy job will finish with the *Warning* status. The countdown starts from the moment when the required backup is finished and ready to be copied.

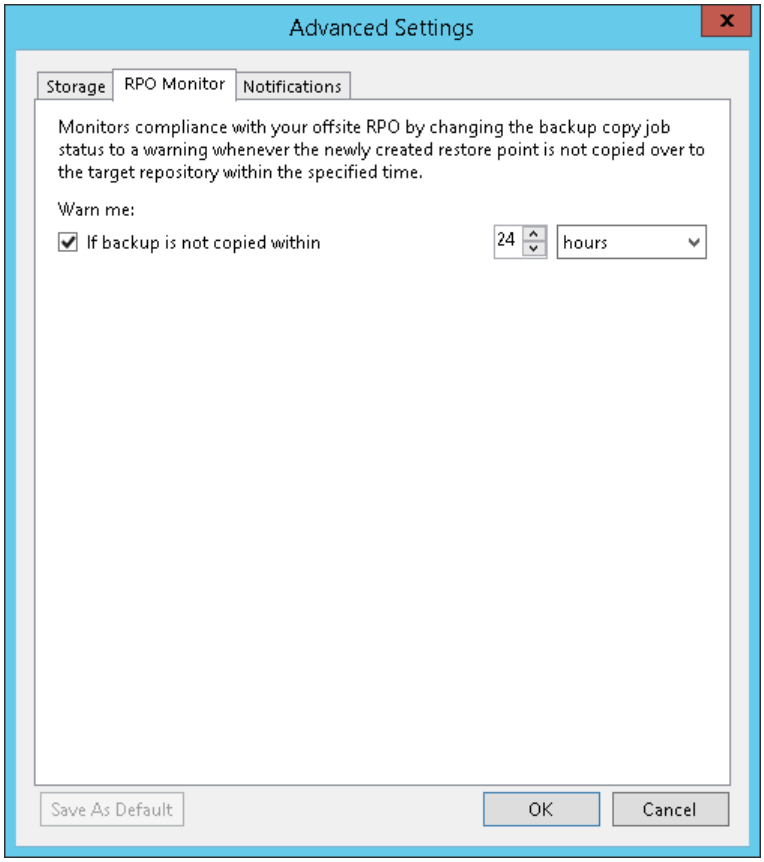

Notifications

At the **Notifications** tab, to specify notification settings for the backup copy job:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully. 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 by email in case of job failure or success. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon.
5. Veeam Backup & Replication sends a consolidated email notification once for the specified backup copy interval. Even if the synchronization process is started several times within the interval, for example, due to job retries, only one email notification will be sent.
6. 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.

7. At the **Send at** field, specify the time when you want to receive notifications. Note that you will receive a notification on a status of the job status once a day.

8. 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%*, *%VmCount%* (number of machines in the job) and *%Issues%* (number of machines in the job that have been processed with the **Warning** or **Failed** status).

     ii. Select the **Notify on success**, **Notify on warning** and **Notify on error** check boxes to receive email notification if data processing within the backup copy interval completes successfully, fails or completes with a warning.

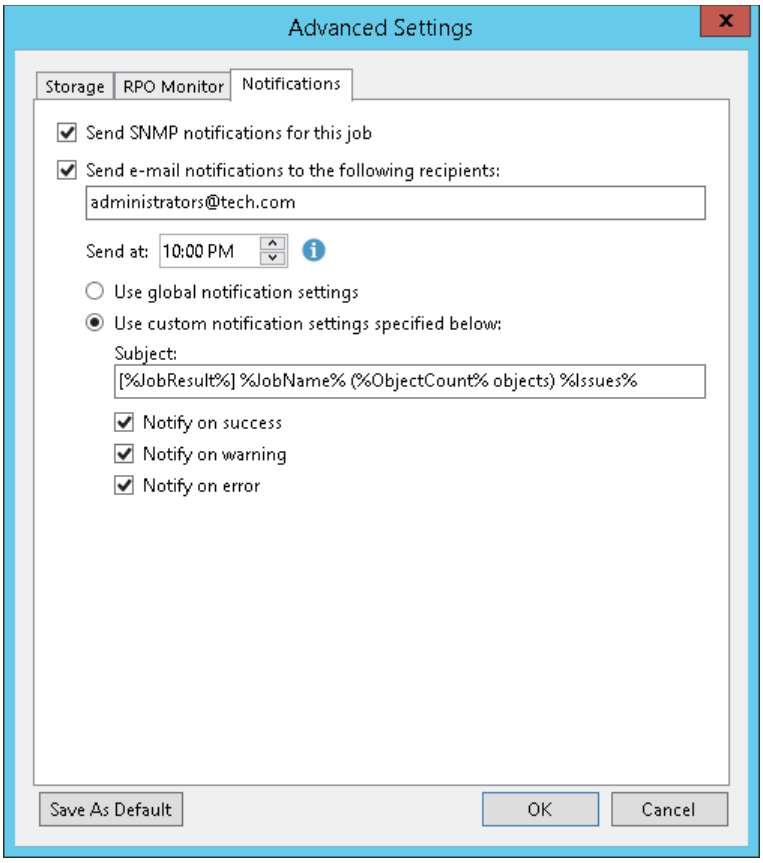

Step 6. Define Backup Copy Schedule

At the **Schedule** step of the wizard, define a time span in which the backup copy job must not transport data between source and target backup repositories. For more information, see Backup Copy Window.

To define a backup window for the backup copy job:

1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enable** and **Disable** options to mark the selected area as allowed or prohibited for the backup copy job.

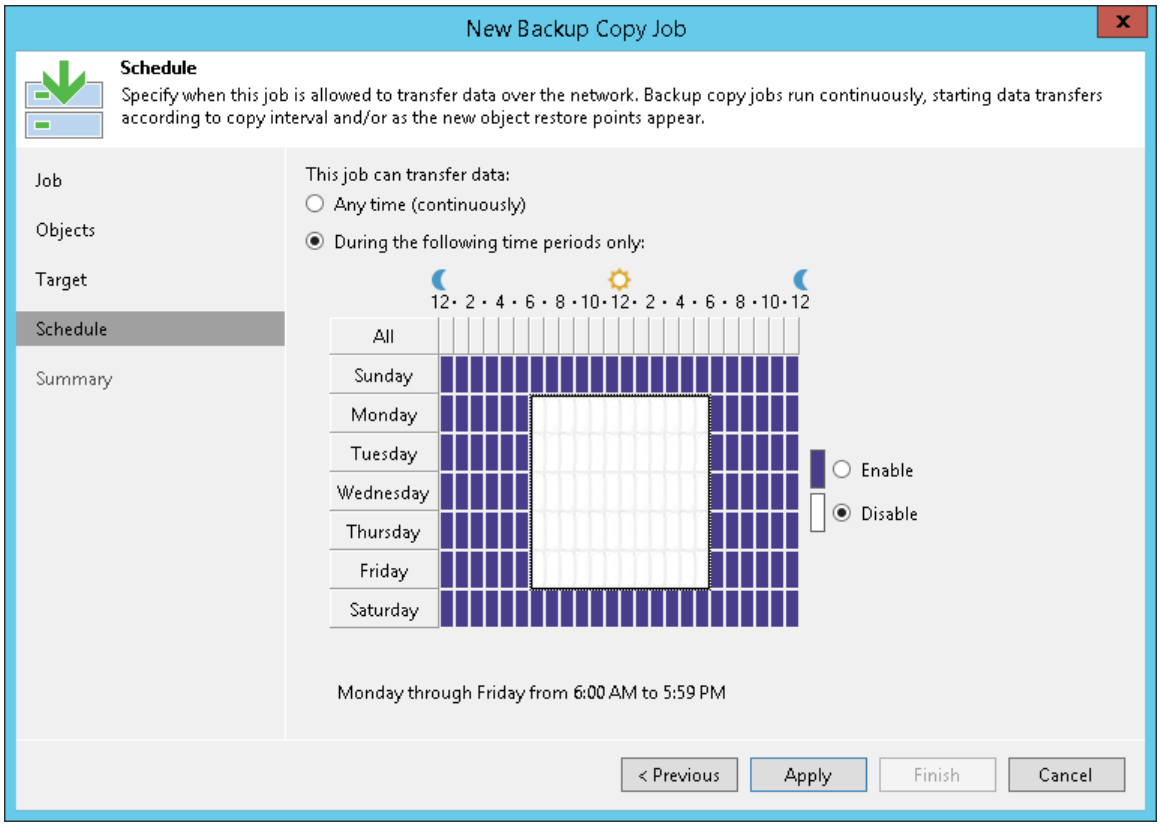

Step 7. Finish working with the wizard

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

1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

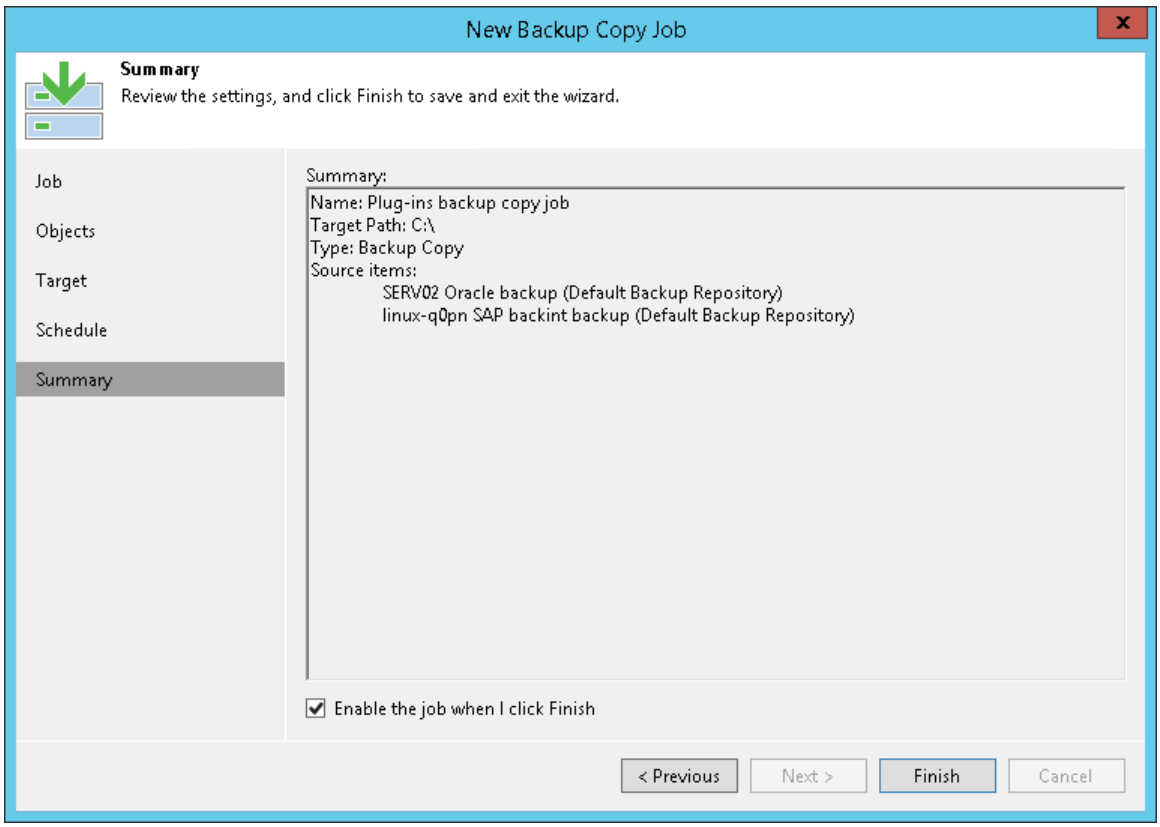

Converting Backup Copy to Backup

If you have imported Veeam Plug-in backup copies from another server, you can convert them into regular backup files. When you convert a backup copy to a backup, Veeam Plug-in creates a backup job with the converted backup. You can use this backup job to continue a backup chain and use the converted backup as a restore point.

You can convert unbind Veeam Plug-in backups into regular Veeam Plug-in backup files in the following cases:

- If you have deleted a backup copy job which created the backup copy.
- If you have excluded a backup job from a backup copy job that used multiple backup jobs as a source.
- If you imported a Veeam Plug-in backup copy from another host.

> **NOTE**
>
> If you want to restore from a backup copy, you don't need to convert the backup copy to backup. For details, see Restore from Backup Copy.

Procedure

To convert a backup copy to a primary backup, use the **--promote-backup-copy-to-primary** parameter as shown below:

1. Run the **OracleRMANConfigTool** with the **--promote-backup-copy-to-primary** parameter and type a backup copy number from the list of available backup copies.

```
OracleRMANConfigTool --promote-backup-copy-to-primary
Backup copies available for promotion to the primary backup target:
1. Backup Copy Job 1\ORCLSERV01 Oracle backup (Default Backup Repository)
Select a backup copy: 1
Changes to be applied to the RMAN configuration
CONFIGURE CHANNEL DEVICE TYPE SBT_TAPE
PARMS 'SBT_LIBRARY=C:\PROGRA~1\Veeam\VEEAMP~1\ORACLE~2.DLL'
FORMAT '94a7ac5a-2cb5-418b-8395-fb362d3aa182/RMAN_%I_%d_%T_%U.vab';
CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE SBT_TAPE TO '%F_RM
AN_AUTOBACKUP.vab';
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE SBT_TAPE TO 1;
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE SBT_TAPE TO 1;
RMAN configuration to be applied:
SQL "alter system set backup_tape_io_slaves=false deferred scope=both";
```

2. Converting a backup copy into regular backup file, requires changes in the RMAN configuration. You can allow the command to change RMAN configuration automatically, or you can change it manually. Select one of the options:

```
Proceed with the action?
1. Promote backup copy destination to the primary backup target and apply
required configuration to RMAN automatically
2. Promote backup copy destination to the primary backup target and export
required RMAN configuration (RMAN will have to be configured manually)
3. Cancel
Enter selection: 1
Promoting backup copy destination
Configuring RMAN
Done
```

Logs and Support

If you have any questions or issues with Veeam Plug-in for SAP on Oracle or Veeam Backup & Replication, you can search for a resolution on Veeam Community Forums or submit a support case in the Veeam Customer Support Portal.

When you submit a support case, we recommend you attach necessary logs related to Veeam Plug-in operations.

For instructions on how to collect logs, see this Veeam KB.

Veeam Plug-in for SAP on Oracle

Veeam Plug-in for SAP on Oracle is a SAP-certified backup tool for SAP applications running on Oracle Database. Veeam Plug-in integrates with SAP BR*Tools and transfers database and log backups to repositories connected to Veeam Backup & Replication.

> **TIP**
>
> If you want to protect the SAP server itself, you can use the image-level backup functionality of Veeam Backup & Replication or Veeam Agent for Linux.

How Veeam Plug-in for SAP on Oracle Works

Veeam Plug-in for SAP on Oracle functions as an agent between SAP BR*Tools and Veeam backup repositories. After you install and configure Veeam Plug-in, you can perform all backup and restore operations with BR*Tools. Veeam Plug-in compresses, deduplicates database backups and transfers them to a backup repository connected to Veeam Backup & Replication.

Veeam Plug-in for SAP on Oracle supports the following tools:

- `brtools`
- `brbackup`
- `brrestore`
- `brarchive`

For details about these tools, see the SAP Database Guide: Oracle.

How Backup Operations Work

After you configure Veeam Plug-in for SAP on Oracle, SAP BR*Tools performs a backup in the following way:

1. When you launch a database backup, BRTOOLS starts the services of Veeam Plug-in.
2. Veeam Plug-in connects to the Veeam Backup & Replication server and creates a backup job object (if it has not been created before). Veeam Backup & Replication administrators can use this backup job object to monitor the backup process, manage backup files and copy the database backup to secondary repositories.
3. BRTOOLS launches the backint BRBACKUP tool that uses the Veeam Plug-in configuration file as an initialization profile.
4. Veeam Plug-in starts Veeam Data Mover on the SAP server and on the backup repository. Veeam Data Movers create communication channels for each backup data stream. Depending on the number of channels specified in Veeam Plug-in settings, there can be 1 or up to 32 parallel channels.

5. Veeam Data Movers transport database backup files to the backup repository.

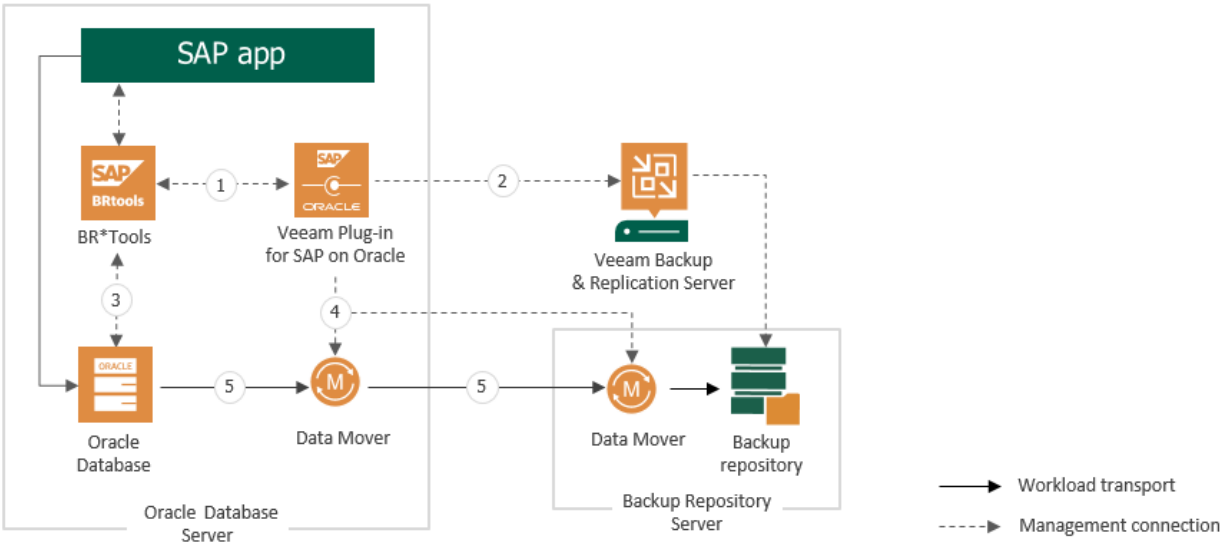

IMPORTANT

Some backup operations, such as backing up of profiles, log files, control files and performing incremental backups of databases can be performed only with RMAN_UTIL. For details, see the RMAN Backup Strategies.

For detail on how Veeam Plug-in for SAP on Oracle functions along with RMAN see, SAP on Oracle Backup Using RMAN_UTIL.

Planning and Preparation

Before you start to use Veeam Plug-in for SAP on Oracle, read the environment planning recommendations and make sure that your environment meets system requirements.

- System Requirements
- Required Permissions
- Used Ports
- Licensing
- Environment Planning
- Veeam Backup Repositories
- Access and Encryption on Repositories

System Requirements

Before you start using Veeam Plug-in for SAP on Oracle, make sure the following requirements are met.

Supported OSes

Veeam Plug-in for SAP on Oracle is supported for the following OSes:

- SUSE Linux Enterprise Server 11, 12, 15 (x86_64)
- Red Hat Enterprise Linux for SAP Applications 6, 7 (x86_64)
- Oracle Linux 6, 7

BR*Tools

Veeam Plug-in for SAP on Oracle supports BR*Tools 7.20 Patch 42 or later.

Oracle DB

Veeam Plug-in for SAP on Oracle supports Oracle Database 11gR2, 12c, 18c, 19c: Standard and Enterprise Edition (Express Edition is not supported).

Veeam Backup & Replication

The version of Veeam Backup & Replication must be the same or newer than version of Veeam Plug-in. Veeam Plug-in for SAP on Oracle version 11 is compatible with Veeam Backup & Replication 11 or later.

Network

Veeam Plug-in should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Plug-in cannot work with the Veeam Backup & Replication server that is located behind the NAT gateway.

Required Permissions

User Rights on SAP on Oracle Server

The account used for installing and updating Veeam Plug-in must have root privileges.

User That Starts BR*Tools Operations

The account used to start BR*Tools backup and restore operations must have permissions described in the Starting BR*Tools section of the SAP Database Guide: Oracle.

Veeam Backup Server User

- The account specified in the Veeam Plug-in configuration settings must be able to authenticate against the Veeam Backup & Replication server. For details, see Configuring Veeam Plug-in for SAP on Oracle.
- The account specified in the Veeam Plug-in configuration settings must be granted access rights on the Veeam backup repository where you want to store backups.

  To learn how to grant permissions on Veeam repositories, see Granting Permissions on Repositories.
- You can work with backups created by Veeam Plug-in only with the account used for creating the backups. If you want to use another account, see required permissions in Configuring Veeam Plug-in for SAP on Oracle.

Used Ports

To enable proper work of Veeam Plug-ins, make sure that the following ports are open.

SAP on Oracle Server

The following table describes network ports that must be opened to ensure proper communication of the SAP on Oracle server and backup infrastructure components.

From	To	Protocol	Port	Notes
SAP on Oracle server where Veeam Plug-in is installed	Veeam Backup & Replication server	TCP	10006	Default port used for communication with the Veeam Backup & Replication server. Note that data between Veeam Plug-ins and backup repositories is transferred directly, bypassing the Veeam Backup & Replication server.
	Backup repository server or gateway server*	TCP	2500 to 3300**	Default range of ports used as data transmission channels. For every TCP connection that a backup process uses, one port from this range is assigned.

* For NFS share, SMB share repositories, and Dell EMC Data Domain, HPE StoreOnce deduplication storage appliances, Veeam Backup & Replication uses an auxiliary backup infrastructure component — gateway server. For details, see the Gateway Server section of the Veeam Backup & Replication User Guide.

** This range of ports applies to newly added backup infrastructure components. If you upgrade to Veeam Backup & Replication 10.0 from earlier versions of the product, the range of ports from 2500 to 5000 applies to the already added components.

Backup Repositories and Gateway Servers

Depending on the type of backup repositories that you use for Veeam Plug-in backups, the following ports must be open to allow communication between backup infrastructure components.

From	To	Protocol	Port	Notes
Veeam Backup & Replication server	Backup repository server or gateway server*	TCP	2500 to 3300**	Default range of ports used as data transmission channels. For every TCP connection that a backup process uses, one port from this range is assigned.
Direct Attached Storage				

From	To	Protocol	Port	Notes
Veeam Backup & Replication server	Linux server used as a backup repository or gateway server	TCP	22	Port used as a control channel from the Veeam Plug-in server to the target Linux host.
	Microsoft Windows server used as a backup repository or gateway server	TCP UDP	135, 137 to 139, 445	Ports used as a management channel from the Veeam Plug-in server to the Repository/Gateway server. Also, the ports are used to deploy Veeam components.
		TCP	6160, 6162	Default ports used by the Veeam Installer Service and Veeam Data Mover Service
Network Attached Storage				
Gateway server (specified in the SMB share repository settings)	SMB server	TCP	445	Default port used by the SMB transport protocol.
		TCP UDP	135, 137 to 139	SMB/Netbios name resolution for the SMB protocol (needed in some cases). For details, see the Used Ports section of the Veeam Backup & Replication User Guide.
Gateway server (specified in the NFS share repository settings)	NFS server	TCP UDP	111, 2049	Standard NFS ports used as a transmission channel from the gateway server to the target NFS share.
Dell EMC Data Domain				
Veeam Backup & Replication server or **Gateway server**	Dell EMC Data Domain For more information, see Dell EMC Documents.	TCP	111	Port used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned.
		TCP	2049	Main port used by NFS. To change the port, you can use the `'nfs set server-port'` command. Note that the command requires SE mode.

From	To	Protocol	Port	Notes
		TCP	2052	Main port used by NFS MOUNTD. To change the port, you can use the `'nfs set mountd-port'` command. Note that the command requires SE mode.
HPE StoreOnce				
Veeam Backup & Replication server or **Gateway server**	HPE StoreOnce	TCP	9387	Default command port used for communication with HPE StoreOnce.
			9388	Default data port used for communication with HPE StoreOnce.
ExaGrid				
Veeam Backup & Replication server	ExaGrid	TCP	22	Default command port used for communication with ExaGrid.
Quantum DXi				
Veeam Backup & Replication server	Quantum DXi	TCP	22	Default command port used for communication with Quantum DXi.

For detailed list of ports used by Veeam Backup & Replication server and backup repositories, see the Used Ports section of the Veeam Backup & Replication User Guide.

Licensing

To use the Veeam Plug-in functionality, you must have a valid Veeam Backup & Replication license. Licenses are installed and managed on the Veeam Backup & Replication server that is connected to the Veeam Plug-in server. If the license is not valid or out of resources, Veeam Plug-in backup jobs fail.

This guide provides information only on specifics of Veeam licenses for Veeam Plug-ins. For terminology and general information about Veeam Licensing, see Veeam Licensing Policy.

See in this section:

- Licensed Objects
- Supported License Types and Packages
- Obtaining and Managing Licenses

Licensed Objects

An SAP on Oracle server is assumed protected if it has been processed by a Veeam Plug-in backup job in the last 31 days.

If you are using any instance-based (Veeam Universal Licensing) license on your Veeam Backup & Replication, you don't need to install any additional licenses. A protected SAP on Oracle server consumes one instance unit from the license. SAP on Oracle servers processed by backup copy jobs are not regarded as protected VMs, these types of jobs provide an additional protection level for VMs that are already protected with Veeam Plug-in backup jobs.

A machine protected by both Veeam Plug-in and Veeam Backup & Replication will consume a license only once. For example, you have an SAP on Oracle server that you back up using Veeam Plug-in. You also back up this server using image-level backup functionality of Veeam Backup & Replication. In this case, only one license will be consumed.

> **NOTE**
>
> [For Perpetual per-socket licenses] If you are using a legacy perpetual per-socket license, a license is required for each hypervisor CPU socket occupied by protected SAP on Oracle servers.
>
> A socket is consumed from the license only if the hypervisor where protected servers reside is added to the Veeam Backup & Replication infrastructure. If the hypervisor is not added to the Veeam Backup & Replication infrastructure, an instance unit will be consumed from the license. To learn how to add a hypervisor to the Veeam Backup & Replication infrastructure, see the Virtualization Servers and Hosts section of the Veeam Backup & Replication User Guide.

Supported License Types and Packages

You can use Veeam Plug-ins with the following license types and packages. Note that this guide contains information on specifics of Veeam license packages only for Veeam Plug-ins. For the full list of license packages, see Pricing and Packaging.

- **For Veeam Universal Licensing**:

  You can use Veeam Plug-ins with all license packages (*Veeam Backup Essentials, Veeam Backup & Replication, Veeam Availability Suite*).

Note that if you use the *Rental* license type, functionality of Veeam Plug-ins is supported only for the *Enterprise Plus* edition of Veeam Backup & Replication.

- **For Perpetual Socket license**:

  Functionality of Veeam Plug-ins is supported only for the *Enterprise Plus* edition of Veeam Backup & Replication.

Obtaining and Managing Licenses

To learn how to install a license and monitor licensed objects, see the Licensing section in the Veeam Backup & Replication User Guide.

Environment Planning

Integration of SAP on Oracle and Veeam Plug-in requires additional environment planning. When you deploy the plug-in, keep in mind the following requirements and limitations.

Scheduling

You can schedule backup processes using `Cron`.

Also, you can schedule and run existing SAP BR*Tools backup scripts within image-level or file-level backup job of Veeam Backup & Replication or Veeam Agent. For details, see the Pre-Freeze and Post-Thaw Scripts section of the Veeam Backup & Replication.

Veeam Backup & Replication Users and Roles

Veeam Plug-in for SAP on Oracle uses the Windows authentication methods of the Veeam Backup & Replication server to establish a connection to this server and to the backup target.

If this user will be later changed manually, the new user must have at least the *Veeam Backup Operator* and *Veeam Restore Operator* rights within the Veeam Backup & Replication user management. To learn how to assign Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication User Guide.

Parallel Data Streams and Backup Repository Task Slots

Any parallel data stream started by SAP Backint will use one backup repository task slot. It is recommended to carefully plan repository task slots, so that SAP Backint can work with multiple channels in parallel.

The following hardware resources are recommended based on tests on Skylake processors:

- **SAP on Oracle server**: 1 CPU core and 200 MB of RAM per currently used channel. Note that resource consumption on the SAP on Oracle server depends on hardware and Oracle settings.
- **Backup repository server**: 1 CPU core and 1 GB of RAM per 5 currently used channels.

  These resources are recommended only if you use a dedicated backup repository for Veeam Plug-in backups. If you use the same backup repository for Veeam Plug-in backups and VM backups created by Veeam Backup & Replication or Veeam Agents, consider adding the hardware resources based on usual load on your backup repository. For details on hardware requirements for a backup repository, see the System Requirements section of the Veeam Backup & Replication User Guide.

  It is recommended to contact your Veeam system engineer to optimize the channel settings and resource allocation.

  It is recommended to use a separate backup repository for Veeam Plug-in backups.
- **Veeam Backup & Replication server**: during manual metadata operations such as import of backup files, the Veeam Backup & Replication server needs additional 15 GB of RAM per 1 million files located in the same backup job folder.

Backup Files

Veeam Plug-in stores backup files in the following formats:

- A .VAB file stores a compressed copy of an Oracle database. Veeam Plug-in creates VAB files for both full and incremental backups.
- A .VASM file stores metadata that contain information about the backup. A .VASM file is created for each .VAB file. The .VASM files are used by Veeam Backup & Replication to get data about Veeam Plug-in backups.
- A .VACM file stores metadata of a backup job object.

Veeam Plug-in generates a name for the .VAB backup file and stores up to 1000 files in backup (FIBs) in one .VAB file. The FIB file names match the external backup IDs (EBID) generated by Veeam Plug-in during the backup.

Also, for each backup file, Veeam Plug-in creates a metadata file that has the same name as the backup file but a different extension (`.VASM`).

Veeam Backup Repositories

Veeam Plug-ins store backup files in repositories added to the Veeam Backup & Replication infrastructure. In this section, you can find the list of supported backup repositories and limitations for Veeam Plug-in backups.

Supported Backup Repositories

Veeam Plug-in for SAP HANA supports integration with the following types of repositories added to the Veeam Backup & Replication infrastructure:

- Windows Server
- Linux Server
- CIFS (SMB) Share
- Dell EMC Data Domain Boost
- Quantum DXi
- ExaGrid
- HPE StoreOnce. If you plan to use HPE StoreOnce as a backup repository for Veeam Plug-in backups, the total number of stored files (data and metadata) must not exceed 3,000,000 per Catalyst store. If necessary, multiple Catalyst stores may be created on the same StoreOnce system.
- NFS File Share

You can also use scale-out backup repositories that contain repositories supported by Veeam Backup & Replication.

Backup Repositories Limitations

- For Veeam Plug-in backups, the warning which indicates that free space on a storage device has reached a specified threshold is configured in the `veeam_config.xml` file of Veeam Plug-in. The warning settings in the Veeam Backup & Replication console does not affect this setting.

  To configure the warning settings, add the following parameter in the `veeam_config.xml` file.

  ```
  <PluginParameters repositoryFreeSpacePercentWarning="10" />
  ```

- The plug-in configuration wizard will not show repositories where the **Encrypt backups stored in this repository** option is enabled. To learn how to disable the encryption option, see Access and Encryption Settings on Repositories.

  If you want to use the same backup target with the repository-based encryption and Veeam Plug-ins, create a second repository in the subfolder for Veeam Plug-in backups.

- Veeam extract utility cannot extract Veeam Plug-in backup files.

Veeam Scale-Out Backup Repositories

If you want to store Veeam Plug-in backups on scale-out backup repositories, mind the following:

- If a scale-out repository is configured in the **Data locality** policy, each time BR*Tools starts a new channel for transferring backup files, Veeam Plug-in checks the free space in the extents and selects a scale-out backup repository extent that has the largest amount of free space. If there are two extents with one slot on each extent, the backup will be launched in two parallel streams (one on each extent).
- For Veeam Plug-in backups and backup copies, the *Performance* policy of a scale-out repository functions differently:

  a. Veeam Backup & Replication checks if there are extents without warning on free space insufficiency. If all extents have the warning, Veeam Backup & Replication uses an extent with the largest amount of free space that has a free task slot.

  b. If there are extents without the warning, Veeam Backup & Replication checks if there are incremental extents with free task slots. If there are no incremental extents with free task slots, Veeam Backup & Replication uses a full extent with the least amount of used task slots.

  c. If there are incremental extents with free task slots, Veeam Backup & Replication will send backup files to an incremental extent with the least amount of used task slots. If the amount of used tasks is the same, an extent with the largest amount of free space.

  To learn more about file placement policies of scale-out repositories, see Backup File Placement section of the Veeam Backup & Replication guide.

- If you want to add a backup repository as an extent to a scale-out backup repository and Veeam Plug-in backups are present on this backup repository, you must do the following:

  a. In the Veeam Backup & Replication console, select Veeam Plug-in backup files that reside in this backup repository and remove them from configuration. For details, see Removing backups from configuration. Note that this action does not delete the backups from the repository.

  b. In the Veeam Backup & Replication console, delete the Veeam Plug-in backup job. For details, see Deleting Jobs.

  c. Add the repository as an extent to the scale-out repository. For details, see Extending Scale-Out Repositories.

  d. Rescan the scale-out repository. For details, see Rescanning Scale-Out Repositories.

  > **NOTE**
  >
  > Names of backup files and paths to backup files must contain only allowed characters:
  >
  > - Alphanumeric characters: `a-zA-Z0-9`
  > - Special characters: `_-.+=@^`
  > - Names of backup files and paths to backup files must not contain spaces.

  e. On the Veeam Plug-in server, set the scale-out repository as the target for backups using the following command:

  ```
  SapBackintConfigTool --set-repository
  ```

f. Map the imported backups using the following command:

```
SapBackintConfigTool --map-backup
```

Capacity Tier

You can configure Veeam Backup & Replication to transfer Veeam Plug-in backup files to a capacity tier. Both policies (Move policy, Copy policy) are supported for Veeam Plug-in backups with the following limitations:

- For Veeam Plug-in backup files, capacity tier does not verify whether data that is being moved is unique and has not been offloaded earlier. Thus, it is highly recommended to check the pricing plans of your cloud storage provider to avoid additional costs for offloading and downloading backup data.
- Capacity tier does not track dependencies of full and incremental Veeam Plug-in backup files. Thus, mind the following:
  - [For the Move policy] When backup files are transferred to the capacity tier, Veeam Backup & Replication takes into account only the creation time of backup files. Make sure that the operational restore window is not longer than the whole backup chain cycle period. Otherwise, you may encounter the scenario when full backup files are transferred to the capacity tier and their increment backup files still remain in the performance tier.
  - The capacity tier immutability expiration date does not have the additional block generation period. The immutability expiration date is based only on the number of days specified in settings of the object storage backup repository.
- If a scale-out repository is down, you cannot restore from Veeam Plug-in backup files stored on a capacity tier. In this case, you can only import the backup files manually and then perform data recovery operations.

Access and Encryption Settings on Repositories

While you configure Veeam Plug-in, you specify an account that must be used to connect to the Veeam Backup & Replication server. To be able to store backups on a backup repository, the specified account must have access permissions on the target backup repository.

To grant access permissions, do the following:

1. In Veeam Backup & Replication, open the **Backup Infrastructure** view.
2. In the inventory pane, click the **Backup Repositories** node or the **Scale-out Repositories** node.
3. In the working area, select the necessary backup repository, click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**.

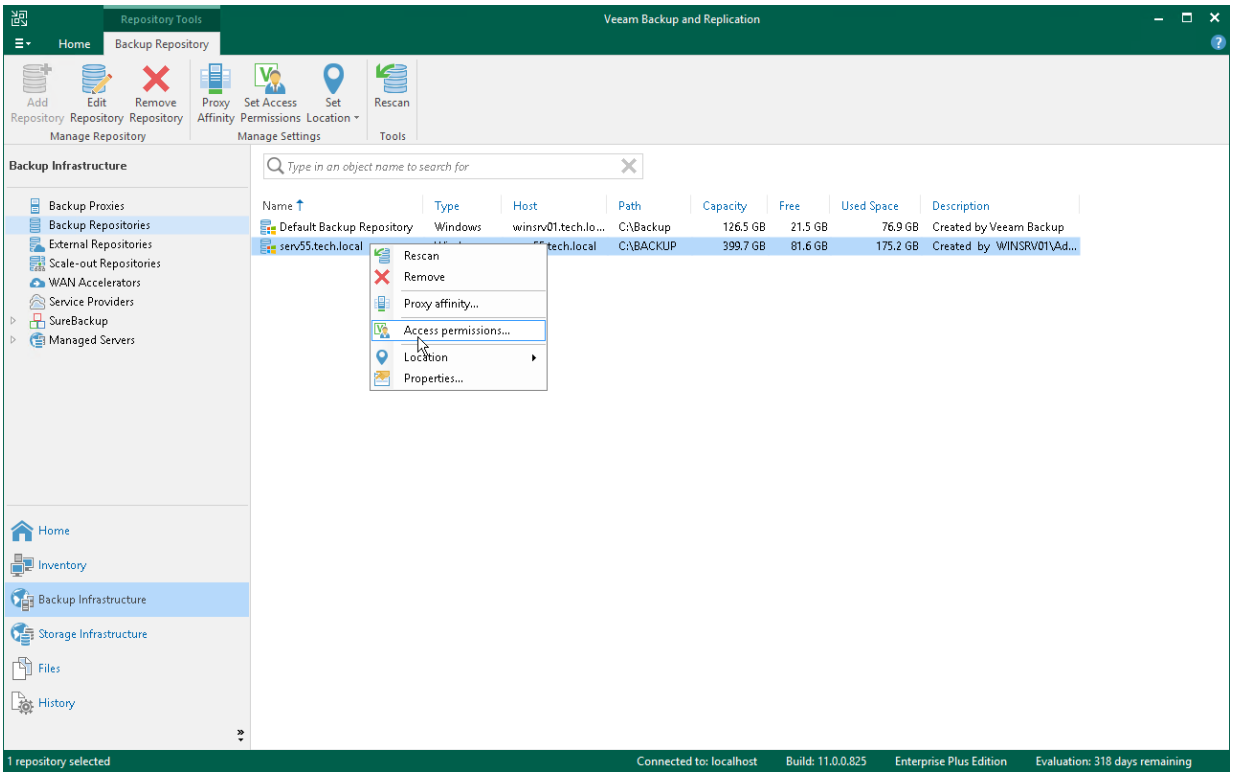

4. In the **Access Permissions** window, specify to whom you want to grant access permissions on this backup repository:
   - *Allow to everyone* – select this option if you want to grant repository access to any user. This option is equal to granting access rights to the *Everyone* group in Microsoft Windows (anonymous users are excluded). For security reasons, the option is not recommended for production environments.
   - *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 repository. Click **Add** to add the necessary users and groups to the list.

5. Veeam Plug-ins cannot send backups or backup copies to a repository where encryption is enabled. Thus, unselect the **Encrypt backups stored in this repository** check box.

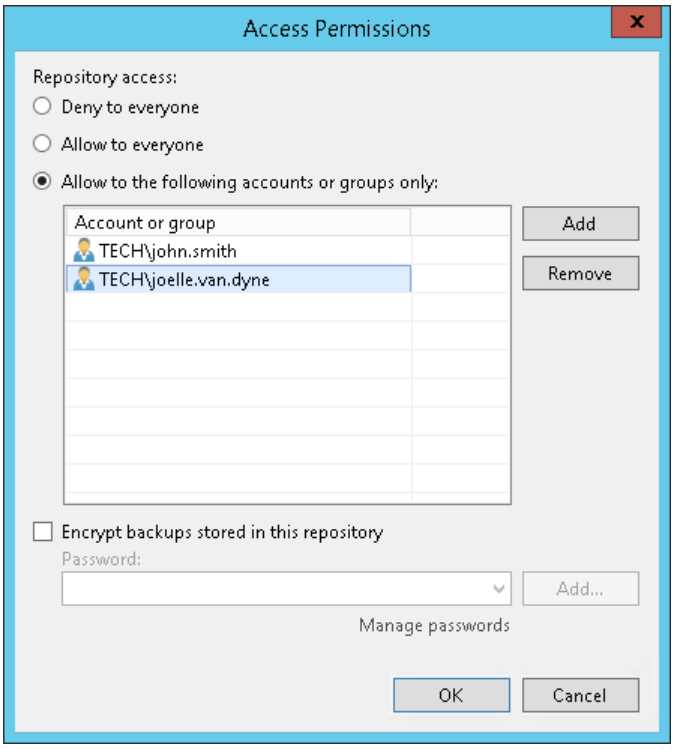

Deployment and Configuration

To deploy Veeam Plug-in, you must install the plug-in on a SAP on Oracle server and configure plug-in integration settings. In this section:

- Installing Veeam Plug-in for SAP on Oracle
- Configuring Veeam Plug-in for SAP on Oracle
- Configuring Parallelism for Redo Logs
- Importing Backups
- Uninstalling Veeam Plug-in for SAP on Oracle

This guide gives instructions on how to deploy Veeam Plug-in assuming that you have already deployed a Veeam Backup & Replication server and configured a backup repository. If you need instructions on how to deploy Veeam Backup & Replication, see the Veeam Backup & Replication User Guide for your platform.

Installing Veeam Plug-in for SAP on Oracle

Veeam Plug-in for SAP on Oracle is an additional component of Veeam Backup & Replication, and the installation package of the plug-in is included in the Veeam Backup & Replication installation ISO file. You must install the plug-in on the Oracle Database server.

You can install Veeam Plug-in using the `.RPM` package or extract the plug-in files from the `.TAR.GZ` archive. Depending on the type of package suitable for your OS, perform steps in one of the following guides:

- Installing Plug-in from .RPM Package
- Unpacking Plug-in from .TAR.GZ Archive

IMPORTANT

Mind the following:

- Veeam Plug-in for SAP on Oracle must be installed on the machine where Oracle Database is deployed.
- The `/opt/veeam` directory must be writable.
- To install the plug-in, use the `sudo` command or a user with root privileges.

Installing Plug-in from .RPM Package

To install Veeam Plug-in using the `.RPM` package, do the following:

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.0.0.XXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk image from the Veeam Backup & Replication: Download page.

2. Open the mounted disk image and go to `/Plugins/SAP on Oracle/x64` directory.
3. Upload the `VeeamPluginforSAPOracle-11.X.X.XXX.x86_64.rpm` file to the Oracle Database server.
4. Run the following command to install the plug-in:

```
rpm -i VeeamPluginforSAPOracle-11.X.X.XXX.x86_64.rpm
```

Unpacking Plug-in from .TAR.GZ Archive

To extract plug-in files from the `.TAR.GZ` archive, perform the following:

1. Mount the Veeam Backup & Replication installation disk (`VeeamBackup&Replication_11.X.X.XXX.iso`).

   If you deploy Veeam backup infrastructure for the first time, you can download the Veeam Backup & Replication installation disk image from the Veeam Backup & Replication: Download page.

2. In the mounted ISO, go to `/Plugins/SAP on Oracle/x64`.
3. Upload the `VeeamPluginforSAPOracle.tar.gz` file to the Oracle Database server.

4. Create the `/opt/veeam` directory.

```
mkdir /opt/veeam
```

5. Unpack the plug-in files from the archive to the `/opt/veeam` directory:

```
tar -xzvf -i VeeamPluginforSAPOracle.tar.gz -C /opt/veeam
```

Configuring Plug-in for SAP on Oracle

When you configure Veeam Plug-in settings, you set up integration settings between the SAP on Oracle server, Veeam Backup & Replication server and backup repositories where backup files will be stored.

Veeam Plug-in uses the **SapOracleBackintConfigTool** wizard to configure the integration settings. The wizard searches for all SAP on Oracle systems deployed on the server and creates the `veeam_initSID.sap` file for each system. Then, this file is used as an initialization profile for `brbackup`, `brrestore` and `brarchive` tools.

Depending on which version of Oracle Database is used, Veeam Plug-in stores the configuration file in the following directories.

- For Oracle 12 and later, the configuration file is stored at `SAPDATA_HOME/sapprof`.
- For Oracle 11, the configuration file is stored at `ORACLE_HOME/dbs`.

IMPORTANT

In some cases when Oracle Database and SAP central instance are running on different hosts, Veeam Plug-in may not be able to detect SAP instances during the configuration. In this case, you must perform the following:

1. Complete all the steps of the Veeam Plug-in configuration wizard described in this section. As a result, Veeam Plug-in will create a new `veeam_initSID.sap` initialization profile and save it in the `/tmp/` directory.
2. Manually copy all lines from the new initialization profile to the default initialization profile. Or you can specify the path to the `veeam_initSID.sap` file in each BRBACKUP, BRRESTORE and BRARCHIVE command.

Veeam Plug-in Configuration

To configure Veeam Plug-in, do the following. Note that the configuration of Veeam Plug-in must be performed by a user with database administrator rights on the Oracle Database server:

1. Log in to the Oracle Database server as a user with database administrator rights and run the following command to launch the Veeam Plug-in configuration tool. You do not need root privileges if you have configured group access as described in the Required Permissions section.

   ```
   SapOracleBackintConfigTool --wizard
   ```

   If you have extracted files form the .TAR.GZ archive, go to the `/opt/veeam/VeeamPluginforSAPOracle` folder and run the following command:

   ```
   ./SapOracleBackintConfigTool --wizard
   ```

2. Specify the DNS name or IP address of the Veeam Backup & Replication server that you want to use.

   ```
   Enter backup server name or IP address: serv02.tech.local
   ```

3. Specify the port which will be used to communicate with the Veeam Backup & Replication server. Default port: *10006*.

```
Enter backup server port number: 10006
```

4. Specify credentials to authenticate against the Veeam Backup & Replication server.

```
Enter username: serv02\administrator
Enter password for serv02\administrator:
```

IMPORTANT

You can work with backup files created by Veeam Plug-in only with the account that was used to create the backups. If you want to use another account, assign the *Veeam Backup Administrator* role or the *Veeam Backup Operator* and *Veeam Restore Operator* roles to the account.

To learn how to assign Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication Guide.

5. Select a backup repository where you want to store backups.

In the wizard dialog, you will see a list of available repositories. Enter the number of the target repository from the list.

```
Available backup repositories:
1. serv10_repo
2. serv07_repo
Enter repository number: 1
```

IMPORTANT

Mind the following:

- The account used to connect to the Veeam Backup & Replication server must have access to the target backup repositories.
- Encryption must be disabled on the target backup repositories.

Otherwise, backup repositories will not be listed as available. To learn how to configure access and encryption settings on repositories, see Access and Encryption Settings on Repositories.

6. Specify the number of parallel data streams for each backup repository.

```
Enter number of data streams (From 1 to 32) to run in parallel: 4
Configuration result:
An auxiliary initialization profile has been successfully created for SAP
system "ODB": /oracle/ODB/sapprof/veeam_initODB.sap
The created profile must be leveraged to perform backup and restore tasks
by BR*Tools.
```

Note that this parallelism setting applies only to backup and restore of Oracle datafiles. If you want to configure parallel channels for backup and restore of redo logs, see Configuring Parallelism for Redo Logs.

Configuration Tool Commands

Apart from running a configuration wizard, you can use the **SapOracleBackintConfigTool** tool to change a specific parameter in the `veeam_config.xml` file or enable/disable Veeam Plug-in features.

See the list of available commands for **SapOracleBackintConfigTool**:

Command	Description
--help	Shows the list of tool parameters.
--show-config	Shows configuration parameters.
--wizard	Starts the wizard to configure the plug-in settings. This wizard edits the `veeam_config.xml` file or creates a new one if the configuration file was removed from the `/opt/veeam/VeeamPluginforSAPOracle/` directory.
--set-credentials <"serv\username"> <password>	Specifies credentials to log in to the Veeam Backup & Replication server.
--set-host <hostname>	Specifies the IP address or hostname of the Veeam Backup & Replication server.
--set-port <port_number>	Specifies a port number that will be used to communicate with the Veeam Backup & Replication server.
--set-repository	Launches a wizard to select a backup repository. A backup repository is selected from repositories which are available in the connected Veeam Backup & Replication instance.
--set-restore-server	[for System Copy] Specifies the backup that will be copied.
--set-parallelism <number_of_channels>	Define the number of parallel channels that must be used to transfer Oracle datafiles during the backup and restore operations. You can set up to 32 channels. Note that the parallelism for redo logs is configured separately. For details, see Configuring Parallelism for Redo Logs.
--map-backup	Maps the imported backup.
--set-force-delete	Deletes backup files after specified days.

Command	Description
--configure-restore-from-copy	Enables restore from backup copy. Note that if you enable restore from backup copy, you cannot back up databases with Veeam Plug-in. To revert changes, you must disable restore from backup copy. Note that when you launch the command, the wizard will ask you to reconfigure the catalog backup from backint to disk.
--promote-backup-copy-to-primary	Maps the imported backup copy to a regular Veeam Plug-in backup chain.

Example:

The following example shows how to specify credentials that will be used to log in to the Veeam Backup & Replication server.

```
SapOracleBackintConfigTool --set-credentials "serv02\Administrator" "password"
```

Configuring Parallelism for Redo Logs

To configure backup or restore of redo logs through multiple channels, you can change the parallelism parameter in the `veeam_config.xml` file.

1. In the machine where Veeam Plug-in is installed, go to the `/opt/veeam/VeeamPluginforSAPOracle/` directory and open the `veeam_config.xml` file with a text editor.
2. Set the necessary values for the parallelism values:

   ```
   <PluginParameters Parallelism="4" LogsParallelism="4" />
   ```

   where the `LogsParallelism` parameter value defines the number of parallel channels for backup and restore of redo logs.

   Note that the first `Parallelism` value configures the parallelism for backup and restore of Oracle datafiles. This setting is configured in the Veeam Plug-in configuration wizard.

Importing Backup Files

If the Veeam Backup & Replication server has failed and you have restored it in a new location, you can copy the backup files to a new repository and re-map the Veeam Plug-in backup files.

Limitations and Prerequisites

Mind the following limitations:

- If backup files are not imported according to instructions given in this section, Veeam Plug-in backup and restore operations may fail.
- The repository from which you plan to import backups must be added to the Veeam Backup & Replication infrastructure. Otherwise you will not be able to access backup files.
- If you are importing backup files from a scale-out backup repository, the names of backup files and paths to backup files must contain only allowed characters:
  - Alphanumeric characters: `a-zA-Z0-9`
  - Special characters: `_-.+=@^`
  - Names of backup files and paths to backup files must not contain spaces.

How to Import Veeam Plug-in Backup Files

To import Veeam Plug-in backup files, do the following:

1. Move the folder with the backup file to the required backup repository or create a new backup repository with this folder as a subfolder.

   > **TIP**
   >
   > Starting from version 11, each Veeam Plug-in backup file (.vab) has its own metadata file (.vasm). Make sure you import backup files and all related metadata files. Also, you must import the backup job metadata file (.vacm) which is stored in the same folder.

2. Use the Veeam Backup & Replication console to log in to Veeam Backup & Replication.
3. Open the **Backup Infrastructure** view.
4. In the inventory pane of the **Backup Infrastructure** view, select the **Backup Repositories** node.

5. In the working area, select the required backup repository and click **Rescan** on the ribbon. Alternatively, you can right-click the backup repository and select **Rescan**.

   During the rescan operation, Veeam Backup & Replication gathers information about backups that are currently available in the backup repository and updates the list of backups in the configuration database. After the rescan operation, backups that were not in this configuration database will be shown on the **Home** view in the **Backups > Disk (Imported)** node.

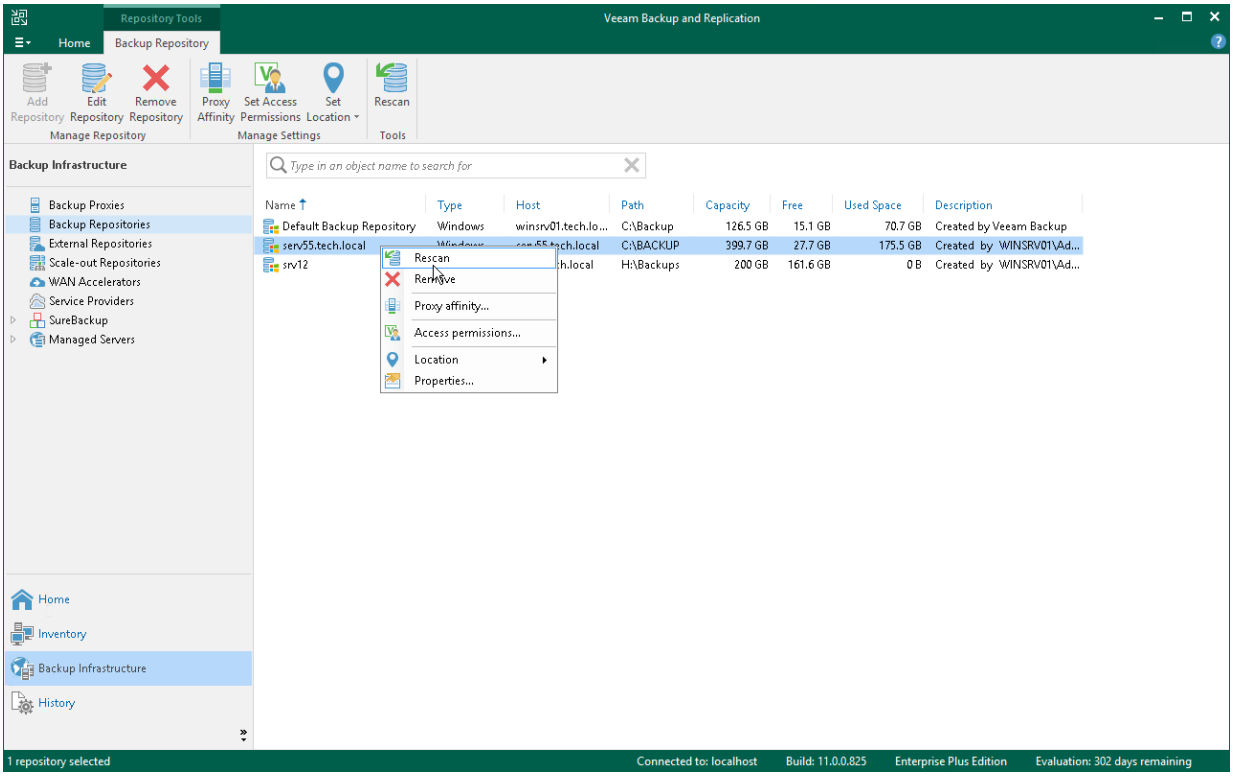

6. On the SAP on Oracle server, change the target backup repository in the Veeam Plug-in settings:

```
sudo SapOracleBackintConfigTool --set-repository
Available backup repositories:
1. serv55.tech.local
2. serv07_repo
Enter repository number: 1
Configuration result:
SID SH2 has been configured
```

7. Run the `--map-backup` command:

```
sudo SapOracleBackintConfigTool --map-backup
```

Uninstalling Plug-in for SAP on Oracle

To uninstall Veeam Plug-in for SAP on Oracle, go to the directory with the Veeam Plug-in installation package and run the following command. Note that the operation requires *root* privileges.

```
rpm -e VeeamPluginforSAPOracle
```

Database Protection

After you configure Veeam Plug-in for SAP on Oracle, you can back up databases using SAP BR*Tools. Veeam Plug-in will automatically transfer backup files to a backup repository and store them in the Veeam proprietary format. The backup process itself is performed by SAP BR*Tools.

Keep in mind that examples in this section are provided only for demonstrating purposes. For details on full backup functionality of SAP BR*Tools, see the BR*Tools for Oracle DBA Guide.

In this section

- Limitations and Considerations
- Backing Up Databases Using Backint
- Backing Up Redo Logs
- Backing Up Databases Using RMAN_UTIL
- Backup Job in Veeam Backup & Replication Console

NOTE

This guide provides examples for SAP BR*Tools commands. Apart from BR*Tools scripts, you can perform backup operations using BRTOOLS interactive wizard and BR*Tools Studio. For details, see the Backing Up the Databases with BR*Tools section of the SAP Database Guide: Oracle.

Limitations and Considerations

Before you start using Veeam Plug-in for SAP on Oracle, mind the following:

- You can make incremental backups only with the `rman_util` parameter. For details, see the RMAN Backup Strategies section of the SAP Database Guide: Oracle. For instructions on how to perform an incremental backup, see Incremental Backup.
- ASM is supported only with the `rman_util` parameter. If you want to back up ASM instances, see SAP on Oracle Backup Using RMAN_UTIL.
- You cannot back up Oracle RAC databases using the BRBACKUP tool.
- Volume backup (`-d util_vol`, `util_vol_online`) is not supported.
- Backup of directories is not supported.

SAP on Oracle Backup Using Backint

Veeam Plug-in for SAP on Oracle integrates with the BRBACKUP tool and transfers backup files to backup repositories connected to Veeam Backup & Replication. The backup process itself is performed by the BRBACKUP tool. To perform the backup you can use BR*TOOLS or BR*TOOLS Studio.

For details on how the backup is performed, see How Veeam Plug-in for SAP on Oracle Works.

To back up Oracle databases, you can use the interactive wizard of BRTOOLS or you can directly run the backup command using BRBACKUP. When you back up the database using the BRBACKUP tool, you must specify the path to the initialization profile file (`$Oracle_HOME/dbs/veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.

> **NOTE**
>
> You can make incremental backups only with the `rman_util` parameter. For details, see the RMAN Backup Strategies section of the SAP Database Guide: Oracle. For instructions on how to perform an incremental backup, see Incremental Backup.

Examples

The following examples are only for demonstration purposes. To see the description of all BRBACKUP parameters, see the Backing Up the Database with BR*Tools for Oracle DBA Guide.

- Full Backup
- Redo Logs Backup

Full Backup

If you want to create a full backup Oracle databases, you can use the BRBACKUP tool. When Veeam Plug-in for SAP on Oracle is configured, the plug-in transfers database backup files to a backup repository connected to Veeam Backup & Replication.

Example 1. Performing Full Database Backup in Offline Mode

```
brbackup -p $Oracle_HOME/dbs/veeam_initSID.sap -d util_file -t offline_force -m
all -u <user>/<password>
```

Run the `brbackup` command with the following parameters:

1. Specify the path to the initialization profile file (`$Oracle_HOME/dbs/veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.
2. Specify `util_file` as the argument for the `-d` (`-device`) parameter. This option defines that a file-by-file backup will be performed using Veeam Plug-in.
3. Specify `offline_force` as the argument for the `-t` (`-type`) parameter. With this option, BRBACKUP shuts down the database and performs an offline backup.
4. Specify the argument for the `-m` (`-mode`) parameter. With the `all` argument, BRBACKUP performs backup of files in all tablespaces, but not the control files and online redo log files. For the full list of arguments for the `-mode` parameter, see SAP Documentation.

5. Specify credentials that will be used to connect to the database as the argument for the `-u (-user)` parameter. For details, see SAP Documentation.

Example 2. Performing Full Database Backup in Online Mode

```
brbackup -p $Oracle_HOME/dbs/veeam_initSID.sap -d util_file_online -t online -m
all -u <user>/<password>
```

Run the `brbackup` command with the following parameters:

1. Specify the path to the initialization profile file (`$Oracle_HOME/dbs/veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.
2. Specify `util_file_online` as the argument for the `-d` (`-device`) parameter.
3. Specify the argument for the `-m` (`-mode`) parameter. With the `all` argument, BRBACKUP performs backup of files in all tablespaces, but not the control files and online redo log files. For the full list of arguments for the `-mode` parameter, see SAP Documentation.
4. Specify credentials that will be used to connect to the database as the argument for the `-u (-user)` parameter. For details, see SAP Documentation.

IMPORTANT

When you use BRBACKUP, you must specify the full directory path to the Veeam Plug-in initialization profile file (`-p $Oracle_HOME/dbs/veeam_initSID.sap`). If the profile file is in the default directory, you can specify only the file name.

Redo Logs Backup

If you want to back up redo log files of Oracle databases, you can use the BRARCHIVE tool. When Veeam Plug-in for SAP on Oracle is configured, the plug-in transfers the redo logs to a backup repository connected to Veeam Backup & Replication.

NOTE

For redo log backup operations, it is recommended to set 4 or less parallel channels. For details on configuring parallel channels, see Configuring Parallelism.

Example: Performing Backup of Archived Redo Logs

To back up redo log files, run the following command.

```
brarchive -p $Oracle_HOME/dbs/veeam_initSID.sap -save -d util_file -u <user>/<p
assword>
```

1. Specify the path to the initialization profile file (`$Oracle_HOME/dbs/veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.
2. Specify the archive function. The `-save` function used in this example archives offline redo log files to a repository.

3. Specify `util_file` as the argument for the `-d` (`-device`) parameter. This option defines that a file-by-file backup will be performed using Veeam Plug-in.
4. Specify credentials that will be used to connect to the database as the argument for the `-u (-user)` parameter. For details, see SAP Documentation.

SAP on Oracle Backup Using RMAN_UTIL

The `rman_util` parameter allows to back up Oracle databases using Oracle RMAN in combination with Veeam Plug-in for Oracle RMAN. BACKINT provides an interface for Veeam Plug-in for Oracle RMAN and is used to back up profiles, log files, control files and to perform incremental backups of databases.

For full description of the `rman_util` parameter, see RMAN Backup with an External Backup Library.

In the Veeam Backup & Replication console, the `rman_util` backup operation will create two backup jobs: Veeam Plug-in for Oracle RMAN backup job for database file backups and another Veeam Plug-in for SAP on Oracle backup job that backups up BR*Tools control data files. Note that if you want to create a backup copy job for SAP on Oracle database, make sure that you have added both jobs to the backup copy job.

Prerequisites

> **IMPORTANT**
>
> - Before you back up the Oracle database with the rman_util parameter, you must install and configure Veeam Plug-in for Oracle RMAN on the SAP on Oracle server.
> - See Limitations and Considerations.

When you perform the backup using the RMAN_UTIL, RMAN_STAGE or RMAN_DISK parameter, by default, BR*Tools creates one backup set for each log or datafile. This means that every backup piece will contain only one file. This results in a large amount of backup files and significantly slows down backup and restore processes. To avoid this problem, do the following:

1. In the SAP on Oracle server, open the `/oracle/ODB/sapprof/veeam_initSID.sap` file using a text editor.
2. Change the default values for the following parameters:

   ```
   rman_filesperset = 10
   rman_filesperset_arch = 100
   ```

   For example: set the `rman_filesperset` value to 10 for datafiles and the `rman_filesperset_arch` value to 100 for logs.

Also, you must add the SBT_LIBRARY directory to the `rman_parms` setting in the `veeam_initSID.sap` file:

1. In the SAP on Oracle server, open the `/oracle/ODB/sapprof/veeam_initSID.sap` file using a text editor.
2. Add the following line in the `veeam_initSID.sap` file:

   ```
   rman_parms = 'SBT_LIBRARY=/opt/veeam/VeeamPluginforOracleRMAN/libOracleRMA
   NPlugin.so'
   ```

How It Works

When you launch the BRBACKUP or BRARCHIVE tool with the RMAN_UTIL parameter, the following happens:

1. SAP BR*Tools launches the RMAN backup script.
2. Oracle RMAN launches Veeam Plug-in for Oracle RMAN services.

3. Oracle RMAN starts the backup process:
   a. Veeam Plug-in compresses, deduplicates database backup files or redo logs and sends them to the target backup repository through one or multiple channels.
   b. Veeam Plug-in for Oracle RMAN connects to Veeam Backup & Replication and creates a backup job object that shows the job progress and logs.
4. BR*Tools launches the Veeam Plug-in for SAP on Oracle services.
5. BR*Tools start the control data files backup:
   a. Control file, BR*Tools logs are compressed and sent to a backup repository.
   b. Veeam Plug-in for SAP on Oracle connects to Veeam Backup & Replication and creates a backup job object that shows the job progress and logs.

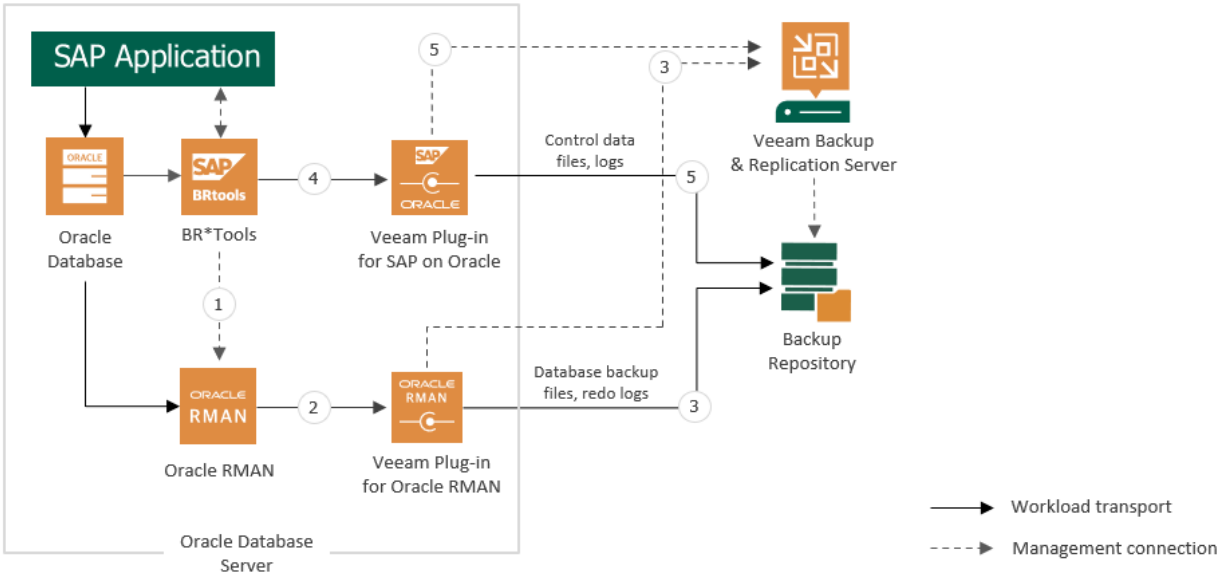

Full Backup

To backup an SAP on Oracle database using RMAN, you must use the `brbackup` tool with the `rman_util` parameter and with the defined directory for SBT library.

```
brbackup -p $Oracle_HOME/dbs/veeam_initSID.sap -t online -d rman_util -m full -
u <user>/<password>
```

1. Specify the path to the initialization profile file (`veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.
2. Specify `rman_util` as the argument for the `-d` (`-device`) parameter. This option defines that the backup will be performed using Oracle RMAN.
3. Specify `online` as the argument for the `-t` (`-type`) parameter. With this option, BRBACKUP performs backup of the database in the online state.
4. Specify `full` as the argument for the `-m` (`-mode`) parameter. With this option, BRBACKUP performs backup of files in all tablespaces, control files and redo log files.
5. Specify credentials that will be used to connect to the database as the argument for the `-u (-user)` parameter. For details, see SAP Documentation.

For the fill list of BRBACKUP parameters, see the Command Options for BRBACKUP section of the SAP Database Guide: Oracle.

Incremental Backup

You can perform an incremental backup by using the BRBACKUP command with the `rman_util` parameter. An incremental backup contains the changed data from the last full backup. Incremental backups use less media and resources than full backups.

Example: Performing Incremental Backup in Online Mode

```
brbackup -p $Oracle_HOME/dbs/veeam_initSID.sap -t online -d rman_util -m incr -
u <user>/<password>
```

Run the `brbackup` command with the following parameters:

1. Specify the path to the initialization profile file (`$Oracle_HOME/dbs/veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.
2. Specify `online_force` as the argument for the `-t` (`-type`) parameter. With this option, BRBACKUP performs backup of the database in the online state.
3. Specify `util_file` as the argument for the `-d` (`-device`) parameter.
4. Set `incr` as the argument for the `-m` (`-mode`) parameter.
5. Specify credentials that will be used to connect to the database as the argument for the `-u (-user)` parameter. For details, see SAP Documentation.

To see all brbackup command options, see the Command Options for BRBACKUP section of the SAP Database Guide: Oracle.

Backup Job in Veeam Backup & Replication Console

After you start a backup process with BRBACKUP, Veeam Backup & Replication creates a backup job object. You can use this job to view the statistics on the backup process, generate backup job reports or you can also disable the backup job.

You cannot launch or edit SAP on Oracle backup job objects in the Veeam Backup & Replication console. You can manage backup operations only on the SAP on Oracle side using BR*Tools.

> **NOTE**
>
> Due to specifics of the SAP on Oracle backup process, the progress bar of a running SAP on Oracle backup job is not available.

To view details of a backup job process, do the following.

1. Open the Veeam Backup & Replication console.
2. In the **Home** view, expand the **Jobs** node and click **Backup**.
3. In the list of jobs, select the BR*Tools backup job to see details of the current backup process or the last backup job session.

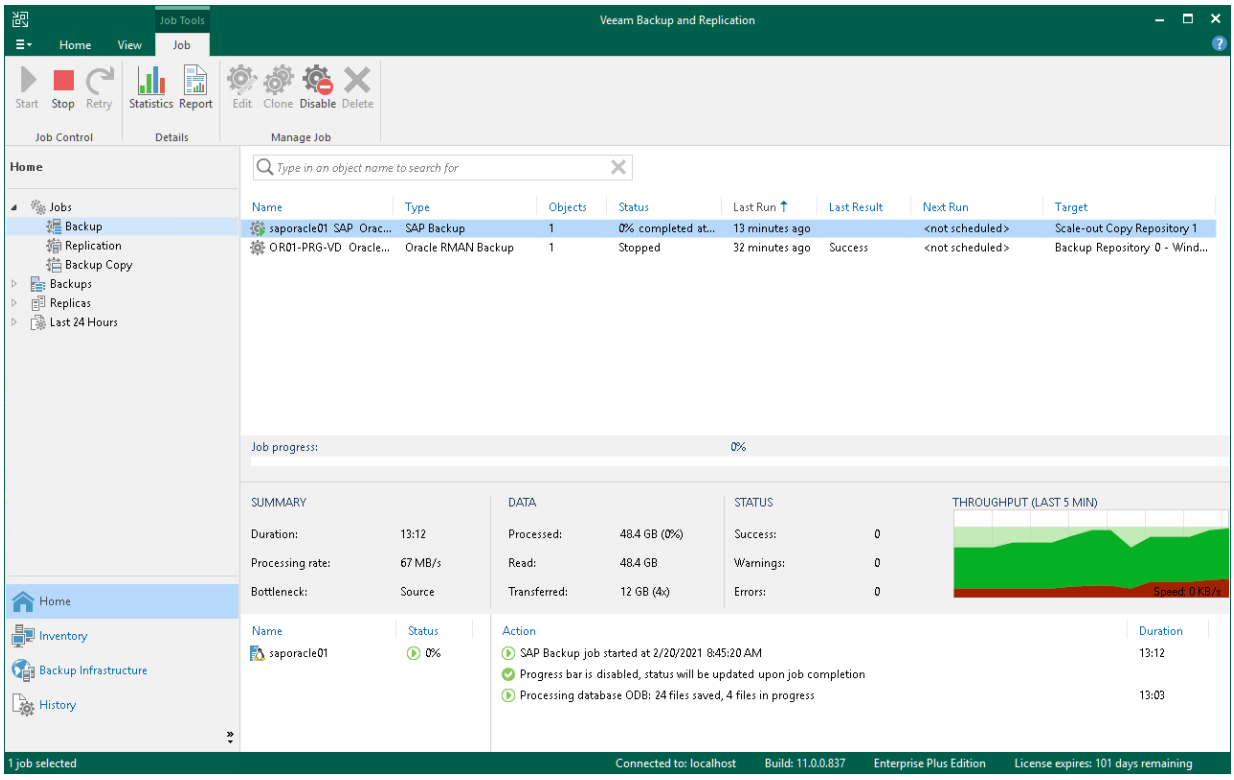

Generating Backup Job Reports

Veeam Backup & Replication can generate reports with details about an BR*Tools backup job session performance. The session report contains the following session statistics: session duration details, details of the session performance, amount of read, processed and transferred data, backup size, compression ratio, list of warnings and errors (if any).

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

Disabling Backup Job

You can disable BR*Tools backup jobs in the Veeam Backup & Replication console. If you disable the job, you will not be able to run BR*Tools backup commands on the SAP on Oracle server.

To disable a backup job:

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

Database Recovery

With the configured Veeam Plug-in you can restore Oracle databases from the backups that reside on backup repositories. All restore operations are performed on the SAP BR*Tools side.

Keep in mind that examples provided in this section are for demostration purposes only. To see the full restore functionality of SAP BR*Tools, see the BR*Tools for Oracle DBA Guide.

To learn how to recover Oracle databases from backups created by Veeam Plug-in for SAP on Oracle, see:

- Restoring Oracle Databases
- Restoring Redo Logs
- Recovering Databases to Other Servers (System Copy)
- Restore from Backup Copy

Oracle Databases Restore

Veeam Plug-in for SAP on Oracle allows to restore databases using the BRRESTORE tool functionality. When you launch the restore, BRRESTORE restores the selected database from backup files stored on the backup repository.

By default, BRRESTORE uses the `initSID.sap` initialization profile. Thus, you must specify the `-p $Oracle_HOME/dbs/veeam_initSID.sap` parameter in the restore commands.

For details on all restore options, see the Command Options for BRRESTORE section of the SAP DATABASE Guide: Oracle.

Example: Performing Full Restore of SAP on Oracle Database

```
brrestore -d util_file -p $Oracle_HOME/dbs/veeam_initSID.sap -b last -m full
```

Run the `brrestore` command with the following parameters:

1. Depending on which backup you want to restore from, use the `util_file` or `rman_util` option as the argument for the `-d` (`-device`) parameter. If the backup was created by Backint, use `util_file`. If the backup was created by RMAN, use `rman_util`.
2. Specify the path to the initialization profile file (`$Oracle_HOME/dbs/veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.
3. Specify `last` as the argument for the `-b` (`-backup`) parameter. With this option, BRRESTORE uses the last successful database backup for the restore.
4. Specify `full` as the argument for the `-m` (`-mode`) parameter. With this option, BRRESTORE performs restore of files in all tablespaces, control files and redo log files.

Redo Logs Restore

If you want to restore redo log files that were backed up with BRARCHIVE, you can use the BRRESTORE tool. For details, see the Names of BRRESTORE Details Logs section of the SAP Database Guide: Oracle.

Example: Performing Restore of SAP on Oracle Redo Logs

```
brrestore -d util_file -p $Oracle_HOME/dbs/veeam_initSID.sap -a 1-100
```

Run the `brrestore` command with the following parameters:

1. Depending on which backup you want to restore from, use the `util_file` or `rman_util` option as the argument for the `-d` (`-device`) parameter. If the backup was created by Backint, use `util_file`. If the backup was created by RMAN, use `rman_util`.
2. Specify the path to the initialization profile file (`$Oracle_HOME/dbs/veeam_initSID.sap`) as the argument for the `-p` (`-profile`) parameter.
3. Specify the log sequence number interval as the argument for the `-a` (`-archive`) parameter. This option defines which log files must be restored.

Database Restore to Another Server (System Copy)

You can restore SAP on Oracle databases from Veeam Plug-in backups to another server. To restore databases to another server, you must reconfigure settings of Veeam Plug-in as shown below.

For security reasons, you can restore databases to another server only in the following condition. The account you use to connect to Veeam Backup & Replication server must be the same account that performed the backup of the source system. If you want to use another account, you can assign the **Veeam Backup Administrator** or **Veeam Restore Operator** roles to the required account. For details on assigning Veeam Backup & Replication roles, see the Roles and Users section of the Veeam Backup & Replication User Guide.

Procedure

To restore databases to another server, you must reconfigure settings of Veeam Plug-in as shown below:

1. Go to `/opt/veeam/VeeamPluginforSAPOracle` and run the following command to select the source server whose backups you want to use during restore.

```
VM2ADM:/opt/veeam/VeeamPluginforSAPOracle> SapOracleBackintConfigTool --se
t-restore-server
```

2. Select the required SAP on Oracle server.

```
Select source SAP on Oracle server to be used as a restore target:
0. To disable this functionality.
1. saporacle01
Enter server number: 1
```

3. Specify a backup repository where the required backup files are stored.

```
Available backup repositories:
1. win_repo02
2. main_repo
Enter server number: 2
```

NOTE

- The account used to connect to Veeam Backup & Replication server must have access permissions on the required repository. Otherwise the repository will not be displayed in the list of available repositories. To learn how to configure access permissions on repositories, see Setting Up User Permissions on Backup Repositories.
- The wizard does not import existing backups from the repository. To perform a System Copy restore from the imported backup, you must map the backup. For details, see Importing Backups.

4. Perform the restore to another server.

5. After the restore, you must revert back the restore-server option of the Veeam Plug-in configuration wizard. Otherwise, you will not be able to restore data from the actual server backup file. If you perform only restore to other server, leave this setting enabled. It will not affect the backups of the actual system.

   To disable the functionality, run the `--set-restore-server` command and enter *0*.

```
SapOracleBackintConfigTool --set-restore-server
Select source SAP on Oracle server to be used as a restore target:
0. To disable this functionality.
1. saporacle01
Enter server number: 0
```

Restore from Backup Copy

You can restore Oracle databases from backups and backup copies. To restore from backup copies, you must enable the restore from backup copy option in the Veeam Plug-in wizard.

> **IMPORTANT**
>
> If the restore from backup copy option is enabled, you cannot back up databases using Veeam Plug-in, and you cannot restore from backups created by primary Veeam Plug-in backup jobs. You can restore only from backup copy files until you disable the restore from backup copy option.

For instruction on how to enable/disable the restore from backup copy option, see the following guides:

- Enabling Restore from Backup Copy
- Disabling Restore from Backup Copy

Enabling Restore from Backup Copy

To be able to restore from backup copies, do the following:

1. In the machine where Veeam Plug-in is installed, open the terminal and run the following command:

```
SapOracleBackintConfigTool --configure-restore-from-copy
```

2. Select the number of the backup copy job you want to use:

```
Select secondary job for failover:
0. Disable
1. Plug-ins backup copy job\linuxq01 SAP Oracle backup <serv10_repo>
Select secondary job for failover:1
```

> **IMPORTANT**
>
> The account used to connect to the Veeam Backup & Replication server must have access permissions on the required repository.

Disabling Restore from Backup Copy

To be able to back up with Veeam Plug-in and restore from backups, disable the restore from backup copies (set the parameter back to **0**):

```
SapOracleBackintConfigTool --configure-restore-from-copy
Select secondary job for failover:
0. Disable
1. Plug-ins backup copy job\linuxq01 SAP Oracle backup <serv10_repo>
Select secondary job for failover:0
```

Retention of SAP on Oracle Backups

To set an automatic removal of old backups, you can use the retention policy of Veeam Plug-in for SAP on Oracle. The `--set-force-delete` command of Veeam Plug-in automatically deletes backup files which are older than specified number of days. For details, see Configuring Retention Policy for Backups.

Also, you can manually delete backups from a backup repository using the Veeam Backup & Replication console and. For details, see Deleting Backups Manually Using Veeam Backup & Replication Console.

Configuring Retention Policy for Backups

Veeam Plug-in for SAP on Oracle has a functionality that automatically deletes backup files which are older than specified number of days. For example, you can use it if a backup repository contains backup files that are no longer in the backup catalog.

1. To enable automatic deletion of backup files, run the following command.

   ```
   SapOracleBackintConfigTool --set-force-delete
   ```

2. Enter the number of days after which Veeam Plug-in will delete backup files on all configured backup repositories.

   ```
   Garbage collector automatically deletes backup files older than the specif
   ied number of days.
   Make sure the number of days value exceeds your retention policy.
   To disable this functionality, set the number of days to 0.
   Enter the number of days to delete backups after, between 7 and 999 [0]:
   ```

   By default, the force delete functionality is disabled (set to *0*).

IMPORTANT

A value for the **number of days** setting must be at least 1 backup generation period longer than the retention period for your Oracle Database backups. Otherwise, Veeam Plug-in will delete earliest backups created within the retention period.

Deleting Backups Manually

Apart from configuring the retention policy, you can delete backups manually from backup repositories using the Veeam Backup & Replication console.

> **NOTE**
>
> If you remove backups from a backup repository manually, the backup catalog will not be updated.

To remove a backup from a backup repository, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the **Inventory** pane, select **Backups**.
3. In the working area, right-click the backup job object name and select **Delete from disk**.

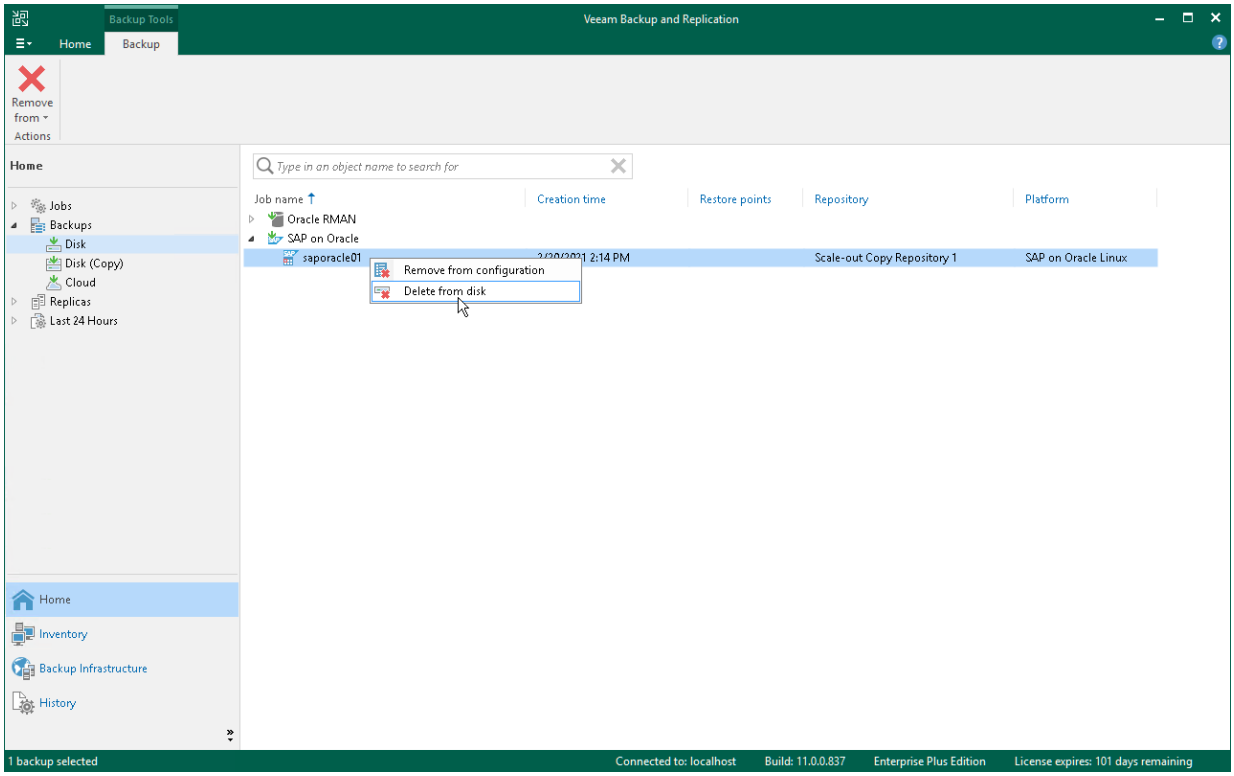

Removing Backups from Configuration

If you want to remove records about backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation.

When you remove a backup from the configuration, backup files (VAB, VASM) remain on the backup repository. You can import backup files later and restore from them.

To remove a backup from the configuration:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, select the backup and click **Remove from** > **Configuration** on the ribbon. You can also right-click the backup and select **Remove from configuration**.

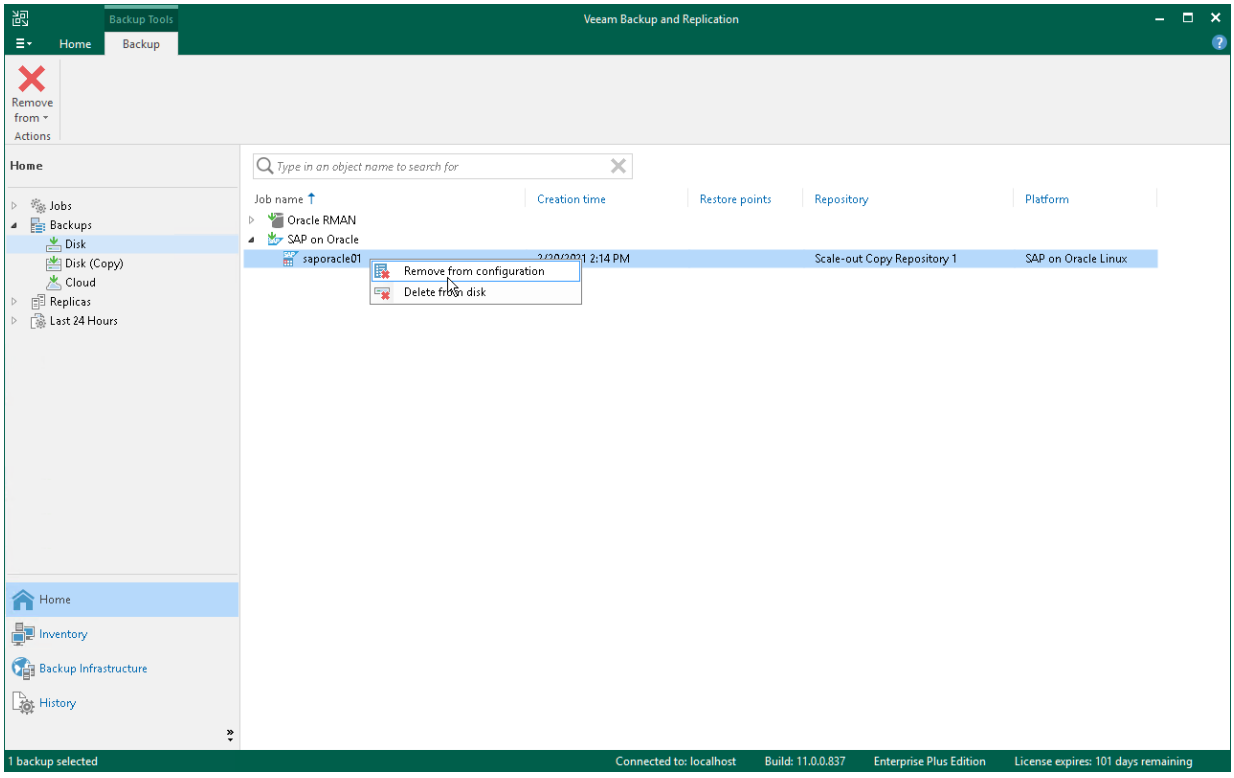

Backup Copy for SAP on Oracle Backups

Having just one backup does not provide the necessary level of safety. The primary backup may get destroyed together with production data, and you will have no backups from which you can restore data.

To build a successful data protection and disaster recovery plan, it is recommended that you follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups.
- 2: You must use at least two different types of media to store the copies of your data, for example, local disk and cloud.
- 1: You must keep at least one backup offsite, for example, in the cloud or in a remote site.

Thus, you must have at least two backups and they must be in different locations. If a disaster takes out your production data and local backup, you can still recover from your offsite backup.

To help you adopt the 3-2-1 rule, Veeam Backup & Replication offers the backup copy functionality that allows you to create several instances of the same backup in different locations, whether onsite or offsite. Backup copies have the same format as those created by backup jobs and you can recover your data from them when you need it.

Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes. Backup copy is a job-driven process. When enabled, the backup copy job for Veeam Plug-in backups runs continuously.

For details on how it works, see the Backup Copy section of the Veeam Backup & Replication User Guide.

In This Section

- Creating Backup Copy Job
- Converting Backup Copy to Backup

Creating Backup Copy Job

To copy backups to a secondary location, you must configure a backup copy job. The backup copy job defines how, where and when to copy backups. One job can be used to process backups of one or more machines.

You can configure a job and start it immediately or save the job to start it later.

Before creating a job, check prerequisites. Then use the **New Backup Copy Job** wizard to configure a backup copy job.

1. Launch Backup Copy Job wizard.
2. Specify a job name and description.
3. Selects backups to process.
4. Define backup copy target.
5. Specify advanced settings.
6. Define backup copy schedule.
7. Finish working with the wizard.

Before You Begin

Before you create a backup copy job, check the prerequisites and limitations:

- Backup infrastructure components that will take part in the backup copy process must be added to the backup infrastructure and properly configured. These include source and target backup repositories between which backups must be copied.
- The target backup repository must have enough free space to store copied backups. To receive alerts about low space on the backup repository, configure global notification settings. For more information, see Specifying Other Notification Settings.
- For Veeam Plug-in backup copy jobs, you cannot select a Veeam Cloud Connect repository as a backup copy target.

> **IMPORTANT**
>
> [For BR*Tools backup with the `rman_util` parameter] To have the whole set of backup files, you must add both jobs to the backup copy job:
>
> 1. The job created by Veeam Plug-in for Oracle RMAN that contains database files or redo logs
> 2. The job created by Veeam Plug-in for SAP on Oracle that contains data control files and BR*Tools logs.

Step 1. Launch Backup Copy Job Wizard

To create a backup copy job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Click the **Backup Copy** tab and select **Application backup**.

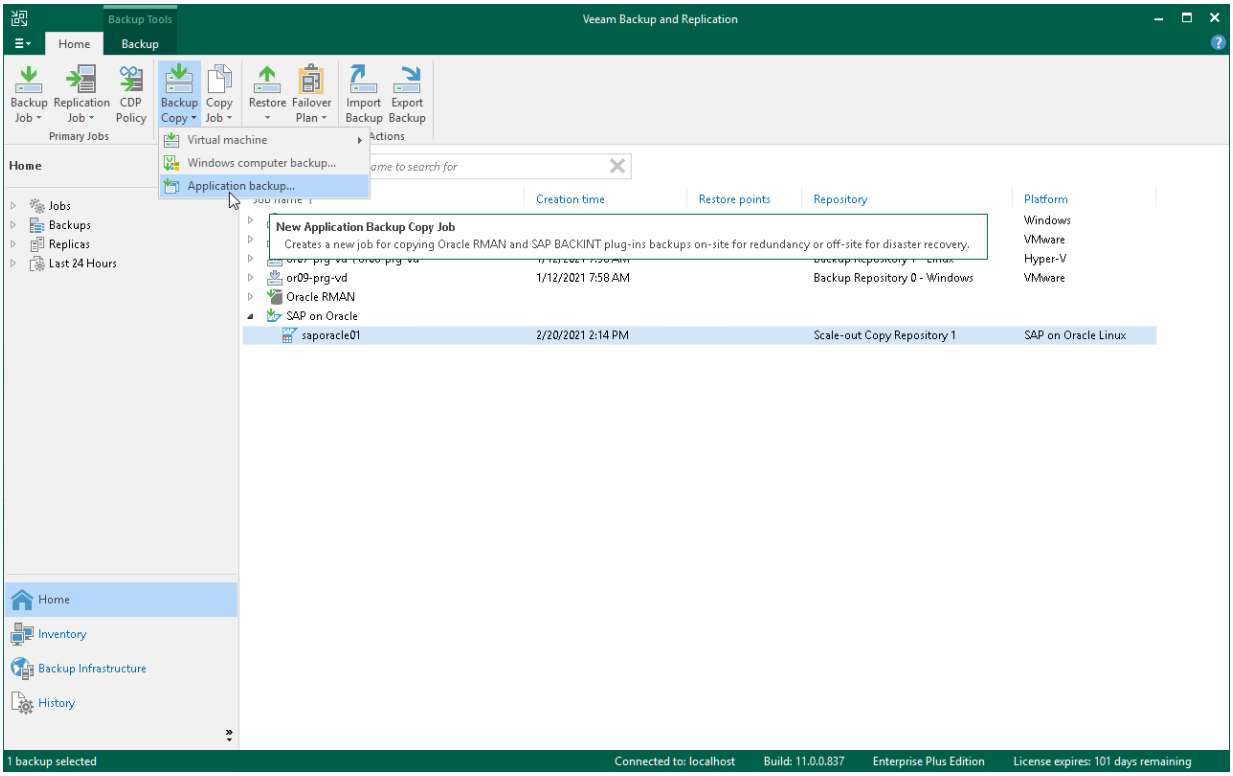

Step 2. Specify Job Name and Description

At the **Job** step of the wizard, specify a name and description for the backup copy job.

1. In the **Name** field, enter a name for the job.
2. In the **Description** field, enter a description for the job. The default description contains information about the user who created the job, date and time when the job was created.

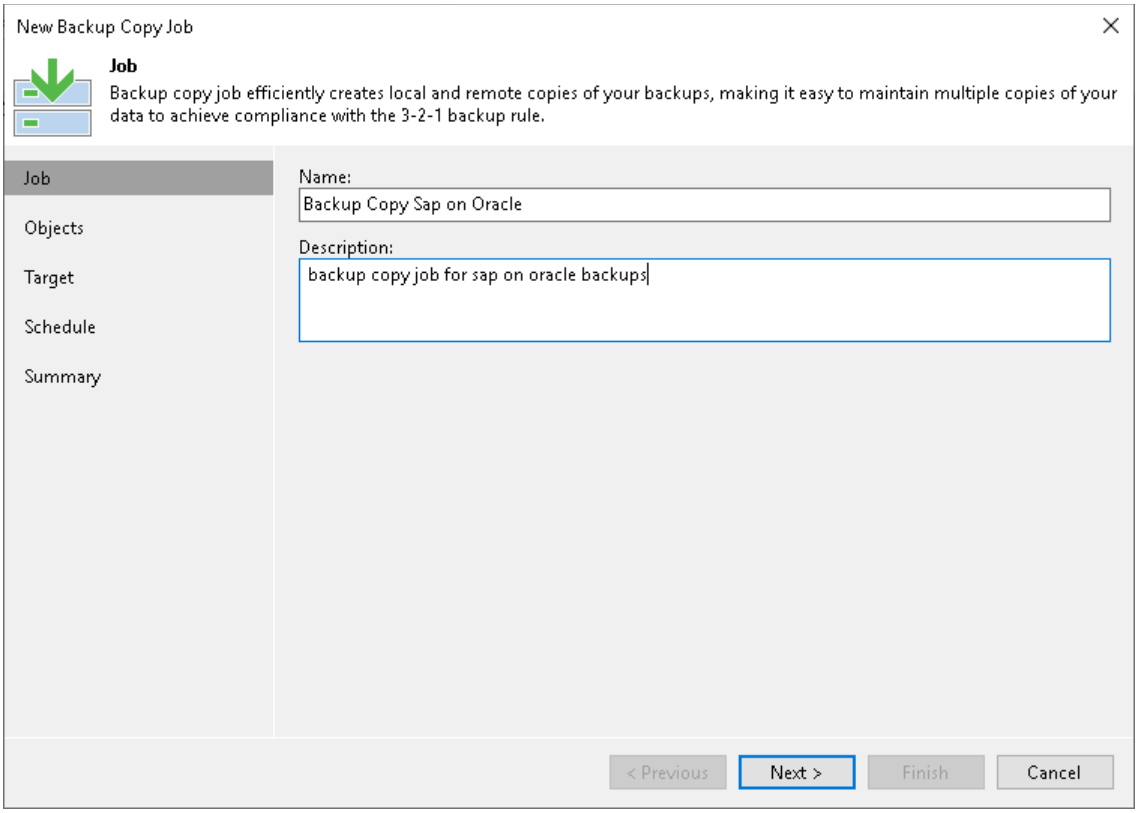

Step 3. Selects Backups to Process

At the **Object** step of the wizard, select machines whose backups you want to copy to the target repository.

1. Click the **Add** button and select from which entity you want to process the machines.
   - **From jobs**: You can select Veeam Plug-in backup jobs. When a backup copy job runs, Veeam Backup & Replication will search for backup files created by selected jobs.
   - **From repositories**: You can select repositories where Veeam Plug-in backups are stored. When a backup copy job runs, Veeam Backup & Replication will search for backup files created by Veeam Plug-in in selected repositories.

     [Optional] If you have added jobs from a repository and want to exclude from processing some of the backup jobs on the selected repository, click **Exclusions** and select the jobs that you want to exclude.
2. Use the **Remove** button if you want to remove selected jobs or repositories from processing.

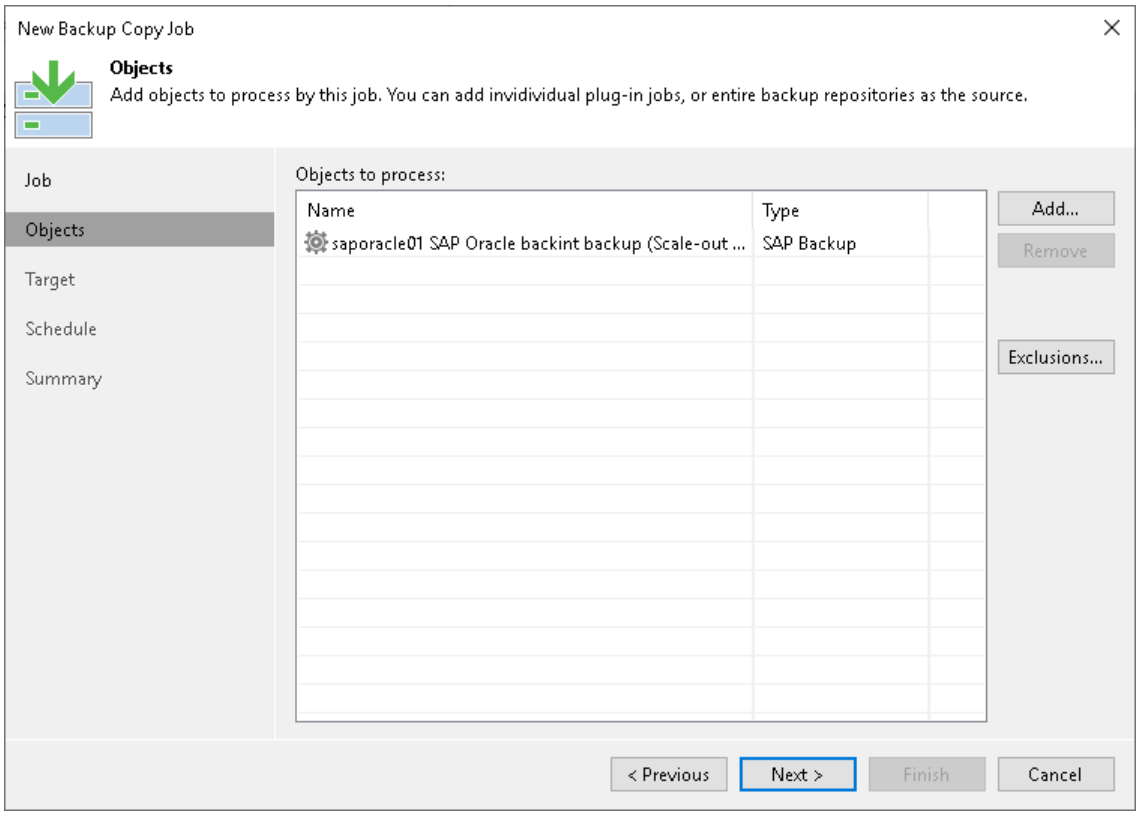

Step 4. Define Backup Copy Target

At the **Target** step of the wizard, configure the target repository settings.

1. From the **Backup repository** list, select a backup repository in the target site where copied backups must be stored. When you select a target backup repository, Veeam Backup & Replication automatically checks how much free space is available on it. Make sure that you have enough free space to store copied backups.

   > **IMPORTANT**
   >
   > For Veeam Plug-in backup copy jobs, you cannot select a Veeam Cloud Connect repository as a backup copy target.

2. [Optional] If the target repository contains a Veeam Plug-in backup that was excluded from the backup copy job, and if you don't want to transfer duplicate data, you can use the mapping feature.

   After you configure mapping, if some of backup files (.VAB) of the source backup are missing in the target backup copy, these files are uploaded to the target backup copy.

   > **NOTE**
   >
   > Veeam Plug-in backup copy jobs does not use WAN accelerators.

   To map a backup copy job to the backup:

   a. Click the **Map backup** link.

   b. Point the backup copy job to the backup in the target backup repository. Backups in the target backup repository can be easily identified by backup job names. To facilitate search, you can use the search field at the bottom of the window.

> **IMPORTANT**
>
> - Used account must have access to Veeam backup repositories that you plan to use.
> - Encryption must be disabled on the repository.
>
> Otherwise, the repositories will not be listed as available. To learn how to configure access permissions and encryption settings on repositories, see Access and Encryption Settings on Repositories.

3. You can specify the number of days after which the backup copy will be deleted from the repository. Note that the countdown starts from the moment when source backup has been created.

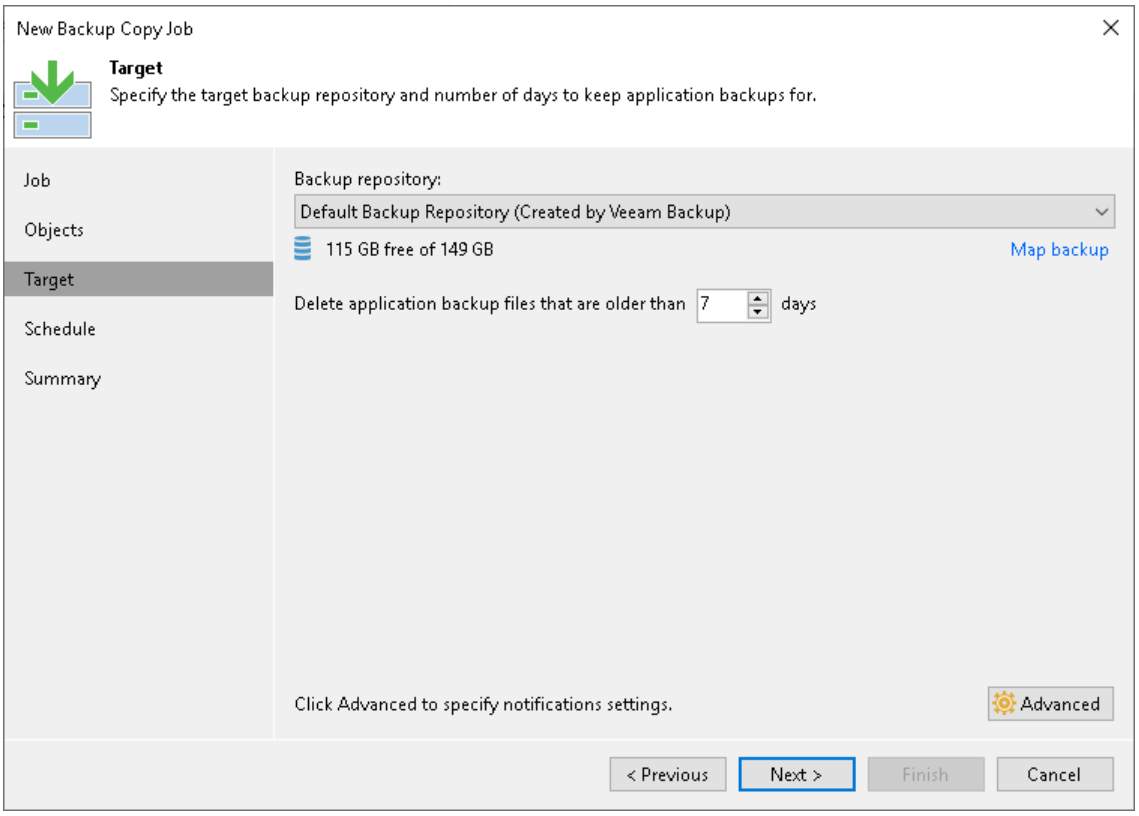

Step 5. Specify Advanced Settings

At the **Target** step of the wizard, click **Advanced** to configure compression, RPO and notifications settings.

- Compression and Deduplication
- RPO
- Notifications

Compression and Deduplication

At the **Storage** tab, define compression and deduplication settings.

By default, Veeam Backup & Replication performs deduplication before storing copied data on the target backup repository. Deduplication provides a smaller size of the resulting backup file but may reduce the job performance.

1. You can disable data deduplication. To do this, clear the **Enable inline data deduplication** check box.
2. From the Compression level list, choose a compression level to be used: **Auto**, **None**, **Dedupe-friendly**, **Optimal**, **High** or **Extreme**. The recommended level of compression for backup copy jobs is **Auto**. In this case, Veeam Backup & Replication uses compression settings of the copied backup files. For more information, see Compression and Deduplication.

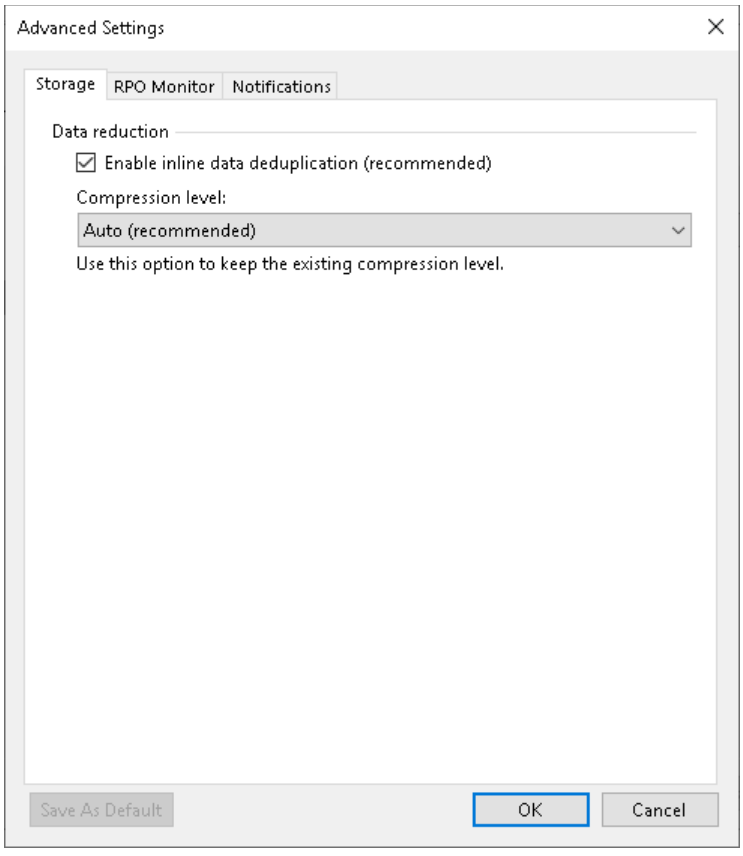

RPO Monitor

At the **RPO Monitor** tab, specify RPO warning settings.

Enable the **Warn me if backup is not copied within** check box and specify the time period in **minutes**, **hours**, or **days**.

If the backup copy is not created within the specified time period, the backup copy job will finish with the *Warning* status. The countdown starts from the moment when the required backup is finished and ready to be copied.

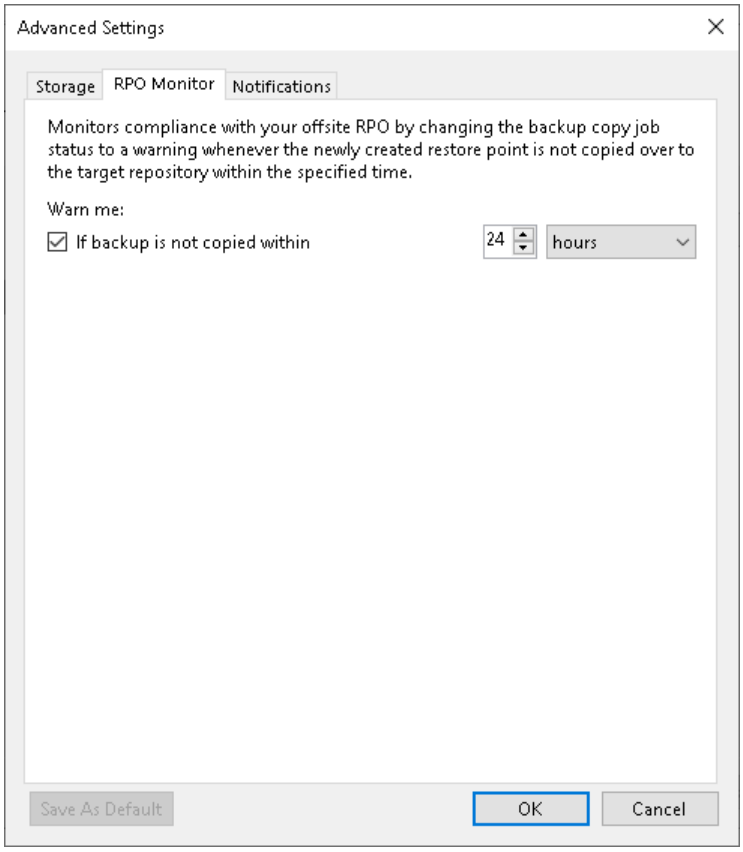

Notifications

At the **Notifications** tab, to specify notification settings for the backup copy job:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully. 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 by email in case of job failure or success. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon.
5. Veeam Backup & Replication sends a consolidated email notification once for the specified backup copy interval. Even if the synchronization process is started several times within the interval, for example, due to job retries, only one email notification will be sent.
6. 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.

7. At the **Send at** field, specify the time when you want to receive notifications. Note that you will receive a notification on a status of the job status once a day.

8. 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%*, *%VmCount%* (number of machines in the job) and *%Issues%* (number of machines in the job that have been processed with the **Warning** or **Failed** status).

     ii. Select the **Notify on success**, **Notify on warning** and **Notify on error** check boxes to receive email notification if data processing within the backup copy interval completes successfully, fails or completes with a warning.

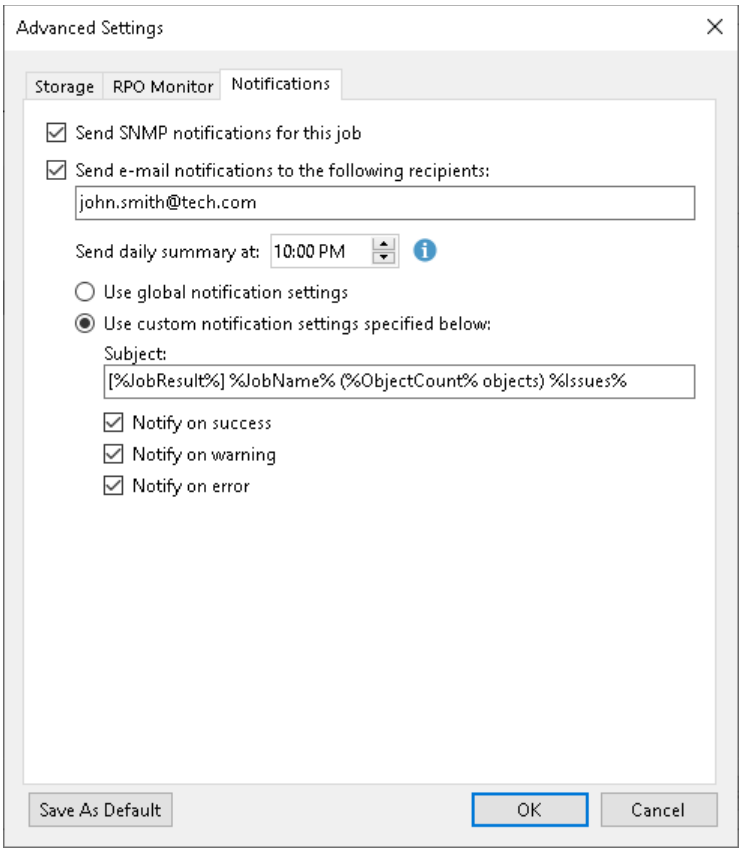

Step 6. Define Backup Copy Schedule

At the **Schedule** step of the wizard, define a time span in which the backup copy job must not transport data between source and target backup repositories. For more information, see Backup Copy Window.

To define a backup window for the backup copy job:

1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enable** and **Disable** options to mark the selected area as allowed or prohibited for the backup copy job.

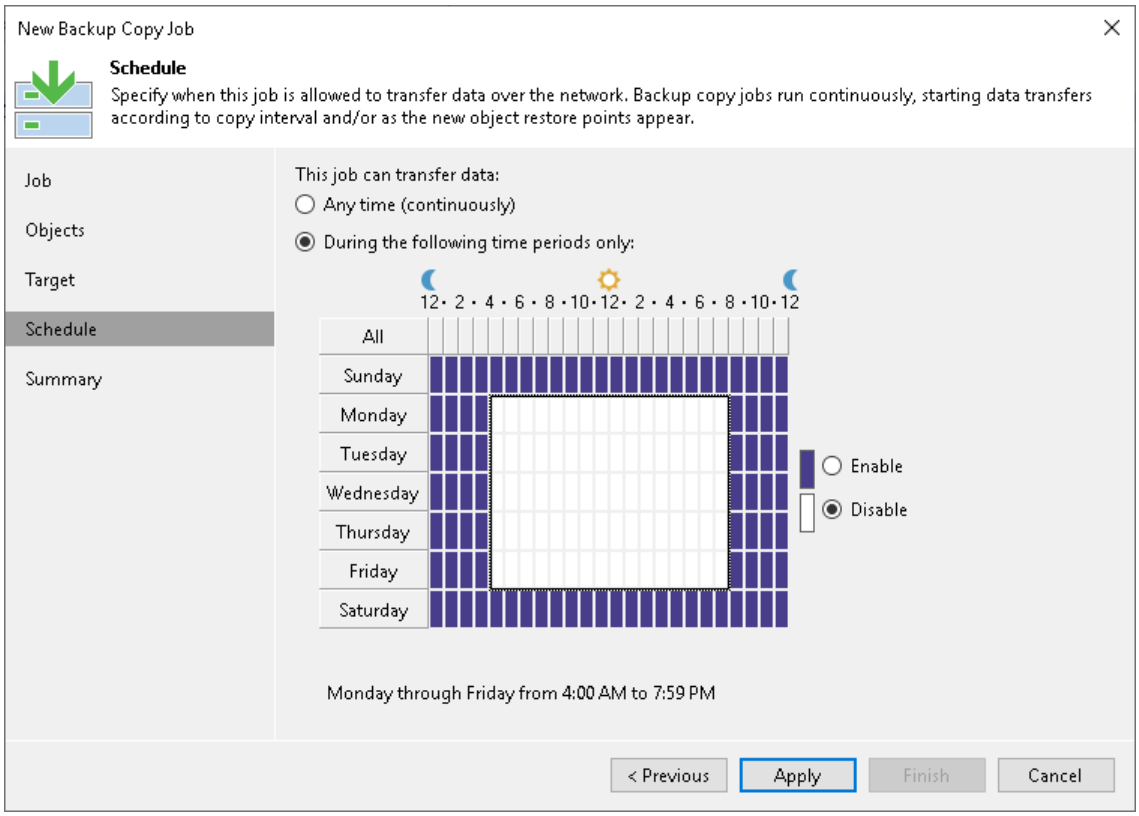

Step 7. Finish Working with the Wizard

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

1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

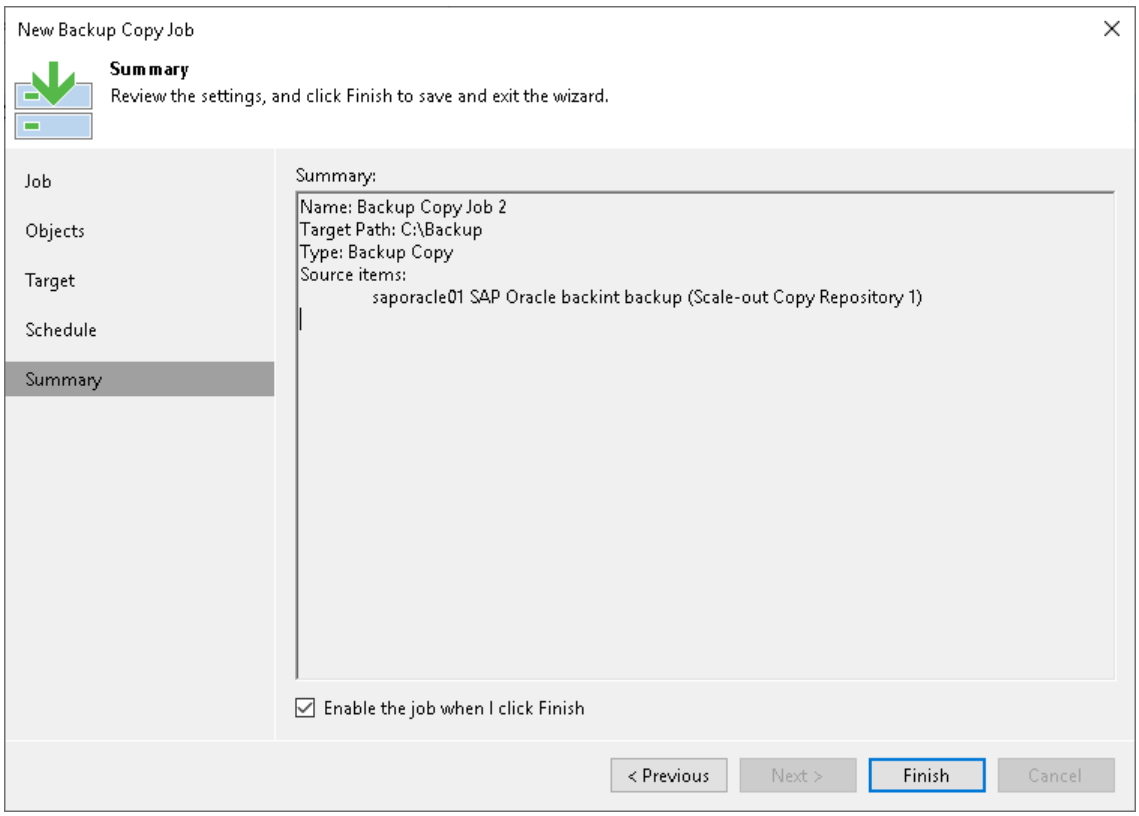

Converting Backup Copy to Backup

If you have imported backup copy files created by a backup copy job from another repository, you can convert them into regular backup files. When you convert backup copy files to regular backup files, Veeam Plug-in creates a backup job and adds attaches the converted backup files to it. You can use this backup job to continue the backup chain and use converted backup files as a restore point.

If you have imported backups created by a backup copy job from another repository, you can convert them into regular backup files. When you convert backup copy files to regular backup files, Veeam Plug-in creates a backup job and adds attaches the converted backup files to it. You can use this backup job to continue the backup chain and use converted backup files as a restore point.

You can convert imported Veeam Plug-in backups into regular Veeam Plug-in backup files in the following cases:

- If you have deleted a backup copy job which created the backup copy.
- If you have excluded a backup job from a backup copy job that used multiple backup jobs as a source.
- If you have imported a Veeam Plug-in backup copy from another repository.

> **NOTE**
>
> If you want to restore from a backup copy, you don't need to convert the backup copy to backup. For details, see Restore from Backup Copy.

Converting Backup Copy to Backup for SAP HANA

To convert a backup copy to a primary backup, use the **--promote-backup-copy-to-primary** parameter as shown below:

```
SapOracleBackintConfigTool --promote-backup-copy-to-primary
Backup copies available for promotion to primary backup:
1. Backup Copy Job 1\saprhel01-localdomain SAP backint backup (Default Backup R
epository)
Select backup: 1
Promotion of backup copy to a primary backup will reconfigure the plug-in to us
e a different repository. Continue? (y/N): y
```

> **IMPORTANT**
>
> [For servers with the `customServerName` option] To avoid failure of conversion of backup copies, the server name must be the same as the name used in the backup copy.

Logs and Support

If you have any questions or issues with Veeam Plug-in for SAP on Oracle or Veeam Backup & Replication, you can search for a resolution on Veeam Community Forums or submit a support case in the Veeam Customer Support Portal.

When you submit a support case, we recommend you attach necessary logs related to Veeam Plug-in operations.

Veeam Plug-in Logs

To export Veeam Plug-in logs, do the following:

1. In the Veeam Backup & Replication server, go to `%PROGRAMDATA%\Veeam\Backup\Plugin`.
2. Copy logs of the required backup or restore process.

SAP Backint Logs

To export SAP Backint logs, in the SAP on Oracle server, go to `/tmp/veeam_plugin_logs/<user_name>/` and copy the following files:

- `SapOracleBackint.log`
- `SapBackintOracleManager.log`
- `Agent.Source.log`

BRTools Logs

The *Detail* and *Summary* logs of BRTools, are stored in the **/oracle/SID/sapbackup** and **/oracle/SID/saparch** directories.

For details, see the following sections of the SAP Database Guide: Oracle: BRBACKUP Logs, BRRESTORE Logs, BRARCHIVE Logs, BRRECOVER Logs.

Veeam Backup & Replication

Version 11

Tape Devices Support Guide

February, 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 6
TAPE DEVICES SUPPORT 7
Supported Devices and Configuration 8
Tape Environment 11
Tape Devices Deployment 13
Connecting Tape Devices 14
Getting Started with Tapes 16
Tape Servers 17
Adding Tape Servers 18
Removing Tape Servers 25
Updating Tape Servers 26
Rescanning Tape Servers 27
Modifying Tape Servers 28
Replacing Tape Servers 29
Tape Libraries 30
Rescanning Tape Devices 31
Renaming Tape Devices 32
Removing Tape Devices 33
Replacing Tape Devices 34
Loading Tapes 35
Inventorying Tapes 39
Cataloging Tapes 40
Automated Drive Cleaning 42
Switching Libraries to Maintenance Mode 44
Tape Drives 45
Tapes 46
Viewing Tape Properties 47
Viewing Data on Tapes 49
Removing Tapes from Catalog 50
Erasing Tapes 51
Marking Tapes as Free 52
Ejecting Tapes 53
Exporting Tapes 54
Verifying Tapes 55
Copying Tapes 57
Restoring Entire Tapes 64

Tapes Availability and Write-Protection 65
WORM Tapes 66
Tape Barcode Labels 67
Relabeling Tapes 69
Media Pools 70
Media Sets 71
Backup Sets 73
Service Media Pools 74
Media Pools 75
GFS Media Pools 87
Moving Tapes to Another Media Pool 98
Modifying Media Pools 99
Removing Media Pools 100
Vaults 101
Creating Vaults 102
Modifying Vaults 106
Removing Vaults 107
Moving Tapes to Vault 108
Managing Tapes in Vault 111
Tape Protection 112
Protecting Tapes 113
Tape Data Retention 114
Virtual Full Backup 116
How Virtual Full Backup Works 118
Virtual Full Intervals 120
Tape Parallel Processing 122
Media Sets Created with Parallel Processing 125
Machines Backup to Tape 127
How Machines Backup to Tape Works 130
Before You Begin 131
Creating Backup to Tape Jobs 132
Linking Backup Jobs to Backup to Tape Jobs 153
GFS Backup to Tape 155
How GFS Backup to Tape Works 156
Creating GFS Tape Jobs 158
Tenant Backup to Tape 159
Viewing Backups on Tape 160
File Backup to Tape 161
How File Backup to Tape Works 162

Before You Begin 163
NDMP Servers Backup to Tape 164
Creating File to Tape Jobs 169
Viewing Files on Tape 181
VM Restore from Tape to Infrastructure 182
How Restoring VM from Tape to Infrastructure Works 183
Before You Begin 185
Restoring VM from Tape to Infrastructure 186
Backup Restore from Tape to Repository 192
How Restoring Backups from Tape to Repository Works 193
Restoring Backups from Tape to Repository 194
Tenant Restore from Tape 200
File Restore from Tape 201
How Restoring Files from Tape Works 202
Before You Begin 203
Restoring Files from Tape 204
Tape States Indicators 214
Tape Alerts 217
Tape Drive Alerts 218
Tape Changer Alerts 228

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

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.

Long-term archiving and compliance are listed as primary reasons for using tape. Tape appears to be one of the most widely used media for offsite backup. Using backup to tape, you can implement the '3-2-1' backup approach (3 copies, 2 types of media, 1 offsite location) considered as a best practice for data protection and disaster recovery. Veeam Backup & Replication offers support for tape devices and allows you to archive your data to tape and restore it from tape whenever needed.

Storing data to tapes provides you with the same data managing options as disk repositories. In particular, you can store full and incremental backups, apply user-defined retention settings to the archived data, select restore points and so on. Compared to disk storage, tape archive requires more time to retrieve files for restore if the tapes are stored offsite.

Supported Devices and Configuration

Veeam Backup & Replication supports Linear Tape-Open tape libraries starting from generation 3 (LTO3) or later:

- Physical libraries, standalone drives and virtual tape libraries.
- Partitions of the physical or virtual tape libraries presented to the Veeam backup server.

Within a library partition, only one generation of LTO drives are supported.

Limitations

VMware does not support tape drives connected directly to ESXi 5.x and later. For more information, see this VMware KB article.

Veeam Backup & Replication will potentially work with such a configuration, but if there is a hardware-related issue, Veeam Support will not provide help to solve it. This functionality is vendor-supported.

Affecting 3rd Party Tape Solutions

Veeam Backup & Replication does not support sharing tape devices with 3rd party tape-recording software. If you plan to run both Veeam Backup & Replication and 3rd party tape-recording software (for example, in your evaluation lab), consider that Veeam Backup & Replication by default will periodically lock the drive to perform rescan, preventing other software from recording.

To share a tape device, configure a partition of the tape library that will be used only by Veeam Backup & Replication.

Mind the following limitations:

- Do not connect partitions presented to the 3rd party tape software to a Veeam tape server.
- Do not install any 3rd party tape-recording software or software components on Veeam tape server.

Driver Installation Mode

Veeam Backup & Replication supports both exclusive and non-exclusive driver installation.

> **NOTE**
>
> If you use standalone tape drives, it is recommended to install drivers in non-exclusive mode.

Path Failover

Veeam Backup & Replication supports path failover for tape devices with multiple drives that manage multiple paths over multiple SANs.

> **NOTE**
>
> Veeam does not supply MPIO drivers. To get proper MPIO drivers, request them from the tape device vendor.

Industry Format

Veeam Backup & Replication uses the MTF (Microsoft Tape Format) industry format to write data to tape.

Supported Connection Types

You can connect the tape device directly or remotely.

- Direct connection:
  - Fibre Channel (FC)
  - Serial Attached SCSI (SAS)
  - SCSI
- Remote connection:
  - iSCSI
  - Switched FC fabric

Data Block Size

Drives use hardware dependent block sizes to read/write the tape data. Generally, the drives support a range of block sizes and report this range to Veeam Backup & Replication.

If you use a tape library with multiple drives or you use multiple standalone drives, Veeam Backup & Replication uses a unified block size to write data to tapes. 256k data block size is used if all drives can support it. If any drive can support only a smaller block size, for example, 64k, then Veeam Backup & Replication uses the largest block supported by all drives. To understand block sizes supported by drives, Veeam Backup & Replication collects the block size ranges reported by each drive and compares them. Note that the reported range is additionally limited by storage controllers settings used in your infrastructure. You can check the resulting range of block sizes supported by Veeam Backup & Replication for a particular drive in the Drives properties. For more information, see Working with Drives.

> **NOTE**
>
> If you connect the tape devices using HBA, Veeam Backup & Replication uses the block size configured for the HBA.

The block size is unified for:

- All drives in one library (if the drives support different block sizes)
- All standalone drives connected to one tape server.

To read data from tape, Veeam Backup & Replication requires that the tape is written with the block size from the supported range (shown in the drive properties).

Mind the block size range when working with the following tapes:

- Tapes with Veeam backups written by another tape library,
- Tapes with Veeam backups written on another tape server,
- Tapes written with other data transfer configuration settings,
- Tapes written on a 3rd party device.

> **IMPORTANT**
>
> You can restore from tapes written with block size that match the block size range set for the tape device.

Unknown Medium Changers

Veeam supports medium changers that have no Microsoft Windows drivers available. Make sure that such device is recognized as an unknown medium changer in the Microsoft Device Manager list.

It is recommended that you use tape devices with original equipment manufacturer (OEM) drivers.

Tape Environment

To administer tape backup and restore procedures, you must add the tape server to your Veeam Backup & Replication console. Adding the tape server to the backup infrastructure makes possible to transfer data between the disk storage and the tape archive.

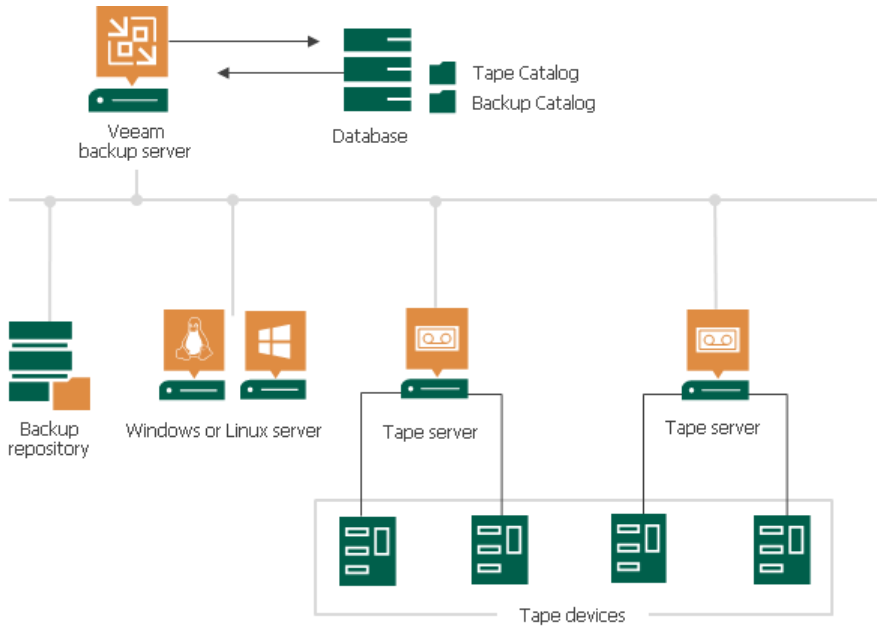

The tape backup and restore process involves the following components:

Veeam Backup Server

The Veeam backup server is the server running Veeam Backup & Replication. It is the core component in the backup infrastructure. The Veeam backup server tasks include the following operations:

- Recognizing all tape devices that are connected to it and managing all drives, slots and tapes.
- Administering tape archiving: scheduling and triggering tape jobs and initiating restore tasks.
- Communicating with the Veeam backup database.

Tape Device

When connected to the Veeam backup server, tape devices provide reading and writing capabilities, while all administration is performed by the Veeam backup server. You can connect multiple tape devices.

Tape Servers

To connect to the tape devices, Veeam Backup & Replication uses tape servers.

The tape server is a backup infrastructure component transferring data between data source and tape device. The tape server runs Veeam Data Mover that creates a communication point over which the data between the tape device and backup repositories or file servers is transferred.

The tape server deployment is fully automated: to create a tape server, you must assign the tape server role to a necessary server in the backup infrastructure. You can create a dedicated tape server or assign the tape server role to the Veeam backup server.

Veeam Backup Database

Veeam Backup & Replication catalogs information about all archived data and stores this information in the Veeam backup database.

The registered tapes stay in the database until you remove the information about them. You can always view details for each tape, for example, information about backups written to it, even if the tape is not inserted in the library.

The catalog lets quickly detect location of the required items on tape. The catalog correlates the archived files and their restore points to the names of the corresponding tapes, both online or offline, and the names of the media sets within which the data was written. When you start restoring data, Veeam Backup & Replication prompts for the tapes you need to bring online. As a result, you can restore data from tape much quicker when necessary.

Veeam Backup & Replication uses the following catalogs for storing the tape-related data:

- **Tape Catalog** stores information about files/folders archived to tape media with file to tape jobs, as well as backup files produced by backup to tape jobs. The content of the Tape Catalog can be examined in the **Files** view.
- **Backup Catalog** stores information about machines whose backups are archived to tape media with backup to tape jobs. The content of the Backup Catalog can be examined under the **Backups > Tape** node in the **Backup & Replication** view.

Backup Repositories and File Servers

When retrieving data to back up, or restoring data to disk, Veeam Backup & Replication can connect to any machine that has been added as a managed server to the Veeam Backup & Replication console.

Tape Devices Deployment

To connect tape devices to Veeam Backup & Replication, you need to deploy a tape server. Tape servers are network appliances that connect tape libraries to the Veeam backup server and manage traffic between tape devices and Veeam backup server. The connected tape devices are recognized by the Veeam Backup & Replication automatically.

Data Movers

With Veeam Backup & Replication, the data transfer during archiving and restore processes is enabled with Veeam Data Movers. Veeam Data Movers run on tape servers and other components of backup infrastructure. They receive tasks from the Veeam backup server and communicate to each other to transfer the data.

Veeam Data Movers are light-weight services that take a few seconds to deploy. Deployment is fully automated: when you assign a tape server role to a server, Veeam Backup & Replication installs the necessary components on this server and starts the required services on it.

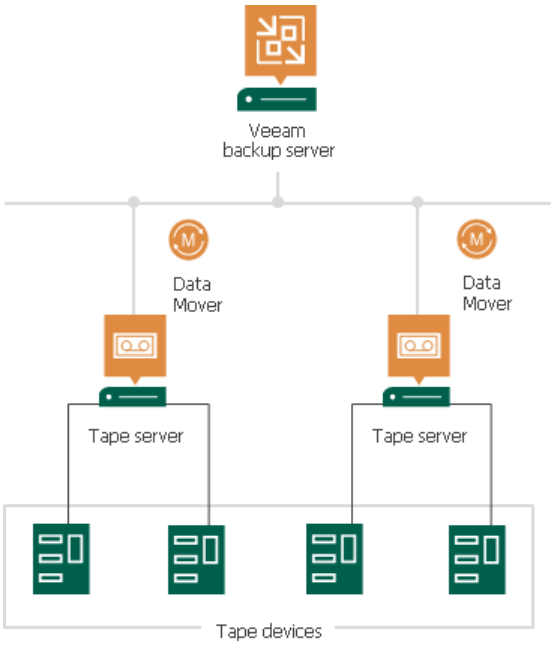

Connecting Tape Devices

To deploy a tape device, you must perform the following steps:

1. Choose a machine that will act as the tape server. You can choose between the following options:
   - Use the Veeam backup server as the tape server. This variant is sufficient for small environments with low traffic workloads.
   - Deploy a dedicated tape server. As a tape server, you can use any Microsoft Windows server, physical or virtual, accessible by Veeam Backup & Replication. A dedicated tape server is recommended for installations with intense data transferring. A dedicated tape server allows you to manage remote tape libraries.
2. Connect the tape device to the machine that will perform the role of the tape server. You can connect the tape device in two ways:
   - Directly over Fibre Channel (FC), Serial Attached SCSI (SAS), SCSI
   - Remotely with iSCSI (you can use Microsoft iSCSI initiator to establish the connection).
3. Install an appropriate device driver on the tape server.
4. If the medium changer driver is not available, check that the tape device uses native SCSI commands.

   Open the **Tape Infrastructure** view, expand the **Libraries** node and select the needed library. Click **Properties** on the ribbon. You can also right-click the necessary library in the working area and select **Properties**. Select the **Use native SCSI commands instead of Windows driver** check box.

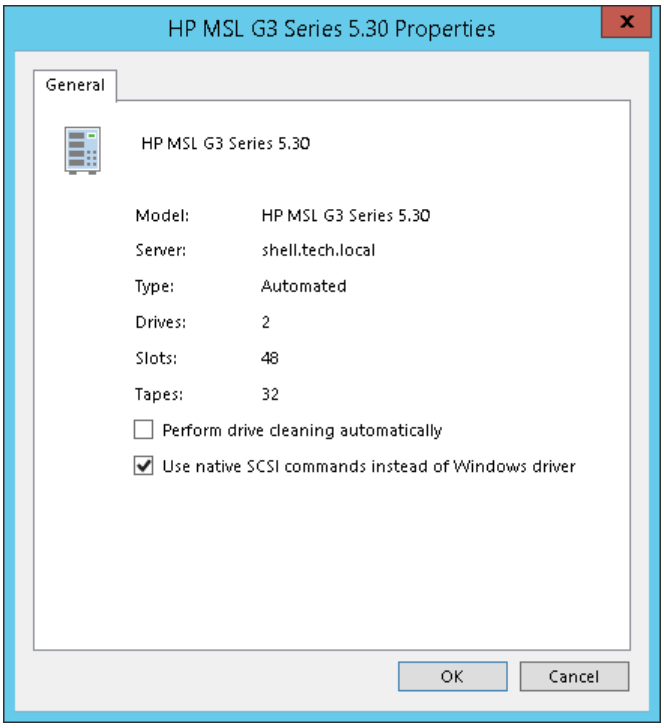

> **IMPORTANT**
>
> Only media changers can use SCSI commands. Tape drives must use native OEM drivers.

5. Add the tape server to Veeam Backup & Replication. For more information, see Adding Tape Servers. Veeam Backup & Replication will automatically scan connected tape devices and display all discovered tape libraries and tape drives under the **Libraries** node in the **Tape Infrastructure** view. Afterward, the auto-discovery process will be performed periodically every 3 minutes.

You can connect multiple tape devices to each tape server. The connected tape devices will appear as list under the **Libraries** node in the **Tape Infrastructure** view.

Veeam Backup & Replication allows deploying multiple tape servers. However, if you have multiple Veeam backup servers, you cannot connect one tape server to several Veeam backup servers simultaneously.

IMPORTANT

We strongly recommend not to connect each tape device to multiple tape servers at a time, as it may lead to conflicts.

TIP

If you have several Veeam backup servers, you can easily reconnect a tape server to another Veeam backup server without reconfiguring the tape device: Veeam backup server will recognize the library settings automatically. Note that when you reconnect the tape server, the tape jobs will not run with another Veeam backup server unless you copy the configuration.

Reconnecting Tape Devices to Another Tape Server

You can reconnect your tape devices to another tape server. For example, you can reconnect the tape devices from the Veeam backup server to a dedicated tape server. To do this, choose a tape server and perform steps for connecting tape devices above. When you add the new tape server to Veeam Backup & Replication, Veeam Backup & Replication will recognize the tape devices and all tape device settings.

Getting Started with Tapes

After you have connected tape devices, you need to complete the following steps:

1. Load tapes to the tape device (if not yet loaded) and run the importing procedure for them.
2. Create one or more media pool or GFS media pool that will be used as targets for tape jobs.
3. Configure one or more media vaults to conveniently organize offline tapes.
4. Configure and run backup to tape or file to tape jobs.
5. Restore virtual machines to infrastructure or both virtual and physical machines to repository, or restore files from tape.

Tape Servers

When you have deployed a tape server and connected the tape device to it, you need to add the tape server to the Veeam backup server. To do so, you must assign the role of the tape server to a Windows server that is already added to the list of managed servers.

You can add multiple tape servers to the same installation of Veeam Backup & Replication. But mind that you cannot add one tape library to two different tape servers and write data to it from both tape servers.

Adding Tape Servers

To add a tape server, follow the next steps:

Step 1. Launch New Tape Server Wizard

To launch the **New Tape Server** wizard:

1. Open the **Tape Infrastructure** view.
2. Do one of the following:
   - In the inventory pane, right-click **Tape Servers** and select **Add tape server**.
   - In the inventory pane, select **Tape Servers** and click **Add Tape Server** on the ribbon.
   - In the inventory pane, select the **Tape Infrastructure** node and click **Add Tape Server** in the working area.

Step 2. Choose Server

At the **Server** step of the wizard, choose a physical or virtual server to which the tape devices are connected and that you want to add as the tape server.

1. From the **Choose server** list, select the server that you want to add as the tape server.
   - If the devices are connected to the Veeam backup server, select **This server**. The tape server role will be assigned to the backup server.
   - If the devices are connected to a separated server, enter the server name or IP address. The tape server must run Microsoft Windows. If the server is not added to Veeam Backup & Replication yet, you can click **Add New** to open the **New Windows Server** wizard.
2. In the **Description** field, provide a description for future reference.

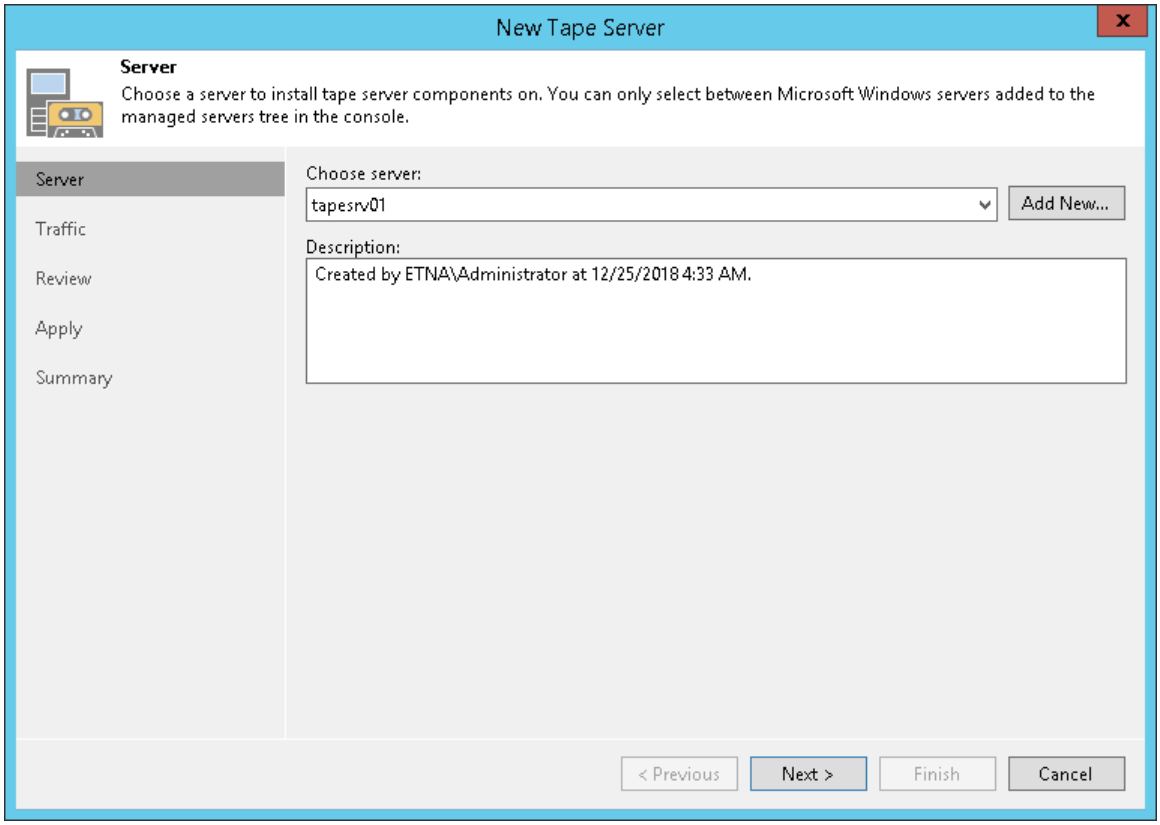

Step 3. Configure Traffic Throttling Rules

At the **Traffic** step of the wizard, configure throttling rules to limit the outbound traffic rate for the tape server. Throttling rules will help you manage bandwidth usage and minimize the impact of tape jobs on network performance. For more information, see the Veeam Backup & Replication User Guide, section *Enabling Traffic Throttling*.

The list of throttling rules contains only those rules that are applicable to the tape server you are adding. The rule is applied to the tape server if its IP address falls under the source IP range of the rule. To view the rule settings, select it in the list and click the **View** button on the right.

You can also open global throttling settings and modify them directly from the wizard by clicking **Manage network traffic throttling rules** at the bottom of the window.

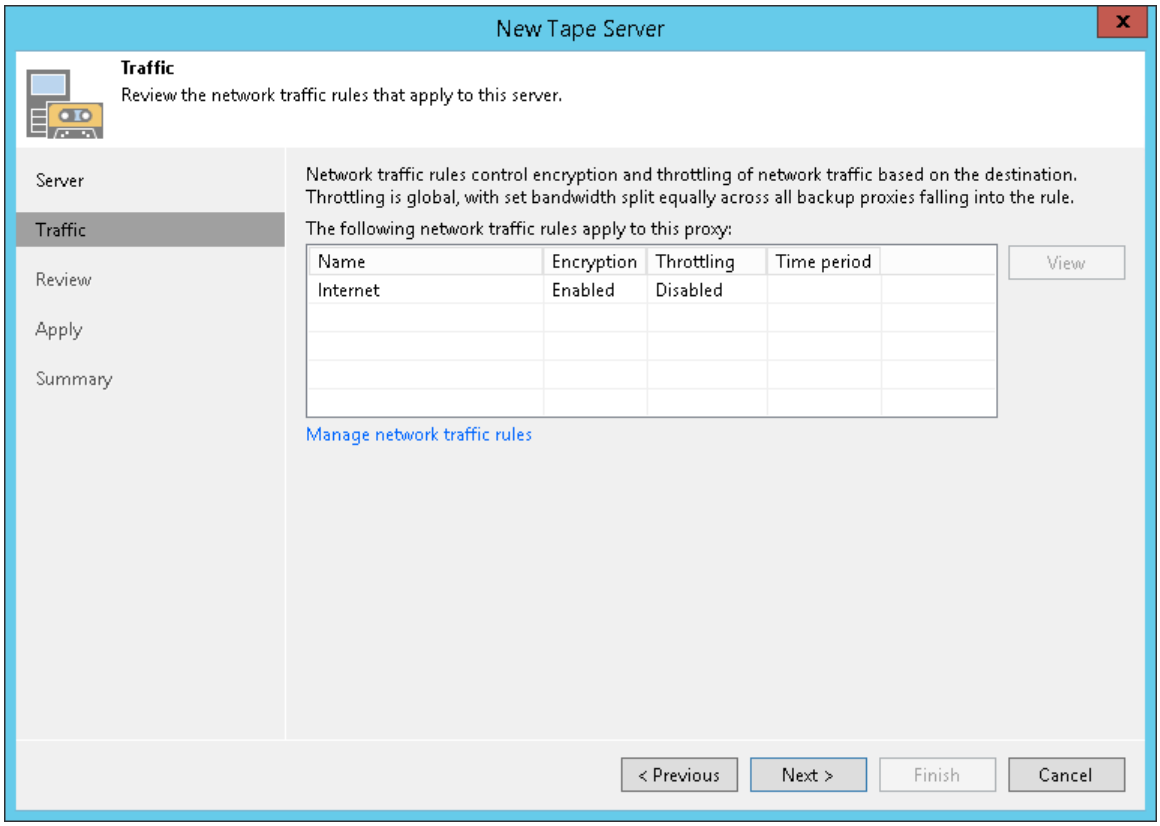

Step 4. Review Components

At the **Review** step of the wizard, Veeam Backup & Replication will display the list of components required for work of the tape server:

- Veeam Transport
- Veeam Tape service

If any of them is missing, Veeam Backup & Replication will automatically install them on the selected server.

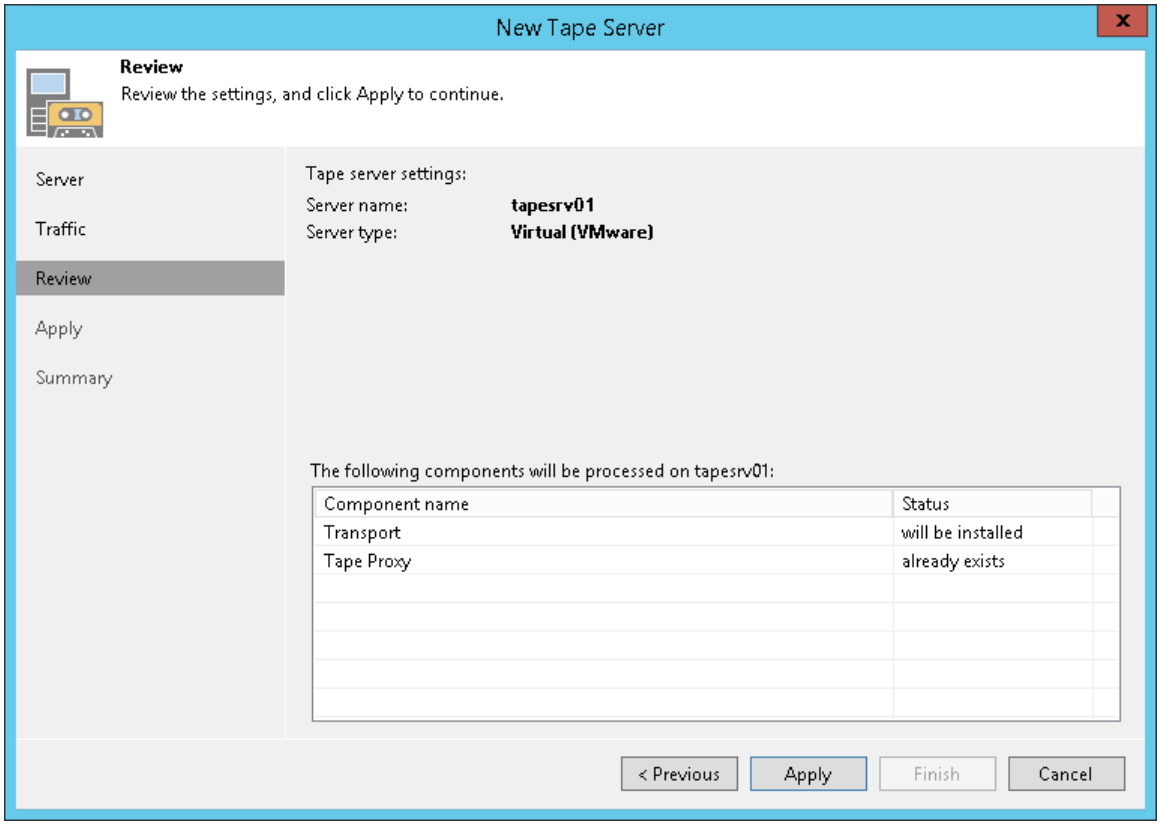

Step 5. Assess Results

At the **Apply** step of the wizard, Veeam Backup & Replication will add the tape server to the backup infrastructure in the real time mode.

If Veeam Backup & Replication detects an unknown media changer, the message will appear in the **Message** area.

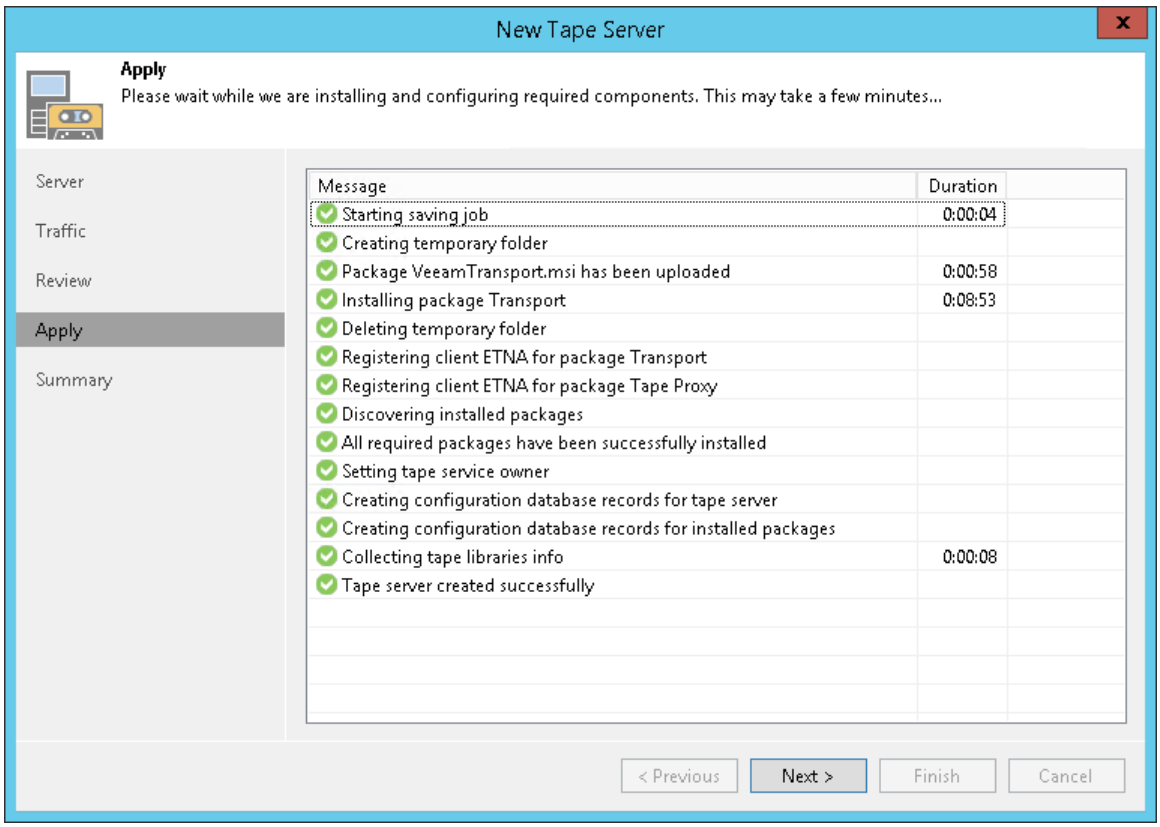

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, finalize the procedure of the tape server configuration. Review details of the added tape server.

To inventory tape libraries connected to the tape server, select the **Start tape libraries inventory when I click Finish** check box. Veeam Backup & Replication will start the inventory process when you finish working with the wizard. If you do not enable this option, you will have to inventory tape libraries manually. For more information, see Inventorying Tapes.

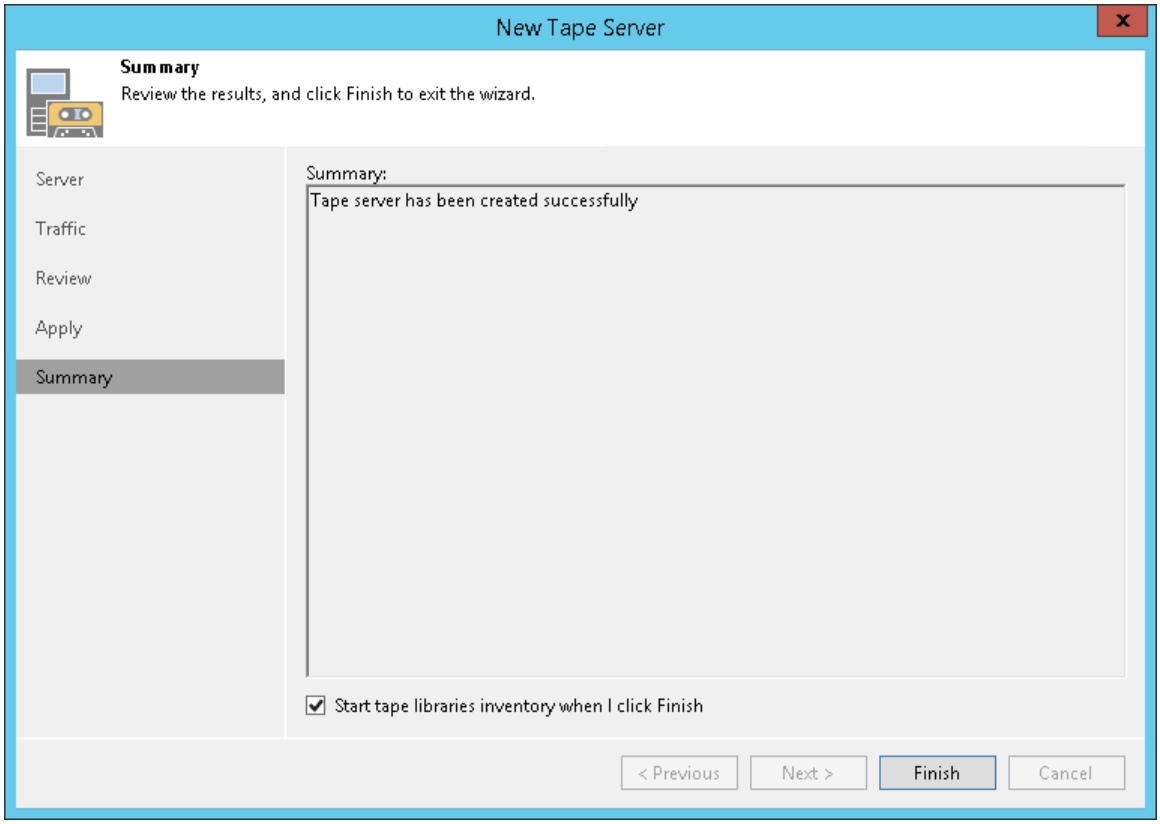

Removing Tape Servers

To remove a tape server:

1. Open the **Tape Infrastructure** view.
2. Select the **Servers** node in the inventory pane.
3. Select the tape server in the working area and click **Remove Tape Server** on the ribbon. You can also right-click the necessary tape server in the working area and select **Remove**.

When you remove a tape server, Veeam Backup & Replication unassigns the tape server role from it, so it is no longer used as a tape server. The actual server remains connected to Veeam Backup & Replication.

Limitations for Removing Tape Servers

You cannot remove a tape server while tape jobs are running. Before you remove the tape server, stop all tape jobs.

Updating Tape Servers

Every time you launch Veeam Backup & Replication, it automatically checks if the components installed on managed servers are up to date. You will be unable to work with tape libraries that are connected to the outdated servers.

If there is a later version of a component available (usually, if you have upgraded Veeam Backup & Replication), the **Components Update** window will be displayed, prompting you to update components on managed tape servers.

You can also open the **Components Update** window by doing one of the following:

- Open the main menu and select **Upgrade** from the main menu.
- Open the **Tape Infrastructure** view, select the **Servers** node and click **Upgrade** on the ribbon.
- Open the **Tape Infrastructure** view, right-click the necessary tape server in the working area and select **Upgrade**.

If components on all managed servers are up to date, the **Upgrade** item will be disabled.

The **Components Update** section lists tape servers that have outdated components deployed. To see the current and the latest available versions for deployed components, select a server in the list and click **Details**. Select check boxes next to servers for which you want to upgrade components and click **Apply**.

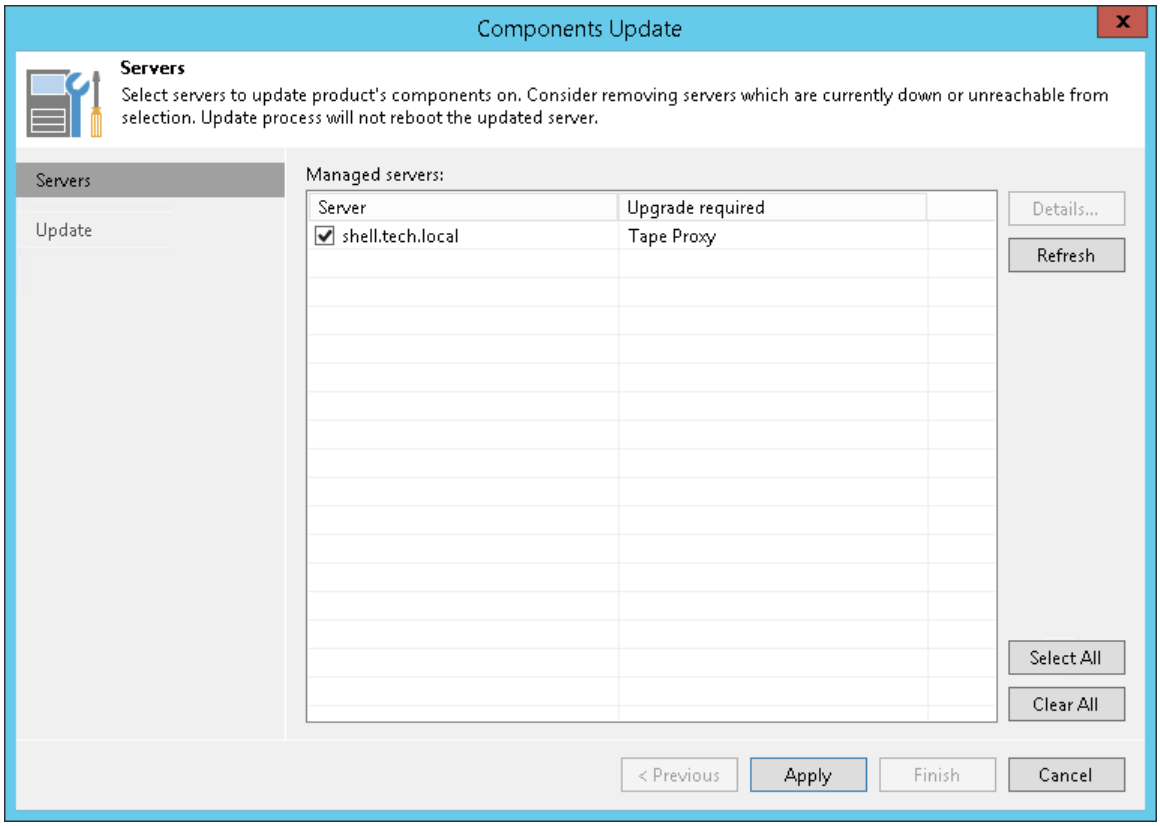

Rescanning Tape Servers

Veeam Backup & Replication automatically rescans the tape servers every 3 minutes. During the rescan, Veeam Backup & Replication updates the list of managed tape devices. You can perform manual rescan of a tape server, for example, if you want to implement the changes made to the tape infrastructure immediately.

To rescan a tape server:

1. Open the **Tape Infrastructure** view.
2. To rescan all tape servers, right-click the **Servers** node and select **Rescan**.
3. To rescan a selected tape server, select the **Servers** node in the inventory pane. Select the necessary server in the working area and click **Rescan** on the ribbon. You can also right-click the necessary server in working area and select **Rescan**.

Modifying Tape Servers

To edit settings of an added tape server:

1. Open the **Tape Infrastructure** view.
2. Select the **Servers** node in the inventory pane.
3. Select the necessary server in the working area and click **Edit Tape Server** on the ribbon. You can also right-click the necessary server in working area and select **Properties**.
4. Then edit the tape server settings as required.

Replacing Tape Servers

You can replace the tape server with another tape server and continue to use the existing libraries and tapes with it.

To correctly replace a tape server:

1. Physically disconnect libraries from the old tape server or disable the tape changer device and drives in the Device Manager on this tape server.
2. Remove the old tape server from the tape infrastructure, as described in Removing Tape Servers.
3. Add a new tape server, as described in Adding Tape Servers.
4. Connect your existing tape devices to the new tape server, as described in Connecting Tape Devices.

   All the information about existing tapes will be preserved after switching to the new tape server.

Tape Libraries

All tape libraries and standalone tape drives managed by Veeam Backup & Replication are shown in a list under the **Libraries** node in the **Tape Infrastructure** view. All connected devices are discovered automatically during the rescan procedure. When you add a new tape device to the tape server, it appears in your console after rescanning.

To view properties of a tape device, open the **Tape Infrastructure** view, expand the **Libraries** node and select the needed item. Click **Properties** on the ribbon. You can also right-click the necessary device in the working area and select **Properties**.

Select the **Perform drive cleaning automatically** check box if you want Veeam Backup & Replication to manage the tape library drive cleaning. For more information about automated drives cleaning, see Automated Drive Cleaning.

Select the **Use native SCSI commands instead of Windows driver** check box if your library is an unknown media changer. For more information about unknown media changers, see Supported Devices and Configuration.

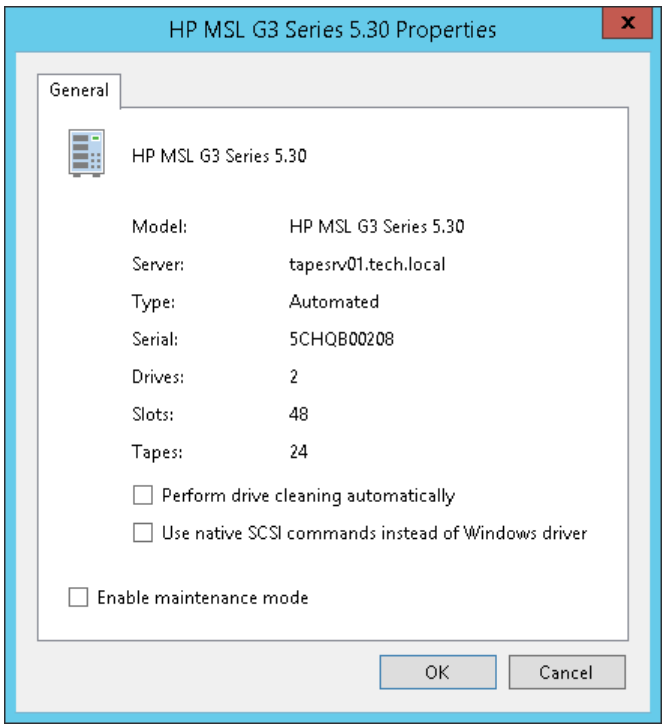

Rescanning Tape Devices

You can rescan a selected tape library or standalone tape drive to update its network status.

To rescan a tape device:

1. Open the **Tape Infrastructure** view.
2. Select the **Libraries** node in the inventory pane.
3. Select a necessary item in the working area and click **Rescan Library** on the ribbon. You can also right-click the necessary device in the working area and select **Rescan**.

Renaming Tape Devices

The Veeam backup console shows all tape libraries and standalone tape drives under the list of libraries. By default, the connected devices are shown with their production model names. For clearer visibility, you can rename them.

To rename a tape device:

1. Open the **Tape Infrastructure** view.
2. Select the **Libraries** node in the inventory pane.
3. Select a necessary item in the working area and click **Rename Library** on the ribbon. You can also right-click the necessary device in the working area and select **Rename**.

Removing Tape Devices

If you want to stop using a tape library or standalone tape drive, you can remove it from your Veeam backup console.

Mind the following limitations:

- Only devices in the offline status can be removed. To remove a tape device, you need to physically disconnect it from the tape server first. Otherwise, the tape device will be discovered during the next rescan and reappear in the **Libraries** list. Alternatively, you can remove the tape server to which the tape device is connected.
- You cannot remove the tape device that is used in a media pool. Remove the device from all media pools first.

To remove a tape device:

1. Disconnect the tape device or the tape server.
2. Open the **Tape Infrastructure** view.
3. Select the **Libraries** node in the inventory pane.
4. Select a necessary item in the working area and click **Remove Library** on the ribbon. You can also right-click the necessary device in the working area and select **Remove Library**.
5. In the displayed dialog box, click **OK** to confirm deletion.

> **TIP**
>
> If you remove the tape server, the tape devices remain connected to it. You can connect the tape server back to your console. In this case, Veeam Backup & Replication will recognize all tape devices that were previously added to it with tapes and media pools (in case they remain in the Veeam database).
>
> You can also connect the tape server to another Veeam backup server. Veeam Backup & Replication will recognize the tape devices automatically, however you will need to create the media pools anew.

Replacing Tape Devices

You may want to stop using a tape library or a standalone tape drive, for example, because of breakdown or low efficiency, and replace it with a new tape device. To replace the device seamlessly and continue using the tapes within the same media pool, follow the instructions below.

To replace a tape device:

1. Connect the new tape device to the tape server. For more information, see Connecting Tape Devices.
2. Add the tape device to the media pool to which the old tape device was added. Open the media pool settings and go to the **Tapes** step of the wizard. For more information, see Modifying Media Pools. If the old tape device was added to several media pools, add the new tape device to the same media pools.
3. Remove the old tape device from the media pools. If the old tape device was added to several media pools, remove the tape device from all of these media pools.
4. Disconnect the old tape device from the tape server.
5. Remove the old tape device from Veeam Backup & Replication. For more information, see Removing Tape Devices.
6. Make sure that you do not remove the offline tapes from Veeam Backup & Replication.
7. Offload the tapes from the old tape device and load them to the new tape device.
8. Rescan the tape device.

To replace a standalone tape drive:

1. Offload the tape from the old standalone tape drive if it was loaded.
2. Add a new standalone tape drive.
3. Disconnect the old tape drive.
4. Add the tape device to the media pool to which the old tape device was added. Open the media pool settings and go to the **Tapes** step of the wizard. For more information, see Modifying Media Pools. If the old tape device was added to several media pools, add the new tape device to the same media pools.
5. Remove the old tape device from the media pools. For more information, see Removing Tape Devices.

   Veeam Backup & Replication prompts to select a new tape drive.
6. Select the new tape drive.

   The tapes in the media pool will now be associated with the new tape drive.
7. Repeat the procedure for each media pool where the old tape drive was used.

As a result, the tapes remain in the same media pools and the data is intact.

Loading Tapes

When you need to add free tapes, or if you need to restore data from offline tapes, you need to load these tapes into your tape device.

When you load tapes into the tape device, Veeam does not recognize them automatically. To start using tapes, you need to introduce them to the Veeam Backup & Replication by running the importing procedure (**Tape Infrastructure** view > **Libraries** node > right-click *LibraryName* > **Import tapes**). The importing procedure conveys the tapes from the mail slot to the library working slots, scans them and sends information about them to Veeam backup database. After importing, Veeam Backup & Replication can use the tapes to read or write data.

Importing runs differently for tapes that are registered in the local Veeam backup database and those that are not. Choose an appropriate procedure for the following tapes:

- Loading Empty Tapes
- Loading Tapes Written on This Backup Server
- Loading Tapes Written on Another Veeam Server

Related Topics

Loading Tapes Written with 3rd-Party Backup Solution

Loading Empty Tapes

To load empty tapes, perform the following steps:

1. Load the tapes into the mail slot or directly into the tape library magazine.
2. If you loaded the tapes to the mail slot, import the tapes. To do this, open the **Tape Infrastructure** view, expand the **Libraries** node and select the library to which you have loaded the tapes. Click **Import Tapes** on the ribbon. You can also right-click the necessary library in the working area and select **Import Tapes**.
3. If you loaded tapes into the tape library magazine, rescan the tape library. To do this, click **Rescan Library** on the ribbon. You can also right-click the necessary library in the working area and select **Rescan**.
4. The tapes will appear in the **Unrecognized** media pool.
5. Right-click the tapes and select **Inventory tape.**
6. The tapes will be moved to the **Free** media pool.

If you want to refill the free tape resources in your tape device, you can overwrite tapes containing outdated Veeam archives. For more information, see Importing Tapes Written on This Backup Server.

> **TIP**
>
> To understand which of the tapes can be overwritten, you can check their expiration date. In the **Tape Infrastructure** view, open the **Media Pools** node and select the media pool that you want to replenish. Check the **Expires in** field in the working area. You can sort the tapes by the expiration date and get the list of tapes that can be reused.

Loading Tapes Written on This Backup Server

If the tapes that you load were written on the same Veeam backup server, the Veeam backup database has information about them. The procedure is different for tape libraries depending on whether they support barcodes or not.

Loading Tapes with Barcodes

To load tapes with barcodes written in the same tape library:

1. Load the tapes into the mail slot or directly into the tape library magazine.
2. If you loaded the tapes to the mail slot, import the tapes. To do this, open the **Tape Infrastructure** view, expand the **Libraries** node and select the library to which you have loaded the tapes. Click **Import Tapes** on the ribbon. You can also right-click the necessary library in the working area and select **Import Tapes**.
3. If you loaded tapes into the tape library magazine, rescan the tape library. To do this, click **Rescan Library** on the ribbon. You can also right-click the necessary library in the working area and select **Rescan**.
4. Veeam Backup & Replication will automatically move the tapes to their parent media pools. You can start working with them.

Moving Tapes with Barcodes Between Libraries

We strongly recommend that you do not move tapes with barcodes written in one library to another library, as that can result in various errors and return messages about barcode duplicates.

If you cannot avoid moving tapes with barcodes between the libraries, make sure you perform all the following steps:

1. Unload the tapes from the mail slot or from the magazine of the tape library.
2. Rescan the library that you remove the tapes from. To do this, open the **Tape Infrastructure** view, expand the **Libraries** node and select the library that you remove the tapes from. Click **Rescan Library** on the ribbon. You can also right-click the necessary library in the working area and select **Rescan**.
3. Remove the tapes from the catalog. To do this, navigate to the list of tapes either under the **Media Pools** node or under the **Libraries** > *LibraryName* node > **Media** > **Offline** node. Select offline tapes you want to remove from the catalog and click **Remove from Catalog** on the ribbon. Alternatively, you can right-click selected tapes and choose **Remove from Catalog** from the shortcut menu. In the opened dialog box, click **Yes** to confirm removal.
4. Load the tapes into the mail slot or directly into the magazine of the tape library that you move the tapes to.
5. Catalog the loaded tapes in the new library. To do this, navigate to the list of tapes either under the **Media Pools** node or under the **Libraries** > *LibraryName* node > **Media** node. Select the necessary tapes in the list and click **Catalog** on the ribbon. Alternatively, you can right-click the selected tapes and choose **Catalog Tape**.
6. Veeam Backup & Replication adds the tapes to the new library. You can start working with them.

Loading Tapes without Barcodes

To load tapes without barcodes:

1. Load the tapes into the mail slot or directly into the tape library magazine.

2. If you loaded the tapes to the mail slot, import the tapes. To do this, open the **Tape Infrastructure** view, expand the **Libraries** node and select the library to which you have loaded the tapes. Click **Import Tapes** on the ribbon. You can also right-click the necessary library in the working area and select **Import Tapes**.
3. If you loaded tapes into the tape library magazine, rescan the tape library. To do this, click **Rescan Library** on the ribbon. You can also right-click the necessary library in the working area and select **Rescan**.
4. The tapes will appear in the **Unrecognized** media pool.
5. Right-click the tapes and select **Inventory tape**.
6. Veeam Backup & Replication will move the tapes to their parent media pools. You can start working with them.

Note that if the tapes are expired, they will be overwritten by tape rotation scheme set for the tapes' parent media pool.

Loading Tapes Written on Another Veeam Server

Veeam Backup & Replication supports restoring data from tapes that were recorded on another Veeam backup server. To read data from such tapes, Veeam Backup & Replication must first catalog the tapes and store the information about them to the Veeam backup database.

To load tapes written on another Veeam backup server, follow the next steps:

1. Load the tapes into the mail slot or directly into the tape library magazine.
2. If you loaded the tapes to the mail slot, import the tapes. To do this, open the **Tape Infrastructure** view, expand the **Libraries** node and select the library to which you have loaded the tapes. Click **Import Tapes** on the ribbon. You can also right-click the necessary library in the working area and select **Import Tapes**.
3. If you loaded tapes into the tape library magazine, rescan the tape library. To do this, click **Rescan Library** on the ribbon. You can also right-click the necessary library in the working area and select **Rescan**.
4. The tapes will appear in the **Unrecognized** media pool.
5. Right-click the tapes you want to import and select **Catalog Tape**.
6. The tapes will be moved to the **Imported** media pool. The **Imported** media pool will be created automatically. When the tapes are imported, you can view and restore data written to them.

Loading Tapes Written with 3rd-Party Backup Solution

Veeam Backup & Replication does not support backups written with another backup solution even if they are written in the supported format.

> **IMPORTANT**
>
> You cannot restore data written with 3rd-party tape recording solutions.

If the tapes written with another backup solution contain valuable data, do not run any procedures against them, for example, inventory or catalog. Veeam Backup & Replication recognizes such tapes as writable and may use them for Veeam tape jobs.

If you do not need the data on these tapes, you can erase them and use for writing data with Veeam Backup & Replication.

To erase and re-use tapes containing non-Veeam data:

1. Load the tapes into the mail slot or directly into the tape library magazine.
2. If you loaded the tapes to the mail slot, import the tapes. To do this, open the **Tape Infrastructure** view, expand the **Libraries** node and select the library to which you have loaded the tapes. Click **Import Tapes** on the ribbon. You can also right-click the necessary library in the working area and select **Import Tapes**.
3. If you loaded tapes into the tape library magazine, rescan the tape library. To do this, click **Rescan Library** on the ribbon. You can also right-click the necessary library in the working area and select **Rescan**.
4. The tapes will appear in the **Unrecognized** media pool.
5. Select tapes you want to erase and click **Erase** on the ribbon. Alternatively, you can right-click selected tapes and choose **Erase Tape**. Choose **Short erase** and click **OK**. For more information, see Erasing Tapes.
6. The tapes will be moved to the **Free** media pool.

Inventorying Tapes

When you load new tapes in your tape device for the first time, these tapes are presented to Veeam Backup & Replication as **Unrecognized**. To identify unrecognized tapes, you need to run tape inventory job against them.

Tape inventory job is a relatively fast process of reading metadata written on tape with the aim of detecting name of the media set and the sequence number for the tape. Tape inventory jobs helps Veeam Backup & Replication identify empty tapes and detect non-empty tapes belonging to specific media set.

You can perform tape inventory for a whole tape library or run the job against selected tapes only.

Inventorying Tape Libraries

To inventory a whole tape library:

1. Open the **Tape Infrastructure** view.
2. Expand the **Libraries** node and select the library you want to inventory.
3. Click **Inventory Library** on the ribbon. Alternatively, you can right-click the tape library and select **Inventory Library** from the shortcut menu. Veeam Backup & Replication will perform inventory for all online tapes in the library.

Inventorying Tapes

To inventory selected tapes:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under the **Media Pools** or under the **Libraries >** *LibraryName* node **> Media**.
3. Select the necessary tapes in the list and click **Inventory** on the ribbon. Alternatively, you can right-click the selected tapes and choose **Inventory Tape**.

The inventory log will display job session results. To access the inventory session details, you can open the **History** view and locate the necessary session under the **Jobs** > **Tape** node.

Inventory Results

After the inventory, Veeam Backup & Replication places tapes in predefined pools:

- Empty tapes are moved to the **Free** pool. You can use these tapes for archiving backups and files
- Tapes that contain data written on another Veeam backup server remain in the **Unrecognized** pool. Veeam Backup & Replication displays the detected media set name and sequence number for these tapes. If you want to restore data from a specific media set, you need to run the tape catalog job for all tapes in this media set.

Cataloging Tapes

Cataloging tapes is required for tapes that contain some data but are not registered in the Veeam database. For example, cataloging is required for tapes that were written on another Veeam backup server. After the cataloging procedure, you can restore data from the tapes.

When a catalog job is performed, Veeam Backup & Replication first performs tape inventory, reads the information about backup contents on tape from the tape catalog information, scans tape contents and updates the database with details of new detected backup sets.

You can perform tape catalog job for a whole tape library or for selected tapes.

Cataloging Tape Libraries

To catalog a whole tape library:

1. Open the **Tape Infrastructure** view.
2. Expand the **Libraries** node and select the library you want to catalog.
3. Click **Catalog Library** on the ribbon. Alternatively, you can right-click the tape library and select **Catalog Library**.

Cataloging Tapes

> **NOTE**
>
> We recommend to catalog in one session all the tapes written within the same backup set. Otherwise (if cataloging not all the tapes, or cataloging them in different sessions), not all the data may be correctly imported into the database and Veeam Backup & Replication may require cataloging the same tapes again.

To catalog selected tapes:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under the **Media Pools** or under the **Libraries** > *LibraryName* node > **Media**.
3. Select the necessary tapes in the list and click **Catalog** on the ribbon. Alternatively, you can right-click the selected tapes and choose **Catalog Tape.**

The catalog log will display job session results. To access the catalog session details, you can open the **History** view and locate the necessary session under the **Jobs** > **Tape** node.

Cataloging Tips

Rescanning of tapes during the catalog job may take a lot of time. To speed up the cataloging process, you might do one of the following:

- First, run inventory job to identify tape media sets and decide which tapes you want to catalog. Run the catalog job only against tapes in the necessary media set.
- If you work with a tape library, you can run the catalog job against the whole media set at once.
- If you work with a standalone drive, start cataloging from the last tape in the media set (as this tape usually stores on tape catalog information).

System Requirements for Large Number of Files on Tape

If the tape backup contains a large quantity of files, for example, more than 1 000 000 files in 1 000 folders, you must provide the following system resources:

- Backup server: 1.3 GB RAM per each 1 000 000 files
- Tape server: 800 MB RAM per each 1 000 000 files.

Automated Drive Cleaning

You can instruct Veeam Backup & Replication to automatically clean the tape library drives.

Assigning the automated cleaning to Veeam Backup & Replication prevents possible overlapping of cleaning tasks and tape jobs. Such overlapping may cause tape jobs failures.

To instruct Veeam Backup & Replication to automatically clean the drives:

1. Open the **Tape Infrastructure** view.
2. Expand the **Libraries** node and select the needed library. Click **Properties** on the ribbon. You can also right-click the necessary library in the working area and select **Properties**.
3. In the **Properties** window, select the **Perform drive cleaning automatically** option.

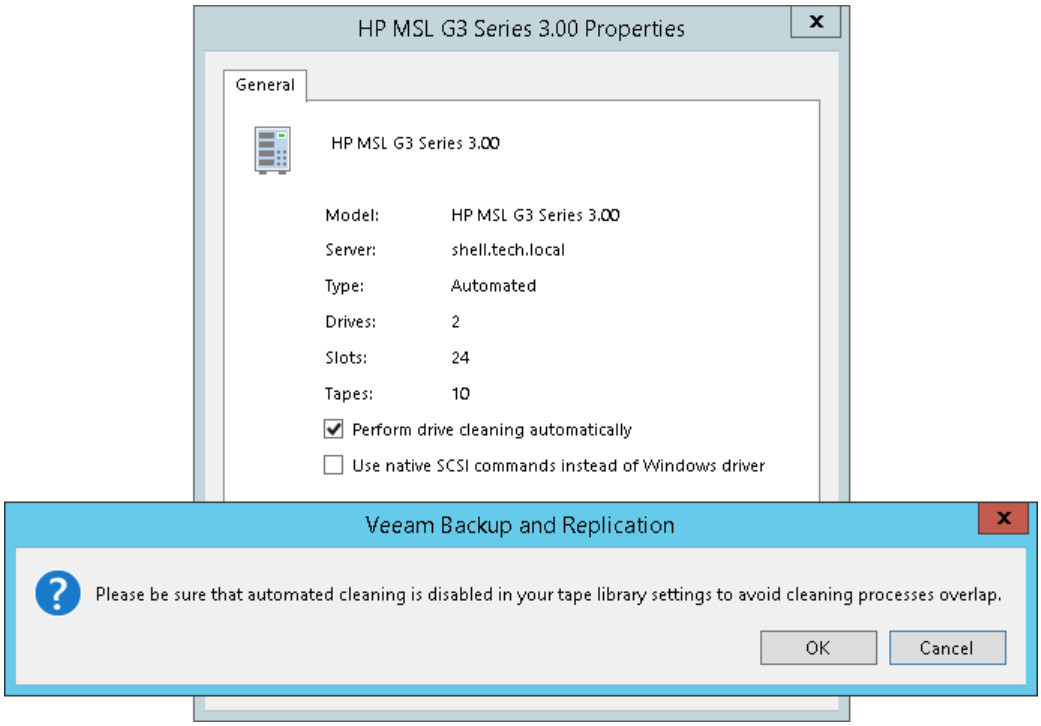

IMPORTANT

If you enable the automated drive cleaning option in Veeam Backup & Replication, make sure that you disabled the drive cleaning tasks on your tape library device.

How It Works

Veeam Backup & Replication cleans the drives at the beginning of backup to tape jobs or file to tape job run. The cleaning is not performed during other tape operations such as, for example, cataloging or export.

To clean the drives automatically, Veeam Backup & Replication performs the following actions:

1. The tape library alerts Veeam Backup & Replication on a drive that requires cleaning.
2. Veeam Backup & Replication waits for a tape job to start.
3. When the tape job locks necessary drives for writing data, Veeam Backup & Replication checks which of them requires cleaning.

4. Veeam Backup & Replication ejects the tape from the drive, inserts a cleaning tape and performs the cleaning.
5. Veeam Backup & Replication ejects the cleaning tape and inserts the tape that was reserved for the tape job.
6. The tape job writes the data on tape.

The cleaning process usually takes several minutes.

The cleaning tapes are located in the **Unrecognized** media pool. The worn-out cleaning tapes are moved to the **Retired** media pool automatically.

If a tape job locks multiple drives simultaneously for parallel processing, and one or more drives require cleaning, all drives wait until the cleaning is finished. After cleaning, all drives start writing simultaneously.

The automated drive cleaning does not affect creation of media sets.

Limitations for Automated Drive Cleaning

Mind the following limitations for automated drive cleaning:

- You cannot enable the automated drive cleaning on standalone tape drives.
- You cannot start the drive cleaning manually with Veeam Backup & Replication. The drive cleaning is fully automated.

Switching Libraries to Maintenance Mode

There can be situations where you need to lock a tape library or a standalone tape drive from being used by Veeam Backup & Replication, for example, to provide maintenance services to them.

Before switching the tape library to the maintenance mode, mind the following:

- You cannot enable the maintenance mode for the tape library if jobs that use it are running.
- After you enable the maintenance mode for the tape library, all jobs that use it will fail with an error when starting.

To switch a tape library to the maintenance mode:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tape libraries under **Libraries**.
3. Select a tape library you want to switch to the maintenance mode and click **Maintenance** on the ribbon. Alternatively, you can right-click the selected tape library and select **Maintenance** from the shortcut menu. Or you can open the tape library properties and select the **Maintenance mode** check box.

The tape library will stay in the maintenance mode until you manually disable it by performing the same actions like you did when enabling it.

After you disable the maintenance mode, Veeam Backup & Replication automatically starts a discovery job. However, it does not apply to standalone tape drives. To discover standalone drives, manually rescan the tape server, as described in Rescanning Tape Servers.

Tape Drives

Drives of each connected library or standalone tape device are displayed in the **Tape Infrastructure** view under the **Libraries** > *LibraryName* node > **Drives**.

To view the drives properties, select **Drives** under the needed library. View the list and a short description of the drives in the working area. To view detailed information on a drive or to change its settings, click **Properties** on the ribbon. You can also right-click the necessary drive in the working area and select **Properties**.

To set block size for the drive manually, select a value from the **Block size** list in the drive properties.

> **NOTE**
>
> Veeam Backup & Replication uses uniform block size. Drive that is set manually to a particular block size may fail to read tape that was written with another block size. For more information on tape block size, see Supported Devices and Configuration.

To disable a drive, clear the **Enabled** option in the drive properties. To set the drive into working mode again, select the **Enable** option. You can disable a drive, for example, for maintenance. When you disable a drive, the library does not use this drive for read/write operations. If the library has multiple drives, other drives will be used for the read/write operations.

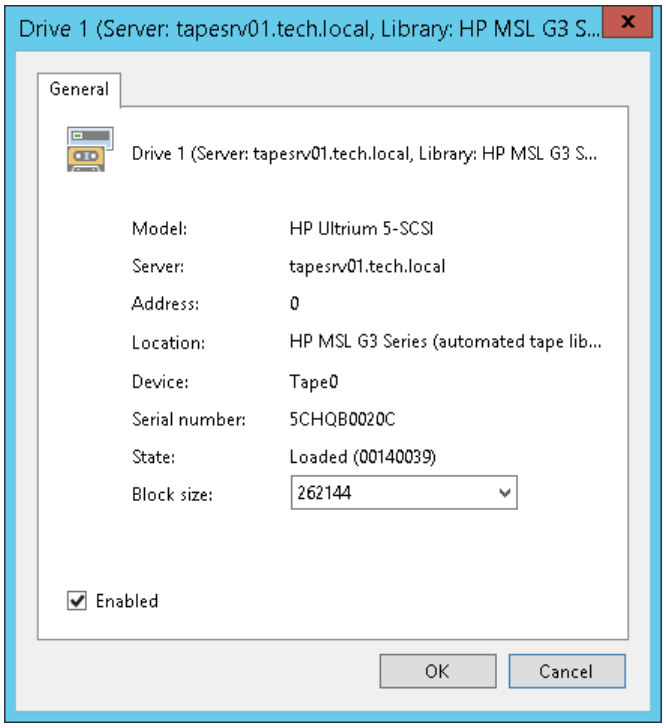

Tapes

All tapes managed by Veeam Backup & Replication belong to one of media pools. Generally, the new tapes are in the **Free** media pool. Tapes that were written by tape jobs stay in the media pools that are targets to these tape jobs.

If you offload tapes from you tape device, their status changes to Offline, but they stay visible in the console.

Tape media in Veeam Backup & Replication are displayed in the **Tape Infrastructure** view, either under the **Media Pools** or under the **Libraries >** *LibraryName* node **> Media**. You can work with both online and offline tapes:

- Tapes that are currently loaded to the tape device are available under the **Online** node.
- Tapes that have been unloaded from the tape device are shown under the **Offline** node.

> NOTE
>
> Before you start to use tapes loaded into the mail slot and tapes recently created in the virtual tape library (VTL), import them. To do that, open the **Tape Infrastructure** view, in the inventory pane click the **Libraries** node, select the required library, right-click it and select **Import tapes**.

All tapes are grouped to service or user-created media pools available under the **Tape Infrastructure** > **Media Pools** node. You can also view tapes under the **Tape Infrastructure** > **Libraries >** *LibraryName* node **> Media.**

Veeam Backup & Replication can use only online tapes for backup to tape and file to tape jobs.
If you work with a standalone tape drive and all its tapes are offline, Veeam Backup & Replication will display a message informing that you need to insert a tape into the drive. At the same time, Veeam Backup & Replication will hint what tape has been recently used for archiving.
You can insert any tape into the tape drive:

- If the tape you have inserted is registered in the Veeam backup database and the current media set can still be used, Veeam Backup & Replication will continue writing to this media set and append the new content to the content recently written on the tape.
- If the tape you have inserted is not registered in the Veeam backup database, it will be processed as follows:
    - A new empty tape will be placed to the **Free** media pool and used for writing data.
    - A tape containing any data written on another Veeam backup server or with another tape backup solution will be placed to the **Unrecognized** media pool. You need to catalog such tapes.

Viewing Tape Properties

You can view properties of each tape registered in Veeam database. The tapes may be online or offline. The **Properties** window shows location, capacity, name and other properties of the selected tape.

Tape Capacity

The capacity of the tape is identified by the tape device where the tape is located. You can also view the capacity of a tape in your tape device console (or in the Microsoft Windows Device Manager for tapes that are in drives).

Tape Name

The name of the tape is generated automatically during tape importing. If your tape device supports barcodes, Veeam Backup & Replication will use the barcodes to identify the tapes and display the barcode as the **Name** property of the tape. If a tape does not have a barcode, Veeam Backup & Replication will automatically create the tape name using the pattern '*Tape N*'.

> **IMPORTANT**
>
> "CLN" and "CLR" prefixes are reserved for cleaning tapes. Veeam Backup & Replication regards tapes with such prefixes as cleaning tapes.
>
> Make sure that the regular tapes intended for writing data do not have barcodes starting with "CLN" or "CLR". If they do, rename the tapes and inventory the tape library.

If necessary, you can change the name of a tape:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries** > *LibraryName* node > **Media** > **Online/Offline**.
3. Select the tape you want to rename and click **Properties** on the ribbon. Alternatively, you can right-click the tape and select **Properties**.
4. Change tape name and description as required.
5. Click **OK** to save changes.

> **NOTE**
>
> If you manage several tape libraries on the same Veeam backup server and use barcodes to identify tapes in these libraries, make sure that all barcodes are unique. If some tapes have identical barcodes, Veeam Backup & Replication may fail to manage them correctly.

Tape Usage Statistics

Tapes wear off gradually every time they are rewritten. When a tape reaches the maximum times of reuse (set by manufacturer), the tape becomes retired. Veeam Backup & Replication automatically moves retired tapes to the **Retired** media pool. For more information, see Service Media Pools.

Veeam Backup & Replication creates tape usage statistics for each tape based on read and write activity and the size of the tape. When Veeam Backup & Replication selects a tape for a job, it considers the tape usage statistics among other criteria. For a job, a tape with minimal usage is selected.

You can view the tape usage statistics in the **Wear** field of the **Tape Properties** window.

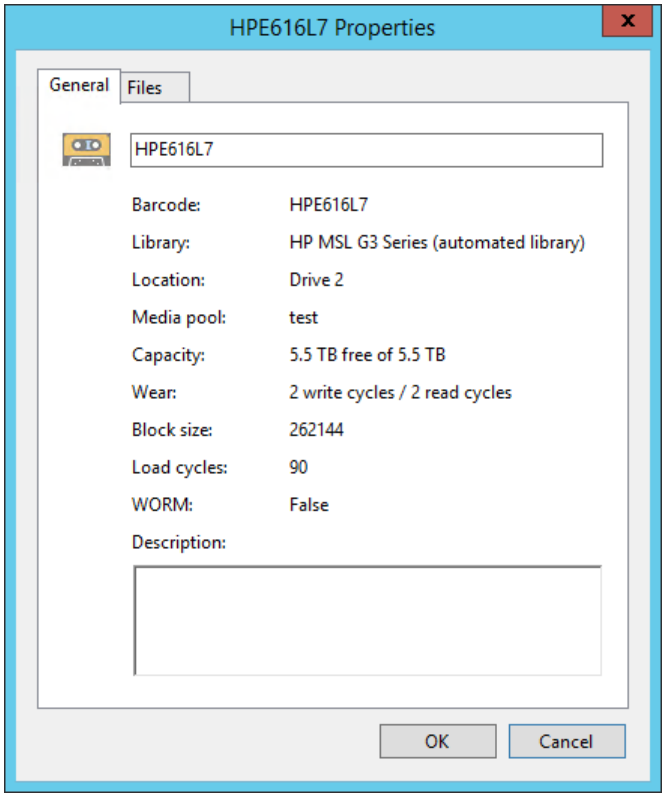

For cleaning tapes, you can view the remaining cleaning cycles in the **Wear** field of the **Tape Properties** window.

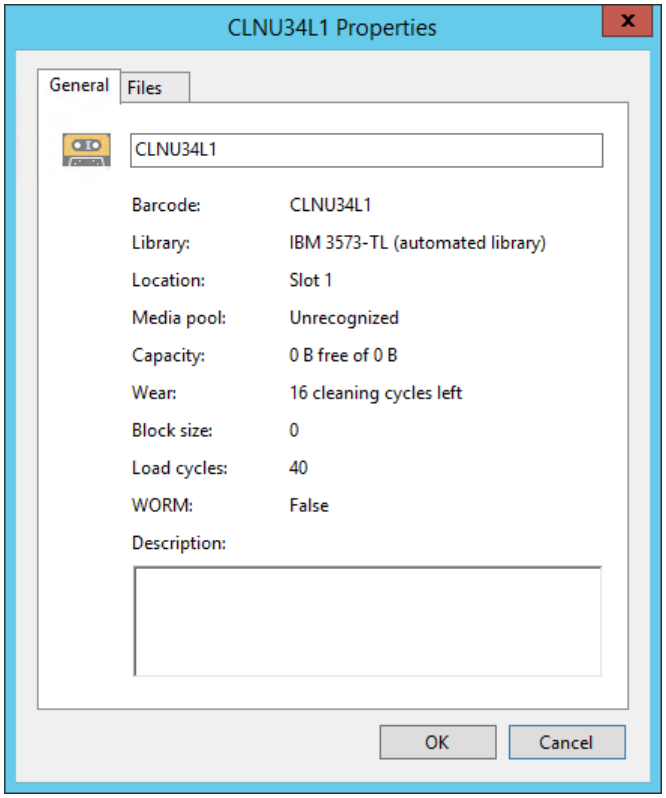

Viewing Data on Tapes

To view data contents that is currently stored on tape:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries** > *LibraryName* node > **Media** > **Online**.
3. Select the necessary tape and click **Properties** on the ribbon. Alternatively, you can right-click the tape and select **Properties**.
4. Open the **Files** tab.

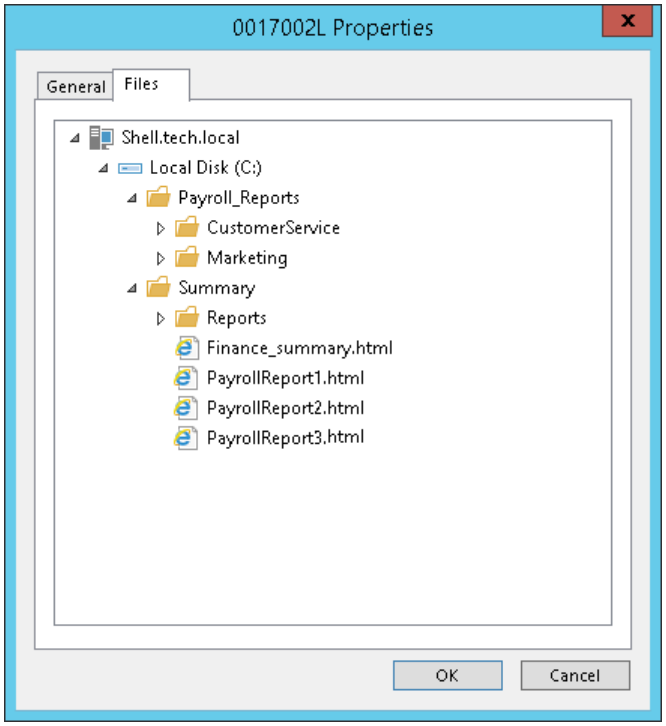

Removing Tapes from Catalog

If you do not want to store information about a tape and contents on this tape in Veeam Backup & Replication database, you can remove the tape from the catalog. For example, you can perform removal from catalog if a physical tape no longer exists, and the tape should no longer be displayed among media in the Veeam Backup & Replication console.

You can remove only tapes that are offline.

To remove one or more tapes from the catalog:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries >** *LibraryName* node **> Media > Offline**.
3. Select offline tapes you want to remove from the catalog and click **Remove from Catalog** on the ribbon. Alternatively, you can right-click selected tapes and select **Remove from catalog** from the shortcut menu.
4. In the opened dialog box, click **Yes** to confirm removal.

> **IMPORTANT**
>
> You cannot remove protected tapes from the catalog. To remove such tapes, you need to switch the protection off first.

Erasing Tapes

If you do not need the contents stored on tape, you can erase tapes. Veeam Backup & Replication supports two options for erasing data:

- **Short erase (fast)** – use this option to speed up the erase process. The short erase operation does not physically erase data written on the tape. It simply loads the tape to the drive and wipes the tape header. Note that short erase is not supported by some tape devices.
- **Long erase (slow)** – use this option to clear all data written to tape. The long erase operation loads the tape to the drive, rewinds the tape and physically erases all data written to the tape.

To erase tapes:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries >** *LibraryName* node **> Media > Online**.
3. Select tapes you want to erase and click **Erase** on the ribbon. Choose the type of erase and click **OK**. Alternatively, you can right-click selected tapes and select **Erase tape**. Next, choose how the tape should be erased and click **OK**.

> **IMPORTANT**
>
> You cannot erase protected tapes. To erase such tapes, you need to switch the protection off first.

Marking Tapes as Free

Instead of erasing tapes, you can mark tapes as free. During this operation, Veeam Backup & Replication deletes from backup and tape catalogs information about backup contents stored on tape. Data written to tape remains intact.

To mark tapes as free:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries >** *LibraryName* node **> Media > Online**.
3. Select tapes you want to mark as free and click **Mark as Free** on the ribbon.
   Alternatively, you can right-click selected tapes and select **Mark as free** from the shortcut menu.
4. In the displayed dialog box, click **Yes**.

After a tape is marked as free, Veeam Backup & Replication removes from the catalog information about contents on this tape.

If you want Veeam Backup & Replication to automatically reuse the tape for storing new backups, instead of just marking it as free, move it to the **Free** media pool, as described in Moving Tapes to Another Media Pool.

> **IMPORTANT**
>
> You cannot mark protected tapes as free. To mark such tapes as free, you need to switch the protection off first.
>
> You cannot mark WORM tapes as free either.

Ejecting Tapes

Ejecting moves the tape that is currently in drive to a library slot. The tape stays in the library and is online. Ejecting prevents tapes, for example, from being occasionally overwritten.

You can eject a tape, or eject a drive. Eject command is enabled only for tapes that are in drives or for drives with loaded tapes.

To eject a tape:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries >** *LibraryName* node > **Media > Online**.
3. Select a tape located in drive that you want to eject and click **Eject** on the ribbon. Alternatively, you can right-click selected tape and select **Eject tape**.

To eject a drive:

1. Open the **Tape Infrastructure** view.
2. Navigate to **Libraries >** *LibraryName* node > **Drives**.
3. Select a drive from where you want to eject the tape and click **Eject Tape** on the ribbon. Alternatively, you can right-click selected tape and select **Eject**.

> **TIP**
>
> You can instruct your tape jobs to eject tape from drive after the job completes. To do so, navigate to the **Options** step of backup to tape job or file to tape job and select the **Eject media upon job completion** check box.

Exporting Tapes

If you want to take the tapes out of the tape device, you need to export them first. Exporting ejects the tapes from the slots in the tape device and moves them to the I/E slot (or Mail slot). When they are in the I/E slot, you can take them from the library and, for example, carry to a storing location.

> **NOTE**
>
> Export is available only for the devices that support corresponding operations and include I/E slot.

To export a tape:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries >** *LibraryName* node **> Media > Online**.
3. Select a tape you want to export and click **Export** on the ribbon. Alternatively, you can right-click selected tape and select **Export** from the shortcut menu.

> **TIP**
>
> You can instruct your tape jobs to export tapes automatically. To do so, navigate to the **Options** step of backup to tape job or file to tape job and select the **Export current media set upon job completion** check box.

Verifying Tapes

Veeam Backup & Replication offers a simple, yet efficient means to verify that your tapes are correctly read and consequently you can restore data stored on them in case of emergency.

> **NOTE**
>
> Mind the following:
>
> - During the tape verification job, Veeam Backup & Replication only checks if the data on the tape can be read. It does not check the application consistency or infrastructure.
> - Verifying tapes in a virtual tape library makes no sense as in this case Veeam Backup & Replication will retrieve no check sums for data blocks.

Reasons for Tape Verification

There can be different reasons to verify your LTO tapes:

- You want to check tapes before copying their data to new tapes.
- LTO has great durability, but it is still limited, and after some time you need to replace the tape media with a new one.
- Compliance reasons.

How Tape Verification Works

Veeam Backup & Replication performs tape verification as follows:

1. Veeam Backup & Replication instructs the tape drive to sequentially read data blocks from the tape and compares their content with the Veeam Backup & Replication database: it checks offsets of backup files, the correspondence of the offsets with correct file headers and the length of full backup files.
2. If the data on the tape does not correspond to the data in the Veeam Backup & Replication database, the latter is reported as corrupted (in the session log or separate .CSV file).
3. The tape drive calculates the checksums for data blocks on the tape and compares them with the checksums already stored for these data blocks on the tape.
4. If the drive detects an error in some data block (checksums differ or the data block could not be read), Veeam Backup & Replication reports all files in the data block as corrupted (in the session log or separate .CSV file).

Prerequisites for Tape Verification

Before starting a tape verification job, check the following prerequisites:

- The tapes to verify must be online.
- The tapes to verify must be cataloged. Tapes in the **Unrecognized** media pool are not available for tape verification. For more information, see Cataloging Tapes.
- All tapes to verify must be decrypted. Encrypted tapes cannot be read. For more information on decryption, see the Veeam Backup & Replication User Guide, section *Decrypting Tapes with Password*.

Besides, mind that you cannot run tape verification for cleaning tapes and free tapes (tapes in the **Free** media pool).

Running Tape Verification

To start the tape verification job:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries** > *LibraryName* node > **Media** > **Online**.
3. Select a tape you want to verify and click **Verify Tape** on the ribbon. Alternatively, you can right-click the selected tape and select **Verify** from the shortcut menu.

If you want to simultaneously use several drives for tape verification, successively start separate tape verification jobs for different tapes.

When writing backup files to tapes, there can be situations where a file does not fit on a tape. Then it is divided into parts and written to several tapes, which are then considered dependent. When verifying tapes, Veeam Backup & Replication checks data on selected tapes only. It does not check data on dependent tapes if they are not selected for verification.

To get an email report about the results of the tape verification job processing, enable notifications as described in the Veeam Backup & Replication User Guide, section *Configuring Global Email Notification Settings*. If there are too many issues reported after verification, the job session log in Veeam Backup & Replication UI displays only first 10 files with errors. For a full list of read errors with detailed paths and tape IDs, open the detailed report in the CSV format stored in the folder where tape verification job logs are stored. For convenience, the path to this folder is specified in the job session log in UI and in the email report.

Copying Tapes

You can use the tape copy feature to migrate backups or files stored on previous generation tapes to newer ones.

You can also use this feature to make copies of tapes.

Before starting a tape copy job, check the following prerequisites:

- The tape server must have at least 1 GB of RAM per each tape drive used during a tape copy session.
- The tapes to copy must not be occupied by any other operations, such as backup, restore, cataloging. During the tape copy procedure, the tapes will be locked.
- The tapes to copy must be cataloged. Tapes in the **Unrecognized** media pool are not available for copying. For more information, see Cataloging Tapes.
- All tapes to copy and machine backups on them must be decrypted. Encrypted tapes cannot be copied. Encrypted backups are skipped from copying. For more information, see Add Optional Media Pool Settings.
- The source and target tape libraries must have at least one free tape drive each. If the tape copying is performed within one tape library, it must have at least two free drives.

> **NOTE**
>
> Due to hardware peculiarities, the speed of copying tapes is on average 15% lower than the speed of writing backups to tapes from the backup repository.

To copy content of a tape to another one, use the **Copy Tapes** wizard.

Step 1. Launch Copy Tapes Wizard

To launch the **Copy Tapes** wizard, do the following:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** > *MediaPoolName* or under **Libraries** > *LibraryName* node > **Media** > **Online**.

> **NOTE**
>
> The following tapes cannot be copied:
>
> - tapes in the **Unrecognized** media pool
> - free tapes (tapes in the **Free** media pool)
> - encrypted tapes
> - cleaning tapes

3. Select the necessary tape and click **Copy** on the ribbon. Alternatively, you can right-click the tape and select **Copy**.

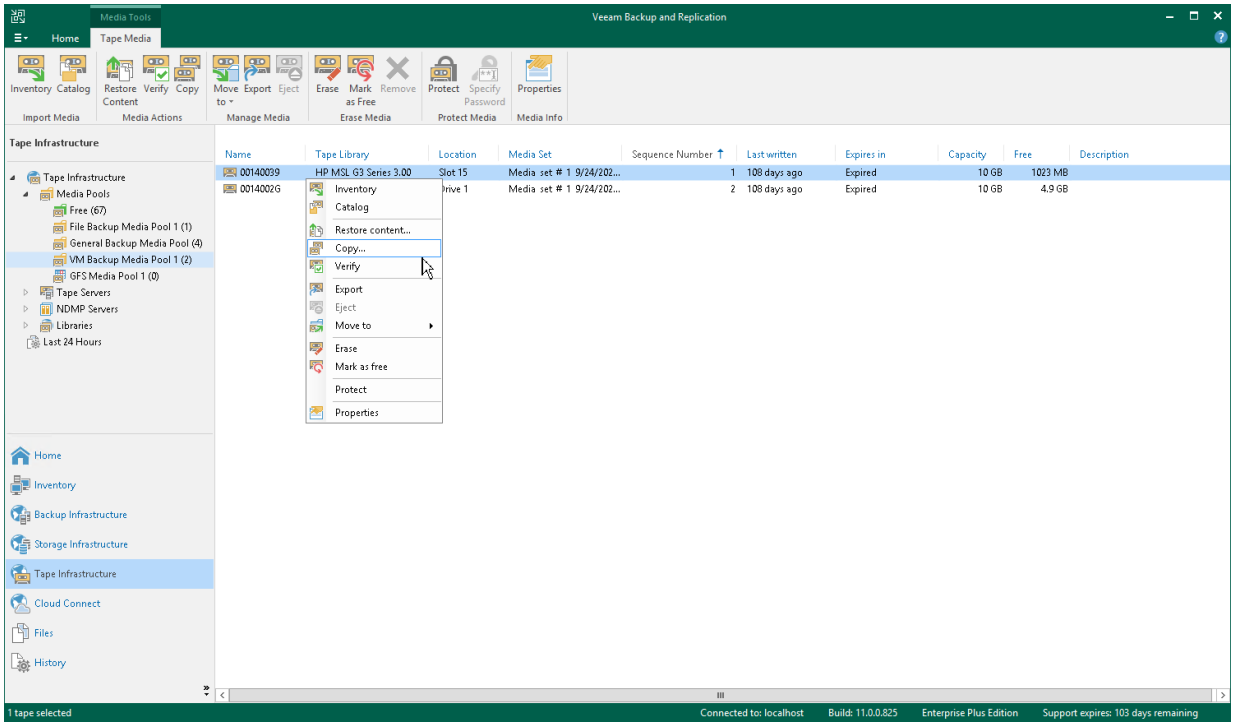

Step 2. Choose Source Tapes to Copy

At the **Source Tapes** step of the wizard, select source tapes to copy. The tape that you selected when launching the wizard is added to the list of tapes to copy by default. To add other tapes to copy, do the following:

1. Click **Add**. Veeam Backup & Replication opens the **Select Tapes** window listing all media pools and tapes in them available for copying.
2. In the **Select Tapes** window, navigate to the required media pool and select the tape you want to copy.

   Within one tape copy session, you can copy tapes of one type only: either tapes from a regular media pool, or tapes from a GFS media pool. You cannot combine them.

When writing backup files to tapes, there can be situations where a file does not fit on a tape. Then Veeam Backup & Replication divides it into parts and writes to several tapes, which are then considered dependent. Veeam Backup & Replication tracks this. When you are copying data from a tape that contains only one part of the backup file, Veeam Backup & Replication detects that. When you click **Next**, Veeam Backup & Replication informs you which tapes contain other parts of the file and prompts you to add these tapes to the list for copying complete content. You can agree to add these tapes to the job. Then Veeam Backup & Replication will copy these tapes marked as **Partially cloned** in the **Copy option** column only partially: it will copy only data required for cloning full content of the initially selected tapes in the consistent manner. Or you can refuse to add dependent tapes for copying. Then Veeam Backup & Replication will not copy files that are not complete from initially selected tapes. The **Used size** column shows the total size of the backups/files on each tape that will be copied. It does not include encrypted backups that are skipped when copying tapes or file parts that do not have accompanying file parts stored on dependent tapes.

To remove tapes from the copy list, select the tapes and click **Remove**.

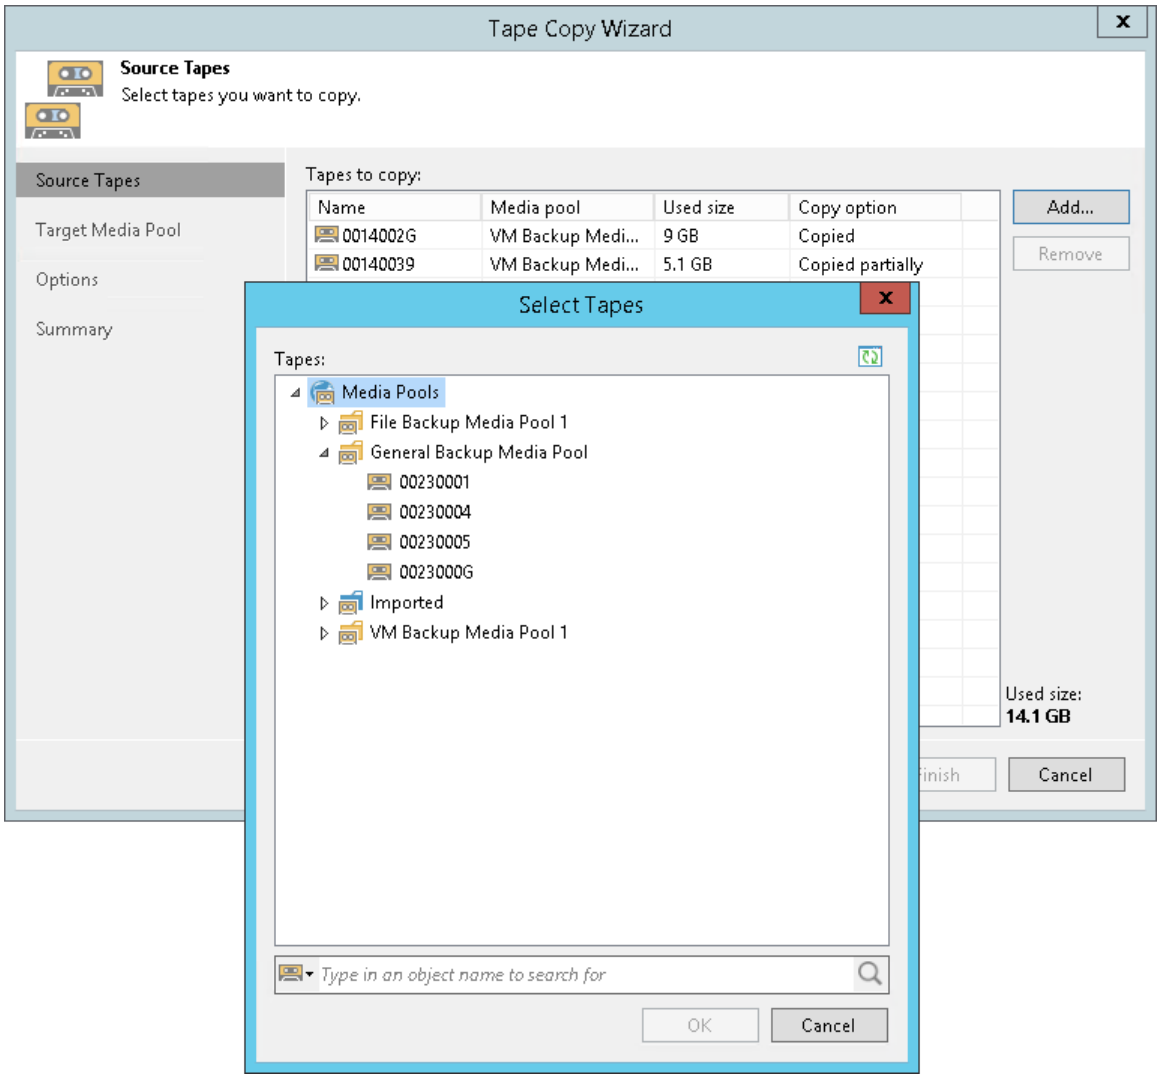

Step 3. Choose Target Media Pool

At the **Target Media Pool** step of the wizard, specify where you want to copy the source tapes to.

> **NOTE**
>
> Mind that the target media pool must be of the same type as the source media pool: they both must be either regular media pools, or GFS media pools.

1. From the **Target library** drop-down list, select the tape library that will store copies of source tapes.
2. From the **Target media pool** drop-down list, select a target media pool to copy source tapes to.

   If you have configured a media pool in advance, select it from the drop-down list. Otherwise, click **Add New** and add a media pool as described in Creating Media Pools.

   If the target media pool is the same as the source media pool, Veeam Backup & Replication prompts you to first close the existing media set, start a new media set and then copy the tape there.

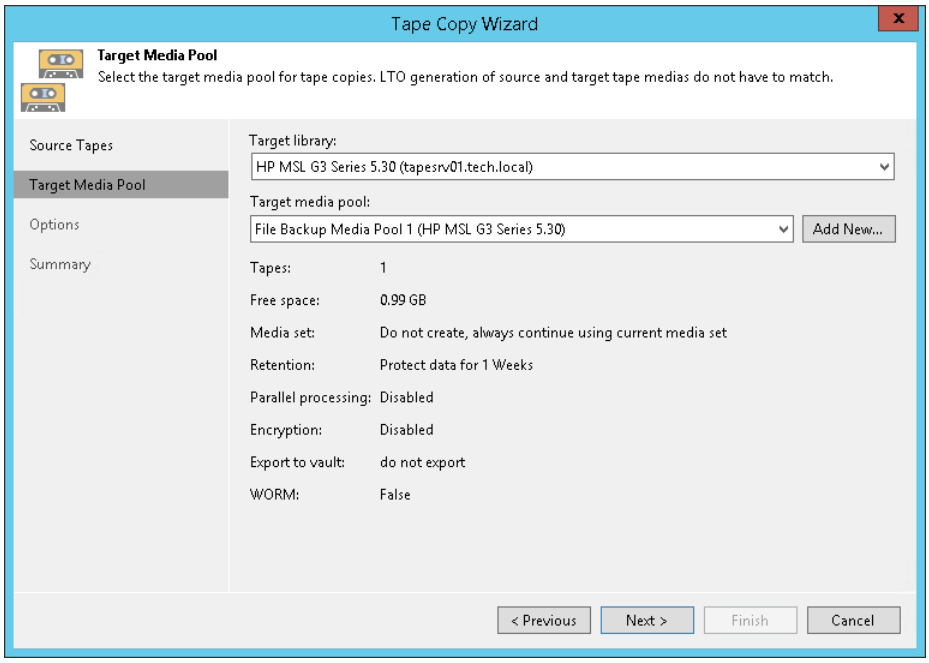

Step 4. Specify Tape Copy Options

At the **Options** step of the wizard, specify tape copy and tape export options.

- If you want to copy retention settings of the source tapes to the target tapes, select the **Preserve source media retention policy** check box.

  If you are copying several tapes with different retention settings, Veeam Backup & Replication will set retention policy for the target tapes to the maximum retention among the source tapes.

- If you want to enable hardware compression while writing new tapes, select the **Enable hardware compression for target tapes** check box.

- If you want to export source media, target media, or both after the tape copying finishes, change the default **Do not export tapes** option in the **Export options** section to the one that most suits your case. For details on tape export, see Exporting Tapes.

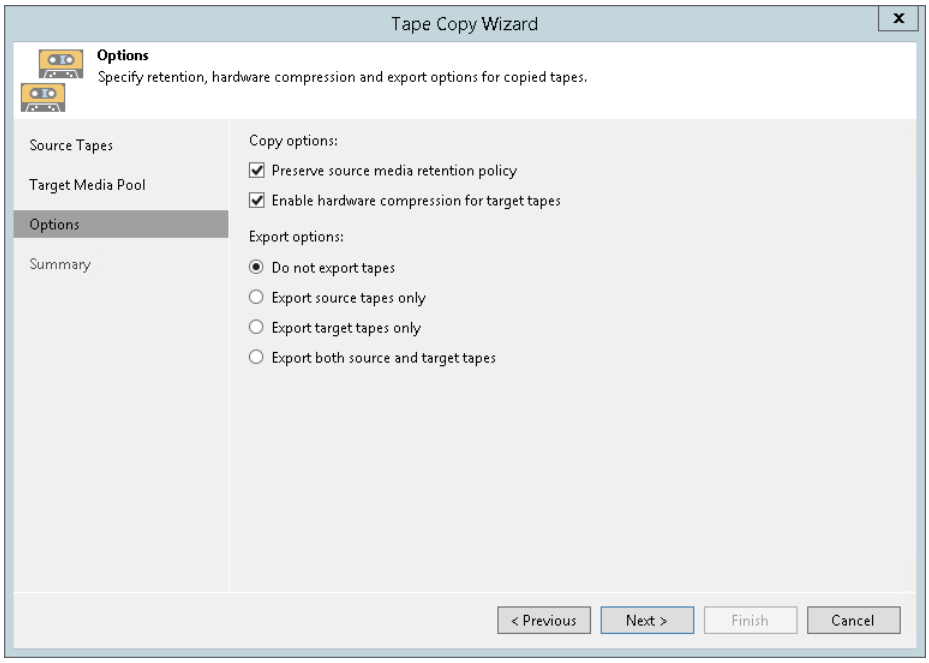

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review the configuration and click **Finish**.

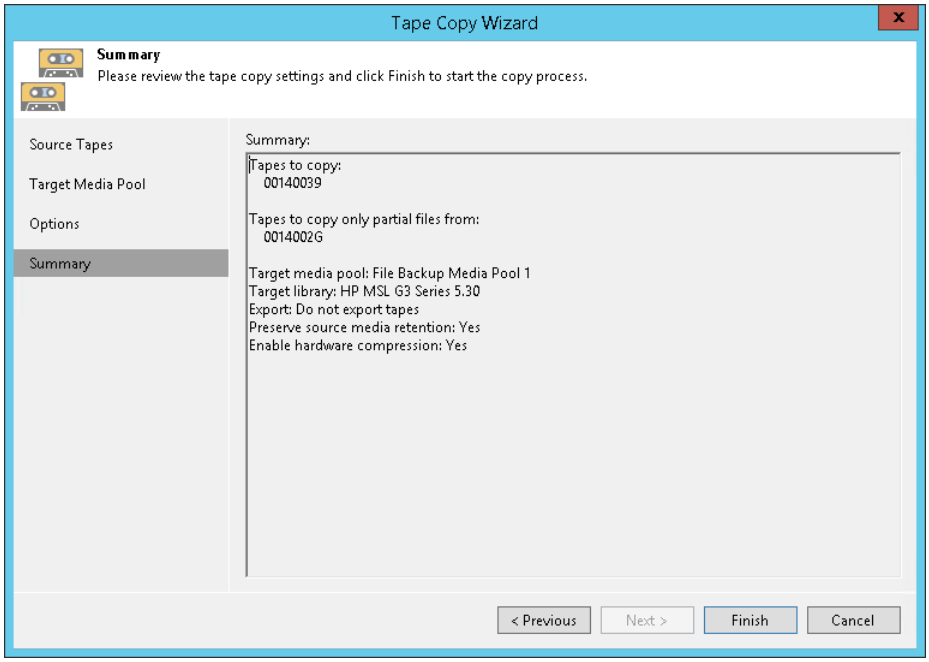

After that, the tape copy job will start. You can track its progress in the displayed **Tape Copy** window.

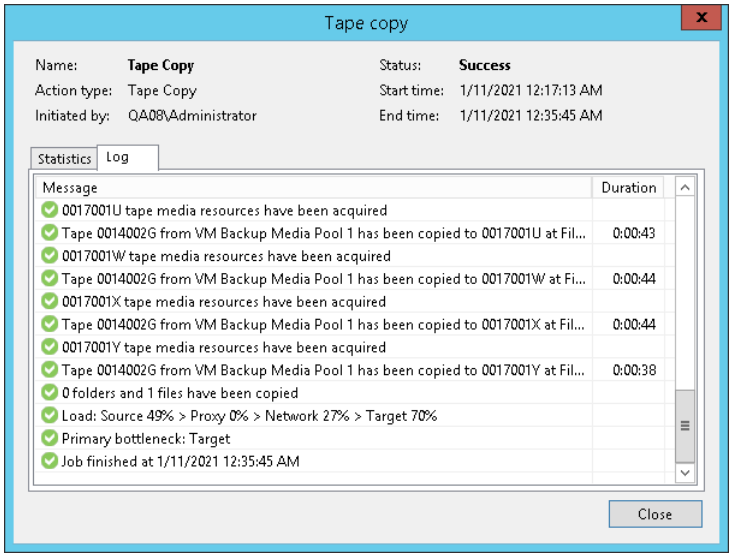

To get an email report about the results of the tape copy job processing, enable notifications as described in the Veeam Backup & Replication User Guide, section *Configuring Global Email Notification Settings*. If there are too many issues reported after the tape copy job, the job session log in Veeam Backup & Replication UI displays only first 10 files that failed to be copied. For a full list of errors with detailed paths and tape IDs, open the detailed report in the CSV format stored in the folder where tape verification job logs are stored. For convenience, the path to this folder is specified in the job session log in UI and in the email report.

Restoring Entire Tapes

A special case of file restore from tape described in Restoring Files from Tape is restoring content of entire tapes. That can be useful in various situations. For example:

- If a media set is damaged, you can use the entire tape restore to recover all available data stored on tapes in this media set.
- You want to check what content is stored on a particular tape.
- Audit requires checking the content of the tapes.

Before restoring tapes, mind the following:

- Restore of the content of entire tapes can be performed to a non-original location only.
- NDMP data is skipped from restore.
- The tapes to restore must be online.
- The tapes to restore must be cataloged. Tapes in the **Unrecognized** media pool are not available for restore. For more information, see Cataloging Tapes.
- Backup files on tapes to restore must be decrypted. Encrypted backup files are skipped from restore. For more information, see Add Optional Media Pool Settings.
- When restoring entire tapes, Veeam Backup & Replication creates a separate folder for each host whose files are backed up on the tapes.
- The entire restore tape job preserves the folder structure and restores all file and folder versions as is. If the tape stores several file versions of the same file or folder, older versions get date postfixes in the names.
- If some files on the tapes to restore have dependent file parts on other tapes, Veeam Backup & Replication will require these tapes to restore the full content of the selected tapes in the consistent manner.
- If you are restoring data from tapes that were previously copied, as described in Copying Tapes, Veeam Backup & Replication will use the most recent online copy as a source for restore.

For information on how to restore entire tapes, see Restoring Files from Tape.

Tapes Availability and Write-Protection

Before archiving data to tape media, Veeam Backup & Replication detects available tapes in the target media pool. Unavailable tapes are filtered out and are not used for writing data. A tape may be unavailable for one of the following reasons:

- The tape cannot be used according to the retention policy set for the media pool.
- The tape is offline.
- The tape has the write-protect switch set.

If your tape is write-protected and you want to write data to this tape, you must eject the tape from the drive and drag the write protection switch off. After you insert the tape back to the drive, you must inventory the tape to mark it as writable in the Veeam Backup & Replication database.

WORM Tapes

Veeam Backup & Replication supports WORM (Write Once Read Many) tapes for archiving backups to tape.

WORM Media Pools

As the WORM tapes cannot be overwritten, you need to treat them differently when using with Veeam Backup & Replication. These tapes cannot have retention period, and for this reason you cannot place them to the same media pool with the overwritable tapes.

For WORM tapes, you need to create separate media pools. In these media pools, only WORM tapes are used and they are not mixed with other tapes. The following media pools support the WORM media:

- WORM media pools
- WORM GFS media pools.

In WORM media pools, retention settings are disabled. For more information, see Media Pools and GFS Media Pools.

You can append new data to the same WORM tape cartridge until it is full as long as the media set containing this media is open. For more information about cases when Veeam Backup & Replication closes the current media set and starts a new one, see Media Sets.

Identifying WORM Tapes

Veeam Backup & Replication identifies WORM tapes and marks them as WORM in the Veeam database. WORM tapes are never mixed with standard tapes.

To identify WORM and standard tapes, run the inventory job against them. For more information, see Inventorying Tapes.

Requirements for WORM Tapes

Tape libraries must support reading/writing data on WORM tape cartridge memory chip.

Tape Barcode Labels

Library tapes can be marked with barcode labels that simplify identification of tapes and their types. Although the main source of information about the tape for Veeam Backup & Replication is a cartridge memory chip (LTO-CM), it can be read only when the media is inserted into the drive. Thus, correct labeling of tapes turns very useful for management of tape media.

Each barcode label used for marking LTO tapes supported by Veeam Backup & Replication consists of eight characters: a six-character volume serial ID and a two-character media ID. The media ID defines the tape type. The labels use the uppercase letters A through Z and the numbers 0 through 9, for example: `ABC123L8`, where `ABC123` is the volume serial ID, `L8` is the media ID indicating LTO of generation 8.

LTO data tapes used in devices supported by Veeam Backup & Replication may have the following media ID:

Media ID	Description
L3	Generation 3
L4	Generation 4
L5	Generation 5
L6	Generation 6
L7	Generation 7
L8	Generation 8
L9	Generation 9
LT	Generation 3 WORM
LU	Generation 4 WORM
LV	Generation 5 WORM
LW	Generation 6 WORM
LX	Generation 7 WORM
LY	Generation 8 WORM
LZ	Generation 9 WORM

Media ID	Description
CU	Universal cleaning

Cleaning and diagnostic tapes use unique labels to distinguish them from data tapes. The first three alphanumeric characters in the volume ID indicate the tape type:

- Cleaning tapes use `CLNxxx` as volume ID plus the cleaning-specific media ID, for example: `CLN001CU`, where `CLN001` is the cleaning media ID, `CU` is the indicator of the cleaning media.
- Diagnostic tapes use `DG[space]xxx` as diagnostic tape serial ID plus the media ID, for example: `DG 001L8`, where `DG 001` is the diagnostic tape serial ID, `L8` is the media ID.

NOTE

If tapes have incorrect barcode labels, for example, labels on ordinary (non-WORM) tapes indicate that they are of WORM or cleaning type, Veeam Backup & Replication may incorrectly identify such tapes.

Relabeling Tapes

To change a barcode label on a tape, you must correctly perform tape relabeling. To do that:

1. Physically remove the tape from the library.
2. From the Veeam Backup & Replication console, rescan the library, as described in Rescanning Tape Devices. Alternatively, you can wait for 3 minutes for the scheduled rescan to run.

   After the library rescan, the tape becomes offline. That is indicated by a black and white icon for this tape.
3. Remove the tape from Veeam Backup & Replication, as described in Removing Tapes from Catalog.
4. Physically change the label on the tape.
5. Insert the tape into the IO slot or into the tape library magazine and give the library a minute to discover the changes.
6. Import the tapes or rescan the library:
   - If you loaded the tapes into the IO slot, import the tapes. To do this, open the **Tape Infrastructure** view, expand the **Libraries** node and select the library into which you have loaded the tapes. Click **Import Tapes** on the ribbon. Alternatively, you can right-click the necessary library in the working area and select **Import Tapes**.
   - If you loaded tapes into the tape library magazine, rescan the tape library. To do this, click **Rescan Library** on the ribbon. Alternatively, you can right-click the necessary library in the working area and select **Rescan**.
7. [Optionally] If you need data written to the inserted tape, catalog the appeared tape.

Media Pools

Media pools are logical containers created in Veeam Backup & Replication to organize and administer tapes.

Veeam Backup & Replication has the following types of media pools:

- Service Media Pools are created automatically and are used for automatic tape administration.
- Media Pools are created by user and contain tapes for backups.
- GFS Media Pools are created by user and contain tapes for GFS backups.

All Veeam media pools are global. They span tapes that belong to multiple libraries. In a global media pool, you can move tapes between the libraries without need to catalog them.

> **NOTE**
>
> Mind the hardware encryption. If the tapes are encrypted by native hardware means of a tape library, you cannot read them at tape device with another encryption standard or with hardware encryption turned off.

Placing Tapes to Media Pools

Veeam Backup & Replication works only with tapes that are placed to a media pool. To introduce tapes to Veeam Backup & Replication, you need to run importing procedure against all new tapes you load to your tape devices. Importing registers tapes in the Veeam Backup database after which Veeam Backup & Replication places tapes to one of the media pools. For more information, see Loading Tapes.

Media Sets

Media set is a set of tapes used for continuously writing backup data. Media set is one of the parameters in media pool configuration.

A new media set always starts with a free tape. Within one media set, the new data block is appended to a previous one on a tape. Veeam Backup & Replication stores information about all tapes that belong to each media set. You can view the list of names or barcodes of tapes that are associated with a particular media set.

Media sets are used to explicitly distinguish data recorded to particular sets of tapes. You can choose between the following configuration options:

- **Always continuing one media set**. Use this option if you do not need to split your tape archives into separate sets of tapes.
- **Creating a new media set for each backup session**. In this case, Veeam Backup & Replication will produce a separate set of tapes for each tape backup session.
- **Starting a new media set for a certain period of time**, for example, each week. As a result, you will have a separate set of tapes containing all backups that have been written to tape during a week.

When planning media sets configuration, you should balance between convenient viewing of data and efficiency of using the tapes. Media sets help create distinct packs of tapes, which will be convenient, for example, if you need to bring certain tapes from an offsite storage location to restore data. However, it may use the tapes ineffectively as the new media sets always require a new tape. This option is helpful when you have a lot of backups stored offsite.

Always continuing one media set option is most efficient in terms of tape capacity usage, for the new data would be written to the same tape until it is full. However, in this case the data will be split across tapes and you may require a larger number of tapes to get a particular backup set for restore. This option is usually used in environments where the tapes are not exported from libraries.

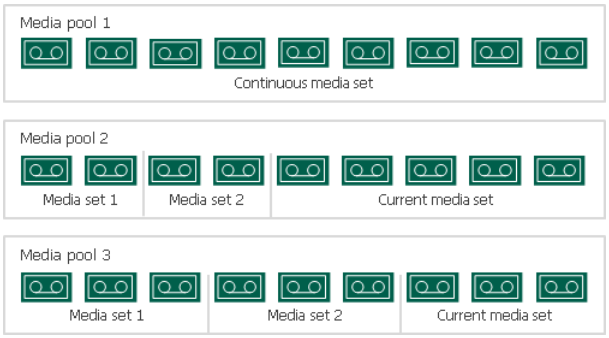

In some cases, Veeam forcibly starts a new media set, even if media pool settings do not instruct Veeam to do so. It happens when a job starts, but the tape that was planned for writing data cannot be used. In this situation, Veeam takes a new tape, and starts a new media set on it. For example, the situation may be as following:

1. The required tape is offline. If you bring such tape online before the tape job starts, it will be used as planned.
2. A backup to tape job is scheduled to export tapes from the media pool upon job completion. Exporting closes the current media set and starts a new one on the next job run.
3. The required tape is hardware or software protected.
4. In library failover on the following events: library is offline and all drives are busy. When a tape library is failed over on these events, Veeam Backup & Replication starts a new media set even if the global media pools are configured to always continue one media set.

Note that if a new media set is started, scheduled or forced, the previous media set cannot be resumed.

Media Set Names

Consistent naming of media sets makes the work with media pools easier. Therefore, we recommend to pay attention to configuring media set naming patterns when configuring media pools and GFS media pools.

To automatically create media set names, use the following variables in the **Media set name** field:

- *%date%* — the date of the media set creation, default variable
- *%id%* — the number of the media set, default variable
- *%day%* — the numeric day in month, on which the media set is created
- *%dayofweek%* — the day of week, on which the media set is created
- *%job%* — the name of job, for which the media set is created
- *%month%* — the month of creation of the media set (the month is shown as a name, for example, 'January')
- *%monthnumeric%* — the month of creation of the media set (the month is shown as a number, for example, '01' for January).
- *%time%* — the time of creation of the media set
- *%year%* — the year of creation of the media set

Backup Sets

When a tape job runs, it analyzes disk storage and spots files that match the tape job parameters. The tape job queues the files and writes them to tape. The set of files that are archived to tape within one tape job session is a backup set.

Depending on size and number of archived files, a backup set may require different amounts of tape space. For example, if the amount of data is large, it may take several tapes.

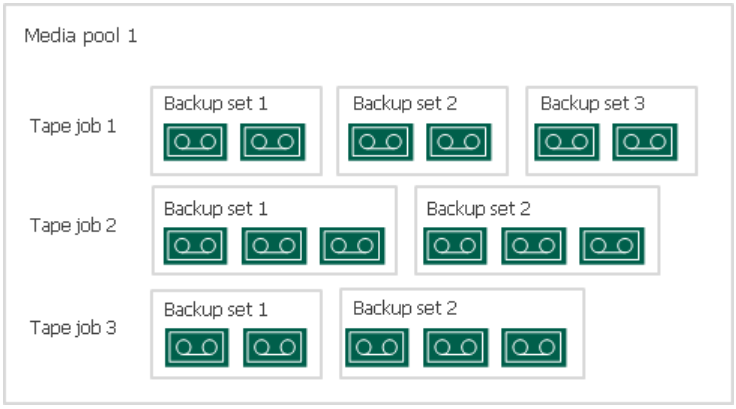

Service Media Pools

Service media pools sort tapes and manage their free capacity for writing data. Service media pools are automatically created by Veeam Backup & Replication when needed. The following predefined media pools can be created:

- **Free** – a media pool containing empty tapes. You can use this media pool to replenish user-created pools with new tapes when needed.
- **Imported** – a media pool containing non-empty tapes. These include tapes identified by the tape catalog job.
- **Retired** – a media pool containing retired tapes that reached the maximum number of re-writes. This media pool may also contain tapes with some hardware malfunction.
- **Unrecognized** – a media pool containing tapes that were loaded to tape device. They need further identification by user that can be done by running the inventory or catalog job. This media pool also contains cleaning tapes.

You cannot create, modify or delete service media pools.

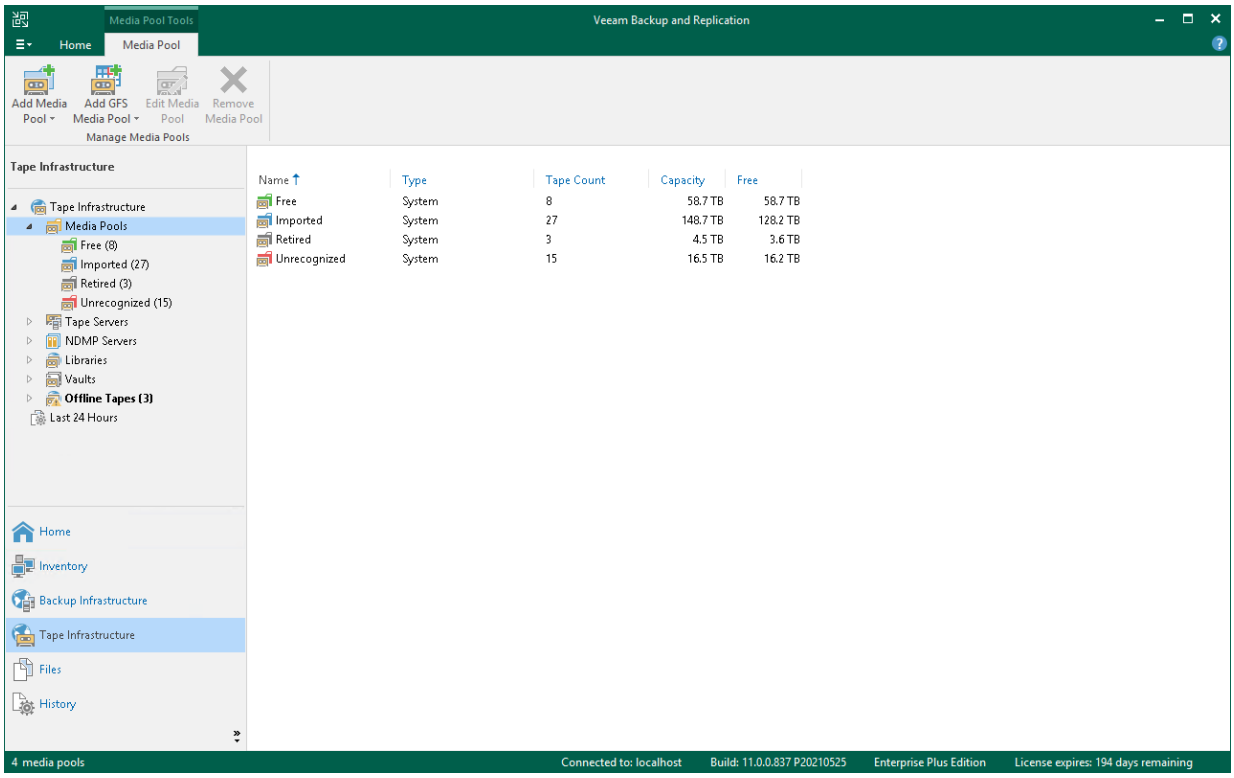

Media Pools

Media pools are target destinations for backup to tape and file to tape jobs. The media pools manage empty (or available for overwriting) tapes to allow the tape jobs write data. You can create as many media pools as you need. One media pool can be target for unlimited number of tape jobs. The GFS jobs require GFS media pools. For more information, see GFS Media Pools.

Tapes for Media Pools

For writing data, you can use rewritable tapes or WORM (Write Once Read Many) tapes. For more information, see WORM Tapes. For using WORM tapes, you need to create separate media pools. The following types of media pools are available:

- Media pools
- WORM media pools.

For more information on creating WORM media pools, see Creating Media Pools.

Media Pools Configuration

For each media pool, you can configure rules. Such rules are further applied to every tape that belongs to this media pool.

You can apply the following rules:

- Add tape libraries: select and manage tape libraries that will be used by the media pool.
- Configure tape replenishment: you can allocate some particular tapes or let the media pool take a free tape when required.
- Create media sets: media sets allow you to create sets of tapes with data for a particular time period.
- Set the data retention period: this setting lets you choose the period for which the data on tapes will be protected from overwriting (not applicable for WORM tapes).
- Enable parallel processing: you can allow the media pool to use multiple tape drives simultaneously for writing data to tape. With parallel processing enabled, the media pool can process simultaneously several tape jobs or split the data within one tape job across tape drives.
- Encrypt the data written to tapes.
- Export tapes to vault.

A tape allocated to a media pool will always be tied to its media pool. When you bring a tape with expired data online, Veeam Backup & Replication places it automatically to the media pool where the tape was written.

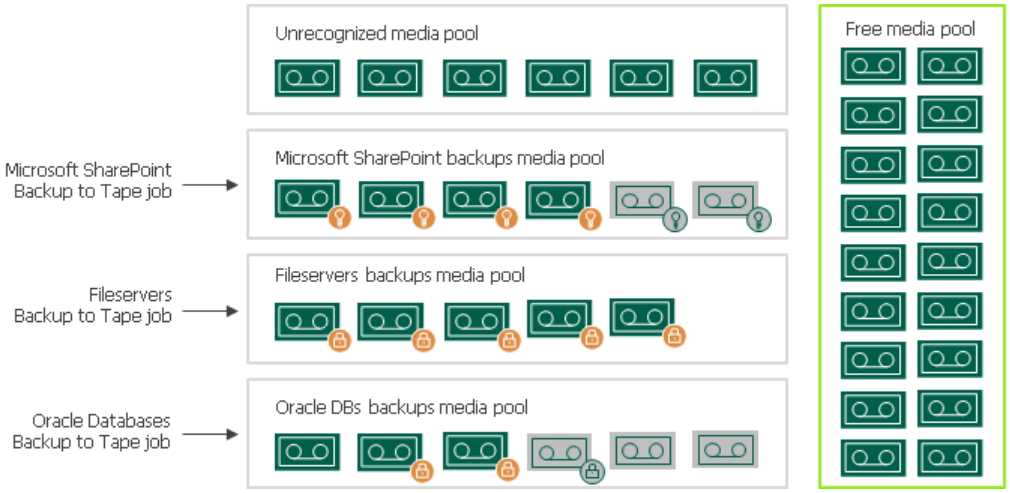

Creating Media Pools

To create a media pool, use the **New Media Pool** wizard or **New WORM Media Pool** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

Step 1. Launch New Media Pool Wizard

To run the **New Media Pool** wizard:

1. Open the **Tape Infrastructure** view and click **Add Media Pool** on the ribbon. Alternatively, you can open the **Tape Infrastructure** view, right-click the Media **Pools** node and choose **Add Media Pool**.
2. Select **Standard** if you want to create a media pool for standard tapes or **WORM** if you want to create a media pool for WORM (Write Once Read Many) tapes.

> **TIP**
>
> You can also launch the **New Media Pool** wizard when configuring tape jobs (that is, directly from the **New Backup to Tape Job** wizard and **New File to Tape Job wizard**. For more information, see Creating Backup to Tape Jobs and Creating File to Tape Jobs.

Step 2. Specify Media Pool Name

At the **Name** step of the wizard, define basic description for the new media pool.

1. In the **Name** field, enter a name for the created media pool.
2. In the **Description** field, enter a description of the new media pool. The default description contains information about the user who created the media pool, date and time when the media pool was created.

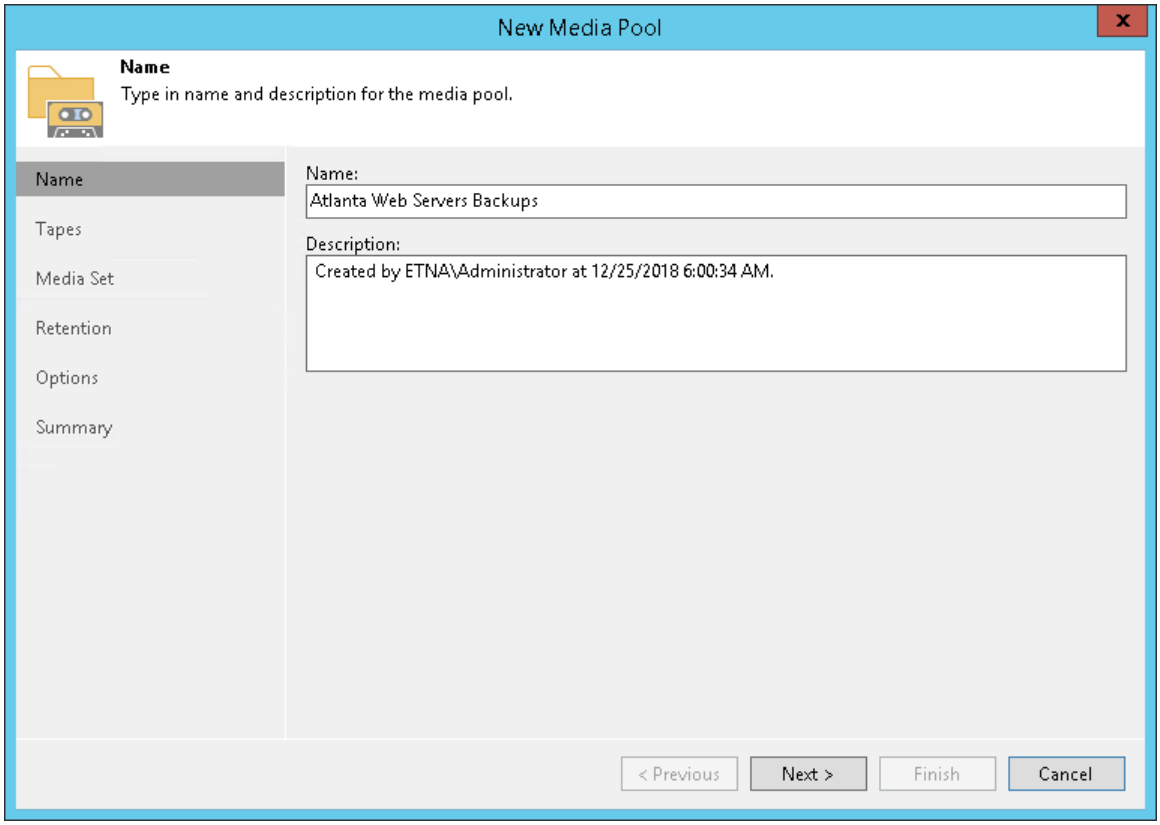

Step 3. Add Tapes to Media Pool

At the **Tapes** step of the wizard, you can select tape libraries and allocate tapes for the pool.

1. From the list in the **Tape library** field, select the tape device that you want to use in this media pool. The list contains all tape libraries managed by this Veeam backup server.
2. To add more libraries, click **Manage**. For one media pool, you can select any tape devices that are managed by this Veeam backup server even if they are connected to different tape servers.

   Click **Add** to add other tape libraries.

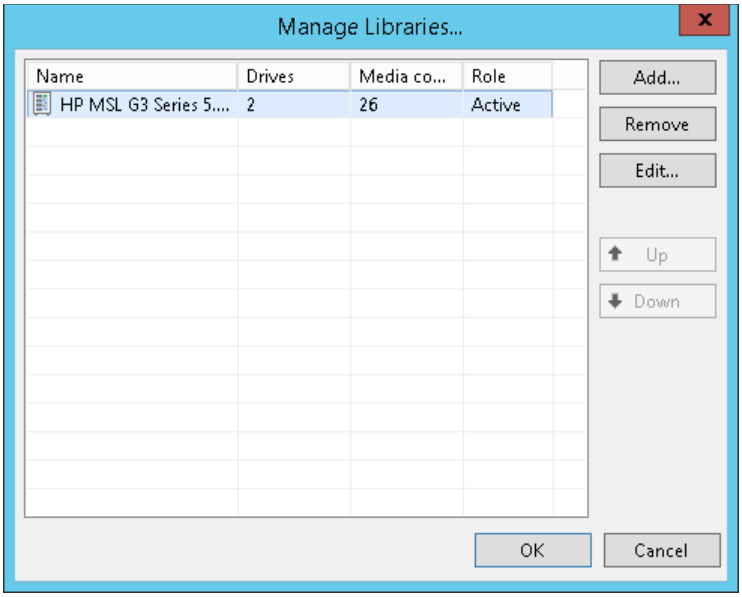

   Configure the multiple libraries mode:

   - Parallel processing: in this mode, tape jobs use drives in all libraries parallelly. For more information, see Tape Parallel Processing.

     To enable the parallel processing mode, set all libraries role to **Active**. To do this, select a library and click **Edit**. Use **Up** and **Down** buttons to position the libraries in order they must be used.

   - Failover: in this mode, the first library in the list is active and is primarily used for writing data. Other libraries added to this media pool are passive and are used for failover in the following events:

     - The primary library is unavailable.
     - The primary library has no media available.

     To enable the failover mode, set the first library role to **Active** and all other libraries roles to **Passive**. To do this, select a library from the list and click **Edit**. Use **Up** and **Down** buttons to position the libraries in order they must fail over. Also ensure that you select at least one failover event.

     On the next backup session, Veeam Backup & Replication will try to switch back to the primary library. If it is still not available, the media pool fails over to the next library as defined.

When using the failover mode, mind the risk of excessive usage of tapes, especially if you use WORM media pools. During the next backup session after the failover, switching back to the primary library may leave a tape in the reserve library not fully written. In the long run, that may result in excessive usage of tapes, which may be crucial if using WORM tapes.

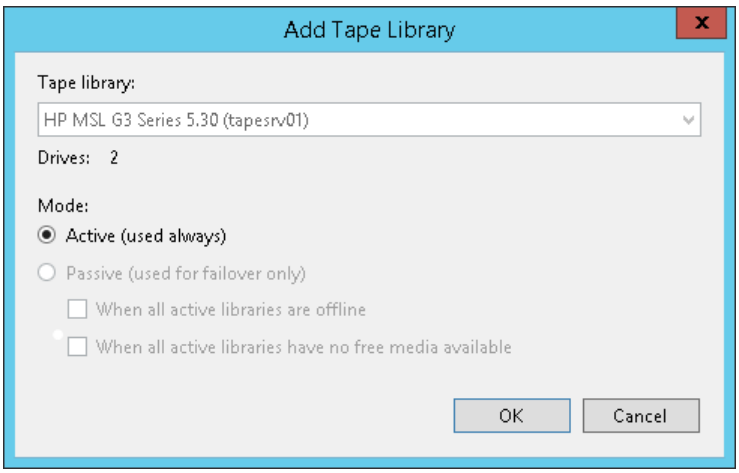

NOTE

If you enable the failover mode, notifications can be sent when Veeam Backup & Replication switches to a passive tape library during failover. Notifications are sent under either of the following conditions:

- The **Notify on failure** check box is selected in the Veeam Backup & Replication global notification settings. For more information, see the Veeam Backup & Replication User Guide, section *Configuring Global Email Notification Settings*.
- The **Notify on error** check box is selected in the tape job notification settings. For more information, see Notifications Settings.

If you enable the parallel processing mode, no email notifications are sent when Veeam Backup & Replication switches from one library to another.

3. To allocate specific tapes from the library, click the **Add** in the **Tapes** field and select tapes that should be added to the media pool. Allocated tapes will be reserved for this media pool; other media pools will not be able to use these tapes. The capacity and free space on the allocated tapes will be displayed in the bottom right corner.

   Allocate standard tapes to the standard media pool and WORM tapes to the WORM media pool. Do not mix standard and WORM tapes in one media pool.

4. To make the media pool replenishable, select the **Add tapes from Free media pool automatically when more tapes are required** check box.

   With this option enabled, additional tapes will be allocated from the Free media pool when needed. That is, when a tape job uses all available tapes from this media pool, Veeam Backup & Replication will automatically add the required number of tapes from the Free media pool to let the job complete. If the option is disabled, the job will pause and prompt the backup administrator to add new tapes to the media pool.

To a standard media pool, Veeam Backup & Replication will add only standard tapes; to a WORM media pool Veeam Backup & Replication will add only WORM tapes.

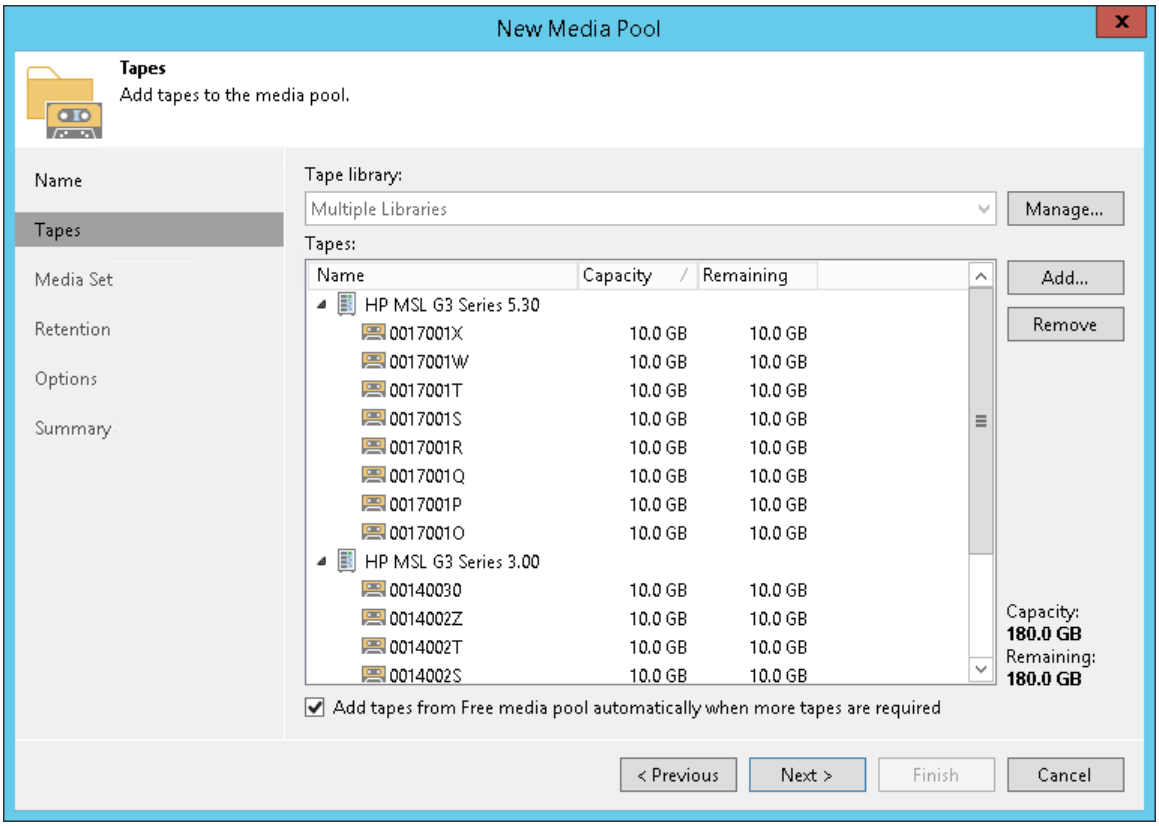

Step 4. Specify Media Set Options

At the **Media Set** step of the wizard, specify rules for creating media sets. For more information, see Media Sets.

In the **Media set name** field, define the pattern according to which created media set(s) will be named. For more information, see Media Set Names.

In the **Automatically create new media set** section, specify conditions for creating new media sets on tapes allocated to the media pool. The following options are available:

- **Do not create, always continue using current media set**. If this option is selected, each subsequent backup session will write its backup set to the existing media set: it will append backup content to the content that was written to tape with a previous backup session. If, however, a backup set is started with a new tape, Veeam Backup & Replication will create a new media set for it.
- **Create new media set for every backup session**. If this option is selected, a new media set will be created for each new backup session. Each backup session will write its backup set starting with a new tape.
- **Daily at**. If this option is selected, you can specify day and time when new media sets should be created. For example, if at the end of the week you send weekly media sets to offsite storage, you can schedule creation of new media sets at the beginning of each week.

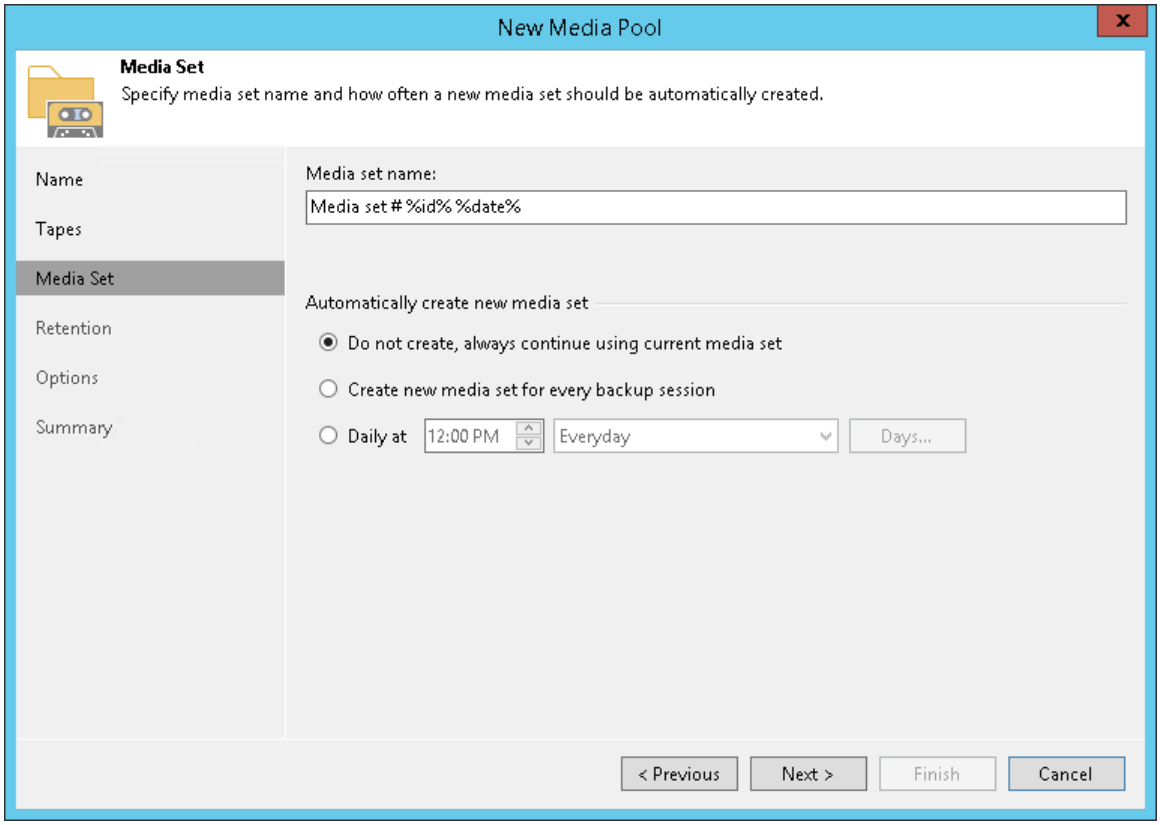

Step 5. Specify Retention Settings

At the **Retention** step of the wizard, specify the retention settings and offline media tracking options for the tapes in this media pool. For more information, see Tape Data Retention.

For WORM media pools, retention option is inapplicable; retention settings are inactive.

For standard media pools, you can select one of the following options:

- **Do not protect data (cyclically overwrite tapes as required)**. If this option is selected, tapes allocated to the pool will be overwritten, starting with the tape that stores the oldest archive.
- **Protect data for (time interval)**. If this option is selected, data on tapes will be preserved for the specified period. When this period is over, tapes will be overwritten, starting with the tape that stores the oldest archive. This setting must accord with the retention policy specified for the backup chain that you plan to archive to tape. For more information on disk retention, see Data Retention.
- **Never overwrite data.** If this option is selected, data on tapes will not be overwritten. If there is not enough tape capacity for the archiving job to complete, Veeam Backup & Replication will pause the job and prompt the backup administrator to add new tapes to the media pool.

> **NOTE**
>
> If at the previous step you selected the **Do not create, always continue using current media set** option, keep in mind that the retention period for such media sets starts after the first tape in the media set becomes full.

In the **Offline media tracking** section, you can set automatic moving tapes to a media vault when they are brought offline. For more information on vaults, see Media Vaults.

To enable this option:

1. Select **Move all offline tapes from this media pool into the following media vault** check box.
2. From the **Vault** list, select the media vault where you want to automatically move the tapes.

If the vault is not added yet, click **Add New** to open the **New Media Vault** wizard.

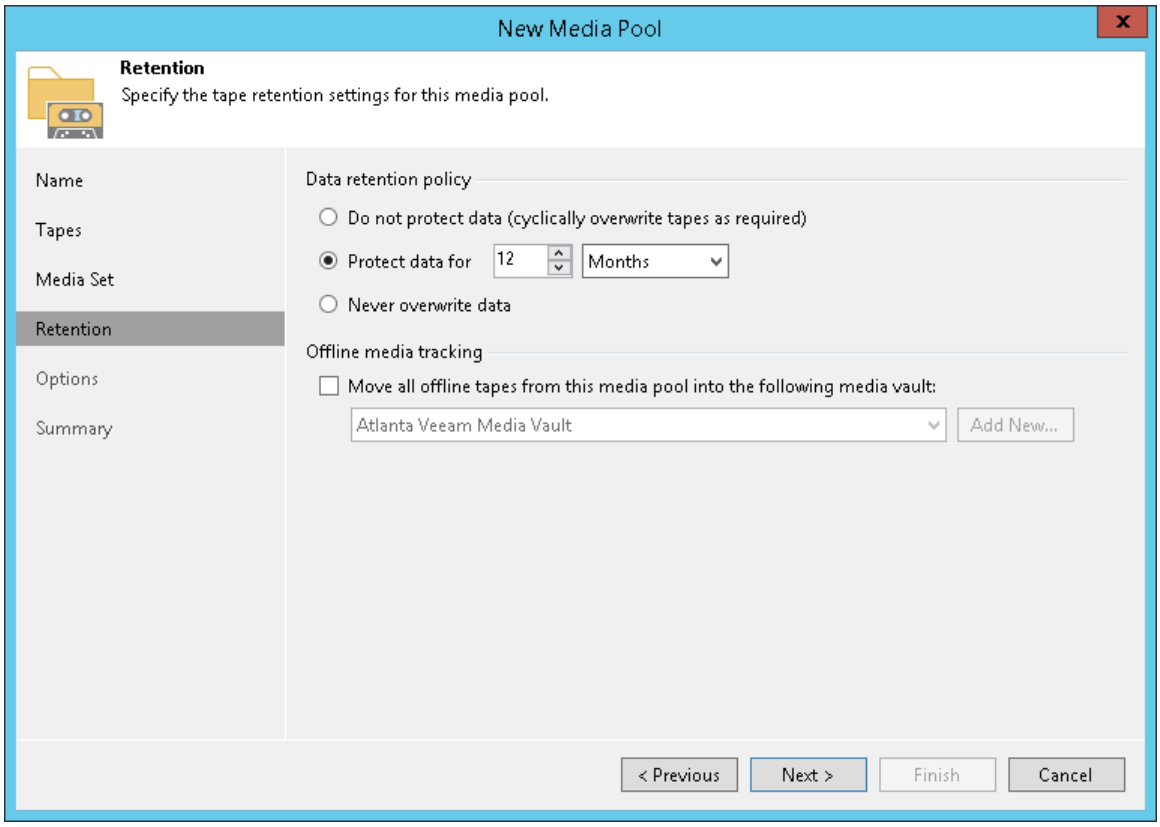

Step 6. Add Optional Media Pool Settings

At the **Options** step of the wizard, you can enable parallel processing and encryption.

1. Select the **Enable parallel processing for tape jobs using this media pool** check box if you want to process multiple tape jobs simultaneously. Set the maximum number of drives that the media pool can use in parallel.

   If you want to use multiple drives simultaneously to process multiple source backup chains within the tape jobs, select the **Enable parallel processing of backup chains within a single job** check box. For more information, see Tape Parallel Processing.

2. To encrypt the content of files or backups archived to tapes in the media pool, select the **Use encryption** check box. Then, select a password you want to use to encrypt data written to tape from the **Password** field. You can also click **Add** or use the **Manage passwords** link to add a password. For more information, see section Tape Encryption in the Veeam Backup & Replication User Guide.

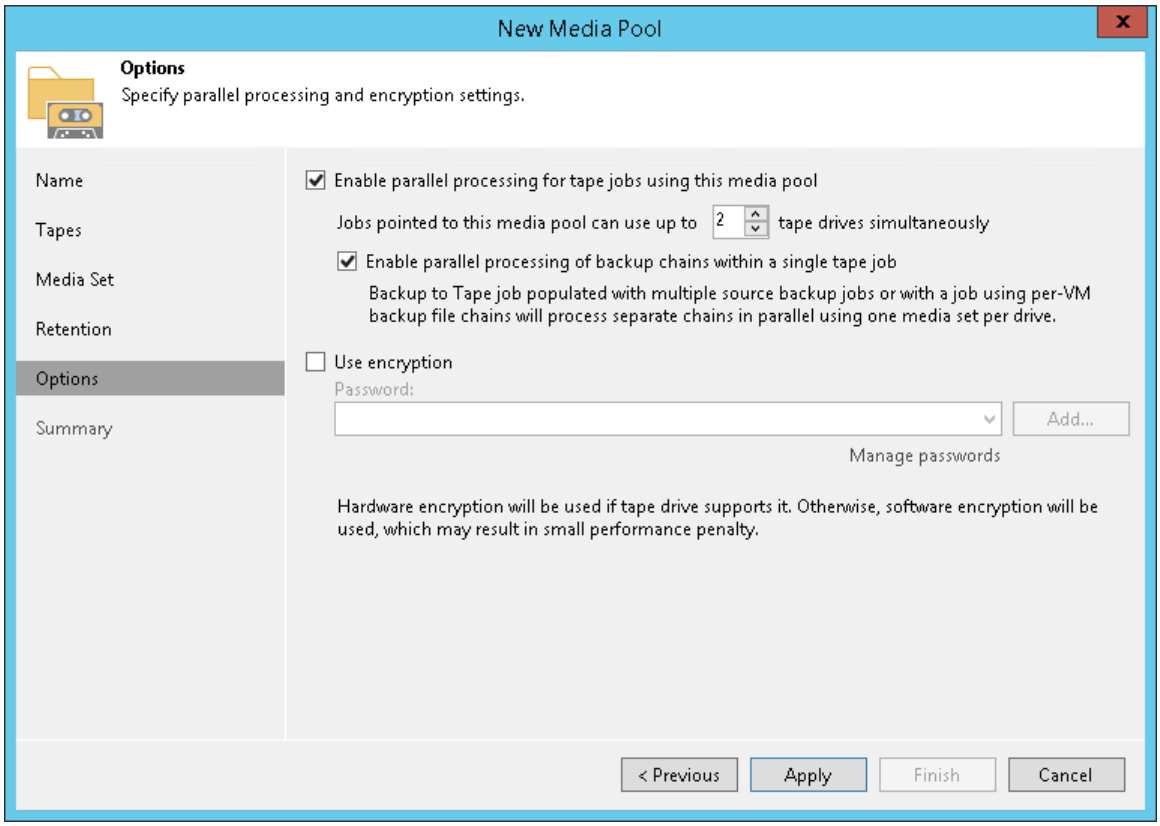

Step 7. Finish Working with Wizard

Review the media pool settings and click **Finish** to complete the wizard.

The new media pool will be available under the **Tape Infrastructure** > **Media Pools** node in the **Tape Infrastructure** view.

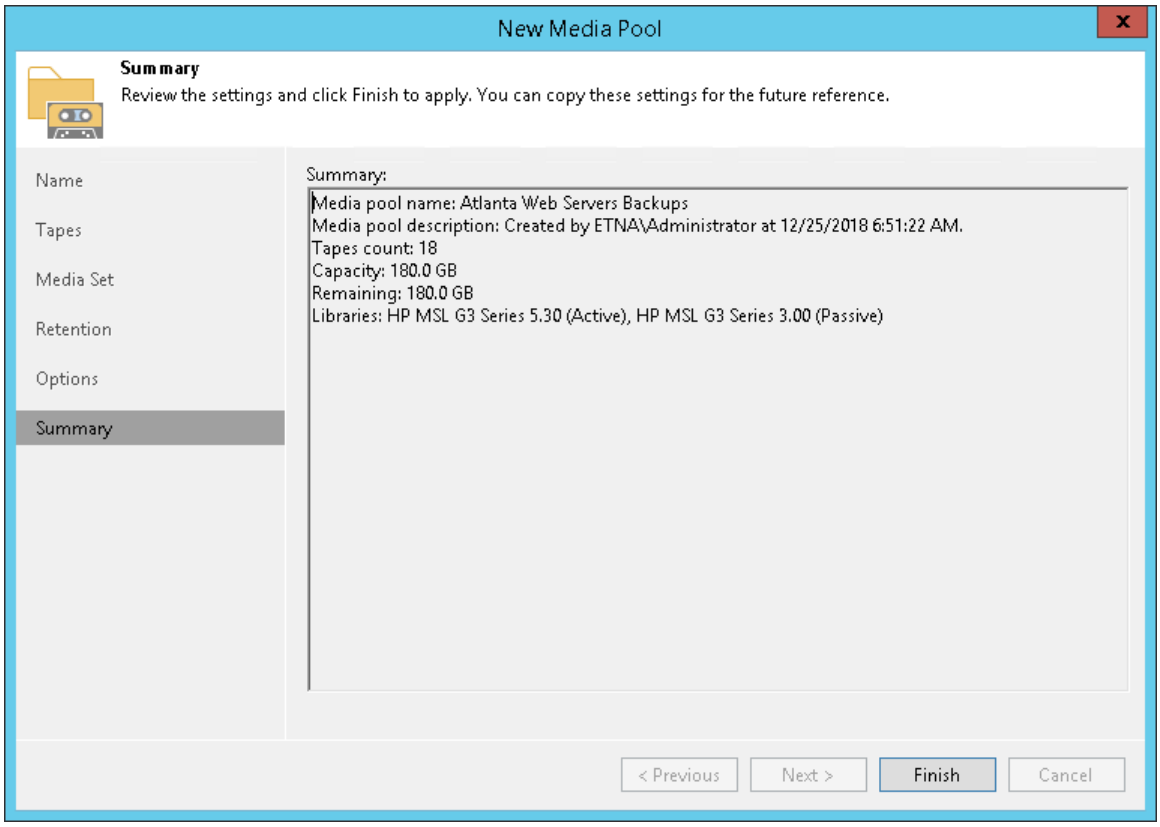

GFS Media Pools

The GFS media pools are special media pools that store data to tape according to the GFS, or Grandfather-Father-Son rotation scheme. The GFS media pools are targets for GFS tape jobs. For more information, see GFS Backup to Tape.

Tapes for GFS Media Pools

For writing data, you can use rewritable tapes or WORM (Write Once Read Many) tapes. For more information, see WORM Tapes. For using WORM tapes, you need to create separate media pools. The following types of media pools are available:

- GFS media pools.
- WORM GFS media pools.

For more information on creating WORM media pools, see Creating GFS Media Pools.

GFS Media Sets

GFS media pool keeps five predefined media sets for storing the machine backups with a tiered retention policy scheme:

- Daily
- Weekly
- Monthly
- Quarterly
- Yearly

You can disable a media set or sets if you do not need them.

GFS Tape Jobs

A GFS media pool can be target for unlimited number of GFS tape jobs.

Creating GFS Media Pools

To create a GFS media pool, use the **New GFS Media Pool** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

Step 1. Launch New GFS Media Pool Wizard

To run the **New GFS Media Pool** wizard:

1. Open the **Tape Infrastructure** view and click **Add GFS Media Pool** on the ribbon. Alternatively, you can open the **Tape Infrastructure** view, right-click the **Media Pools** node and choose **Add GFS Media Pool**.
2. Select **Standard** if you want to create a media pool for standard tapes or **WORM** if you want to create a media pool for WORM (Write Once Read Many) tapes.

> **TIP**
>
> You can also launch the **New GFS Media Pool** wizard when configuring backup to tape jobs (that is, directly from the **New Backup to Tape Job** wizard. For more information, see Creating Backup to Tape Jobs.

Step 2. Specify Media Pool Name

At the **Name** step of the wizard, define basic description for the new GFS media pool.

1. In the **Name** field, enter a name for the created media pool.
2. In the **Description** field, enter a description of the new media pool. The default description contains information about the user who created the media pool, date and time when the media pool was created.

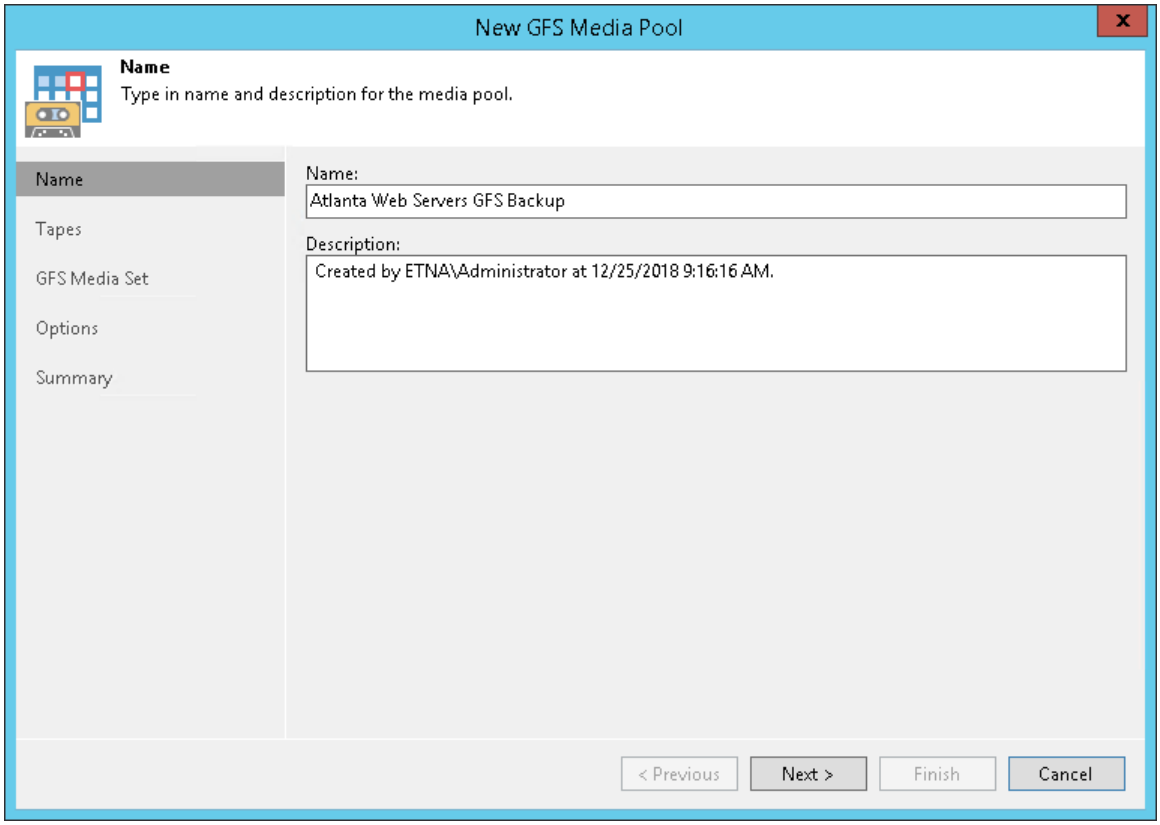

Step 3. Add Tapes to Media Pool

At the **Tapes** step of the wizard, you can select tape libraries and allocate tapes for the pool.

1. From the list in the **Tape library** field, select the tape device that you want to use in this media pool. The list contains all tape libraries managed by this Veeam backup server.
2. To add more libraries, click **Manage**. For one media pool, you can select any tape devices that are managed by this Veeam backup server even if they are connected to different tape servers.

   Click **Add** to add other tape libraries.

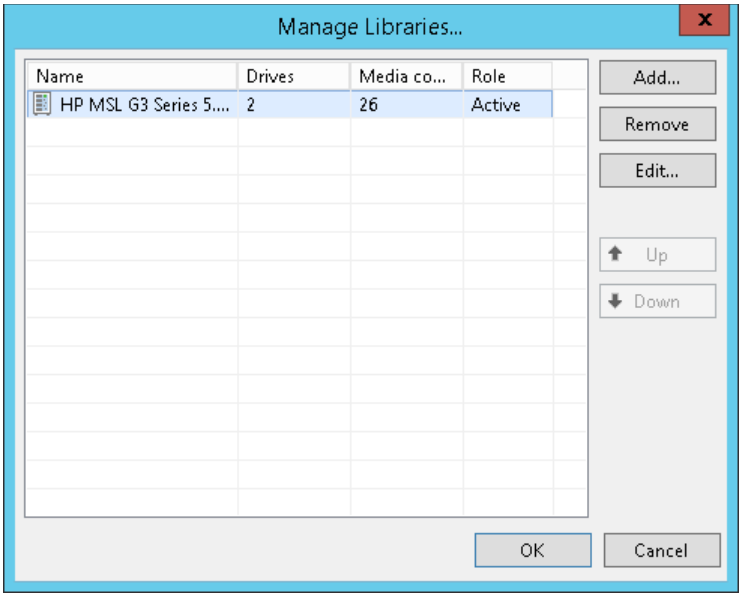

   Configure the multiple libraries mode:

   - Active: in this mode, tape jobs use drives in all libraries parallelly. For more information, see Tape Parallel Processing.

     To enable the parallel processing mode, set all libraries role to **Active**. To do this, select a library and click **Edit**. Use **Up** and **Down** buttons to position the libraries in order they must be used.

   - Passive: in this mode, the first library in the list is active and is primarily used for writing data. Other libraries added to this media pool are passive and are used for failover in the following events:

     - The primary library is unavailable.
     - The primary library has no media available.

     To enable the failover mode, set the first library role to **Active** and all other libraries roles to **Passive**. To do this, select a library from the list and click **Edit**. Use **Up** and **Down** buttons to position the libraries in order they must fail over. Also ensure that you select at least one failover event.

On the next backup session, Veeam Backup & Replication will try to switch back to the primary library. If it is still not available, the media pool fails over to the next library as defined.

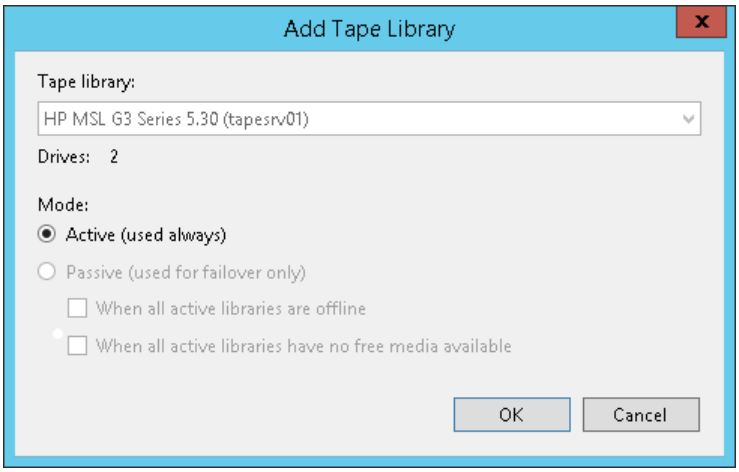

NOTE

If you enable the failover mode, email notifications can be sent when Veeam Backup & Replication switches to a passive tape library during failover. Notifications are sent under either of the following conditions:

- The **Notify on failure** check box is selected in the Veeam Backup & Replication global notification settings. For more information, see the Veeam Backup & Replication User Guide, section *Configuring Global Email Notification Settings*.
- The **Notify on error** check box is selected in the tape job notification settings. For more information, see Notifications Settings.

If you enable the parallel processing mode, no email notifications are sent when Veeam Backup & Replication switches from one library to another.

3. To allocate specific tapes from the library, click the **Add** in the **Tapes** field and select tapes that should be added to the media pool. Allocated tapes will be reserved for the created media pool; other media pools will not be able to use these tapes. The capacity and free space on the allocated tapes will be displayed in the bottom right corner.

   Allocate standard tapes to the standard media pool and WORM tapes to the WORM media pool. Do not mix standard and WORM tapes in one media pool.

4. To make the media pool replenishable, select the **Add tapes from Free media pool automatically when more tapes are required** check box.

   With this option enabled, additional tapes will be allocated from the Free media pool when needed. That is, when a backup to tape or file to tape job uses all available tapes from this media pool, Veeam Backup & Replication will automatically add the required number of tapes from the Free media pool to let the job complete. If the option is disabled, the job will pause and prompt the backup administrator to add new tapes to the media pool.

To a standard media pool, Veeam Backup & Replication will add only standard tapes; to a WORM media pool Veeam Backup & Replication will add only WORM tapes.

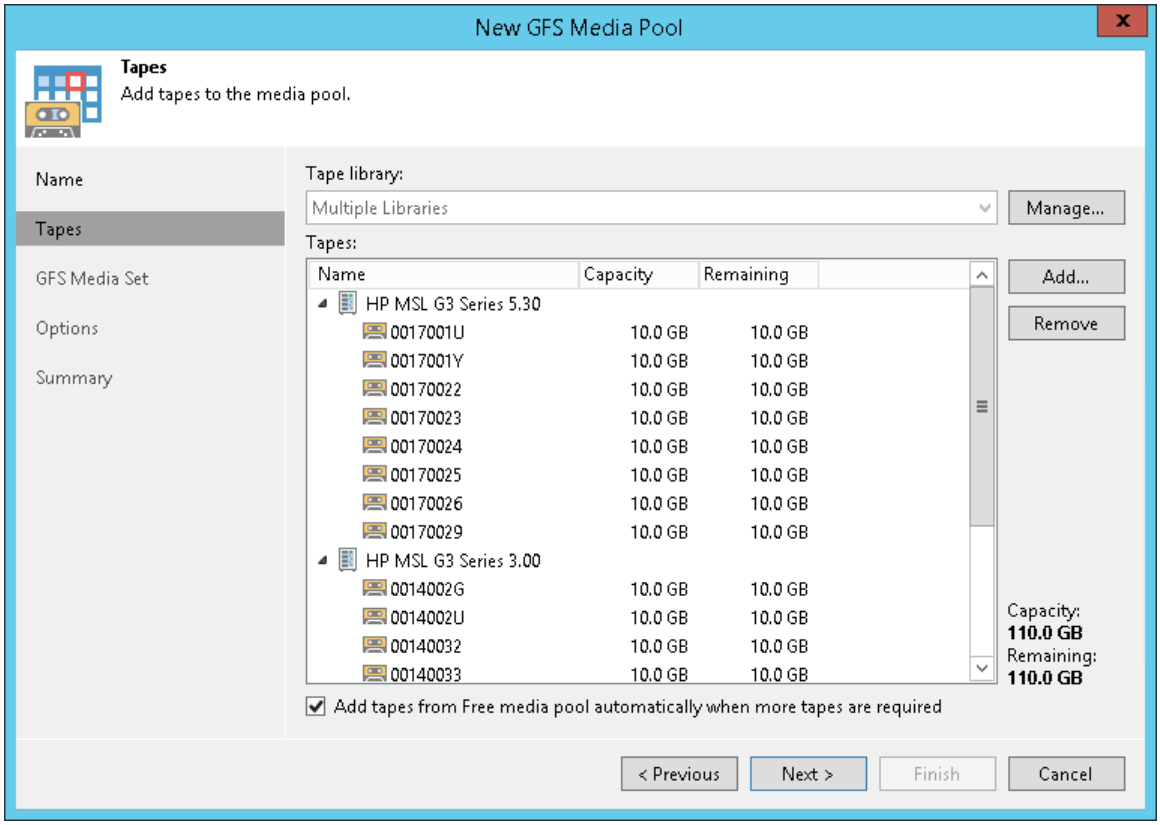

Step 4. Specify Media Set Options

In the GFS media pool, a separate media set is created for each backup cycle: daily, weekly, monthly, quarterly and yearly.

For each media set, you can set the data retention period. The retention period indicates for how long the data is protected from overwriting.

Usually, daily media sets do not contain full backups as daily backups are mostly used to keep day-to-day incremental changes. To complete the backup chain they use the periodic full backup stored in the weekly media set. Due to this, the daily media set creation cannot be enabled unless the weekly media set creation is enabled.

Click **Advanced** to configure the advanced media set options.

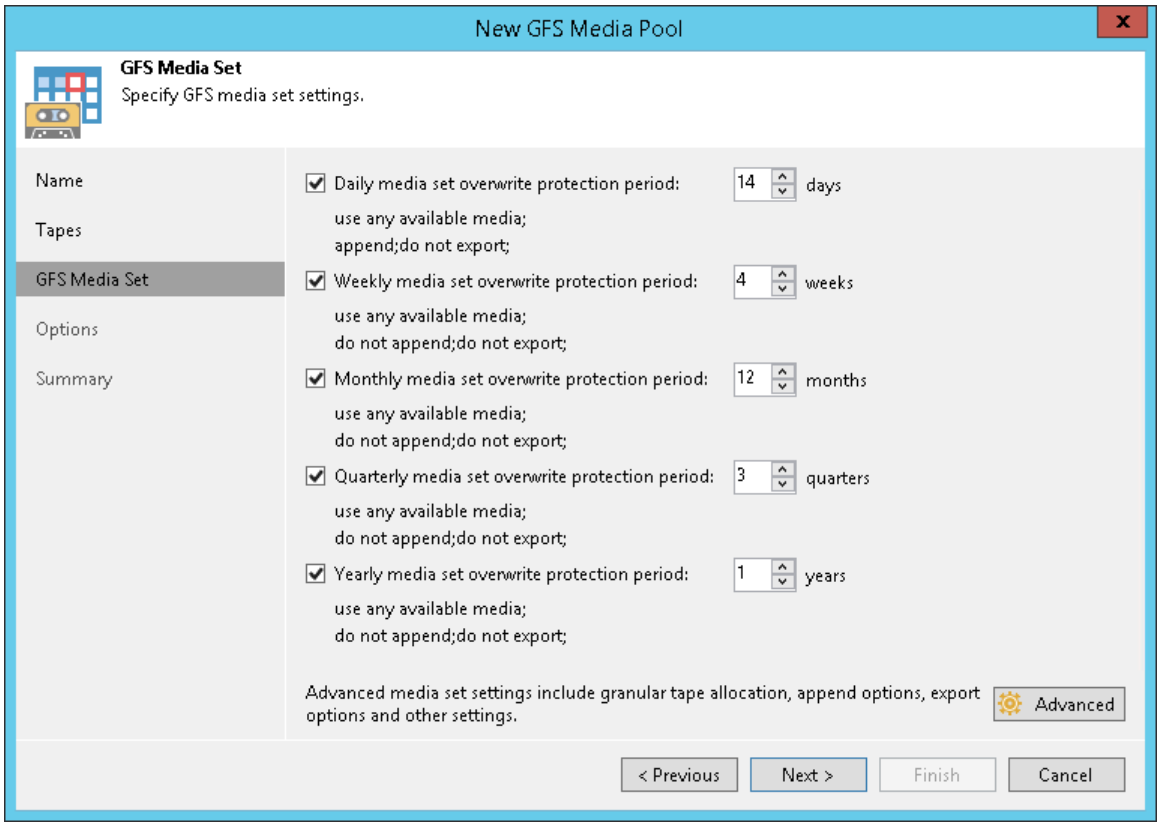

Step 5. Specify Advanced Media Set Options

You can granularly configure rules for each media set.

At the **GFS Media Set** step of the wizard, click **Advanced**. Select the **Daily** tab to configure the daily media set.

1. If you want the **Daily** media set to span several selected tapes, click **Add** and select tapes that should be added to the media set. You can select from tapes added to the media pool at the **Tapes** step of the wizard. The selected tapes will be displayed in the **Daily media set** field.

   This option is useful if you want the media set to use some particular user-defined tapes. Otherwise, you can use the **Add tapes from media pool automatically** option.

2. In the **Media set name** field, define the pattern according to which the created media set will be named. For more information, see Media Set Names.

3. Select the **Add tapes from media pool automatically** check box to allow the media set to take a tape from the media pool when required.

4. Select **Append backup files to incomplete tapes** if you want to write the next backup set to the tape where the previous backup set was written. Otherwise, the GFS tape job will use a new tape for each backup set.

5. You can set automatic moving tapes to a media vault when they are brought offline. To enable this option:

   a. Select **Move all offline tapes from this media pool into the following media vault** check box.

   b. From the list of vaults, select the media vault where you want to automatically move the tapes.

   c. If the vault is not added yet, click **Add New** to open the **New Media Vault** wizard. For more information, see Creating Vaults.

> **NOTE**
>
> Tapes written in this media set will always be used by this media set. Other media sets will not be able to use these tapes even after the data expires.

Select the **Weekly**, **Monthly**, **Quarterly** of **Yearly** tabs to set the advanced media set options for weekly, monthly, quarterly and yearly media sets.

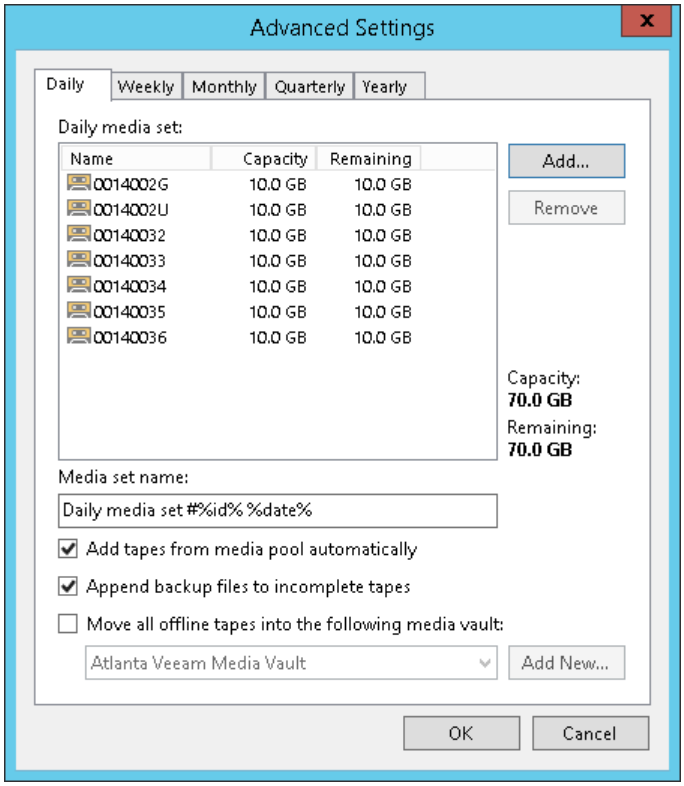

Step 6. Add Optional Media Pool Settings

At the **Options** step of the wizard, you can enable parallel processing and encryption.

1. Select the **Enable parallel processing for tape jobs using this media pool** check box if you want to process multiple tape jobs simultaneously. Set the maximum number of drives that the media pool can use in parallel.

   If you want to use multiple drives simultaneously to process multiple source backup chains within the tape jobs, select the **Enable parallel processing of backup chains within a single job** check box. For more information, see Tape Parallel Processing.

2. To enable encryption, select the **Use encryption** check box. This option allows you to encrypt the content of files or backups archived to tapes in the media pool.

   From the **Password** field, select a password you want to use to encrypt data written to tape. You can also click **Add** or use the **Manage passwords** link to add a password. For more information, see the Veeam Backup & Replication User Guide, section *Tape Encryption*.

   If the source job produces encrypted backups, and you enable encryption in the properties of the media pool, the data on tape will be encrypted twice. To decrypt such tape backups, you will need to specify two passwords:

   - Password for the GFS media pool.
   - Password for the source backup job.

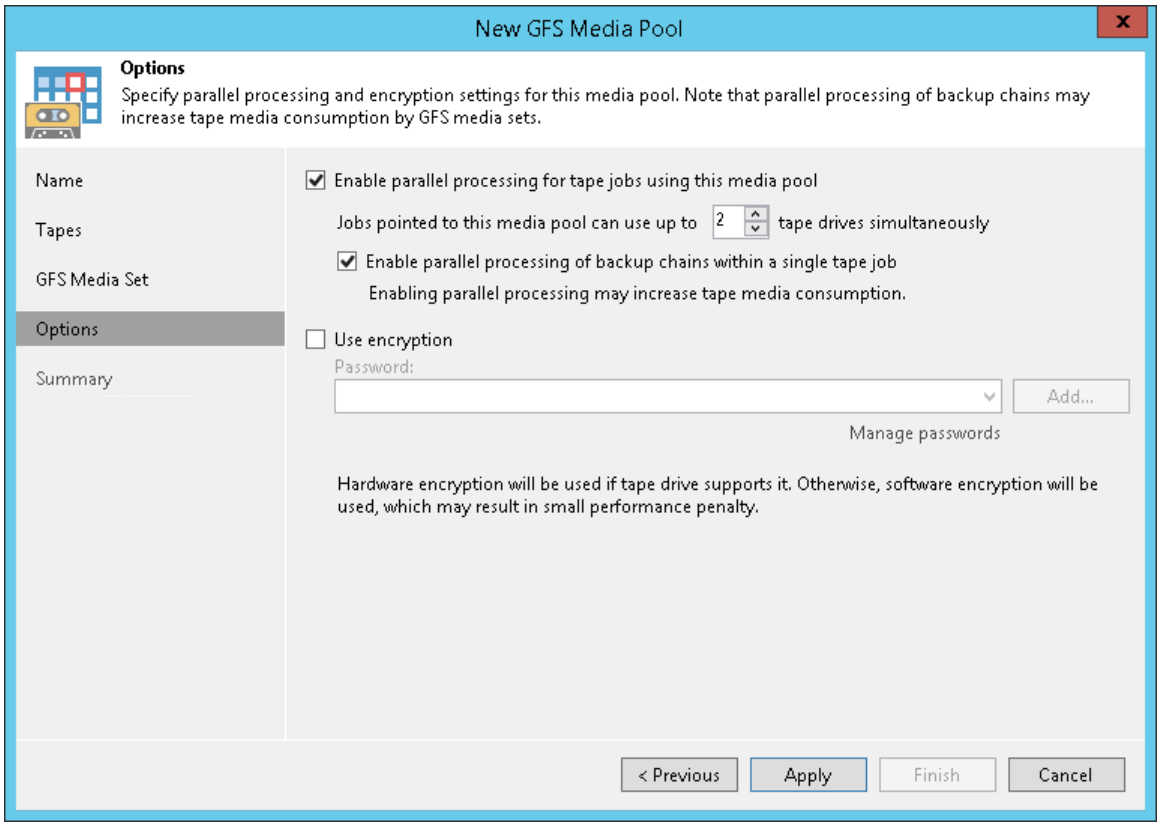

Step 7. Finish Working with Wizard

Review the media pool settings and click **Finish** to complete the wizard.

A new media pool will be available under the **Tape Infrastructure > Media Pools** node in the **Tape Infrastructure** view.

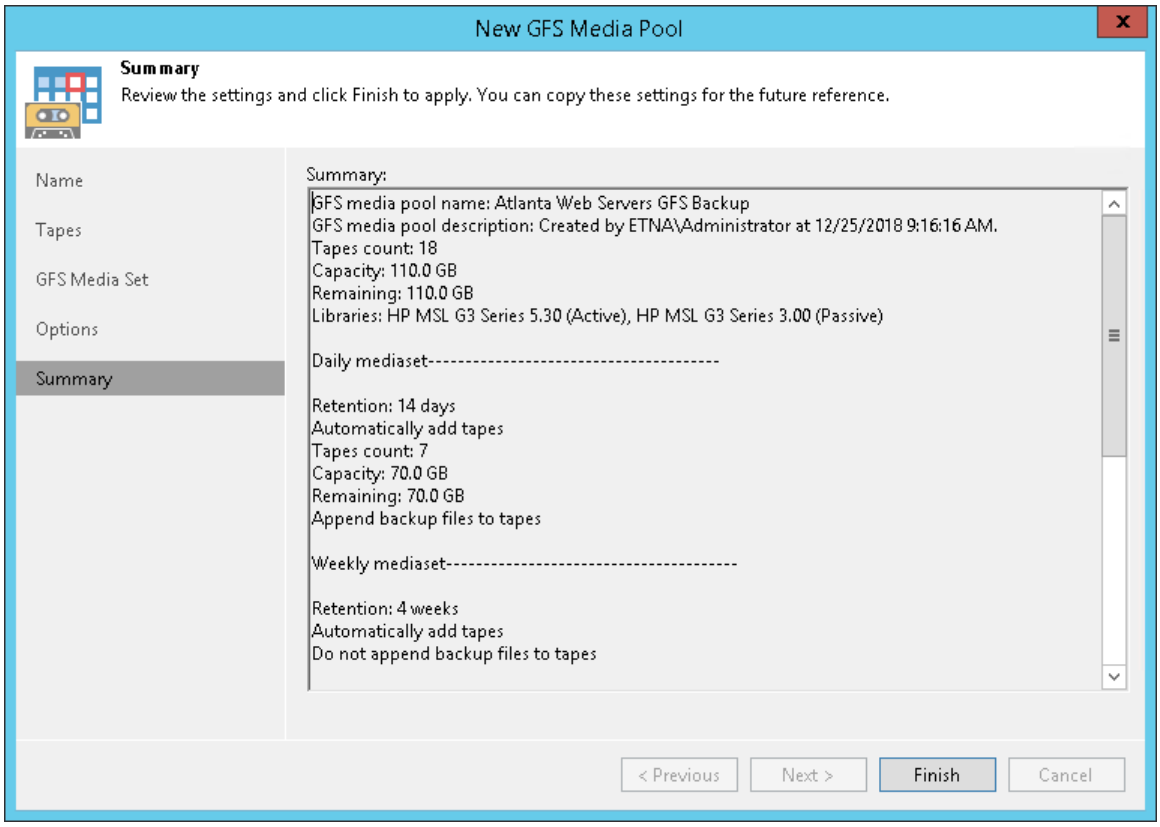

Moving Tapes to Another Media Pool

Veeam Backup & Replication allows you to move tapes between media pools. You can move tapes between media pools and GFS media pools, and from and to the **Free** media pool. Veeam Backup & Replication uses tapes from the **Free** media pool to automatically add them to other media pools and write backups to them if there are no more free tapes available in those media pools.

You can move tapes between media pools only if the media pools span the same tape libraries.

You cannot manually move tapes to **Imported**, **Unrecognized** or **Retired** media pools.

> **IMPORTANT**
>
> When you move a tape to any media pool, Veeam Backup & Replication marks this tape as free.

To move tapes from one media pool to another:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under the **Media Pools** or under the **Libraries >** *LibraryName* node **> Media > Online**.
3. Select tapes you want to move and click **Move to** on the ribbon. Choose the target media pool from the list.
   Alternatively, you can right-click selected tapes and choose **Move to Media Pool**. Next, choose the target media pool from the list.

> **NOTE**
>
> You cannot move tapes that are protected. To move such tapes, you need to switch the protection off first.

Modifying Media Pools

If necessary, you can modify settings of a media pool.

To modify media pool settings:

1. Open the **Tape Infrastructure** view.
2. Navigate to the **Media Pool** node.
3. Right-click the necessary media pool and choose **Properties**. Alternatively, select a media pool and click **Edit Media Pool** on the ribbon.
4. Go through the **Edit Media Pool** wizard to change the necessary settings.
5. Apply changes.

Limitations for Modifying Media Pools

You can modify only regular or GFS media pools; service media pools cannot be modified.

Removing Media Pools

If you no longer need a media pool or a GFS media pool, you can remove it:

1. Open the **Tape Infrastructure** view.
2. Navigate to the **Media Pool** node.
3. Right-click the necessary media pool and choose **Remove Media Pool** from the shortcut menu. Alternatively, select a media pool and click **Remove Media Pool** on the ribbon.
4. In the displayed dialog box, click **OK** to confirm deletion.

When you remove a media pool, its tapes are moved to the **Imported** media pool. The data on the tapes is not erased. You can use the tapes in the **Imported** media pool for restore operations.

The **Imported** media pool does not have any retention settings. The tapes in the **Imported** media pool are never overwritten as they are not used by any jobs. The media set number and the media set expiration date are inherited and remain for user information. The data will remain on tapes until you choose to erase it manually.

If you do not need the data on tapes, you can move the imported tapes **to the Free** media pool or any other media pool. In this case, the tapes will be marked as free. For more information, see Moving Tapes to Another Media Pool.

Limitations for Removing Media Pools

- You can remove only regular media pools or GFS media pools; service media pools cannot be removed.
- You cannot remove a media pool that is used in a backup to tape or file to tape job. To remove such media pool, you need to point corresponding jobs to other media pools or GFS media pools.

Vaults

Vaults are logical containers that help you to additionally organize offline tapes. When the tapes are recorded and moved offsite, you need to have easy ways to manage data on such tapes.

Vaults visualize information about tapes. They store information about original location of tapes, such as library and media pool, data written to tapes and retention terms for it. Vaults are not limited to any original location of tapes: you can place tapes from different libraries and media pools to one vault.

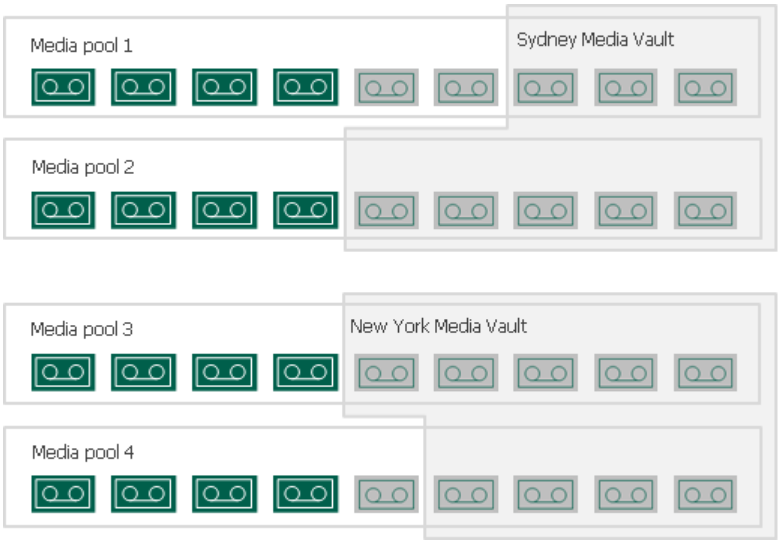

You can virtually move tapes to the vaults grouping them by any criteria. For example, if you have several offsite storage locations for physical tapes, you can create a vault for each of them. When you transport the physical tapes to the offsite storage, you can accordingly move the offline tapes to a vault in your Veeam backup console mirroring the physical storage. This will allow you to have convenient representation of the list of tapes in each of your physical storage locations.

Tapes can be moved to a vault automatically if you instruct media pools to move tapes to a selected vault after the tapes go offline. You can also move the tapes from one vault to another. Tapes will not be shown in their vault when you bring them online. The tapes that have been moved to a vault stay visible in their media pools and under the **Offline** node. When you remove a tape from the Veeam backup server, it is automatically deleted from the vault.

To start working with vaults, you need to create one or more of them. Vaults are virtual container, and do not require any physical resources.

> **NOTE**
>
> You can work with tape vaults if Veeam Backup & Replication Enterprise license or higher is installed on the Veeam backup server.

Creating Vaults

Created vaults are shown under the **Vaults** node. If you have no vaults created, the **Vaults** node will not be shown.

To create a new vault, use the **New Media Vault** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

To add a vault, follow the next steps:

Step 1. Launch New Media Vault Wizard

To run the **New Media Vault** wizard, do either of the following:

- Open the **Tape Infrastructure** view and select the **Tape Infrastructure** node. Click **Add Media Vault** on the ribbon.
- Open the **Tape Infrastructure** view, right-click the **Tape Infrastructure** node and choose **Add Media Vault**.

> **TIP**
>
> You can also launch the **New Media Vault** wizard when configuring media pools (that is, directly from the **New Media Pool** or **New GFS Media Pool** wizard. For more information, see Creating Media Pools or Creating GFS Media Pools.

Step 2. Specify Media Vault Name

At the **Name** step of the wizard, you should define basic description for the new vault.

1. In the **Name** field, enter a name for the created vault.
2. In the **Description** field, enter a description of the new vault. The default description contains information about the user who created the vault, date and time when the vault was created.
3. If you want to set overwrite protection to all tapes in this vault, select the **Protect all tapes moved to this media vault automatically** check box. For more information, see Tape Protection.

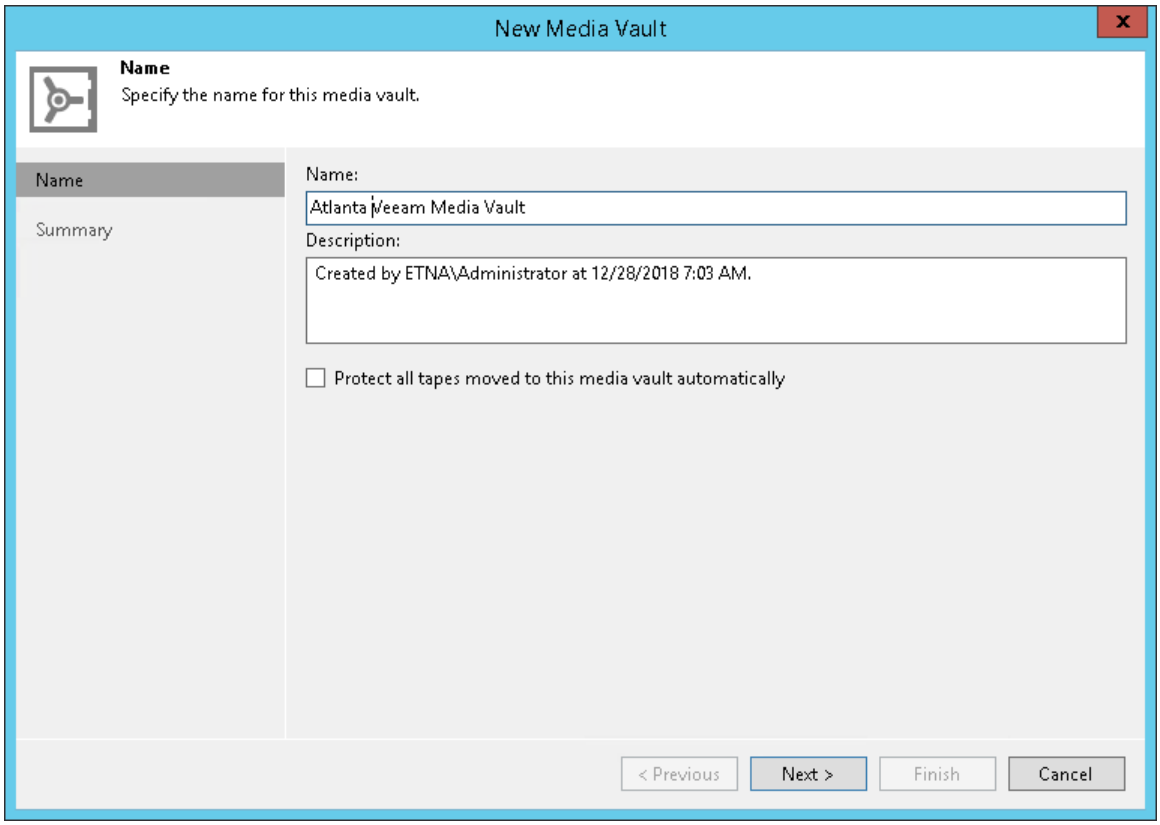

Step 3. Finish Working with Wizard

Review the media vault settings and click **Finish** to complete the wizard.

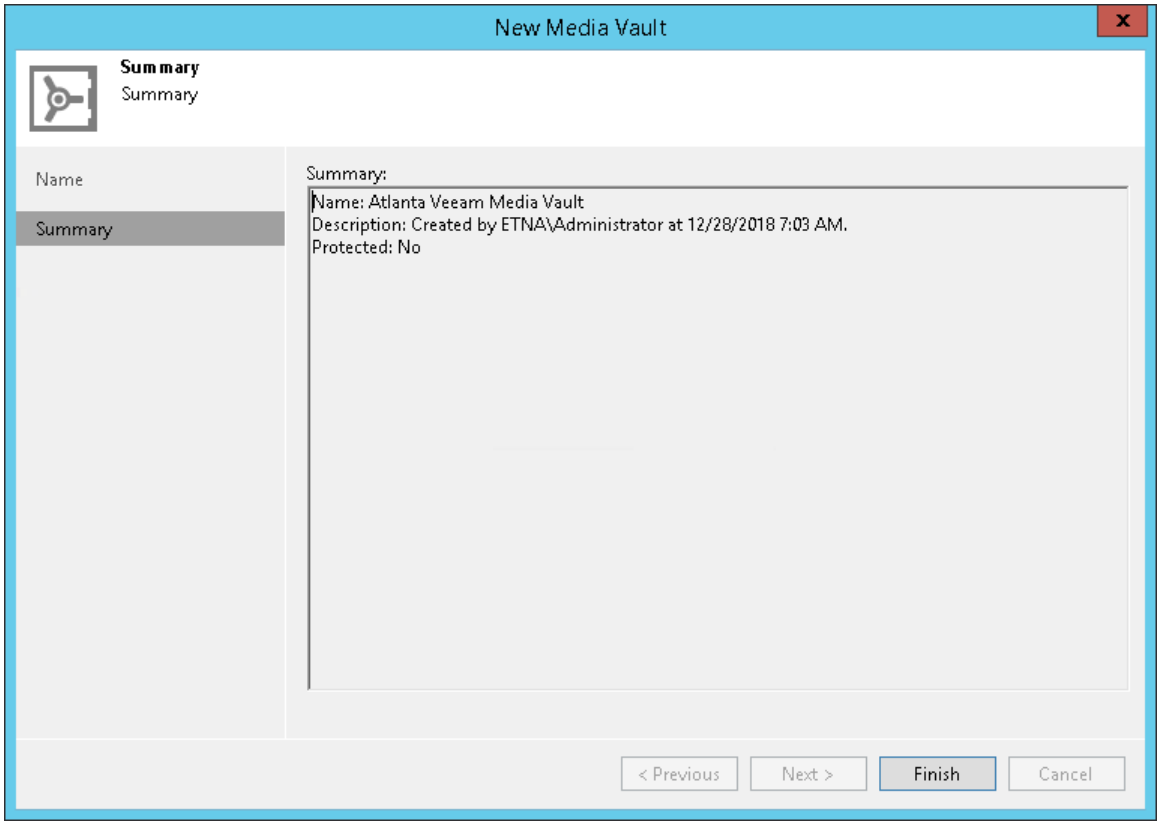

Modifying Vaults

To modify vault settings:

1. Open the **Tape Infrastructure** view.
2. Navigate to the **Vaults** node.
3. Right-click the necessary vault and choose **Properties**. Alternatively, select a vault and click **Edit Vault** on the ribbon.
4. Go through the **Edit Media Vault** wizard to change the necessary settings.
5. Apply changes.

Removing Vaults

If you no longer need a vault, you can remove it. Mind the following limitation:

- You cannot remove a vault that contains tapes. To be able to remove such vault, first move tapes from this vault to other vaults or remove them from the vault. When you remove a tape from the vault, you can see it in its media pool or under **Media** > **Offline** node.
- You cannot remove a vault that has a reference to any media pool. To be able to remove such vault, first edit settings of the corresponding media pool and remove a reference to this vault.

To remove a vault:

1. Open the **Tape Infrastructure** view.
2. Navigate to the **Vaults** node.
3. Select a vault and click **Remove Vault** on the ribbon. Alternatively, you can right-click the necessary vault in the working area and choose **Remove Vault**.
4. In the displayed dialog box, click **OK** to confirm deletion.

Moving Tapes to Vault

Recorded tapes can be moved to a vault. In the vault, you can view tape information, including data written to tape and data retention settings.

If the tapes are placed to a vault, you can still view them under the **Offline** node or in the media pool.

You can move the tapes manually or instruct the relevant media pool to place the tapes to a vault automatically.

In This Section

- Moving tapes to vault manually
- Automatic vaulting for media pools
- Automatic vaulting for GFS media pools

Moving Tapes to Vault Manually

You can move only tapes that contain data and that are offline. To move the tapes manually:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under the **Media Pools** or under the **Libraries** > *LibraryName* node > **Media** > **Offline**.
3. In the working area, right-click the tapes you want to move and select **Move to Vault**. Choose the vault you need.

Automatic Vaulting for Media Pools

To set the automatic placing of tapes to a vault, you need to configure the media pool where the target tapes are recorded. To configure a media pool:

1. Open the **Tape Infrastructure** view
2. Navigate to the **Media Pools** node
3. Select the media pool you want to customize.
4. Click **Edit Media Pool** on the ribbon. Alternatively, you can right-click the media pool and select **Properties**.
5. Go to the **Retention** step of the **Edit Media Pool** wizard. In the **Offline media tracking** field, select **Move all offline tapes from this media pool into the following media vault**. Select the needed vault from the drop-down list.

Note that moving tapes to a vault automatically will not be available if you select **Do not protect data (cyclically overwrite tapes as required)** option as a data retention policy.

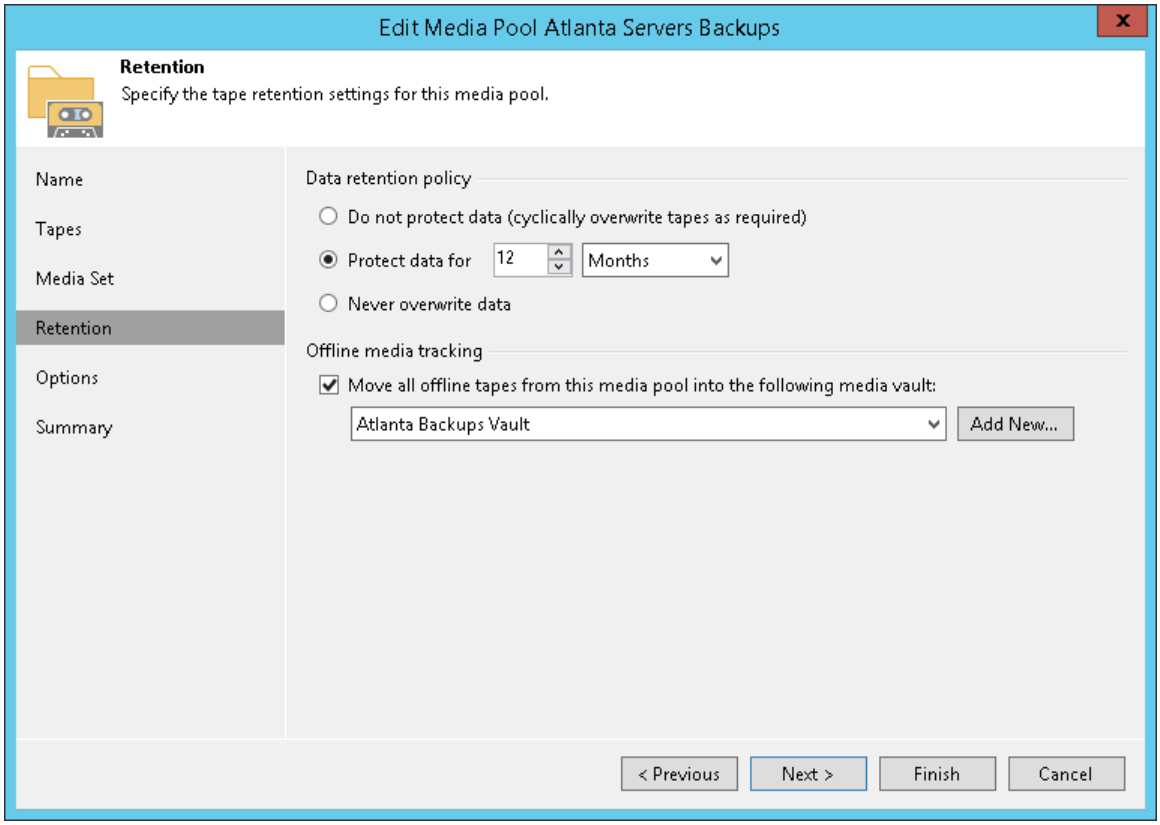

TIP

If you have not previously created a vault, you can click the **Add New** button and create a new vault without closing the job wizard. For more information, see Creating Vaults.

Automatic Vaulting for GFS Media Pools

Vaulting of GFS tapes is set for media sets individually. To change a media set configuration:

1. Open the **Tape Infrastructure** view.
2. Navigate to the **Media Pools** node. Select the GFS media pool you want to customize.
3. Click **Edit Media Pool** on the ribbon. Alternatively, you can right-click the media pool and select **Properties**.

4. Go to the **GFS Media Set** step of the **Edit Media Pool** wizard and click **Advanced**. Select a tab with the needed media set and select **Move all offline tapes into the following media vault**. Select the needed vault from the drop-down list.

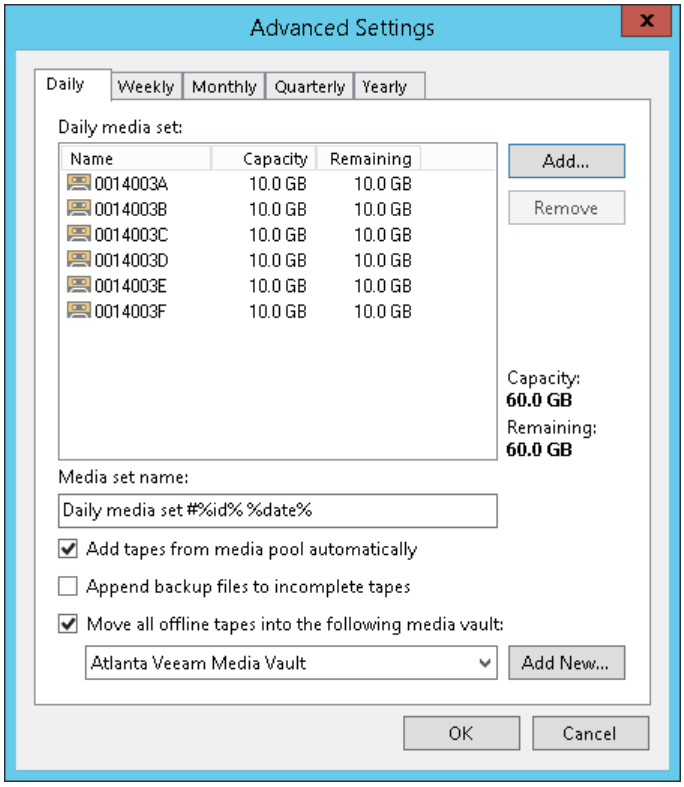

TIP

If you have not previously created a vault, you can click the **Add New** button and create a new vault without closing the job wizard. For more information, see Creating Vaults.

Managing Tapes in Vault

When you have tapes placed to vaults, you can move them to another vault or remove from vault. When you remove tapes from a vault, they are still available in the **Offline** list or in their media pool.

To move tapes to another vault:

1. Open the **Tape Infrastructure** view and select the **Vaults** node. Select the vault where the needed tapes are stored.
2. In the working area, right-click the tapes you want to move and select **Move to Vault**. Select the target vault.

To remove tape from vault:

1. Open the **Tape Infrastructure** view and select the **Vaults** node. Select the vault you need.
2. In the working area, right-click the tapes you want to remove and select **Remove from Vault**.

The tapes are removed from vault automatically in the following situations:

- A tape comes online. It will re-appear in the vault when it goes offline if it was not erased or overwritten.
- A tape is removed from catalog.
- A tape is marked as free.

Tape Protection

Tape protection is a software option that prohibits overwriting, erasing or appending data to protected tape. You can protect individual tapes that, for example, contain particularly valuable data.

Tape protection overrides retention settings of the media pool. The retention of other tapes in the media pool is not modified. You can remove the protection at any time. The tape will return to the retention period set by the media pool. If the media pool retention expired during the time the tape was protected, the tape will be queued for overwriting.

You can set protection for any tape, online or offline, that contains data. If the offline tape is in a media vault, the media vault will update the **Protected** status for this tape automatically.

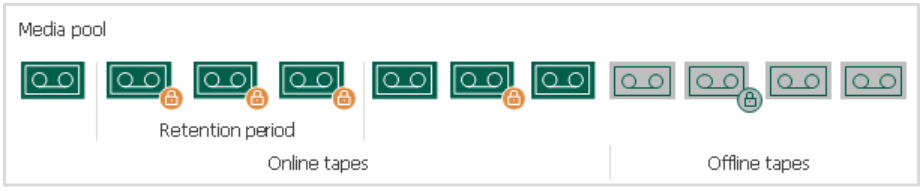

When the tapes are protected, the following operations are prohibited for them:

- Appending data to tapes
- Erasing tapes
- Marking tapes as free
- Removing tapes from catalog.

To perform these operations, you need to switch the protection off first.

Protecting Tapes

If you have tapes for which you need to change the retention to 'never overwrite', you can protect them. Protection option sets lifelong retention period for the selected tapes overriding the retention settings of the media pool.

You can set protection for both online or offline tapes that contain data.

> **IMPORTANT**
>
> You can protect only tapes that contain data. You cannot protect free (empty) tapes.

Switching On Protection

To enable protection:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under the **Media Pools** or under the **Libraries >** *LibraryName* node **> Media > Online**.
3. Select tapes you want to protect and click **Protect** on the ribbon.

   Alternatively, you can right-click selected tapes and choose **Protect** from the shortcut menu.

If the tape is a part of a media set, Veeam Backup & Replication will prompt you to protect the other tapes in this media set.

If the tapes are protected, they cannot be erased or marked as free by Veeam Backup & Replication or manually. To erase such tape, you need to switch the protection off first.

Switching Off Protection

You can switch off the protection at any time. The retention will return to the value set for the media pool.

Tape Data Retention

Data retention period is a period of time when data written to tapes is protected from overwriting.

Retention period is set by user for media pool and is applied to all tapes in this media pool. To set the retention policy, you can choose between the following options:

- Never overwrite data
- Define a particular time period to protect data
- Not to protect data at all.

During the retention period, Veeam Backup & Replication will not overwrite data on tape. If a tape contains several backup sets, it will expire when the backup set with the longest retention period expires.

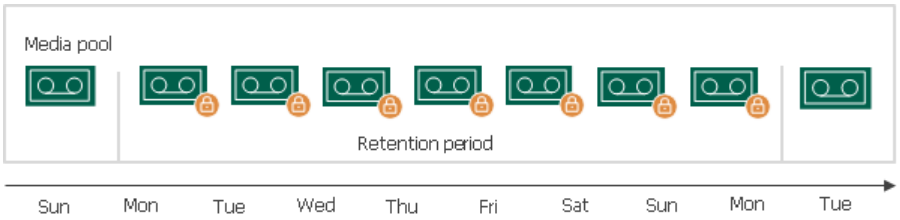

You can change the retention at any time. When you change the retention policy, you can select if this modification works for tapes that will be written after you apply this change or also to tapes that are already recorded. In the latter case, the retention settings of the recorded tapes will change immediately and will be applied to tapes that are both online and offline.

> **NOTE**
>
> If you choose to set a shorter retention period and apply the retention settings to all tapes, some of your tapes may immediately become outdated. Such tapes will be queued for overwriting. Be careful when applying a new retention policy in order not to lose any data you need, or use the protection option. For more information, see Tape Protection.

Tape Retention and Disk Retention

For forward incremental and reverse incremental backup chains, it is strongly recommended that you set the retention period for the tape archive at least twice longer than the retention period for source backups on disk. The tape jobs analyze the existing tape archives and synchronize them with disk backups. If you have some restore points on disk and the tape archive misses them (for example, if the media pool retention allowed overwriting of these tapes), the tape job will re-write all missing restore points. For more information, see the Veeam Backup & Replication User Guide, section *Retention Policy*.

For example, there are 14 backup files in the backup repository that are kept for 14 days. The backup to tape job archives files once a week. The retention policy for the media pool is set to 7 days. In this case, Veeam Backup & Replication will first write 14 backup files from the backup repository to tape. After a 7-day interval, Veeam Backup & Replication will start recording the whole set of backup files from the backup repository to tape anew, overwriting backup files on tape with their copies from the backup repository.

When the source job produces a forever forward incremental backup chain, you can set any retention period for the tape archive. To back up such chains, the virtual full mechanism is enabled automatically. Keep in mind that the virtual full backup must be always newer than the full backup on the disk. Otherwise, the tape job will copy not only the full backup from the disk, but also the whole chain of incremental backups. That may result in excessive use of additional tapes and significant increase in the tape job processing time. For more information, see Virtual Full Backup.

Managing Outdated Tapes

Veeam Backup & Replication will place outdated tapes to the media pools (the same media set for GFS media pools) where the tapes were originally written.

The tapes containing outdated data are handled in the following way:

- **If the expired tape is online**, it will be overwritten next time a tape job requires a free tape. The expired tape can be used only by the same media pool (or the same media set for GFS media pools) unless you erase the tape manually or move it to another media pool. The tapes are rewritten by the FIFO method.
- **If the expired tape is offline**, you can re-load it back to the library. Veeam Backup & Replication will place tapes to the media pools where the tapes were originally written. You can move tapes to another media pool or erase the tape manually.

You can erase tapes or move tapes to another media pool manually. For more information, see Moving Tapes to Another Media Pool.

Virtual Full Backup

The virtual full method allows you to transform forever forward incremental backup chains created by source backup jobs to save them to tapes. The forever forward incremental chain always keeps on disk one full backup followed by a fixed number of increments. The full backup is constantly rebuilt: as new increments appear, the older ones are injected into the full.

Unlike disk backups, tape archives are static: tape jobs cannot rebuild backups once they are written to tape. Also, the standard backup to tape scheme (archiving new restore points during each tape session) cannot be used: the tape archive would have one full backup and an endless chain of increments all of which would be required for restore. To adapt the forever forward incremental chains to tapes, Veeam Backup & Replication uses the virtual full.

The virtual full mechanism creates a periodic synthesized full backup on tape. The periodic fulls split the forever incremental backup chain into shorter series of files that can be effectively stored to tapes. Each series contains one synthesized full backup and a set of increments. Such series are convenient for restore: you will need to load to the tape device only those tapes that are part of one series.

The virtual full does not require additional repository disk space: it is synthesized directly on tape on the fly, when the tape job runs. To build such full backup, Veeam Backup & Replication uses backup files that are already stored in the backup repository.

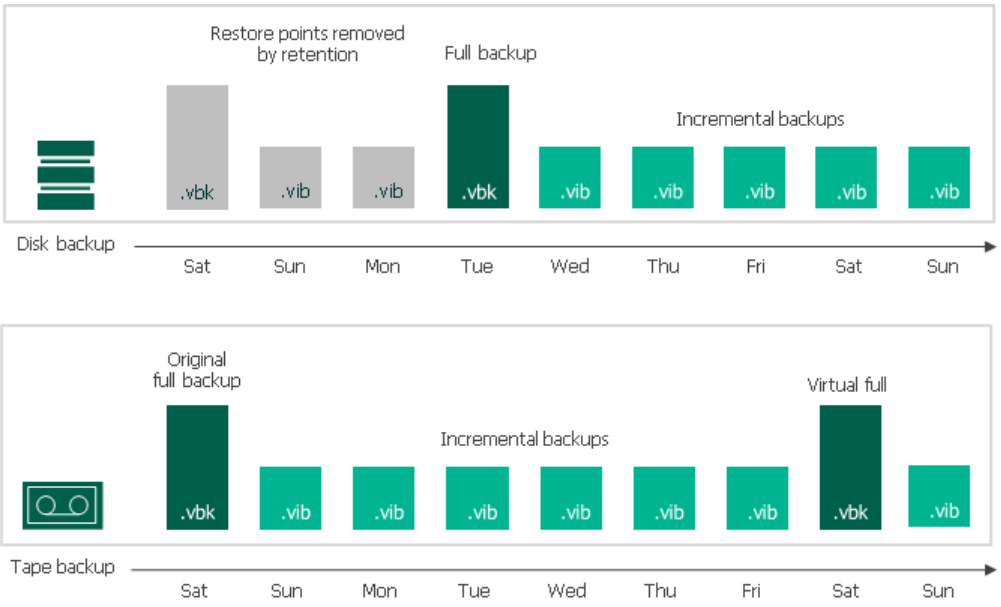

The virtual full is enabled automatically if the source job has no scheduled full backups. You can schedule the virtual full on a specific day (or days) of week. This day indicates the day for which you want to synthesize full backup of the source machines.

The virtual full does not necessarily need to be scheduled on the day when the tape job runs. When the tape job starts, it synthesizes the source machines state as they were on the chosen day. For example, you can set Friday as the virtual full day, and schedule the tape job to run on Saturday. The tape job will not run on Friday. On Saturday, the tape job will copy blocks of data to reconstruct the state of the machines as they were on Friday (or, if no backup is available for Friday, the closest day preceding Friday).

You can customize the virtual full schedule according to your needs, although you cannot disable the virtual full if you back up the forever forward incremental backup chain. Also, you cannot create more than one virtual full a day.

NOTE

The source job configuration must not have scheduled synthetic or active full backups. If it does, the virtual full will be disabled even if configured in the tape job.

How Virtual Full Backup Works

When the tape job archives the forever incremental backup chain, it copies the full backup from disk only once. After that, it copies new increments and creates virtual fulls regularly, or, if archiving of increments is disabled, creates only virtual fulls.

To create a virtual full backup, Veeam Backup & Replication uses a small temporary file of the VSB (Veeam Synthetic Backup) format. The VSB file does not contain backup data; it contains pointers to data blocks inside files of the backup chain on disk that are required to build a full backup as for the needed day. Using these pointers, the tape job detects these blocks and writes them to tape.

To create a virtual full backup, the tape job performs the following operations:

1. On the day the tape job runs, it creates a VSB file and stores it in the backup repository, next to restore points in the backup chain.
2. According to the VSB file pointers, the tape job detects what backup chain and what data blocks on disk are required to synthesize a full backup, and writes these blocks to tape as a full backup file (.vbk).
3. At the end of the tape job session, the VSB file is removed from the backup repository.

For example:

- The source backup job creates an increment every day.
- The tape job is scheduled on Saturday with a virtual full on Friday.

On Saturday, the tape job picks data blocks from the full backup on Tuesday and the subsequent increments up to Friday. Using these blocks, the tape job synthesizes a virtual full that represents the machine state as of Friday.

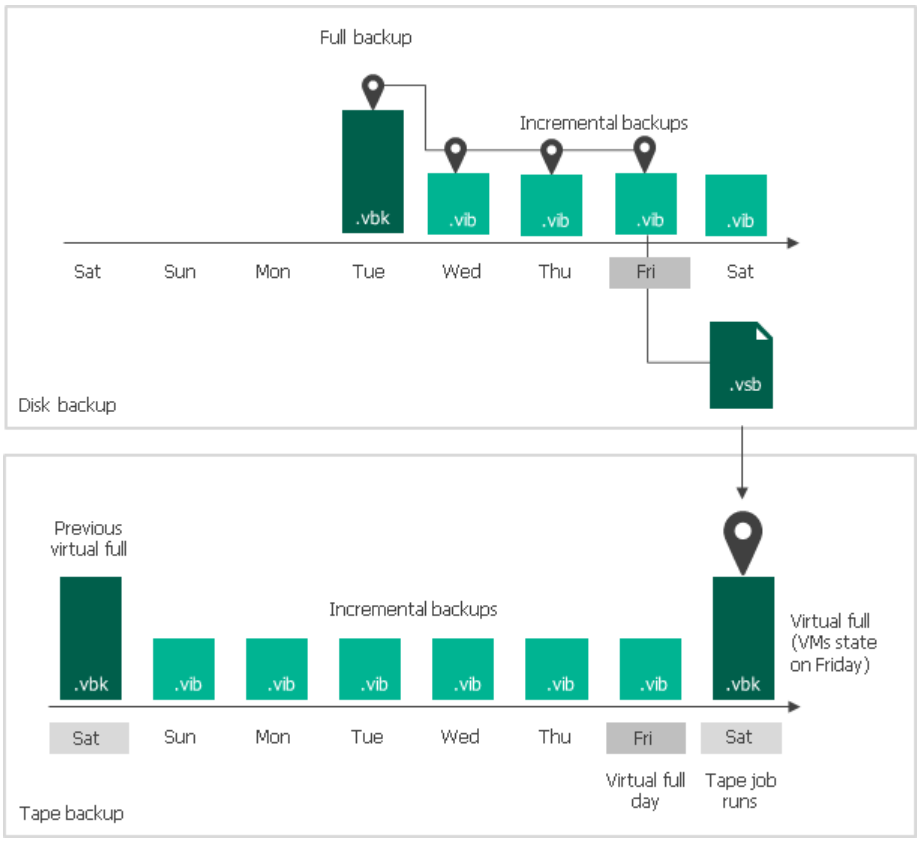

> **NOTE**
>
> If the source backup chain contains a full backup that was created on the virtual full day (for example, an active full backup created manually), the tape job will copy it and will not create a virtual full.

Virtual Full Intervals

When the tape job needs to create a new virtual full, it analyzes the tape archive and detects the date of the last virtual full on tape. To create a new virtual full, the tape job needs restore points that were created in the following interval:

- The start of the interval is midnight after the last virtual full was created (the end of tape writing session).
- The end of the interval is midnight after the virtual full day.

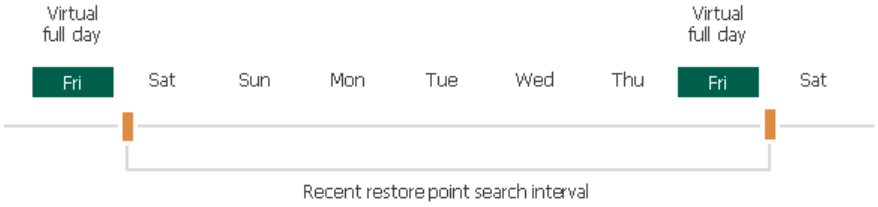

To synthesize a new virtual full, the tape job needs at least one increment created in this interval. If there are many increments, the tape job builds a virtual full up to the most recent increment within the interval.

Missing Source Backups

The source chain may not contain an increment on the virtual full day. In this case, the tape job uses a backup that is closest to the virtual full day within the virtual full interval. For example, if the source backup closest to the virtual full day was on Tuesday, the tape job will synthesize a virtual full that represents the machine state as of Tuesday.

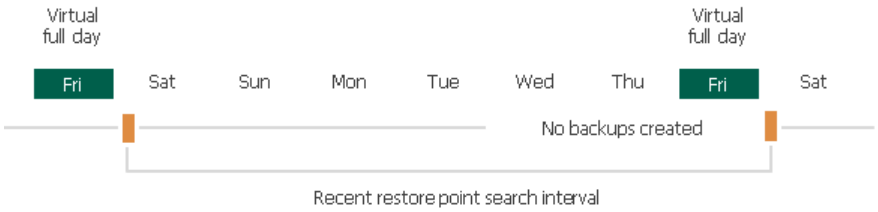

The source chain may have missing backups because of the following reasons:

- The source job does not run on the virtual full day.
- The tape job has multiple source jobs scheduled on different days. Some of the source jobs does not create a backup on the virtual full day.
- The source job fails to produce increments for several days, for example, because of a hardware breakdown, network inaccessibility or if the source job was disabled by the user for some time.
- Some machines in the source job have corrupted backups. The tape job will use the last valid backups for these machines and create the virtual full up to the latest valid restore point.

No Backups in the Virtual Full Interval

The source job may not produce any backups after the previous full backup day. In this case, the tape job will not create a virtual full for this period.

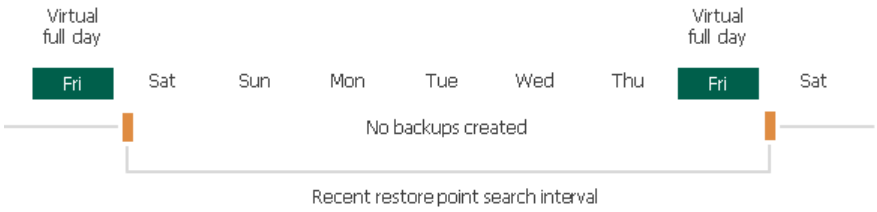

IMPORTANT

The virtual full cannot be forced before its scheduled day. For example, if you start the tape job manually off-schedule, it will not create the virtual full if the virtual full day is in future.

Tape Parallel Processing

You can use multiple drives simultaneously for writing data to tape. This option is useful if you have a lot of tape jobs running at the same time or you have a lot of data that must be written to tape in a limited backup window. You can configure multiple drives in the following ways:

- Use multiple drives of one tape library
- Manage multiple tape libraries and use several or all drives across all managed libraries.

Parallel processing is enabled in optional media pool settings. To use drives of multiple libraries, you must enable the parallel processing mode for the libraries that are managed by the media pool.

To process the tape data parallelly, you can split the data across drives in two ways:

- Parallel processing for tape jobs.
- Parallel processing for source chains of one (or more) tape jobs.

Tape jobs may simultaneously use less drives than configured in optional media pool settings. That may happen under one of the following conditions:

- The number of tapes available is less than the number of drives assigned for this media pool.
- The number of drives available is less than the number of drives assigned for this media pool.

Processing Tape Jobs Simultaneously

When you process tape jobs parallelly, the media pool assigns a drive to each running tape job. For example, if you set three drives as the maximum, you can process up to three tape jobs at the same time. If you have more jobs running at the same time, they are queued. When one of the jobs finishes and releases its drive, the first queued job takes the drive.

To enable this mode, select the **Enable parallel processing for tape jobs using this media pool** check box in optional media pool settings. If this mode is enabled in media pool settings, for machine backup to tape jobs you can limit the number of drives the job can use, as described in Specify Archiving Options.

This option is available for backup to tape and file to tape jobs.

For example:

- You set the maximum number of drives to three.
- Four tape jobs start at the same time.

The tape jobs start, and jobs A, B and C occupy three drives to write data to tape. Tape job D is queued and waits. When one of the jobs finishes and releases its drive, tape job D takes the drive and starts writing data.

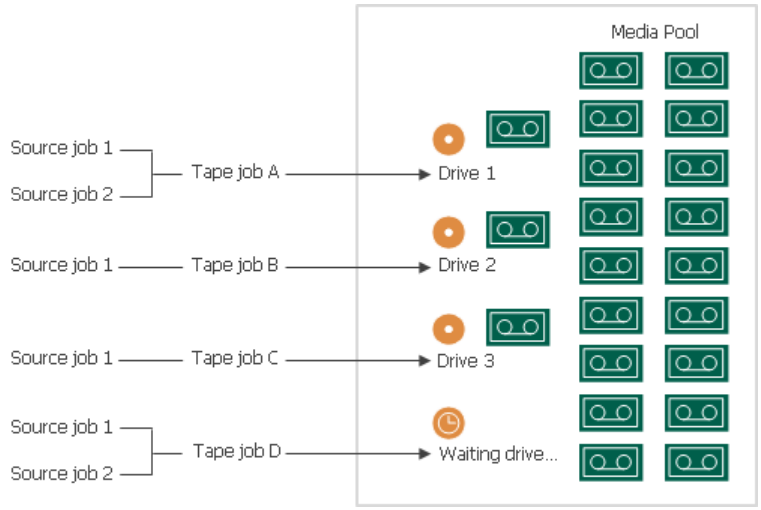

Processing Backup Chains Simultaneously

When you select processing backup chains parallelly, the media pool processes several source jobs simultaneously.

To enable this mode, select the **Enable parallel processing for tape jobs using this media pool** and **Enable parallel processing of backup chains within a single job** check boxes in optional media pool settings.

This option is available for backup to tape jobs only.

For example:

- You set the maximum number of drives to three.
- Tape job A has one source job; tape job B has three source jobs. Both tape jobs start at the same time.

Tape job A occupies drive 1. Tape job B occupies two other available drives and processes two source jobs. Tape job A finishes and releases drive 1, the third source job from tape job B occupies drive 1.

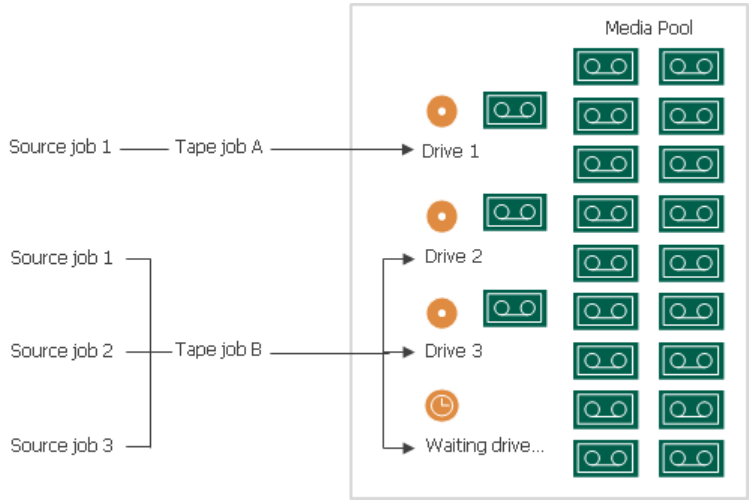

If a source job produces multiple per-machine backups, the media pool processes several per-machine backup chains simultaneously. For more information about per-machine backup files, see the Veeam Backup & Replication User Guide, section *Per-Machine Backup Files*.

For example:

- You set the maximum number of drives to three.
- Tape job A has one source job, the source job processes four VMs.

Tape job A starts and occupies all three available drives. It writes each separate VM of the source job by a separate drive. After tape job A finishes writing one of the current VMs and releases one of the three available drives, it occupies this drive with writing VM4. In other words, if a job with per-machine chains is configured to process more VMs than there are drives available, it occupies the entire media pool.

The order in which VMs are queued for processing is defined by the order of VMs within each source job, by the order of source jobs within each tape job and the schedule configured for running tape jobs.

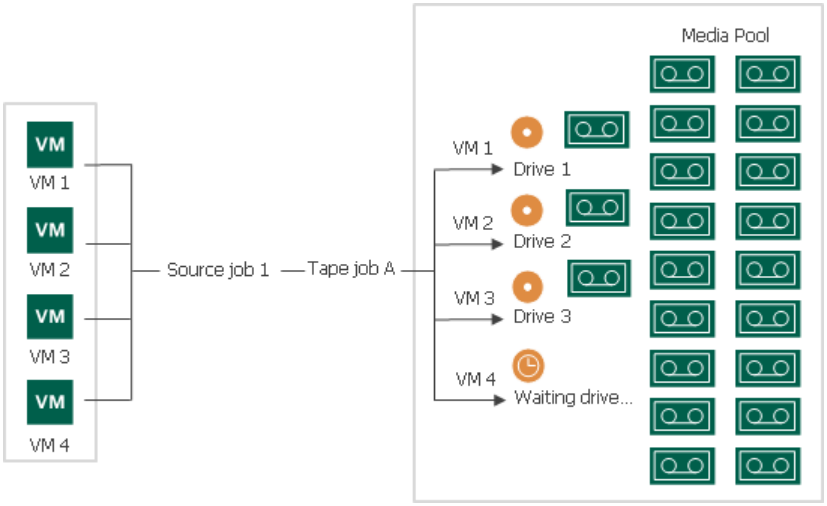

Similarly, if the source job is an agent job with the configured backup policy (the **Managed by agent** option is selected) and produces backups for multiple machines/protection groups, the media pool processes backups for several computers simultaneously. For more information about agent backup jobs, see the Veeam Agent Management Guide, section *Working with Veeam Agent Backup Jobs*.

Media Sets Created with Parallel Processing

With parallel processing, a separate media set is opened per each drive used.

> **NOTE**
>
> The media pool must have an available tape for each drive.

The media sets are opened according to media set options configured for the media pool:

- Parallel processing with Do not create, always continue using current media set option
- Parallel processing with Create new media set for every backup session option
- Parallel processing with Daily at option

You can edit the media pool settings at any time, and increase or reduce the maximum number of drives:

- If you reduce the number of drives, the media pool will use fewer drives for the next writing session. The media sets that become excessive will be closed.
- If you increase the number of drives, the media pool will open new media sets per each added drive. Further, the media pool will use these media sets according to rules described above.

> **TIP**
>
> When multiple media sets open simultaneously, they may have identical sequence number and time of creation. To distinguish between the media sets easily, use the *%id%* variable in the media set name. This variable is added to the media set name by default.

Parallel Processing with Do Not Create, Always Continue Using Current Media Set Option

When the media pool is set to the **Do not create, always continue using current media set** option, a media set is opened per each drive during the first tape session. The next time a tape job starts, it chooses a tape that has most free space, and appends data to it.

For example, you set the maximum number of drives to 3. The tape jobs that are first to run start 3 media sets. The following tape sessions continue these 3 media sets appending data to them.

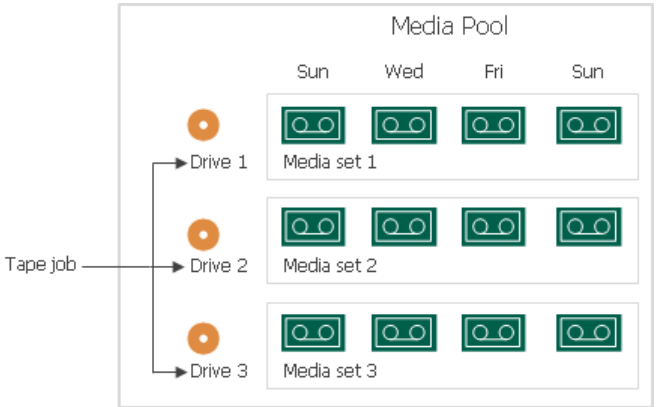

Parallel Processing with Create New Media Set for Every Backup Session Option

When the media pool is set to the **Create new media set for every backup session** option, a media set is opened per each drive for every writing session. When the jobs run the next time, they open several media sets anew.

Each media set uses a free tape. Note that this configuration is the most tape-consumptive.

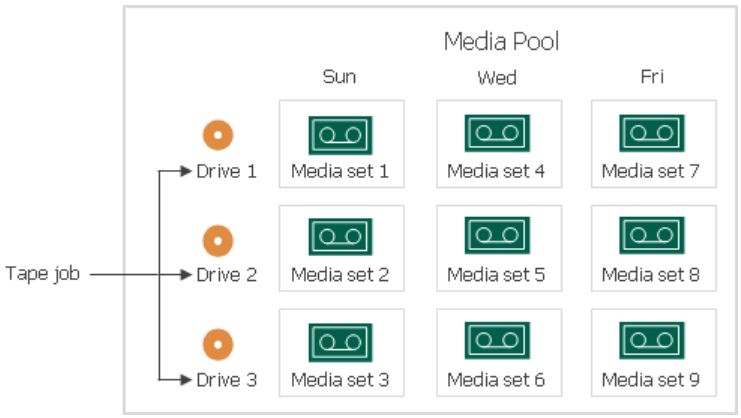

Parallel Processing with Daily at Option

When the media pool is set to **Daily at** option, several media sets (per each drive) are opened on the scheduled day. If tape jobs run twice or more before the scheduled day, they continue these media sets — in this case, the jobs use tapes that have most free space first.

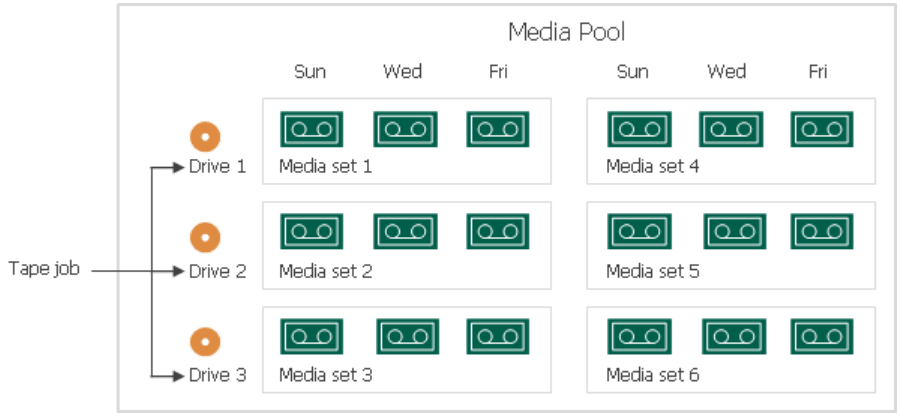

Machines Backup to Tape

To back up data to tape, you need to create and run tape jobs. The backup to tape job is a dedicated job that archives to tape Veeam backups that were produced by Veeam backup jobs. Mind that you cannot directly backup virtual machines to tapes.

> **NOTE**
>
> Support of the machine backup to tape is available in the Enterprise and Enterprise *Plus* editions of Veeam Backup & Replication. For details, see Veeam Editions Comparison.

You can archive the following data to tape:

- VM backups
- Veeam Agent backups
- Physical machine backups
- Backup repositories
- [For Veeam Cloud Connect] Tenants.

When a backup to tape job runs, it does not create new backups: it locates already existing backups and copies them from backup repository to tape. You need to set the source of the tape job: jobs and/or backup repositories.

Jobs as Source

The following jobs can be source for tape jobs:

- VMware/Hyper-V/Nutanix AHV/RHV backup job
- VMware/Hyper-V/Nutanix AHV/RHV backup copy job
- Windows/Linux/Unix/Mac Veeam Agent backup job configured in Veeam Agent operating in the standalone mode
- Windows/Linux/Unix/Mac Veeam Agent standalone backup copy job
- Windows/Linux/Unix/Mac Veeam Agent backup job configured in Veeam Backup & Replication
    a. Veeam Agent backup job managed by the backup server
    b. Veeam Agent backup job managed by Veeam Agent (backup policy)
- Windows/Linux/Unix/Mac Veeam Agent backup copy job.
- Kasten K10 policy*
- Backup copy job for Kasten K10 policy*

* This version is supported starting from Veeam Backup & Replication <%BLDVER%> with Cumulative Patch P20211211.

> **NOTE**
>
> Backup to tape jobs can process only agent backup jobs that are targeted to a Veeam backup repository.

When the tape job starts on its schedule, it picks the restore points that were produced by the source jobs in period since the last tape job run. If you change the configuration of the source jobs, the tape job is updated automatically: it adds new machines to the list of machines to archive or stops archiving machines that were removed from source jobs.

The tape job will process its source jobs in alphabetical order.

Source Jobs Backup Methods

The source jobs may use any backup method:

- Forever forward incremental backup method

  To back up the forever forward incremental chains to tape, the tape job uses the virtual full. The virtual full creates a synthetic full backup on tape regularly (for example, once a week) and splits the chain into short series of tapes which is more convenient for restore. For more information, see Virtual Full Backup.

  If the source job is a backup copy job in the periodic copy mode, keep in mind that the last restore point of the backup copy job stays active until the next restore point is created. The tape job does not copy such active points, because they may be updated. For this reason, the backup chain on tape will be always one restore point shorter than on disk.

- Forward incremental backup method

  When the tape job backs up the forward incremental chain to tape, it creates a copy of the disk backup chain.

- Reverse incremental backup method

  The last backup in the reverse incremental backup chain is always the full backup. If the source backup chain is reverse incremental, the tape job will copy the full backup each time the tape job runs. The increments are skipped.

Backup Repositories as Source

When you add a repository as source to tape job, the tape job constantly scans the selected repository (or repositories) and writes the newly created backups to tape. The tape job monitors the selected repository in a background mode. You can set explicit backup windows for the tape job. In this case, the tape job will start on the set time and archive all new restore points that were created in period since the last job run.

If you create or remove backup jobs that use this repository, or if you change the configuration of such backup jobs, you do not need to reconfigure the tape job that archives the repository.

Tenants as Source

This option is available for Veeam Cloud Connect service providers. Service providers can back up tenants backups from cloud repositories to tape. As the source to the backup to tape job, you can add tenants or cloud repositories. For more information, see the Veeam Cloud Connect guide.

For tenants backups, only GFS media pools are used. For more information, see GFS Media Pools.

Mixed Jobs

To one tape job, you can link an unlimited number of sources. You can mix source jobs of different type: backup and backup copy, and of different platform (VMware, Hyper-V, Windows Agent or Linux Agent). You can add jobs and repositories as source to the same tape job.

Also, you can add an unlimited number of tenants to one tape job. However, you cannot mix tenants with backup jobs and backup repositories.

> **IMPORTANT**
>
> The tape job looks only for the Veeam backups that are produced by backup jobs running on your console. Other files will be skipped.

Note that to back up files, you need to configure file to tape job. For more information, see File Backup to Tape.

Linking Source Jobs

You can add source jobs to tape jobs at any moment: when you create a tape job, or later. Adding source jobs is not obligatory when you create a tape job: you can create an "empty" job and use it as a secondary destination target. When you link jobs, the tape job processes them in the same way as the jobs added with the Tape Job Wizard. For more information, see Linking Backup Jobs to Backup to Tape Jobs.

How Machines Backup to Tape Works

When Veeam Backup & Replication executes a backup to tape job (started manually or on schedule), it performs the following operations:

1. The backup to tape job addresses the Backup Catalog in the Veeam Backup & Replication database to detect backups that match the job criteria.
2. The files are queued for archiving.
3. Veeam Backup & Replication connects to Veeam Data Movers and starts the data transfer process. The source Veeam Data Mover retrieves data from the backup repository and target Veeam Data Mover sends data to tape.
4. The tape job addresses the media pool that is set for this job as target. The media pool allots tapes for writing data according to the following configuration options:
   - Tapes consumption
   - Media sets
   - Tape retention.
5. While tape recording is performed, Veeam Backup service updates data in the Backup Catalog and Tape Catalog in Veeam Backup database. The Veeam Backup console displays refreshed information about backups archived to tape and shows job statistics.

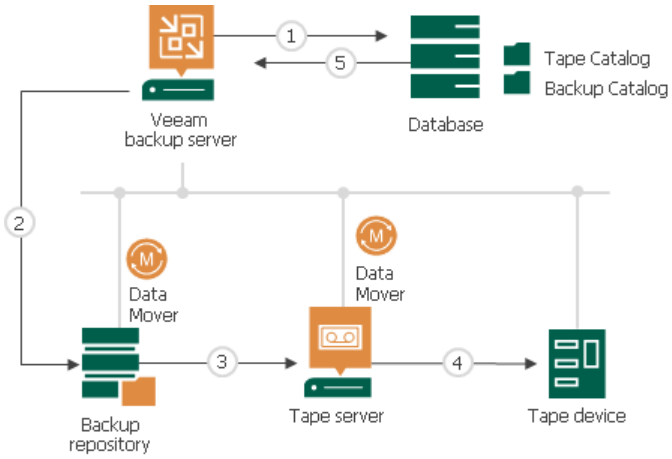

Before You Begin

Before you configure a backup to tape job, complete the following prerequisites:

- You must have Veeam Backup & Replication Enterprise license or higher installed on the Veeam backup server.
- [For Veeam Cloud Connect service providers] To back up tenants to tape, you must have Veeam Cloud Connect service provider license installed on the Veeam backup server.
- You must check configuration of source backup job(s). Tape jobs have the following requirements for the restore points to keep setting:
  - A source job with **forever forward incremental** chain must keep not less than 3 restore points on disk.
  - A source job with **forward incremental** chain must keep not less than 3 restore points on disk.
  - A source **backup copy job in the immediate copy mode** must keep not less than 3 restore points on disk.
  - A source **backup copy job in the periodic copy mode** must keep not less than 4 restore points on disk.
  - A **reverse incremental** chain has no required minimum.
- You must configure one or more media pools with the necessary media set and retention settings.
- You must load tapes to the tape device and configure the target media pool so that it has access to them. If the media pool has no available tape, the tape job will wait for 72 hours and then terminate.

Mind the following limitations:

- The backup to tape job processes only VBK (full backups) and VIB files (forward incremental backups).
- If you back up to tape a reverse incremental chain, the tape job will always copy the full backup. Reverse incremental backups (VRB) are skipped from processing.
- Microsoft SQL Server log files (VLB) are skipped from processing.
- [For Veeam Cloud Connect service providers] The backup to tape job that backs up tenants to tape will not process backups created with previous versions of Veeam Backup & Replication or Veeam Agents. To avoid this, you must upgrade the tenants' backup server or agent machines to the last version. After the upgrade, the tenants' jobs must run at least once.
- If a job is unable to complete within 21 days period, it will be stopped with the 'Failed' status.

Creating Backup to Tape Jobs

To archive backups to tape, you should create a backup to tape job using the **New Backup to Tape Job** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

With this wizard, you can create a backup to tape job or a GFS job.

To create a backup to tape job, follow the next steps:

Step 1. Launch New Backup to Tape Job Wizard

To run the **New Backup to Tape Job** wizard, do either of the following:

- On the **Home** tab, click **Tape Job** and select **Backups**.
- Open the **Home** view, right-click the **Jobs** node and select **Tape Job** > **Backups to Tape**.

Step 2. Specify Job Name and Description

At the **Name** step of the wizard, you should define basic settings for the backup to tape job.

1. In the **Name** field, enter a name for the job.
2. In the **Description** field, enter a description of the job. The default description contains information about the user who created the job, date and time when the job was created.

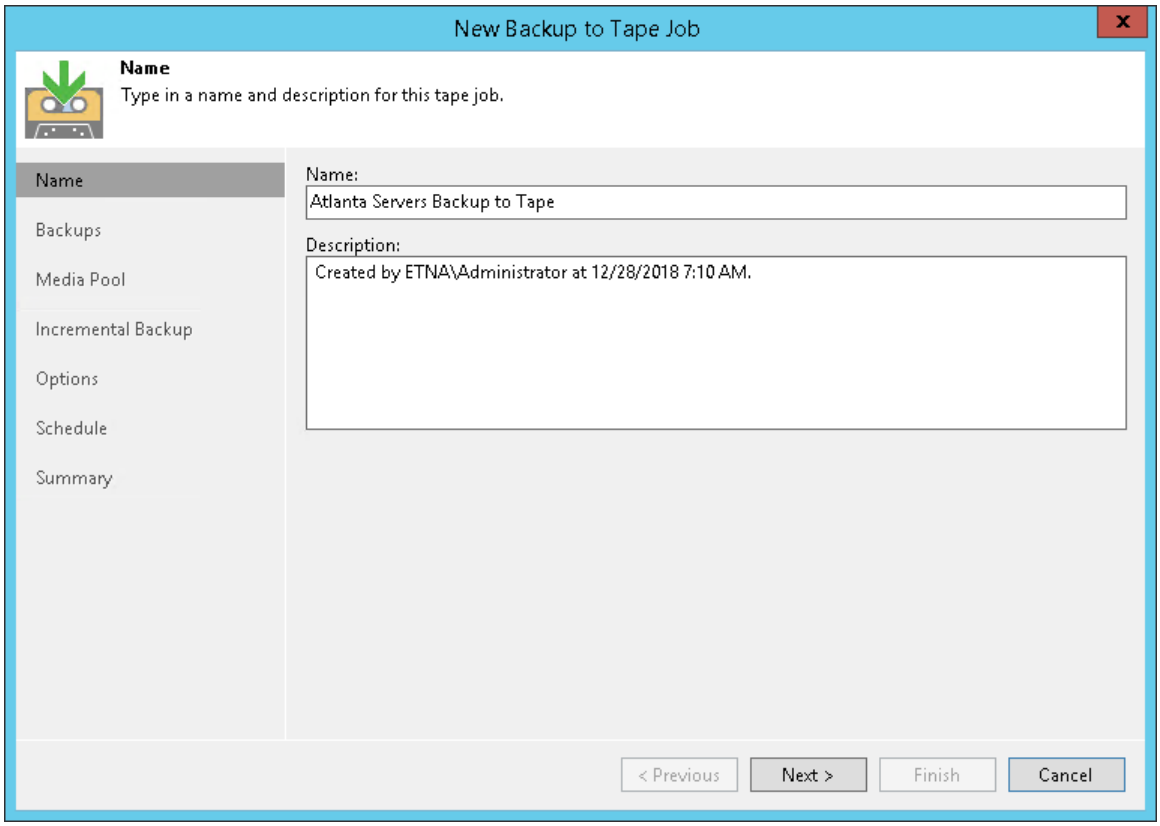

Step 3. Choose Backups to Archive

At the **Backups** step of the wizard, select backups that you want to write to tape with the tape job.

Click **Add** and select the necessary backups. You can choose backups from the following sources:

- **Backup jobs**. You can select backups from available backup, backup copy or agent backup jobs. When a backup to tape job runs, Veeam Backup & Replication searches and archives restore points from the backups created by the selected backup jobs.
- **Backup repositories**. Using this option, you can select whole backup repositories. When a backup to tape job runs, Veeam Backup & Replication searches and archives restore points from all backups stored in the chosen backup repositories.

  If you choose to archive data from backup repositories, the backup to tape job will process only the backups that were created with backup jobs configured on this Veeam backup server. Imported backups and configuration backups will be skipped from processing.
- **Tenants**. This option is available for Veeam Cloud Connect service providers. With this option, you can select backups of your tenants or cloud repositories.

To learn more about sources to tape jobs, see Machines Backup to Tape.

To remove a source from the list, select it and click **Remove** on the right.

The sources of the backup to tape job are processed in order the displayed here. To move a source up or down in the list, use the **Up** and **Down** buttons on the right.

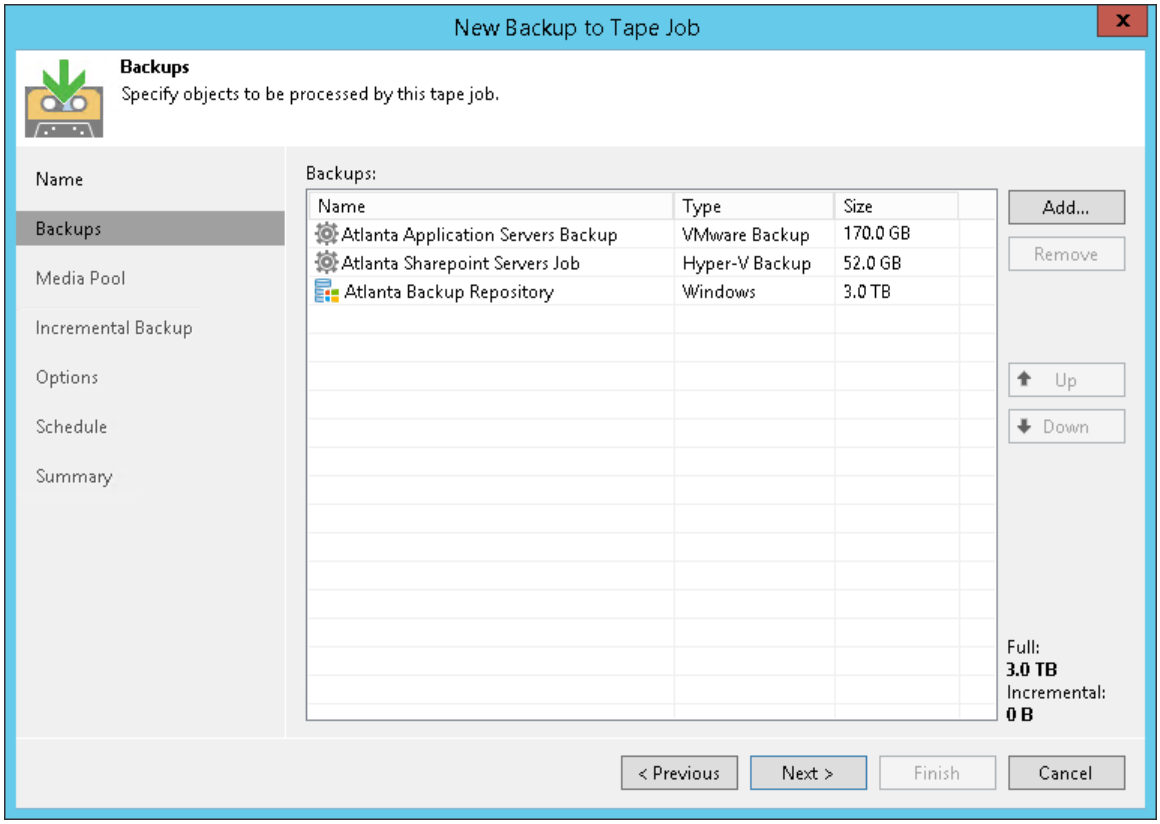

The total size of full backups added to the backup to tape job is displayed in the **Full** field. The total size of incremental backups added to the backup to tape job is displayed in the **Incremental** field.

> **TIP**
>
> This step is optional. You can create a backup to tape job without source. You can add the source later, or you can link this job to a backup job. For more information, see Linking Backup Jobs to Backup to Tape Jobs.

Selecting Backup Chains to Archive

On the first run, the tape job can copy all restore points or only the latest backup chain. This option is useful if the source backup job or the source backup repository has several backup chains. The backup chain here means a full backup file plus a set of incremental backup files following it.

To select what backup files you want to archive:

1. Click **Add** and select **Backup jobs or Backup repositories**.
2. Select the necessary backup job or repository you want to add to the tape job.
3. Veeam Backup & Replication will display a warning informing about backup chain options. Click **Latest** to archive only the latest backup chain to tape. This option lets you have only the latest restore points written to tape and use less tape for data archiving. Click **All** to archive all available backup files.

This option selects what backup chain will be copied on the first run of the job. You can instruct the job to copy only the latest backup chain on every tape job run. This option is available in the advanced settings of the job. For more information, see Advanced Settings.

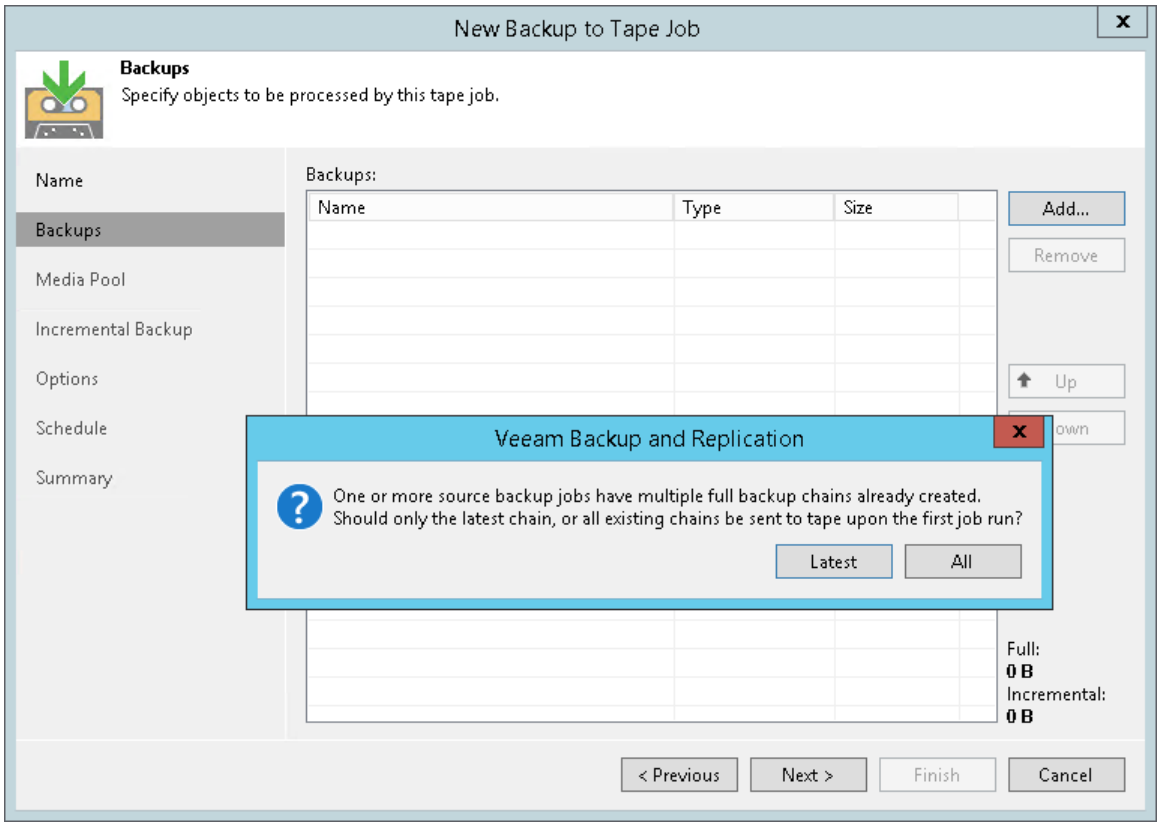

Step 4. Choose Media Pool for Full Backups

At the **Media Pool** step of the wizard, choose media pool for full backups.

1. From the **Media pool for full backups** list, choose a media pool that will be used for archiving full backup files.

   You can select a media pool or a GFS media pool. Depending on the selected type, the job schedule will change to regular schedule or GFS schedule accordingly.

   > **NOTE**
   >
   > For tenants backups, you can select only GFS media pools.

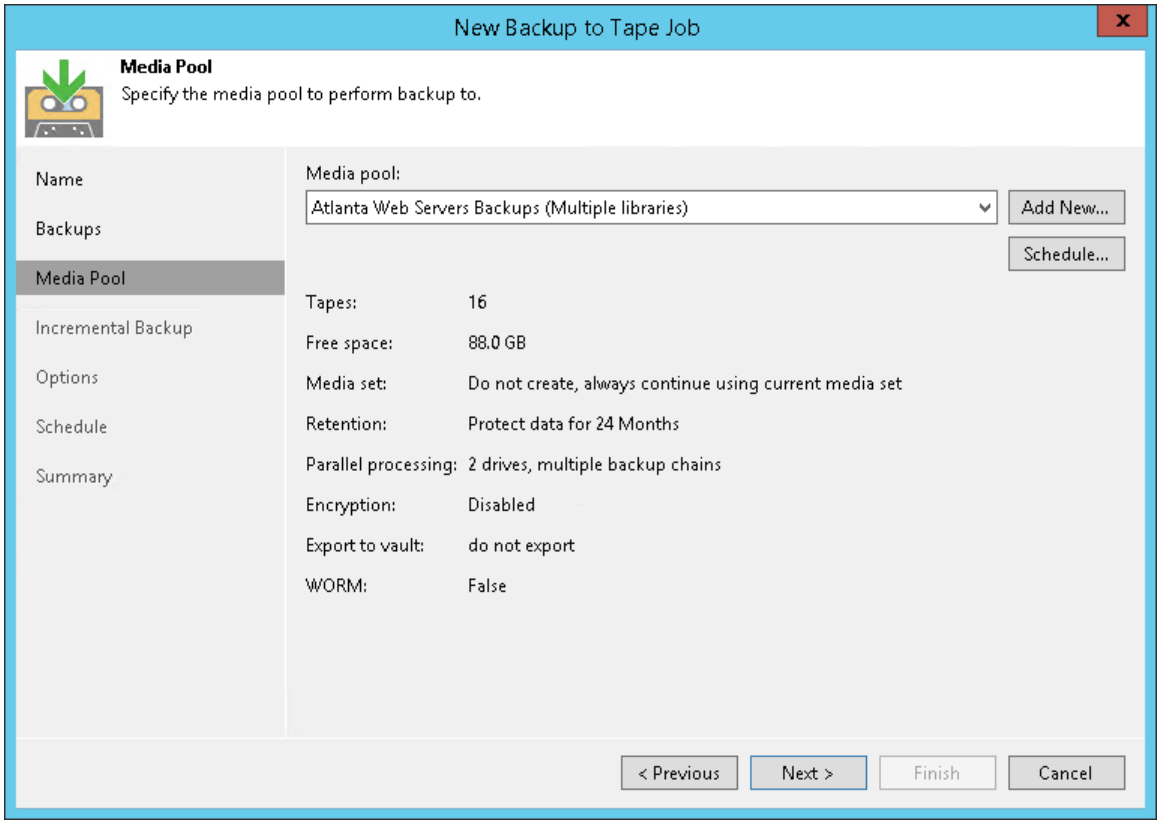

   > **TIP**
   >
   > If you have not previously created a media pool with the required settings, you can click the **Add New** button and create a new media pool or a GFS media pool without closing the job wizard. For more details, see Creating Media Pools or Creating GFS Media Pools.

2. You can configure schedule for virtual full backup. This option is available only if you selected a regular media pool.

   If the source job produces a forever incremental backup chain or is a backup copy job without GFS option enabled, Veeam Backup & Replication will periodically create a virtual full backup. You can configure the full backup schedule at this step of the wizard. The virtual full cannot be switched off; however, it is disabled automatically if the source job periodically creates active full or synthetic full backups.

   The virtual full does not depend on the job settings for incremental backups. If you enable the virtual full for the job, it will be created in any case, no matter whether you enable or do not enable incremental backups. For more information, see Virtual Full Backup.

NOTE

If the source job is a backup copy job with the GFS option enabled, Veeam Backup & Replication will not create a virtual full backup on tapes periodically, as this source job is considered forward incremental. For more information on methods of creating archive full backups, see the Veeam Backup & Replication User Guide, section *Methods for Archive Backups Creation*.

To configure the virtual full schedule, click **Schedule** and select the necessary scheduling options.

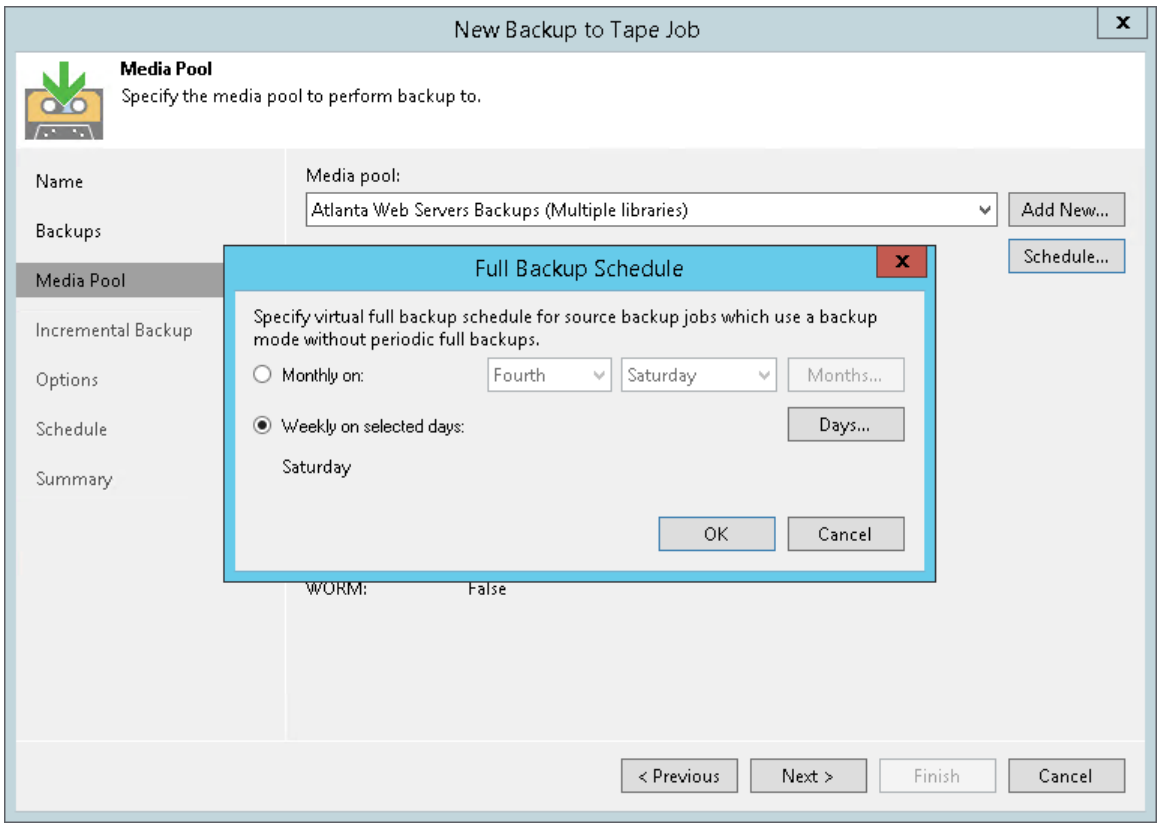

Step 5. Choose Media Pool for Incremental Backups

This step is available only if you selected a regular media pool as target for full backups.

At the **Incremental Backup** step of the wizard, you can enable or disable incremental backups processing and choose media pool for incremental backups.

If you want to back up incremental backups to tape, select the **Archive incremental backups to tape** check box. If this option is disabled, the backup to tape job will archive only VBK files and will skip VIB files from processing.

From the **Media pool for incremental backups** list, select a media pool that will be used for incremental backups.

Veeam Backup & Replication allows you to select different media pools for full backups and incremental backups. This can be required if you use different media set or retention settings for archiving full backups and increments.

You cannot select a GFS media pool to copy incremental backups to tape. Instead you can configure a daily media set in the GFS media pool to archive incremental points along with GFS points, combining normal and GFS functionality within a single job. For more details, see Creating GFS Media Pools.

> **TIP**
>
> If you have not previously created a media pool with the required settings, you can click the **Add New** button and create a new media pool without closing the job wizard. For more details, see Creating Media Pools.

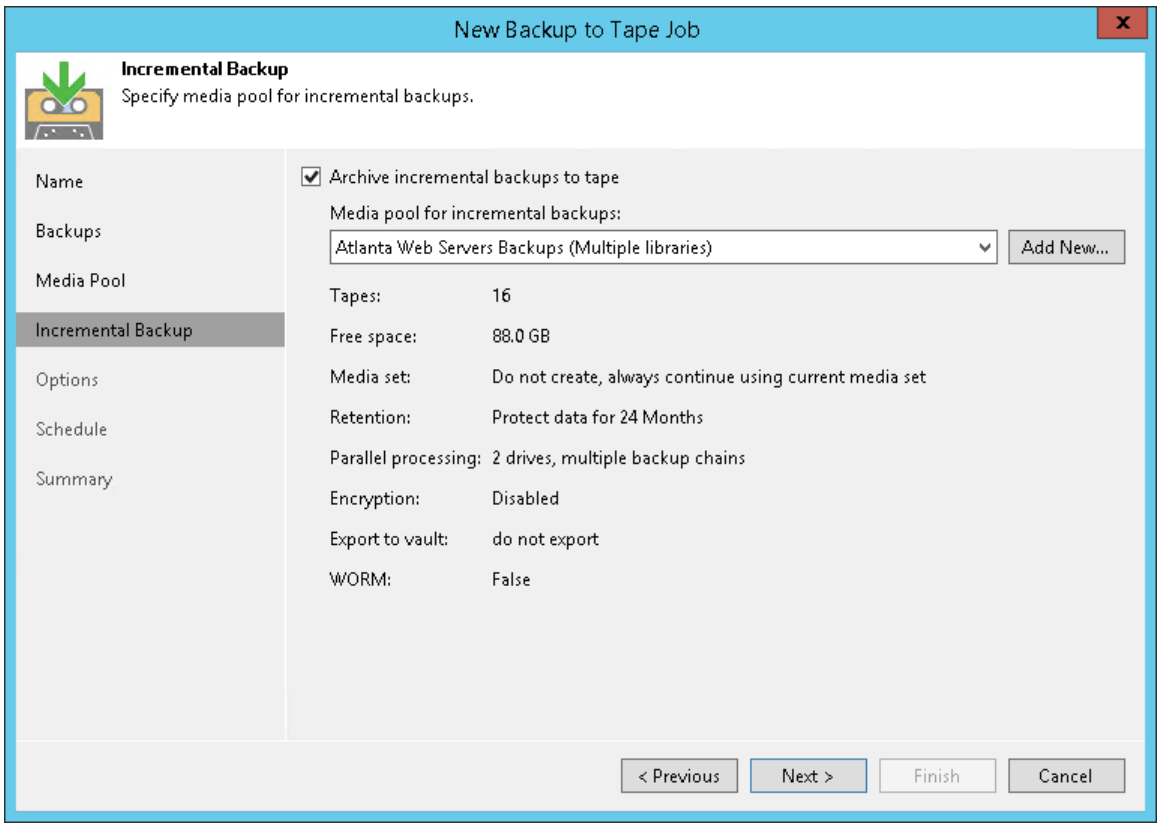

Step 6. Specify Archiving Options

At the **Options** step of the wizard, specify archiving and media automation options. The options are different for regular backup to tape jobs and GFS tape jobs:

- Options for Backup to Tape Job
- Options for GFS Tape Job

Options for Backup to Tape Job

This step of the wizard is available if you selected a regular media pool at the **Full Backup** step of the wizard.

In the **Media Automation** section, specify options for automatic operations after finishing the job:

- Select the **Eject media upon job completion** check box if the tape should be automatically ejected from the tape drive after the job successfully finishes. The ejected tapes are placed into a free tape device slot. Note that if the job started and failed, the tape will remain in the drive.

  This option does not prevent the tape job from appending data to this tape. If not configured otherwise in media pool settings, this tape will be placed into a drive on the next tape job run.

- Select the **Export current media set upon job completion** check box if you want to pull out the tapes with the current media set from the tape device, for example, to move to a storage location. The tape device will eject the tapes that belong to the current media set when the media set is closed.

  If you want to export tapes on specific days only (for example, every Saturday), click **Days** and schedule export on the necessary days.

  Note that with this option selected, a new media set is started after each export.

If the media pool, which is the target for the tape job, uses parallel processing, you can limit the number of drives to use for processing this tape job. To enable this, select the **Limit the number of drives this job can use to N tape drives** check box and specify the number of drives to use. For more information on how to enable parallel processing for the media pool, see Add Optional Media Pool Settings.

Click **Advanced** to configure the advanced job options.

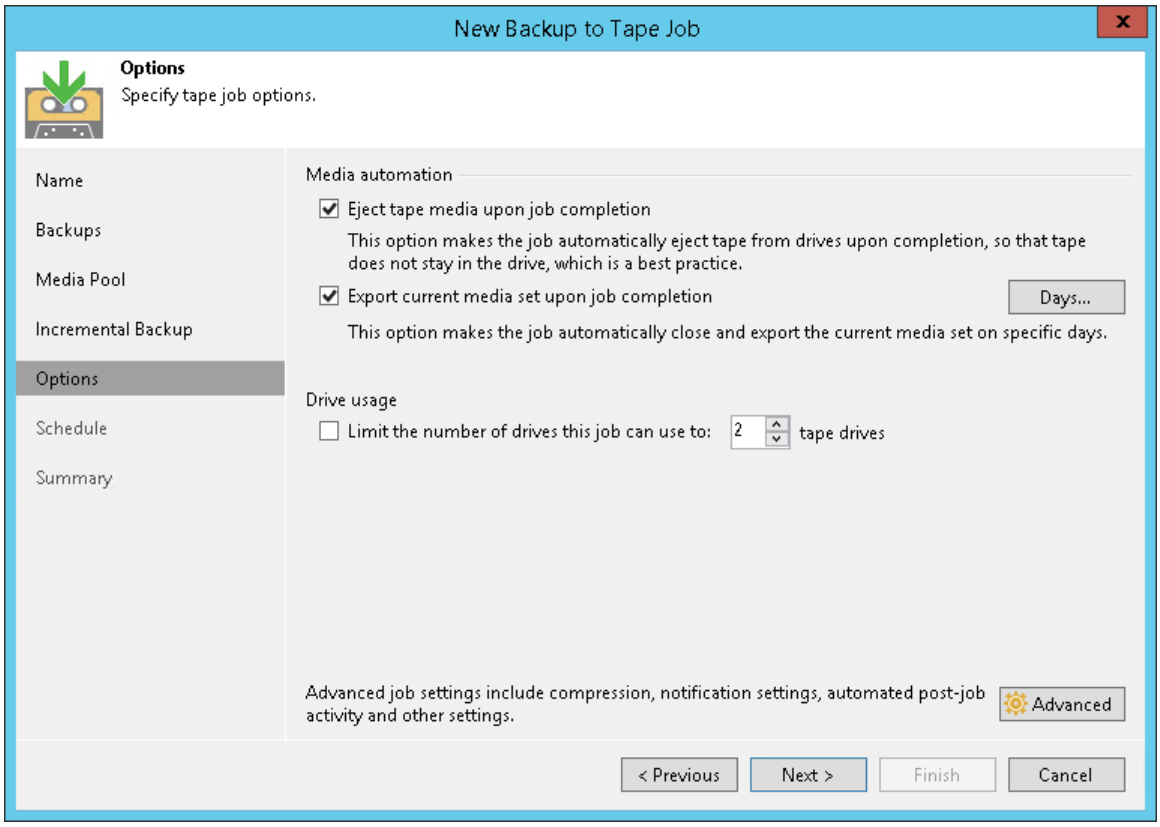

Options for GFS Job

This step of the wizard is available if you selected a GFS media pool at the **Full Backup** step of the wizard.

In the **Media Automation** section, specify options for automatic operations after finishing the job:

- Select the **Eject media upon job completion** check box if the tape should be automatically ejected from the tape drive and placed into a free tape device slot when the job finishes.

  This option does not prevent the tape job from appending data to this tape. If not configured otherwise in media pool settings, this tape will be placed into a drive on the next tape job run.

- Select the **Export the following media sets upon job completion** check box if you want to pull out the tapes with some media sets from the tape device, for example, to move to a storage location. The tape device will eject the tapes that belong to the selected media set.

  Click **Media sets** and select the media sets that you want to export.

If the media pool, which is the target for the tape job, uses parallel processing, you can limit the number of drives to use for processing this tape job. To enable this, select the **Limit the number of drives this job can use to N tape drives** check box and specify the number of drives to use. For more information on how to enable parallel processing for the media pool, see Add Optional Media Pool Settings.

Click **Advanced** to configure the advanced job options.

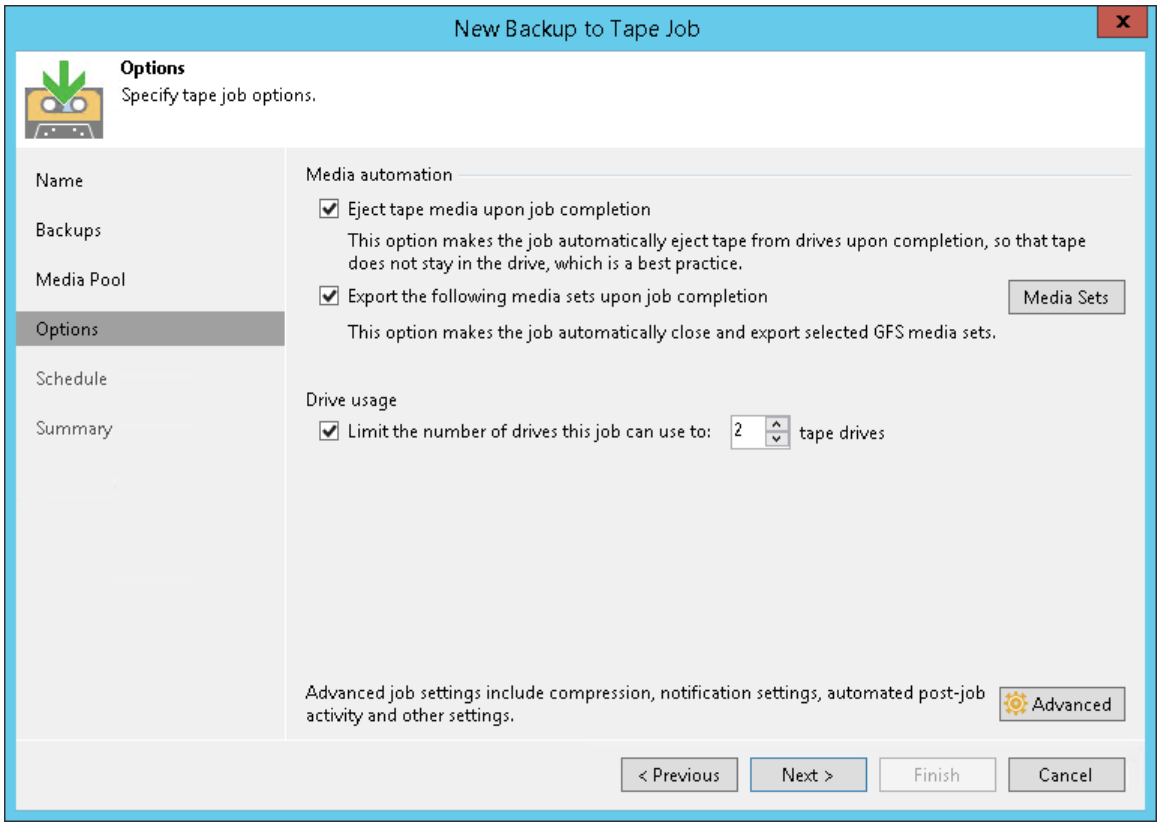

Step 7. Specify Advanced Job Settings

You can specify the following advanced settings for the backup to tape job:

- Notifications Settings
- Advanced Settings

Notifications Settings

At the **Notifications** tab, you can specify notification settings for the backup to tape job.

At the **Options** step of the wizard, click **Advanced**. Then select the **Notifications** tab.

Select the **Send email notifications to the following recipients** check box if you want to receive notifications about tape job status. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon.

> **IMPORTANT**
>
> To receive notifications about tape jobs status, you must enable general email notifications in Veeam Backup & Replication. For more information, see the Veeam Backup & Replication User Guide, section *Specifying Email Notification Settings*.

You can choose between the following options:

- **Use global notification settings**: Veeam Backup & Replication will notify you according to global email notification settings specified for the Veeam backup server.
- **Use custom notification settings specified below**: you can specify notification settings for tape jobs.

  In the **Subject** field, specify a notification subject. You can use the following variables in the subject:

  - *%Time%*: the time when the tape job finished,
  - *%JobName%*: the name of the tape job,
  - *%TapeCount%*: the number of tapes used for the tape job session,
  - *%JobResult%*: the job result,
  - *%VmCount%*: the number of machines in the job,
  - *%Issues%*: the number of machines in the job processed with Warning or Failed status.

Select the occasions on which you want to receive email notifications:

- **Notify on success**: Veeam Backup & Replication will notify you if the tape job finishes successfully,
- **Notify on warning**: Veeam Backup & Replication will notify you if the tape job finishes with a warning,
- **Notify on error:** Veeam Backup & Replication will notify you if the tape job fails,

- **Notify when waiting for tape**: Veeam Backup & Replication will notify you if the tape job cannot start because there are no available tapes.

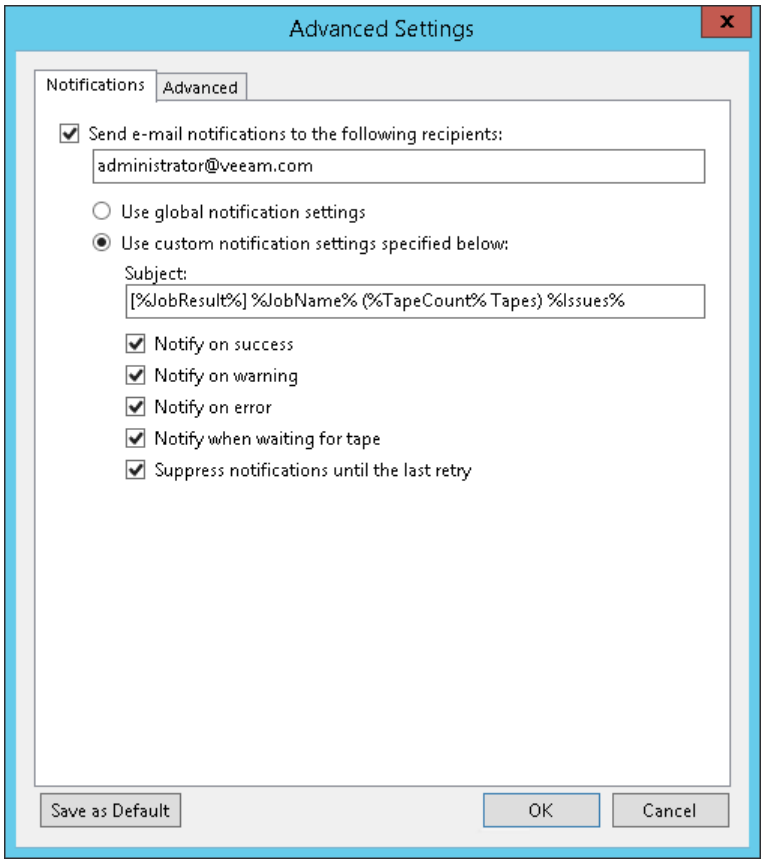

TIP

After you specify necessary settings for the tape 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 to tape job, Veeam Backup & Replication will automatically apply the default settings to the new job.

Advanced Settings

At the **Advanced** tab, you can select what backups from the backup chain to copy, enable hardware compression or specify custom scripts that you want to run before and/or after the backup to tape job.

The advanced options are different for regular backup to tape jobs and GFS tape jobs:

- Advanced Settings for Tape Job
- Advanced Settings for GFS Tape Job

Advanced Settings for Backup to Tape Job

This step of the wizard is available if you selected a regular media pool at the **Full Backup** step of the wizard.

1. Select the **Process latest full backup chain only** check box if you want to copy only the last backup chain with each tape job run. The source backup chain consists of a full backup (active or synthetic) and subsequent increments. If you disable this option, the tape job will back up all restore points that are not on tape.

2. Disable the **Use hardware compression when available** check box.

   Usually image-level backups of VMs are produced by Veeam Backup & Replication by using software compression. We strongly do not recommend to use the hardware compression when writing already compressed data to tapes.

3. Select the **Run the following script before the job** and/or **Run the following script after the job** check boxes and click **Browse** to choose executable file(s).

   You can select to execute pre- and post-job actions after a number of job sessions or on specific week days.

   - If you select the **Run every... backup session** option, specify the number of the job sessions after which the script(s) must be executed.
   - If you select the **Run on selected days only** option, click **Days** and specify week days on which the script(s) must be executed.

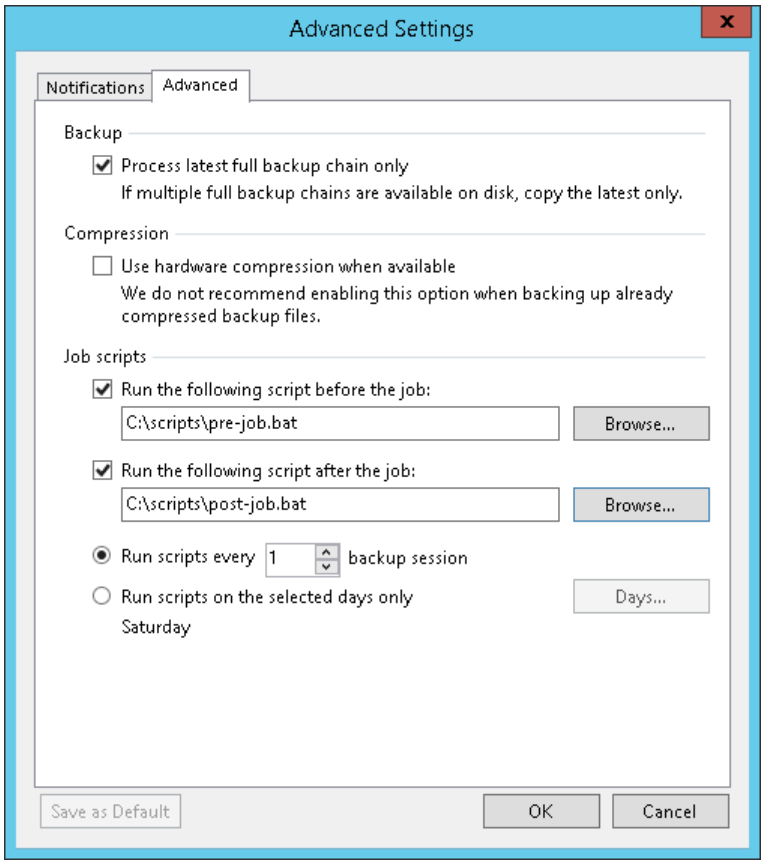

TIP

After you specify necessary settings for the tape 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 to tape job, Veeam Backup & Replication will automatically apply the default settings to the new job.

Advanced Settings for GFS Tape Job

This step of the wizard is available if you selected a GFS media pool at the **Full Backup** step of the wizard.

1. Sometimes when the tape job starts there is no new daily backup on disk. Select the **Process the most recent restore point instead of waiting** check box if you want the tape job to copy the latest restore point instead of waiting. If this option is unselected, the tape job will wait for a new restore point.
2. Disable the **Use hardware compression when available** check box.

   Usually image-level backups of VMs are produced by Veeam Backup & Replication by using software compression. We strongly do not recommend to use the hardware compression when writing already compressed data to tapes.
3. Select the **Run the following script before the job** and/or **Run the following script after the job** check boxes and click **Browse** to choose executable file(s).

   You can select to execute pre- and post-job actions after a number of job sessions or on specific week days.

   - If you select the **Run every... backup session** option, specify the number of the job sessions after which the script(s) must be executed.
   - If you select the **Run on selected days only** option, click **Days** and specify week days on which the script(s) must be executed.

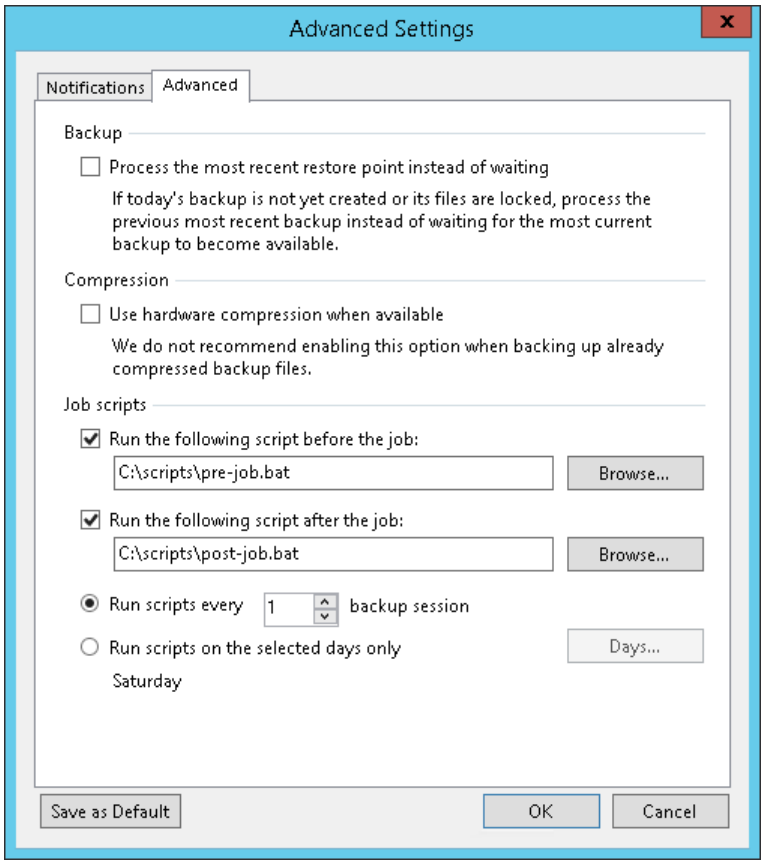

> **TIP**
>
> After you specify necessary settings for the tape job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup to tape job, Veeam Backup & Replication will automatically apply the default settings to the new job.

Step 8. Define Job Schedule

At the **Schedule** step of the wizard, you can define a schedule for the job to run on a regular basis. The schedule settings are different for regular backup to tape jobs and GFS tape jobs:

- Schedule for Backup to Tape Job
- Schedule for GFS Tape Job

Schedule for Backup to Tape Job

This step of the wizard is available if you selected a regular media pool at the **Full Backup** step of the wizard.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be started manually.

You can define the following scheduling settings for the job:

- **Daily at this time**: the tape job will run at specific time on defined week days.
- **Monthly at this time**: the tape job will run at specific time monthly.
- **After this job**: the tape job will start when a corresponding backup job completes. Choose the preceding backup job from the list.

> **NOTE**
>
> Mind the following:
>
> - This option will only start the tape or backup copy job if the source job is started automatically by schedule. If the source job is started manually you will choose to start it chained or not.
> - This option is not available for agent backup jobs managed by agent or for standalone agent backup jobs.

- **As new backup files appear**: the tape job will monitor the source jobs. As soon as the source job creates a new backup, the job will write this backup to tape.

  If necessary, you can limit the time when the backup to tape job can start with a schedule. You can restrict start of the backup to tape jobs when backup repositories are busy with other tasks (backup jobs writing to repositories or backup copy jobs reading from repositories). Schedule is set for backup to tape jobs start. Jobs that have started during the allowed time interval will continue working on restricted hours. To define prohibited time for the backup to tape job start, click the **Schedule** button and define the time when the job is allowed and prohibited to start.

If you have scheduled the job to run at the specific time daily or monthly, consider configuring wait timeout value. Select the **If some linked backup jobs are still running, wait for up to** check box and specify the new timeout. When a backup to tape job starts, Veeam Backup & Replication checks the status of the source jobs. If a source job is still writing data to the source repository, the backup to tape job will wait for the specified time interval.

If the timeout is disabled, the backup to tape job will terminate without waiting for the backup job to complete.

The timeout option is unavailable if you schedule the backup to tape job to run after a backup job or if you schedule the backup to tape job to start when new backups appear.

Sometimes, the source job may start when the tape job is still running. This may cause a conflict if the source job needs to perform the following operations:

- Merge of backup files

- Reverse incremental backup transformation
- Compact of full backup file
- Retention Policy maintenance.

By default, the source job has priority. In this case, the tape job terminates with error and no data is written to tape. Select the **Prevent this job from being interrupted by source backup jobs** option if you want to give the tape job a higher priority. If this option is selected, the tape job will not terminate and will finish writing the data. If the source job needs to perform the listed above operations, it will wait for the tape job to finish. Note that in this case the source job may finish with a significant delay.

> **NOTE**
>
> If the source job creates per-machine backups, the tape job gets priority for each per-machine restore point separately. In this case, the source job will be able to lock other per-machine restore points and perform the needed operations. Make sure that you set sufficient timeout in the **If some linked backup jobs are still running, wait for up to** window to avoid the tape job failing.

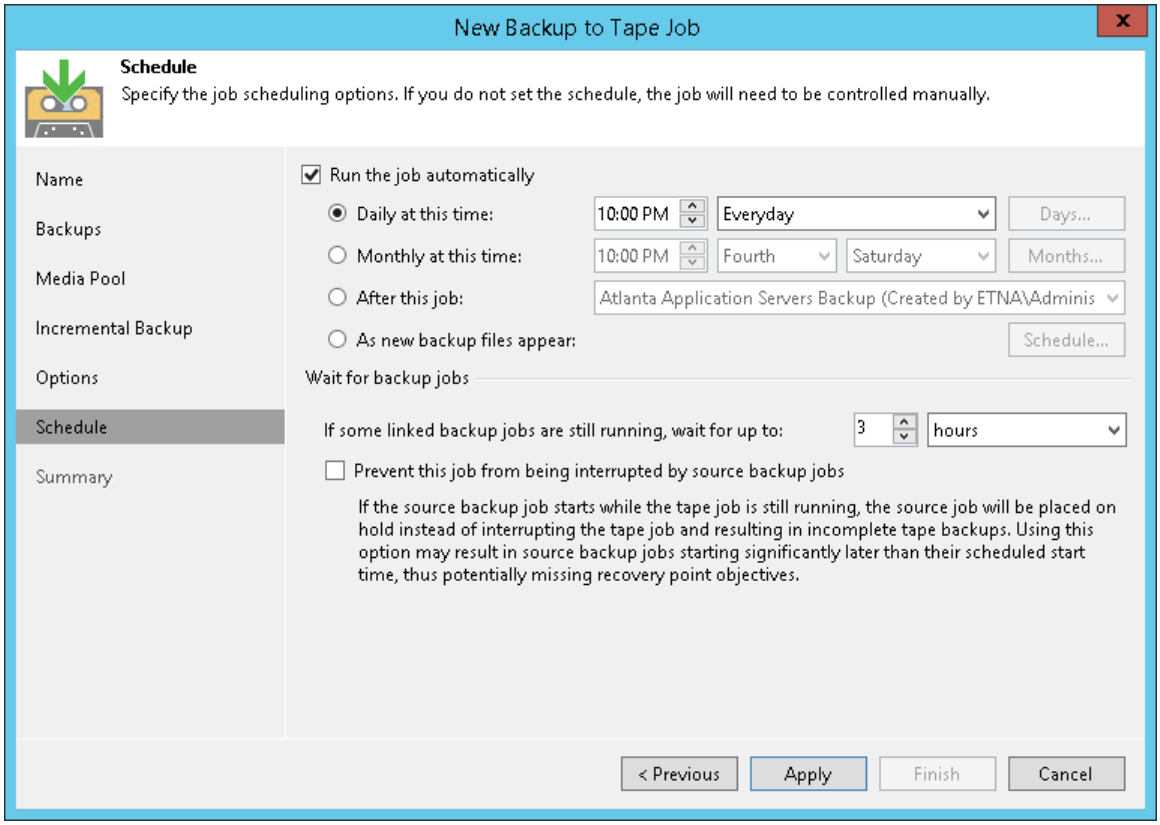

Schedule for GFS Job

This step of the wizard is available if you selected a GFS media pool at the **Full Backup** step of the wizard.

Click **Schedule** to select days for each media set.

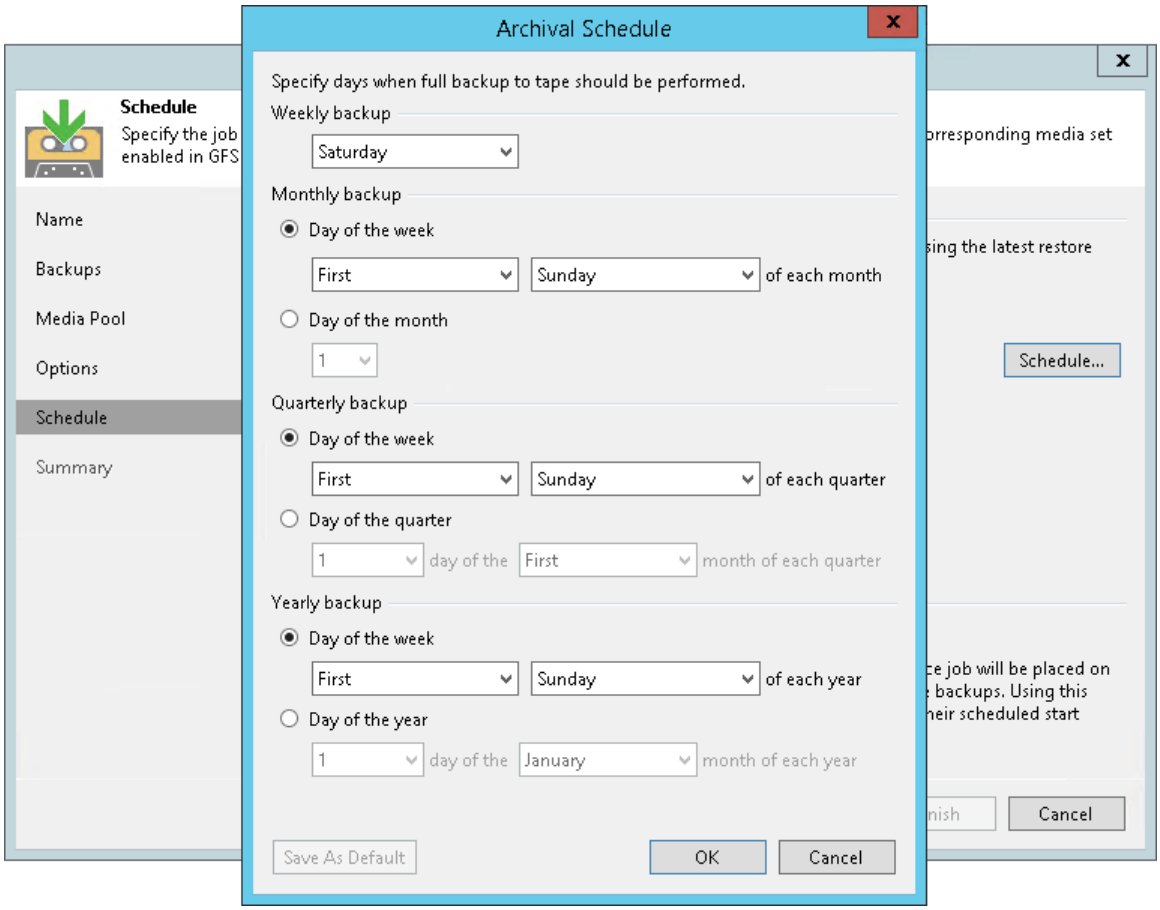

In the **Perform GFS scan daily at** field, specify the time when the GFS job must start.

> **TIP**
>
> After you specify necessary schedule settings, you can save them as default. To do this, click **Save as Default** at the bottom left corner of the **Archival Schedule** window. When you create a new GFS job, Veeam Backup & Replication will automatically apply default settings to the new job schedule.

Sometimes, the source job may start when the tape job is still running. By default, the source job has priority. In this case, the tape job terminates with error and no data is written to tape. Select the **Prevent this job from being interrupted by source backup jobs** option if you want to give the tape job a higher priority. If this option is selected, the source job will wait until the tape job finishes. Note that the source job may start with a significant delay.

For tenant to tape jobs, this option is active by default and cannot be disabled.

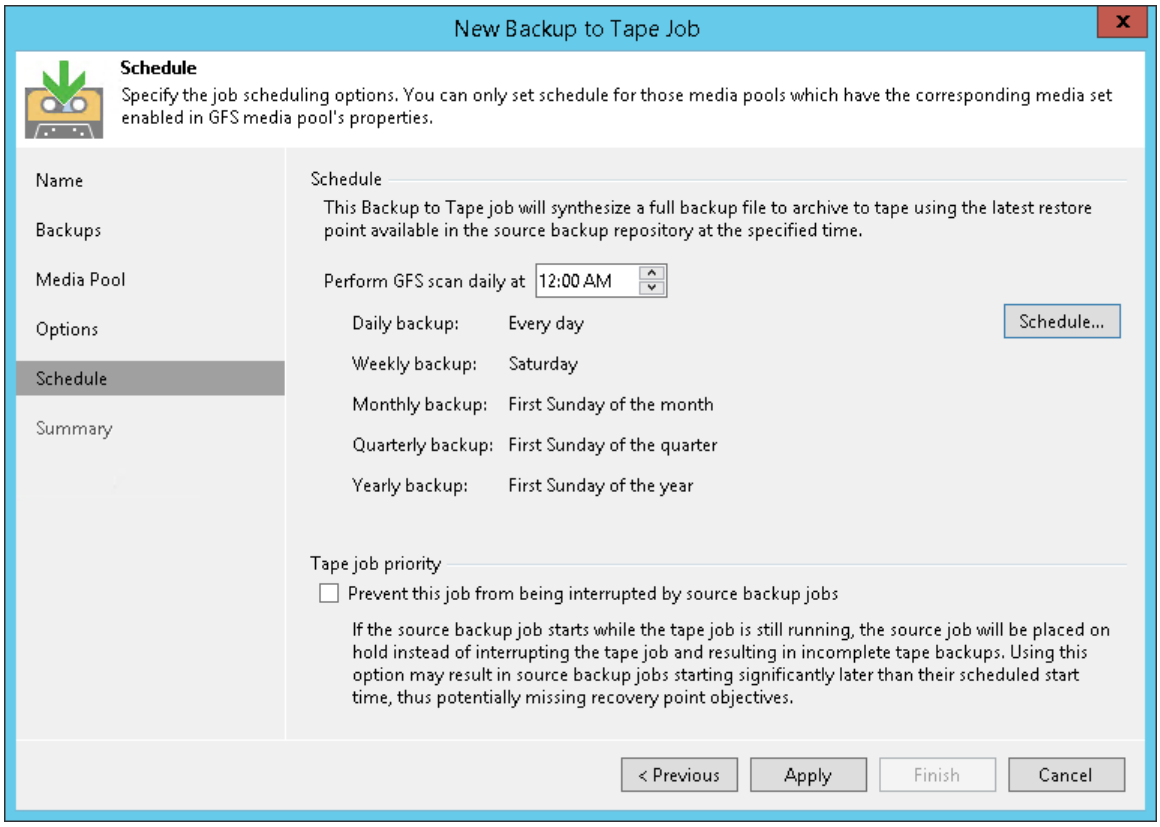

Step 9. Finish Working with Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start archiving backups to tape job right after you complete working with the wizard. Click **Finish** to close the wizard.

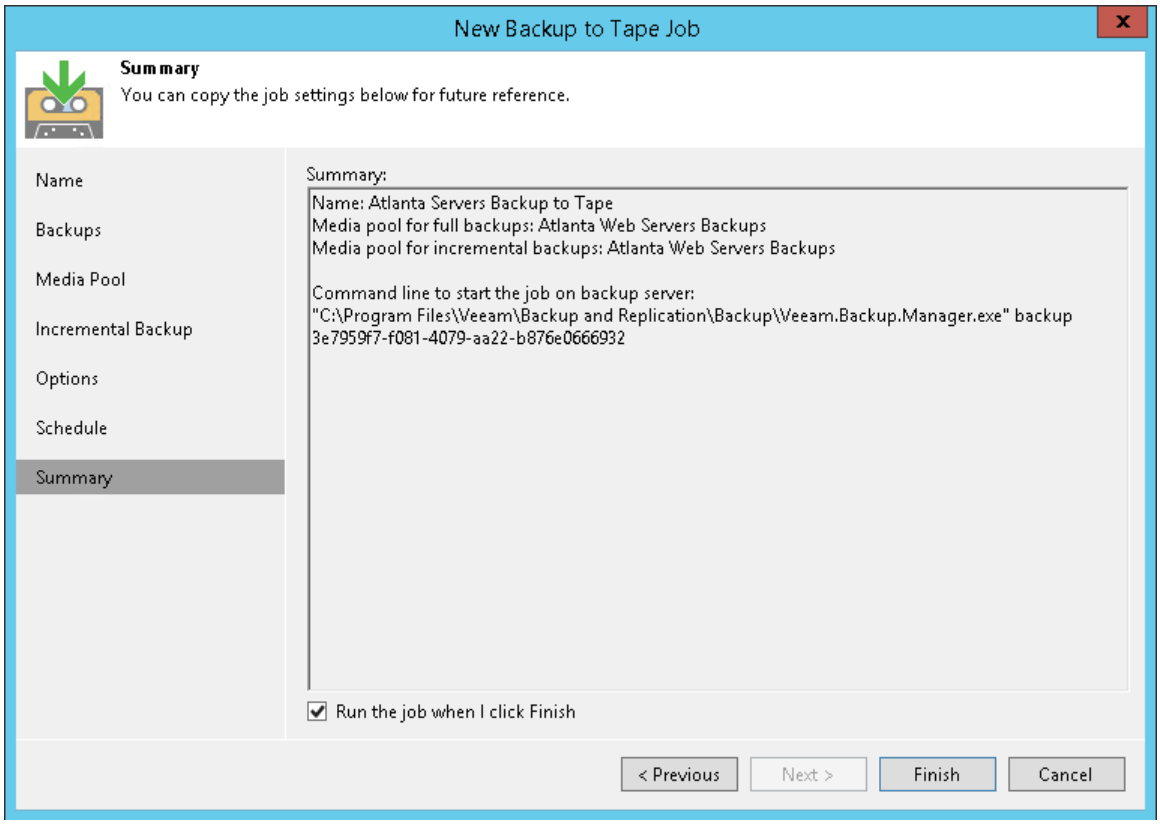

Linking Backup Jobs to Backup to Tape Jobs

Veeam Backup & Replication provides two options for linking backup jobs to backup to tape jobs:

- If you already have backup jobs configured, you can choose the necessary jobs in the **Backup to Tape Job Wizard**. For details, see Creating Backup to Tape Jobs.
- Alternatively, you can point a backup job to an existing backup to tape job using the **Backup Job** wizard.

To point a backup job to an existing backup to tape job, perform the following steps:

1. Open the backup job settings and navigate to the **Storage** step. Select the **Configure secondary destination for this job** check box.

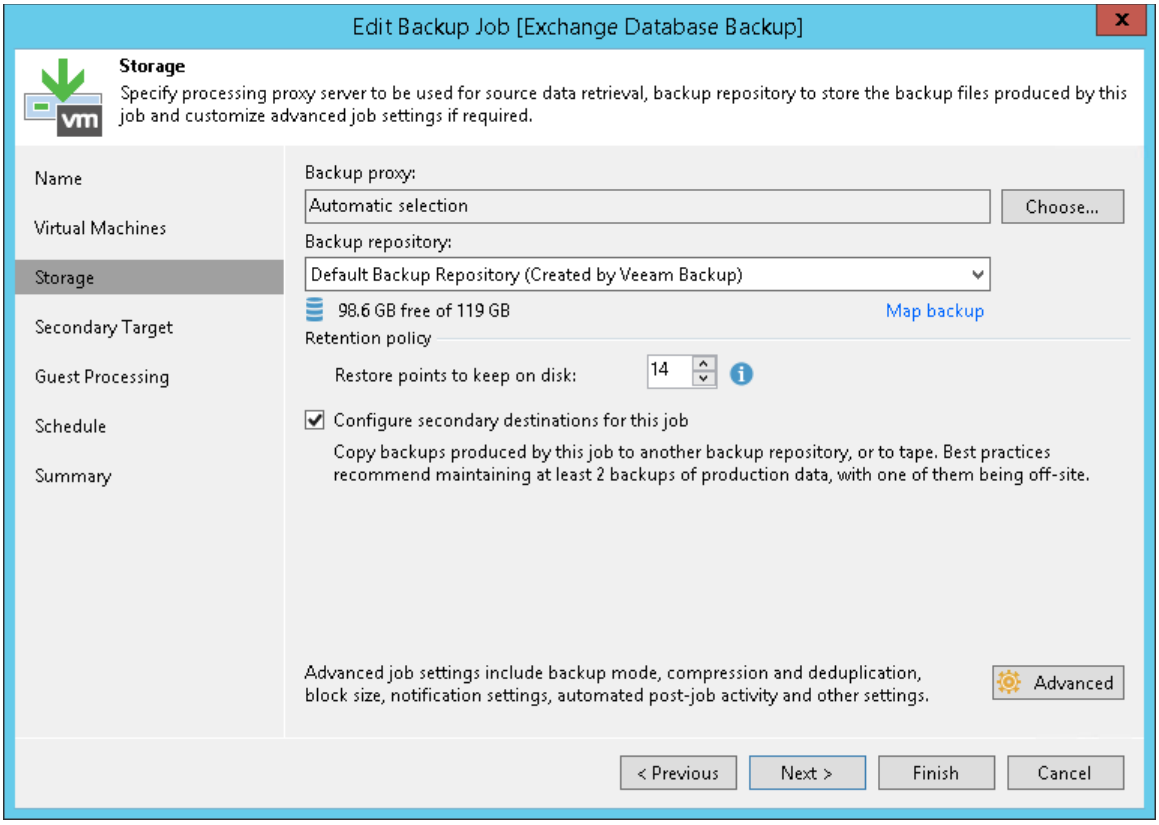

2. At the **Secondary Target** step, click **Add** and choose a backup to tape job to which the backup job should be linked.

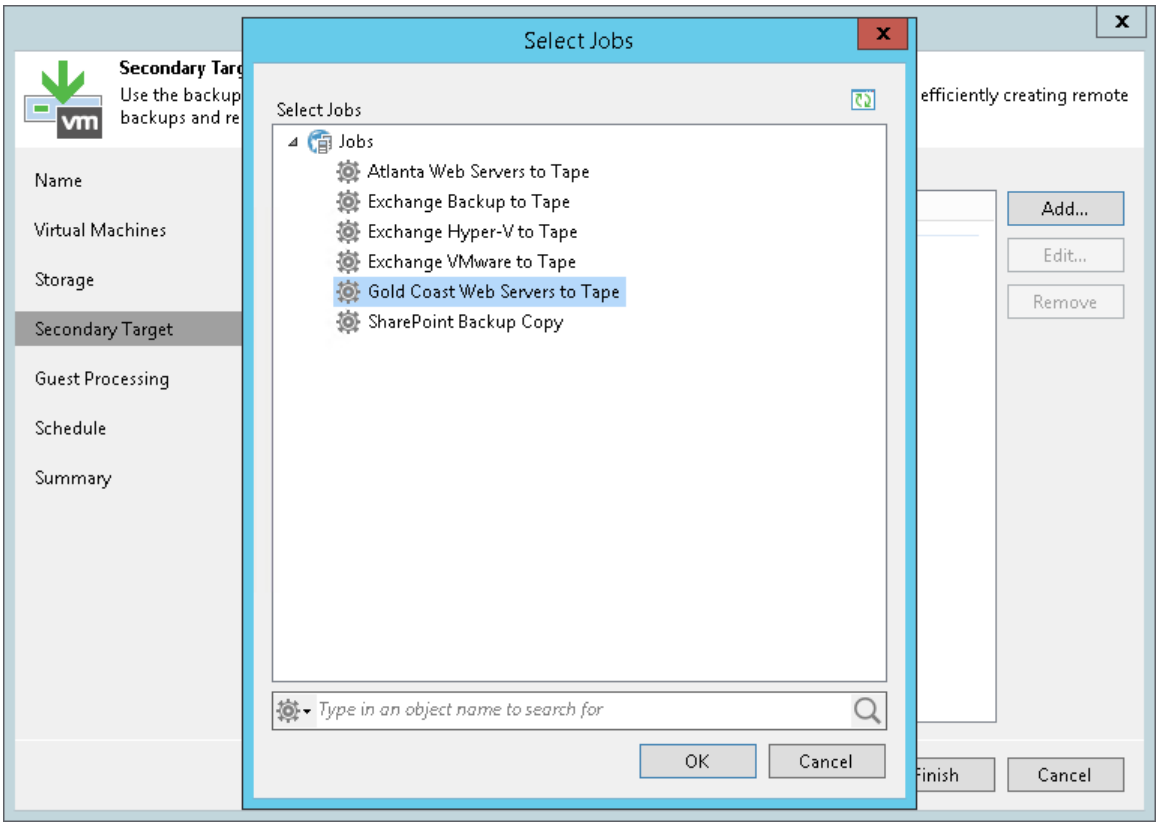

3. Save settings. Veeam Backup & Replication will automatically update backup files settings of the chosen backup to tape job.

GFS Backup to Tape

The GFS tape job creates yearly archive for source machines by the GFS (Grandfather-Father-Son) scheme. The GFS archive includes one yearly backup and several quarterly, monthly, weekly and daily backups.

You can create GFS backups for VMs and Veeam agent computers (physical machines) on tape.

The GFS tape job does not process source machines. The job uses backups created by machine-to-disk jobs.

GFS Media Sets

To distinguish between the backup cycles, the GFS media pool has 5 pre-defined media sets: yearly, quarterly, monthly, weekly and daily.

Media Set	Restore Point Type
Yearly	Full backup or synthesized virtual full backup
Quarterly	
Monthly	
Weekly	
Daily	As on disk (for details, see How GFS Backup to Tape Works)

Individual Media Sets Configuration

You can configure each media set individually. You can set retention period, appending data to tape or other settings for each media set separately.

For example, the weekly backups must be stored for 4 weeks, data is appended, tapes may be overwritten. The yearly backups must be never overwritten, each backup on an individual tape (or set of tapes), tapes must be moved to vault.

Manual Active Full

You can manually create a restore point for any media set at any time with the **Active full** option.
Veeam Backup & Replication will copy or synthesize the most recent full backup possible for the selected period. The manual active full is not available for the **Daily** media set.

Schedule Overlap

If the schedule for backup periods overlap, only the backup for the eldest media set will be archived. The media sets have the following priority: **Yearly**, **Quarterly**, **Monthly**, **Weekly**, **Daily**. For example, if "the first Sunday of January" is the date for the monthly, the quarterly and the yearly backup, the GFS job will create only the yearly backup; the monthly and the quarterly backups will be skipped.

How GFS Backup to Tape Works

To create a GFS archive, Veeam Backup & Replication performs the following operations:

1. The GFS job starts according to the GFS job schedule and checks which media set is scheduled for today.
2. For GFS jobs with the **Daily** media set disabled:
   - If the source job creates a full backup, the GFS job copies it.
   - If the source job creates an incremental backup, the GFS job synthesizes a virtual full backup.

   If the source job created several restore points, the tape job will copy only the first one.
3. For GFS jobs with the **Daily** media set enabled:
   - If today some elder media sets are also scheduled, for example, the weekly media set, the GFS job archives the first restore point to the **Weekly** media set:
     - If the source job creates a full backup, the GFS job copies it.
     - If the source job creates an incremental backup, the GFS job synthesizes a virtual full backup.

     If there is more than 1 restore point on disk available, the GFS job copies the rest of them to the **Daily** media set.

     If only the **Daily** media set is due today, the GFS job copies all available restore points to the **Daily** media set.
4. If no backup appears on the disk till the end of the day, the GFS job waits for the source job to start. If the source job starts during the day, the GFS job waits till the source job finishes and copies the created backup to tapes. If the source job does not start, the GFS job stops at midnight. It starts again according to the configured schedule.

Consider the following:

- No restore point will be archived to the **Daily** media set when the daily backup overlaps elder media sets dates, and only one source restore point is available.
- Daily increment restore points are linked to a full restore point in any media set. For example, a full restore point from the **Weekly** media set is used. If there is no full restore point that can be used for linking the daily increments, the tape job will create a virtual full backup and put it to the **Daily** media set.
- The GFS job copies only new restore points. If the source job did not create a new point after the last GFS job run, the GFS job will not copy anything to tape. For example, if you want weekly backups in the tape GFS archive, make sure that the source job runs at least once a week.
- When you link a source job to a GFS job, the GFS job does not copy old restore points. The GFS job will copy only the restore point created on the day when the GFS job was linked to the source job or new restore points created later.
- If the restore point is available but locked (for example, if the source job is a backup copy job), the GFS job will wait for this restore point to become available for up to 7 days. If the restore point is still locked after 7 days, the GFS job will use the most recent restore point that is not locked.

  If you do not want to wait, you can instruct the tape job to copy the most recent restore point instead of waiting. To do this, select the **Process the most recent restore point instead of waiting** check box in the settings of the tape job. For more details, see Advanced Settings for GFS Tape Job.
- If you stop the tape job manually, it will start again the next day at the time scheduled.

- If you use the **Daily** media set, after creating or enabling a tape job it will be started within an hour even if the scheduled start time has already passed.
- If there is a technical problem, for example, the tape library is offline, or the backup server is down, the GFS job will not run. When the problem is fixed, the GFS job will archive one most recent missed full backup for each media set.
- If the GFS tape job was disabled manually for some period, it will skip all backups for this period after it is resumed.
- If the GFS job failed for some reason, it will try to restart every hour for 24 hours.
- If a job is unable to complete within a period of 21 days, it will be stopped with the 'Failed' status.

Creating GFS Tape Jobs

To create a GFS archive on tape, you need to create a GFS media pool and target a backup to tape job to it. Technically, a GFS tape job is a variant of backup to tape job. When you select a GFS media pool as target, the job schedule automatically changes to the GFS mode. For details, see Creating Backup to Tape Jobs.

As a source, you can use any backup job with any backup method: forever forward incremental, forward incremental or reversed incremental.

> **NOTE**
>
> If you use a job with enabled GFS retention as a source for a GFS tape job, the GFS tape job writes restore points to tapes according to its own schedule. The schedule of the source job is not taken into account.

Tenant Backup to Tape

This option is available for Veeam Cloud Connect service providers.

Service providers can back up their tenants' backups from cloud repositories. This allows for better protection of the tenants or for regulations compatibility. For more information, see the Veeam Cloud Connect guide.

To back up a tenant to tape, run a backup to tape job. For more information, see Machines Backup to Tape.

Viewing Backups on Tape

After the backup to tape job completes, you can view the created archive on tape:

1. Open the **Home** view.
2. Expand the **Backups** > **Tape** node and locate the backup archive under the name of a corresponding job.

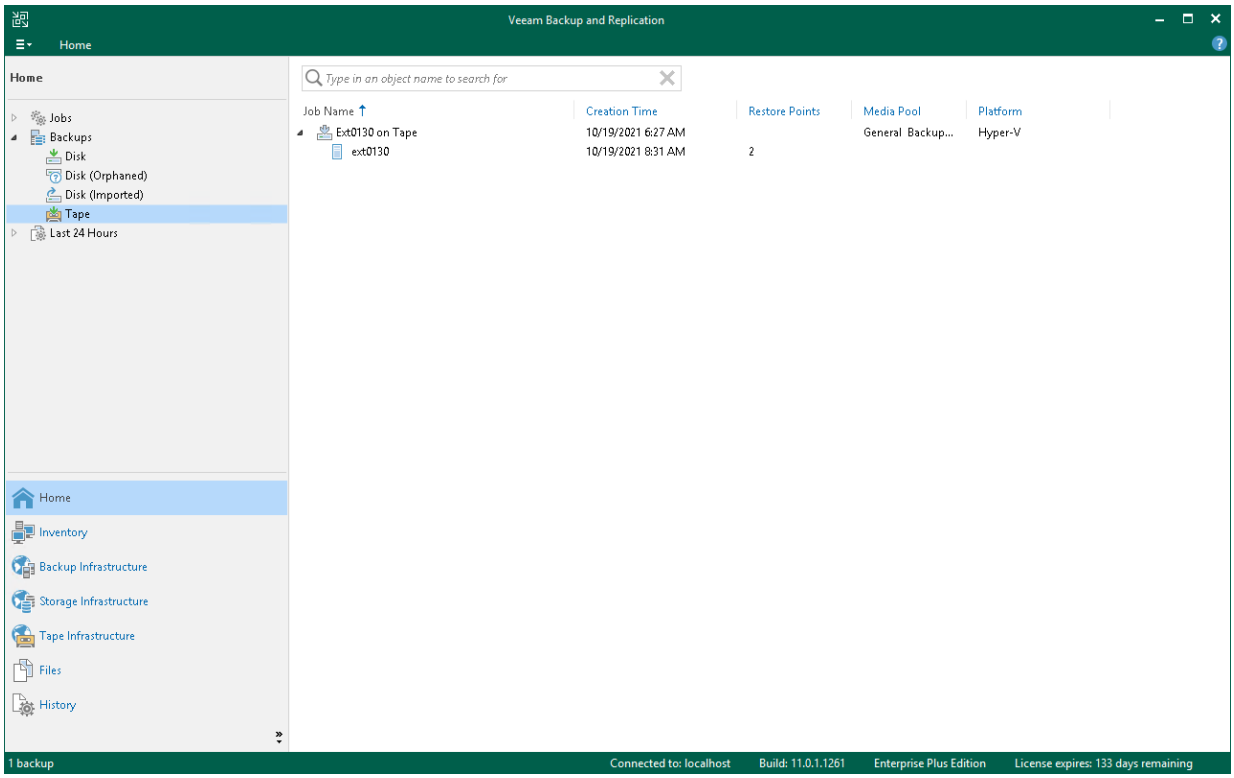

File Backup to Tape

File to tape job allows you to back up to tape any Microsoft Windows or Linux files. Also, you can back up volumes of storage devices that support the NDMP protocol. For more information, see NDMP Servers Backup to Tape.

You can back up files to tape from the following storing devices:

- Windows servers, including physical boxes.
- Linux servers, including physical boxes.
- NAS devices (by specifying the path to the SMB (CIFS) or NFS file share).
- NDMP servers.

The file to tape job compares the source files to the files stored in tape archive and copies the changes to tape. You can create both full and incremental backups of files on tape. Veeam Backup & Replication supports file backup from any server which has been added as a managed server to the Veeam Backup console (that is, Windows or Linux server, including physical boxes). You can also archive files residing on NAS devices. For more information on backup of NAS file shares, see the Veeam Backup & Replication User Guide, section *NAS Backup Support*.

To back up Veeam backup files, you can use backup to tape jobs that are specially intended for this and offer more possibilities. However, you can archive backups as files using file to tape job.

When planning file to tape jobs, consider that the job performance depends more on the number of files to back up then on the amount of data. For example, writing a large number of small files with overall size of 10GB with one job will take more time than writing one 10GB file. If your job contains an extra-large number of files (like millions of files) with one job, the job performance will be affected significantly. To improve performance, consider creating several file to tape jobs.

> **NOTE**
>
> If the file to tape job fails to complete in 3 weeks, it is terminated by timeout.

How File Backup to Tape Works

When Veeam Backup & Replication executes a file to tape job (started manually or on schedule), it performs the following operations:

1. The file to tape job detects files that match the job criteria.
2. The files are queued for archiving. The following scenarios are used:
   - If it is a first job run or a scheduled full backup, all selected files are queued for archiving.
   - If it is an incremental backup run, Veeam Backup & Replication addresses the Tape Catalog in the Veeam Backup & Replication database to detect if any data has been modified since the latest backup. Detected changes are queued for archiving.
3. Veeam Backup & Replication connects to Veeam Data Movers and starts the data transfer process. The source Data Mover retrieves data from the source servers and target Data Mover sends data to tape.
4. The tape job addresses the media pool that is set for this job as target. The media pool allots tapes for writing data according to the following configuration options:
   - Tapes consumption
   - Media sets
   - Tape retention
5. While tape recording is performed, Veeam Backup service updates data in the Tape Catalog in Veeam Backup database. The Veeam Backup console displays refreshed information about files archived to tape and shows job statistics.

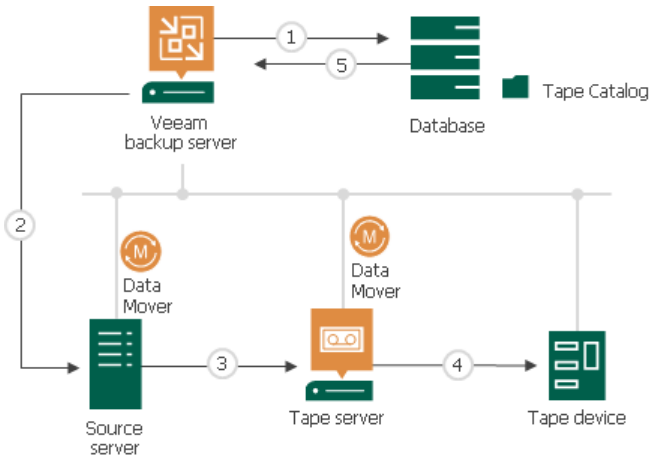

Before You Begin

Before you configure a file to tape job, complete the following prerequisites:

- You must configure one or more regular media pools with the necessary media set and retention settings.
- You must load tapes to the tape device and configure the target media pool so that it has access to them. If the media pool has no available tape, the tape job will wait for 72 hours and then terminate.
- Make sure that files and folders that you want to back up do not contain certain Unicode symbols, including but not limited to the symbols with the following codes: U+D800 – U+DFFF, U+FFFE, U+FFFF (55296 – 57343, 65534, 65535), U+007C (124). Also the file and folder names must contain only symbols that are supported by the source file system. Files and folders with unsupported names will be skipped from processing.
- Make sure that files and folders that you want to back up are not symlinks. Backup of symlinks is not supported.
- For SMB and NFS file share backup to tapes, make sure that the shared source does not contain folders with the same name entered in different letter case (for example, **Documents** and **documents**).
- To back up NDMP volumes, add a storage device that supports the NDMP protocol as an NDMP server. For more information, see Adding NDMP Servers.

Keep in mind that if a job is unable to complete within 21 days period, it will be stopped with the 'Failed' status.

System Requirements for Large Number of Files in Job

If the file to tape job will process large quantities of files, for example, more than 1 000 000 files in 1 000 folders, you must provide the following system resources:

- Backup server: 1.5 GB RAM per each 1 000 000 files
- Tape server: 800 MB RAM per each 1 000 000 files.

You must provide these resources in addition to general system requirements. For more information, see the Veeam Backup & Replication User Guide, section *System Requirements*.

> **IMPORTANT**
>
> Do not back up to tape large number of files (over 1 000 000 files in 1 000 folders) in one tape job if your Veeam backup server uses Microsoft SQL Server Express edition. Processing a large number of files will result in considerable performance penalty. Also bear in mind that the usage of Microsoft SQL Server Express Edition is limited by the database size up to 10 Gb. If you plan to have larger databases, use other editions of Microsoft SQL Server.

NDMP Servers Backup to Tape

If your NAS device supports the NDMP protocol, you can back up data from it to tape. For this, you need to utilize the file to tape job.

> **NOTE**
>
> Support of NDMP servers backup to tape is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise Plus edition of Veeam Backup & Replication is required. For details, see Veeam Editions Comparison.

To back up data from NAS devices by the NDMP protocol, you need to add the device as an NDMP server. For details, see Adding NDMP Servers.

> **IMPORTANT**
>
> Even if the NAS device supporting NDMP protocol is already added to Veeam Backup & Replication, you need to add the NDMP server as a separate procedure. Otherwise you will not be able to perform file backup to tape.

Requirements for NDMP Servers

NDMP version 4 or higher is supported.

Number of Restore Points

For all NDMP devices, the backup chain stored on tapes will consist of 10 restore points maximum. On the 11th run, Veeam Backup & Replication will force an active full.

Limitations for NDMP Servers

- Only backup of whole volume is supported.
- Both the NetApp "SVM-scoped NDMP" mode and the NetApp Cluster Aware Backup (CAB) extension of the NDMP protocol are supported.

  In SVM-scoped configuration, add each cluster node to the backup user interface on their own data LIF IP/DNS address. This is also known as "Node-scoped NDMP".

  To support backup of NetApp CAB extension, add the NDMP server with IP address of the cluster management interface or with the NAS server name resolving to this IP address.
- The NDMP server performs its own backup consistency check for the NDMP volumes. To provide for the correct calculation of the delta of changes and to ensure the backup chain integrity, Veeam Backup & Replication does not allow adding one NDMP volume to more than one file in tape jobs.

See Also

File Restore from Tape

Adding NDMP Servers

To add an NDMP server, follow the next steps:

Step 1. Launch New NDMP Server Wizard

To launch the **New NDMP Server** wizard:

1. Open the **Tape Infrastructure** view.
2. Do one of the following:
   - In the inventory pane, right-click the **NDMP Servers** node and select **Add Server**.
   - In the inventory pane, select the **NDMP Servers** node and click **Add Server** on the ribbon.
   - In the inventory pane, select the **Tape Infrastructure** node and click **Add NDMP Server** on the ribbon.

Step 2. Specify NDMP Server Name and Location

At the **NDMP Server** step of the wizard, enter the name of the NDMP server and connection details.

1. In the **NDMP server name and port** field, enter the name or the IP address of the NDMP server you want to connect and the port for the connection. Check your NAS device settings for details.
2. From the **Credentials** list, select credentials for the account that has administrator privileges on the NDMP server. If you have not set up credentials beforehand, click the Manage accounts link or click **Add** on the right to add the credentials. For more information, see Managing Credentials.

   Veeam Backup & Replication will use the provided credentials to deploy its components on the added server.
3. In the **Gateway server** section, specify settings for the gateway server:
   - If a network connection between the NDMP server and gateway server is fast, choose **Automatic selection**. In this case, Veeam Backup & Replication will automatically select a gateway server.
   - If you perform backup over WAN or slow connections, choose **The following server**. From the list below, select a Microsoft Windows server on the target site that you want to use as a gateway server. The selected server must have a direct access to the NDMP server and must be located as close to the NDMP server as possible.

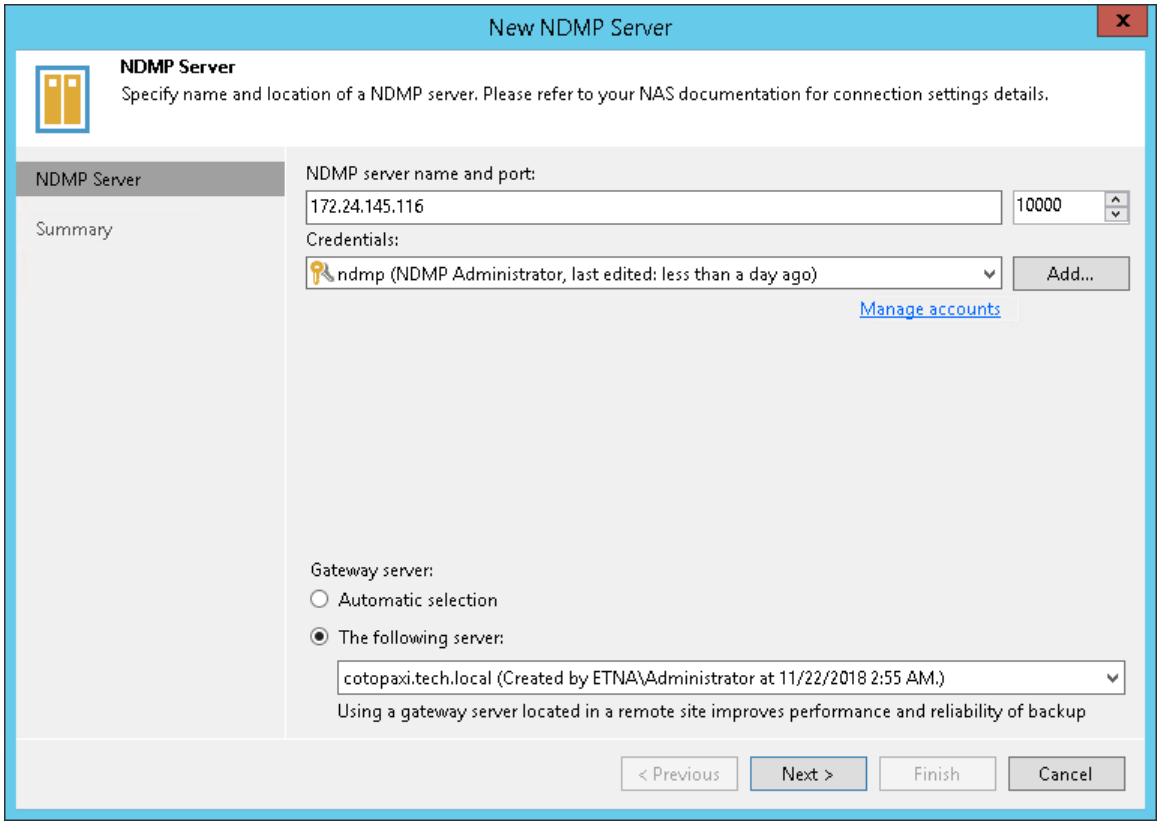

Step 3. Finish Working with the Wizard

Review the NDMP server settings and click **Finish** to complete the wizard.

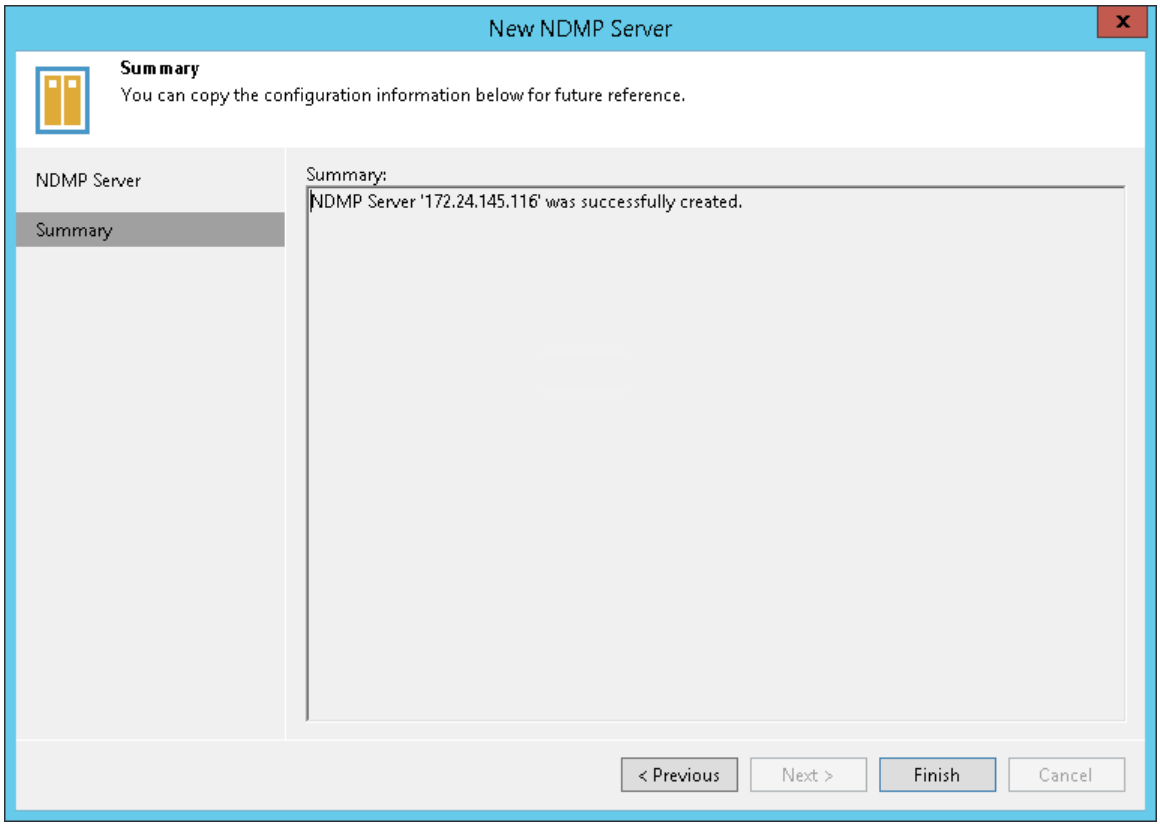

Creating File to Tape Jobs

To write files to tape, you should create a file to tape job using the **New File to Tape Job** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

To create a file to tape job, follow the next steps:

Step 1. Launch New File to Tape Job Wizard

To run the **New File to Tape Job** wizard, do either of the following:

- On the **Home** tab, click **Tape Job** and select **Files**.
- Open the **Home** view, right-click the **Jobs** node and select **Tape Job** > **Files to Tape.**

Alternatively, you can:

- Open the **Files** view, browse to the necessary files, select the files and choose **Add to Tape Job** > **New job** from the ribbon menu.
- Open the **Files** view, browse to the required files, right-click the necessary files and choose **Add to Tape Job** > **New job**.

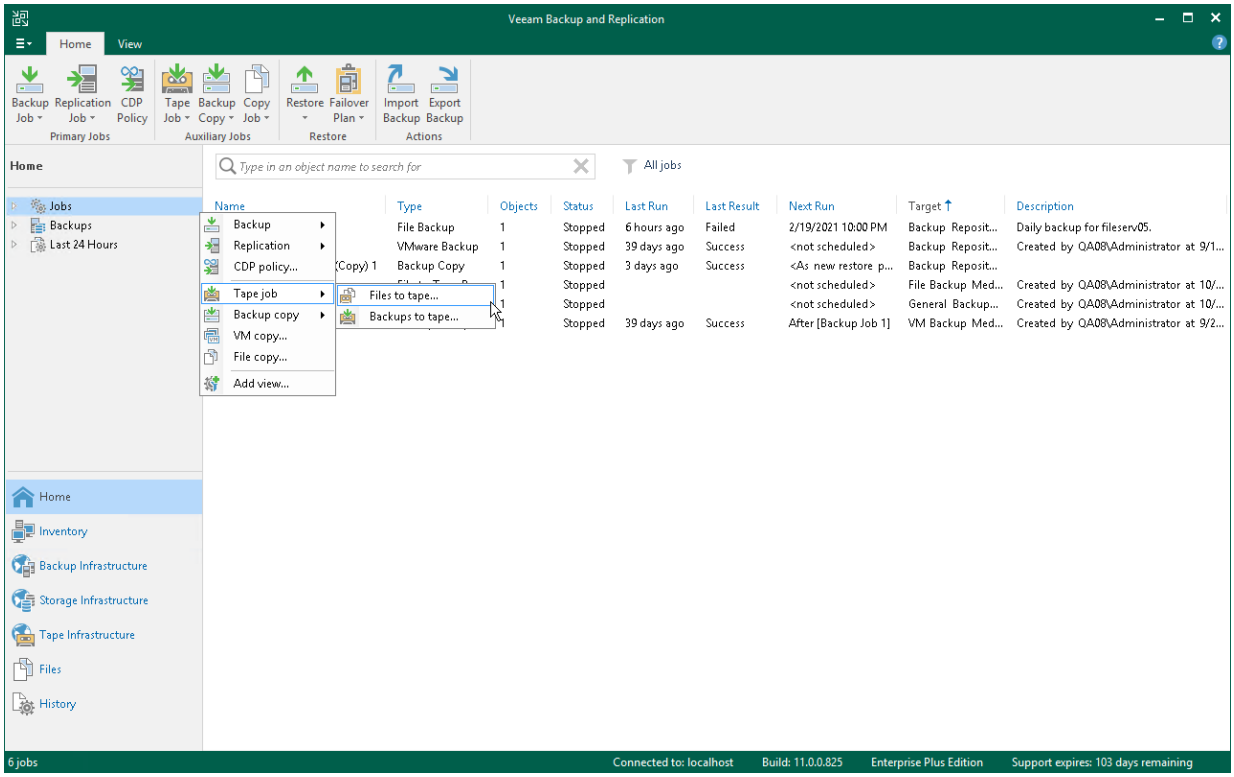

Step 2. Specify Job Name and Description

At the **Name** step of the wizard, you should define basic settings for the created file to tape job.

1. In the **Name** field, enter a name for the created job.
2. In the **Description** field, enter a description of the created job. The default description contains information about the user who created the job, date and time when the job was created.

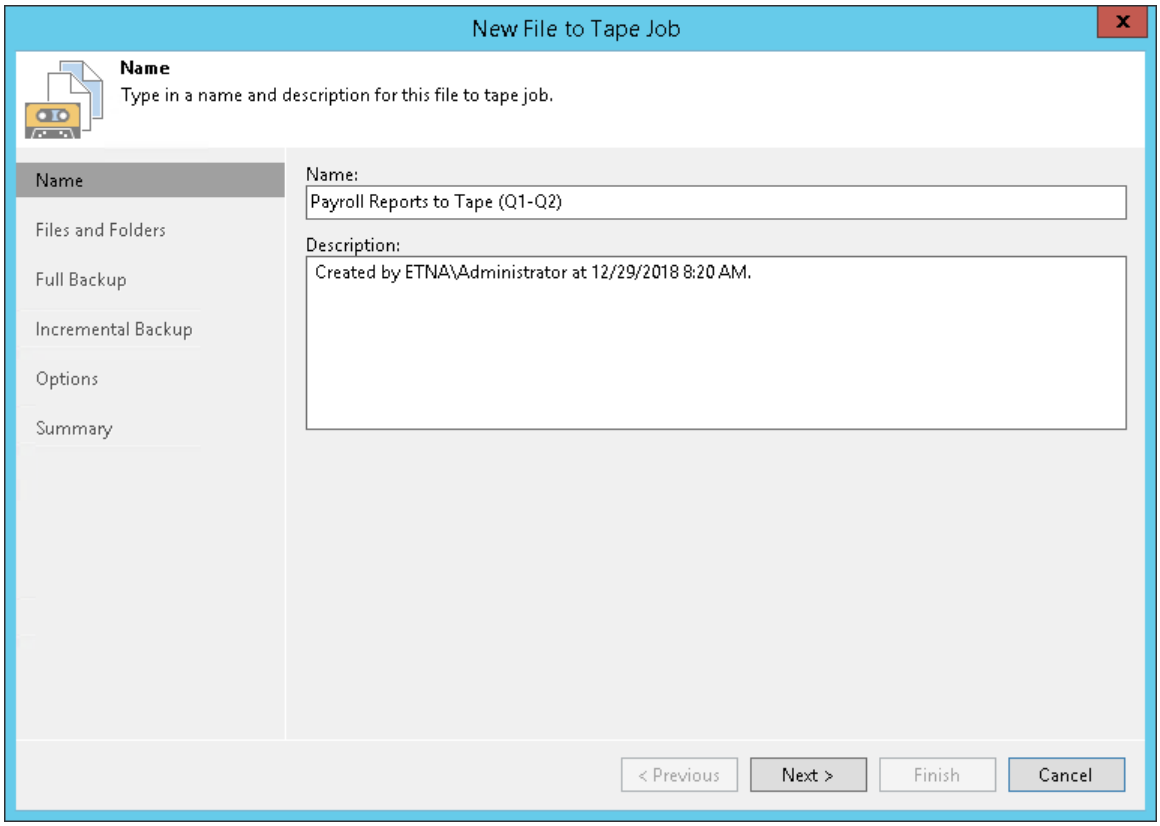

Step 3. Choose Files and Folders to Archive

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** drop-down list, select a data source where the necessary files or folders reside:
   - If you plan to back up files and folders from a Windows or Linux server, make sure it is added as a managed server to the backup infrastructure. For details, see the Veeam Backup & Replication User Guide, sections *Adding Microsoft Windows Servers* and *Adding Linux Servers*.
   - If you plan to back up files and folders from an SMB or NFS share, make sure it is added as a file share to the inventory. For details, see the Veeam Backup & Replication User Guide, section *Adding File Share*.
   - If you plan to back up files and folders from a device supporting the NDMP protocol, make sure it is added as an NDMP server to the tape infrastructure. For details, see Adding NDMP Servers.

   The account you use to access the target server must have read permissions for the folder and all its child folders.
3. In the **Folders** tree, select folders you want to backup.

   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.

   > **NOTE**
   >
   > On an NDMP server, you can select only volumes. You cannot back up separate files or folders. Also, you cannot use file masks.

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.

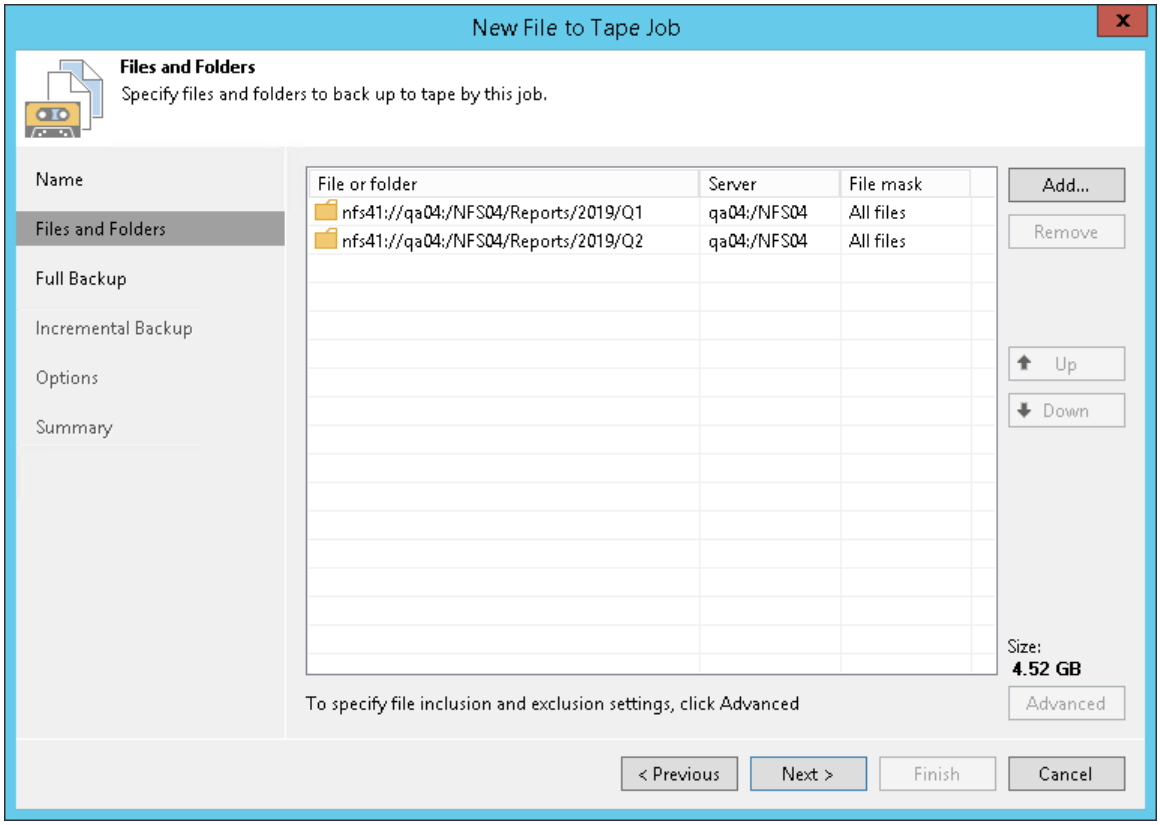

If you add a folder to the job, all the folder contents will be processed. If necessary, you can choose only specific files from the added folder. To do so, select a folder in the list and click **Advanced**. In the **File Filters** window, use **Include masks** and **Exclude masks** fields to filter the folder contents. You can use exact file names or create name masks (for example, ***.EVT** or ***.PDF**).

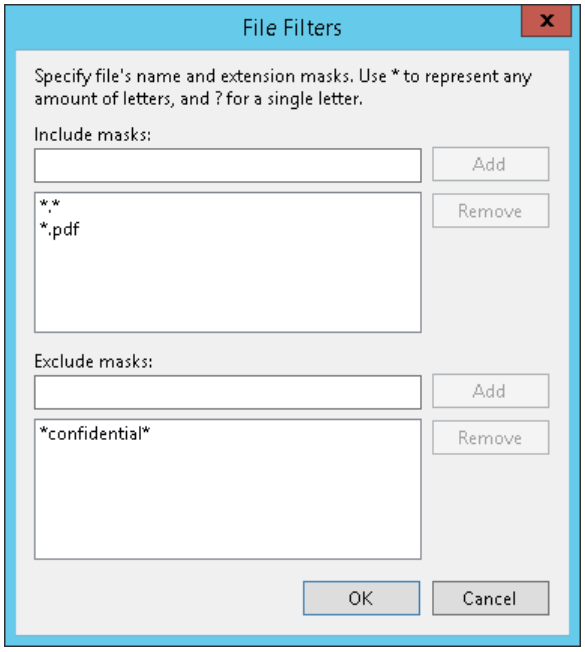

Step 4. Choose Media Pool for Full Backup

At the **Full Backup** step of the wizard, choose a media pool that will be used for archiving full backups of the selected files and create a schedule for full file backups.

1. From the **Media pool for full backups** list, choose a media pool that will be used for full file backups.
2. To schedule periodic creation of full file backups, select the **Run the full backup automatically** check box and specify the schedule according to which the job will create full file backups. If this option is disabled, you will need to start the job manually to create full backups of files.

TIP

If you have not previously created a media pool with the required settings, you can click the **Add New** button and create a new media pool without closing the job wizard. For more details, see Creating Media Pools.

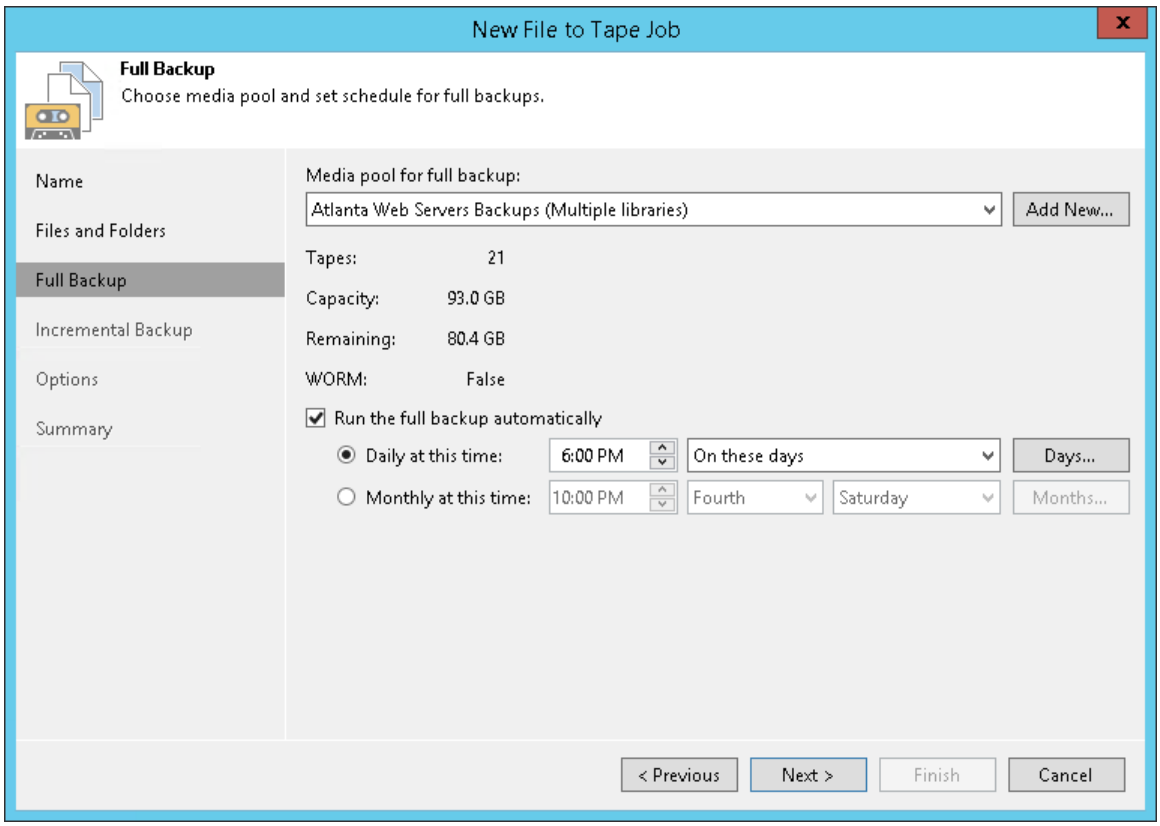

Step 5. Specify Media Pool for Increments

At the **Incremental Backup** step of the wizard, choose a media pool that will be used for archiving incremental backups of the selected files and create a schedule for incremental file backups.

1. From the **Media pool for incremental backup** list, choose a media pool that will be used for incremental file backups.
2. To schedule periodic creation of incremental file backups, select the **Run incremental backup automatically** check box and specify the schedule according to which the job will create incremental file backups. If this option is disabled, you will need to start the job manually to create incremental backups of files.

> **TIP**
>
> If you have not previously created a media pool with the required settings, you can click the **Add New** button and create a new media pool without closing the job wizard. For more details, see Creating Media Pools.

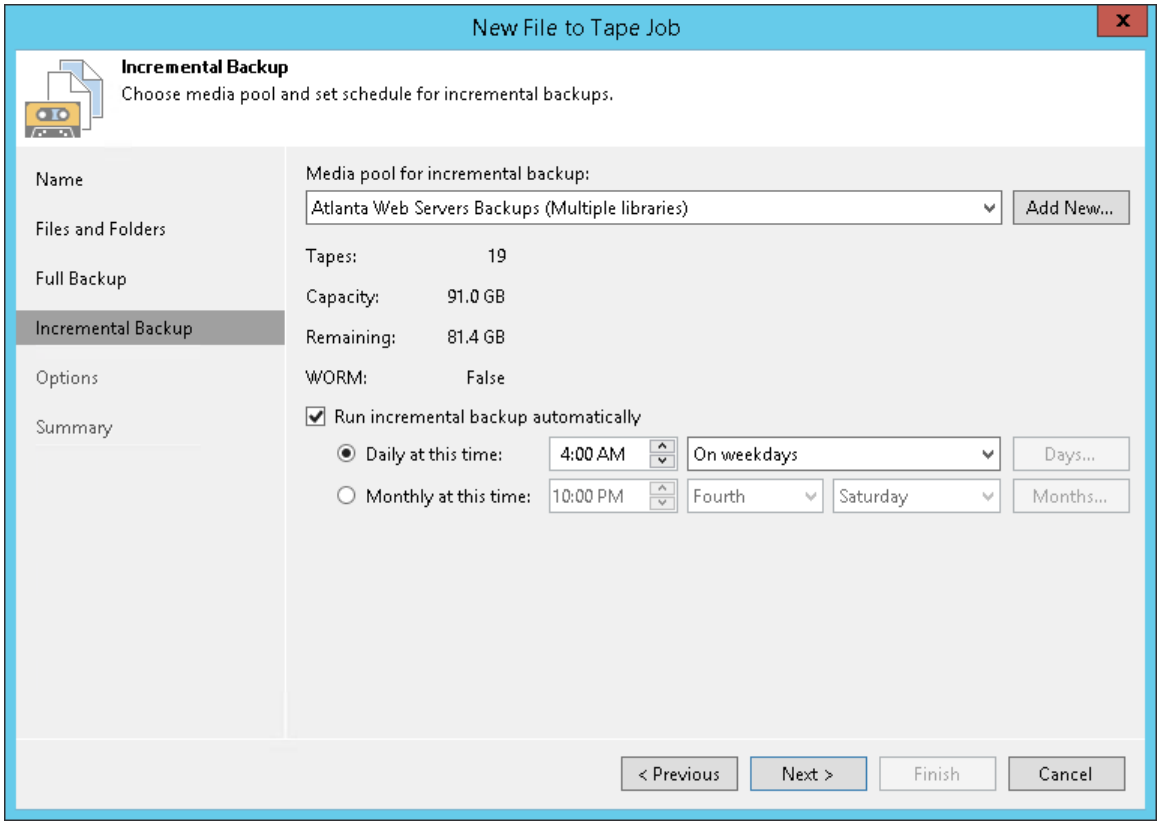

Step 6. Specify Archiving Options

At the **Options** step of the wizard, specify archiving and media automation options:

1. Select the **Use Microsoft volume shadow copy (VSS)** check box to enable backup of files with the help of Microsoft shadow volume copies. This option enables backup of files locked by application and provides file-level quiescence. This possibility can only be applied for files from servers running under Windows Server 2003 or later Windows-family OSs.
2. Select the **Eject media upon job completion** check box if the tape should be automatically ejected from the tape drive and placed into a slot when the job finishes.
3. Select the **Export current media set upon job completion** check box if you want to pull out the tapes with the current media set from the tape device, for example, to move to a storage location. The tape device will eject the tapes that belong to the current media set when the media set is closed.

   If you want to export tapes on specific days only (for example, every Saturday), click **Days** and schedule export on the necessary days.

   Note that with this option selected, a new media set is started after each export.

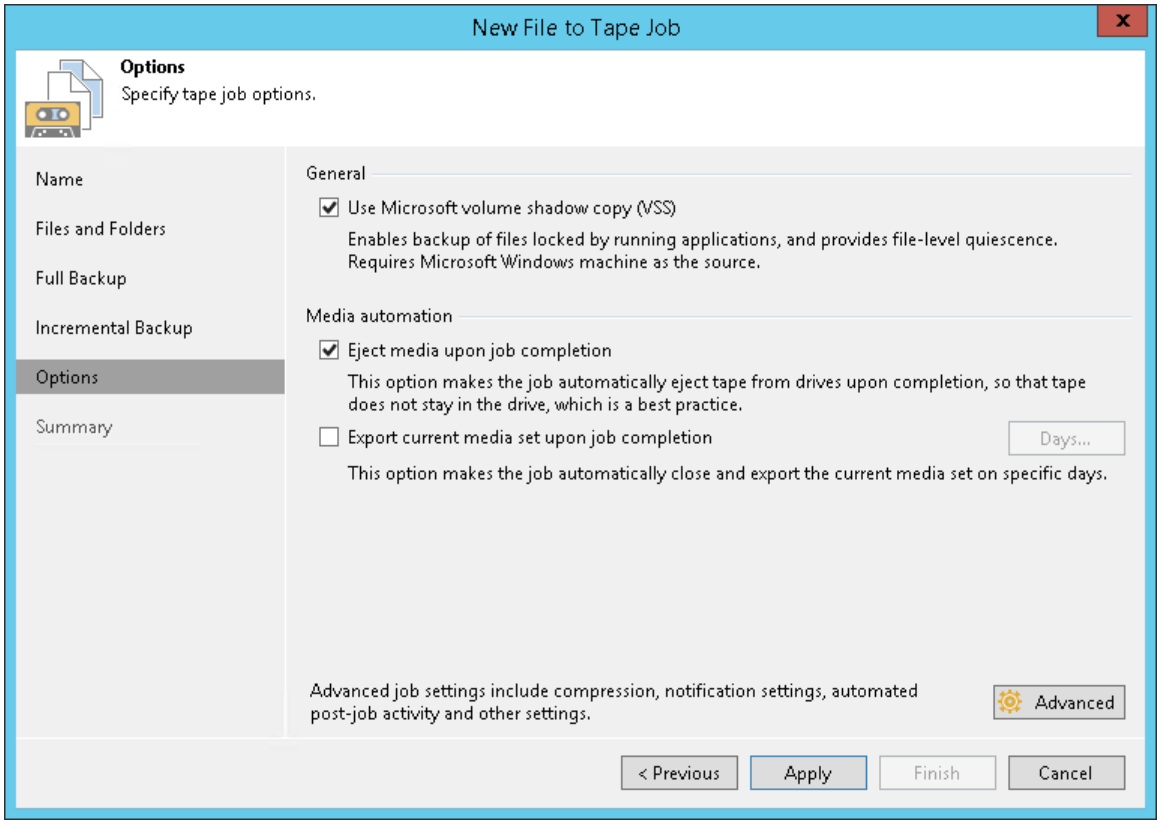

Step 7. Specify Advanced Job Settings

You can specify the following advanced settings for the file to tape job:

- Notifications Settings
- Advanced Settings

> **TIP**
>
> After you specify necessary settings for the file to tape 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 file to tape job, Veeam Backup & Replication will automatically apply default settings to the new job.

Notifications Settings

At the **Notifications** tab, you can specify notification settings for the file to tape job.

At the **Options** step of the wizard, click **Advanced**. Then select the **Notifications** tab.

Select the **Send email notifications to the following recipients** check box if you want to receive notifications about tape job status. In the field below, specify a recipient's email address. You can enter several addresses separated by a semicolon.

> **IMPORTANT**
>
> To receive notifications about tape jobs, you must enable general email notifications in Veeam Backup & Replication. For more information, see the Veeam Backup & Replication User Guide, section *Specifying Email Notification Settings*.

You can choose between the following options:

- **Use global notification settings**: Veeam Backup & Replication will notify you according to global email notification settings specified for the Veeam backup server.
- **Use custom notification settings specified below**: you can specify notification settings for tape jobs.

  In the **Subject** field, specify a notification subject. You can use the following variables in the subject:

  - *%Time%*: the time when the tape job finished,
  - *%JobName%*: the name of the tape job,
  - *%TapeCount%*: the number of tapes used for the tape job session,
  - *%JobResult%*: the job result,
  - *%VmCount%*: the number of machines in the job,
  - *%Issues%*: the number of machines in the job processed with Warning or Failed status.

Select the occasions on which you want to receive email notifications:

- **Notify on success**: Veeam Backup & Replication will notify you if the tape job finishes successfully,
- **Notify on warning**: Veeam Backup & Replication will notify you if the tape job finishes with a warning,
- **Notify on error:** Veeam Backup & Replication will notify you if the tape job fails,

- **Notify when waiting for tape**: Veeam Backup & Replication will notify you if the tape job cannot start because there are no available tapes.

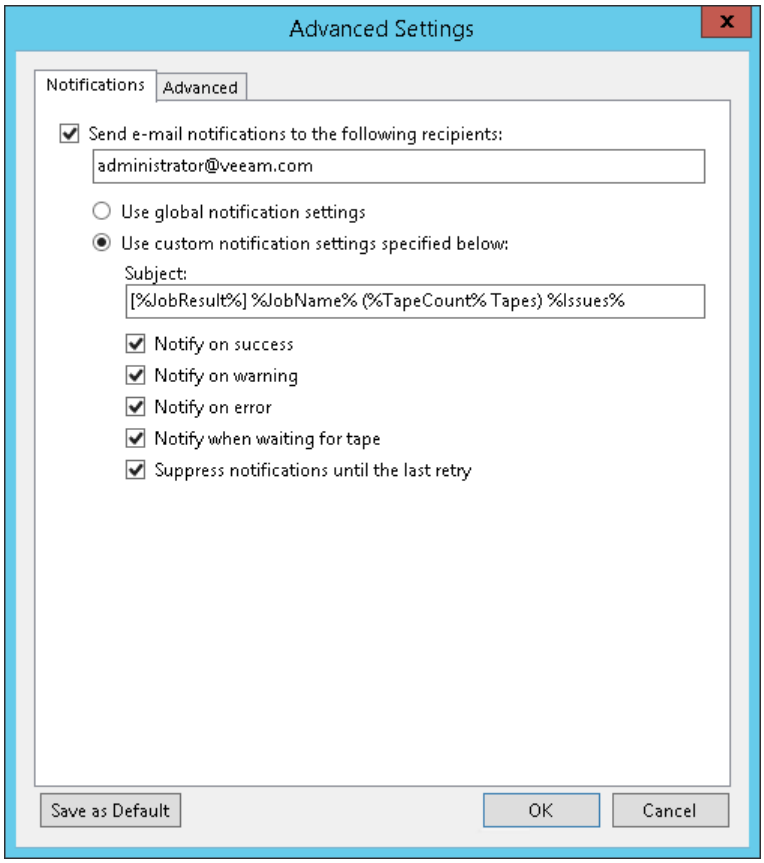

Advanced Settings

At the **Advanced** tab, you can enable hardware compression or specify custom scripts that you want to run before and/or after the file to tape job.

At the **Options** step of the wizard, click **Advanced**. Select the **Advanced** tab.

1. Select the **Use hardware compression if available** check box if the tape drive should compress data before writing it to tape.

   Most manufacturers add the hardware compression feature to their tape drives. This allows performing data compression by a drive. Usually the declared compression ratio for magnetic tape drives is 2:1 or 2.5:1. But in practice this ratio depends on the type of the raw data. Regrettably, Veeam Backup & Replication cannot display the compressed tape capacity. You can learn the size of the compressed data written to tapes only by checking how much free space left on the tape after the tape job finished.

   > **IMPORTANT**
   >
   > Enable the hardware compression only if your tape library provides support for hardware compression.

2. Select the **Run the following script before the job** and/or **Run the following script after the job** check boxes and click **Browse** to choose executable file(s). You can select to execute pre- and post-job actions after a number of job sessions or on specific week days.

   - If you select the **Run every... backup session** option, specify the number of the job sessions after which the script(s) must be executed.

- If you select the **Run on selected days only** option, click **Days** and specify week days on which the script(s) must be executed.

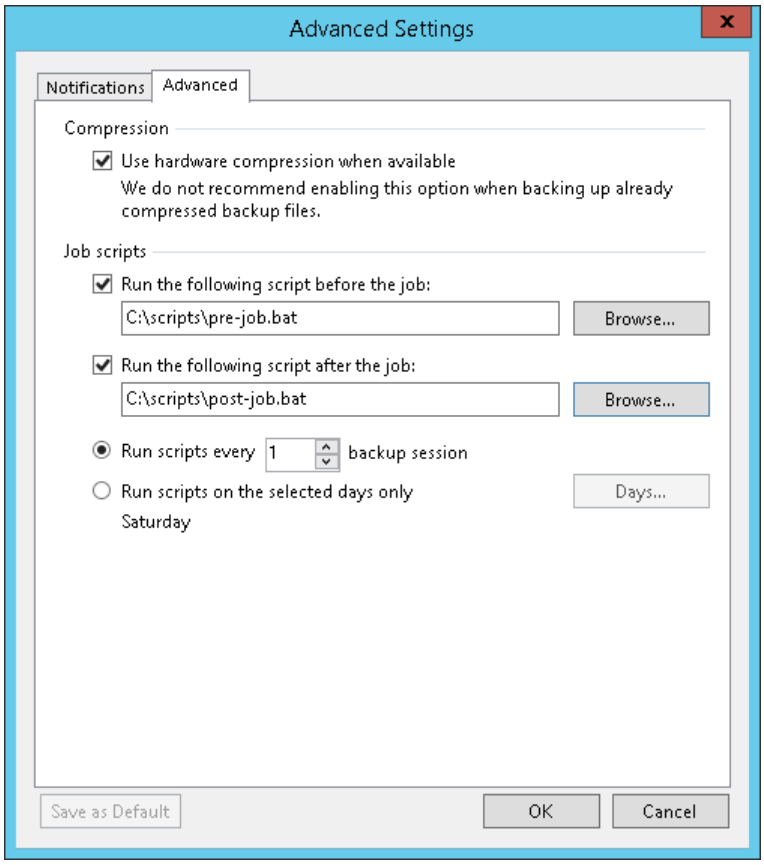

Step 8. Finish Working with Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start archiving file to tape job right after you complete working with the wizard. Click **Finish** to close the wizard.

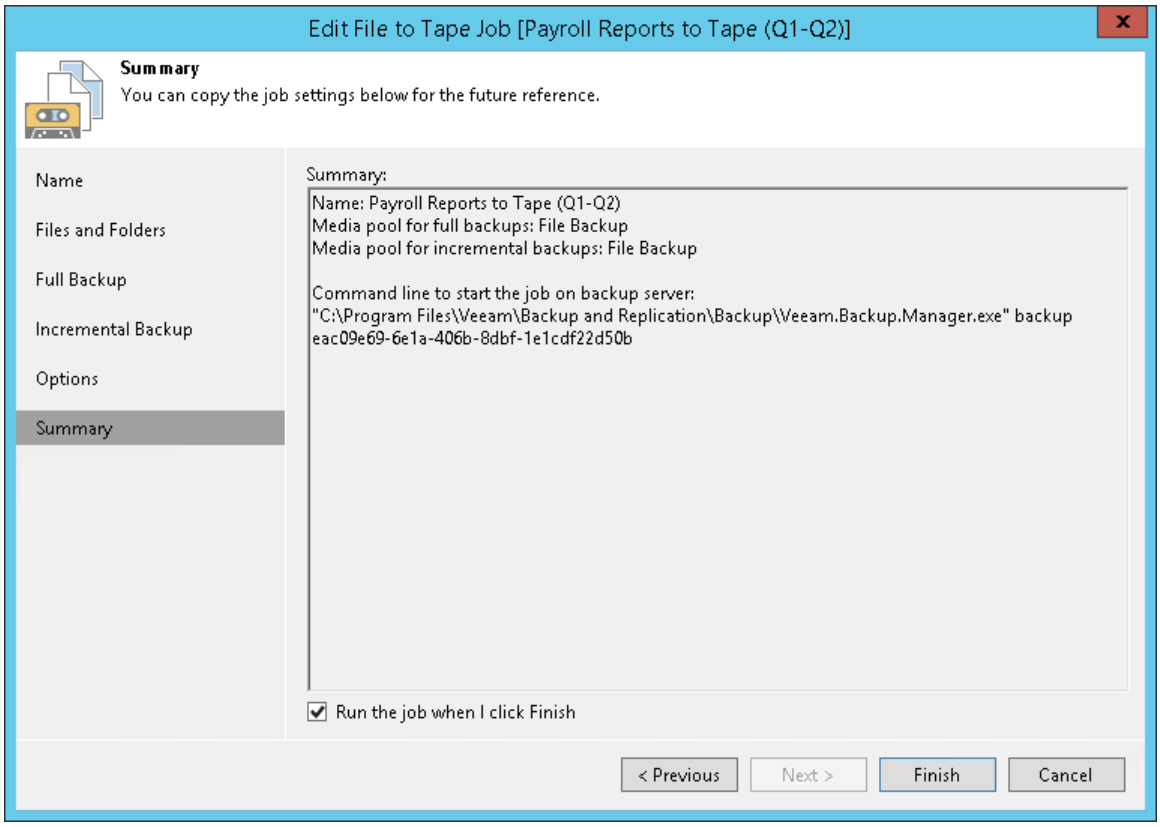

Viewing Files on Tape

After the file to tape job completes, you can view the created archive on tape:

1. Open the **Files** view and press F5 to refresh it.
2. Expand the **Tape** node and locate the files archive.

For Microsoft Windows and Linux file backups, Veeam Backup & Replication preserves the source hierarchy of folders for archived files.

For NMDP backups, you can view only volumes.

You can use the created archives for file recovery. For more information, see File Restore from Tape.

> **NOTE**
>
> The **Files on Tape** view displays only first 1000 items. For archive folders with more than 1000 items, you can use the search field to display the items you are looking for.

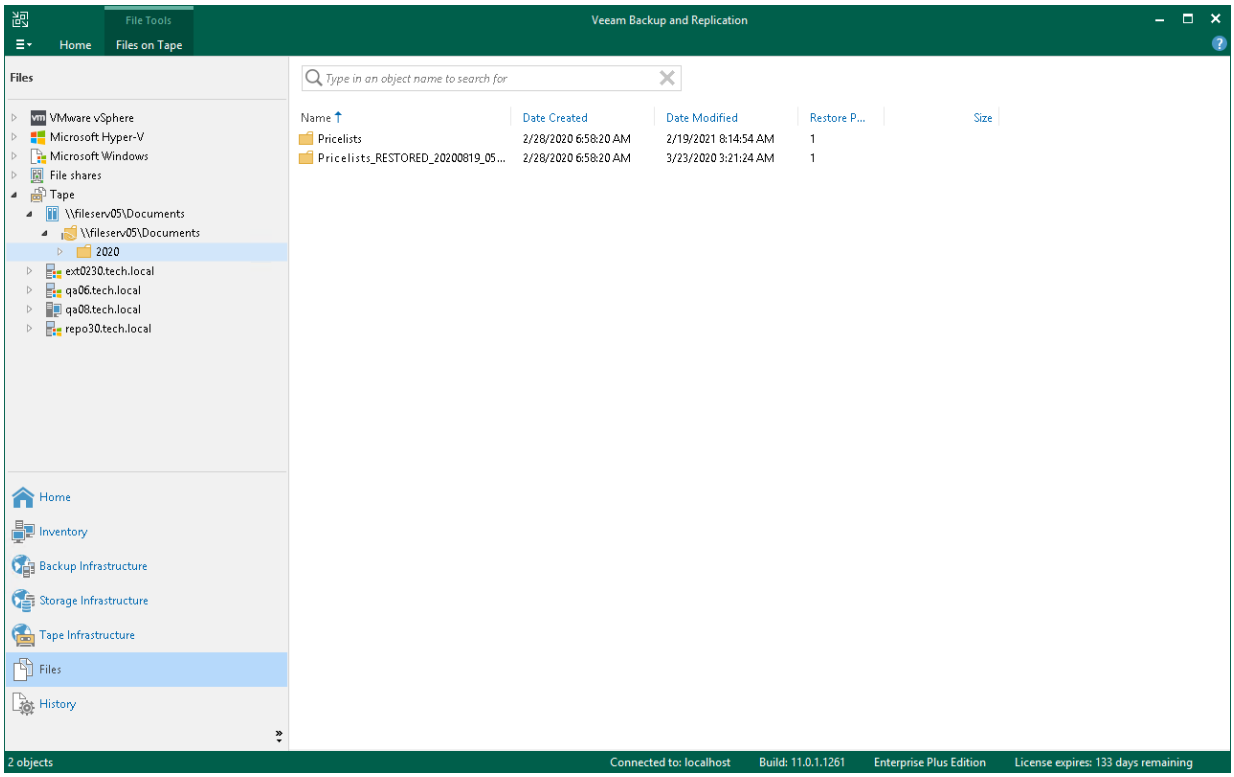

VM Restore from Tape to Infrastructure

Restoring a VM from tape with Veeam Backup & Replication is a lot like restoring a VM from disk. For example, you can choose a desired restore point, select the target location or change the configuration of the restored VM.

To restore a VM from tape, you can choose between the following options:

- Restore VMs directly to infrastructure
- Restore VMs through a staging repository

To choose the needed option, select **Restore directly to the infrastructure** or **Restore through the staging repository** at the **Backup Repository** step of the **Full VM Restore** wizard. For more information, see Choose Backup Repository.

You can restore only backups of VMware and Hyper-V virtual machines from tape to infrastructure. For backups of virtual machines running on other platforms and physical machine backups, use Backup Restore from Tape to Repository with further VM restore to infrastructure/VM file restore if necessary.

Restore Directly to Infrastructure

When you restore VMs from tape directly to the infrastructure, the restore process publishes the VMs to the virtual infrastructure copying the VM data directly from tape. This option is recommended if you want to restore one VM or a small number of VMs from a large backup that contains a lot of VMs. In this case, you do not need to provide a staging repository for a large amount of data most of which is not needed to you at the moment.

This option is not recommended in the following cases:

- If you want to restore a large number of VMs
- If you used parallel processing to record the VM backup.

In these cases, the restore may be very slow because the required blocks of data may be located on tapes randomly. The restore process may require a lot of rewinding of tapes as tapes do not provide random access to data.

> **NOTE**
>
> If the list of proxies to use for direct restore contains a Linux proxy, only the Appliance mode (HotAdd) is available for the source VM disks that will be processed using this Linux proxy. The Appliance mode (HotAdd) has a number of requirements and limitations described in this KB article.
>
> If disks cannot be processed in the Appliance mode (HotAdd), the VM restore fails.

Restore Through Staging Repository

When you restore VMs from tape through a staging repository, the restore process temporarily copies the whole restore point to a backup repository or a folder on disk. After that Veeam starts a regular VM restore.

This option is recommended if you want to restore a lot of VMs from a backup as the disk provides a much faster access to random data blocks than tape.

How Restoring VM from Tape to Infrastructure Works

Restore Directly to Infrastructure

For each VM, Veeam starts a separate restore job. However, you can select multiple VMs to restore: in this case, Veeam will queue the restore jobs and run them one after the other.

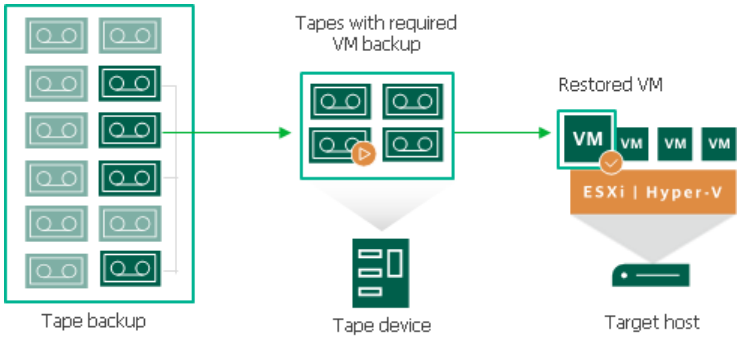

For restoring a VM directly to the infrastructure, Veeam Backup & Replication performs the following steps:

1. The VM from tape restore job checks the Backup Catalog in the Veeam Backup & Replication database to discover the tapes containing the needed backup files. If the tapes are offline, Veeam Backup & Replication prompts the user to insert the required tapes.

   Veeam needs to load and read the selected tapes for 3 times.

2. Veeam loads the tapes for the 1st time.
3. Veeam reads metadata and caches it on the tape server. This requires approximately 100Mb of disk space per 1TB of tape data for the default data block size.
4. Using the cached data, Veeam builds a map of data blocks. The map contains referenced to the data blocks of the VM you restore: the VM configuration file and the VM disks data.
5. Veeam loads the tapes for the 2nd time.
6. Veeam restores the VM configuration:
   - To restore to original location, Veeam synchronizes the VM configuration file up to the backup state.
   - To restore to another location, Veeam copies the VM configuration file and registers the VM on the target host.
7. Veeam loads the tapes for the 3rd time.
8. Veeam restores the VM disks. Multiple disks are restored parallelly. Veeam reads the tape consequently and, using the map of data blocks, copies the VM disks data.

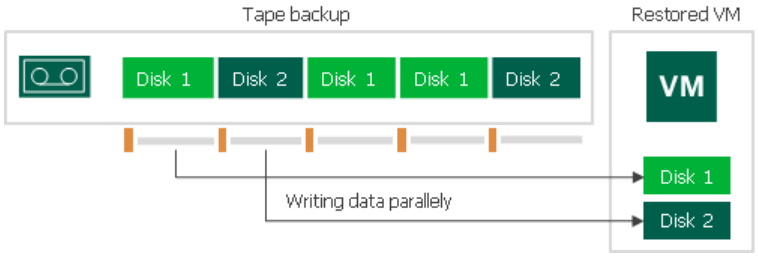

Restore Through Staging Repository

For restoring through a staging repository, Veeam Backup & Replication performs the following steps:

1. Veeam starts 2 jobs at background:
   - Backup restore job: temporarily restores the backup to the repository.
   - VM restore job: restores the VM to the infrastructure.

   The jobs start simultaneously. The VM restore job starts and remains pending until the backup restore completes.
2. The backup restore job checks the Backup Catalog in the Veeam Backup & Replication database to discover the tapes containing the needed backup. If the tapes are offline, Veeam Backup & Replication prompts the user to insert the required tapes.
3. The backup restore job reads the tapes consecutively and copies the backup to the selected staging repository or folder.
4. When the backup is copied, Veeam registers it temporarily as an imported backup.
5. The backup restore job finishes.
6. The VM restore job receives information about the backup restored successfully and launches a standard VM restore process. For more information, see the Veeam Backup & Replication User Guide, section *Full VM Recovery*.
7. When the VM is successfully restored, the VM restore process finishes.
8. Veeam deletes the backup from the staging repository and from disk.

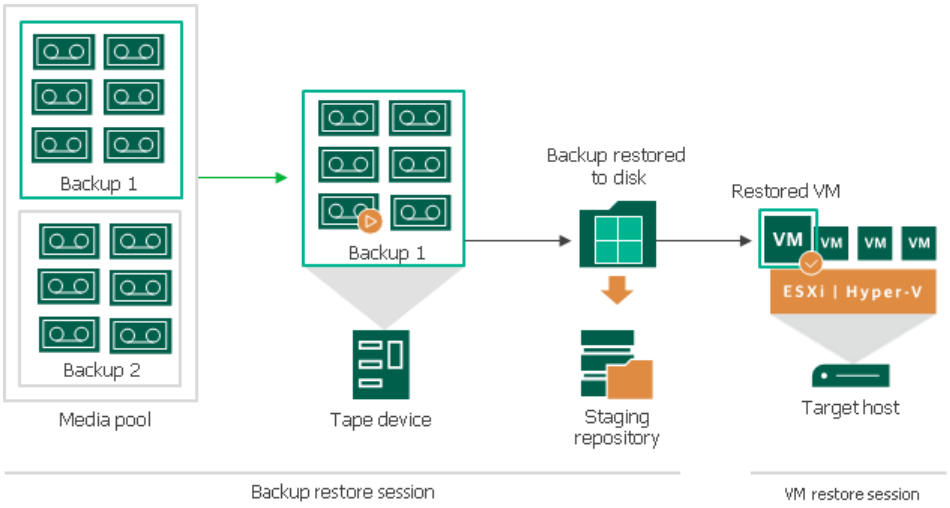

Before You Begin

For restoring VMs from tape, you must provide the following system resources:

- Tape server: 200 MB RAM for each restored VM disk.

Keep in mind that if a job is unable to complete within 21 days period, it will be stopped with the 'Failed' status.

Restoring VM from Tape to Infrastructure

To restore VMs from tape to virtual infrastructure, use the **Full VM Restore** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

If the tapes are online in your tape device, Veeam Backup & Replication will get the needed data automatically. If the tapes are offline, Veeam Backup & Replication will list the names of the required tapes.

To restore VMs from tape, follow the next steps:

Step 1. Launch Full VM Restore Wizard

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

- Open the **Home** view, expand the **Backups** > **Tape** node. Choose and expand the tape backup you need. Select the necessary VMs in the tape backup and click **Restore Entire VM** on the ribbon.
- Open the **Home** view, expand the **Backups** > **Tape** node. Choose and expand the tape backup you need. Right-click the necessary VMs in the tape backup and choose **Restore entire VM**.

Step 2. Choose Virtual Machines to Restore

At the **Virtual Machines** step of the wizard, review VMs that should be restored. To add one or more VMs to the list, click **Add VM** and select where to browse for the machines:

- **From infrastructure** – this option is available for VMs only. Browse the virtual environment and select VMs to restore. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list. To quickly find a VM, use the search field at the top of the list: enter the VM name or a part of it and click the search button on the right or press **[ENTER]**.

  Make sure that VMs you select from the virtual environment have been successfully archived to tape at least once.

- **From backup** – browse existing backups and select VMs under backup to tape jobs. To quickly find VMs, use the search field at the bottom of the **Select Objects** window: enter a VM name or a part of it and click the **Start search** button on the right or press **[ENTER]**. Full Restore from backups

> **NOTE**
>
> You can add VMs only of the same virtualization platform.

To remove a VM, select it in the list and click **Remove** on the right.

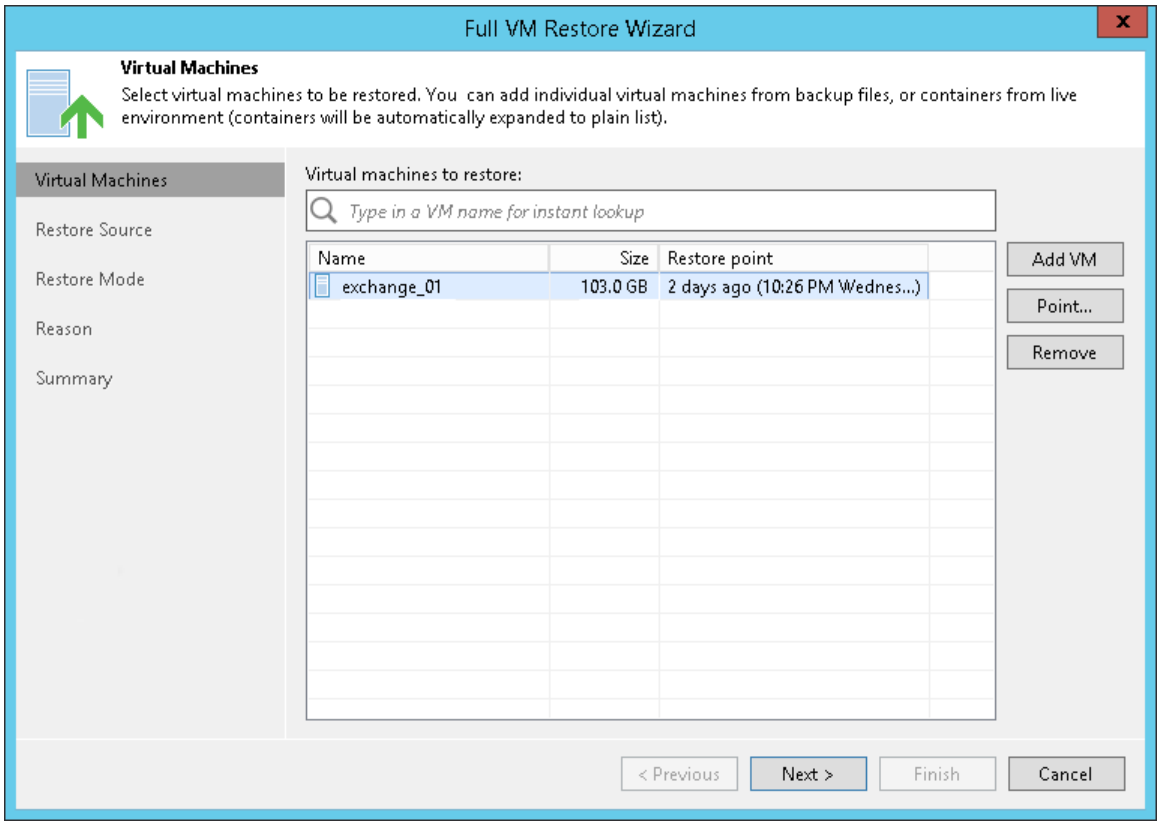

Step 3. Select Restore Point

By default, Veeam Backup & Replication will restore VMs to their latest state archived to tape. However, if you want to restore a VM to an earlier state, select a VM in the list and click Point on the right. In the **Restore Points** section, select a restore point that should be used for full VM recovery.

If you have chosen to restore multiple VMs, you can select a different restore point for every VM specifically.

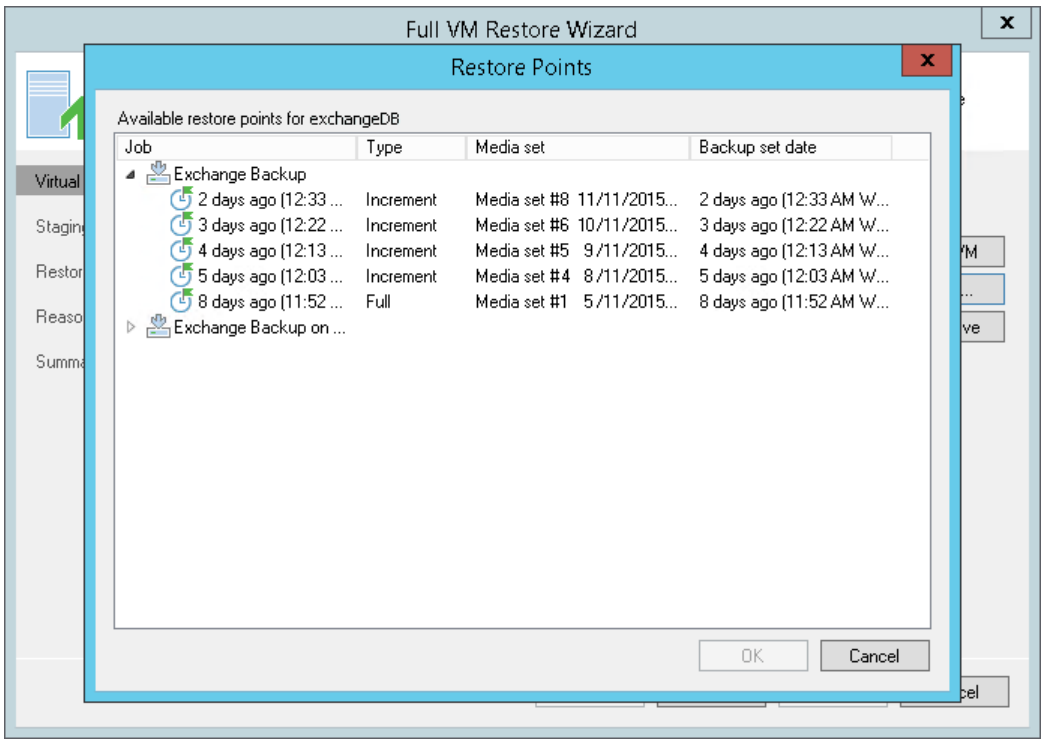

Step 4. Choose Restore Source

At the **Restore Source** step of the wizard, choose the restore mode:

- Select **Restore directly from tape** if you want to restore a VM without a staging repository. In this case, the restore job will read data directly from tape.

  Note that restoring VMs directly to the infrastructure is recommended for restore of 1 VM at a time. The process may be slow if you restore a lot of VMs simultaneously.

- Select **Restore through a staging repository** if you want to temporarily copy the tape backup to a repository or a folder first. The restore job will copy the backup to the selected destination and launch a standard restore VM from disk process.

  From the **Backup repository** list, select the repository that should be used as a temporary storage (staging area) for machine backup before the VM is restored to the virtual infrastructure.

  You can also select a target folder on any server connected to Veeam Backup & Replication. To do so, choose the **Select folder** option from the **Backup repository** list and choose the location to which backups should be restored before full VM recovery.

  If you choose to restore files to a shared folder, make sure that the account under which Veeam Backup Service runs has write permissions to the target folder. If the account does not have sufficient permissions, Veeam Backup & Replication will prompt you to enter credentials for the account that can be used for writing to the target folder.

NOTE

If you plan to restore a VM directly from tape using NBD transport mode, make sure that the VM has 28 disks maximum. To restore a VM with more than 28 disks, use HotAdd mode or restore through a staging repository.

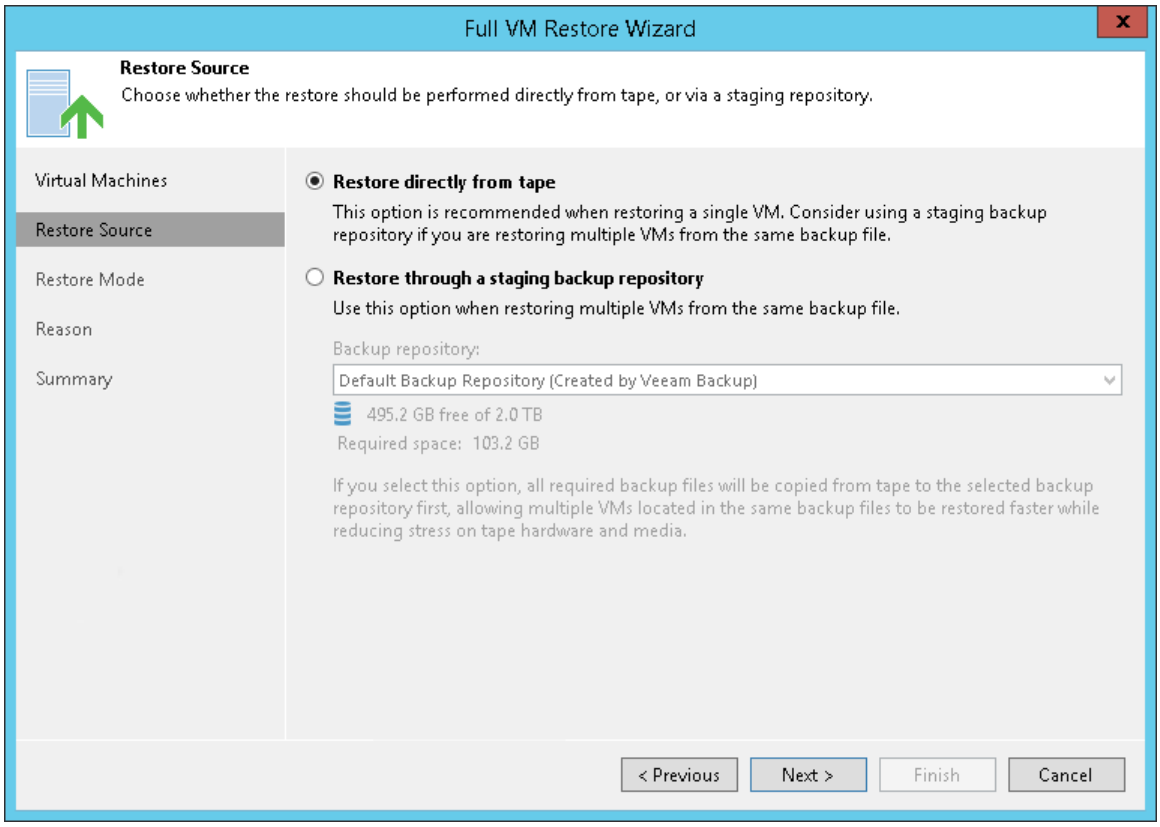

Step 5. Specify Restore Mode and Other Recovery Options

Go through the remaining steps of the **Full VM Restore** wizard. The procedure is identical to full VM recovery. For more information, see the Veeam Backup & Replication User Guide, section *Performing Full VM Restore*.

Backup Restore from Tape to Repository

This option allows you to copy machines backups from tape to repository. This is helpful if you need some backups on disk for later use, or also for machines guest OS files restore. You can restore full backups or incremental backups to a repository or any location of your choice. The restored backup is registered in the Veeam Backup & Replication console as an imported disk backup so that you can use it for any restore from disk scenario later on. For one restore session at a time, you can choose one restore point available on tape.

> **NOTE**
>
> Restore of databases from backups restored from tapes to the repository is not supported.

How Restoring Backups from Tape to Repository Works

For restoring machine backups to a repository or a folder on disk, Veeam Backup & Replication performs the following steps:

1. Veeam checks the Backup Catalog in the Veeam Backup & Replication database to discover the tapes containing the required backup. If the tapes are offline, Veeam Backup & Replication prompts the user to insert the required tapes.
2. The backup restore job locks one tape drive to process the restore of the whole backup.
3. The tape drive loads the required tapes one by one, reads them and copies the backup data to the selected repository or folder.
4. When the backup is copied, Veeam registers it as an imported backup.

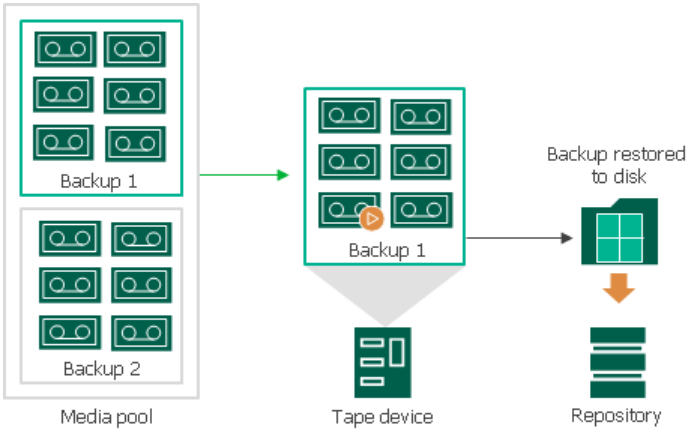

Restoring Backups from Tape to Repository

To restore backups from tape, use the **Restore Backup from Tape to Repository** wizard. This section will guide you through all steps of the wizard and provide explanation on available options.

If the tapes are online in your tape device, Veeam Backup & Replication will get the needed data automatically. If the tapes are offline, Veeam Backup & Replication will list the names of the required tapes.

Keep in mind that if a job is unable to complete within 21 days period, it will be stopped with the 'Failed' status.

To restore files from tape, follow the next steps:

Step 1. Launch Restore Backup from Tape to Repository Wizard

To run the **Restore Backup from Tape to Repository** wizard, on the **Home** tab, click **Restore** and choose **Tape** > **Restore Backups.**

Alternatively, you can:

- Open the **Home** view, expand the **Backups** > **Tape** node. Select the necessary machines in backup and click **Restore to Repository** on the ribbon.
- Open the **Home** view, expand the **Backups** > **Tape** node. Right-click the necessary machines in backup and choose **Restore backup from tape to repository**.

Step 2. Choose Machines to Restore

At the **Source** step of the wizard, select one or more physical or virtual machines for which backup files should be restored. If you have chosen machines to restore from archives on tape, the list of objects to restore will be populated with selected machines.

To add one or more machines to the list, click **Add VM** and select where to browse for the machines:

- **From vSphere Infrastructure** – this option is available for VMs only. Browse the virtual environment and select VMs to restore. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list. To quickly find a VM, use the search field at the top of the list: enter the VM name or a part of it and click the search button on the right or press **[ENTER]**.
  Make sure that VMs you select from the virtual environment have been successfully archived to tape at least once.
- **From backups** – browse existing backups on tape and select machines under backup to tape jobs. To quickly find machines, use the search field at the bottom of the **Select Objects** window: enter a virtual or physical machine name or a part of it and click the **Start search** button on the right or press **[ENTER]**.

To remove a machine, select it in the list and click **Remove** on the right.

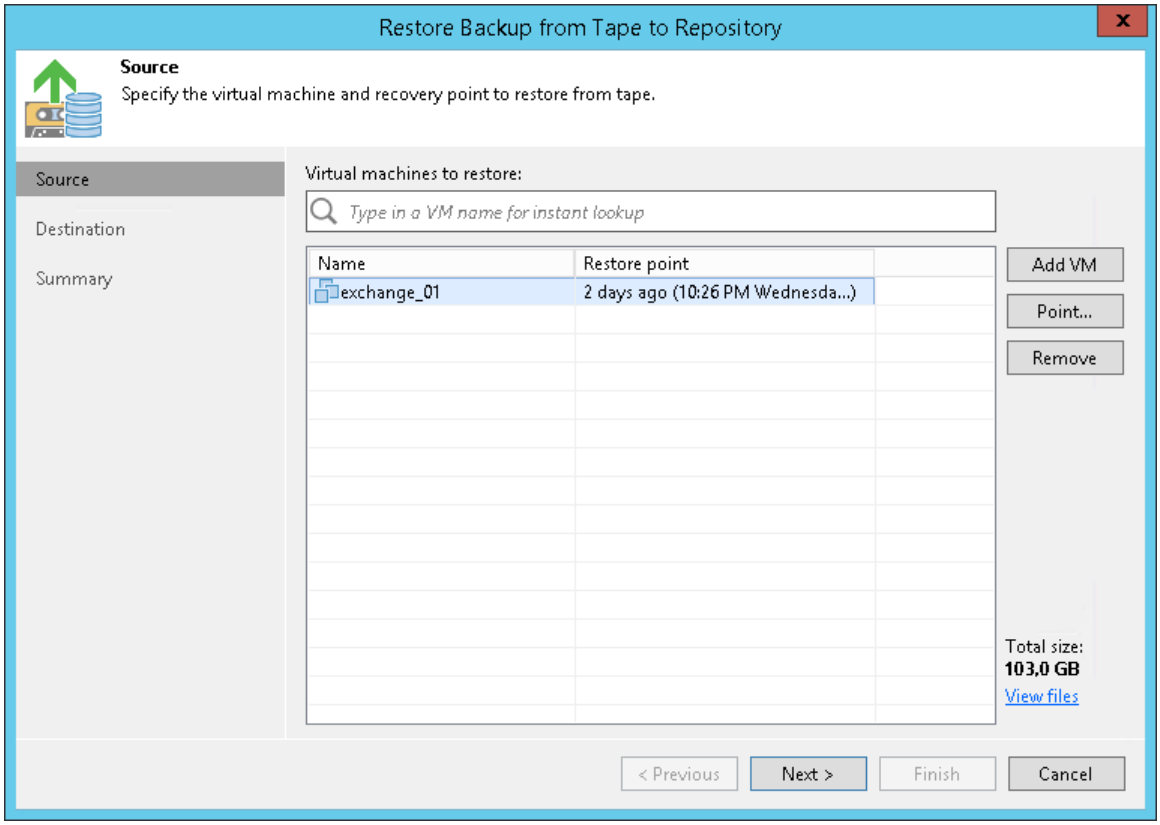

Step 3. Select Restore Point

By default, Veeam Backup & Replication will restore backup with the latest state of the archived machine. However, if you want to restore a backup for the machine to an earlier state, select a machine in the list and click **Point** on the right. In the **Restore Points** section, select a restore point that should be used to restore machine backup.

> **NOTE**
>
> If you choose a full backup point in the list, Veeam Backup & Replication will restore only this full backup. If you choose an increment, Veeam Backup & Replication will restore a chain consisting of a full backup and forward increments, necessary to restore machines to the required point-in-time.

If you have chosen to restore multiple machines, you can select a different restore point for every machine specifically.

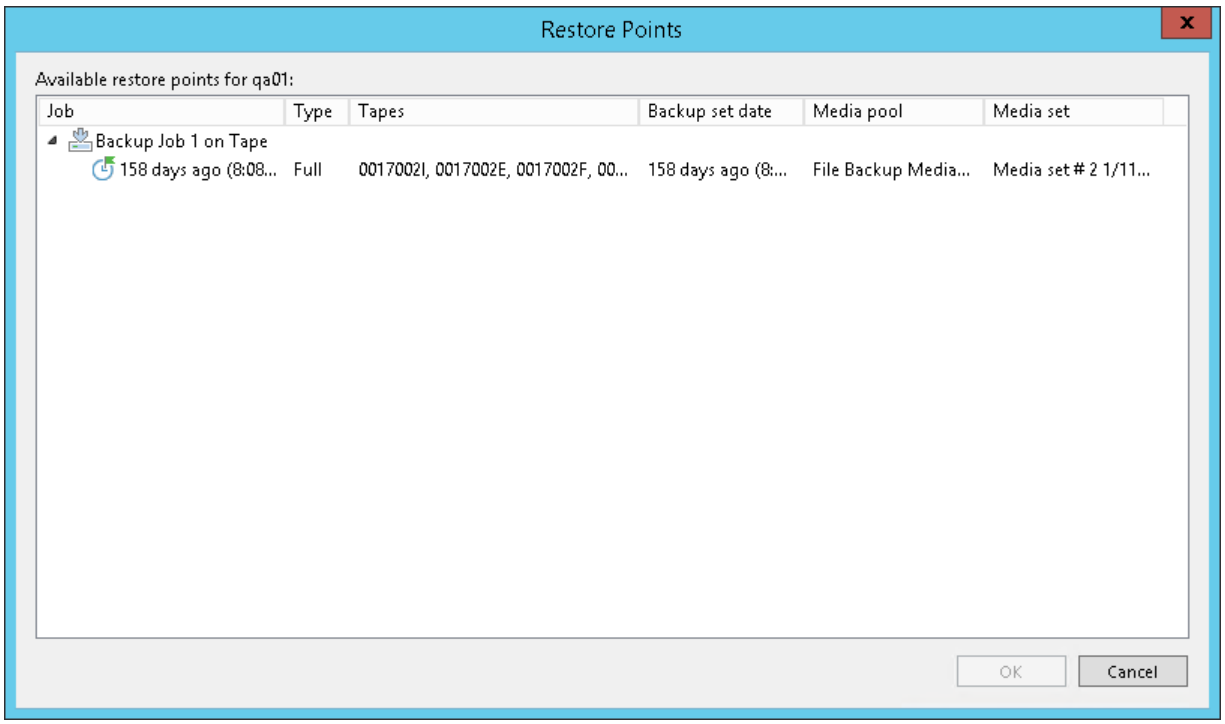

Step 4. Choose Backup Destination

At the **Destination** step of the wizard, select where the backup files for the selected machines should be restored:

- To restore machine backup files to a repository, select the **Backup repository** option and choose the necessary repository from the list.
- To restore machine backup files to the Veeam backup server, shared folder or to any Microsoft or Linux server connected to Veeam backup server, select the **Server** option. Choose the necessary server from the list and specify path to the target folder in the **Path to folder** field.

  If you choose to restore files to a shared folder, make sure that the account under which Veeam Backup Service runs has write permissions to the target folder. If the account does not have sufficient permissions, Veeam Backup & Replication will prompt you to enter credentials for the account that can be used for writing to the target folder.

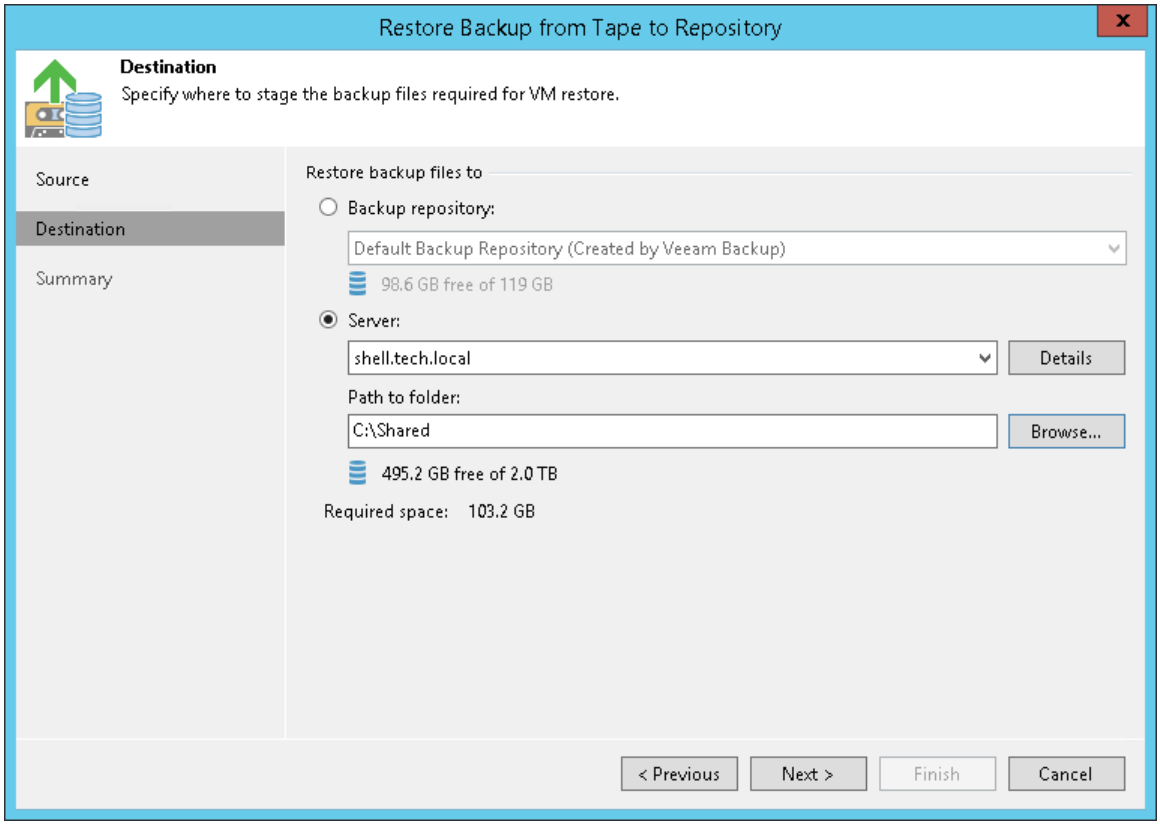

Step 5. Finish Working with Wizard

Review the settings and click **Finish** to restore selected machines backups.

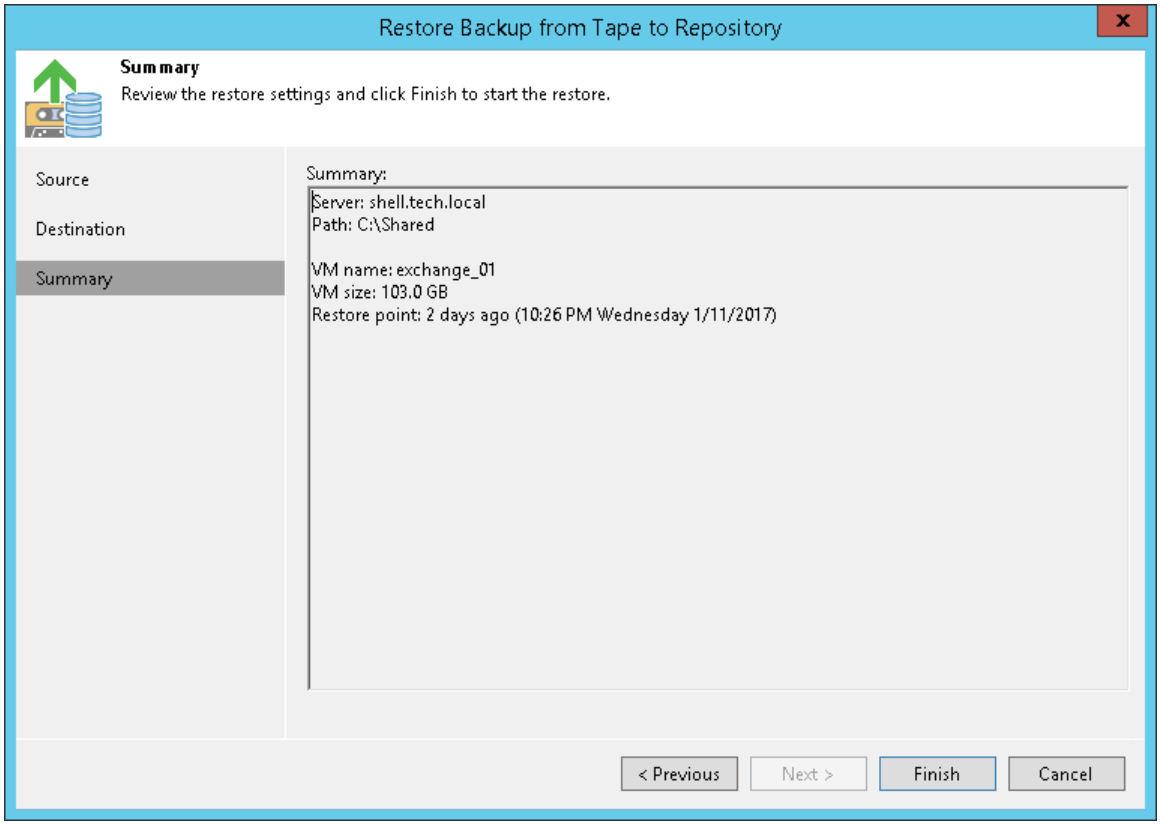

Viewing Restored Backups

After backups are restored from tape, they are displayed as imported backups in the **Home** view > **Backups** > **Imported**. You can use the restored to disk backup for regular data recovery, including full recovery, recovery of machine files, guest OS files recovery and other.

Tenant Restore from Tape

This option is available only for Veeam Cloud Connect service providers.

Service providers can restore their tenants' backups from tape. For more information, see the Veeam Cloud Connect guide.

File Restore from Tape

You can restore data previously archived with file to tape jobs. You can restore the following data:

- Microsoft Windows files and folders.
- Linux files and folders.
- Files and folders on SMB (CIFS) and NFS file shares.
- NDMP volumes.
- Content of entire tapes.

The file restore process allows you to restore files to any restore point available on tape.

Limitations for NFS File Share Restore

If you restore files to the NFS file share running the NFS protocol version different from the NFS version of the source file share, Veeam Backup & Replication does not restore security attributes.

Limitations for NDMP Volumes Restore

- Only restore of whole volume is supported.
- Only restore to original location or to a similar NDMP server is supported.

How Restoring Files from Tape Works

When Veeam Backup & Replication restores files archived to tape, it performs the following operations:

1. Veeam Backup & Replication checks the Tape Catalog in the Veeam Backup & Replication database to discover the tapes containing the needed restore point of the files. If the tapes are offline, Veeam Backup & Replication prompts the user to insert the required tapes.
2. Veeam Backup & Replication connects to Veeam Data Mover deployed on a tape server.
3. Veeam Data Mover copies the relevant files from the tapes.
4. The copied files are transferred to a chosen target location, same or new one.
5. Veeam Backup & Replication updates the Tape Catalog in the Veeam Backup database.

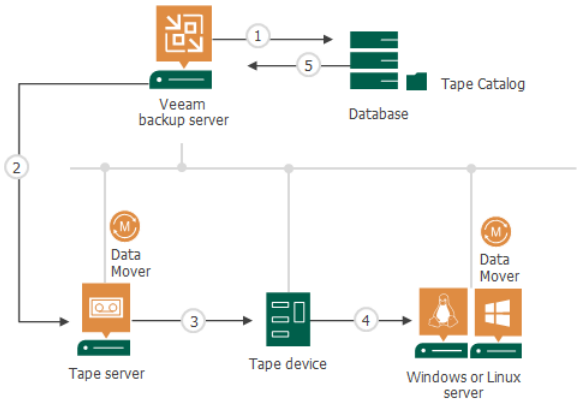

Before You Begin

If the tape backup contains a large quantity of files, for example, more than 1 000 000 files in 1 000 folders, you must provide the following system resources for restore: backup server — 2.6 GB RAM per each 1 000 000 files.

Keep in mind that if a job is unable to complete within 21 days period, it will be stopped with the 'Failed' status.

Restoring Files from Tape

To restore files or NDMP volumes backed up to tape, use the **Files from Tape Restore** wizard.

If the tapes are online in your tape device, Veeam Backup & Replication will get the needed data automatically. If the tapes are offline, Veeam Backup & Replication will list the names of the required tapes.

> **NOTE**
>
> The **Files from Tape Restore** wizard allows you to restore files and folders archived to tape. You cannot restore VM guest OS files using this wizard. To restore VM guest OS files, restore a machine backup from tape to a backup repository and perform VM guest OS files restore. For more information, see Backup Restore from Tape to Repository.

Restoring Entire Tapes

For information about restoring entire tapes, see Restoring Entire Tapes.

This section will guide you through all steps of the wizard and provide explanation on available options for restoring files from tapes and for restoring entire tapes.

To restore files from tape or to restore entire tapes, follow the next steps:

Step 1. Launch Restore from Tape Wizard

The wizard may work in one of the two following modes:

- Restoring files from tape
- Restoring entire tapes

Restoring Files from Tape

To run the **Restore from Tape** wizard, do either of the following:

- Open the **Home** tab and click **Restore** on the ribbon. Choose **Tape > Restore Files**.
- Open the **Files** view, expand the **Tape** node and browse to the necessary files on tape. Select the files and click **Restore Files from Tape** on the ribbon. Alternatively, you can right-click the files and choose **Restore files from tape**.

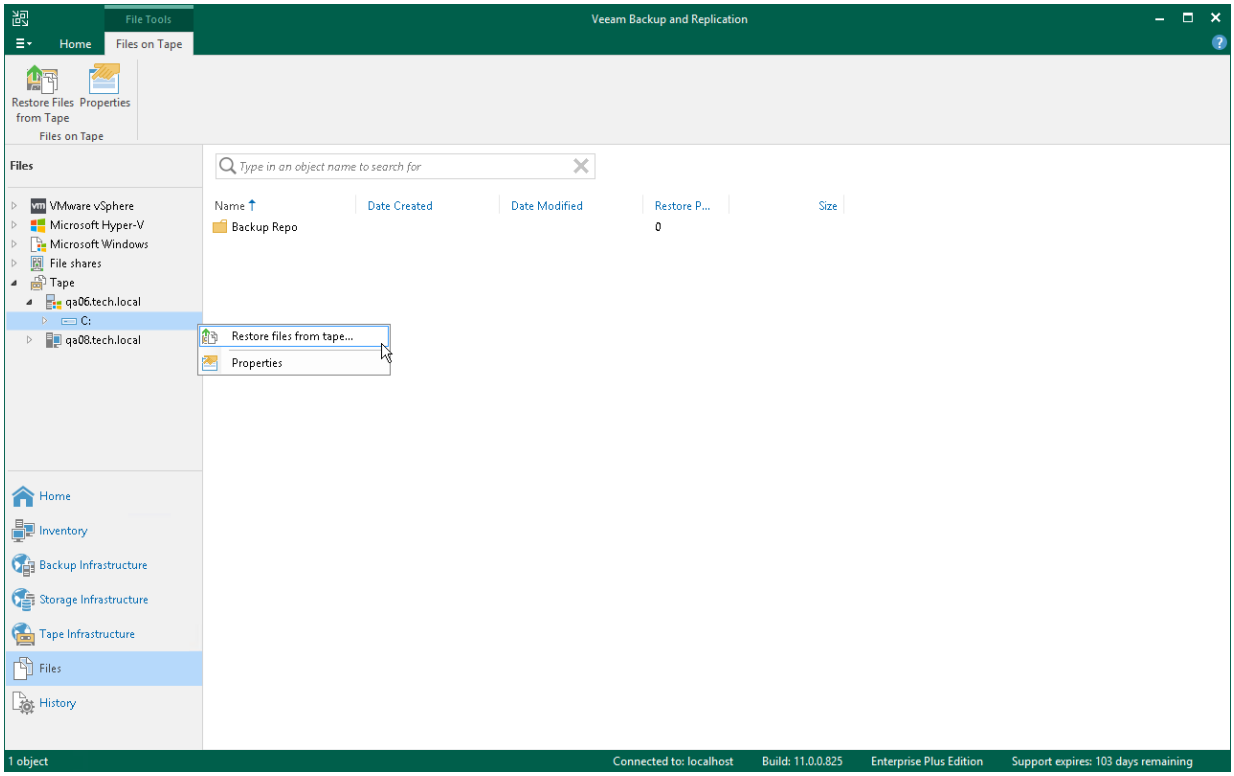

Restoring Entire Tapes

To launch the **Restore from Tape** wizard for restoring entire tapes, do the following:

1. Open the **Tape Infrastructure** view.
2. Navigate to the list of tapes either under **Media Pools** or under **Libraries** > *LibraryName* node > **Media** > **Online**.

> **NOTE**
>
> Tapes in the **Unrecognized** media pool are not available for data restore.

3. Select the necessary tapes and click **Restore Content** on the ribbon. Alternatively, you can right-click the tape and select **Restore content**.

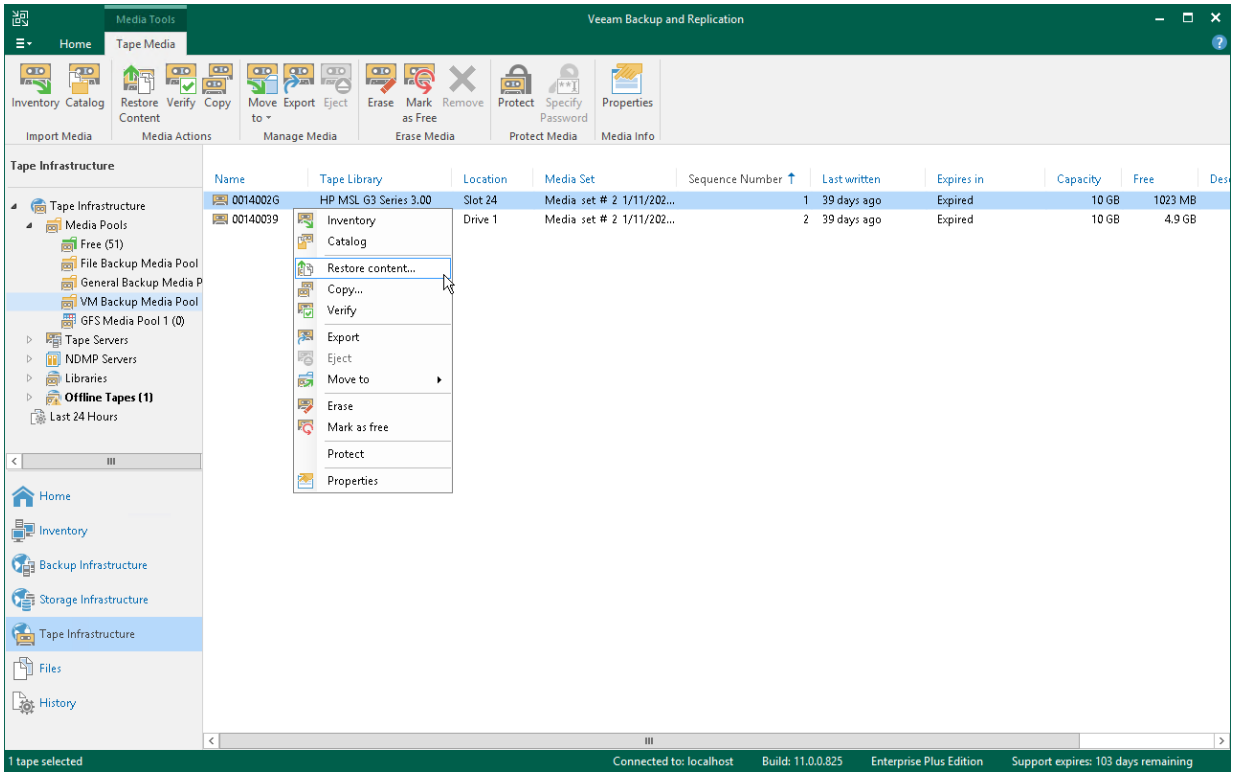

Step 2. Choose Files to Restore

The objects to restore depend on the mode that the wizard is working in:

- Restoring files from tape
- Restoring entire tapes

Restoring Files from Tape

At the **Objects to Restore** step, choose data that you want to restore.

Click **Add** and browse to the file or folder, NDMP volume, or entire host (except for Linux) that you want to restore. The selected item will be added to the list. To quickly find file, folder or volume, use the search field at the top of the list: enter an object name or a part of it and click the search button on the right or press **[ENTER]**.

If you have chosen files to restore in the **Files** view, the list of objects to restore will be populated with selected files.

To remove a file or folder from the list, select it and click **Remove**.

> **NOTE**
>
> For one restore session, you can select only Microsoft Windows and Linux files, or NDMP volumes.

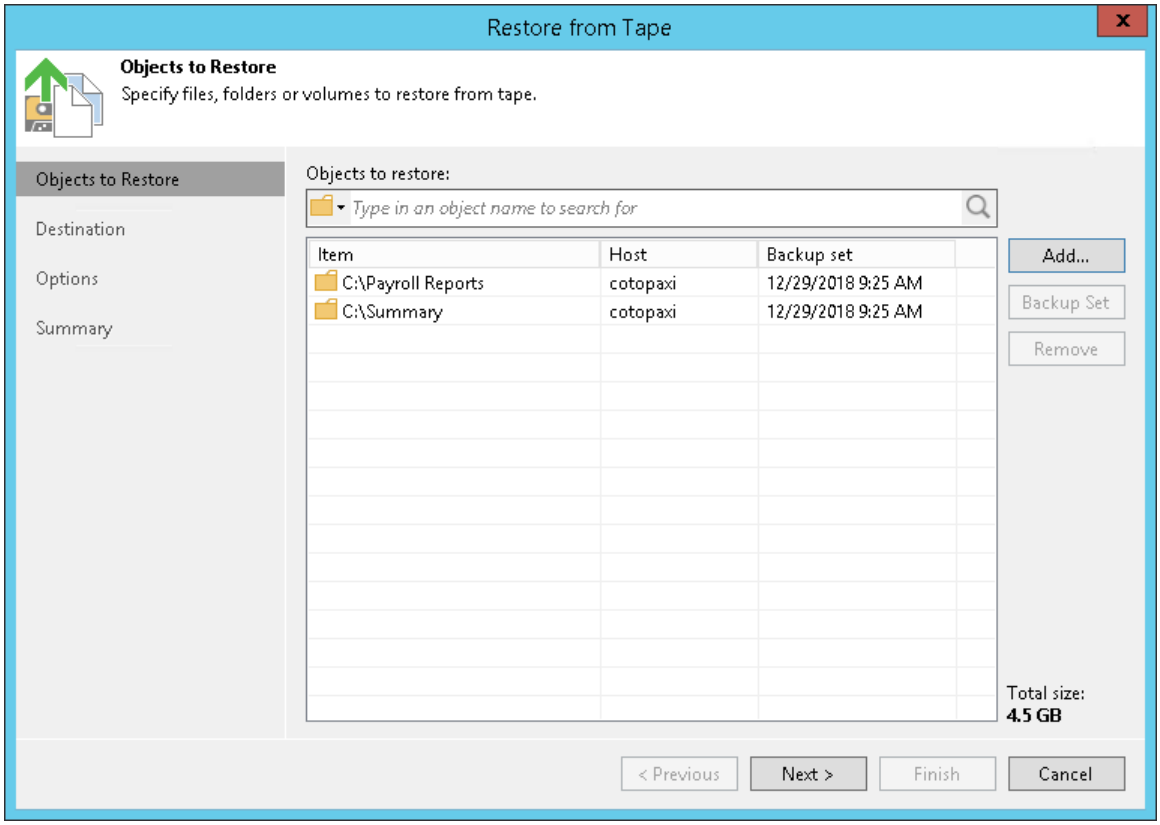

By default, Veeam Backup & Replication will restore the latest version of files available on tape. If you want to restore files from another restore point, select the necessary file and click **Backup Set**. In the list of available backup sets, select the necessary archiving session and click **OK**.

NOTE

For backups created with Veeam Backup & Replication 9.5 Update 3a and earlier: when you restore an entire folder from tape, Veeam Backup & Replication restores all files that have ever existed in the folder and been archived to tape. To restore files in the folder, Veeam Backup & Replication scans the selected backup set and backup sets that were created previously. If a file is not found in the selected backup set, Veeam Backup & Replication will restore the most recent version of the file from a backup set preceding the selected backup set.

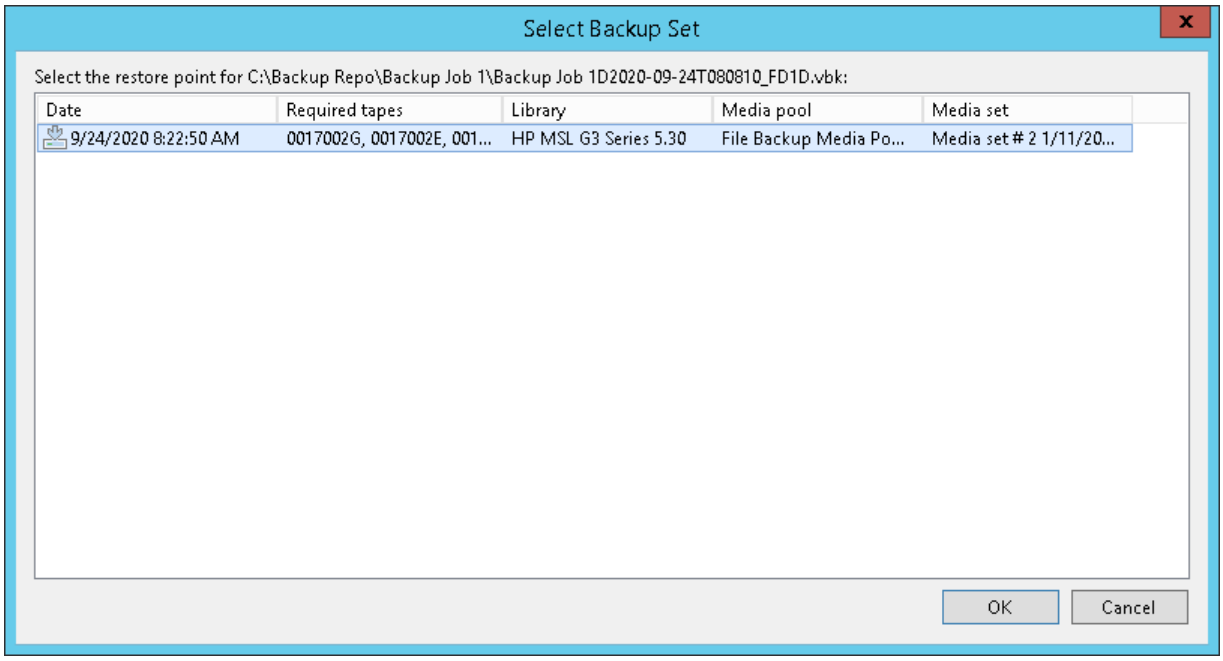

NOTE

Files written to tapes can be further removed from the source backup repository. Information about these removed files (about the fact that they are removed) is also written to tapes. If the tapes that store this information are rewritten as expired or erased by some reason, during the file restore from tape, Veeam Backup & Replication tries to restore the files removed from the source backup repository. Therefore, Veeam Backup & Replication can require even older tapes that store the copy of these removed files.

Restoring Entire Tapes

At the **Objects to Restore** step of the wizard, select tapes to restore. The tape, which you selected when launching the wizard, is added to the list of tapes to restore by default. To add other tapes to restore, do the following:

1. Click **Add**. Veeam Backup & Replication opens the **Select Tapes** window listing all media pools and tapes in them available for restore.
2. In the **Select Tapes** window, navigate to the required media pool and select the tapes you want to restore.

When writing backup files to tapes, there can be situations when a file does not fit on a tape. Then it is divided into parts and written to several tapes, which are then considered dependent. Veeam Backup & Replication tracks this. When you are restoring data from a tape that contains only one part of the backup file, Veeam Backup & Replication detects that. It also informs you what tapes contain other parts of the file and prompts you to add these tapes to the list for restoring the full content. Mind that tapes marked as **Restored partially** in the **Restore Options** column will be restored by Veeam Backup & Replication only partially: it will restore only data required for restoring full content of the initially selected tapes in the consistent manner.

To remove tapes from the restore list, select the tapes and click **Remove**.

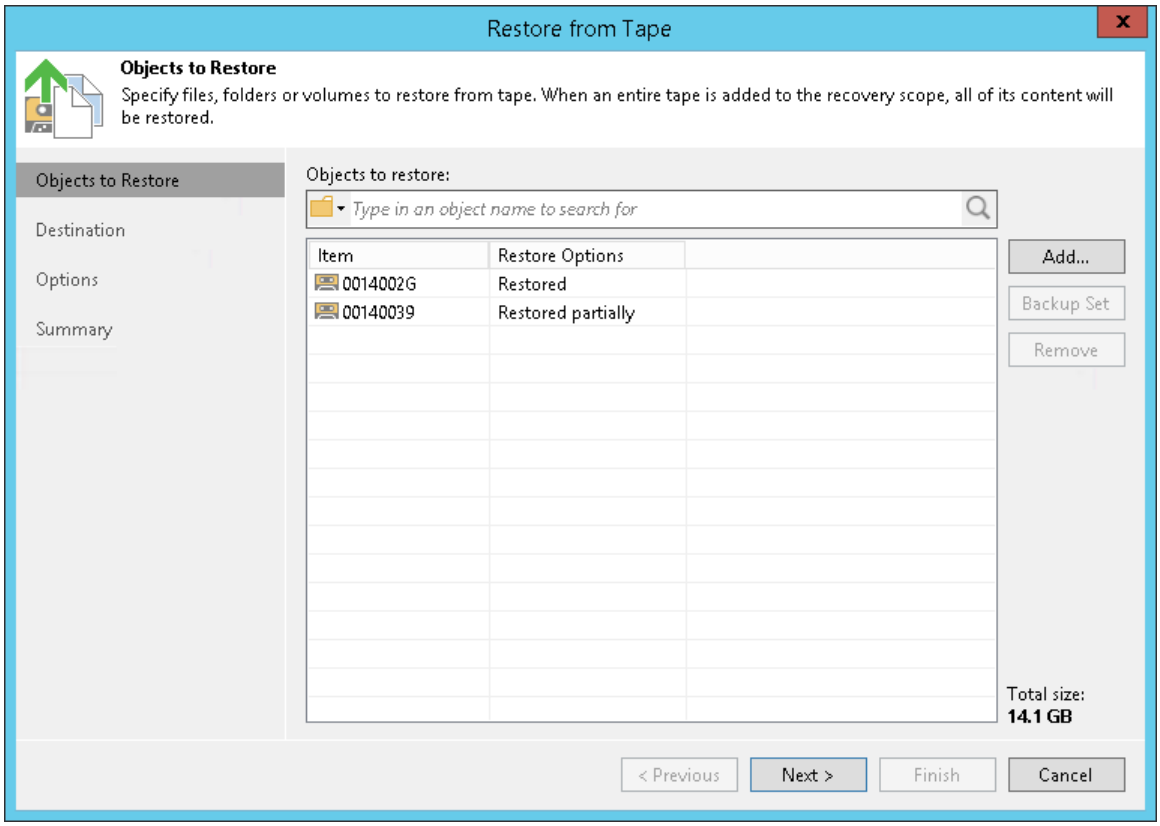

Step 3. Specify Restore Destination

At the **Destination** step of the wizard, specify destination where the archived files will be restored. You can use one of the following options:

- **Original location**. Use this option to restore data to the location where the original file resides (or resided): Microsoft Windows server, Linux server, or NMDP server. This type of restore is only possible if the original machine is connected to Veeam Backup & Replication and powered on.

  > **NOTE**
  >
  > If you are restoring content of the entire tapes, the **Original location** option is not available.

- **This server**. Use this option if you want to restore data to another server. The following options are available:

  - Microsoft Windows or Linux files: you can restore files to the Veeam backup server, shared folder or to any other machine added to Veeam Backup & Replication. From the server list, choose a machine to which files should be restored and specify path to the target folder.

    If you choose to restore files to a shared folder, make sure that the account under which Veeam Backup Service runs has write permissions to the target folder. If the account does not have sufficient permissions, Veeam Backup & Replication will prompt you to enter credentials for the account that can be used for writing to the target folder.

  - NDMP volumes: you can restore data to another NDMP server added to Veeam Backup & Replication. From the server list, select an NDMP server added to Veeam Backup & Replication. In the **Path to folder** field, specify the volume to which you want to restore the archived volume.

  - NAS file shares: you can restore data to an SMB (CIFS) or NFS file share. From the server list, select a NAS file share added to Veeam Backup & Replication. In the **Path to folder** field, specify the folder to which you want to restore the archived files and folders.

When restoring files to a location other than original, Veeam Backup & Replication preserves the folder hierarchy. To restore files to the specified target folder without keeping the folder structure, clear the **Preserve folder hierarchy** check box. This option is unavailable for restore of NDMP volumes.

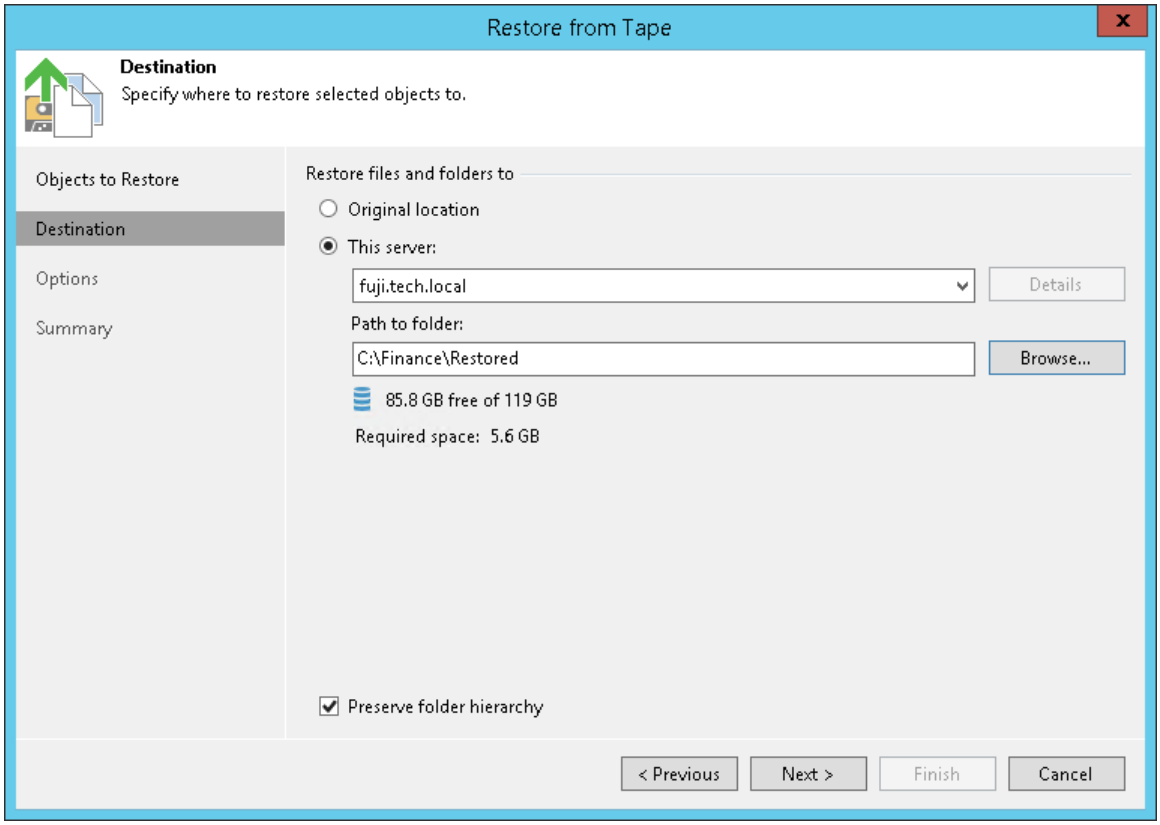

Step 4. Specify Restore Options

At the **Options** step of the wizard, specify overwrite options in case the file already exists in the target folder:

- **Leave the existing file**. Select this option if you do not want to overwrite the existing file with the restored one.
- **Overwrite the existing file if older than the backed up file**. Select this option if you want to overwrite the existing file only if it is older than the restored file.
- **Always overwrite the existing file**. Select this option if you want to overwrite the existing file with the restored file in all cases.

> **NOTE**
>
> Overwrite options are not available when restoring content of entire tapes.

Select the **Restore file and folder security** check box if you want the restored files to keep their original ownership and security permissions. In the opposite case, 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, while access permissions will be inherited from the target folder to which the objects are restored.

> **NOTE**
>
> This step is inapplicable for NMDP restore. If you restore NMDP volumes, these options will be skipped. The NMDP volumes will be restored according to the hardware settings.

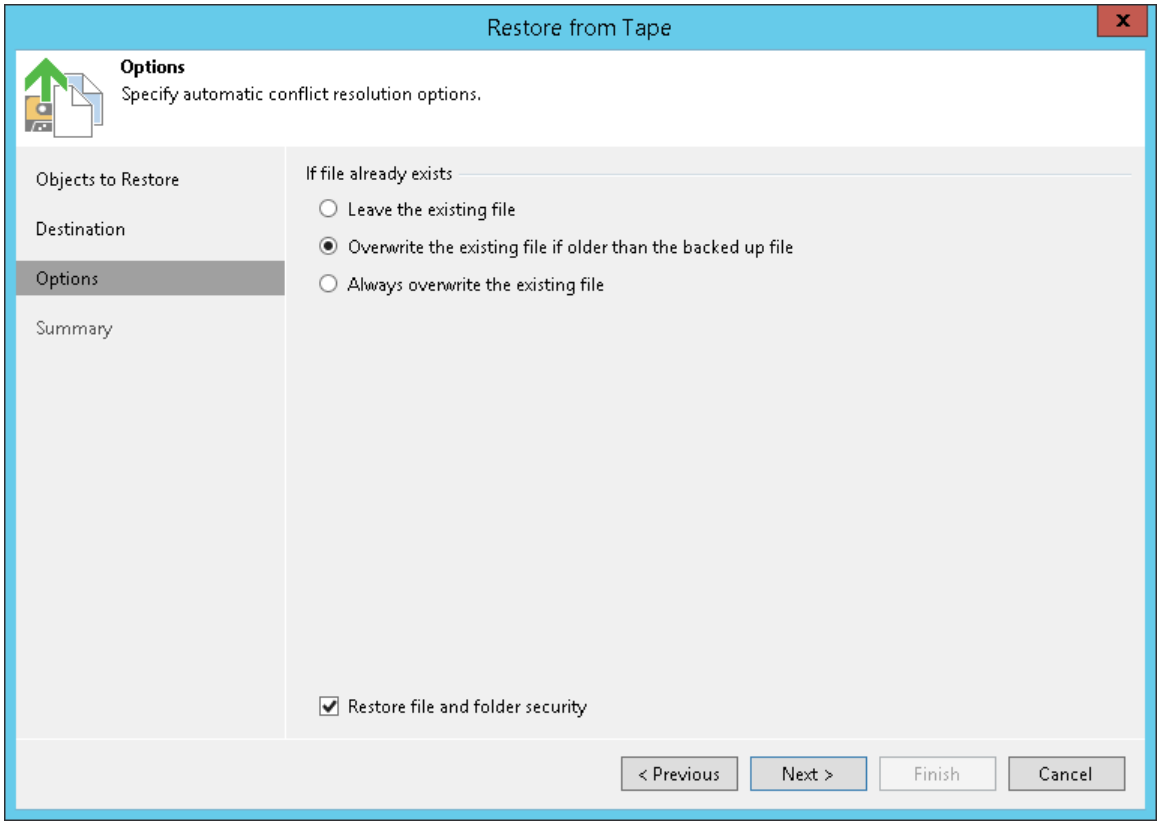

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review the restore settings and click **Finish** to restore selected files, folders, or tapes.

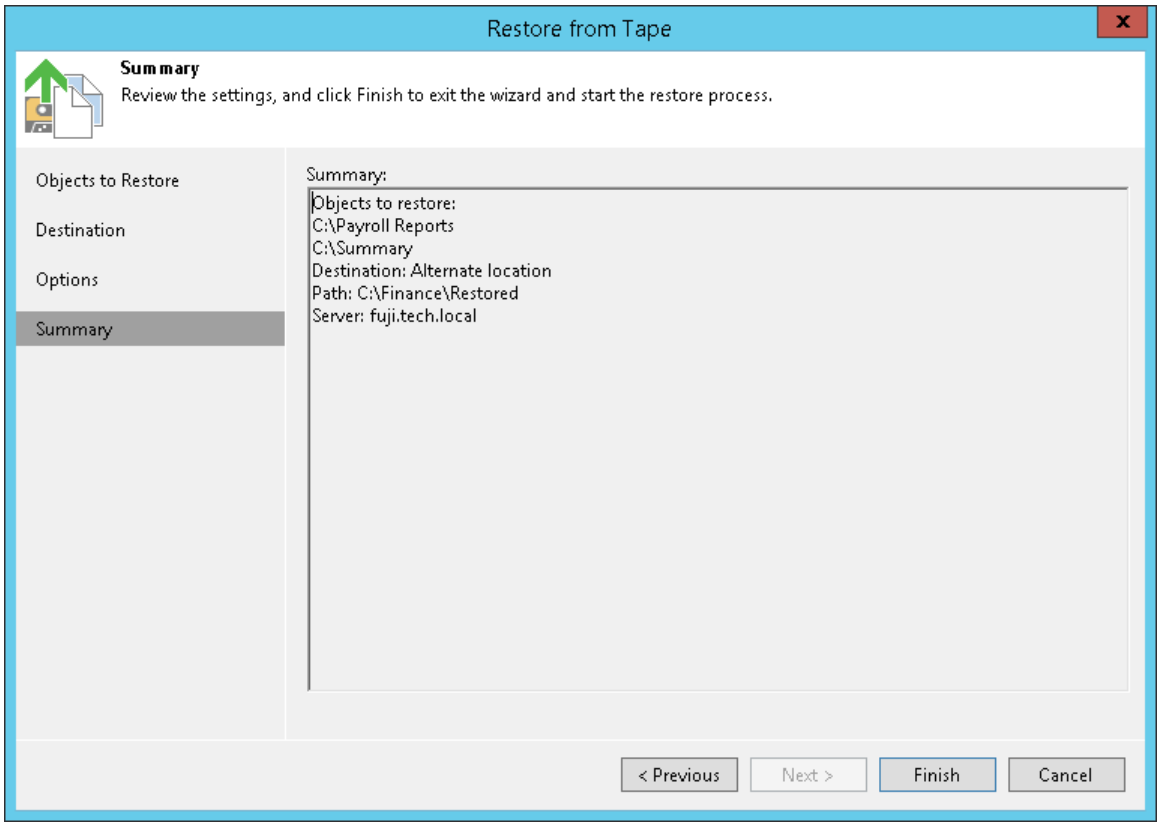

Tape States Indicators

The user interface indicates on states of the tape infrastructure components.

Tape Libraries

Icon	Component state
	Connected tape library
	Tape library connected over SCSI
	Tape library in the maintenance mode

Drives

Icon	Component state
	Enabled drive
	Disabled drive
	Drive in the maintenance mode

Media Pools

Icon	Component state
	Free media pool
	Imported media pool
	Unrecognized media pool
	Retired media pool
	Media pool

Icon	Component state
	GFS media pool
	WORM media pool
	WORM GFS media pool

Online Tapes

Icon	Component state
	Tape
	Retired tape
	Encrypted tape
	Software protected tape
	Hardware protected tape
	Software protected and encrypted tape
	Hardware protected and encrypted tape

Online WORM Tapes

Icon	Component state
	WORM tape
	Retired WORM tape
	Encrypted WORM tape
	Software protected WORM tape
	Hardware protected WORM tape

Icon	Component state
	Software protected and encrypted WORM tape
	Hardware protected and encrypted WORM tape

Offline Tapes

Offline indication does not differ for standard and WORM tapes.

Icon	Component state
	Tape
	Retired tape
	Encrypted tape
	Hardware protected tape
	Software protected tape
	Software protected and encrypted tape
	Hardware protected and encrypted tape

Cleaning Tape

Icon	Component state
	Online cleaning tape
	Offline cleaning tape
	Retired online cleaning tape
	Retired offline cleaning tape

Tape Alerts

Tape drive and tape changer alerts are informational, warning or critical notifications designed to provide real-time monitoring of the hardware operation status. They ensure quick detection and troubleshooting of various issues with devices.

In This Section

- Tape Drive Alerts
- Tape Changer Alerts

Tape Drive Alerts

There are 3 types of error flags:

Severity	Urgent User Intervention	Risk of Data Loss	Explanations
Critical	X	X	
Warning		X	X
Information			X

You may get one of the following tape drive alerts when working with Veeam Backup & Replication:

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
1	0x01	Read warning	Warning	The tape drive is having problems reading data. No data has been lost, but there has been a reduction in the performance of the tape.	The drive is having severe trouble reading.
2	0x02	Write warning	Warning	The tape drive is having problems writing data. No data has been lost, but there has been a reduction in the capacity of the tape.	The drive is having severe trouble writing.
3	0x03	Hard error	Warning	The operation has stopped because an error has occurred while reading or writing data that the drive cannot correct.	The drive had a hard read or write error.
4	0x04	Media error	Critical	Your data is at risk: 1. Copy any data you require from this tape. 2. Do not use this tape again. 3. Restart the operation with a different tape.	Media can no longer be written/read, or performance is severely degraded.
5	0x05	Read failure	Critical	The tape is damaged or the drive is faulty. Call the tape drive supplier helpline.	The drive can no longer read data from the tape.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
6	0x06	Write failure	Critical	The tape is from a faulty batch or the tape drive is faulty: 1. Use a good tape to test the drive. 2. If problem persists, call the tape drive supplier helpline.	The drive can no longer write data to the tape.
7	0x07	Media life	Warning	The tape cartridge has reached the end of its calculated useful life: 1. Copy data you need to another tape. 2. Discard the old tape.	The media has exceeded its specified life.
8	0x08	Not data grade tape	Warning	The tape cartridge is not data-grade. Any data you back up to the tape are at risk. Replace the cartridge with a data-grade tape.	The drive has not been able to read the MRS stripes.
9	0x09	Attempted backup to write-protected tape	Critical	You are trying to write to a write-protected cartridge. Remove the write-protection or use another tape.	Write command is attempted to a write protected tape.
10	0x0a	Tape drive is in use	Information	You cannot eject the cartridge because the tape drive is in use. Wait until the operation is complete before ejecting the cartridge.	Manual or s/w unload attempted when prevent media removal on.
11	0x0b	Attempted backup/restore to cleaning tape	Information	The tape in the drive is a cleaning cartridge.	Cleaning tape is loaded into drive.
12	0x0c	Unsupported cartridge type	Information	You have tried to load a cartridge of a type which is not supported by this drive.	Attempted loaded of unsupported tape format, for example, DDS2 in DDS1 drive.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
13	0x0d	Detected snapped tape	Critical	The operation has failed because the tape in the drive has experienced a mechanical failure: 1. Discard the old tape. 2. Restart the operation with a different tape.	Tape snapped/cut in the drive where media can be ejected.
14	0x0e	Detected snapped tape	Critical	The operation has failed because the tape in the drive has experienced a mechanical failure: 1. Do not attempt to extract the tape cartridge. 2. Call the tape drive supplier helpline.	Tape snapped/cut in the drive where media cannot be ejected.
15	0x0f	Cartridge memory chip failure	Warning	The memory in the tape cartridge has failed, which reduces performance. Do not use the cartridge for further write operations.	Memory chip failed in cartridge.
16	0x10	Forced eject	Critical	The operation has failed because the tape cartridge was manually de-mounted while the tape drive was actively writing or reading.	Manual or forced eject while drive actively writing or reading.
17	0x11	Read-only format	Warning	You have loaded a cartridge of a type that is read-only in this drive. The cartridge will appear as write-protected.	Media loaded that is read-only format.
18	0x12	Tape directory corrupted in the cartridge memory	Warning	The tape directory on the tape cartridge has been corrupted. File search performance will be degraded. The tape directory can be rebuilt by reading all the data on the cartridge.	Tape drive powered down with tape loaded, or permanent error prevented the tape directory being updated.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
19	0x13	Nearing media life	Information	The tape cartridge is nearing the end of its calculated life. It is recommended that you: 1. Use another tape cartridge for your next backup. 2. Store this tape in a safe place in case you need to restore data from it.	Media may have exceeded its specified number of passes.
20	0x14	Clean now	Critical	The tape drive needs cleaning: 1. If the operation has stopped, eject the tape and clean the drive. 2. If the operation has not stopped, wait for it to finish and then clean the drive. Check the tape drive user's manual for device specific cleaning instructions.	The drive thinks it has a head clog, or needs cleaning.
21	0x15	Clean periodic	Warning	The tape drive is due for routine cleaning: 1. Wait for the current operation to finish. 2. Then use a cleaning cartridge. Check the tape drive user's manual for device specific cleaning instructions.	The drive is ready for a periodic clean.
22	0x16	Expired cleaning media	Critical	The last cleaning cartridge used in the tape drive has worn out: 1. Discard the worn out cleaning cartridge. 2. Wait for the current operation to finish. 3. Then use a new cleaning cartridge.	The cleaning tape has expired.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
23	0x17	Invalid cleaning cartridge	Critical	The last cleaning cartridge used in the tape drive was an invalid type: 1. Do not use this cleaning cartridge in this drive. 2. Wait for the current operation to finish. 3. Then use a new cleaning cartridge.	Invalid cleaning tape type used.
24	0x18	Request retention	Warning	The tape drive has requested a retention operation.	The drive is having severe trouble reading or writing, which will be resolved by a retention cycle.
25	0x19	Dual-port interface error	Warning	A redundant interface port on the tape drive has failed.	Failure of one interface port in a dual-port configuration, for example, Fibrechannel.
26	0x1a	Cooling fan failure	Warning	A tape drive cooling fan has failed.	Fan failure inside tape drive mechanism or tape drive enclosure.
27	0x1b	Power supply	Warning	A redundant power supply has failed inside the tape drive enclosure. Check the enclosure user's manual for instructions on replacing the failed power supply.	Redundant PSU failure inside the tape drive enclosure or rack subsystem.
28	0x1c	Power consumption is outside the specified range	Warning	The tape drive power consumption is outside the specified range.	Power consumption of the tape drive is outside specified range.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
29	0x1d	Preventative maintenance is required	Warning	Preventive maintenance of the tape drive is required. Check the tape drive user's manual for device specific preventive maintenance tasks or call the tape drive supplier helpline.	The drive requires preventative maintenance (not cleaning).
30	0x1e	Drive hardware fault	Critical	The tape drive has a hardware fault: 1. Eject the tape or magazine. 2. Reset the drive. 3. Restart the operation.	The drive has a hardware fault that requires reset to recover.
31	0x1f	Drive hardware fault	Critical	The tape drive has a hardware fault: 1. Turn the tape drive off and then on again. 2. Restart the operation. 3. If the problem persists, call the tape drive supplier helpline.	The drive has a hardware fault which is not read/write related or requires a power cycle to recover.
32	0x20	Host interface fault	Warning	The tape drive has a problem with the application client interface: 1. Check the cables and cable connections. 2. Restart the operation.	The drive has identified an interfacing fault.
33	0x21	Eject media request	Critical	The operation has failed: 1. Eject the tape or magazine. 2. Insert the tape or magazine again. 3. Restart the operation.	Error recovery action.
34	0x22	Firmware download fault	Warning	The firmware download has failed because you have tried to use the incorrect firmware for this tape drive. Obtain the correct firmware and try again.	Firmware download failed.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
35	0x23	Humidity specification exceeded	Warning	Environmental conditions inside the tape drive are outside the specified humidity range.	Drive humidity limits exceeded.
36	0x24	Temperature specification exceeded	Warning	Environmental conditions inside the tape drive are outside the specified temperature range.	Drive temperature limits exceeded.
37	0x25	Voltage specification exceeded	Warning	The voltage supply to the tape drive is outside the specified range.	Drive voltage limits exceeded.
38	0x26	Tape device predicted to fail	Critical	A hardware failure of the tape drive is predicted. Call the tape drive supplier helpline.	Predictive failure of drive hardware.
39	0x27	Drive hardware fault	Warning	The tape drive may have a hardware fault. Run extended diagnostics to verify and diagnose the problem. Check the tape drive user's manual for device specific instructions on running extended diagnostic tests.	The drive may have had a failure which may be identified by stored diagnostic information or by running extended diagnostics (for example, Send Diagnostic).
40	0x28	Autoloader communication fault	Critical	The changer mechanism is having difficulty communicating with the tape drive: 1. Turn the autoloader off then on. 2. Restart the operation. 3. If problem persists, call the tape drive supplier helpline.	Loader mechanism is having trouble communicating with the tape drive.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
41	0x29	Detected stray tape in autoloader	Critical	A tape has been left in the autoloader by a previous hardware fault: 1. Insert an empty magazine to clear the fault. 2. If the fault does not clear, turn the autoloader off and then on again. 3. If the problem persists, call the tape drive supplier helpline.	Stray tape left in loader after previous error recovery.
42	0x2a	Autoloader mechanism fault	Warning	There is a problem with the autoloader mechanism.	Loader mechanism has a hardware fault.
43	0x2b	Autoloader door open	Critical	The operation has failed because the autoloader door is open: 1. Clear any obstructions from the autoloader door. 2. Eject the magazine and then insert it again. 3. If the fault does not clear, turn the autoloader off and then on again. 4. If the problem persists, call the tape drive supplier helpline.	Tape changer door open.
44	0x2c	Autoloader hardware fault	Critical	The autoloader has a hardware fault: 1. Turn the autoloader off and then on again. 2. Restart the operation. 3. If the problem persists, call the tape drive supplier helpline. Check the autoloader user's manual for device specific instructions on turning the device power on and off.	The loader mechanism has a hardware fault that is not mechanically related.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
45	0x2d	Autoloader cannot operate without magazine	Critical	The autoloader cannot operate without the magazine: 1. Insert the magazine into the autoloader. 2. Restart the operation.	Loader magazine not present.
46	0x2e	Autoloader predicted to fail	Warning	A hardware failure of the changer mechanism is predicted. Call the tape drive supplier helpline.	Predictive failure of loader mechanism hardware.
47 - 49	0x2f - 0x31			These codes are reserved for future use.	
50	0x32	Media statistics is lost	Warning	Media statistics have been lost at some time in the past.	Drive or library powered down with tape loaded.
51	0x33	Tape directory is corrupted	Warning	The tape directory on the tape cartridge just unloaded has been corrupted. File search performance will be degraded. The tape directory can be rebuilt by reading all the data.	Error prevented the tape directory being updated on unload.
52	0x34	Tape system area write failure	Critical	The tape just unloaded could not write its system area successfully: 1. Copy data to another tape cartridge. 2. Discard the old cartridge.	Write errors while writing the system log on unload.
53	0x35	Tape system area read failure	Critical	The tape system are could not be read successfully at load time: 1. Copy data to another tape cartridge. 2. Discard the old cartridge.	Read errors while reading the system area on load.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
54	0x36	Start of data not found	Critical	The start or data could not be found on the tape: 1. Check you are using the correct format tape. 2. Discard the tape or return the tape to your supplier.	Tape damaged, bulk erased, or incorrect format.
55	0x37	Load failure	Critical	The operation has failed because the media cannot be loaded and threaded: 1. Remove the cartridge, inspect it as specified in the product manual, and retry the operation. 2. If the problem persists, call the tape drive supplier help line.	Media could not be loaded/threaded.
56	0x38	Medium unload failure	Critical	The operation has failed because the medium cannot be unloaded: 1. Do not attempt to extract the tape cartridge. 2. Call the tape driver supplier help line.	Unrecoverable unload failed.
57	0x39	Interface failure	Critical	The tape drive has a problem with the automation interface: 1. Check the power to the automation system. 2. Check the cables and cable connections. 3. Call the supplier help line if problem persists.	Automation interface failure.
58	0x3a	Firmware failure	Warning	The tape drive has reset itself due to a detected firmware fault. If problem persists, call the supplier help line.	Firmware failure.

Tape Changer Alerts

There are 3 types of error flags:

Severity	Urgent User Intervention	Risk of Data Loss	Explanations
Critical	X	X	
Warning		X	X
Information			X

You may get one of the following tape changer alerts when working with Veeam Backup & Replication:

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
0	0x00		Information	Unknown changer alert.	
1	0x01	Library Hardware A	Critical	The library mechanism is having difficulty communicating with the drive: 1. Turn the library off then on. 2. Restart the operation. 3. If the problem persists, call the library supplier help line.	Changer mechanism is having trouble communicating with the internal drive.
2	0x02	Library Hardware B	Warning	There is a problem with the library mechanism. If problem persists call the library supplier help line.	Changer mechanism has a hardware fault.
3	0x03	Library Hardware C	Critical	The library has a hardware fault: 1. Reset the library. 2. Restart the operation. 3. Check the library users manual for device specific instructions on resetting the device.	The changer mechanism has a hardware fault that requires reset to recover.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
4	0x04	Library Hardware D	Critical	The library has a hardware fault: 1. Turn the library off then on again. 2. Restart the operation. 3. If the problem persists, call the library supplier help line. 4. Check the library users manual for device specific instructions on turning the device power on and off.	The changer mechanism has a hardware fault that is not mechanically related, or requires a power cycle to recover.
5	0x05	Library Diagnostics Required	Warning	The library mechanism may have a hardware fault: 1. Run extended diagnostics to verify and diagnose the problem. 2. Check the library users manual for device specific instructions on running extended diagnostic tests.	The changer mechanism may a hardware fault which would be identified by extended diagnostics (for example, SCSI Send Diagnostic).
6	0x06	Library Interface	Critical	The library has a problem with the host interface: 1. Check the cables and connections. 2. Restart the operation.	The library has identified an interfacing fault.
7	0x07	Predictive Failure	Warning	A hardware failure of the library is predicted. Call the library supplier help line.	Predictive failure of library hardware.
8	0x08	Library Maintenance	Warning	Preventive maintenance of the library is required. Check the library users manual for device specific preventative maintenance tasks, or call your library supplier help line.	Library preventative maintenance required.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
9	0x09	Library Humidity Limits	Critical	General environmental conditions inside the library are outside the specified humidity range.	Library humidity limits exceeded.
10	0x0a	Library Temperature Limits	Critical	General environmental conditions inside the library are outside the specified temperature range.	Library temperature limits exceeded.
11	0x0b	Library Voltage Limits	Critical	The voltage supply to the library is outside the specified range. There is a potential problem with the power supply or failure of a redundant power supply.	Library voltage limits exceeded
12	0x0c	Library Stray Tape	Critical	A cartridge has been left inside the library by a previous hardware fault: 1. Insert an empty magazine to clear the fault. 2. If the fault does not clear, turn the library off and then on again. 3. If the problem persists, call the library supplier help line.	Stray cartridge left in library drive after previous error recovery.
13	0x0d	Library Pick Retry	Warning	There is a potential problem with the drive ejecting cartridges or with the library mechanism picking a cartridge from a slot. No action needs to be taken at this time. If the problem persists, call the library supplier help line.	Operation to pick a cartridge from a slot had to perform an excessive number of retries before succeeding.
14	0x0e	Library Place Retry	Warning	There is a potential problem with the library mechanism placing a cartridge into a slot. No action needs to be taken at this time. If the problem persists, call the library supplier help line.	Operation to place a cartridge into a slot had to perform an excessive number of retries before succeeding.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
15	0x0f	Library Load Retry	Warning	There is a potential problem with the drive or the library mechanism loading cartridges, or an incompatible cartridge.	Operation to load a cartridge into a drive had to perform an excessive number of retries before succeeding.
16	0x10	Library Door	Critical	The library has failed because the door is open: 1. Clear any obstructions from the library door. 2. Close the library door. 3. If the problem persists, call the library supplier help line.	Changer door open prevents library functioning.
17	0x11	Library Mailslot	Critical	There is a mechanical problem with the library media import/export mailslot.	Mechanical problem with import/export mailslot
18	0x12	Library Magazine	Critical	The library cannot operate without the magazine: 1. Insert the magazine into the library. 2. Restart the operation.	Library magazine not present.
19	0x13	Library Security	Warning	Library security has been compromised.	Library door opened then closed during operation.
20	0x14	Library Security Mode	Information	The library security mode has been changed. The library has either been put into secure mode, or the library has exited the secure mode. This is for information purposes only. No action is required.	Library security mode changed.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
21	0x15	Library Offline	Information	The library has been manually turned offline and is unavailable for use.	Library manually turned offline.
22	0x16	Library Drive Offline	Information	A drive inside the library has been taken offline. This is for information purposes only. No action is required.	Library turned internal drive offline.
23	0x17	Library Scan Retry	Warning	There is a potential problem with the bar code label or the scanner hardware in the library mechanism. No action needs to be taken at this time. If the problem persists, call the library supplier help line.	Operation to scan the barcode on a cartridge had to perform an excessive number of retries before succeeding.
24	0x18	Library Inventory	Critical	The library has detected an inconsistency in its inventory. 1. Redo the library inventory to correct inconsistency. 2. Restart the operation. 3. Check the applications users manual or the hardware users manual for specific instructions on redoing the library inventory.	Inconsistent media inventory.
25	0x19	Library Illegal Operation	Warning	A library operation has been attempted that is invalid at this time.	Illegal operation detected.
26	0x1a	Dual-Port Interface Error	Warning	A redundant interface port on the library has failed.	Failure of one interface port in a dual-port configuration, for example, Fibrechannel.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
27	0x1b	Cooling Fan Failure	Warning	A library cooling fan has failed.	One or more fans inside the library have failed. Internal flag state only cleared when all fans are working again.
28	0x1c	Power Supply	Warning	A redundant power supply has failed inside the library. Check the library users manual for instructions on replacing the failed power supply.	Redundant PSU failure inside the library subsystem.
29	0x1d	Power Consumption	Warning	The library power consumption is outside the specified range.	Power consumption of one or more devices inside the library is outside specified range.
30	0x1e	Pass-through mechanism failure.	Critical	A failure has occurred in the cartridge pass-through mechanism between two library modules.	Error occurred in pass-through mechanism during self test or while attempting to transfer a cartridge between library modules.
31	0x1f	Cartridge in pass-through mechanism	Critical	A cartridge has been left in the pass-through mechanism from a previous hardware fault. Check the library users guide for instructions on clearing this fault.	Cartridge left in the pass-through mechanism between two library modules.

Alert Code (Dec)	Alert Code (Hex)	Flag	Severity	Description	Cause
32	0x20	Unreadable bar code labels	Information	The library was unable to read the bar code on a cartridge.	Unable to read a barcode label on a cartridge during library inventory/scan.

Veeam Backup for Nutanix AHV

Version 3.0

User Guide

October, 2021

© 2021 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 7
ABOUT THIS DOCUMENT 8
OVERVIEW 9
BACKUP INFRASTRUCTURE COMPONENTS 11
PLANNING AND PREPARATION 12
System Requirements 13
Required Permissions 14
Used Ports 15
Licensing 18
DEPLOYMENT 19
Step 1. Configure Veeam Backup & Replication Server 20
Install Veeam Backup & Replication Server 21
Configure Backup Repository 23
Step 2. Install Nutanix AHV Plug-In 24
Step 3. Perform Initial Configuration 26
Troubleshooting AHV Backup Proxy Deployment 27
Configuring Hostname Resolution 28
Disabling HTTP Proxy 30
Upgrading to Version 3.0 31
Uninstalling Nutanix AHV Plug-In 33
CONFIGURATION OF BACKUP INFRASTRUCTURE 34
Adding Nutanix AHV Cluster 35
Step 1. Launch Add AHV Cluster Wizard 36
Step 2. Specify Cluster Name or Address 38
Step 3. Specify Credentials 39
Step 4. Specify Helper Appliance Settings 40
Step 5. Finish Working with Wizard 42
Managing Nutanix AHV Clusters 44
Rescanning Nutanix AHV Cluster 45
Removing AHV Clusters from Backup Infrastructure 47
Editing AHV Cluster Settings 48
Adding VMs to Protection Domain (Auto-Protection) 49
Adding AHV Backup Proxy 51
Step 1. Launch New Nutanix Proxy Wizard 52
Step 2. Select Deployment Mode 53
Step 3. Specify VM Configuration 54

Step 4. Specify Network Settings 56
Step 5. Specify Credentials 58
Step 6. Grant Permissions 59
Step 7. Finish Working with Wizard 60
Managing AHV Backup Proxies 62
Editing Nutanix AHV Backup Proxies 63
Rescanning Nutanix AHV Backup Proxies 64
Removing Nutanix AHV Backup Proxies 65
Enabling SSH on AHV Backup Proxy VM 67
Managing TLS Certificates 68
Generating Self-Signed Certificates 69
Importing Certificates from Certificate Store 72
Importing Certificates from PFX Files 74
Using Certificate Signed by Internal CA 76
WEB CONSOLE OF AHV BACKUP PROXY 80
Accessing Web Console 81
Getting to Know UI of Web Console 82
Monitoring Events 83
GENERAL SETTINGS 84
Configuring Network Settings 85
Configuring Administrator Account Settings 86
Configuring Time Zone Settings 87
Updating AHV Backup Proxy 89
Checking for Updates 90
Installing Updates 91
Configuring E-mail Settings 93
Configuring Notifications 95
Backing Up Configuration Settings 96
Restoring Configuration Settings 98
Rebooting or Shutting Down Proxy VM 102
Protecting AHV Backup Proxy VM 103
DATA PROTECTION 105
Backup Jobs 106
How Backup Jobs Work 107
Backup Chain and Retention Policy 108
Creating Backup Jobs 110
Creating Active Full Backup 127
Snapshot-Only Jobs 128
Nutanix AHV Snapshot Types 129

Snapshot Chain and Retention Policy 130
Before You Begin 131
Creating PD Snapshot Jobs 132
Creating Snapshot Jobs 138
VeeamZIP 146
Creating VeeamZIP Backups 147
Backup Copy Jobs 149
Backup to Tape Jobs 151
Managing Jobs 152
Starting Jobs Manually 153
Stopping Jobs 155
Editing Job Settings 157
Deleting Jobs 158
Disabling Jobs 160
Managing Backups and Snapshots 162
Viewing Backup Properties 163
Importing Backups 164
Exporting Backups 166
Deleting Backups 167
Deleting Snapshots 168
DATA RECOVERY 169
Entire VM Restore 170
Restoring VMs Using Veeam Backup & Replication Console 171
Restoring VMs Using AHV Backup Proxy Console 183
VM Disks Restore 198
Before You Begin 199
Launch Virtual Disk Restore Wizard 200
Step 1. Select Virtual Machine 201
Step 2. Select Restore Point 202
Step 3. Configure Mapping Settings 203
Step 4. Specify Reason for Restore 205
Step 5. Finish Working with Wizard 206
Instant Recovery 207
Performing Instant Recovery of Workloads to Nutanix AHV 208
Performing Instant Recovery of Workloads to VMware vSphere 222
Performing Instant Recovery of Workloads to Hyper-V 223
VM Guest OS Files Restore 224
Application Items Restore 227
VM Disk Export 229

Restore to Amazon EC2 230
Restore to Microsoft Azure 231
LOGS AND SUPPORT 232
Exporting Logs in AHV Backup Proxy 233
Exporting Logs in Veeam Backup & Replication 234
APPENDICES 235
Appendix A. Manual Deployment 236
Step 1. Upload AHV Backup Proxy Virtual Disk Image 237
Step 2. Create VM for AHV Backup Proxy 239
Step 3. Log in to AHV Backup Proxy Web Console 244
Step 4. Install AHV Backup Proxy 247
Step 5. Add AHV Cluster to AHV Backup Proxy Infrastructure 252
Step 6. Add Veeam Backup & Replication Server 254
Step 7. Restoring Configuration Settings During Deploy 256
Appendix B. Custom Scripts 262
Step 1. Preparing Custom Script 263
Step 2. Enabling Custom Scripts Functionality 264
Step 3. Uploading Custom Script 266
Step 4. Adding Custom Script to Job 267
Step 5. Checking Custom Script Logs 269

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 guide provides information about main features of Veeam Backup for Nutanix AHV. The document describes how to deploy, configure and use Veeam Backup for Nutanix AHV.

Intended Audience

This guide is designed for IT specialists who plan to use Veeam Backup for Nutanix AHV to protect VMs residing on AHV clusters.

Overview

Veeam Backup for Nutanix AHV is an additional component of Veeam Backup & Replication responsible for protection and disaster recovery tasks in the Nutanix AHV environment.

Veeam Backup for Nutanix AHV uses agentless approach to back up Nutanix AHV VMs and volume groups. The solution creates image-based backups of VMs at a specific point in time. Veeam Backup for Nutanix AHV retrieves the VM data block by block, compresses, deduplicates it, and forwards it to a backup repository.

> **NOTE**
>
> Veeam Backup for Nutanix AHV is an updated version of Veeam Availability for Nutanix AHV 1.0.

Functionality

To manage protection and disaster recovery operations for Nutanix AHV VMs, you can use two consoles: Veeam Backup & Replication console and AHV Backup Proxy web console. Note that each console supports different functionality.

- **With the AHV Backup Proxy** web console, you can:
  - Create backups of Nutanix AHV VMs and store them in Veeam backup repositories.
  - Create VM snapshots and protection domain snapshots.
  - Restore VMs from backups and snapshots to the original Nutanix AHV cluster.
  - Restore VM disks and attach them to VMs running on Nutanix AHV clusters.
- **With the Veeam Backup & Replication** console, you can:
  - Create VeeamZIP backups of Nutanix AHV VMs.
  - Create backup copies of backups created by AHV Backup Proxy.
  - Perform Instant Recovery of AHV VMs to Nutanix AHV, VMware vSphere and Microsoft Hyper-V.
  - Restore AHV VMs from backups or from snapshots to AHV, VMware vSphere, Microsoft Hyper-V, Microsoft Azure Blob, Amazon EC2 and Google CE.
  - Restore VMware ESXi and Microsoft Hyper-V VMs from backups to AHV clusters.
  - Restore from backups of physical machines created by Veeam Agents to AHV clusters.
  - Restore VM guest OS files.
  - Restore application items (Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Oracle Database and Microsoft SQL Server).
  - Export VM disks to VMDK, VHD and VHDX formats.

> **NOTE**
>
> Use Veeam Backup for Nutanix AHV to back up VMs hosted on Nutanix AHV cluster nodes. If you want to back up and restore VMs hosted on Nutanix systems running VMware vSphere and Microsoft Hyper-V, you can use native capabilities of Veeam Backup & Replication.

Related Topics

- Backup Infrastructure Components
- How It Works

Backup Infrastructure Components

The backup infrastructure of Veeam Backup for Nutanix AHV includes the following components:

1. **Veeam Backup & Replication**: a starting point in the deployment process. In the Veeam Backup & Replication console, you can add components required to protect Nutanix AHV VMs.
2. **Nutanix AHV Plug-in**: an installer component that deploys AHV Backup Proxy services on the Veeam Backup & Replication server. After you install Nutanix AHV Plug-in, you can add AHV clusters and AHV Backup Proxy instances to the Veeam Backup & Replication infrastructure.
3. **AHV Backup Proxy**: the main component that manages backup and restore processes. AHV Backup Proxy allows you to create VM backups, snapshots and restore VMs from backups to an AHV cluster.

   AHV Backup Proxy is deployed as a VM on the AHV cluster. For one Nutanix AHV cluster, one AHV Backup Proxy is required. To back up and restore VMs on multiple clusters, you will need to deploy an AHV Backup Proxy appliance on each cluster.

   AHV Backup Proxy includes the following additional components:

   - **Data Mover**: a service that transports VM backup data from AHV cluster to backup repositories.
   - **AHV Backup Proxy web console**: a web console that allows you to manage backup and restore operations, configure notifications and configure settings of AHV Backup Proxy.
4. **Backup repository**: a storage location where AHV Backup Proxy stores backup files. The backup repository also includes a Data Mover service that is responsible for receiving VM backup data from the AHV cluster.

   You can add multiple repositories and select the required one within settings of a specific backup job. For the list of supported repository types, see Configure Backup Repository.

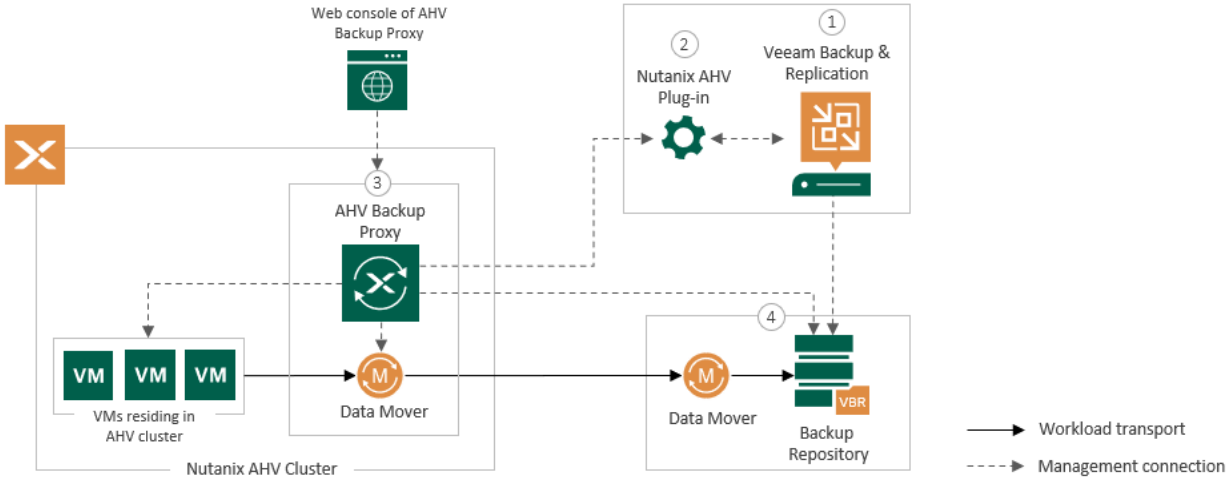

Related Topics

- How It Works
- Nutanix AHV Snapshot Types

Planning and Preparation

Before you start deploying Veeam Backup for Nutanix AHV, you must make sure that the virtual environment and machines that you plan to use as backup infrastructure components meet product hardware recommendations and system requirements.

In This Section:

- System Requirements
- Required Permissions
- Used Ports
- Licensing

System Requirements

Before you start deploying Veeam Backup for Nutanix AHV, make sure the Nutanix AHV cluster and backup infrastructure components meet system requirements listed below.

Platform

AHV Backup Proxy supports the following versions of the Nutanix AOS:

- Nutanix AOS versions 5.5.X up to 6.0.X.

  Instant Recovery to Nutanix AHV is supported only for Nutanix AOS 6.0 and later.

- Community Edition: CE 2019.11.22 or later.

> **IMPORTANT**
>
> To integrate the AHV cluster with Veeam Backup for Nutanix AHV, you must define an IP address for iSCSI data services in settings of the AHV cluster. For more information, see the Modifying Cluster Details section of the Prism Web Console Guide.

Virtual Hardware for AHV Backup Proxy VM

During the deployment process, AHV Backup Proxy is deployed as an independent Linux-based VM on a Nutanix AHV cluster that you want to protect. The AHV Backup Proxy VM must meet the following requirements.

- *CPU*: x86-64 processors. 4 CPU cores plus 1 CPU core for each additional concurrent task.*
- *Memory*: 4 GB RAM plus 1 GB RAM for each additional concurrent task.*
- *Disk Space*: 1 GB for product installation.

*AHV Backup Proxy creates one task per every VM processed by the job. CPU and hardware resources for the AHV Backup Proxy VM are calculated automatically based on specified number of concurrent tasks. VMs with multiple CPU cores can handle more concurrent tasks. For example, for 4 concurrent tasks, we recommend you to specify minimum 4 CPU cores and 4GB RAM, for 8 concurrent tasks — 8 CPU cores and 8 GB RAM. When defining the number of concurrent tasks, keep in mind the network traffic throughput in the virtual infrastructure.

Veeam Backup & Replication and Nutanix AHV Plug-in

As described in Backup Infrastructure Components, Veeam Backup for Nutanix AHV consists of 3 main components: AHV Backup Proxy, Nutanix AHV Plug-in and Veeam Backup & Replication.

AHV Backup Proxy 3.0.775 supports integration with **Nutanix AHV Plug-in 11.0.3.819** installed on the **Veeam Backup & Replication 11.0.1.1261** server.

> **IMPORTANT**
>
> The hostname of the Veeam Backup & Replication server must be resolved on the AHV Backup Proxy VM. Otherwise, AHV Backup Proxy cannot connect to Veeam Backup & Replication. The issue is also valid if AHV Backup Proxy is running behind NAT.
>
> For more information, see Solving Hostname Resolution Problem.

Required Permissions

Make sure the user accounts that you plan to use have permissions described in the following sections.

Installing and Configuring Nutanix AHV Plug-in

The accounts used to install and configure Nutanix AHV Plug-in must have the following permissions.

Account	Required Permission
Setup Account	The account used to install Nutanix AHV Plug-in must have the local Administrator permissions on the target machine.
Veeam Backup & Replication User Account	The account used to run Veeam Backup & Replication services including the Veeam AHV Service must be a LocalSystem account or must have the local Administrator permissions on the server where Nutanix AHV Plug-in is installed.

Adding Nutanix AHV Cluster to Veeam Backup & Replication Infrastructure

When you add an AHV cluster to the Veeam Backup & Replication infrastructure, the specified account must have privileges of the Prism Admin role or higher. For more information, see the Controlling User Access section of the Prism Central Guide.

Used Ports

When you add new backup infrastructure components, Veeam Backup & Replication automatically creates firewall rules for the required ports on this servers. These rules allow communication between the components. In this topic, you can find the full list of required ports.

> **IMPORTANT**
>
> For Debian, SLES, RHEL/CentOS distributions, you must create firewall or security rules manually. For more information, see this Veeam KB article.

AHV Backup Proxy Server

The following table describes network ports that must be opened to ensure proper communication of the AHV Backup Proxy with backup infrastructure components.

From	To	Port	Protocol	Notes
Browser	AHV Backup Proxy server	443 or 8100	HTTPS	Port used by web console of AHV Backup Proxy.
AHV Backup Proxy	Nutanix REST API	9440	HTTPS	Port used to communicate with Nutanix REST API.
	Veeam Backup & Replication server	10006	TCP	Port used to connect to Veeam Backup & Replication.
	Nutanix AHV server (Cluster Virtual IP, Cluster ISCSI Data Services IP, Cluster CVM IPs)	3205, 3260	TCP/iSCSI	Ports used for connections to disks on Nutanix AHV VMs.
	Veeam backup repository (or gateway server)	2500-3300*	TCP	Default range of ports used as transmission channels for jobs and restore sessions. For every TCP connection that a job uses, one port from this range is assigned.
	Ubuntu Security and OS Update repository (*security.ubuntu.com*, *archive.ubuntu.com*)	80	HTTP	Required to get OS security updates.
	OS Update repository (*packages.microsoft.com*)	443	HTTPS	Required to get OS updates.

From	To	Port	Protocol	Notes
	Veeam Updater repository (*repository.veeam.com*, *cloudfront.net*)	443	TCP	Required to download AHV Backup Proxy packages and information on available product updates.
	Nginx repository (*nginx.org/packages/*, *nginx.org/keys/*)	443	TCP	Required to download nginx packages.

* This range of ports applies to newly added backup infrastructure components. If you upgrade to Veeam Backup & Replication 10.0 from earlier versions of the product, the range of ports from 2500 to 5000 applies to the already added components.

Veeam Backup & Replication Server

The following table describes network ports that must be opened to ensure proper communication of Veeam Backup & Replication with backup infrastructure components.

From	To	Port	Protocol	Notes
AHV Backup Proxy, Veeam Backup & Replication console and Veeam ONE server	Veeam Backup & Replication server	8543	HTTPS	Port used by AHV Platform Service (the main service of Nutanix AHV Plug-in).
Veeam Backup & Replication server	FLR helper appliance	22	TCP	Port used to connect to the file-level restore helper appliance.
	Veeam Backup & Replication server	6172	HTTPS	Port used by AHV Platform Service to connect to a component that enables communication with the Veeam Backup & Replication database.
	Nutanix AHV cluster	9440	HTTPS	Port used by AHV Platform Service to connect to AHV cluster.
	AHV Backup Proxy	443 or 8100	HTTPS	Port used by AHV Platform Service to connect to AHV Backup Proxy.

vPower NFS Service

IMPORTANT

- To see ports used by Veeam Backup & Replication to communicate with backups repositories, see the Used Ports section of the Veeam Backup & Replication User Guide.
- If you perform VM guest OS files restore, see required ports for Proxy Appliance (Multi-OS FLR) and Mount Server.

From	To	Port	Protocol	Notes
Nutanix AHV cluster	Microsoft Windows server running vPower NFS Service	111	TCP UDP	Standard port used by port mapper service.
		1058+ or 1063+	TCP UDP	Default mount port. The number of port depends on where the vPower NFS Service is located: • 1058+: If the vPower NFS Service is located on the backup server. • 1063+: If the vPower NFS Service is located on a separate Microsoft Windows machine. If port 1058/1063 is occupied, the succeeding port numbers will be used.
		2049+	TCP UDP	Standard NFS port. If port 2049 is occupied, the succeeding port numbers will be used.

Licensing

To use Veeam Backup for Nutanix AHV, you must have a valid Veeam License installed on the Veeam Backup & Replication server. If the license is not valid or there are no available license instances, backup jobs of Veeam Backup for Nutanix AHV fail. Without a valid license, you can only create snapshots and restore AHV VMs from snapshots.

This guide provides information only on specifics of Veeam licensing for Veeam Backup for Nutanix AHV. For terminology and general information about Veeam licensing, see the Veeam Licensing Policy article.

Licensed Objects

VMs protected by Veeam Backup for Nutanix AHV are licensed per protected VM. An AHV VM is considered protected if it has a restore point created in the last 31 days.

If you are using any instance-based (Veeam Universal Licensing) license on your Veeam Backup & Replication, you don't need to install any additional licenses.

When you back up one AHV VM, one license instance is consumed from the Veeam Backup & Replication license. AHV VMs processed by backup copy and tape jobs are not regarded as protected VMs, these types of jobs provide an additional protection level for VMs that are already protected with backup jobs. VM templates are licensed as regular VMs.

> **NOTE**
>
> [For Perpetual per-socket licenses] If you use a legacy perpetual per-socket license and you want to deploy AHV Backup Proxy, you must install an additional instance-based (Veeam Universal Licensing) license. The new license will be merged with the existing perpetual socket license. To obtain the additional instance-based license, contact your Veeam sales representative.

Obtaining and Renewing License

You can obtain an Evaluation or paid license for Veeam Backup for Nutanix AHV when you download the product from the Veeam website.

- **Obtaining paid license**: If you want to see the full list of license packages and obtain a paid license, see the Pricing and Packaging page and click **Contact Sales**.
- **Obtaining trial license**: If you want to obtain a free 30-day trial license, do the following:
  a. Sign in to Veeam.com.
  b. Open the Download page.
  c. In the **Get Trial Key** section, click the **Request Trial Key** link to download the Evaluation license.
- **Renewing license**: If you want to renew your maintenance plan, contact Veeam Renewals Team at renewals@veeam.com.

Managing License

Licensing of protected AHV VMs is performed on the Veeam Backup & Replication side. Thus, to learn how to install a License on Veeam Backup & Replication or how to monitor licensed objects, see the Licensing section of the Veeam Backup & Replication User Guide.

Deployment

To start working with Veeam Backup for Nutanix AHV, you must deploy a Veeam Backup & Replication server, install Veeam AHV Plug-in and configure the backup infrastructure for an AHV cluster that you want to protect.

To deploy Veeam Backup for Nutanix AHV and set up the backup infrastructure, perform the following steps:

1. Configure Veeam Backup & Replication server: install Veeam Backup & Replication and configure the backup repository that will be used as a target for VM backups.
2. Install Nutanix AHV Plug-in: install Nutanix AHV Plug-in on the Veeam Backup & Replication server.
3. Perform initial configuration for the cluster that you want to protect:
   a. Add Nutanix AHV cluster.

      In the Veeam Backup & Replication console, add Nutanix AHV cluster to the Veeam Backup & Replication infrastructure.

   b. Deploy AHV Backup Proxy.

      In the Veeam Backup & Replication console, launch the deploy of AHV Backup Proxy on an AHV cluster where VMs that you plan to back up reside.

      > **IMPORTANT**
      >
      > If the DNS server on the Veeam Backup & Replication server is not configured properly and the deployment finishes with the warning that the proxy hostname is not resolved, see Troubleshooting AHV Backup Proxy Deployment.

After performing these steps, you will have a ready-to-use backup infrastructure that allows you to back up and restore VMs and VGs hosted on the Nutanix AHV cluster.

Step 1. Configure Veeam Backup & Replication Server

You must configure Veeam Backup & Replication components that will take part in protection of Nutanix AHV VMs.

1. Install Veeam Backup & Replication

   Alternatively, you can use an already deployed instance of Veeam Backup & Replication version 11 (build 11.0.1.1261) or later.

2. Configure backup repository

   In the Veeam Backup & Replication, add a backup repository where you want to store backups of AHV VMs.

Install Veeam Backup & Replication Server

You must install Veeam Backup & Replication on the virtual or physical machine that meets the system requirements.

To install Veeam Backup & Replication:

1. Download the latest version of Veeam Backup & Replication from the Download Veeam Products page.
2. Mount the installation image to the machine on which you plan to install Veeam Backup & Replication or burn the image file to a flash drive or other removable storage device. If you plan to install Veeam Backup & Replication on a VM, use built-in tools of the virtualization management software to mount the installation image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISOs of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun will open a splash screen with installation options. If Autorun is not available or disabled, run the `Setup.exe` file from the image or disk.
2. In the **Veeam Backup & Replication** section of the splash screen, click **Install** to launch the **Veeam Backup & Replication Setup** wizard.

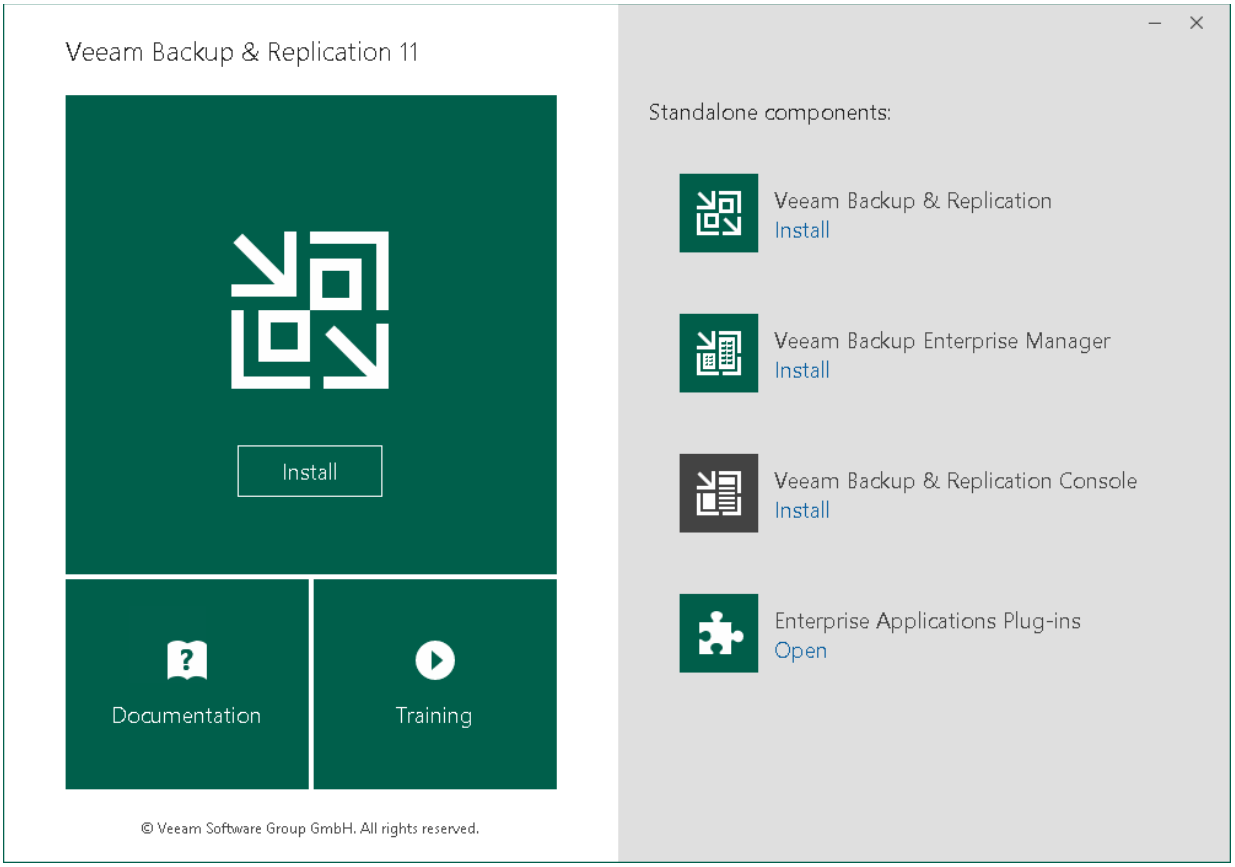

3. Go through the steps of the **Veeam Backup & Replication Setup** wizard. If you have questions on the installation process, see the Installing Veeam Backup & Replication section of the Veeam Backup & Replication User Guide.
4. Once the installation is complete, launch the Veeam Backup & Replication console and go to the next step: Configure Veeam Backup Repository.

What To Do Next

Configure Backup Repository

Configure Backup Repository

Backup repository is a storage location where AHV Backup Proxy keeps backup files. Backup repository is managed in the Veeam Backup & Replication console.

If you leave the default settings, your backups will be stored on the Veeam Backup & Replication server, on the volume with the largest amount of free disk space. If you do not want to store your backups on the Veeam Backup & Replication server, add a new repository.

You can use the following storage types as a repository:

- **Direct attached storage**. You can add virtual and physical machines as backup repositories: Microsoft Windows Servers and Linux Servers
- **Network attached storage**: CIFS (SMB) shares, NFS shares
- **Deduplicating storage appliances**: ExaGrid, Quantum DXi, Dell EMC Data Domain, HPE StoreOnce
- **Object storage**. You can use cloud storage services as backup repositories. For more information, see Object Storage.

You can also use scale-out backup repositories that contain supported repository types.

To learn how to add a backup repository, see the Adding Backup Repositories section of the Veeam Backup & Replication User Guide.

> **NOTE**
>
> Mind the following limitations:
>
> - Veeam Cloud Connect repositories are not supported for backups created by AHV Backup Proxy.
> - [For scale-out backup repositories] Due to specifics of backup jobs for AHV VMs, Veeam Backup for Nutanix AHV always creates a separate backup chain for each VM added to a backup job. Thus, even if you unselect the Use per-VM backup files check box in the advanced settings of a scale-out backup repository, backups of multiple AHV VMs are not stored in a single backup file.

What To Do Next

Install Nutanix AHV Plug-In

Step 2. Install Nutanix AHV Plug-In

Standard installation package of Veeam Backup & Replication does not provide features that allow to protect Nutanix AHV clusters. To be able to back up and restore AHV VMs, you must install Nutanix AHV Plug-in on the Veeam Backup & Replication server.

> **TIP**
>
> If you use a remote Veeam Backup & Replication console, you don't need to install Nutanix AHV Plug-in on the server where the remote Veeam Backup & Replication console is deployed. You must install Nutanix AHV plug-in on the Veeam Backup & Replication server.

To install Nutanix AHV Plug-in, do the following:

1. Download the latest version of the product from the Veeam Backup for Nutanix AHV: Download page.
2. Open the downloaded `VeeamBackupNutanixAHV_3.0.XXX.zip` file and launch the `VeeamBackupNutanixAHV_3.0.XXX.exe` file.
3. The wizard will ask you to install Microsoft .NET Core Runtime 3.1.16 as a prerequisite. Click **OK** to install it.
4. At the **License Agreement** step of the **Nutanix AHV Plug-in for Veeam Backup & Replication** wizard, you must accept the license agreement for Veeam and 3rd party components that Veeam incorporates. If you do not accept the license agreement, you will not be able to pass to next step of the setup wizard.
   a. Read the license agreement.

      To view the license agreement for 3rd party components, click **View**.

   b. Select the **I accept the terms of the Veeam license agreement** check box.

   c. Select the **I accept the terms of the 3rd party components license agreements** check box.

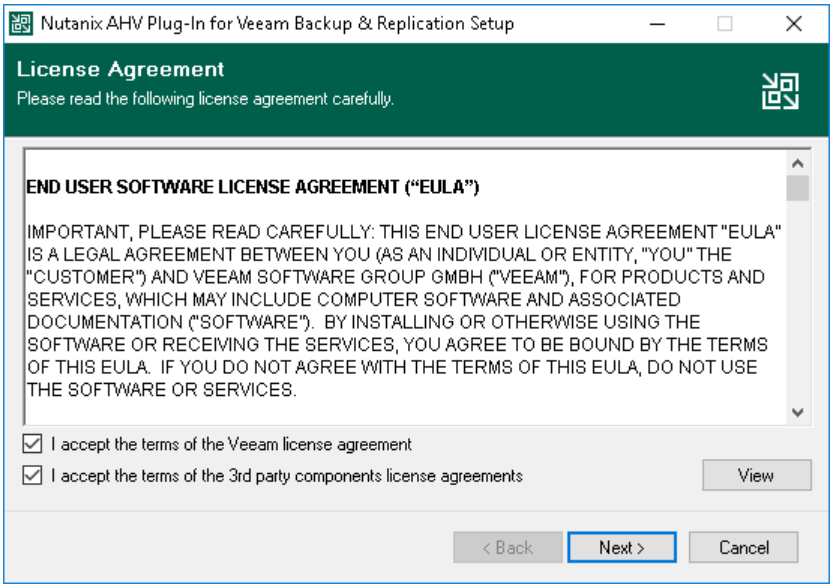

5. At the **Custom Setup** step of the wizard, you can check what components the setup wizard will install on the machine and choose the installation folder.

   To choose the installation folder:

   a. On the right of the **Install to** field, click **Change**.

b. In the **Browse for Folder** window, select the installation folder for the product. The default installation folder is `C:\Program Files\Veeam\Plugins\NutanixAHV\`.

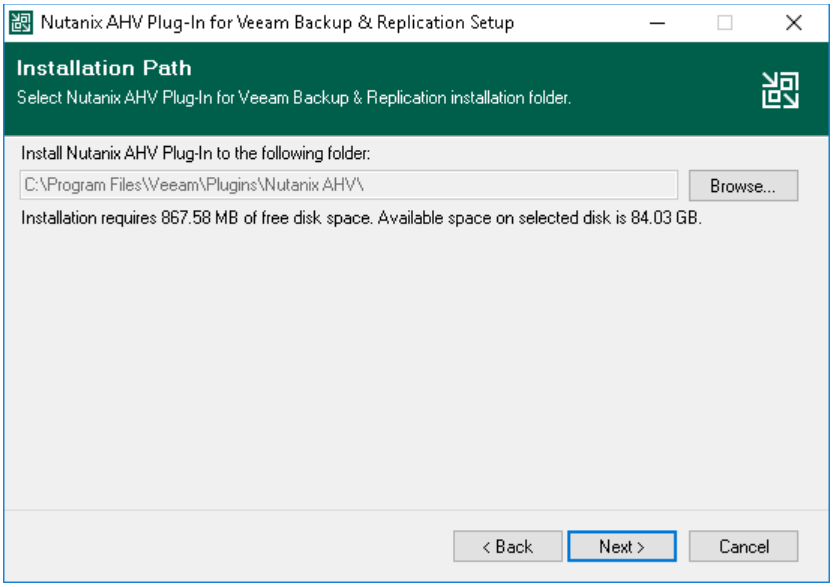

6. Click **Install** to begin the installation.

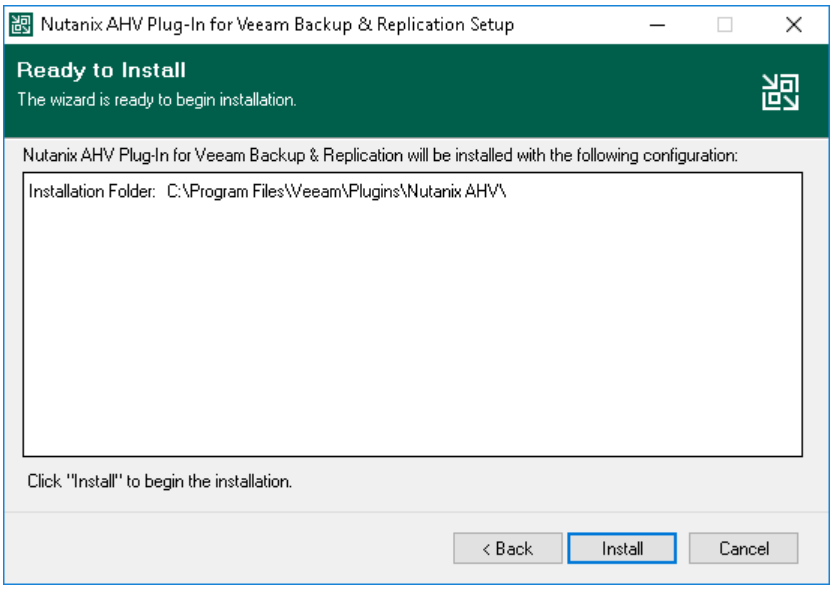

After you install Nutanix AHV Plug-in, you must add Nutanix AHV cluster to the Veeam Backup & Replication infrastructure.

What To Do Next

Perform Initial Configuration

Step 3. Perform Initial Configuration

After configuring Veeam Backup & Replication and installing Nutanix AHV Plug-in, perform initial configuration in the Veeam Backup & Replication console:

1. Add Nutanix AHV cluster

   Add a Nutanix AHV cluster that you want to protect to the Veeam Backup & Replication infrastructure.

2. Deploy AHV Backup Proxy

   Deploy AHV Backup Proxy on an AHV cluster where VMs that you want to back up reside.

Related Topics

Troubleshooting AHV Backup Proxy Deployment

Troubleshooting AHV Backup Proxy Deployment

If the communication between AHV Backup Proxy and Veeam Backup & Replication is disrupted, try the following instructions to resolve the issue:

- Configuring Hostname Resolution
- Disabling HTTP Proxy

Configuring Hostname Resolution

If the DNS server on your network is not configured properly, the deployment of AHV Backup Proxy will finish with the warning that the proxy cannot connect to the Veeam Backup & Replication server. In this case, the AHV Backup Proxy is deployed on the AHV cluster but the Veeam Backup & Replication server is not connected to the AHV Backup Proxy and backup jobs will fail.

The hostname of the Veeam Backup & Replication server must be resolvable on the AHV Backup Proxy server. For instructions on how to configure resolution of hostnames, see the following guides:

- Configuring AHV Backup Proxy server to resolve hostname of Veeam Backup & Replication server
- Configuring Veeam Backup & Replication server to resolve the hostname of AHV Backup Proxy server

> **TIP**
>
> To verify hostname resolution, we recommend to use `ping hostname` commands without the domain name of the target host. It is not recommended to use `nslookup`.

Configuring AHV Backup Proxy server to resolve hostname of Veeam Backup & Replication server

To make the AHV Backup Proxy server correctly resolve the hostname of the Veeam Backup & Replication server, the DNS zone where the AHV Backup Proxy server is located must contain a DNS resource record for the Veeam Backup & Replication server. To add a new DNS resource record, you can use DNS Manager deployed on the DNS server. For instructions, see the Add a DNS Resource Record section of the Microsoft's Networking guide. You can also use the `Add-DnsServerResourceRecordA` PowerShell command.

If you don't want to add a DNS resource record for the Veeam Backup & Replication server in DNS Manager, you can add the host name entry to the `\etc\hosts` file on the AHV Backup Proxy server.

To configure the AHV Backup Proxy server to resolve the hostname of the Veeam Backup & Replication server, do the following:

1. Log in to the Prism Element or Prism Central console.
2. Open the **VM** dashboard.
3. In the list of VMs, select the AHV Backup Proxy VM and click **Launch Console**.
4. Log in to the VM console using credentials, you have specified for AHV Backup Proxy in Specify Credentials.
5. Open the `hosts` file with a text editor:

   ```
   sudo nano /etc/hosts
   ```

6. At the end of the text, add the IP address, FQDN and the hostname for the Veeam Backup & Replication server and the Nutanix CVM.

   ```
   172.11.11.111 vbr_server_hostname.domain.com vbr_server_hostname
   172.11.11.112 cvm_hostname.domain.com cvm_hostname
   ```

> **TIP**
>
> To find the hostname of Nutanix CVM, log in the to server where Nutanix CVM is deployed and run the `hostname` command.
>
> To find the Nutanix CVM IP address, log in the to server where Nutanix CVM is deployed and run the `cat /root/cvm_ip` command.

7. Save the changes and exit the text editor.
8. Restart the networking service.

```
sudo service networking restart
```

After adding the record on the `hosts` file, you will be able to see the proxy in the Veeam Backup & Replication console. Note that you may have to wait several minutes until Veeam Backup & Replication checks the availability of the host.

Configuring Veeam Backup & Replication server to resolve the hostname of AHV Backup Proxy server

Since version 11, Veeam Backup & Replication connects to the AHV Backup Proxy server using the IP address not the hostname. However, if you have upgraded AHV Backup Proxy to version 2.1 and you still use Veeam Backup & Replication version 10, then you need to configure Veeam Backup & Replication server to resolve the hostname of the AHV Backup Proxy server.

To make the Veeam Backup & Replication server correctly resolve the hostname of the AHV Backup Proxy server, the DNS zone where the Veeam Backup & Replication server is located must contain a DNS resource record for the AHV Backup Proxy server. To add a new DNS resource record, you can use DNS Manager deployed on the DNS server. For instructions, see the Add a DNS Resource Record section of the Microsoft's Networking guide. You can also use the `Add-DnsServerResourceRecordA` PowerShell command.

After adding the record on the DNS server, you will be able to see the proxy in the Veeam Backup & Replication console. Note that you may have to wait several minutes until Veeam Backup & Replication checks the availability of the host.

If you don't want to add a DNS resource record for the AHV Backup Proxy server in DNS Manager, you can add the hostname entry to the `%windir%\System32\drivers\etc\hosts` file on the Veeam Backup & Replication server:

1. On the Veeam Backup & Replication server, go to `%windir%\System32\drivers\etc`.
2. Open the `hosts` file with a text editor.
3. In the `hosts` file, add the IP address of the AHV Backup Proxy VM. You can find the IP address of the VM in the Nutanix Prism console.
4. Press [Tab] and add a hostname of the AHV Backup Proxy VM which you have defined in Specify Network Settings.

```
172.28.30.170 ahv-proxy-21
```

5. Save changes in the `hosts` file and wait several minutes until AHV Backup Proxy becomes available in the Veeam Backup & Replication console.

Disabling HTTP Proxy

Communication between AHV Backup Proxy and Veeam Backup & Replication is not possible through an HTTP proxy.

For instructions on how to remove HTTP proxy configured on the Nutanix AHV cluster, see the Configuring an HTTP Proxy section of the Prism Central Guide.

If the WinHTTP proxy is enabled on the Veeam Backup & Replication server, you must disable it to allow communication with the AHV Backup Proxy VM:

1. Check if the WinHTTP proxy is configured on the Veeam Backup & Replication server. Launch CMD as administrator and run the following command:

```
C:\Windows\system32> netsh winhttp show proxy
Proxy Server(s) : http://httpproxy.local:8900
```

2. If the WinHTTP proxy is configured, run the following command:

```
C:\Windows\system32> netsh winhttp reset proxy
Current WinHTTP proxy settings:
Direct access (no proxy server)
```

Upgrading to Version 3.0

If you want to upgrade AHV Backup Proxy from version 2.0 or 2.1 to 3.0, do the following:

1. Make sure the AHV Backup Proxy VM is turned on before you start the upgrade process.
2. Make sure there are no running jobs or restore sessions. We do not recommend to stop running jobs but let them complete successfully. Disable any periodic and backup copy jobs temporarily to prevent them from starting during the upgrade.
3. Upgrade Veeam Backup & Replication to version 11a (build 11.0.1.1261 or later). For more information, see Upgrading to Veeam Backup & Replication 11 section of the Veeam Backup & Replication User Guide.
3. During the upgrading process, the wizard will automatically upgrade Veeam AHV Plug-in to version 3.0.
4. Open the Veeam Backup & Replication console. The upgrade wizard will appear and prompt you to upgrade the product components. Follow the wizard to complete the upgrade process for AHV backup infrastructure components.

> **NOTE**
>
> The update process may take up to one hour.

Upgrading from Veeam Availability for Nutanix AHV 1.0

> **IMPORTANT**
>
> Veeam Availability for Nutanix AHV 1.0 is no longer supported. If you are using version 1.0, we highly recommend to upgrade it to AHV Backup Proxy 3.0.

Before upgrading Veeam Availability for Nutanix AHV to Veeam Backup for Nutanix AHV 3.0, create a configuration backup for rollback in case of necessity. For that doing the following:

1. Click the gear icon on the main menu and select **Appliance Settings**.
2. At the **Summary** tab of the **Appliance Settings** section, click **Configuration Backup**.
3. Select both **Common** and **Events**, input password and click **Download**.

To upgrade from Veeam Availability for Nutanix AHV 1.0 to AHV Backup Proxy 3.0, do the following:

1. Make sure the proxy appliance VM is turned on before you start the upgrade process.
2. Upgrade Veeam Backup & Replication to version 11a (build 11.0.1.1261 or later). For more information, see Upgrading to Veeam Backup & Replication 11 section of the Veeam Backup & Replication User Guide.
2. On the Veeam Backup & Replication server, install Veeam AHV Plug-in.
3. In the Veeam Backup & Replication console, add the AHV cluster to the Veeam Backup & Replication infrastructure.
3. In the Veeam Backup & Replication console, add AHV Backup Proxy. Note that at the **Deployment Mode** step of the **New Nutanix Proxy** wizard, you must select the **Connect proxy** option. With this option, you can register the existing proxy appliance of Veeam Availability for Nutanix AHV 1.0.

   Note that in the Virtual Machine step of the wizard, you must specify the number of vCPU core, RAM size, and the number of concurrent tasks. Otherwise, the AHV Backup Proxy VM will have the default configuration: 4 concurrent tasks, 4 vCPU, 4 GB RAM.

After you add AHV Backup Proxy, it will be automatically upgraded to version 3.0. All the configuration settings (AHV backup proxy settings, jobs, events) will be automatically saved and transferred to the new AHV Backup Proxy.

Uninstalling Nutanix AHV Plug-In

To uninstall Nutanix AHV Plug-in for Veeam Backup & Replication:

1. On the machine where Nutanix AHV Plug-in for Veeam Backup & Replication is installed, open to the **Start** menu and select **Control Panel > Programs > Programs and Features**.
2. In the programs list, right-click **Nutanix AHV Plug-in for Veeam Backup & Replication** and select **Uninstall**.
3. In the opened window, click **Remove**.

Related Topics

- Removing Nutanix AHV Backup Proxies
- Removing AHV Clusters from Backup Infrastructure

Configuration of Backup Infrastructure

After you install Veeam AHV Plug-in, you must configure the backup infrastructure. To be able to back up VMs that reside on an AHV cluster, you must add the Nutanix AHV cluster to the Veeam Backup & Replication infrastructure and deploy AHV Backup Proxy on that cluster. For detailed instructions, see the following guides:

- Adding Nutanix AHV Cluster
- Managing Nutanix AHV Clusters
- Adding Nutanix AHV Backup Proxy
- Managing Nutanix AHV Backup Proxies
- Managing TLS Certificates

Adding Nutanix AHV Cluster

After you install Veeam AHV Plug-in, you can add a Nutanix AHV cluster that you want to protect to the Veeam Backup & Replication infrastructure.

To add a Nutanix AHV cluster to the Veeam Backup & Replication infrastructure, take the following steps.

1. Launch Add AHV Cluster Wizard
2. Specify Cluster Name or Address
3. Specify Credentials
4. Specify Appliance Settings
5. Finish Working with Wizard

Step 1. Launch Add AHV Cluster Wizard

To launch the **New Nutanix AHV Cluster** wizard, do the following:

1. In the Veeam Backup & Replication console, click the **Backup Infrastructure** view.
2. In the inventory pane, right-click **Managed Servers** and click **Add server**.

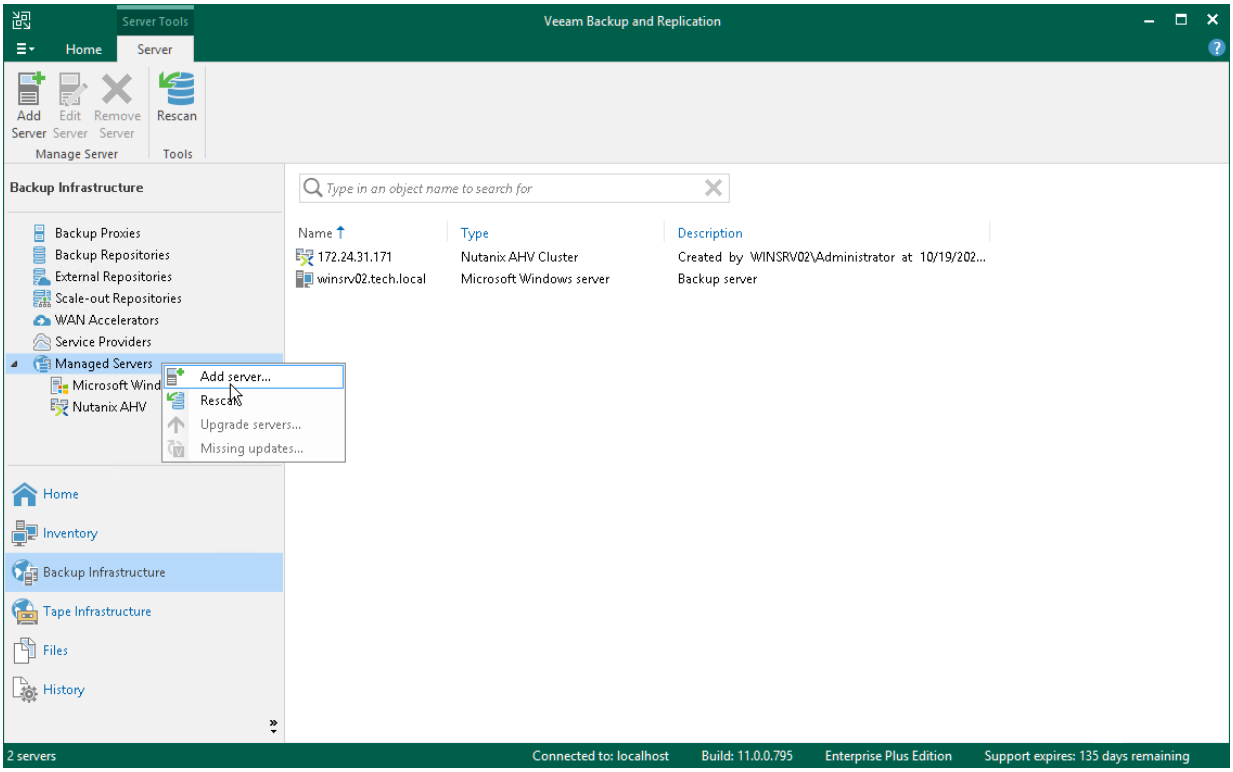

3. In the **Add Server** window, select **Nutanix AHV to launch the Add AHV Cluster wizard.**

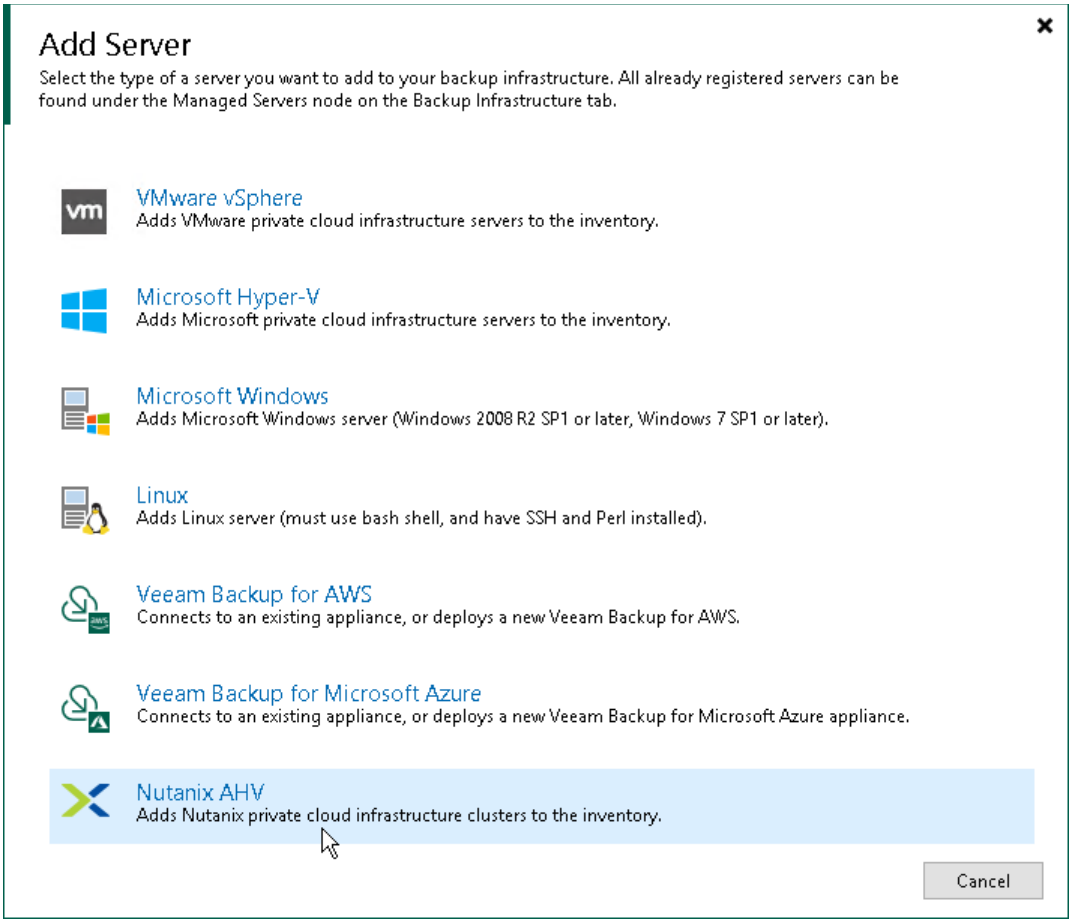

Step 2. Specify Cluster Name or Address

At the **Name** step of the **New Nutanix AHV Cluster** wizard, specify an address and description for the Nutanix AHV cluster.

1. In the **DNS name or IP address** field, specify the FQDN or IP address of the Nutanix AHV cluster.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the cluster, date and time when the cluster was added.

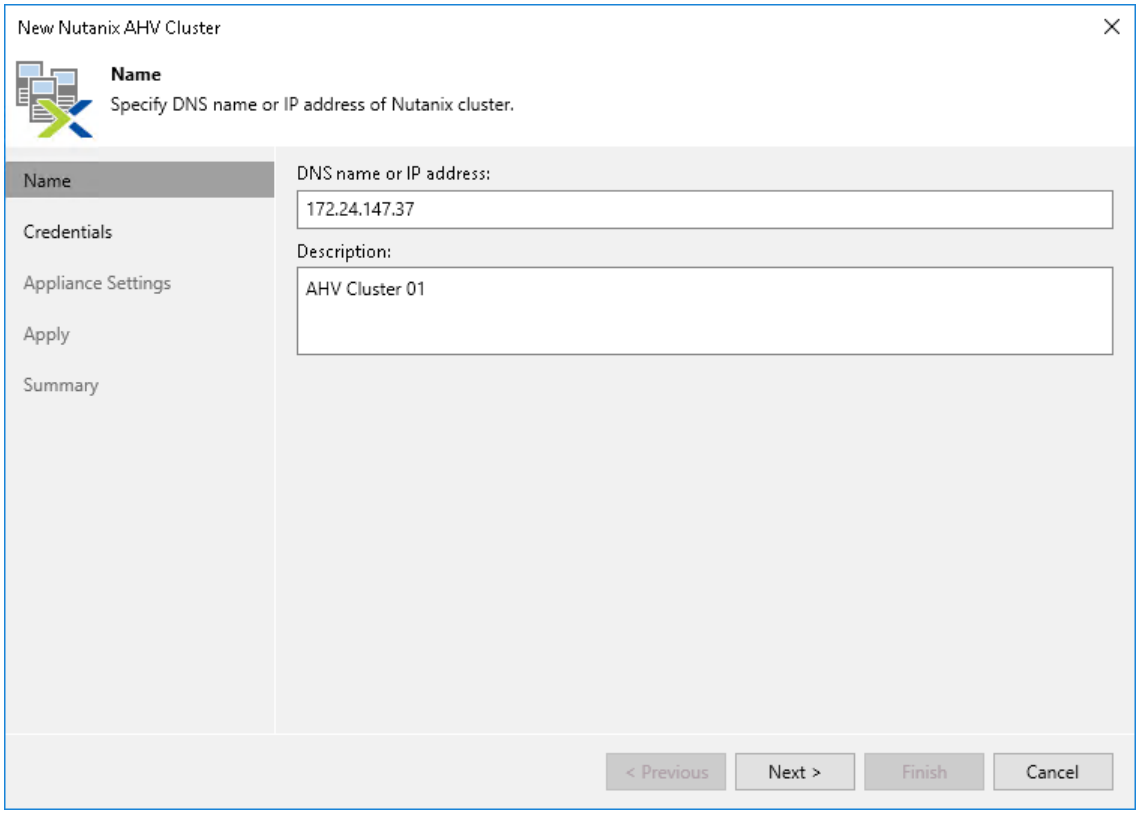

Step 3. Specify Credentials

At the **Credentials** step of the wizard, add an account with access to the target Nutanix AHV cluster.

You can select already added account or use a new one. To add a new account:

1. On the right of the **Credentials** field, click **Add**.
2. Specify the username and password of the account and click **OK**.

   > **IMPORTANT**
   >
   > The specified account must have privileges of the *Prism Admin* role or higher. For more information, see the Access Control Management section of the Prism Central Guide.

3. When adding an AHV cluster, Veeam Backup & Replication checks the TLS certificate installed on the cluster. If the certificate is unknown, Veeam Backup & Replication issues a warning that the certificate is untrusted.
4. In the **Certificate Security Alert** window, click **Continue** to allow Veeam Backup & Replication to connect to the cluster.
5. In the **New Nutanix AHV Cluster** wizard, click **Next**.

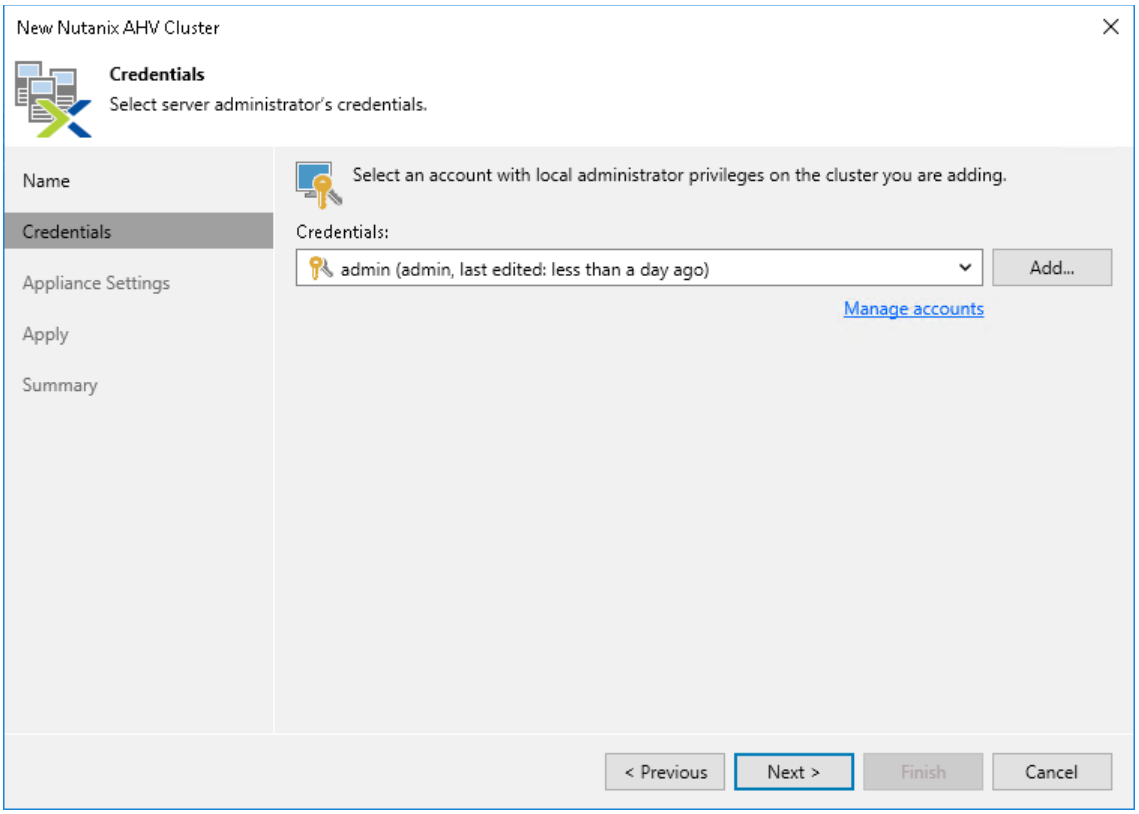

Step 4. Specify Helper Appliance Settings

When you restore guest OS files, application items from a snapshot or Linux VM backup, Veeam Backup & Replication deploys a small helper VM on an AHV cluster. For more information, see Restoring VM Guest OS Files or Application Items Restore.

At the **Appliance Settings** step of the wizard, specify the storage container and network settings for a helper appliance VM which is used for file-level restore.

To define the settings of the helper appliance VM, do the following:

1. [Optional] By default, the helper appliance VM is deployed on a container which has the most free disk space. If you want to change the default container, click **Choose** and select the container where the helper appliance VM will be deployed during the Linux FLR or FLR from a snapshot.
2. [Optional] By default, the helper appliance VM is connected to the first network from the list of networks that Veeam Backup & Replication gets from the AHV cluster. If you want to change the default network, click **Choose** and select the network to which the helper appliance VM will be connected.
3. 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.
4. Set a static IP address for the helper appliance VM or you can select the **Obtain an IP address automatically** check box if DHCP is configured on your server.

   If you set a static IP address for the helper appliance VM, you can perform only one file-level restore session at a time.

5. 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 appliance and download necessary files on their own.

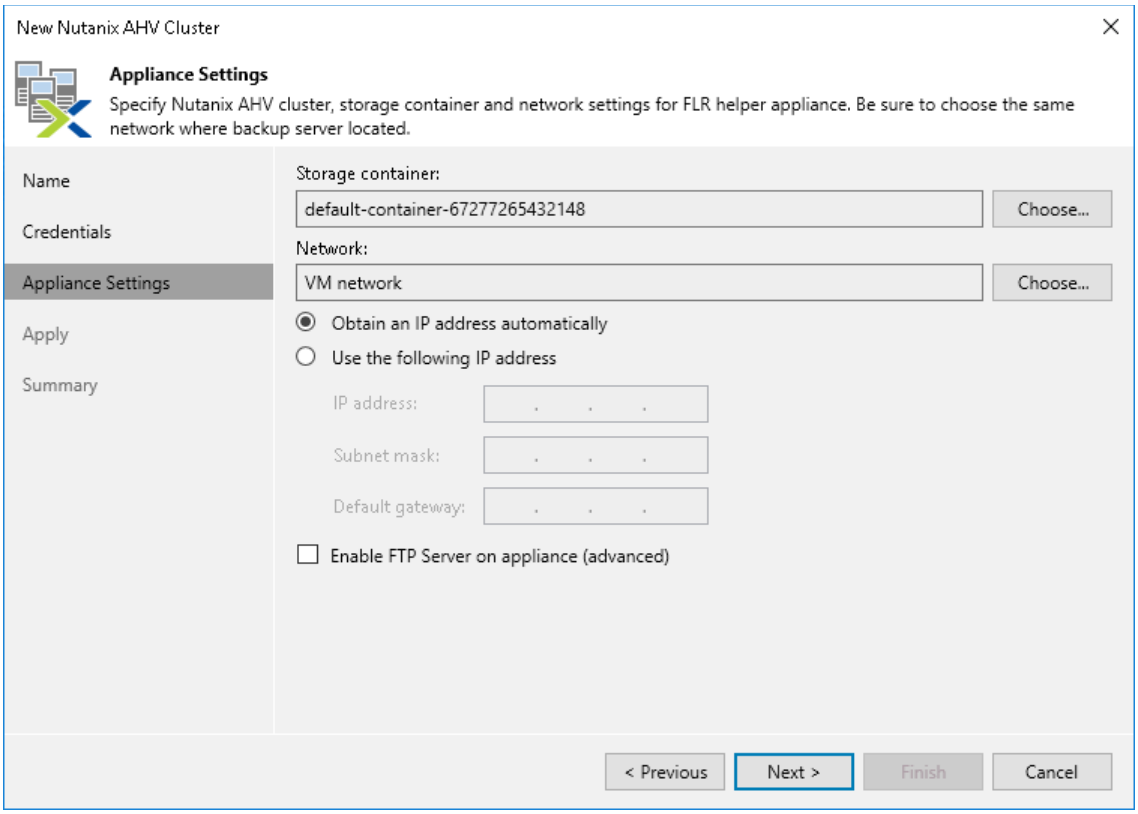

Step 5. Finish Working with Wizard

At the **Apply** and **Summary** steps, finish working with the wizard.

1. At the **Apply** step of the wizard, wait until the cluster is added and click **Next**.

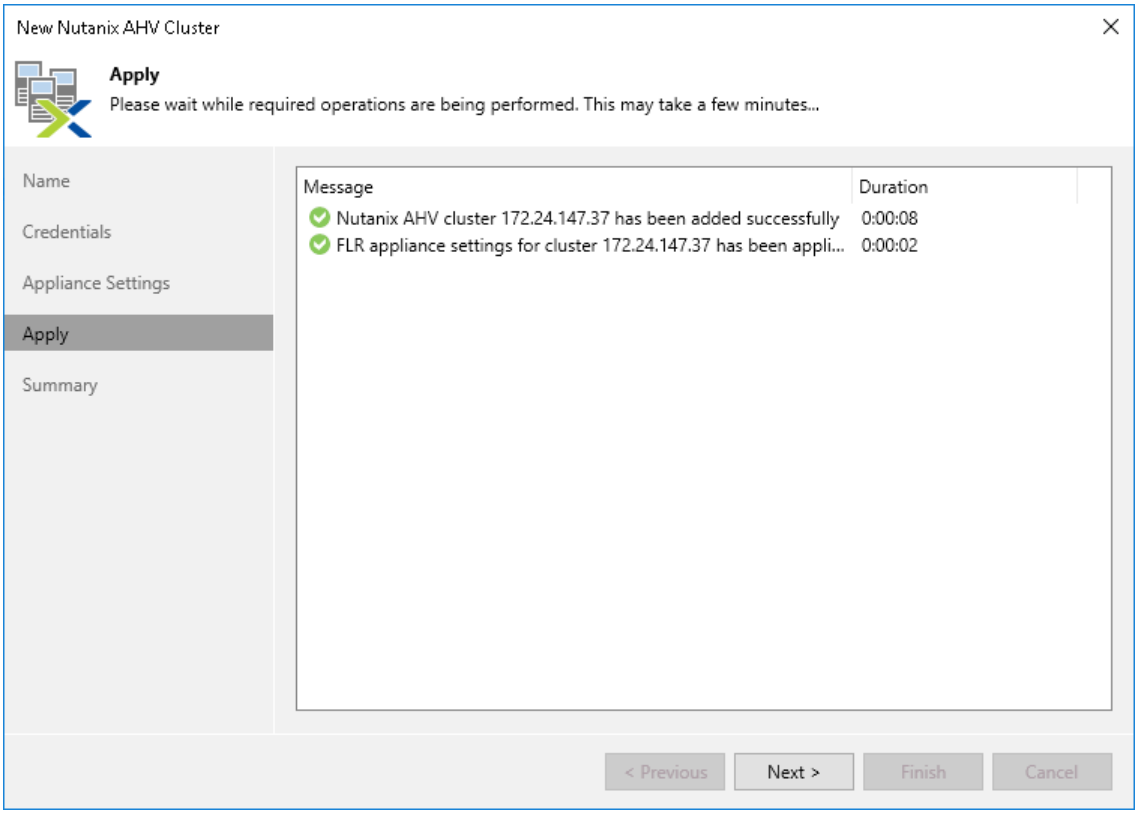

2. At the **Summary** step of the wizard, review the configuration and click **Finish**.

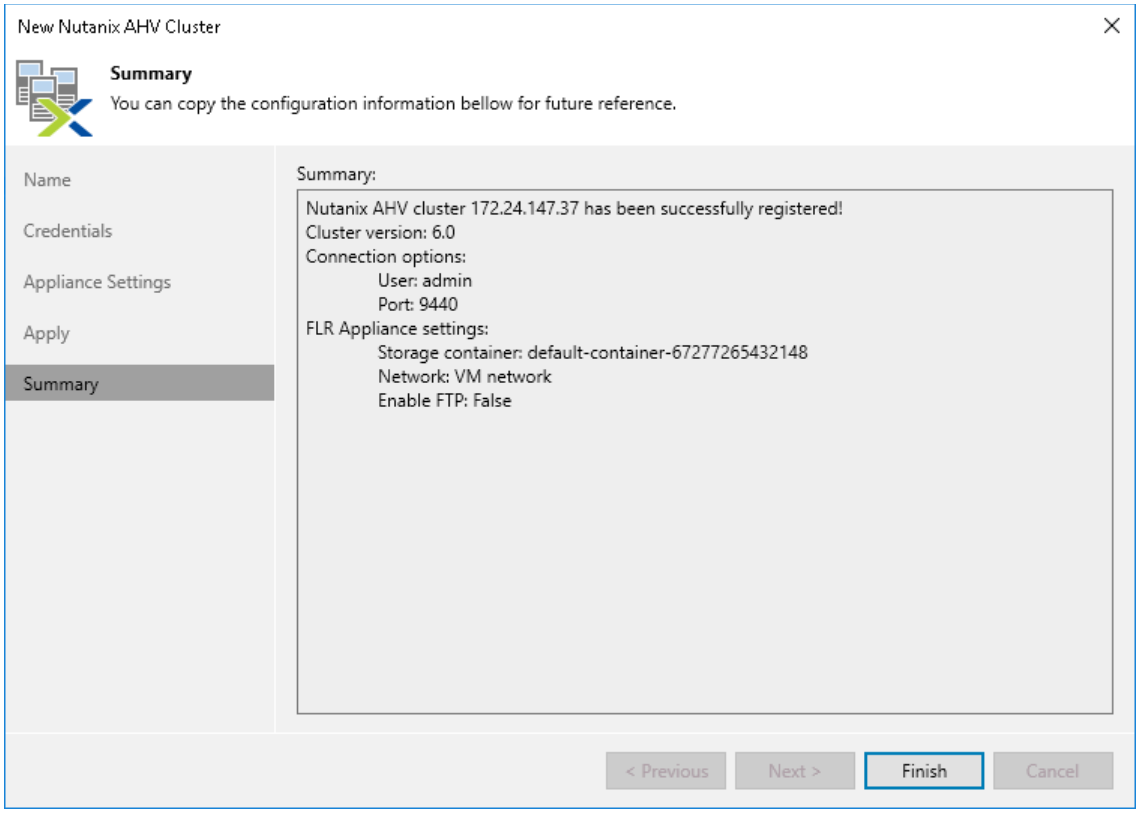

After finishing the wizard, the AHV cluster will be added to the Veeam Backup & Replication infrastructure. To see the added cluster, click the **Backup Infrastructure** view and select **Managed Servers** in the inventory pane.

What To Do Next

Add AHV Backup Proxy

Managing Nutanix AHV Clusters

In the Veeam Backup & Replication console, you can edit, rescan and remove added Nutanix AHV clusters.

In this Section

- Rescanning Nutanix AHV Cluster
- Removing AHV Clusters from Backup Infrastructure
- Editing AHV Cluster Settings
- Adding VMs to Protection Domain (Auto-Protection)

Rescanning Nutanix AHV Cluster

In some cases, you may need to rescan a Nutanix AHV cluster in the backup infrastructure. The rescan operation may be required if you have added or removed new VMs or snapshots to/from the cluster and want to display actual information in the AHV Backup Proxy web console.

During the rescan operation, AHV Backup Proxy retrieves information about the AHV cluster infrastructure and stores this information to the configuration database.

Rescanning AHV Cluster

To rescan an AHV cluster, do the following:

1. At the top right corner of the main menu, click the gear icon and select **Manage AHV Cluster**.

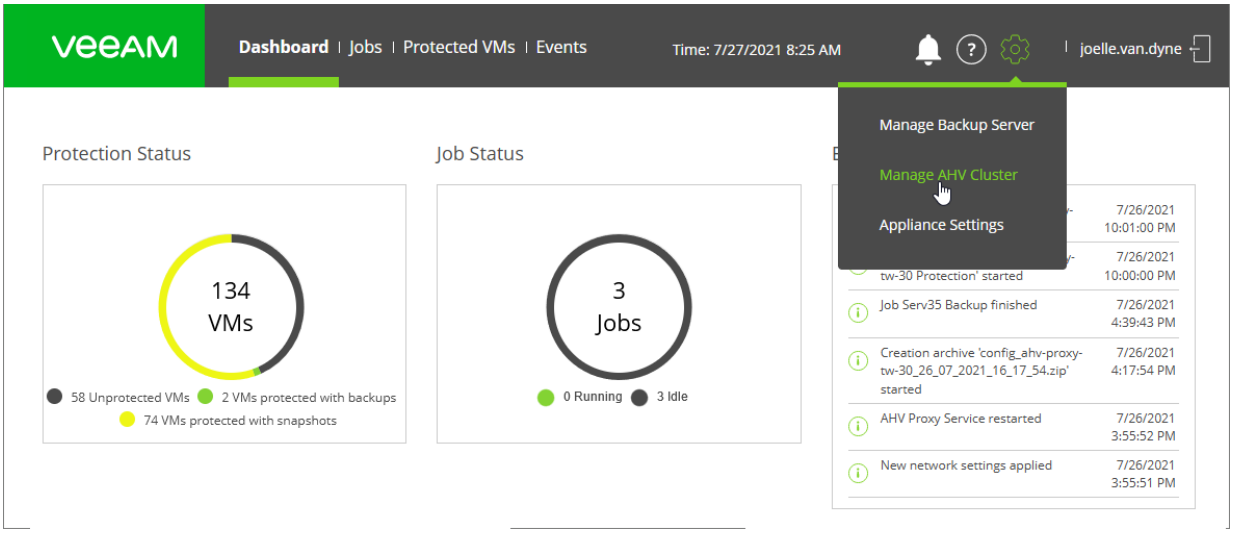

2. Select the AHV cluster and click **Rescan** on the toolbar.

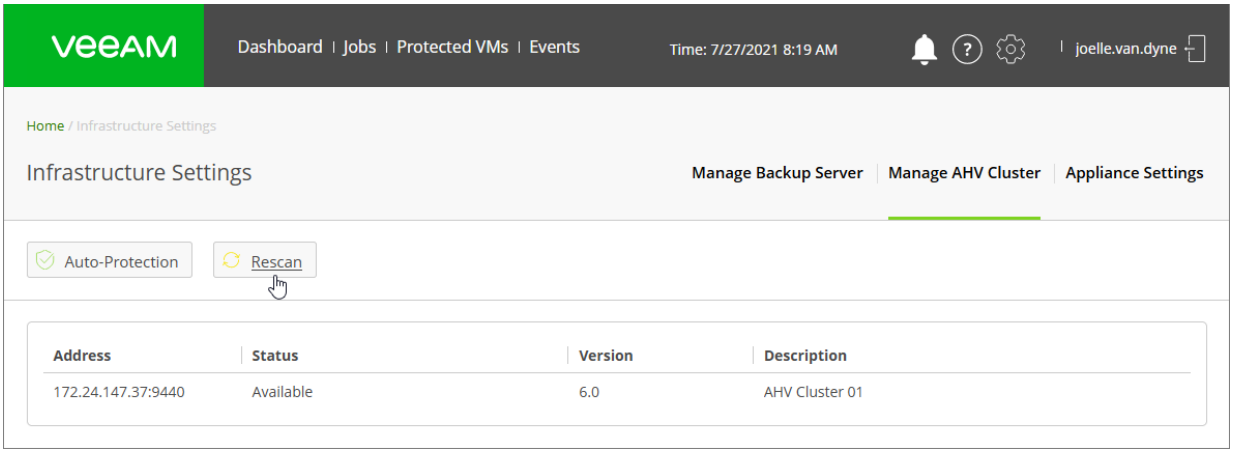

Launching AHV Cluster Rescan in Veeam Backup & Replication Console

To rescan an AHV cluster, do the following:

1. In the Veeam Backup & Replication console, open the **Inventory** view.

2. In the inventory pane, expand the **Virtual Infrastructure** section.
3. Right-click the AHV cluster and click **Rescan**.

Alternatively, you can open the **Inventory** pane, click the AHV cluster and click the **Rescan** tab.

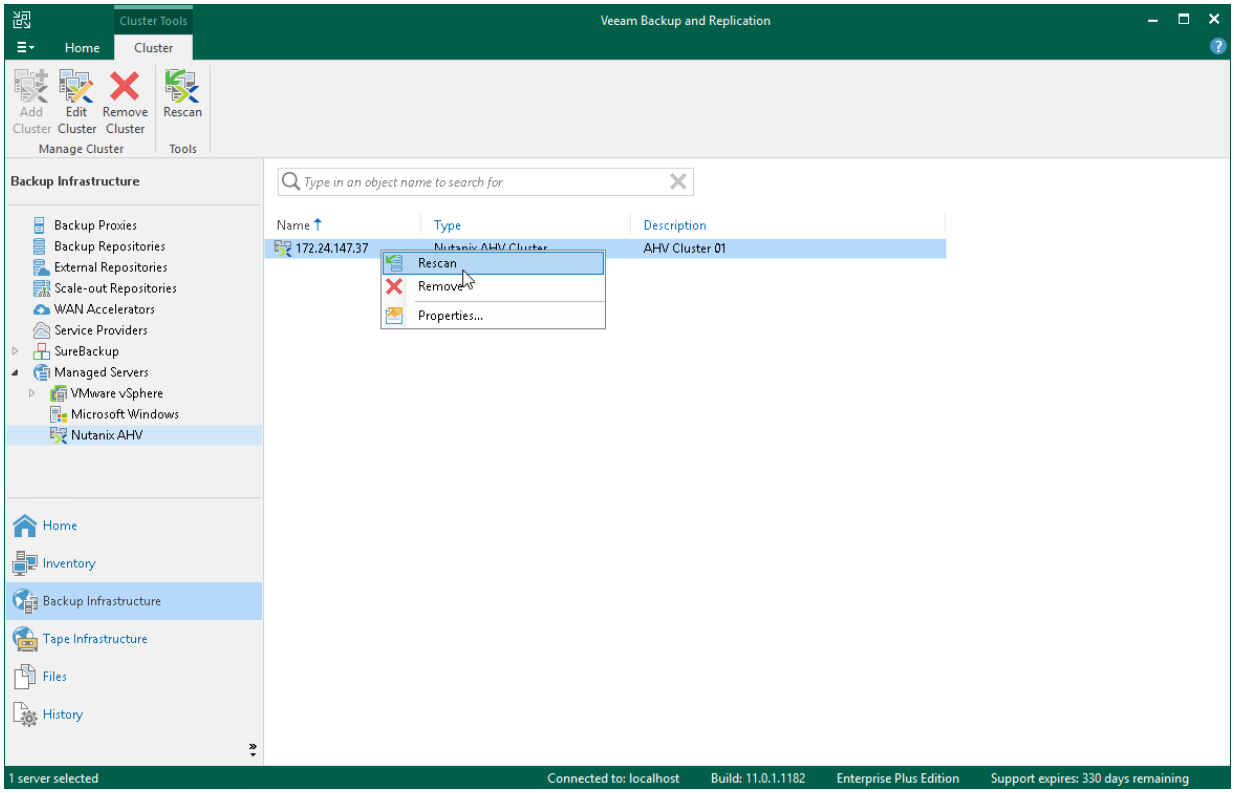

Removing AHV Clusters from Backup Infrastructure

If you do not plan to use an AHV cluster anymore, you can remove it from the Veeam Backup & Replication infrastructure.

You cannot remove a cluster if there is an AHV Backup Proxy instance deployed on that cluster. To remove such cluster, you will need to delete the AHV Backup Proxy instance first.

To remove an AHV cluster from the Veeam Backup & Replication infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers**.
3. In the working area, select the AHV cluster and click **Remove Cluster** on the ribbon or right-click the server and select **Remove**.

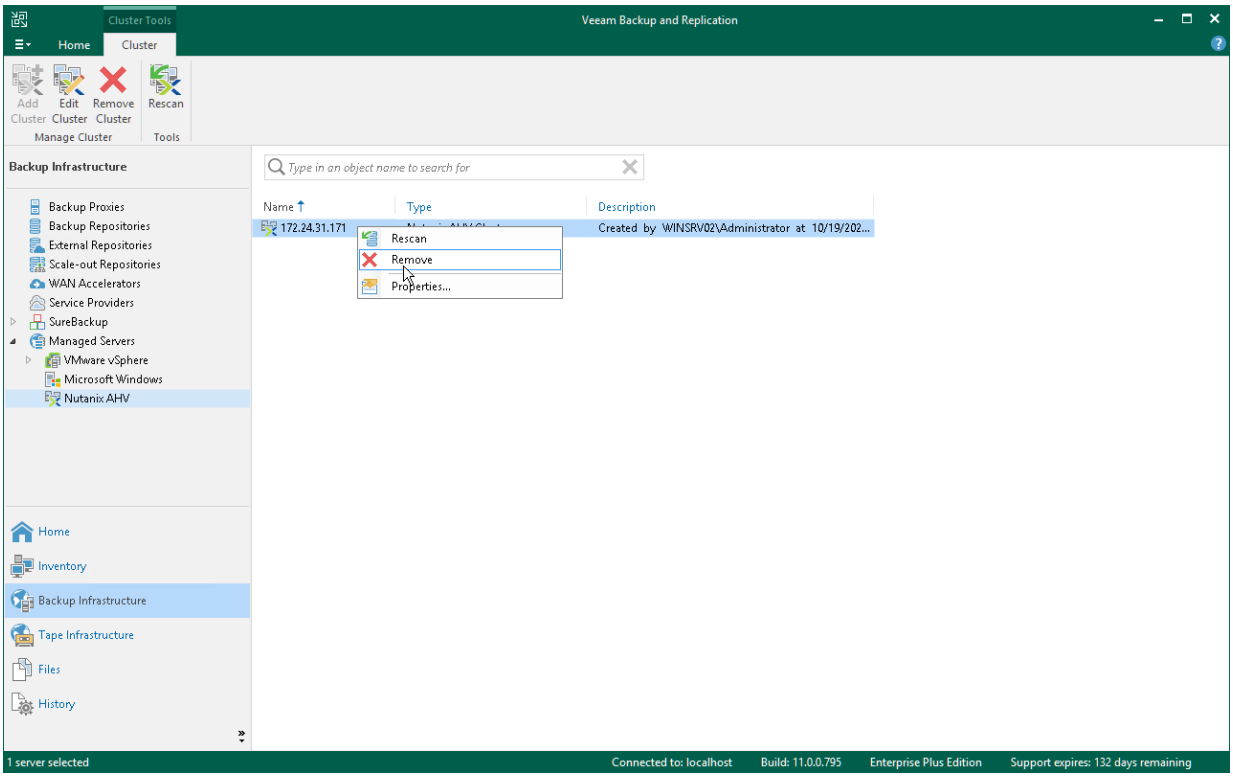

Editing AHV Cluster Settings

To edit settings of an AHV cluster in the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers**.
3. In the working area, select the AHV cluster and click **Edit Cluster** on the ribbon or right-click the cluster and select **Properties**.
4. In the **Edit Nutanix AHV Cluster** wizard, edit cluster settings following the same steps as you have followed when adding the cluster.

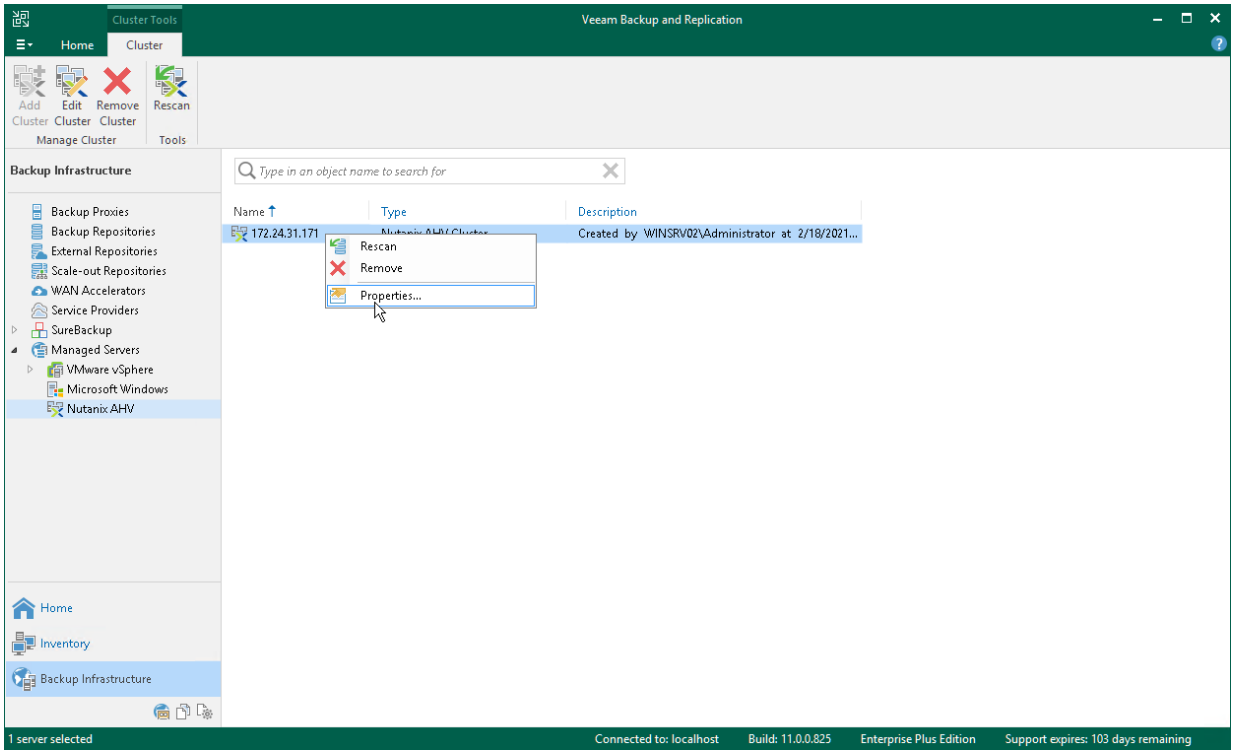

Adding VMs to Protection Domain (Auto-Protection)

You can add all unprotected VMs to a protection domain. To add unprotected VMs to a protection domain, do the following:

1. At the top right corner of the main menu, click the gear icon and select **Manage AHV Cluster**.

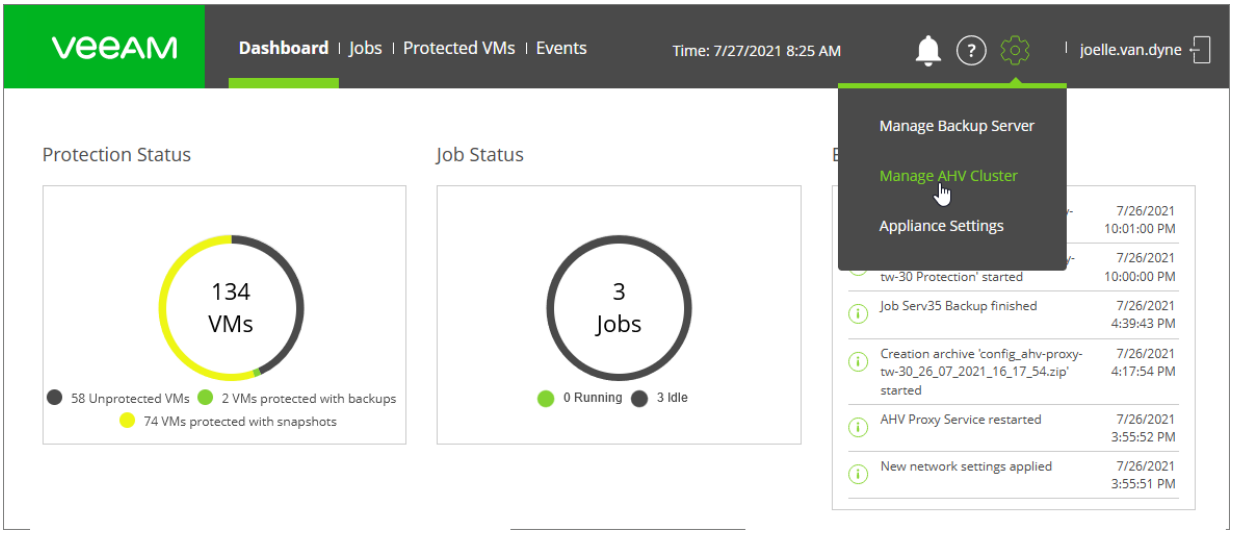

2. Select the AHV cluster and click **Auto-Protection** on the toolbar.

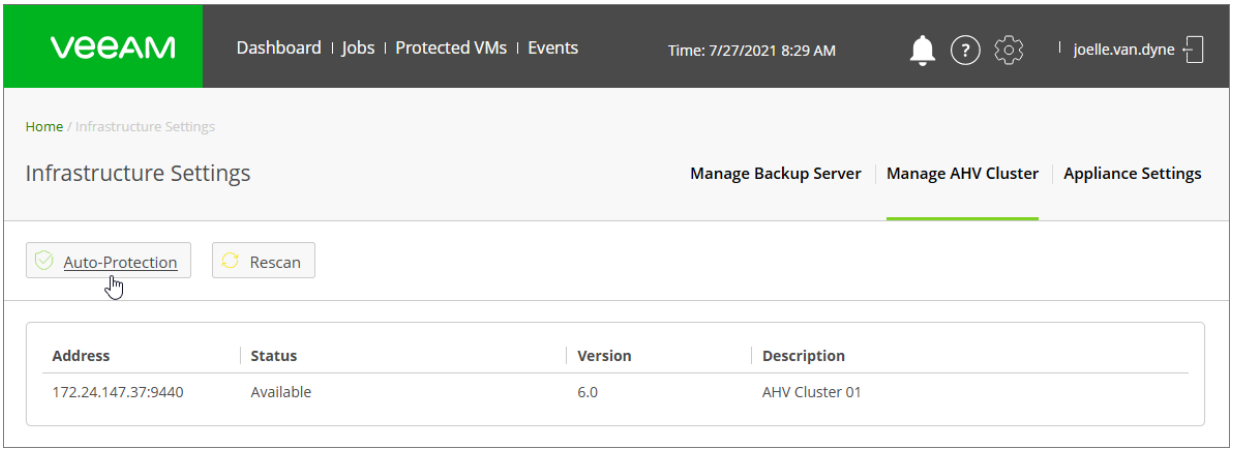

3. In the **VM Auto-Protection** window, switch the **Auto-Protection** slider to the **Enabled** state.

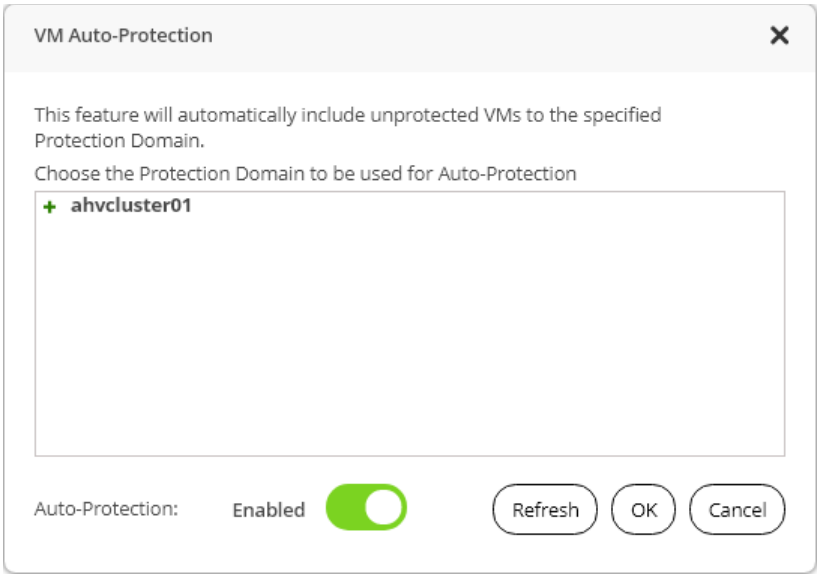

4. Expand the list of protection domains by clicking the **+** icon next to the cluster address.
5. Select a protection domain where you want to add the unprotected VMs.

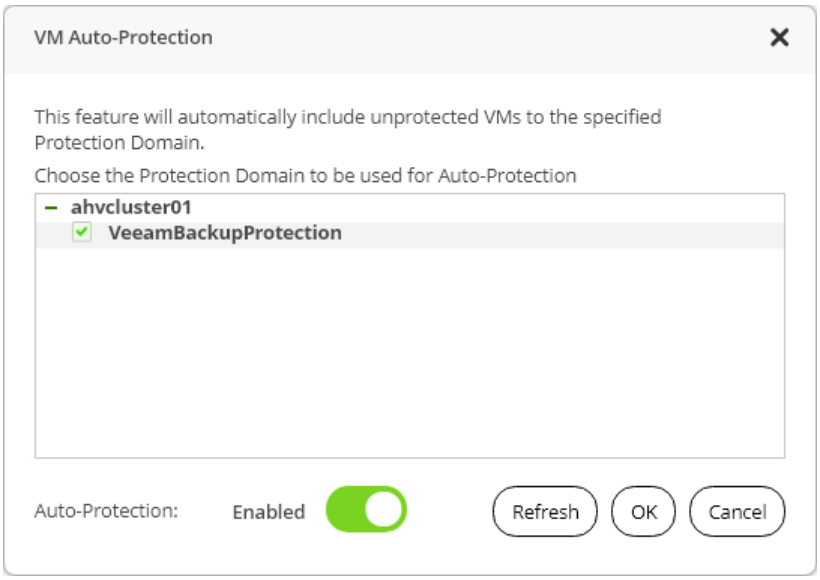

6. Click **Save**.

   All unprotected VMs will be added to the selected protection domain.

Adding AHV Backup Proxy

To be able to back up VMs residing on a Nutanix AHV Cluster, you must add an AHV Backup Proxy to the AHV cluster. AHV Backup Proxy is a virtual module that acts as a coordinator between the AHV cluster on one side and the Veeam Backup & Replication server on the other side. For more information, see How Backup Jobs Work.

> **NOTE**
>
> After the AHV Backup Proxy is successfully deployed, AHV Backup Proxy creates a PD snapshot job and takes a PD snapshot of the AHV Backup Proxy VM. You can use the PD snapshot to restore the VM in case of any problems. For more information, see Protecting AHV Backup Proxy VM.

After you add a Nutanix AHV cluster, you can add an AHV Backup Proxy to the added cluster. To add a new AHV Backup Proxy or register the existing AHV Backup Proxy, take the following steps.

1. Launch New Nutanix Proxy Wizard
2. Select Deployment Mode
3. Specify VM Configuration
4. Specify Network Settings
5. Specify Credentials
6. Grant Permissions
7. Finish Working with Wizard

> **IMPORTANT**
>
> If the DNS server on the Veeam Backup & Replication server is not configured properly and the deployment finishes with the warning that the proxy hostname is not resolved, see Troubleshooting AHV Backup Proxy Deployment.

Step 1. Launch New Nutanix Proxy Wizard

To deploy AHV Backup Proxy on an AHV cluster, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** pane.
2. In the inventory pane, right-click **Backup Proxies** and select **Add Nutanix backup proxy**.

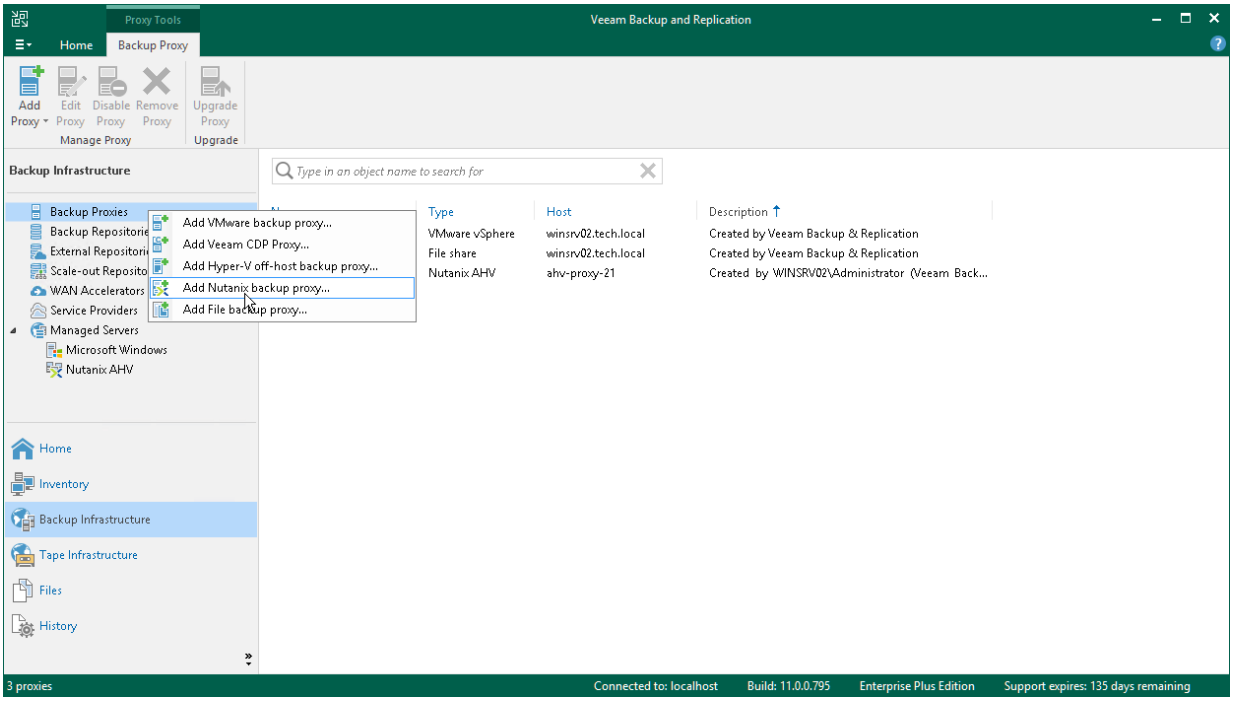

Step 2. Select Deployment Mode

At the **Deployment Mode** step of the **New Nutanix Proxy** wizard, select whether you want to deploy a new AHV Backup Proxy or register an existing one.

- If you want to deploy a new AHV Backup Proxy instance on an AHV cluster, click **Deploy new proxy**.
- If you already have AHV Backup Proxy deployed on an AHV cluster and you want to add this proxy to the Veeam Backup & Replication infrastructure, click **Connect proxy.**

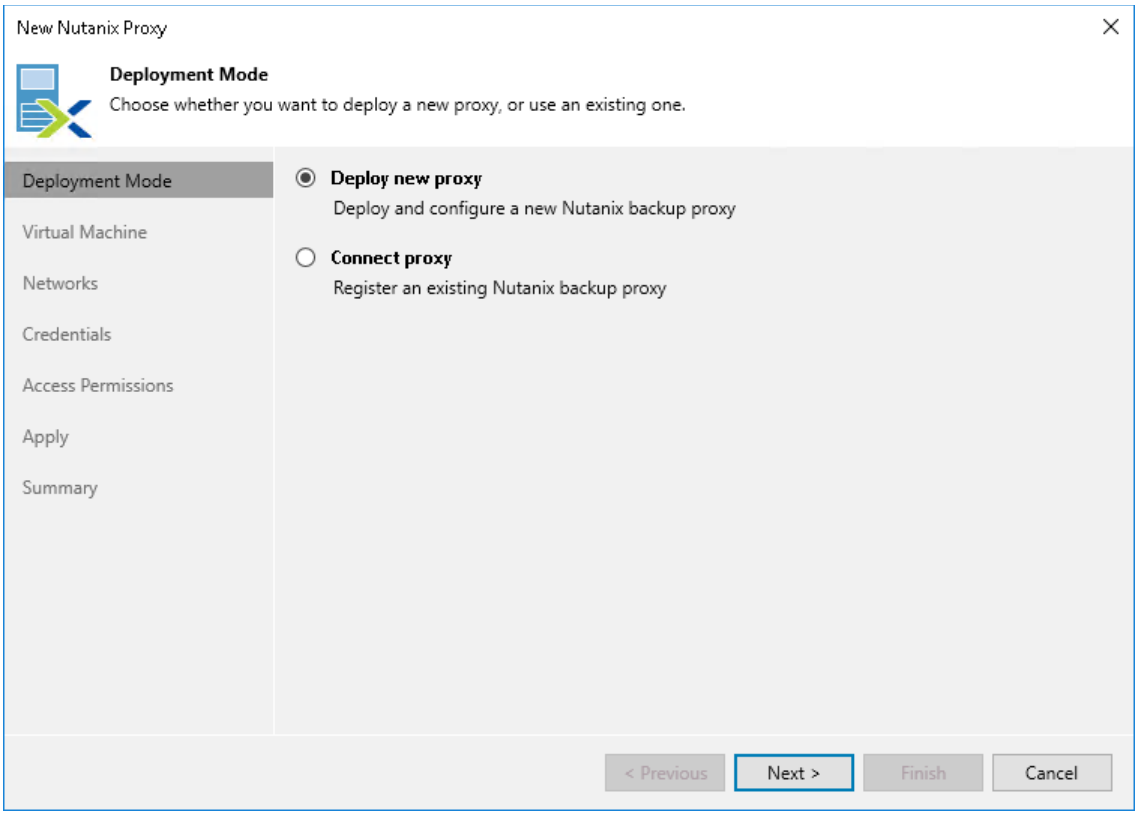

Step 3. Specify VM Configuration

At the **Virtual Machine** step of the wizard, specify the AHV Backup Proxy settings:

1. Click **Choose** and select an AHV cluster where you want to deploy AHV Backup Proxy. By default, the proxy is deployed on a container that has the most free disk space.
2. At the **Name** field, specify a name for the VM of AHV Backup Proxy.
3. By default, AHV Backup Proxy is deployed on a container with the most free space. If you want to change the default container, click **Choose** and select the storage container where AHV Backup Proxy VM files will be stored.
4. In the **Max concurrent tasks** field, specify the number of tasks that AHV Backup Proxy must handle in parallel. If this value is exceeded, the backup proxy will not start a new task until one of current tasks finishes. The default value is set to *4*.

   AHV Backup Proxy creates one task per every VM processed by the job. CPU and hardware resources for the AHV Backup Proxy VM are calculated automatically based on specified number of concurrent tasks. VMs with multiple CPU cores can handle more concurrent tasks. For example, for 4 concurrent tasks, we recommend you to specify minimum 4 CPU cores, for 8 concurrent tasks — 8 CPU cores. When defining the number of concurrent tasks, keep in mind the network traffic throughput in the virtual infrastructure.
5. In the **Proxy description** field, provide a description for future reference. The default description contains information about the user who added the AHV backup proxy, date and time when the proxy was added.
6. [Optional] Click **Advanced** to configure hardware resources for the AHV Backup Proxy virtual machine. By default, the AHV Backup Proxy VM will have 1 CPU, 4 CPU cores, 4GB RAM, and 50GB disk.
   a. In the **Number of vCPU cores** field, specify the number of virtual CPU cores that the AHV Backup Proxy virtual machine will use.
   b. In the **Memory size (GB)** field, specify the memory size that the AHV Backup Proxy virtual machine will use.
   c. Click **OK** to save the configuration.

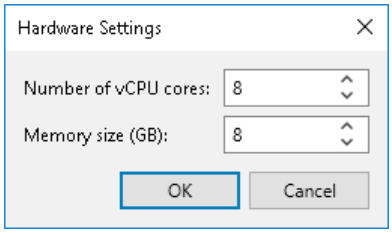

d. In the **New Nutanix Proxy** wizard, click **Next** to continue to the next step of the wizard.

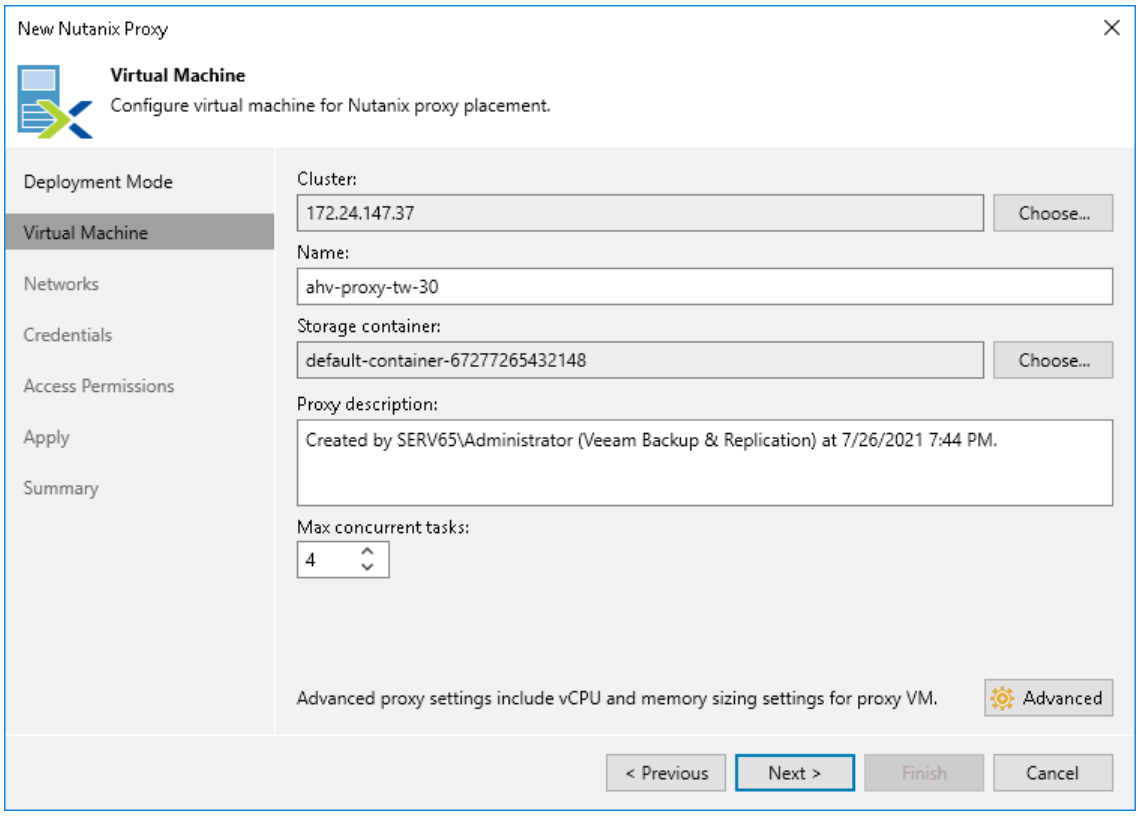

Step 4. Specify Network Settings

At the **Networks** step of the wizard, specify the network settings of the AHV Backup Proxy server.

1. By default, the AHV Backup Proxy VM is connected to the first network from the list of networks that Veeam Backup & Replication gets from the AHV cluster. If you want to change the default network, click **Browse** and select the required network adapter.

   > **IMPORTANT**
   >
   > To provide the best performance, the AHV Backup Proxy VM and Nutanix Controller VM (CVM) must be connected to the same VLAN.

2. In the **Hostname** field, specify a hostname that will be assigned to the AHV Backup Proxy VM.

   The hostname must meet the following requirements:

   a. As the VM OS is based on Linux, the hostname must be in the format of a Linux hostname, not in the format of an FQDN with domain zones.

   b. Hostname must be from 1 to 63 symbols.

   c. Hostname can contain only one type of special character: "-". The "-" character cannot be used as the first or the last symbol in the hostname.

   By default, hostname is the same as the VM name that you assigned in Specify VM Configuration.

   > **IMPORTANT**
   >
   > After deployment of the VM, Veeam Backup & Replication communicates with AHV Backup Proxy using this hostname. Thus, make sure the DNS server on the Veeam Backup & Replication resolves the hostname and IP address of the AHV Backup Proxy VM.
   >
   > If the DNS server on the Veeam Backup & Replication server is not configured properly and AHV Backup Proxy is in the *Not Available* state, see Configuring Hostname Resolution.

3. When you add a new AHV Backup Proxy server, you can continue with default network settings or set them manually.

   If DHCP is enabled for the target network adapter, the IP address and the DNS server for the AHV Backup Proxy VM will be set automatically.

   If DHCP is disabled for the target network adapter or you want to set the IP address manually and/or change DNS server settings, you can specify the network settings manually:

   a. Click **Configure**.

   b. In the **Network Settings** window, specify IP address, subnet mask, default gateway address and DNS server address in the corresponding fields.

> **NOTE**
>
> When you add an existing AHV Backup Proxy server, network settings of the proxy remain unchanged.

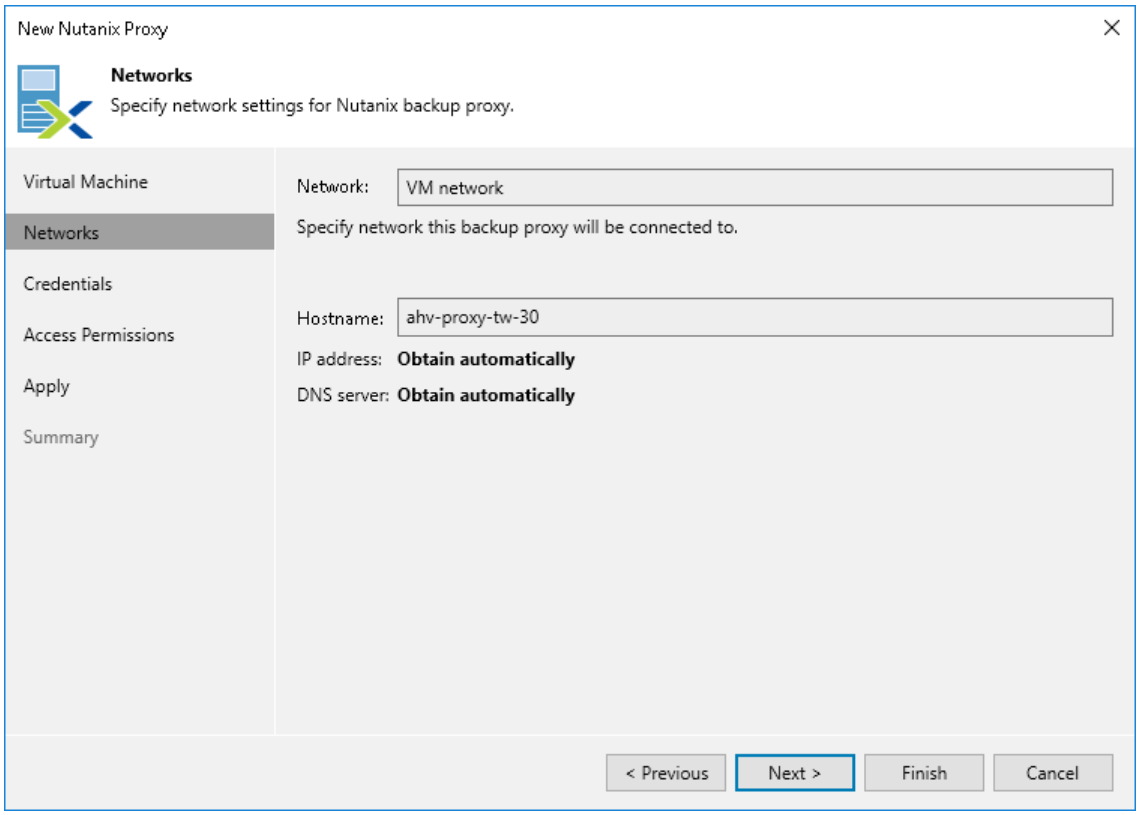

Step 5. Specify Credentials

At the **Credentials** step of the wizard, specify credentials that will be used to access AHV Backup Proxy. The wizard will create a Linux user account on the AHV Backup Proxy VM and the same account credentials will be used to log into the web console of AHV Backup Proxy.

If you have created the account beforehand, you can select the required account from the **Credentials** list. Note that you cannot select the Active Directory accounts.

To create a new account, do the following:

1. Click **Add** on the right of the **Credentials** list.
2. In the **Credentials** window, specify the username and password for the account
3. Click **OK**.

After the AHV Backup Proxy is deployed, use these credentials to connect to the AHV Backup Proxy web console or to the VM where AHV Backup Proxy is deployed.

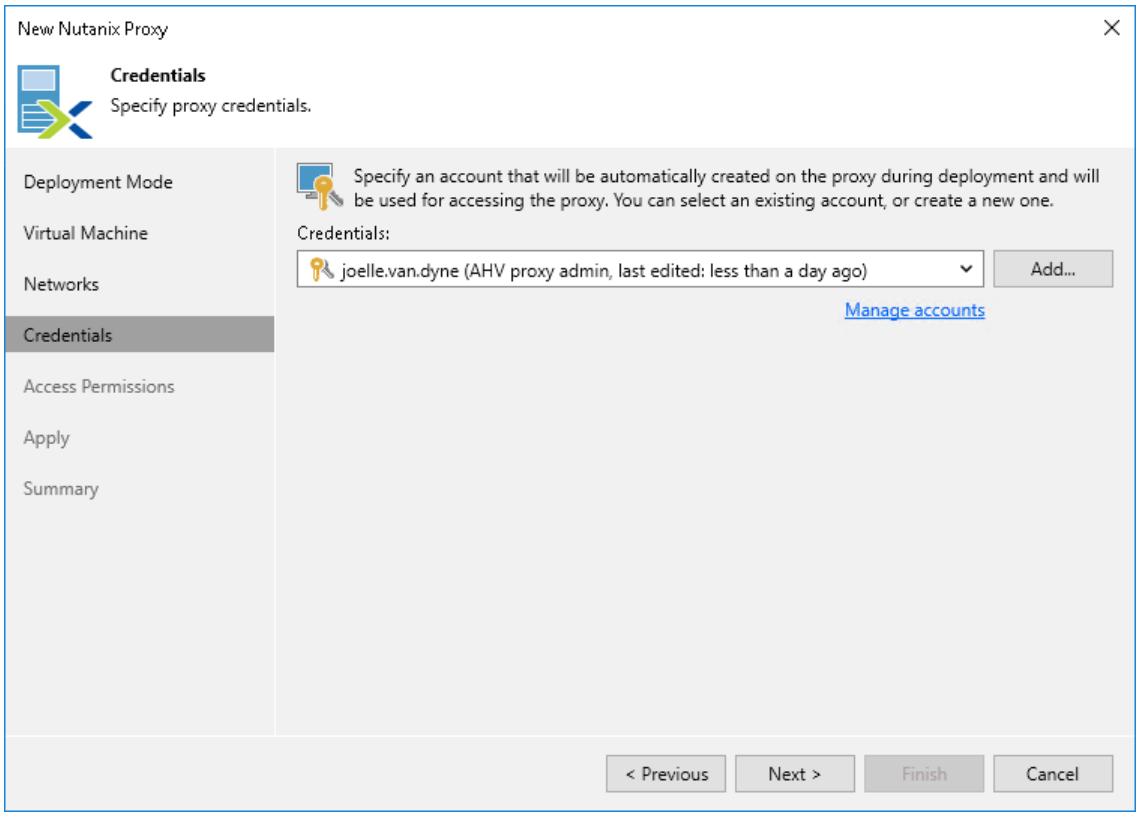

Step 6. Grant Permissions

At the **Access Permissions** step of the wizard, specify backup repositories to which AHV Backup Proxy must have access to.

1. Select whether you want to allow access to all backup repositories or only specific backup repositories.
   - **Allow access to all backup repositories**: AHV Backup Proxy will have access to all backup repositories that are added to the Veeam Backup & Replication infrastructure.
   - **Allow access to the following backup repositories**: AHV Backup Proxy will have access to specified backup repositories. If you select this option, do the following.
2. In the **Repository** field, select backup repositories to which AHV Backup Proxy must have access.

   If you select the second option and don't select any repositories, AHV Backup Proxy will not be able to create backups. You will still be able to create VeeamZIP to folder backups and restore from existing backups using the Veeam Backup & Replication console.

   You can use the **Select All/Clear All** buttons on the right of the **Repository** field to select/unselect the repositories.

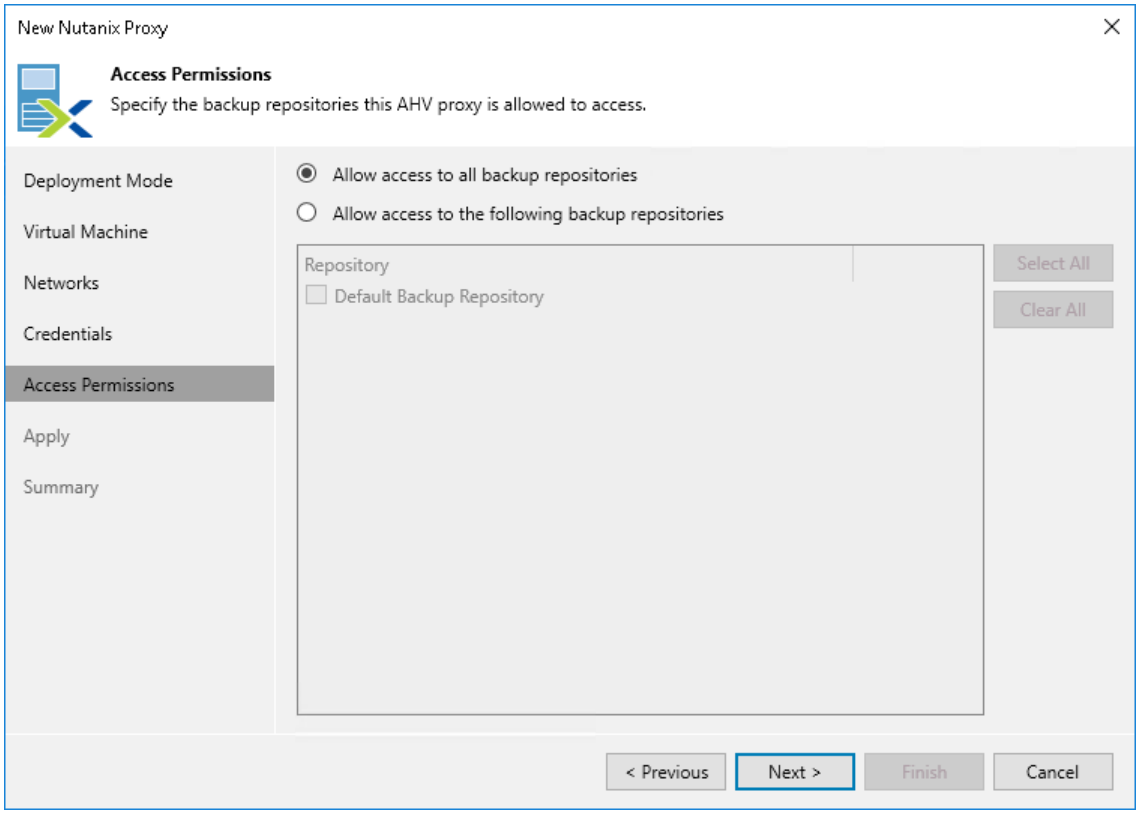

Step 7. Finish Working with Wizard

At the **Apply** and **Summary** steps of the **New Nutanix AHV Proxy** wizard.

1. At the **Apply** step of the wizard, wait until AHV Backup Proxy is deployed and click **Next**.

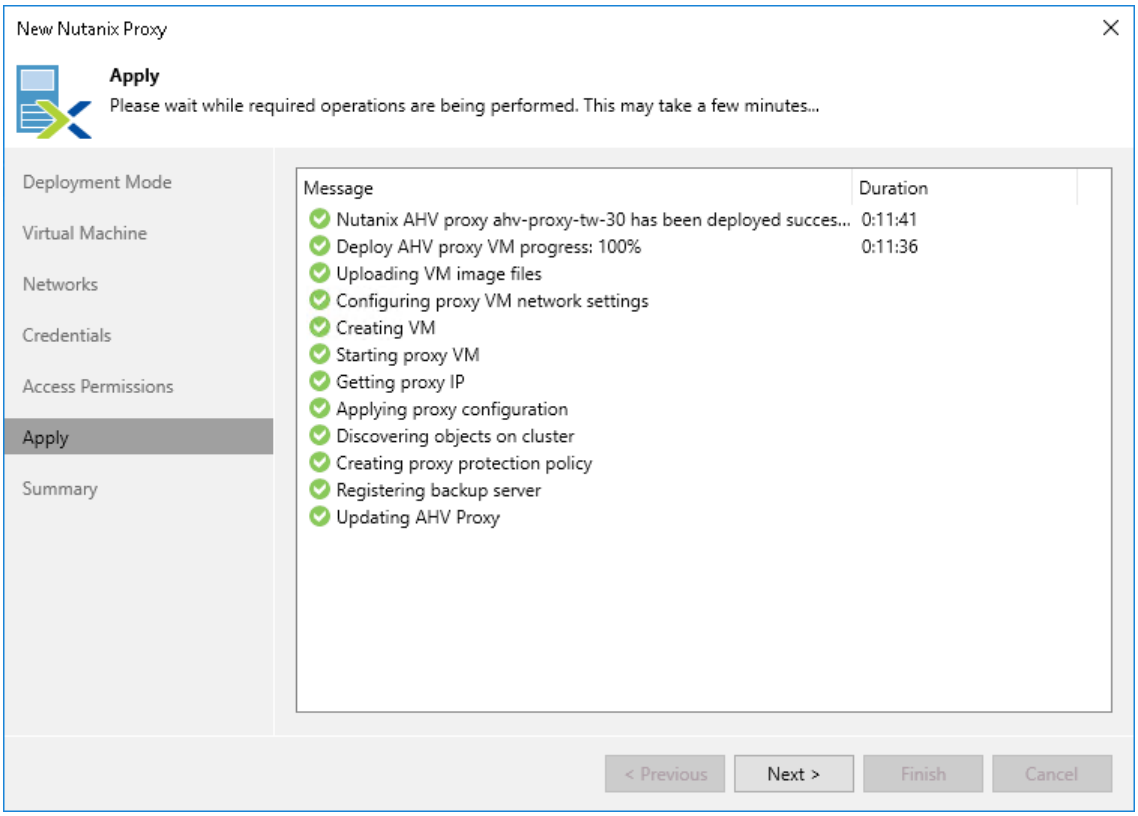

> **IMPORTANT**
>
> If the DNS server on the Veeam Backup & Replication server is not configured properly and the deployment finishes with the warning that the proxy hostname is not resolved, see Troubleshooting AHV Backup Proxy Deployment.

2. At the **Summary** step of the wizard, review settings of the deployed proxy.
   a. At the bottom of the wizard window, you can see the link to the AHV Backup Proxy web console. You can click the link to open the web console in a web browser.

b. Click **Finish** to exit the wizard.

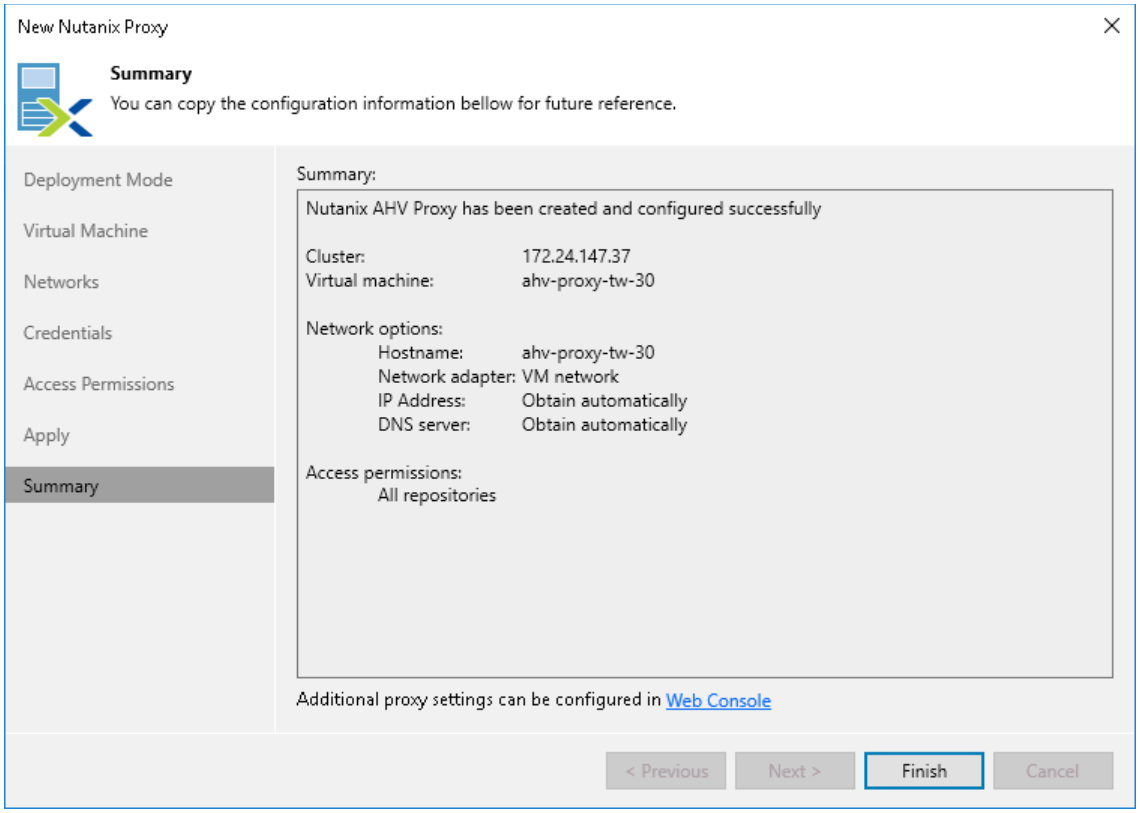

Managing AHV Backup Proxies

In the Veeam Backup & Replication console, you can edit the settings of AHV Backup Proxy, rescan or remove it from the Veeam Backup & Replication infrastructure. For more information, see the following sections:

- Editing Nutanix AHV Backup Proxies
- Rescanning Nutanix AHV Backup Proxies
- Removing Nutanix AHV Backup Proxies

SSH is disabled on the AHV Backup Proxy VM by default. For more information on how to enable SSH, see Enabling SSH on AHV Backup Proxy VM.

Editing Nutanix AHV Backup Proxies

You can edit settings of AHV backup proxies you have added to the backup infrastructure.

> **NOTE**
>
> If you change the number of maximum concurrent tasks in the settings of an AHV backup proxy, the hardware resources (number of CPU cores and RAM) of the AHV Backup Proxy VM will be changed as specified in new settings.

To edit an AHV backup proxy settings, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the AHV backup proxy and click **Edit Proxy** on the ribbon or right-click the AHV backup proxy and select **Properties**.
4. Edit AHV backup proxy settings as required. You will be requested to go through the same steps as you have followed when adding an AHV backup proxy.

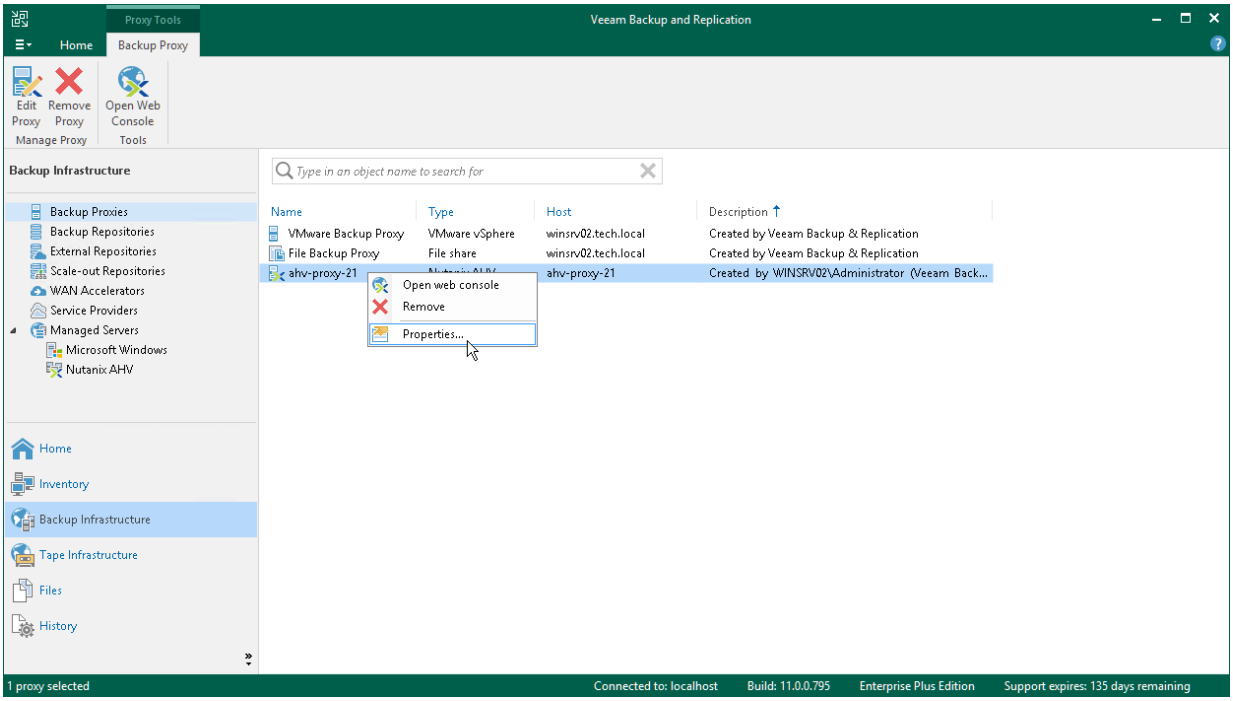

Rescanning Nutanix AHV Backup Proxies

If an AHV Backup Proxy becomes marked as unavailable, you can rescan it to identify inconsistencies between data in the backup console and in the proxy.

During the rescan operation, Veeam Backup & Replication performs full data synchronization between an AHV Backup Proxy and Veeam Backup & Replication. Rescan will erase and re-download AHV Backup Proxy settings and statistics.

To rescan an AHV Backup Proxy, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the AHV backup proxy and click **Rescan Proxy** on the ribbon or right-click the AHV backup proxy and select **Rescan**.

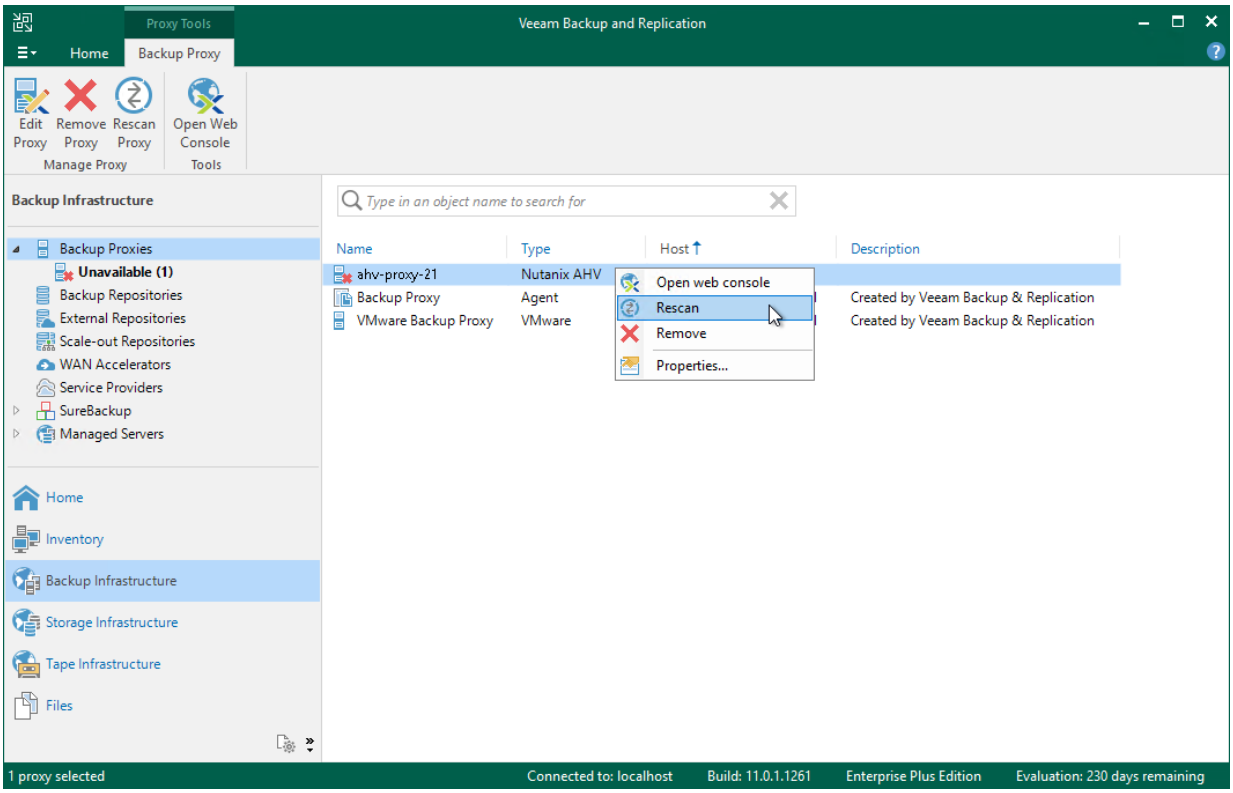

Removing Nutanix AHV Backup Proxies

When you remove an AHV backup proxy from the Veeam Backup & Replication infrastructure, information about AHV VM snapshots, backup/snapshot of the removed AHV backup proxy are deleted from the Veeam Backup & Replication database.

To remove an AHV backup proxy, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the AHV backup proxy and click **Remove Proxy** on the ribbon or right-click the AHV backup proxy and select **Remove**.

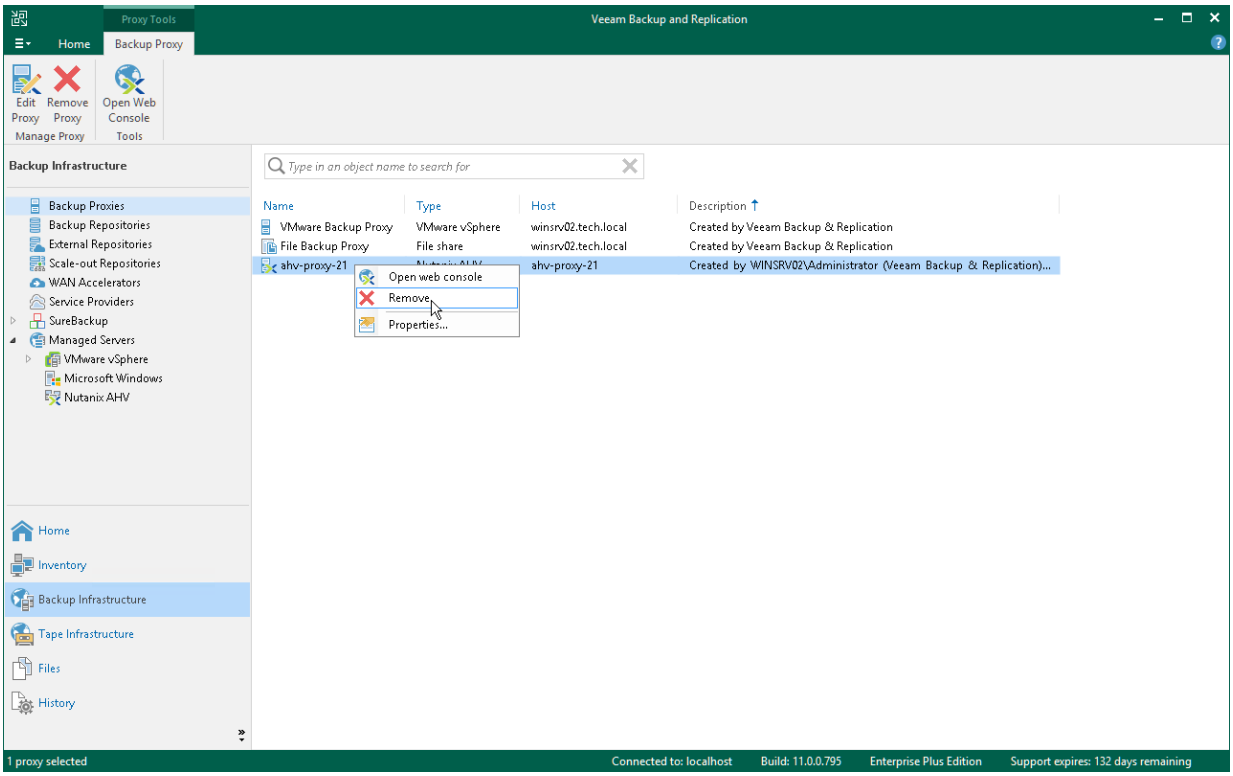

4. In the **Veeam Backup & Replication** window, define whether you want to remove the AHV backup proxy VM from the AHV cluster.

   By default, Veeam Backup & Replication removes information about the AHV backup proxy and its jobs only from the Veeam Backup & Replication database. In this case, AHV Backup Proxy VM and its backup jobs remain on the AHV cluster. This can be useful if you want to connect the AHV Backup Proxy server to another Veeam Backup & Replication server.

If you want to permanently delete the AHV Backup Proxy VM and its backup jobs, select the **Permanently delete the proxy VM from the cluster and all of its backup jobs (irreversible)** check box and click **Yes**.

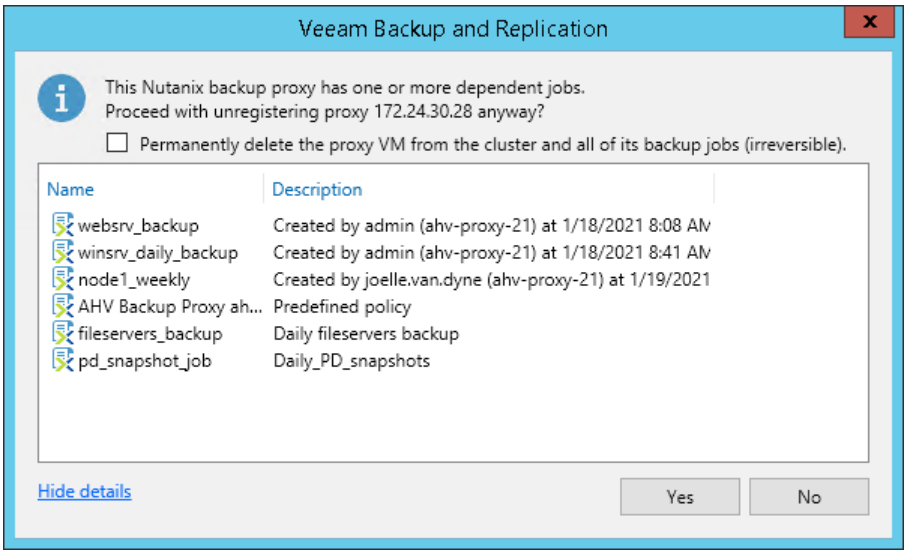

Enabling SSH on AHV Backup Proxy VM

SSH is disabled on the AHV Backup Proxy VM by default. If you want to enable SSH on the AHV Backup Proxy VM, do the following:

1. Log in to the web console of Nutanix Prism that manages the AHV cluster protected with AHV Backup Proxy.
2. From the pull-down list on the left of the main menu, select **VM**.
3. In the **VM** dashboard, click the **Table** view.
4. Select the VM where AHV Backup Proxy is deployed and click **Launch Console** on the toolbar.
5. In the VM console, log in using credentials specified in the Specify Credentials step of the New Nutanix Proxy wizard.
6. In the VM console, do one of the following.
   - To enable SSH for one session, execute the following command.

     ```
     sudo /etc/init.d/ssh start
     ```

   - To enable SSH permanently, execute the following commands.

     ```
     sudo systemctl enable ssh.socket
     sudo /etc/init.d/ssh restart
     ```

Managing TLS Certificates

When you configure the Veeam Backup for Nutanix AHV infrastructure, you can specify what TLS certificate must be used to establish a secure connection between the Veeam Backup & Replication server and the AHV Backup Proxy server. Veeam Backup & Replication offers the following options for TLS certificates:

- You can choose to keep the default self-signed TLS certificate generated by Veeam Backup & Replication.
- You can use Veeam Backup & Replication to generate a new self-signed TLS certificate. For more information, see Generating Self-Signed Certificates.
- You can select an existing TLS certificate from the certificates store. For more information, see Importing Certificates from Certificate Store.
- You can import a TLS certificate from a file in the PFX format. For more information, see Importing Certificates from PFX Files.

If you plan to use a certificate issued by your own Certificate Authority (CA), make sure that the certificate meets the requirements. For more information, see Using Certificate Signed by Internal CA.

Generating Self-Signed Certificates

You can use Veeam Backup & Replication to generate a self-signed certificate for authenticating parties in the Veeam Backup for Nutanix AHV infrastructure.

To generate TLS certificates, Veeam Backup & Replication employs the RSA Full cryptographic service provider by Microsoft Windows installed on the Veeam Backup & Replication server. The created TLS certificate is saved to the *Shared* certificate store. The following types of users can access the generated TLS certificate:

- User who created the TLS certificate
- LocalSystem user account
- Local Administrators group

If you use a self-signed TLS certificate generated by Veeam Backup & Replication in case of initial AHV Backup Proxy deployment, you do not need to take any additional actions to deploy the TLS certificate on an AHV Backup Proxy VM. When you add an AHV Backup Proxy to the Veeam Backup & Replication infrastructure, a matching TLS certificate with a public key is installed on the AHV Backup Proxy VM. During discovery, Veeam Installer Service deployed on the AHV Backup Proxy VM retrieves the TLS certificate with a public key from the Veeam Backup & Replication server and installs a TLS certificate with a public key on the AHV Backup Proxy VM.

If you currently work with deployed AHV Backup Proxy and you try to re-generate self-signed TLS certificate, the connection with Veeam Backup & Replication server will be lost and all subsequent backup jobs fail accordingly. In this case, you can manually re-apply settings for AHV Backup Proxy in the Veeam Backup & Replication console.

> **NOTE**
>
> When you generate a self-signed TLS certificate with Veeam Backup & Replication, you cannot include several aliases to the certificate and specify a custom value in the *Subject* field. The *Subject* field value is taken from the Veeam Backup & Replication license installed on the Veeam Backup & Replication server.

To generate a self-signed TLS certificate:

1. From the main menu of the Veeam Backup & Replication console, select **General Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.

4. At the **Certificate Type** step of the wizard, select **Generate new certificate**.

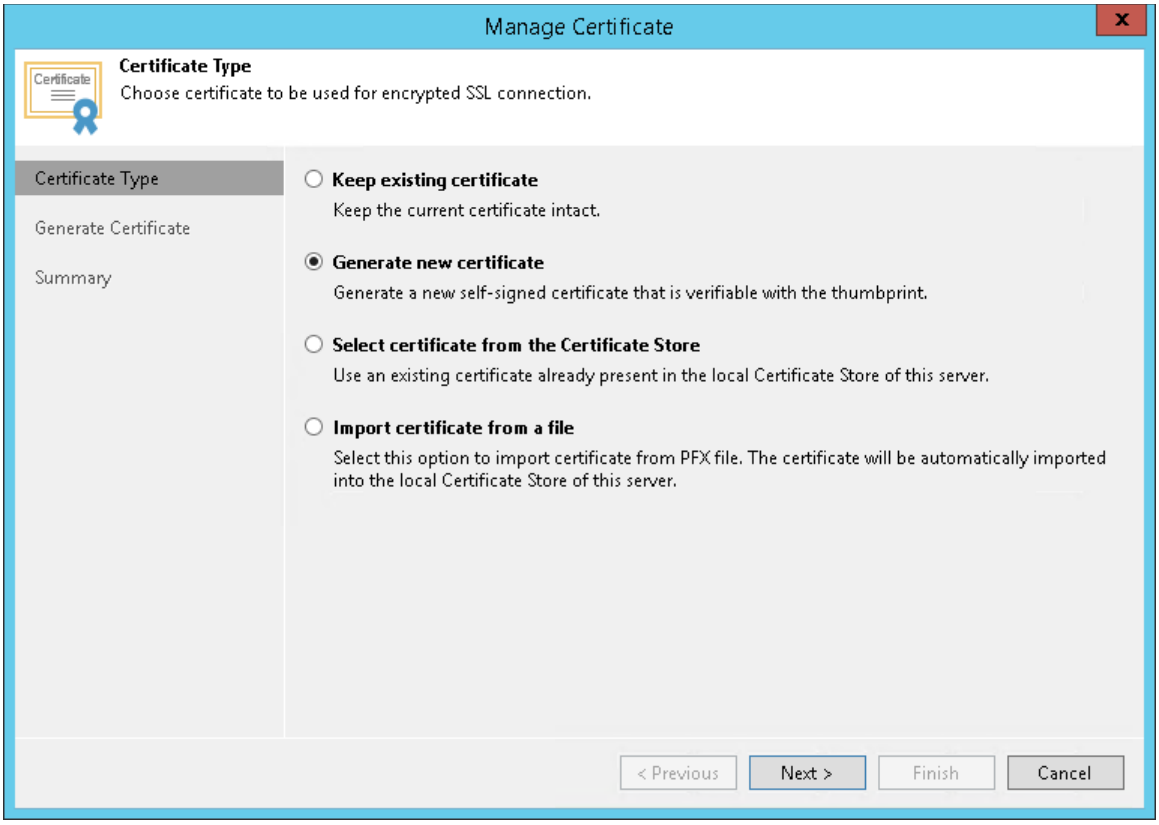

5. At the **Generate Certificate** step of the wizard, specify a friendly name for the created self-signed TLS certificate.

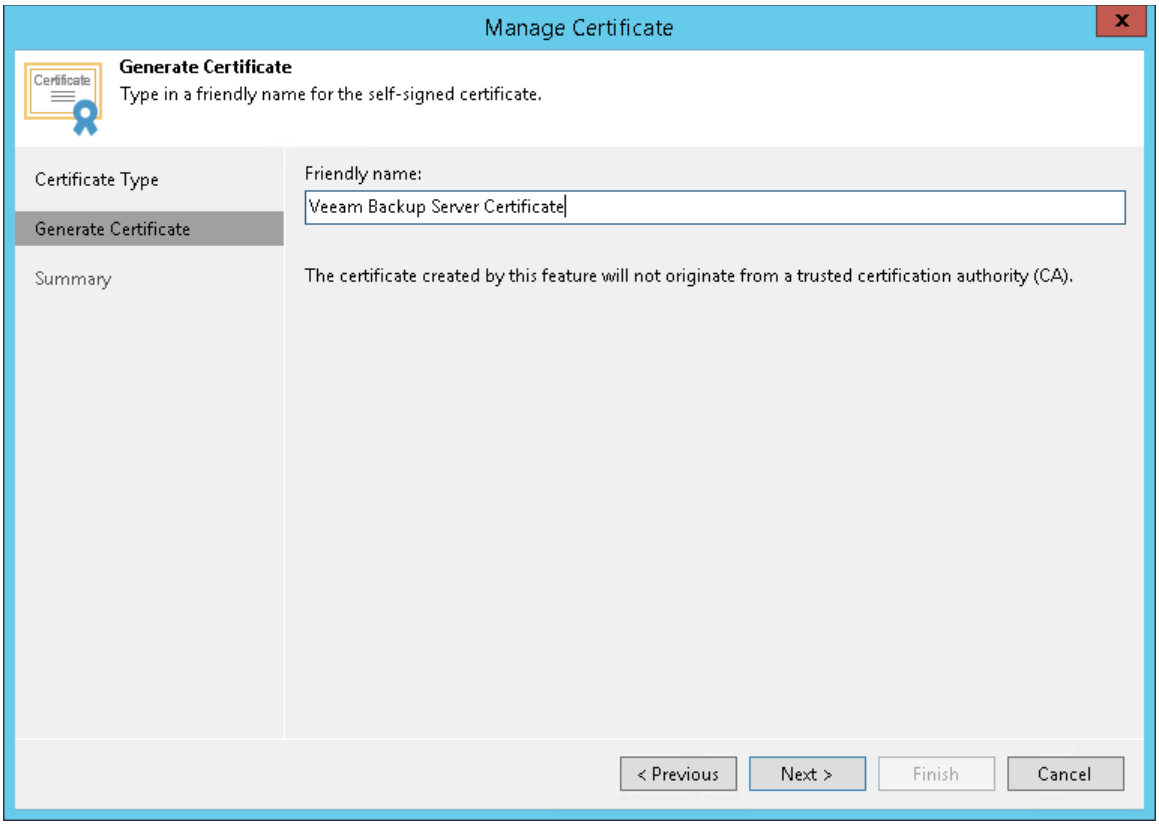

6. At the **Summary** step of the wizard, review the certificate properties. Use the **Copy to clipboard** link to copy and save information about the generated TLS certificate. You will be able to use the copied information to verify the TLS certificate with the certificate thumbprint.

7. Click **Finish**. Veeam Backup & Replication will save the generated certificate in the *Shared* certificate store on the Veeam Backup & Replication server.

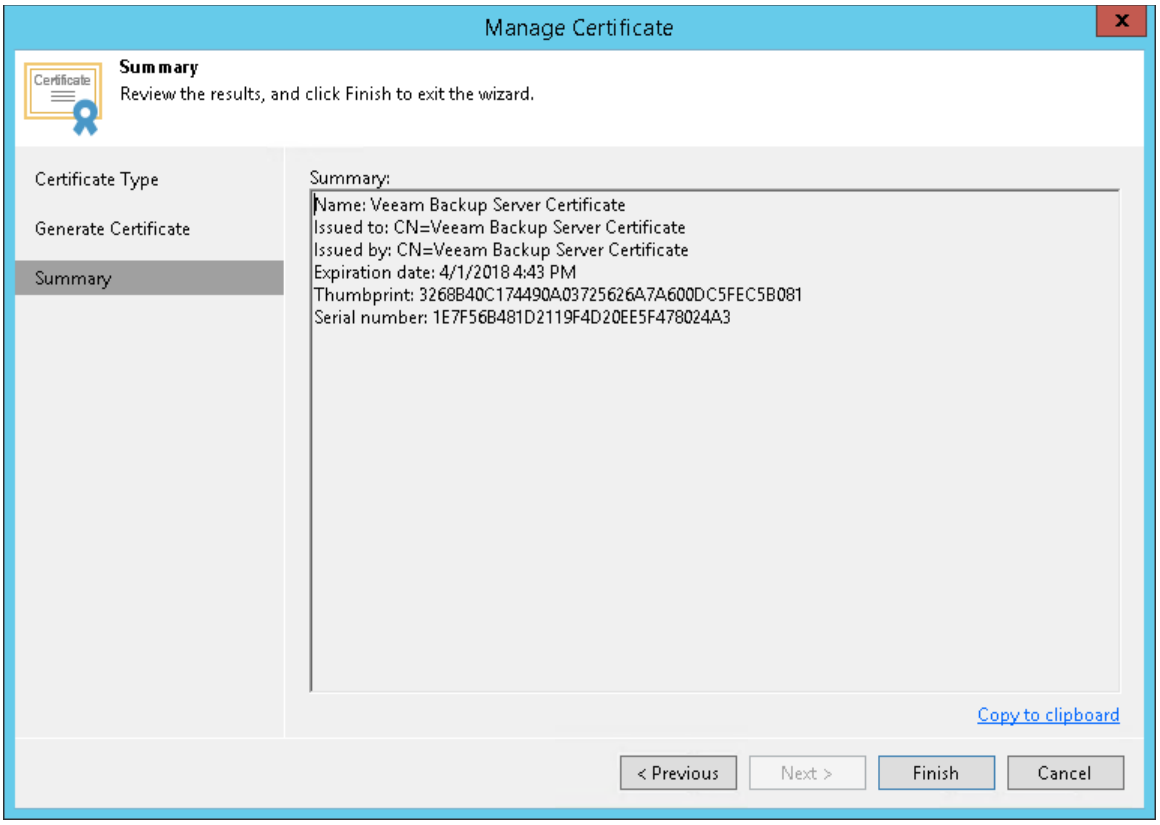

Importing Certificates from Certificate Store

If your organization has a TLS certificate signed by a CA and the TLS certificate is located in the Microsoft Windows Certificate store, you can use this certificate for authenticating parties in the Veeam Backup for Nutanix AHV infrastructure.

To select a certificate from the Microsoft Windows Certificate store:

1. From the main menu of the Veeam Backup & Replication console, select **General Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.
4. At the **Certificate Type** step of the wizard, choose **Select certificate from the Certificate Store**.

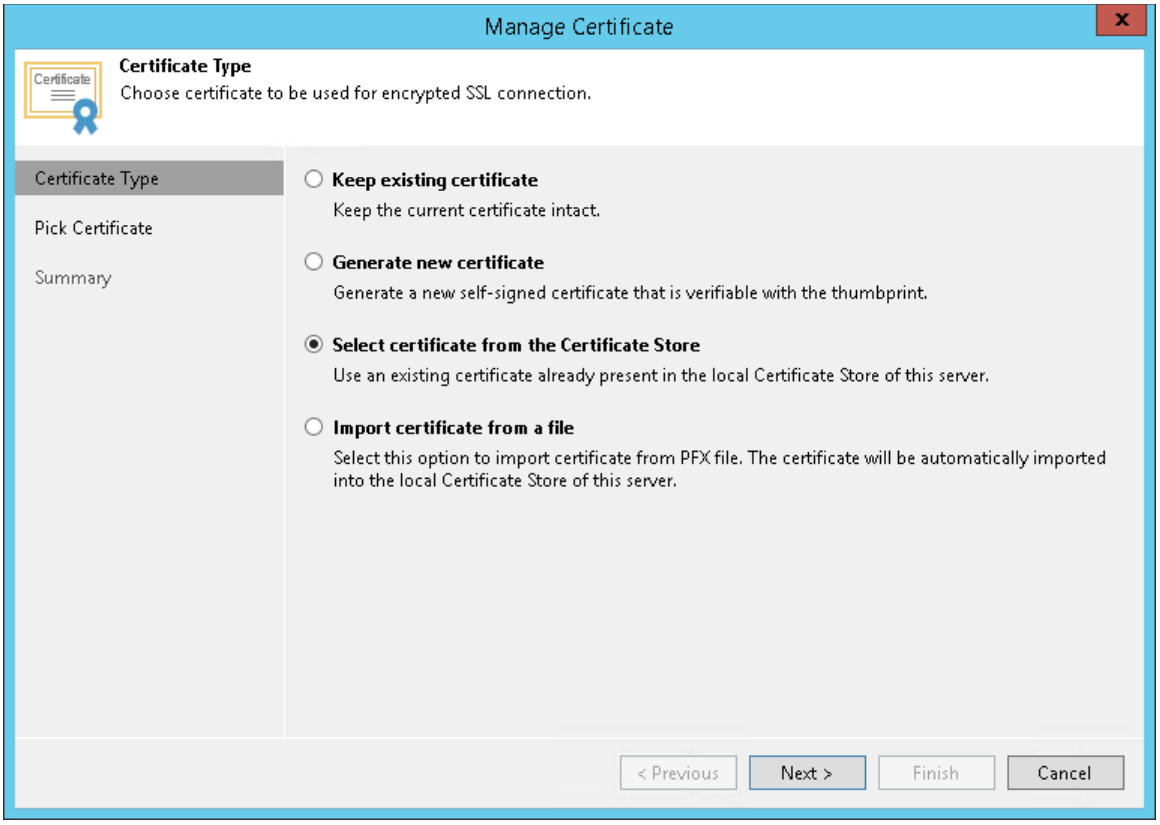

5. At the **Pick Certificate** step of the wizard, select a TLS certificate that you want to use. You can select only certificates that contain both a public key and a private key. Certificates without private keys are not displayed in the list.

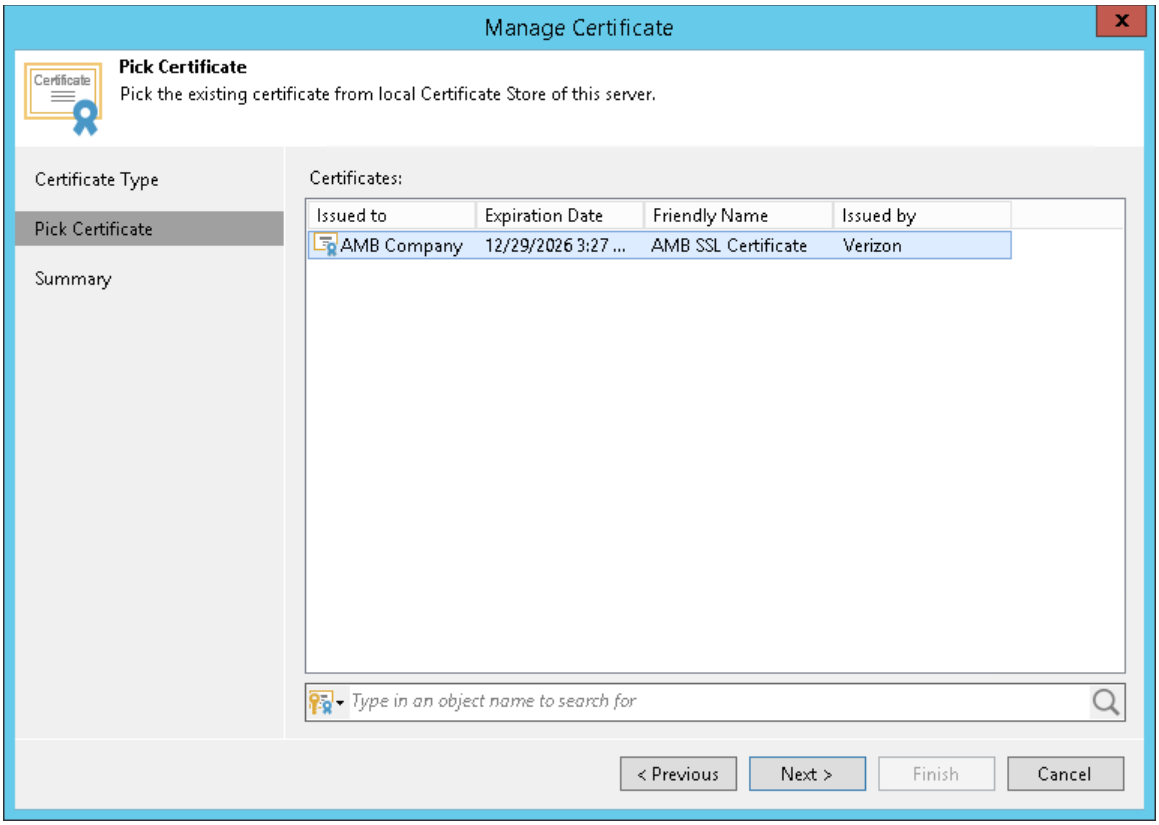

6. At the **Summary** step of the wizard, review the certificate properties.
7. Click **Finish** to apply the certificate.

Importing Certificates from PFX Files

You can import a TLS certificate in the following situations:

- Your organization uses a TLS certificate signed by a CA and you have a copy of this certificate in a file of PFX format.
- You have generated a self-signed TLS certificate in the PFX format with a third-party tool and you want to import it to Veeam Backup & Replication.

> **IMPORTANT**
>
> The TLS certificate must pass validation on the Veeam Backup & Replication server. Otherwise, you will not be able to import the TLS certificate.

To import a TLS certificate from a PFX file:

1. From the main menu of the Veeam Backup & Replication console, select **General Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.
4. At the **Certificate Type** step of the wizard, choose **Import certificate from a file**.

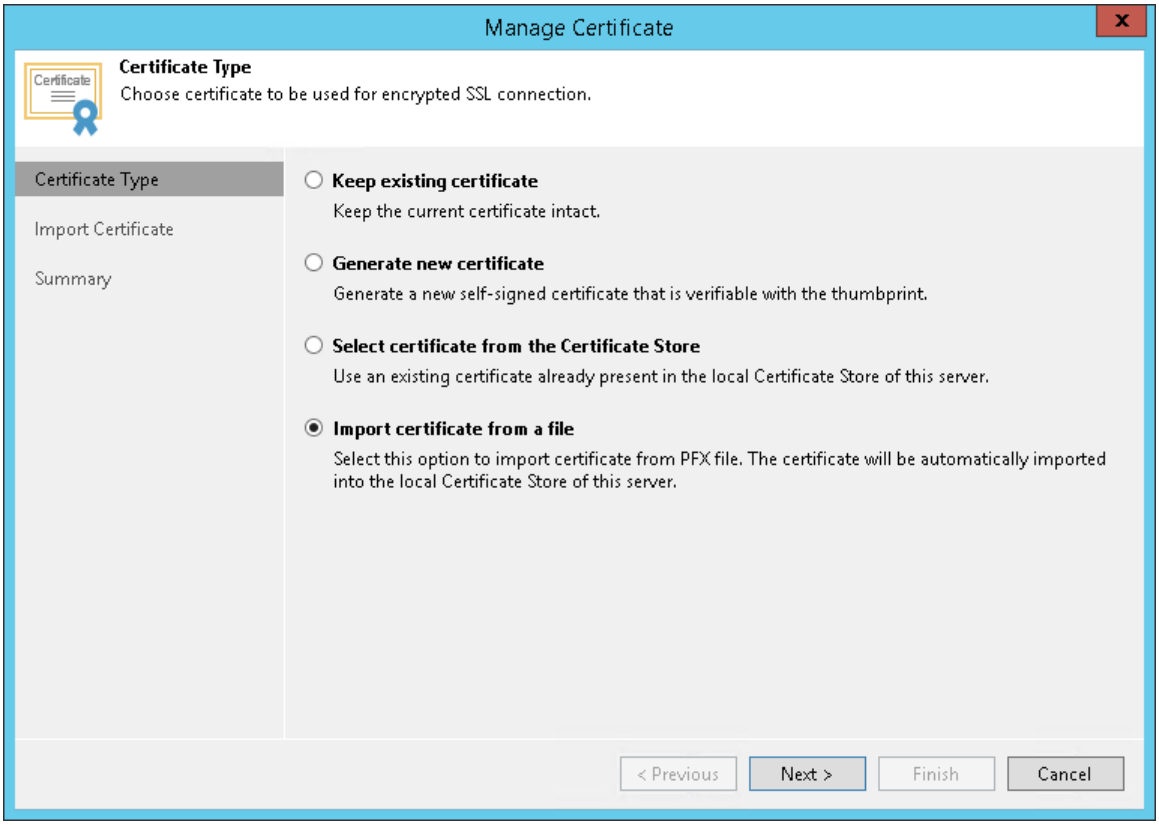

5. At the **Import Certificate** step of the wizard, specify a path to the PXF file.

6. If the PFX file is protected with a password, specify the password in the field below.

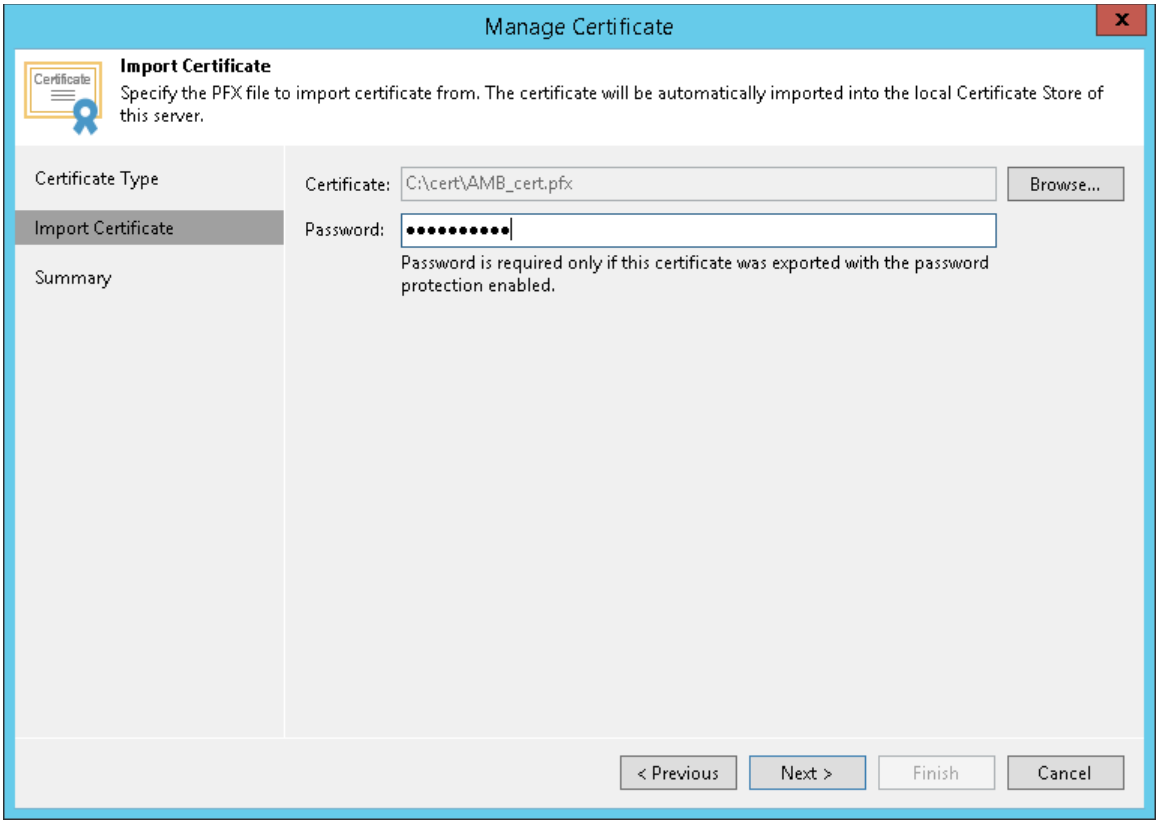

7. At the **Summary** step of the wizard, review the certificate properties. Use the **Copy to clipboard** link to copy and save information about the TLS certificate. You can use the copied information on a AHV Backup Proxy VM to verify the TLS certificate with the certificate thumbprint.
8. Click **Finish** to apply the certificate.

Using Certificate Signed by Internal CA

To establish a secure connection between the backup server and the AHV Backup Proxy VM, Veeam Backup & Replication uses a TLS certificate. By default, Veeam Backup & Replication uses a self-signed certificate. Veeam Backup & Replication generates this certificate when you install the product on the machine.

> **IMPORTANT**
>
> If you want to use a certificate signed by your internal Certification Authority (CA), make sure that the following requirements are met:
>
> - Veeam Backup & Replication server must trust the CA. That means that the Certification Authority certificate must be added to the Trusted Root Certification Authority store on the Veeam Backup & Replication server.
> - Certificate Revocation List (CRL) must be accessible from the Veeam Backup & Replication server.
> - When issuing the certificate, make sure the Subject Alternative Name field contains both the FQDN and the NetBIOS name. You can add multiple DNS entries in the following format: `DNS:vbrserver.domain.local,DNS:vbrserver`

A certificate signed by a CA must meet the following requirements:

1. The certificate subject must be equal to the fully qualified domain name of the Veeam Backup & Replication server. For example: *vbrserver.domain.local*.

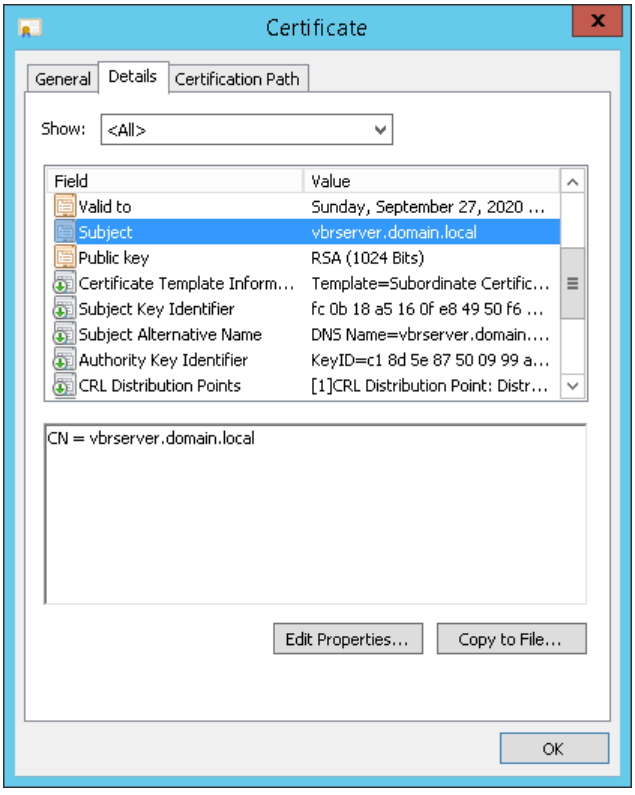

2. The following key usage extensions must be enabled in the certificate to sign and deploy child certificates for the AHV Backup Proxy VM:
   - Digital Signature
   - Certificate Signing
   - Off-line CRL Signing

- CRL Signing (86)

If you use Windows Server Certification Authority, we recommend you to issue a Veeam Backup & Replication certificate based on the built-in "Subordinate Certification Authority" template or templates similar to it.

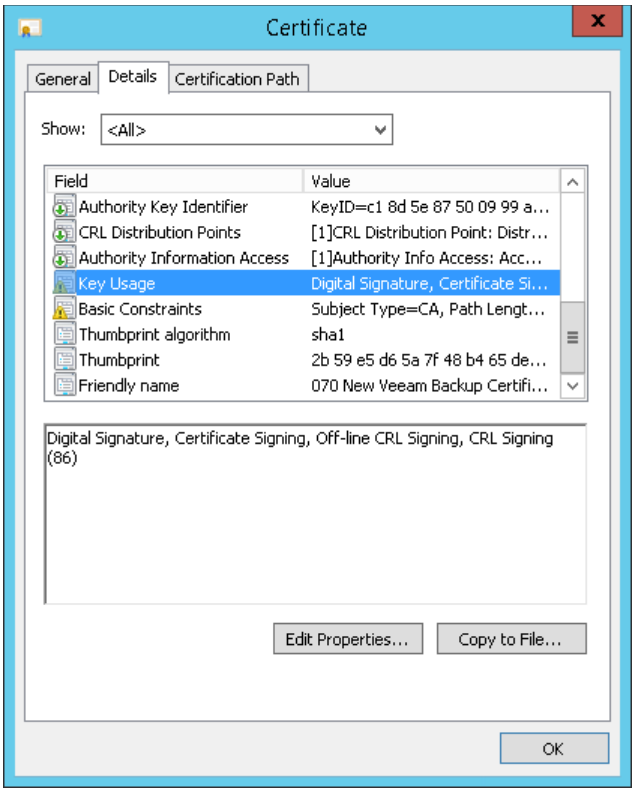

3. The key type in the certificate must be set to *Exchange*.

   If you create a certificate request using the Windows MMC console, to specify the key type, do the following:

   a. At the **Request Certificates** step of the **Certificate Enrollment** wizard, select a check box next to the necessary certificate template and click **Properties**.

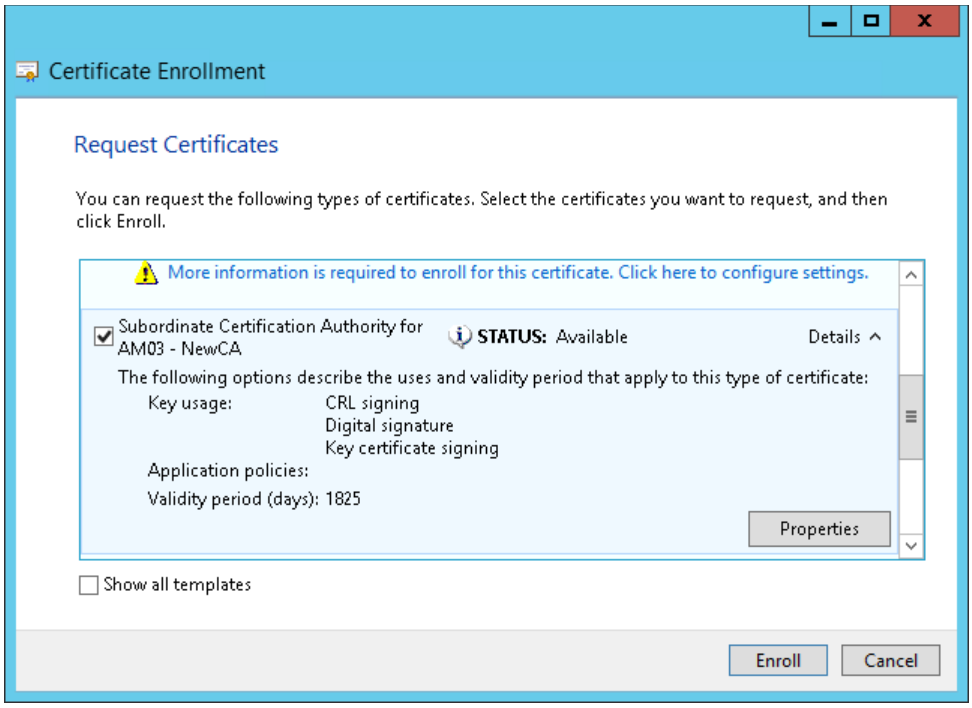

   b. In the **Certificate Properties** window, click the **Private Key** tab.

   c. In the **Key Type** section, select **Exchange**.

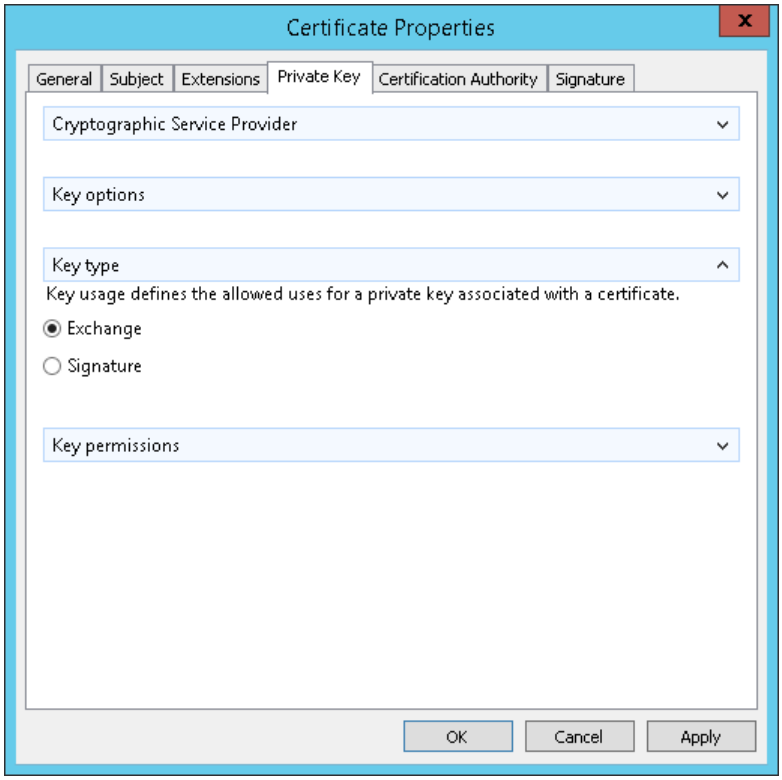

To start using the signed certificate, you must select it from the certificates store on the Veeam Backup & Replication server. To learn more, see Importing Certificates from Certificate Store.

Reconnecting to AHV Backup Proxy

After you specify the signed certificate in Veeam Backup & Replication, AHV Backup Proxy is not able to communicate with the Veeam Backup & Replication server and backup jobs fail. To reconnect the Veeam Backup & Replication server to AHV Backup Proxy, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** pane.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the AHV backup proxy and click **Edit Proxy** on the ribbon or right-click the AHV backup proxy and select **Properties**.
4. In the opened **Edit AHV Backup Proxy** wizard, go through the same steps as you have followed when adding an AHV backup proxy.

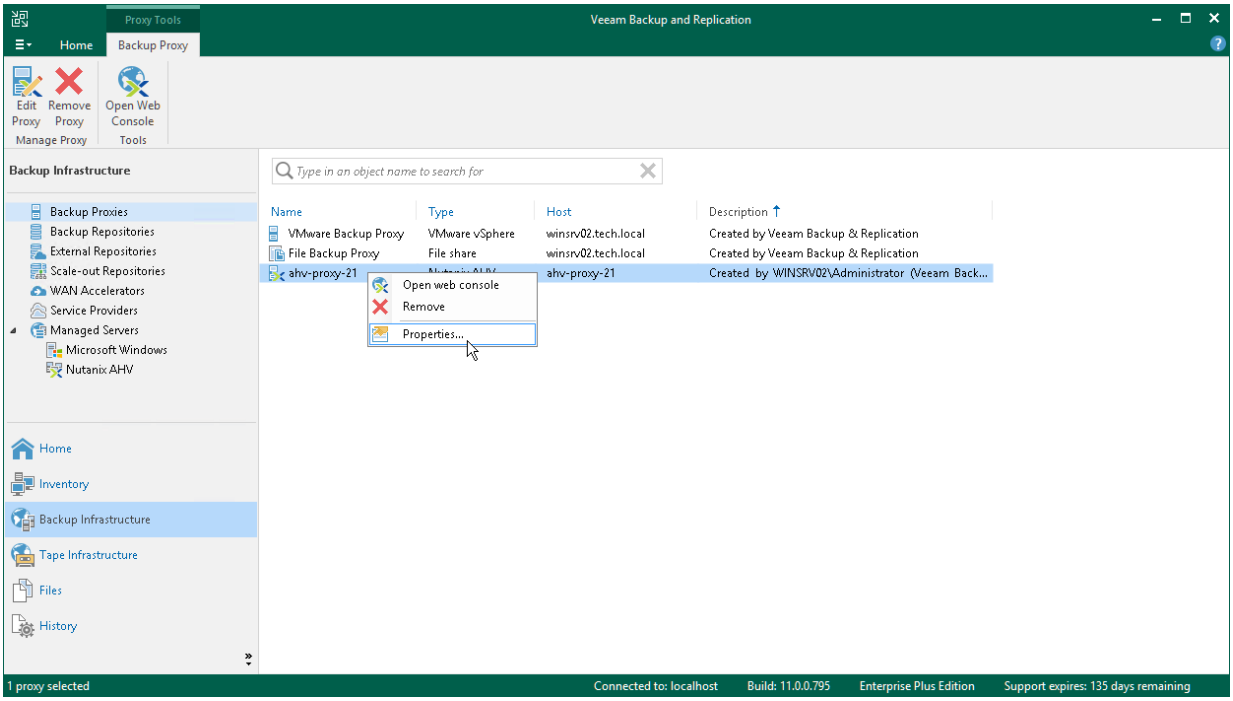

Web Console of AHV Backup Proxy

After you deploy AHV Backup Proxy, you can access the web console of AHV Backup Proxy. The web console allows you to manage backup and restore processes and configure settings and event notifications for AHV Backup Proxy.

In This Section

- Accessing Web Console
- Getting to Know UI of Web Console
- Monitoring AHV Backup Proxy Events

Accessing Web Console

To log in to the web console of AHV Backup Proxy, do the following.

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the AHV backup proxy and click the **Open Web Console** tab.

   Alternatively, you can open `https://<proxy_hostname>:443` in a web browser, where `<proxy_hostname>` is a hostname or IP address of the AHV Backup Proxy VM.

   > **NOTE**
   >
   > Port 8100 is also supported for backward compatibility.

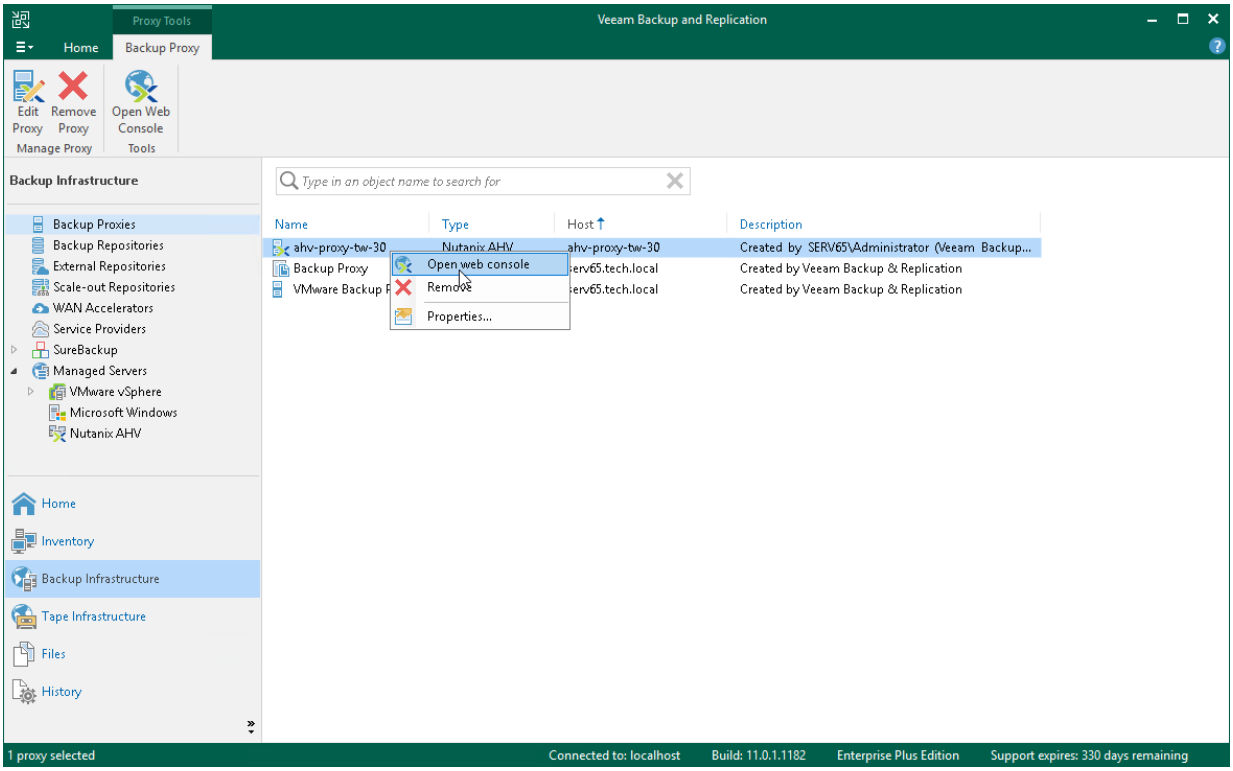

4. At the welcome screen, enter credentials that you have specified in the settings of the AHV backup proxy.

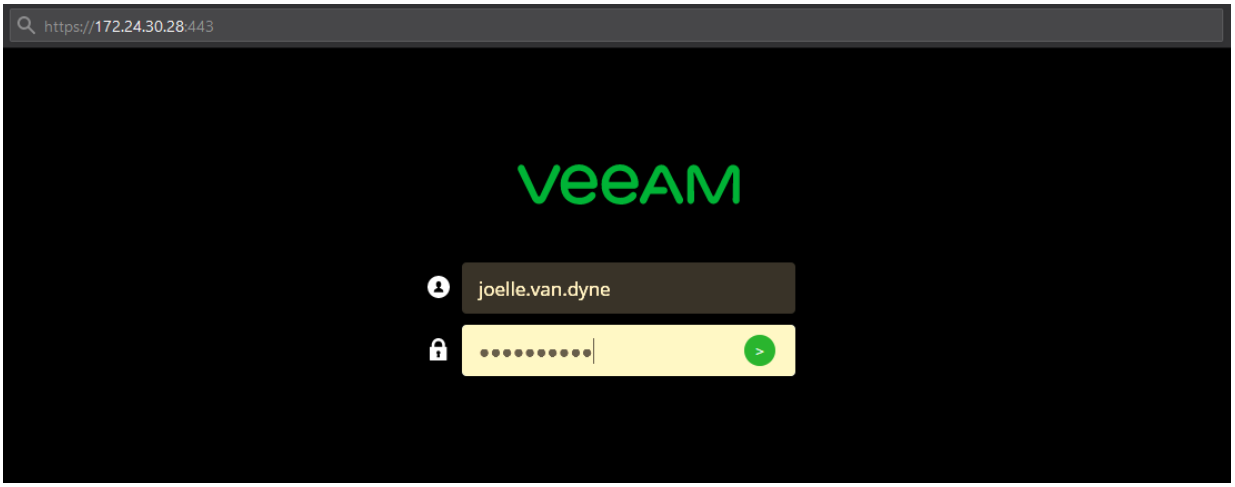

Getting to Know UI of Web Console

After you log in to the AHV Backup Proxy web console, the **Dashboard** tab opens. At the main panel of the web console, you can navigate through tabs to manage backup and restore operations and configure settings of AHV Backup Proxy.

In the AHV Backup Proxy web console, you can perform the following:

- View ongoing statistics for the backup infrastructure components and view the list of event logs using the **Dashboard** tab.
- Create and manage backup and snapshot jobs using the **Jobs** tab.
- Manage backups/snapshots and perform restore from existing backups using the **Protected VMs** tab.
- View the list of internal events including changes in the infrastructure of AHV Backup Proxy, started/stopped jobs, errors, warnings. For more information, see Monitoring Events.
- Configure notifications and AHV Backup Proxy maintenance settings. For more information, see General Settings.
- Check and install updates for Veeam Backup for Nutanix AHV. For more information, see Updating AHV Backup Proxy.

> **TIP**
>
> On every tab and on every wizard window, there is the **?** icon. Click it if you want to open online help for this particular tab or wizard window.

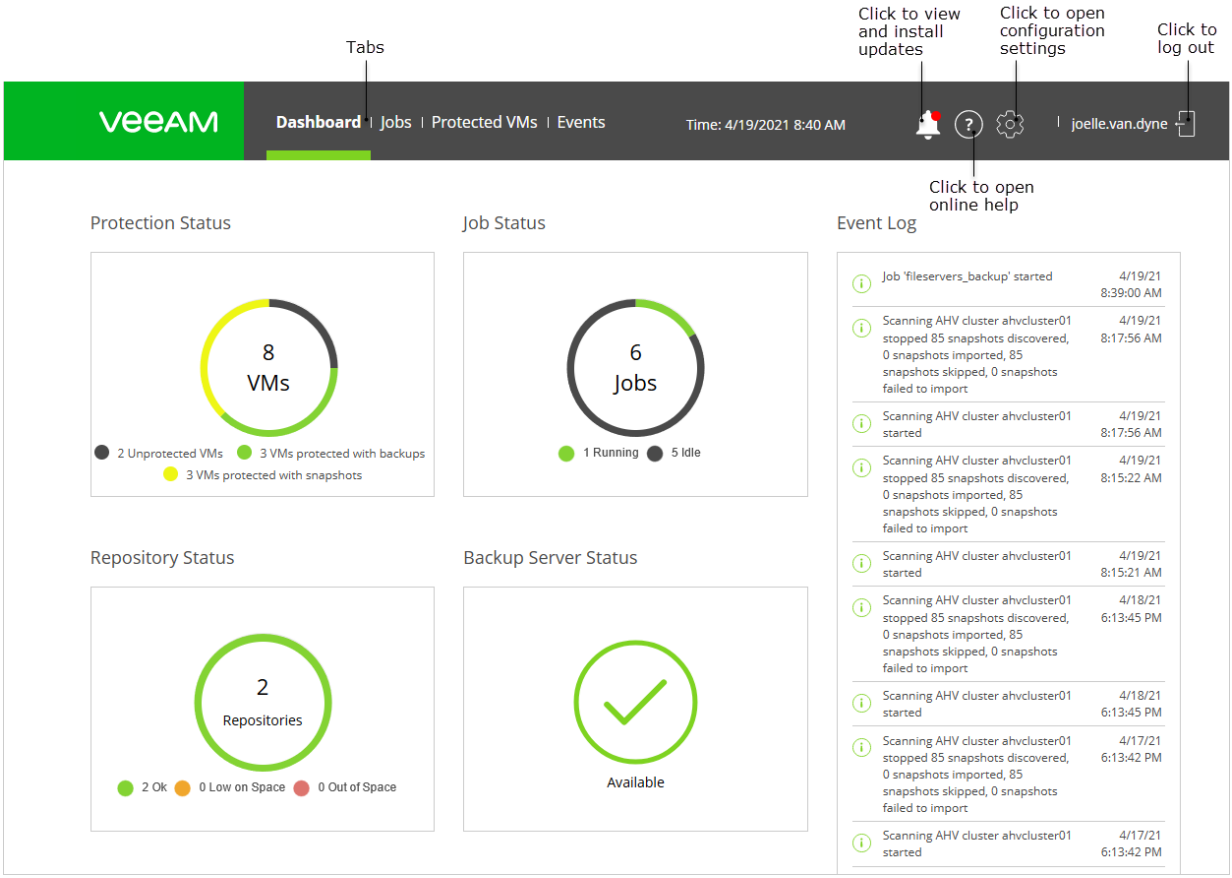

Monitoring Events

At the **Events** tab of the main panel, you can view the list of internal events including changes in the infrastructure of AHV Backup Proxy, started/stopped jobs, errors, warnings. You can use the search field to find a specific event and use the **Resolved** and **Type** filters to display events according to status or event type. Use the following tips for working with events.

- To mark an event as resolved, select the check boxes next to required events and click the **Resolve** button. Note that you can resolve only the ***Error*** or ***Warning*** type events.

  The time of resolution will appear at the **Resolved** column of the table.

- To see session details of a backup or restore operation, click the date of the event.

- Use the **Type** filter to view only certain types of events (*Information*, *Warning*, *Error*).

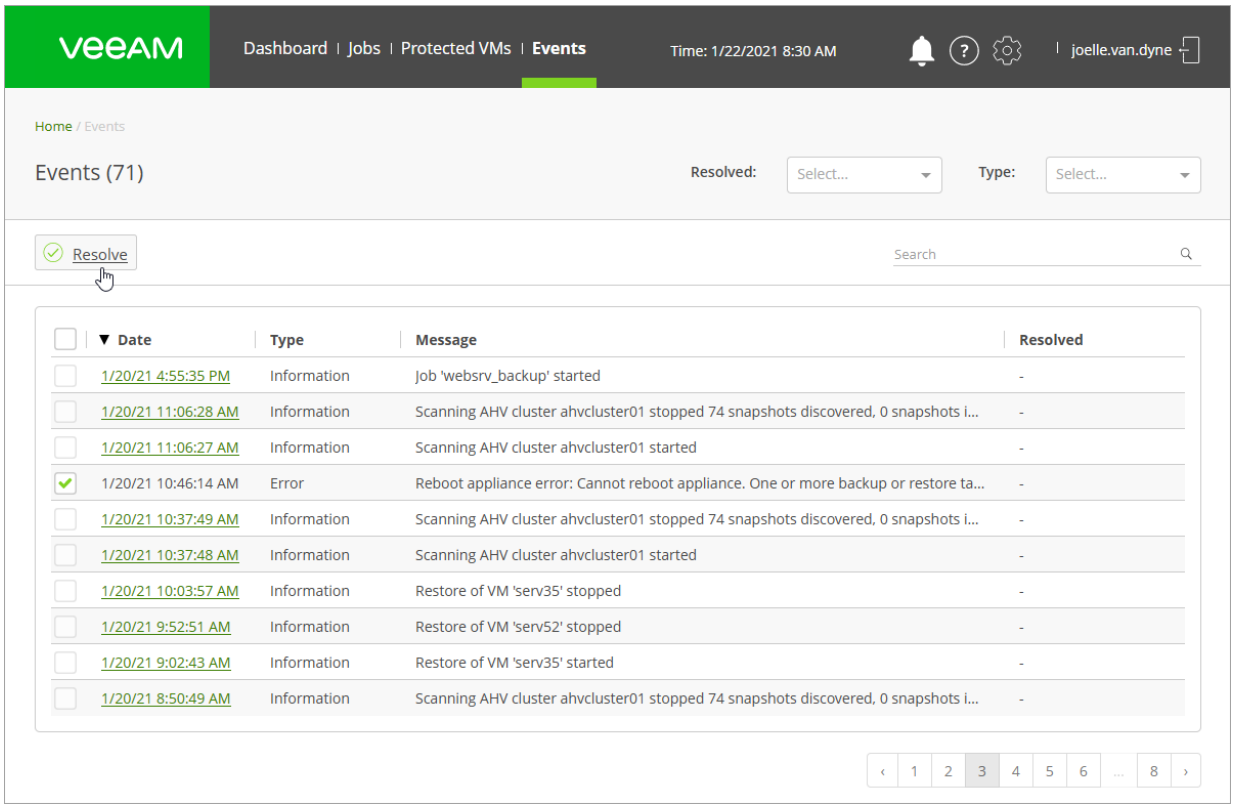

General Settings

In the **Appliance Settings** view of the AHV Backup Proxy web console, you can perform the following tasks:

- Edit network settings.
- Edit the Administrator account settings.
- Select the time zone and configure time synchronization settings.
- Configure e-mail settings.
- Configure notifications.
- Back up configuration settings of AHV Backup Proxy.
- Restore configuration settings of AHV Backup Proxy.
- Reboot or shut down the AHV Backup Proxy VM.

You can also restore the AHV Backup Proxy VM from PD snapshots which are created automatically after adding the proxy. For more information, see Protecting AHV Backup Proxy VM.

For instructions on how to export logs of AHV Backup Proxy, see Logs and Support.

Configuring Network Settings

At the **Network** tab of the **Appliance Settings** view, you can change network settings of the AHV Backup Proxy VM.

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.
2. At the **Network** tab, configure network settings and click **Apply**.

> **NOTE**
>
> If you specify new DNS server, make sure the DNS server resolves the domain name of the AHV Backup Proxy server. If the domain name is not resolved, you must register the existing Nutanix backup proxy in the Veeam Backup & Replication console. For more information, see Add Nutanix Proxy.

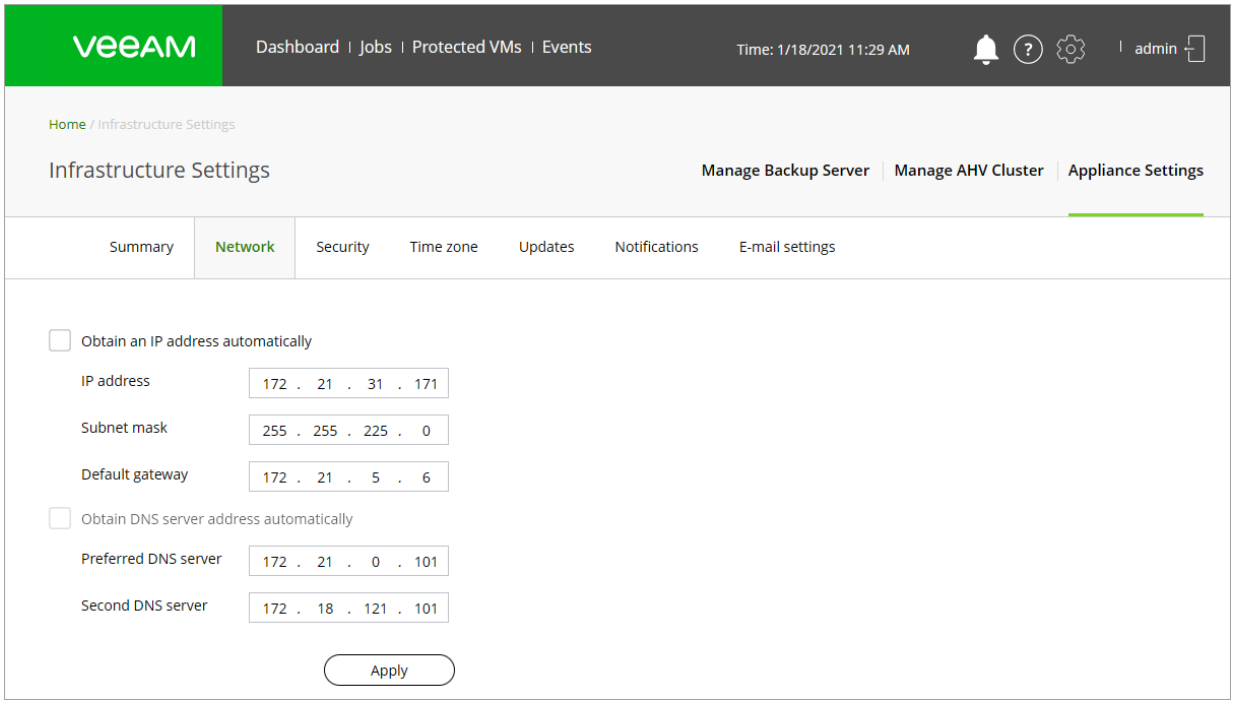

Configuring Administrator Account Settings

At the **Security** tab of the **Appliance Settings** view, you can change the login and password of the administrator account that can access the AHV Backup Proxy web console.

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.
2. In the **Appliance Settings** view, open the **Security** tab.
3. In the **Login** field, enter the account name.
4. In the **Old password** field, enter the current password for the administrator account.
5. In the **New password** field, assign a new password for the administrator account.
6. In the **Confirm new password** field, enter the new password again.
7. Click **Apply** to save the settings.

IMPORTANT

If you change the administrator account credentials, you must also update the credentials specified in the AHV Backup Proxy properties in the Veeam Backup & Replication console. For more information, see Editing and Removing Nutanix AHV Backup Proxies.

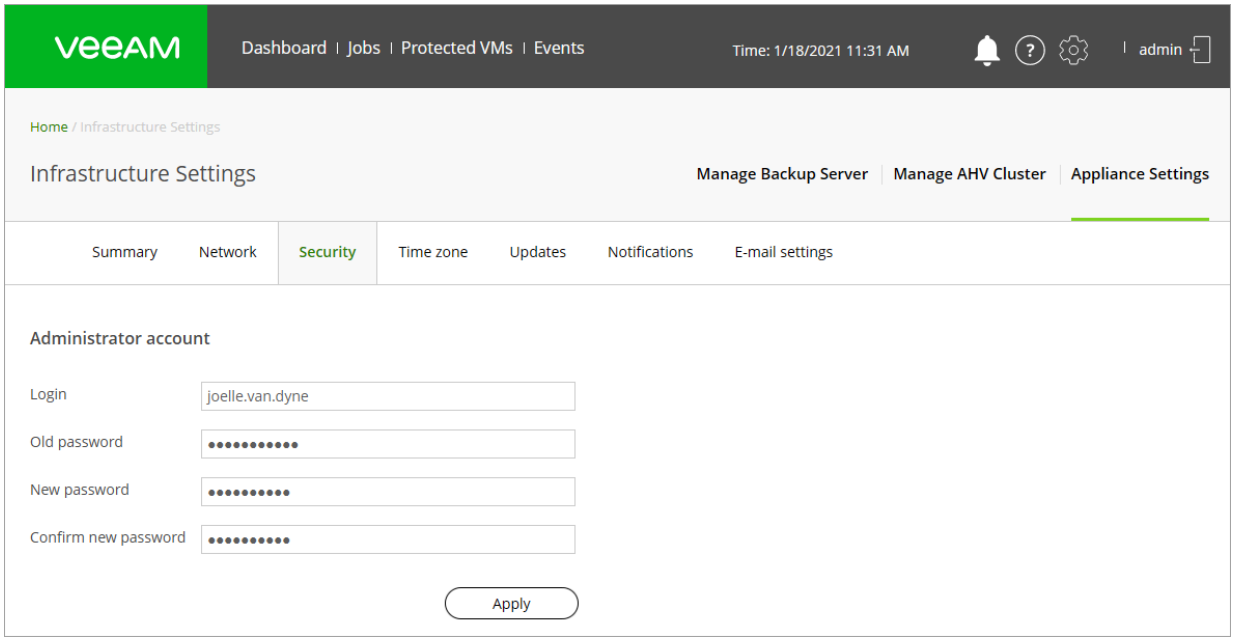

Configuring Time Zone Settings

Veeam Backup for Nutanix AHV performs all data protection and disaster recovery operations according to the time zone set on the AHV Backup Proxy.

Since the AHV Backup Proxy is deployed on VM in Nutanix AHV cluster, the time zone is set to Coordinated Universal Time (UTC) by default. However, you can change the time zone if required. For example, you may want the time on the backup appliance to match the time on the workstation from which you access AHV Backup Proxy.

> **NOTE**
>
> If you change the time zone, AHV Backup Proxy does not recalculate the starting time of scheduled jobs.

To edit time zone settings, perform the following steps:

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.
2. At the **Time zone** tab, click **Settings**.

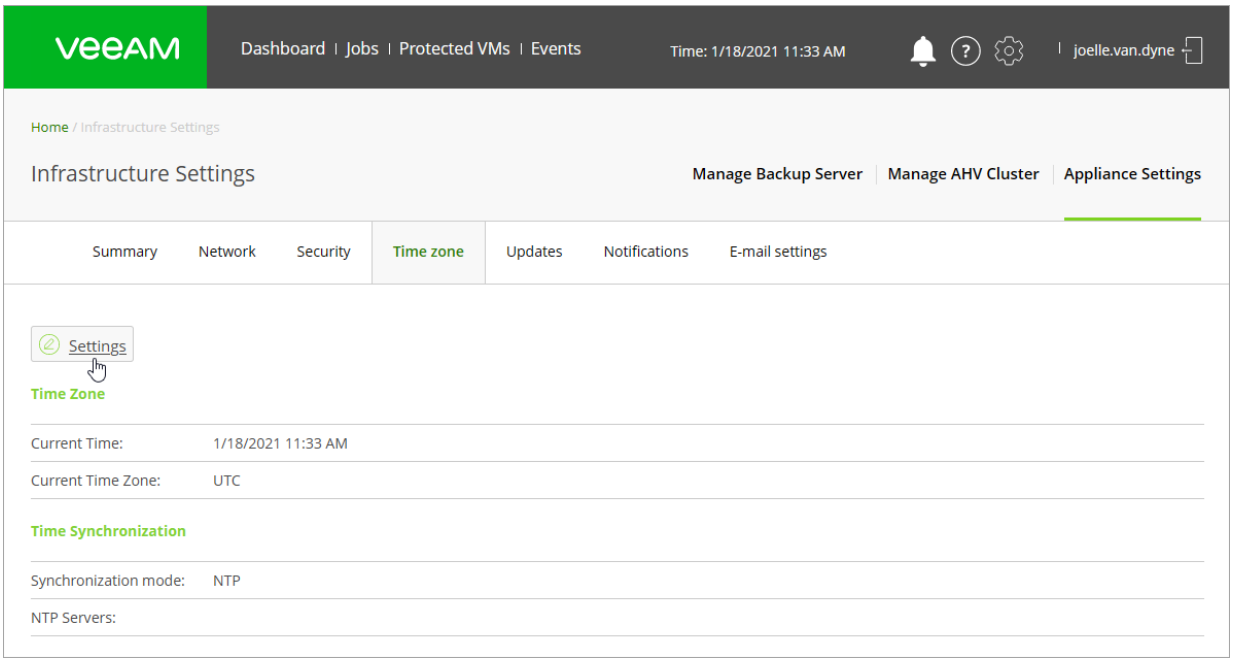

3. In the **Select the time zone** list, choose the time zone of your region.
4. Specify the synchronization mode:
   - **NTP: Time of the proxy appliance will be synchronized with the specified NTP servers.**
   - **Host: Time of the proxy appliance will be synchronized with the time on guest OS of the host where the proxy appliance is installed.**
5. [For NTP mode] Specify required NTP servers in the **NTP Servers** text field.

6. Click **OK** to save the changes.

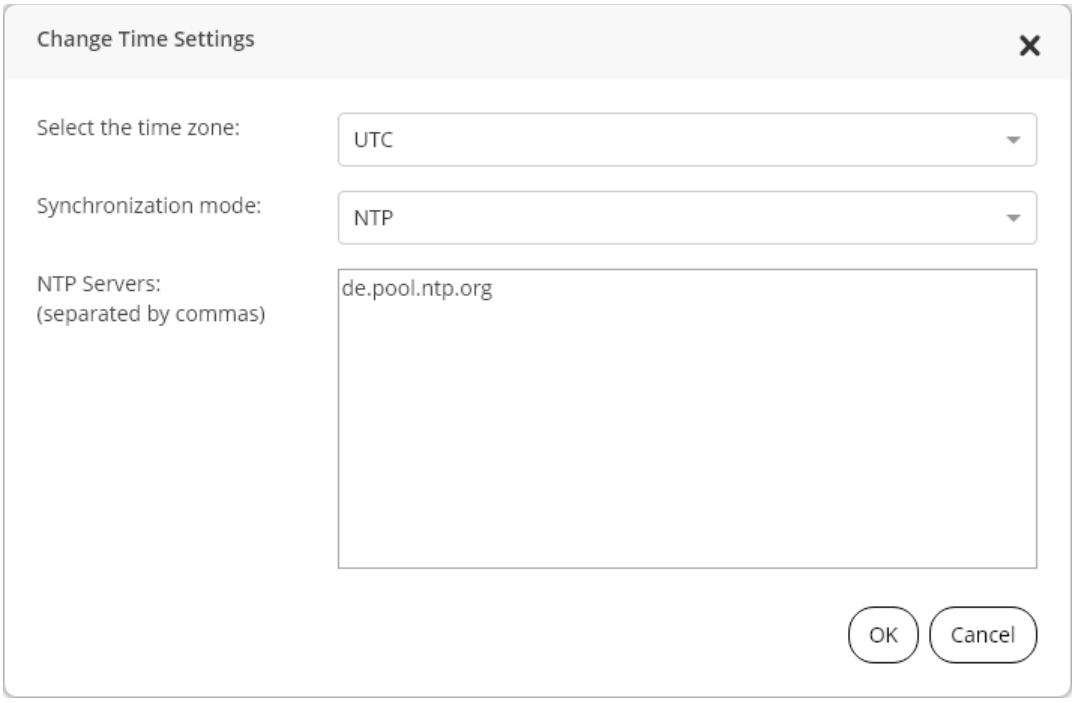

Updating AHV Backup Proxy

Since version 2.1, AHV Backup Proxy has an additional component that is responsible for updates: **Veeam Updater**. You can install and schedule the following packages:

- **AHV Backup Proxy updates**: updates that fix bugs and performance issues in AHV Backup Proxy.
- **Security updates**: updates that fix security issues on the Ubuntu distributive on the kernel of which AHV Backup Proxy is built.

> **NOTE**
>
> Before you install updates for AHV Backup Proxy, make sure all required ports are open. For more information, see Used Ports.

For instructions on how to install and manage updates, see:

- Checking for Updates
- Installing Updates

Checking for Updates

By default, AHV Backup Proxy automatically checks for available minor updates. You will see a red dot on the bell icon at the top-right of the AHV Backup Proxy web console.

To check for available minor updates and security updates, do the following:

1. At the top-right of the main menu, click the gear icon and select **Appliance Settings**.
2. Open the **Updates** tab and click the **Check and view updates** link.

   After you click the link, you will be redirected to the Veeam Proxy AHV Updater portal where you can install or schedule the updates. For more information, see Installing Updates.

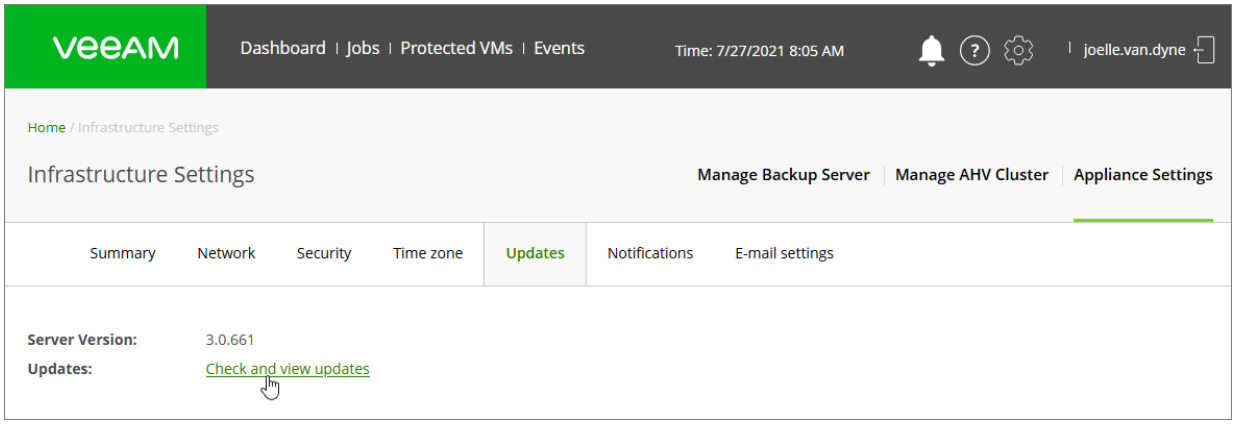

Installing Updates

After you check for updates in the AHV Backup Proxy web console, you will be redirected to the Veeam Proxy AHV Updater portal.

At the Veeam Proxy AHV Updater portal, you can install the AHV Backup Proxy updates and Ubuntu security updates. Also, you have the option to select a specific date and time when the updates must be installed automatically.

1. Open the Veeam Proxy AHV Updater portal:
   a. On the main menu of the AHV Backup Proxy console, click the gear icon and select **Appliance Settings**.
   b. Open the **Updates** tab and click the **Check and view updates** link.
2. At the Veeam Proxy AHV Updater portal, in the **Updates are available for this system** section, select packages that you want to install.
3. In the **Choose action** section, you can select whether you want to install updates now or schedule the installation at the specified time and date.
   - **Install updates now**: select this option if you want to install the updates immediately.

     Note that if you start the update while backup or restore processes are running, these backup or restore sessions will be interrupted and finished with an error.
   - **Schedule updates installation**: select this option if you want to install the updates at the specified date and time.

     For this option, you must specify the date and time when updates must be installed. Make sure backup jobs are not scheduled to run on the specified date at the specified time. Otherwise, the update process will interrupt running activities, which may result in data loss.
   - **Remind me later**: If you have not decided when to install updates, you can set an update reminder for the next day or next Monday.
4. After some updates Veeam Backup for Nutanix AHV has to be rebooted. You can select the **Reboot automatically after install if required** check box if you want to allow the automatic reboot after the update installation. Otherwise, you must reboot Veeam Backup for Nutanix AHV manually. For more information, see Rebooting or Shutting Down Proxy VM.

5. Depending on the option you have chosen, click the **Install Updates**, **Schedule Updates** or **Remind me Later** button.

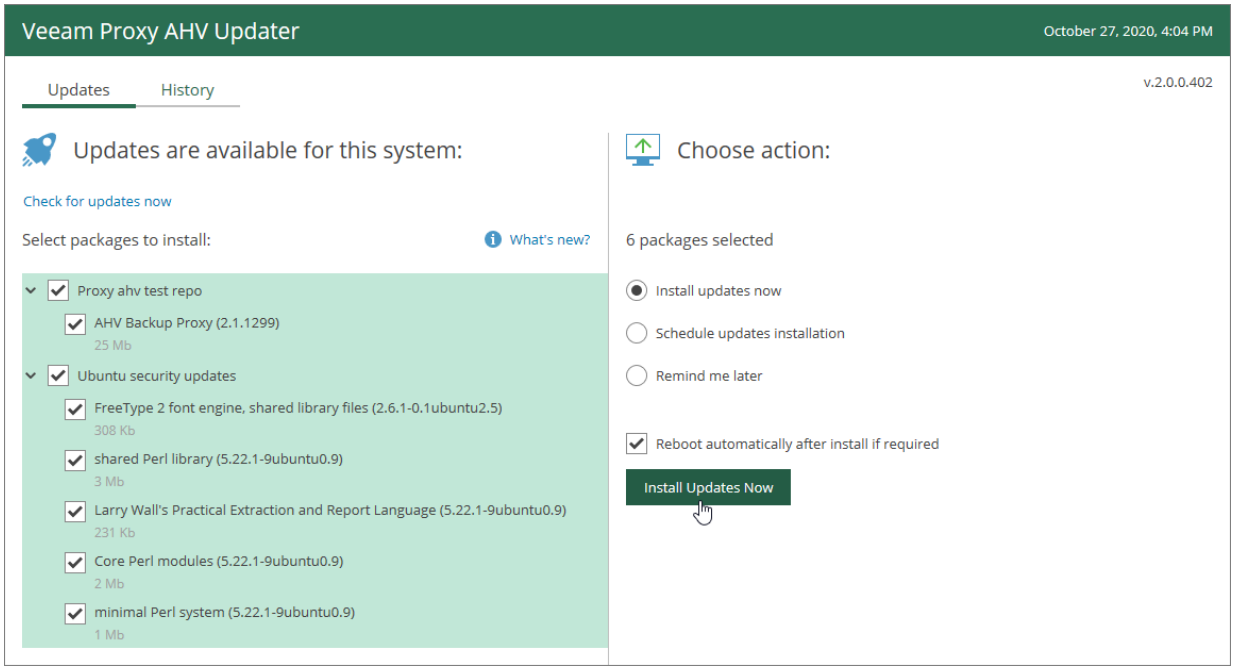

Configuring E-mail Settings

You can configure e-mail settings, to receive e-mail notifications. AHV Backup Proxy sends the following notifications:

- Reports with results of backup, snapshot, and PD snapshot jobs
- Notification that there is a new version of AHV Backup Proxy is available
- Notifications that you need to reboot AHV Backup Proxy due to a new security update has been installed
- Reports with results of a configuration database backup

If you want to receive notifications, configure notification settings:

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.
2. Open the **E-mail settings** tab and select the **Enable e-mail notifications** check box.
3. In the **SMTP server** field, enter a full DNS name or IP address of the SMTP server that will be used for sending email notifications.
4. Click the **Advanced** button to specify advanced SMTP settings.
   a. Specify the port number for the SMTP server.
   b. To use a secure connection for email operations, select the **Connect using SSL** check box.
   c. If you need to connect to the SMTP server using a specific account, select the **This SMTP server requires authentication** check box and specify credentials that will be used to connect to the SMTP server.
   d. Click **OK**.

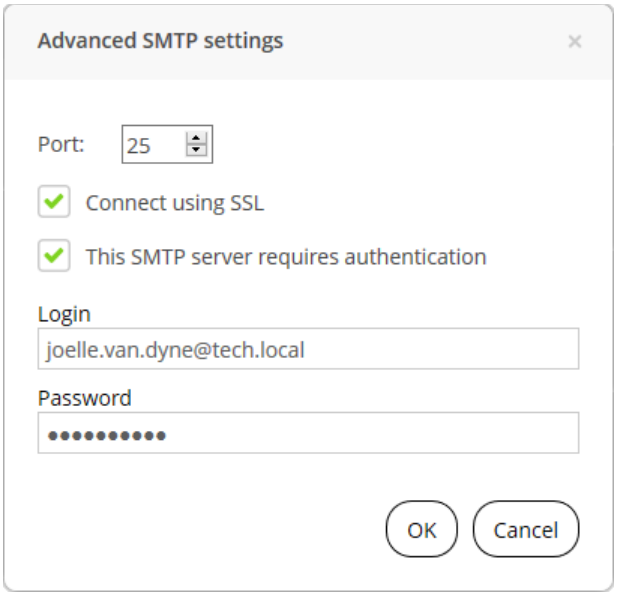

5. In the **From** field, specify an email address from which email notifications must be sent.
7. In the **To** field, specify the recipient addresses. Use a semicolon to separate multiple addresses. Recipients specified in this field will receive notification about every job managed by the backup server.

8. In the **Subject** field, specify a subject for the sent message. You can use the following variables in the subject:
   a. *%JobName%* — name of the job
   b. *%JobResult%* — status with which the job has finished
   c. *%ObjectCount%* — number of objects in the job
9. Select the **Notify on success**, **Notify on warning** and/or **Notify on failure** check boxes to receive email notification if a job is run successfully, not successfully or with a warning.
10. 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, AHV Backup Proxy will send one notification per every job retry.
11. AHV Backup Proxy allows sending a test email to check if all settings have been configured correctly. To send a test email, click **Test Message**.

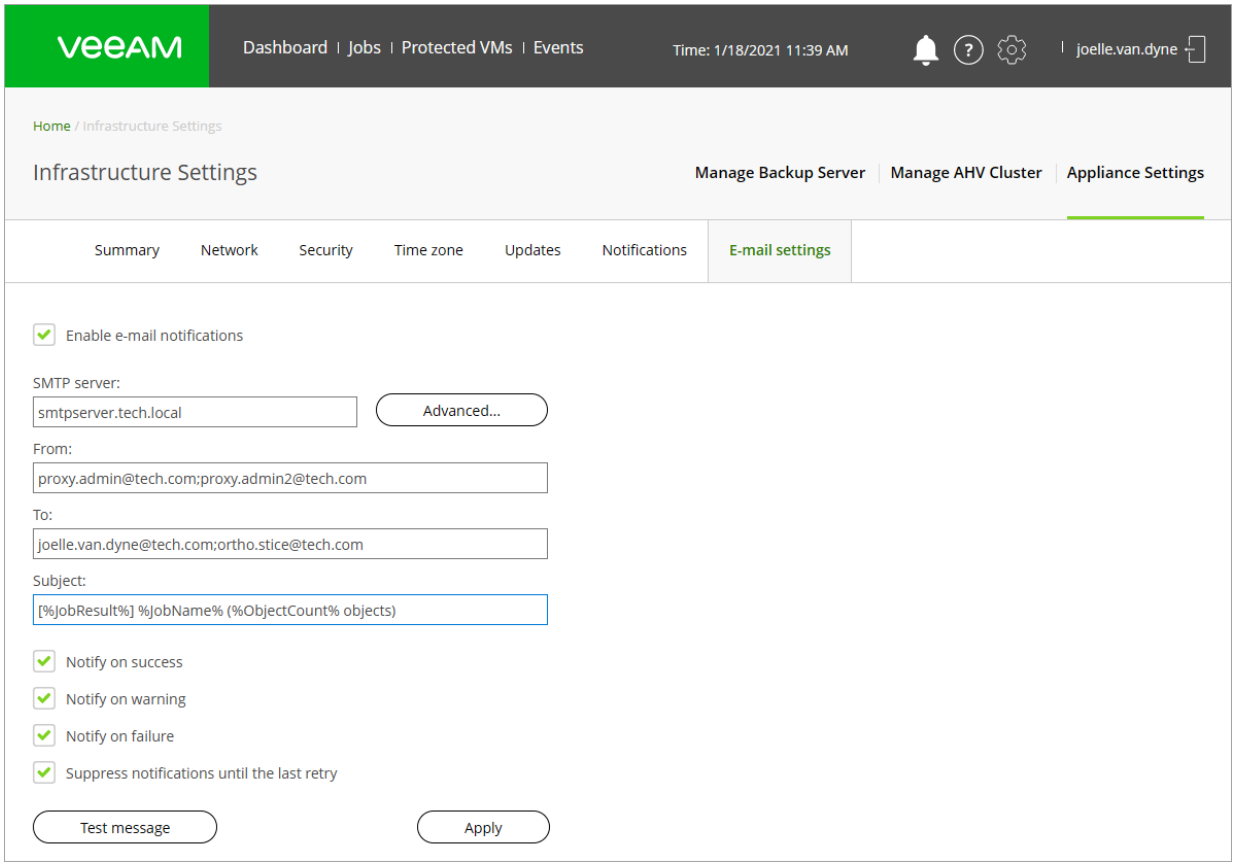

Related Topics

Configuring Notifications

Configuring Notifications

If you want to receive e-mail notifications on available product updates and low storage space, you can enable notifications in the **Notifications** tab. Note that to receive e-mail notifications, you must configure E-mail settings beforehand.

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.
2. At the **Notifications** tab of the **Appliance Settings** view, select the notifications you want to receive.

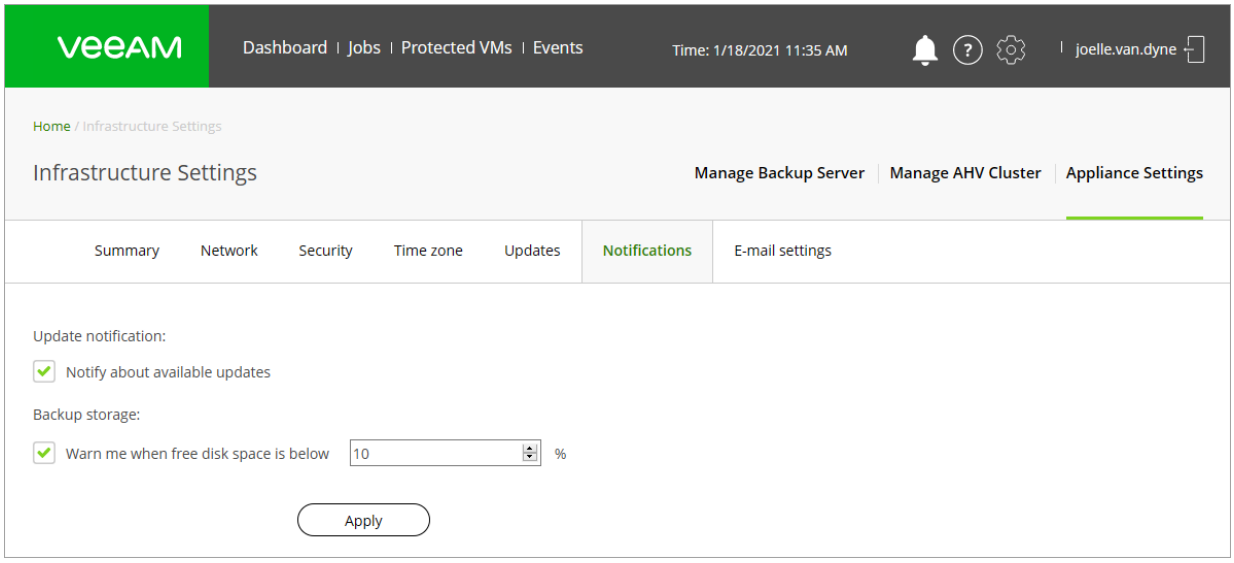

Related Topics

Configuring E-mail Settings

Backing Up Configuration Settings

You can back up and then restore the configuration database of AHV Backup Proxy. If the AHV Backup Proxy fails for some reason, you can re-deploy the AHV Backup Proxy and quickly restore its configuration from the configuration backup file.

When you perform a configuration backup, AHV Backup Proxy retrieves data from the configuration database, writes this data into a set of JSON and DAT files and archives these files to a ZIP file. The default naming scheme for the configuration backup file is `config_AhvProxyName-DD_MM_YYYY_HH_MM_SS.zip`.

AHV Backup Proxy configuration backup stores information about the following objects:

- Backups and snapshots created by AHV Backup Proxy
- Backup and snapshot jobs created by AHV Backup Proxy
- Events: performed backup and restore sessions, configuration changes
- AHV Backup Proxy appliance settings

To back up the configuration database of AHV Backup Proxy, do the following.

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.

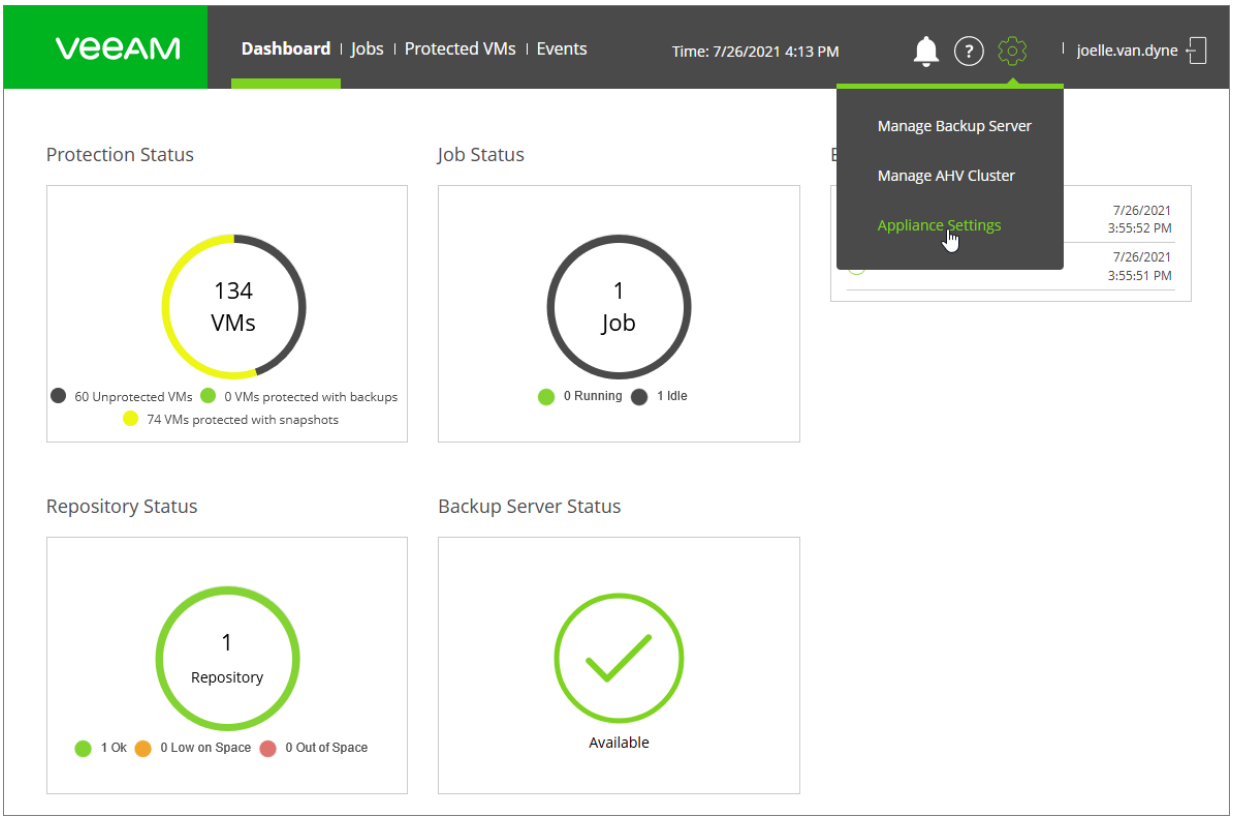

2. At the **Summary** tab of the **Appliance Settings** section, click **Configuration Backup**.

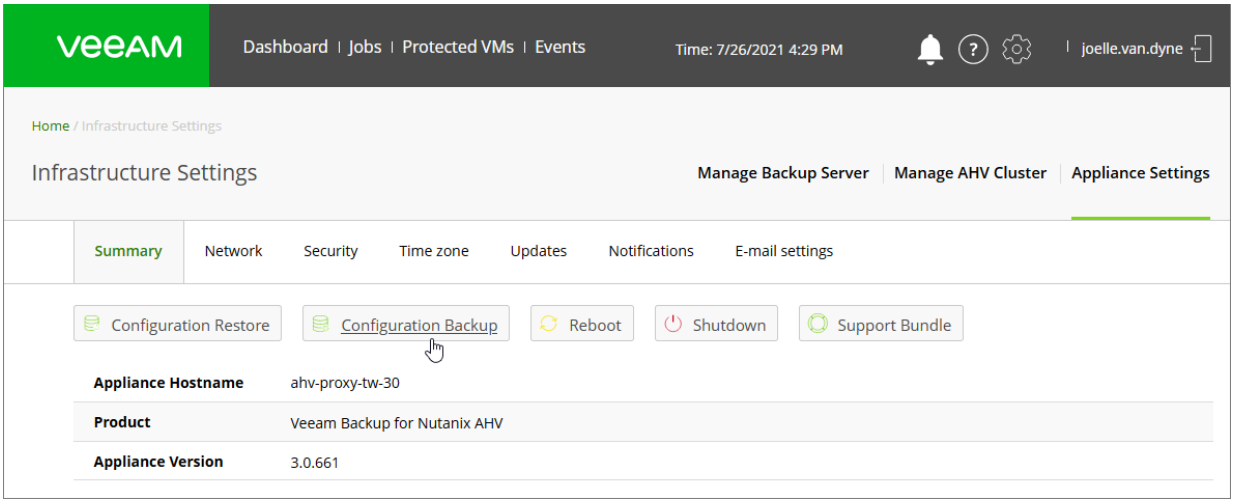

3. In the **Configuration Backup** window, select items that you want to back up:
   - **Common**: configuration settings, added Nutanix AHV clusters, Veeam Backup & Replication servers, backup jobs, and so on.
   - **Events**: list of events and event details.
4. In the **Protect passwords in exported files** section, specify a password which will be used to decrypt saved file.
5. Click **Download**. A configuration backup file will be downloaded to your machine.

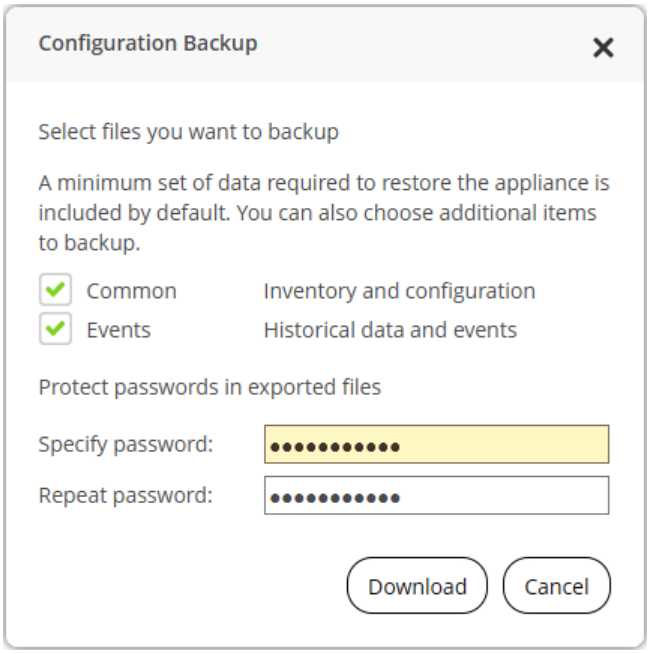

Related Topic

Restoring Configuration Settings

Restoring Configuration Settings

Since release of Veeam Backup for Nutanix AHV version 3.0, you can restore configuration settings of a AHV Backup Proxy instance. To restore configuration settings of AHV Backup Proxy, perform the following steps:

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.

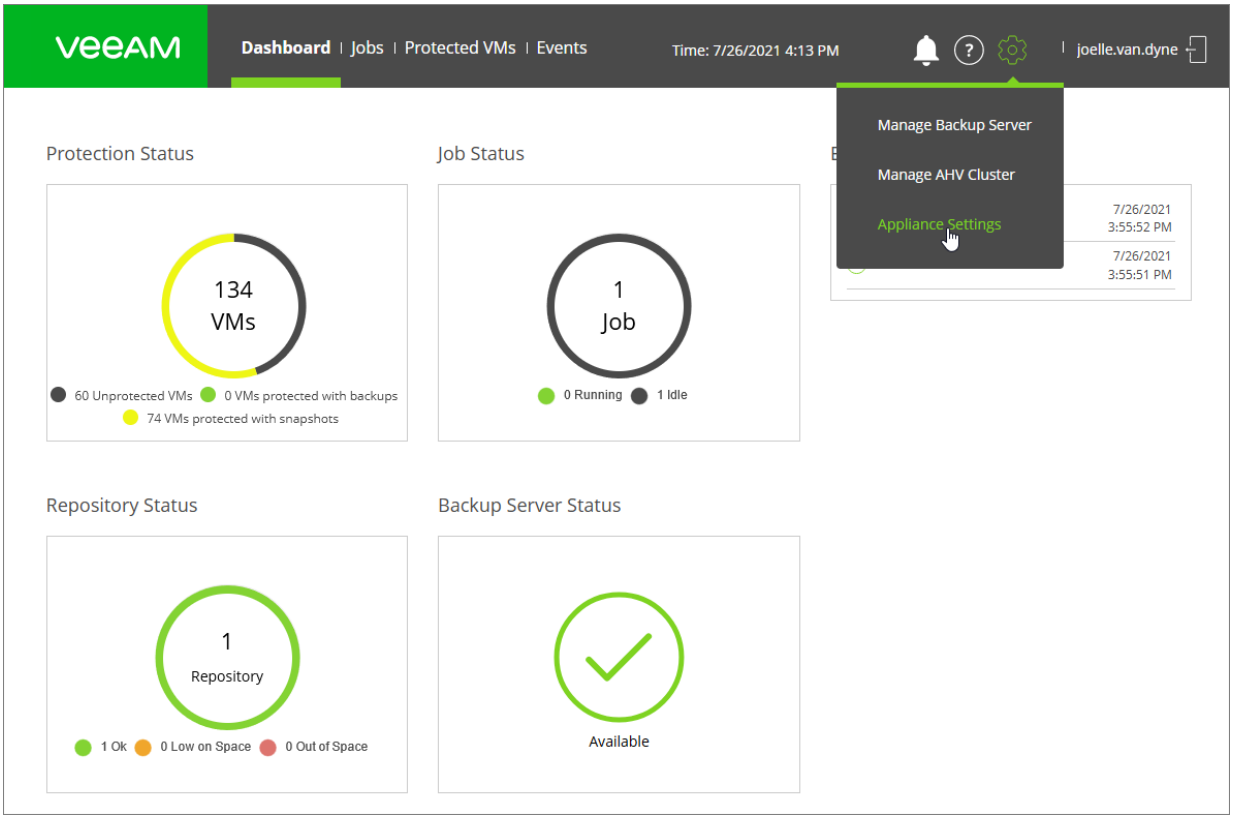

2. At the **Summary** tab of the **Appliance Settings** section, click **Configuration Restore** to launch the **Configuration Restore** wizard.

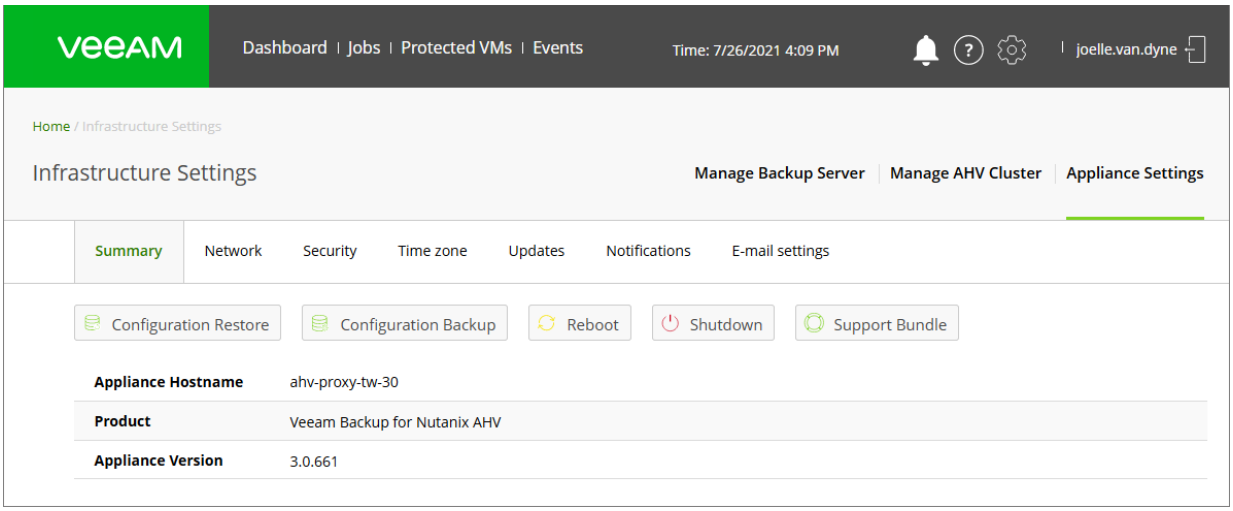

3. At the **Configuration File** step of the wizard, click **Browse** and specify the configuration backup file.

   Before uploading the configuration backup file, the wizard will request you to enter the password that you have specified when backing up the configuration.

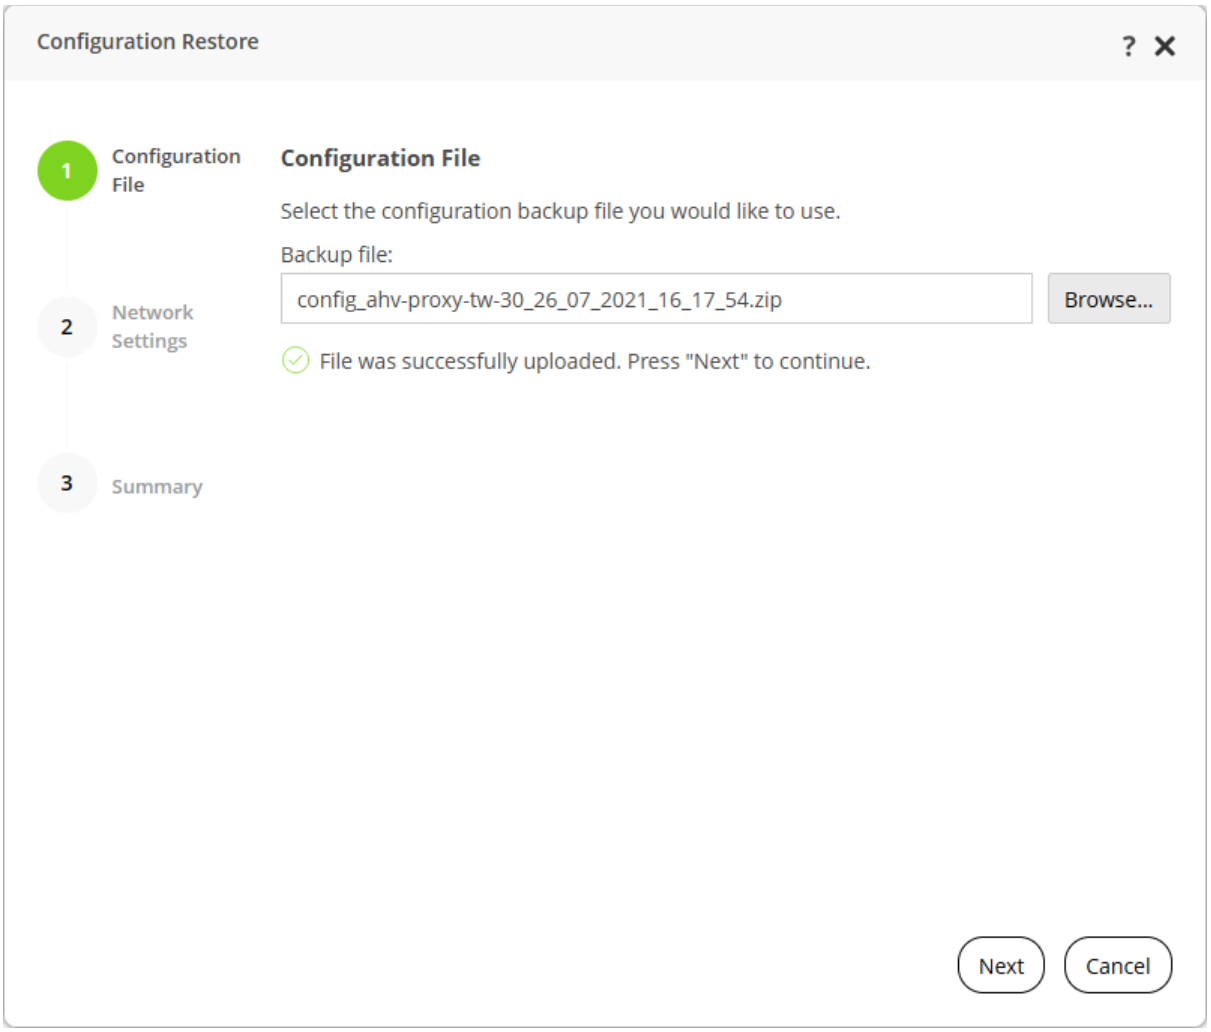

4. At the **Network Settings** step of the wizard, specify network settings for AHV Backup Proxy. You can use the existing settings from the configuration backup or you can configure new settings.

   If you change IP address, Veeam Backup & Replication detects it automatically.

If you change the appliance host name, additionally, edit the AHV backup proxy settings.

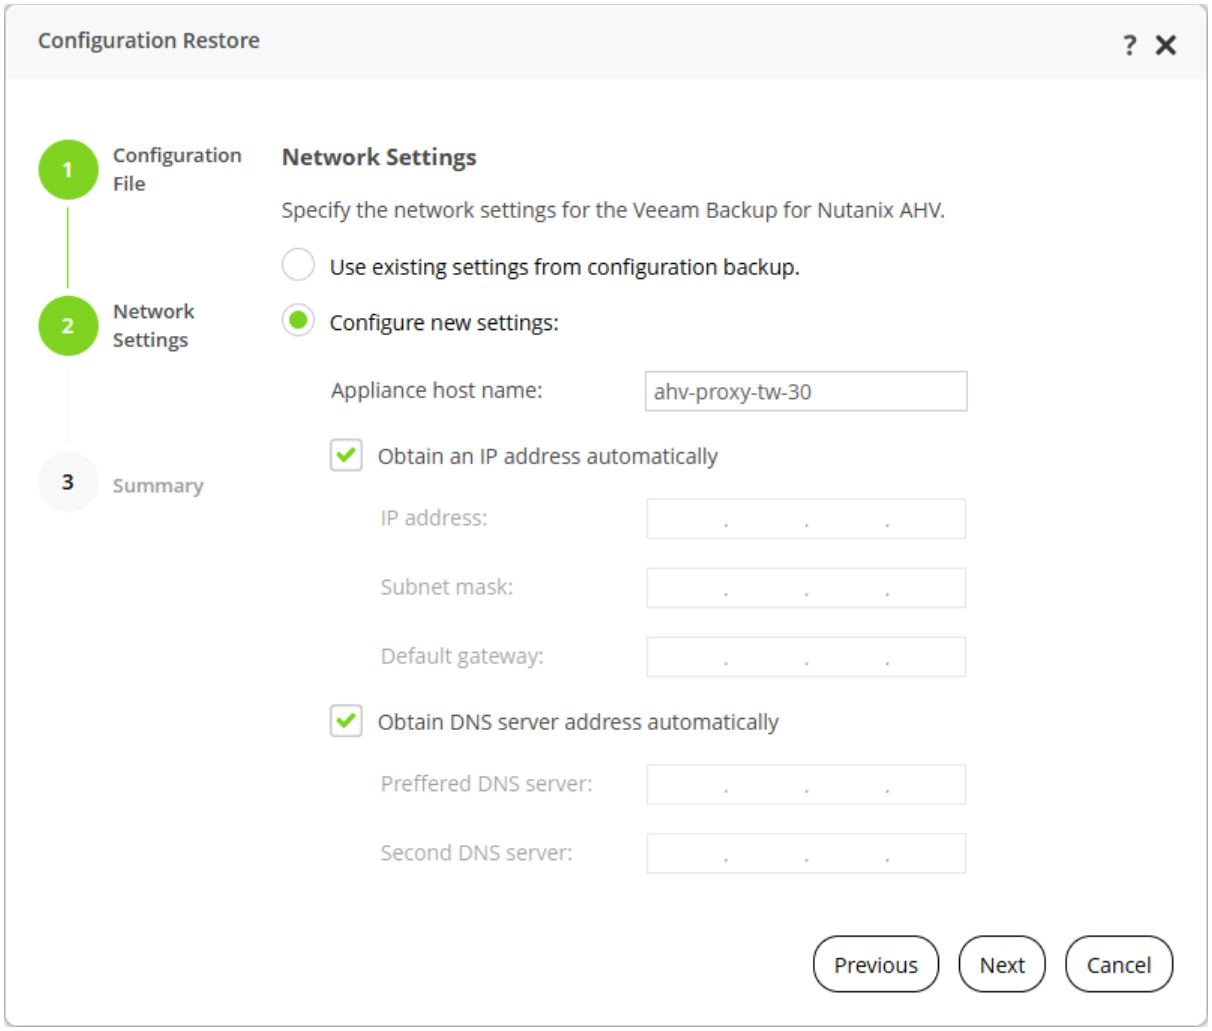

5. At the **Summary** step of the wizard, review the settings and click **Finish** to start the configuration restore process.

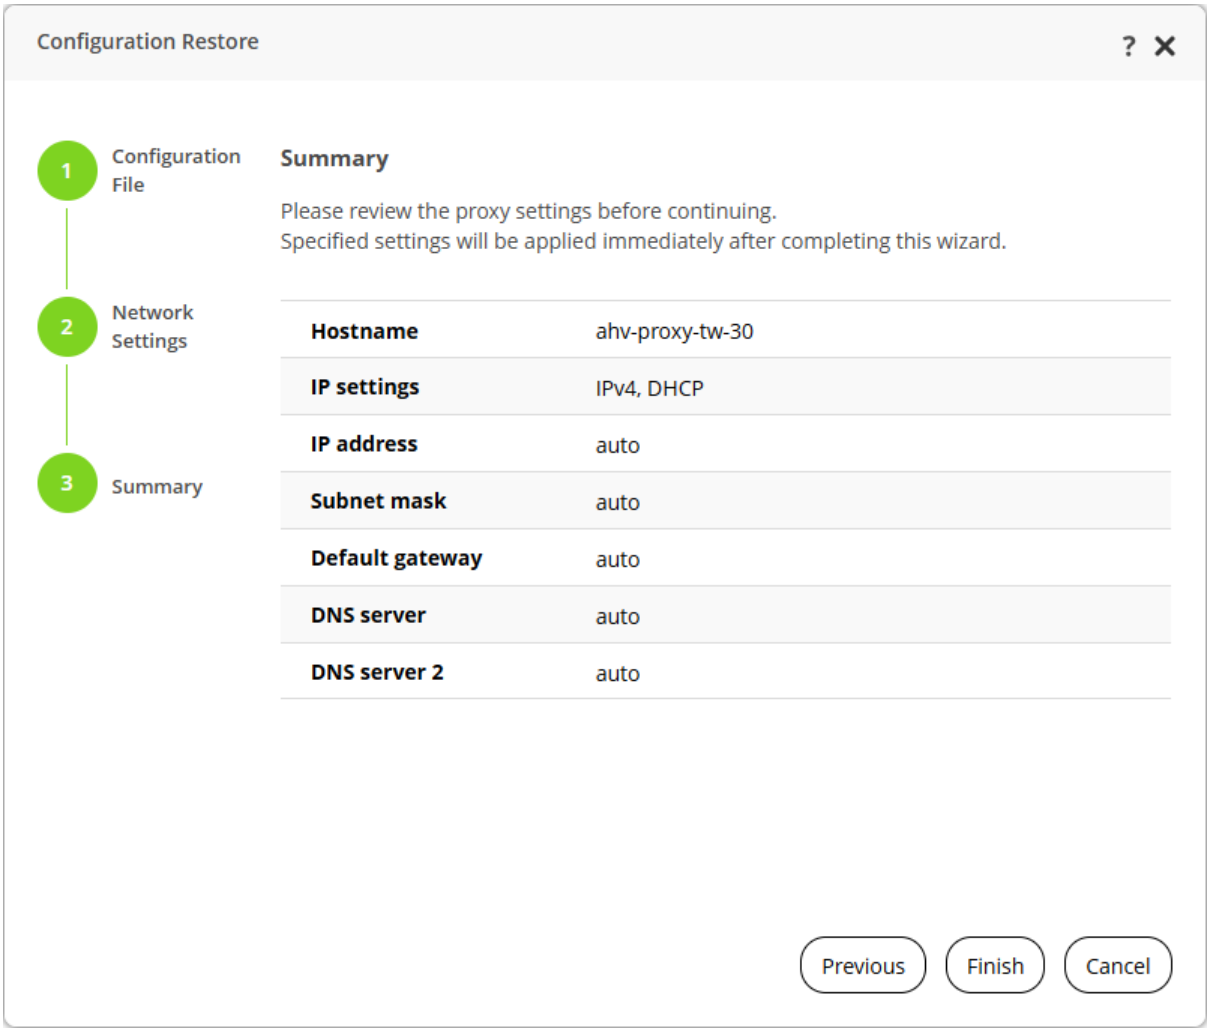

Related Topics

Backing Up Configuration Settings

Rebooting or Shutting Down Proxy VM

To reboot or shut down the AHV Backup Proxy VM, do the following.

1. At the main panel of the web console, click the settings icon and select **Appliance Settings**.
2. At the **Summary** tab of the **Appliance Settings** view, click **Reboot** or **Shutdown**.

   The action will reboot/shut down the VM where AHV Backup Proxy is deployed.

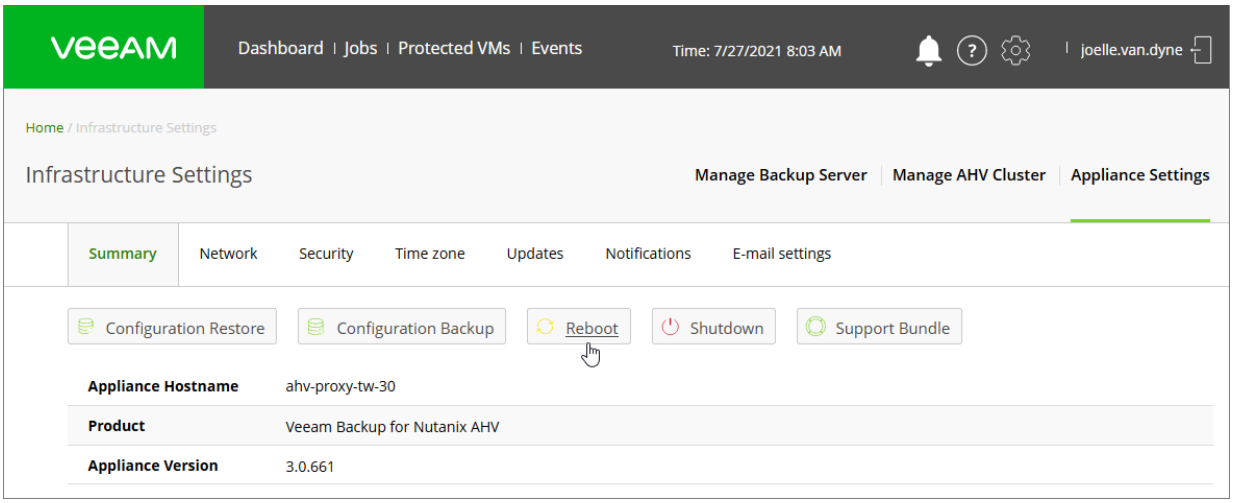

Protecting AHV Backup Proxy VM

After you add a new AHV Backup Proxy, AHV Backup Proxy creates an autoprotection PD snapshot job. This job creates snapshots of the AHV Backup Proxy VM everyday at 22:00. In case of corruption or in case the AHV Backup Proxy VM is deleted, you can use these snapshots to restore the VM.

To restore the AHV Backup Proxy VM, you can use the Prism console or the AHV Backup Proxy console. For instructions, see:

- Restoring AHV Backup Proxy VM using Prism console
- Restoring AHV Backup Proxy VM using AHV Backup Proxy web console

Restoring AHV Backup Proxy VM using Prism console

To restore the AHV Backup Proxy VM using Prism console, do the following:

1. In the main menu of the Nutanix Prism console, select **Data Protection**.
2. In the list of protection domains, select *VeeamBackupProtection*.
3. At the bottom of the console, switch to the **Local Snapshots** tab.
4. Select the PD snapshot from which you want to restore the AHV Backup Proxy VM and click **Restore** on the right.

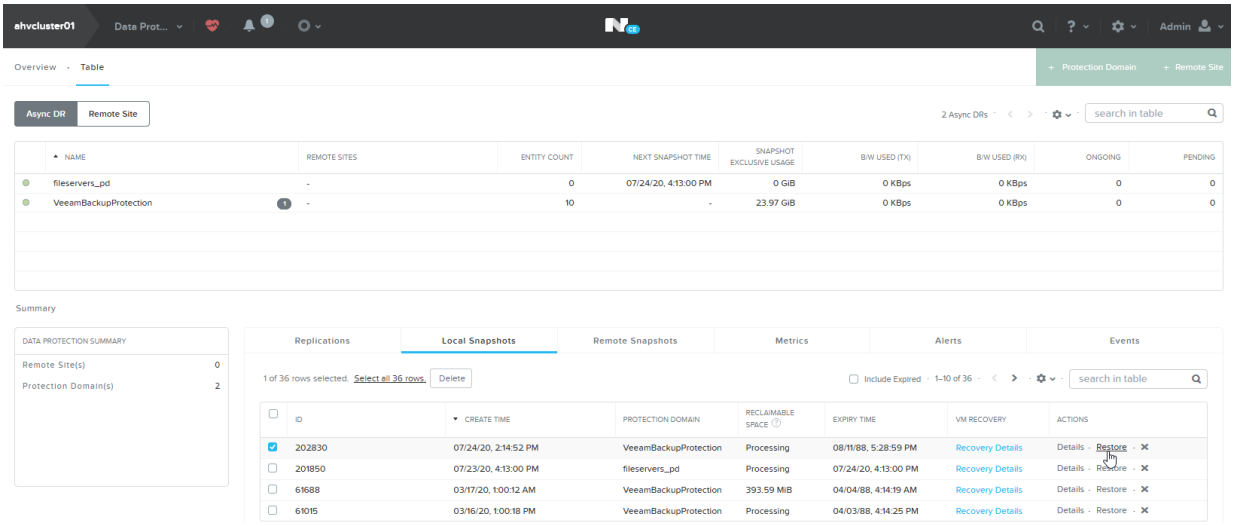

5. In the list of entity names, select the name of the AHV Backup Proxy VM and click **OK**.

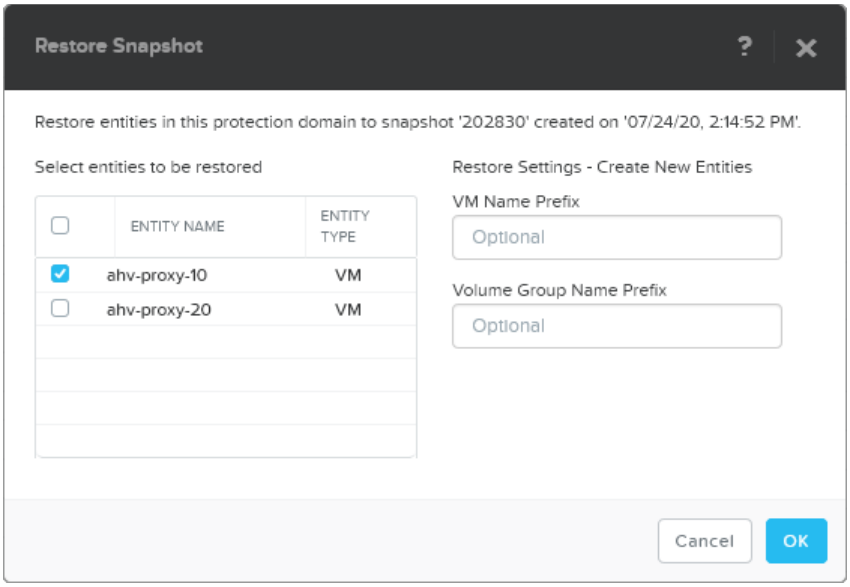

6. After the restore is finished, turn on the restored VM in the Prism console.
7. After the VM is turned on, add the AHV Backup Proxy to the Veeam Backup & Replication infrastructure.

   If you haven't removed the previous AHV Backup Proxy, remove it from the infrastructure before you add the restored proxy.

Restoring AHV Backup Proxy VM using AHV Backup Proxy web console

To restore the VM using the AHV Backup Proxy web console, do the following:

1. Restore the AHV Backup Proxy VM as instructed in Restoring VMs.

   You can restore to a new location or to the original location. If you restore to the original location, AHV Backup Proxy will be synchronized with Veeam Backup & Replication automatically.

   Note that if you select the **Restore to the original location** option at the Restore Mode step of the wizard, the original VM will be replaced by the restored VM and the web console will stop responding until you turn on the VM.
2. After the restore is finished, turn on the restored VM in the Prism console.
3. After the VM is turned on:
   - If you restored the VM to the original location, AHV Backup Proxy will be synchronized with Veeam Backup & Replication automatically.
   - If you restored the VM to a new location, add the AHV Backup Proxy to the Veeam Backup & Replication infrastructure. If you haven't removed the old AHV Backup Proxy, remove it from the Veeam Backup & Replication infrastructure before you add the restored proxy.

Data Protection

With Veeam Backup for Nutanix AHV, you can perform the following data protection tasks:

- Create backups of AHV VMs.
- Create snapshots of AHV VMs (individual snapshots and PD snapshots).
- Create VeeamZIP backups of AHV VMs.
- Copy backups to secondary location.
- Copy backups to tape devices.

> **NOTE**
>
> In Veeam Backup for Nutanix AHV version 2.0 and later, protection operations can be performed in two different consoles:
>
> - In the AHV Backup Proxy web console, you can create backups and snapshots of AHV VMs.
> - In the Veeam Backup & Replication console, you can create VeeamZIP backups and copy backups to secondary location or to tape devices.

Backup Jobs

AHV Backup Proxy creates image-level backups of VMs. AHV Backup Proxy retrieves VM data from the Nutanix AHV cluster at a block level, compresses and deduplicates it, and forwards it to a backup repository in the Veeam proprietary format. You can use image-level backups to restore VMs and VM files.

Backup is a job-driven process. To perform a backup, you need to configure a backup job. Backup job is a configuration unit which allows you to specify VMs that you want to back up, select a backup repository, define schedule and other backup settings.

In this Section

- How Backup Job Works
- Backup Chain and Retention Policy
- Creating Backup Jobs
- Creating Active Full Backup

How Backup Jobs Work

When you launch a backup job in the AHV Backup Proxy web console, the following happens:

1. The web console sends the backup job configuration data to AHV Backup Proxy.
2. AHV Backup Proxy starts a backup job and forwards the backup session data to Veeam Backup & Replication.
3. AHV Backup Proxy connects to the AHV cluster over Nutanix REST API and takes a snapshot of VMs or a PD added to the job.
4. AHV Backup Proxy creates a volume group (VG) on the AHV cluster and mounts VM disks via iSCSI and retrieves VM data on block level.
5. Veeam Data Mover compresses and deduplicates the VM data and forwards it to the backup repository in the Veeam proprietary format.

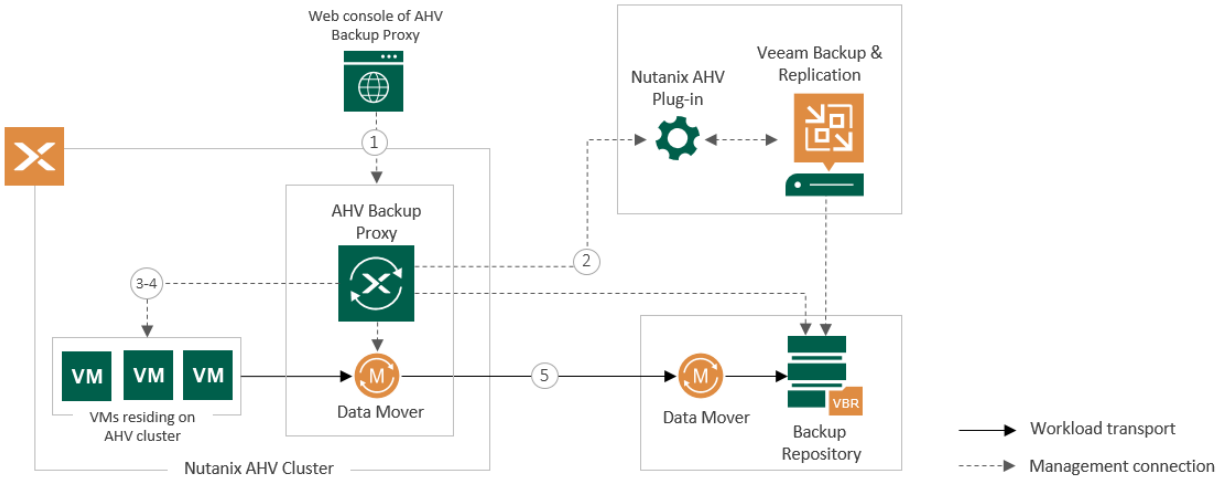

> **NOTE**
>
> This diagram is valid only for the backup process. When you run snapshot jobs, snapshots are not transported to a backup repository. Snapshots are stored on the AHV cluster.

Related Topics

- Backup Infrastructure Components
- Creating Backup Jobs
- Nutanix AHV Snapshot Types

Backup Chain and Retention Policy

See in this section:

- Backup Chain
- Changed Block Tracking
- Retention Policy
- Active Full Backup

Backup Chain

For backups of Nutanix AHV VMs, AHV Backup Proxy implements the forever forward incremental and forward incremental backup method. For more information on backup methods, see the Backup Methods section of the Veeam Backup & Replication User Guide.

To back up Nutanix AHV VMs, you must configure a backup job in the AHV Backup Proxy web console. You can instruct AHV Backup Proxy to run jobs automatically by schedule or start them manually.

During the first backup session, AHV Backup Proxy copies the whole content of the VM and creates a full backup file (VBK) in the repository. The full backup file acts as a starting point of the backup chain.

During the second and further backup sessions, AHV Backup Proxy copies only those data blocks that have changed since the previous backup session. Changed data blocks are stored as an incremental backup file in the repository. Incremental backup files depend on the full backup file and preceding incremental backup files of the backup chain. In the backup job settings, you can also enable periodic creation of active full backups.

One backup job can be used to process one or several VMs. For Nutanix AHV VMs, AHV Backup Proxy creates per-VM backup chains: one backup chain contains data for one VM only. Even if you add several VMs to the backup job, AHV Backup Proxy creates several backup chains on the backup repository, one per every VM processed by the job.

Backup jobs and backup files created by AHV Backup Proxy are visible in the Veeam Backup & Replication console. However, some of the Veeam Backup & Replication functionality is not available for backups of Nutanix AHV VMs. For more information, see Backup Job Limitations.

> **NOTE**
>
> If you want to create application-consistent backups of AHV VMs, follow the guidelines of Nutanix. For more information, see the Application-Consistent Snapshot Guidelines section of the Prism Web Console Guide.

Changed Block Tracking (CBT)

When AHV Backup Proxy performs incremental backup, it needs to know what data blocks have changed since the previous job session. To get the list of changed data blocks, AHV Backup Proxy uses the Nutanix AHV changed block tracking mechanism (CBT). CBT increases the speed and efficiency of incremental backups.

Nutanix AHV CBT uses REST API to determine the changed metadata regions in Nutanix AHV virtual disks. CBT compares two snapshots: a snapshot of the current virtual disks environment with the snapshot that was obtained during the previous backup. Only the changed metadata regions that the API detects are backed up and written to backup repository. After the backup completes, the current snapshot will become the previous snapshot for comparison in the next backup operation. Only one snapshot per VM will be retained for comparison.

Retention Policy

Every successful backup job session creates a new restore point that lets you roll back VM data to an earlier point in time. To control the number of restore points in the backup chain, you must specify retention policy settings. The retention policy defines how many restore points you want to retain on disk and thus how 'far' you are able to roll back. After the allowed number of restore points is exceeded, AHV Backup Proxy automatically removes the earliest restore point from the backup chain.

To define the retention policy for a backup job, you must specify the necessary number of restore points in the **Restore points to keep** field in the backup job settings. By default, AHV Backup Proxy keeps 7 restore points.

For more information on the retention policy for AHV VM backups, see the Forward Incremental Backup Retention Policy section of the Veeam Backup & Replication User Guide.

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, AHV Backup Proxy lets you create active full backups.

The active full backup produces a full backup of a VM, just as if you run the backup job for the first time. AHV Backup Proxy retrieves data for the whole VM from the source, compresses and deduplicates it and stores it to the full backup file — VBK.

The active full backup resets a backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the retention policy.

You can create active full backups manually or schedule a backup job to create active full backups periodically.

- To create an active full backup manually, see Creating Active Full Backup.
- To schedule active full backups, specify scheduling settings in the **Advanced** section of a corresponding backup job. You can schedule active full backups to run weekly, for example, every Saturday, or monthly, for example, every fourth Sunday of a month. For more information, see Specify Repository.

Related Topics

- How Backup Job Works
- Creating Backup Jobs

Creating Backup Jobs

To back up VMs, you must configure a backup job. Backup job defines the source VMs, repository where the backup file will be stored, the job schedule and other settings.

Before creating a backup job, check prerequisites and limitations. Then launch the Add new backup job wizard and configure the backup job setting:

1. Specify general settings.
2. Selects VMs to backup.
3. Specify the backup repository where backups will be stored.
4. Define the job schedule.
5. Finish working with the wizard.

Before You Begin

Before you create a backup job, check the following prerequisites and limitations.

Prerequisites

- A backup repository where you want to store backup files must be added to the Veeam Backup & Replication infrastructure.
- A backup repository where you want to store backup files must have enough space to store created backup files.
- AHV Backup Proxy must have access to the target backup repository. For more information, see Grant Permissions.

Nutanix AHV Cluster Limitations

Mind the following limitations on the Nutanix Controller VM (CVM) side:

- The Nutanix CVM cannot be backed up by AHV Backup Proxy.
- If you raise the number of concurrent backup tasks, backup jobs may fail due to CVM resource limitations. The CVM on each node of the cluster may need additional resources.

Backup Job Functionality Limitations

- AHV Backup Proxy creates application-consistent backups when the source VM meets Nutanix requirements described in the Application-Consistent Snapshot Guidelines section of the Prism Web Console Guide. If a VM does not meet these requirements, AHV Backup Proxy will create a crash-consistent backup.

  Note that AHV Backup Proxy does not take part in guest processing of VMs. The creation of application-consistent snapshots is performed on the Nutanix AHV side. Also, keep in mind that Nutanix AHV does not perform data truncation. To perform data truncation, you can use customized pre-freeze and post-thaw scripts. Veeam Software does not provide support for cases with customized scripts.

  To create application-consistent backups, you can also use Veeam Agent for Linux or Veeam Agent for Windows.

- AHV Backup Proxy does not back up data that resides on Nutanix Files. AHV Backup Proxy backs up only disks that are attached to the VM.
- You cannot back up a VM if the restore process is running for this VM. Wait for the restore process to finish, and then you can start the backup job.
- Backup job compression, deduplication and block size settings cannot be changed. Default settings are the following:
    - 1 MB block size;
    - LZ4 compression;
    - Deduplication is enabled.
- You can configure encryption settings in backup repository settings using the Veeam Backup & Replication console. For more information, see the Access Permissions section of the Veeam Backup & Replication User Guide.
- GDPR locations cannot be set within AHV Backup Proxy backup jobs.
- Health check of backup files is not supported for backups created by AHV Backup Proxy.

Limitations for Backing Up Protection Domains

- Only one protection domain can be added in backup job settings. If you need to protect multiple protection domains, create multiple backup jobs.
- Inactive protection domains are not visible in backup job settings.

Limitations for Backing Up Volume Groups

- You cannot add a Volume Group to a backup job directly. To back up a volume group, you must back up a VM to which volume groups are attached.
- If you add a protection domain as source for a backup job, AHV Backup Proxy does not process individual volume groups added to the protection domain. To back up a volume group within a protection domain, you must back up a protection domain that includes VMs to which the required volume groups are attached. Also, a volume group and the VM to which the volume group is attached must be members of the same consistency group.
- AHV Backup Proxy does not process volume groups where CHAP authentication is enabled.

Backup Chain Limitations

- AHV Backup Proxy creates per-VM backup chains: one backup chain contains data for one VM only. Even if you add several VMs to a backup job, AHV Backup Proxy creates several backup chains on the backup repository, one per every VM processed by the job.
- You cannot disable deduplication and compression of backups.
- AHV Backup Proxy can create forever forward incremental backup chains and forward incremental backup chains with periodic full backups. For more information, see Backup Chain and Retention Policy.

- If a VM has an active backup chain and at some point the UUID of the VM is changed (for example, due to migration to another cluster), Veeam Backup for Nutanix AHV cannot continue the old backup chain for this VM. In this case, you can re-add this VM to the backup job and Veeam Backup for Nutanix AHV will create a new backup chain for the VM with a new UUID.

  Note that you can still use backups from the old backup chain for restore operations.

Backup File and Jobs in Veeam Backup & Replication Console

Jobs and backups created by AHV Backup Proxy are viewable in the Veeam Backup & Replication console. However, some of the Veeam Backup & Replication functionality is not available for these backups and jobs. Mind the following limitations.

- In the Veeam Backup & Replication console, you cannot create or edit jobs of AHV Backup Proxy.
- In the Veeam Backup & Replication console, you can:
  - Disable scheduled jobs of AHV Backup Proxy;
  - Start and stop jobs of AHV Backup Proxy;
  - Use backup jobs as a source for backup copy or backup to tape jobs;
  - Delete unused jobs of AHV Backup Proxy;
  - View statistics of jobs.

  For more information, see Managing Jobs.

Launch New Job Wizard

To launch the **New job** wizard, do the following:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. Click **Add at the toolbar to launch the New job** wizard.

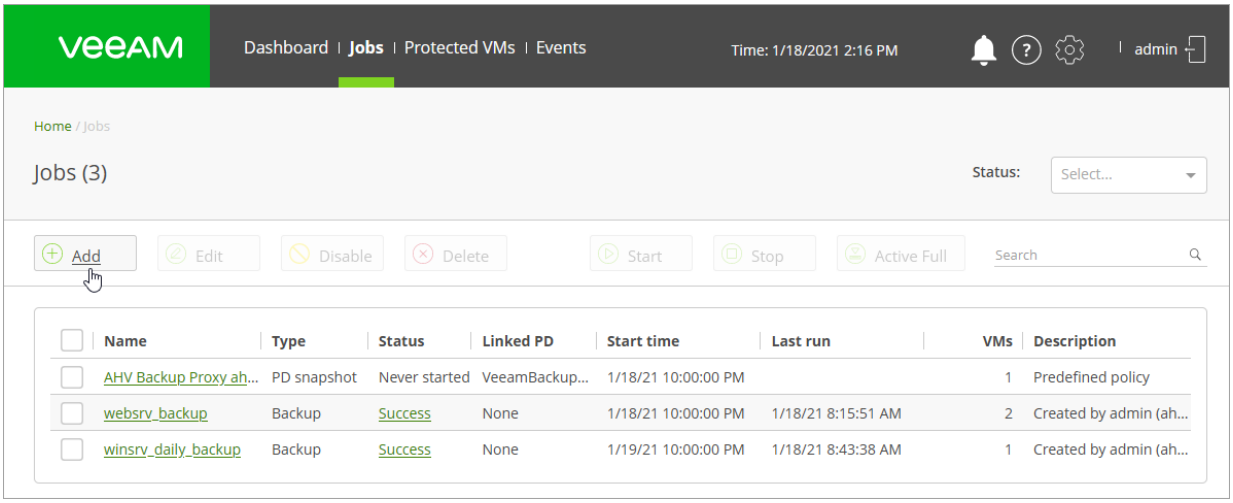

Step 1. Specify General Settings

At the **General Settings** step of the **New job** wizard, specify a name, description for the job and select the job type.

1. In the **Job Name** field, enter a name for the backup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.
3. At the **Select the job type** section, select **Backup job**.

> **NOTE**
>
> Snapshot jobs and PD snapshots jobs are described in the Snapshot-Only Jobs section.

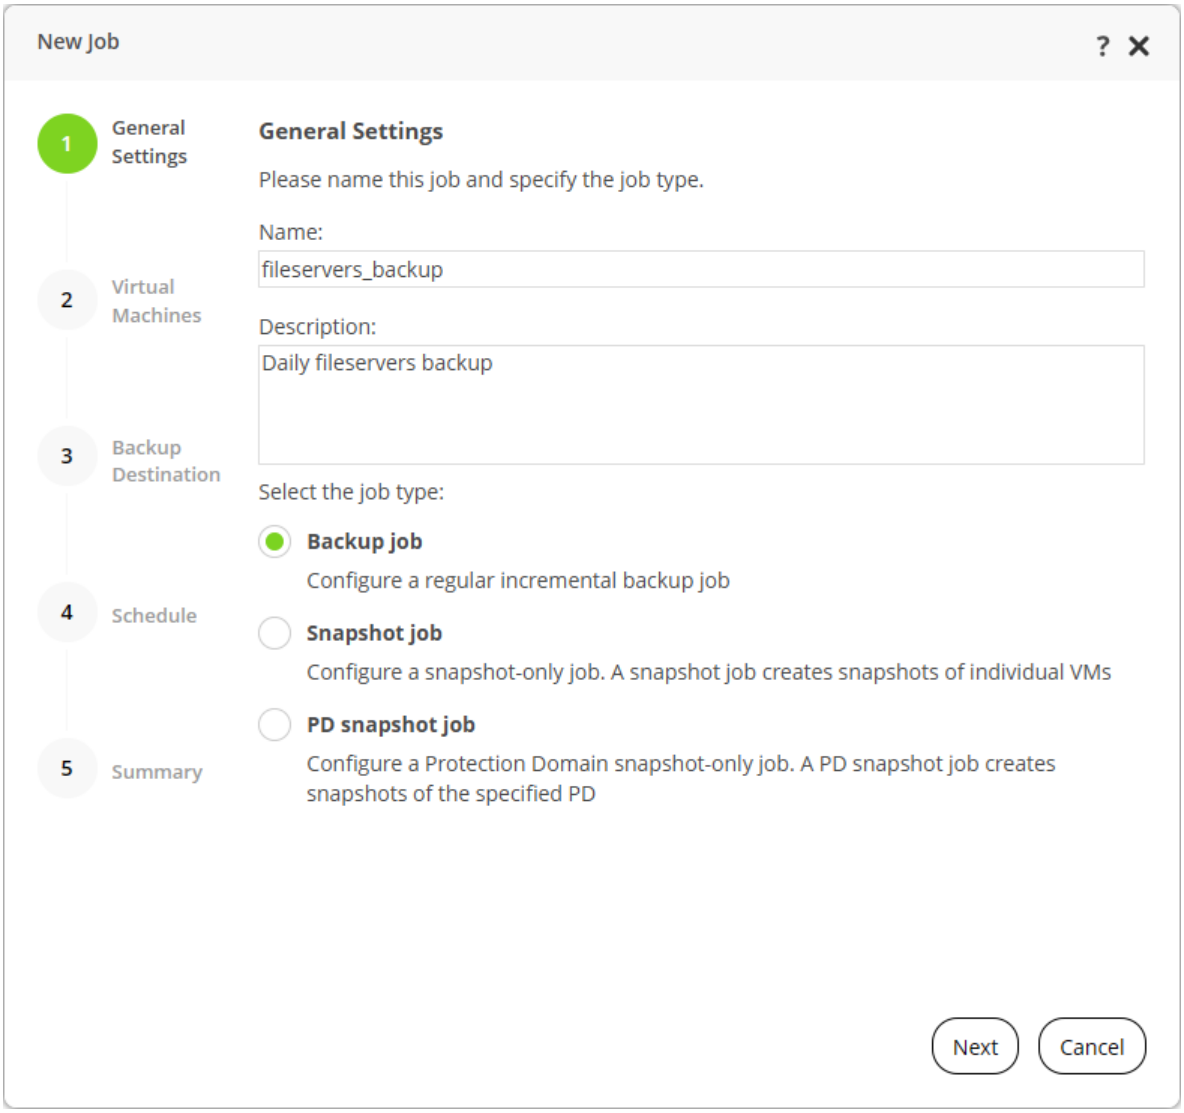

Step 2. Select VMs to Back Up

At the **Virtual Machines** step of the wizard, specify VMs or protection domains that you want to back up.

For instructions on how to add objects to a backup job, see:

- Adding VMs to Backup Job
- Adding Protection Domains to Backup Job

> **NOTE**
>
> You cannot add volume groups to a backup job directly. If you want to back up a volume group, you must backup up a VM to which the volume group is attached. For VMs with attached volume groups, mind the following:
>
> - If you back up a VM that has attached volume groups, AHV Backup Proxy may take snapshots of the VM and the volume group not simultaneously. In this case, the backup will not be crash-consistent.
> - If you want to create a crash-consistent backup of a volume group, you must back up a protection domain which includes the VM to which the volume group is attached. Note that the volume group and the VM must be members of the same consistency group. If the volume group and the VM are not members of the same consistency group, AHV Backup Proxy will process only the VM, the volume group will not be backed up.
>
>   For backup jobs that include a protection domain with consistency groups that contain two or more entities, AHV Backup Proxy does not use Veeam Changed Block Tracking. This means that incremental backup jobs will run longer.

Adding VMs to Backup Job

To add VMs to a backup job, do the following:

1. At the **Virtual Machines** step of the wizard, click **Add on the toolbar**.
2. In the **Add Objects** window, click **+** on the left of the cluster name to expand the list of VMs.

3. From the list of VMs, select VMs that you want to back up and click **Add**.

   To quickly find the necessary VM, you can use the search field at the top of the **Select objects** field.

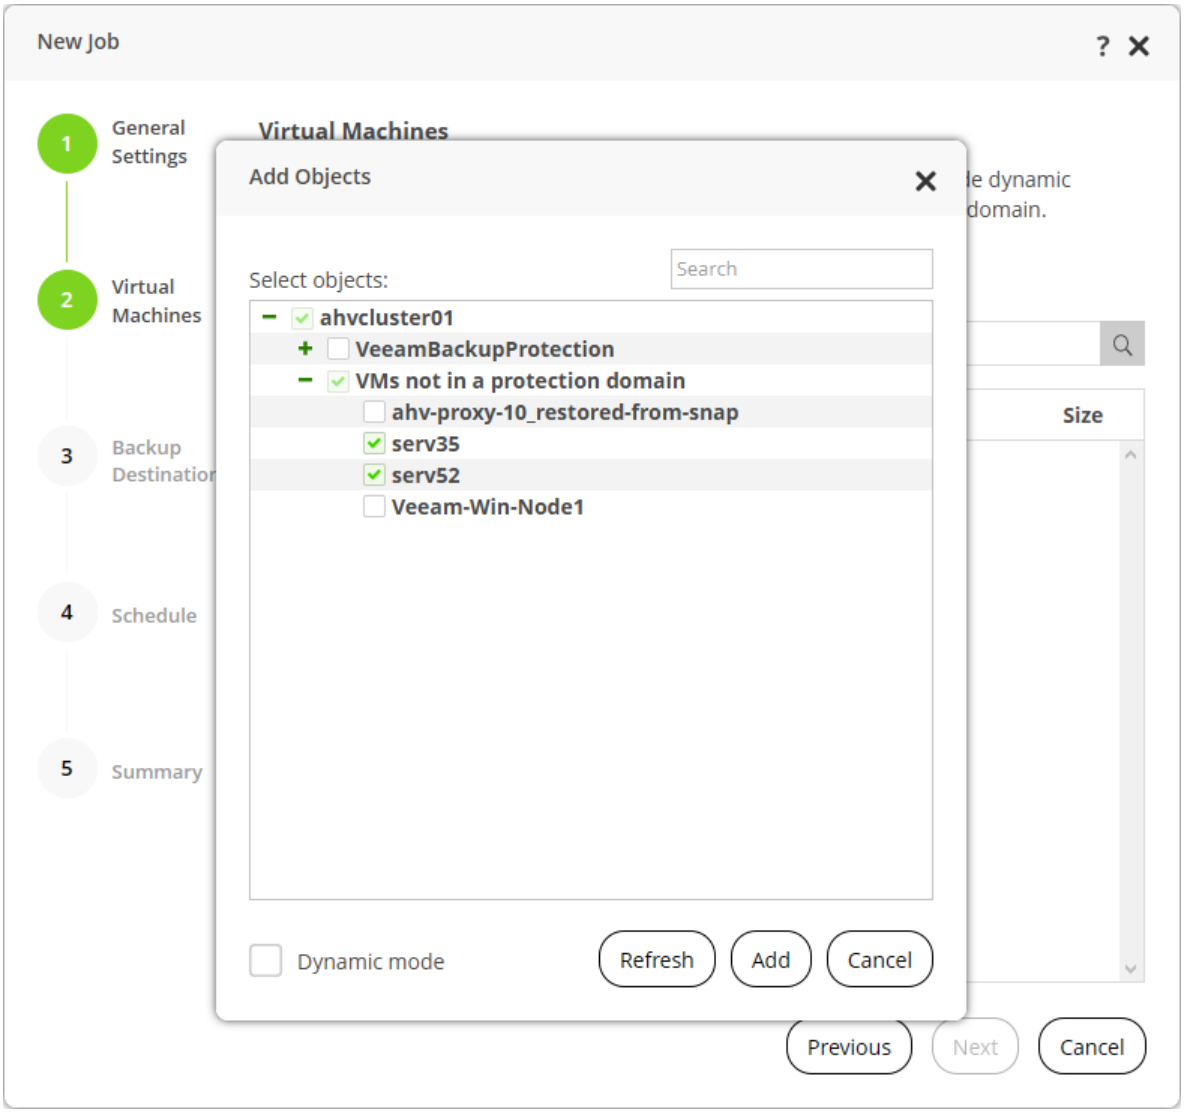

4. [Optional] If there are VM disks that you don't want to back up, you can exclude these disks from processing.

   a. On the wizard window, click **Exclusions**.

b. Select a VM and click **Edit**.

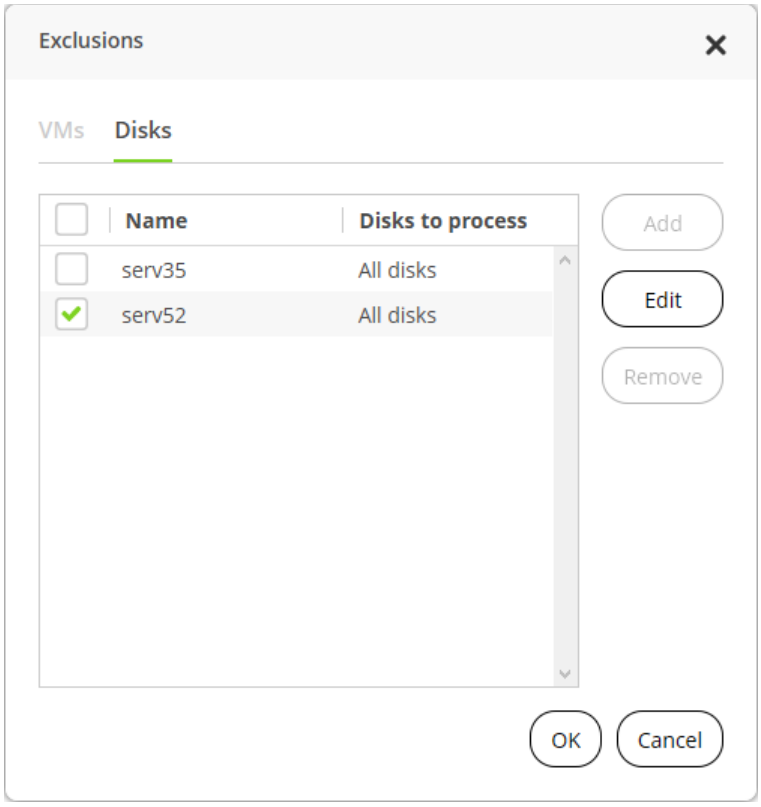

c. In the **Select Disks** window, select **Selected disks** and click **Add**.

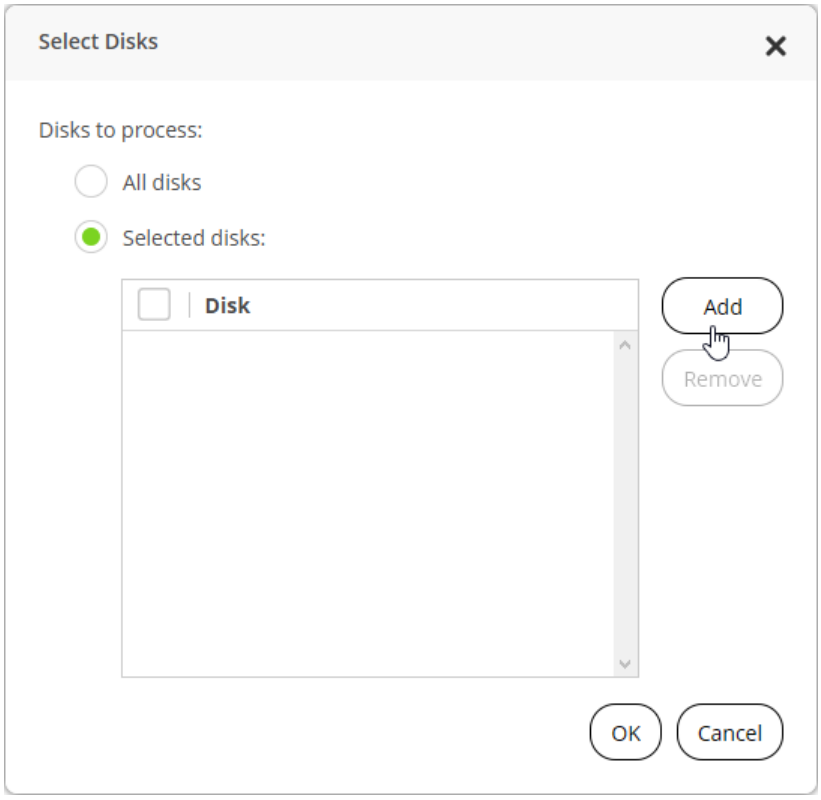

d. In the **Add Disks** window select the bus type or volume groups that you want to include and select the disks.

> **NOTE**
>
> Select the disks that you want to process, not the disks that you want to exclude. AHV Backup Proxy will process the disks that you have selected.

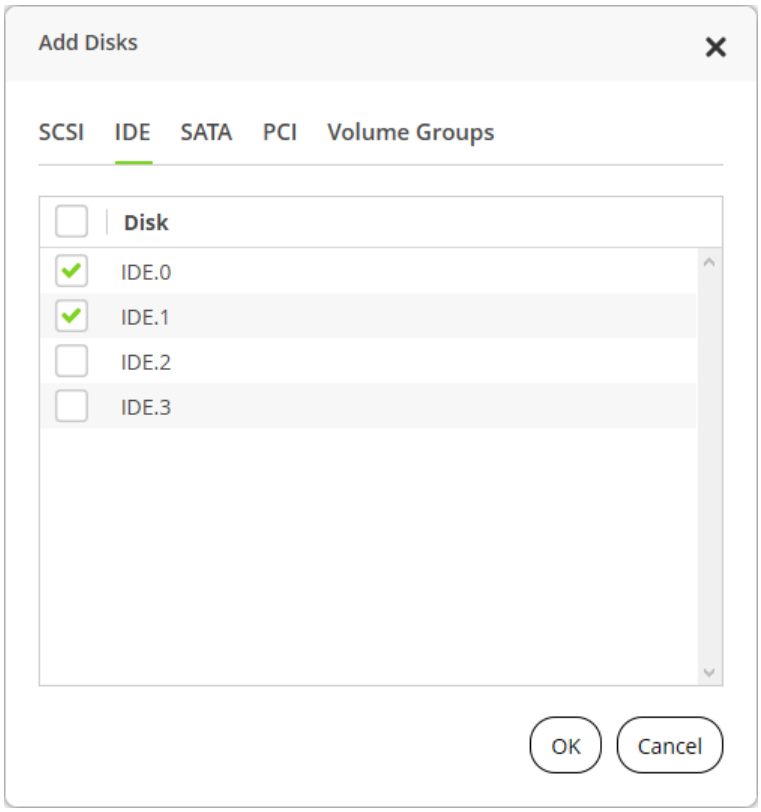

Adding PDs to Backup Job

If you want to select a protection domain instead of individual VMs, do the following:

1. At the **Virtual Machines** step of the wizard, click **Add on the toolbar**.
2. In the **Add Objects** window, select the **Dynamic mode** check box.
3. Click **+** on the left of the cluster name to expand the list of protection domains.

4. From the list of protection domains, select one protection domain and click **Add**.

   To quickly find the necessary protection domain, you can use the search field at the top of the **Select objects** field.

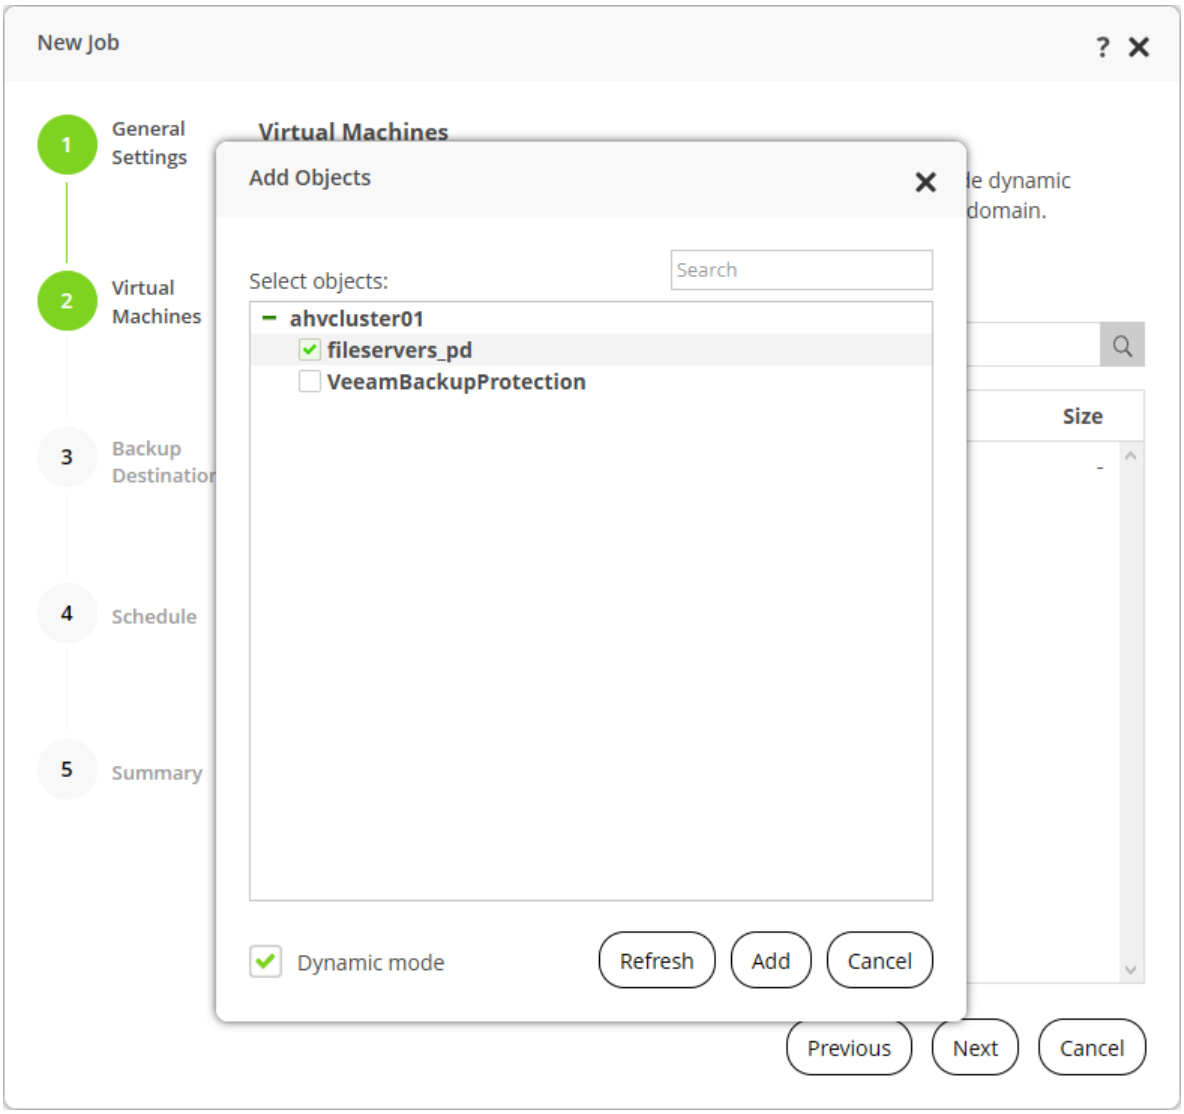

5. [Optional] If you want to exclude some VMs from processing, do the following:

   a. At the **Virtual Machines** steps of the wizard, click **Exclusions on the toolbar**.

   b. In the **VM name contains** field, specify a string value to filter the VMs.

      You can use the following wildcard characters: * and ?. Note that you cannot specify multiple conditions separated by comma, the filter supports only one condition.

c. Click **Ok to save the exclusions setting**.

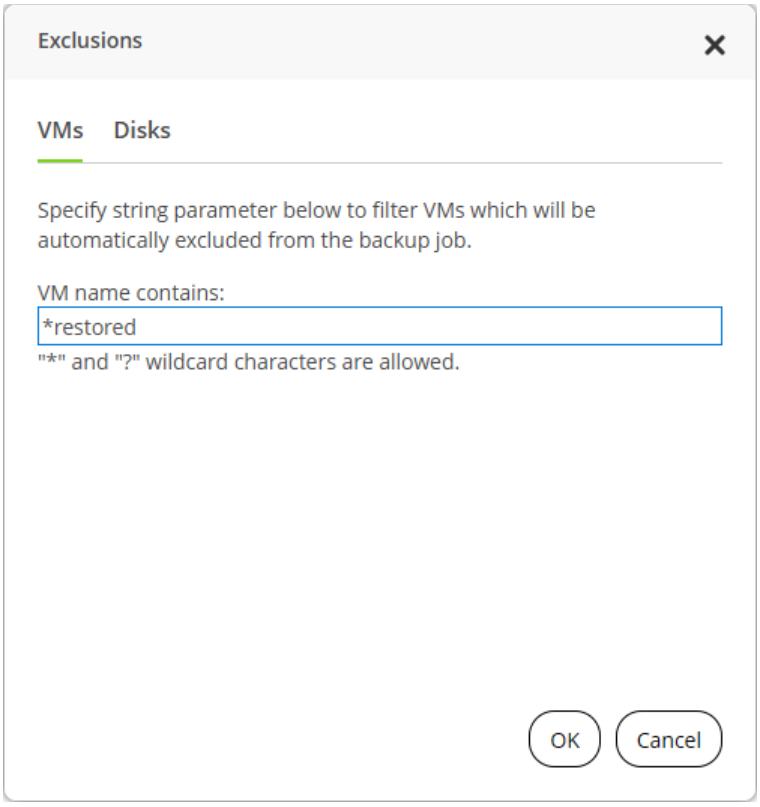

> **NOTE**
>
> If you delete a VM from the protection group added to a backup job, the backup job will not process this VM on next runs.

6. [Optional] If you want to exclude some VM disks from processing, do the following:
   a. On the wizard window, click **Exclusions**.
   b. In the **Exclusions** window, open the **Disks** tab.

c. Select the protection domain and click **Edit**.

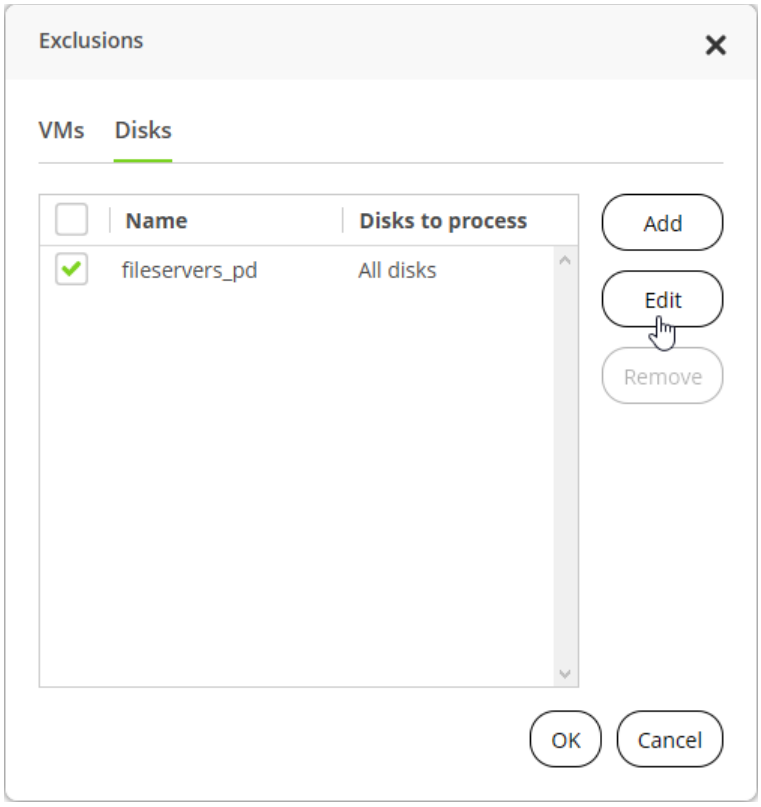

d. In the **Select Disks** window, click **Selected disks** and then click **Add**.

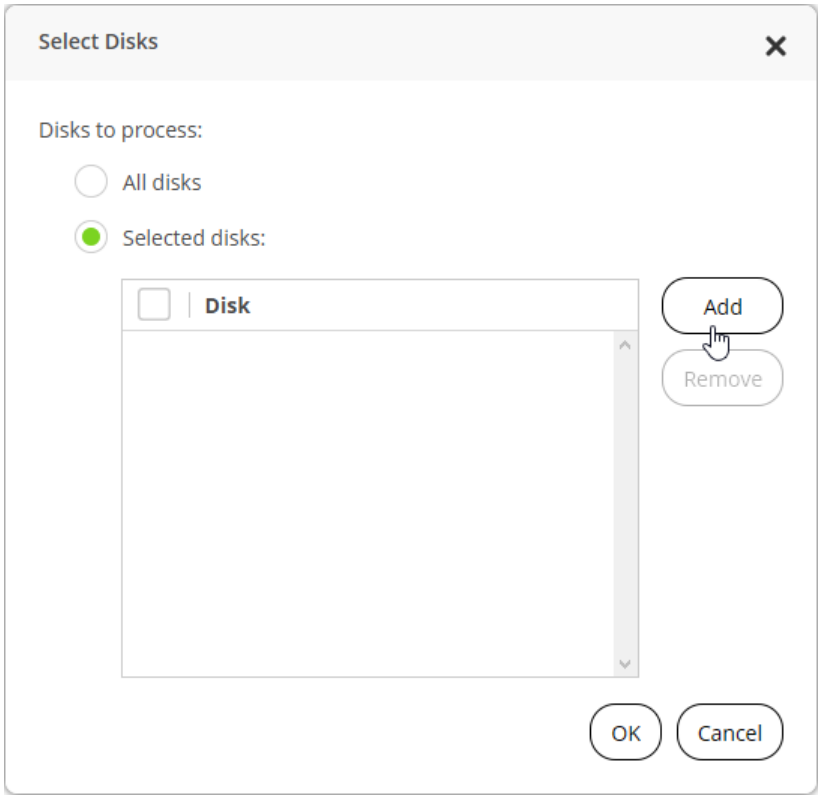

e. In the **Add Disks** window select the bus type or volume groups that you want to include and select the disks.

NOTE

Select the disks that you want to process, not the disks that you want to exclude. AHV Backup Proxy will process the disks that you have selected.

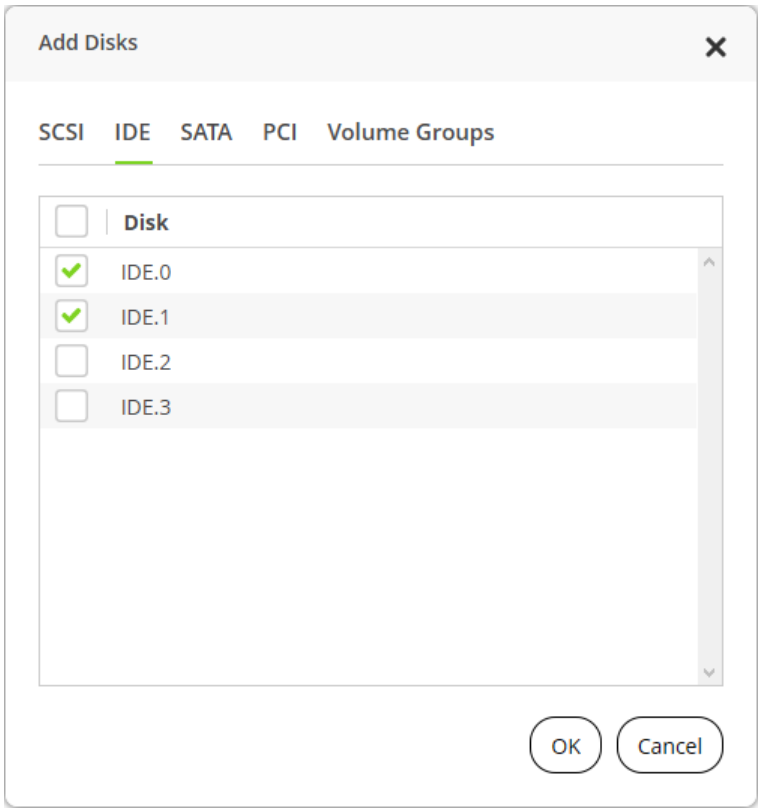

Step 3. Specify Repository

At the **Backup Destination** step of the wizard, specify the repository where you want to store backups.

1. In the **Backup repository** list, select a backup repository where you want to store backups. The wizard displays only the repositories that are managed by the Veeam Backup & Replication server.
2. In the **Restore Points to keep** field, specify the number of restore points that you want to keep. When this number is exceeded, the earliest restore point will be removed from the backup chain. For more information on the retention policy for AHV VM backups, see the Forward Incremental Backup Retention Policy section of the Veeam Backup & Replication User Guide.

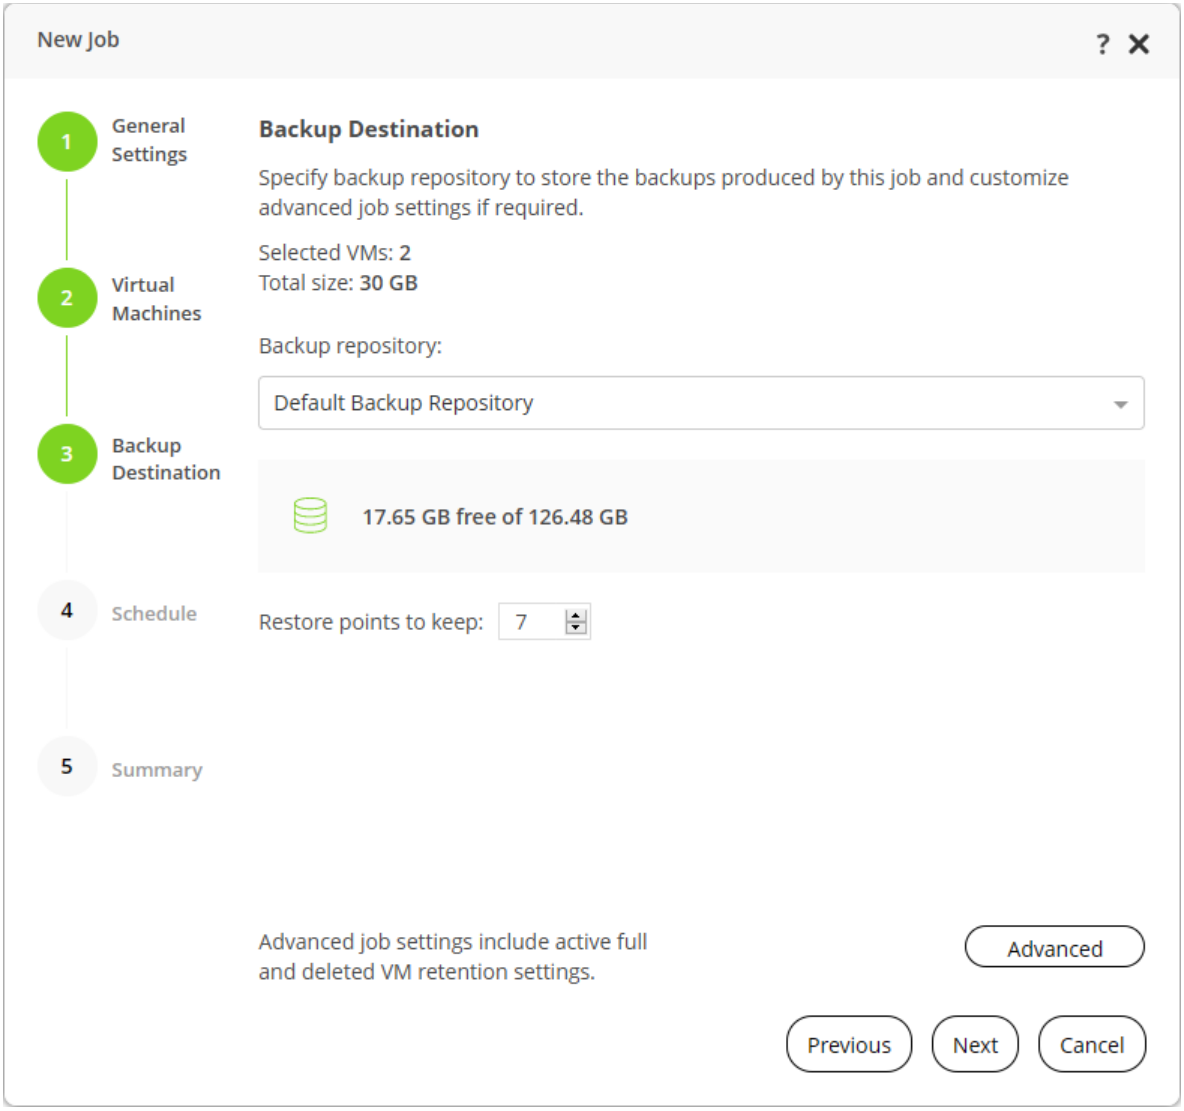

> **NOTE**
>
> You can select only the repositories to which AHV Backup Proxy has access.
>
> You can grant access to repositories at the Grant Permissions step of the **Edit Nutanix Proxy** wizard. To learn how to edit settings of a Nutanix AHV backup proxy, see Editing Nutanix AHV Backup Proxies.

3. Configure a retention policy for backups that are no longer included in the backup job.
   a. At the **Backup Destination** step of the wizard, click **Advanced**.

b. In the **Advanced Settings** window, select the **Remove deleted VMs data after** check box and specify the number of days.

c. Click **Add**.

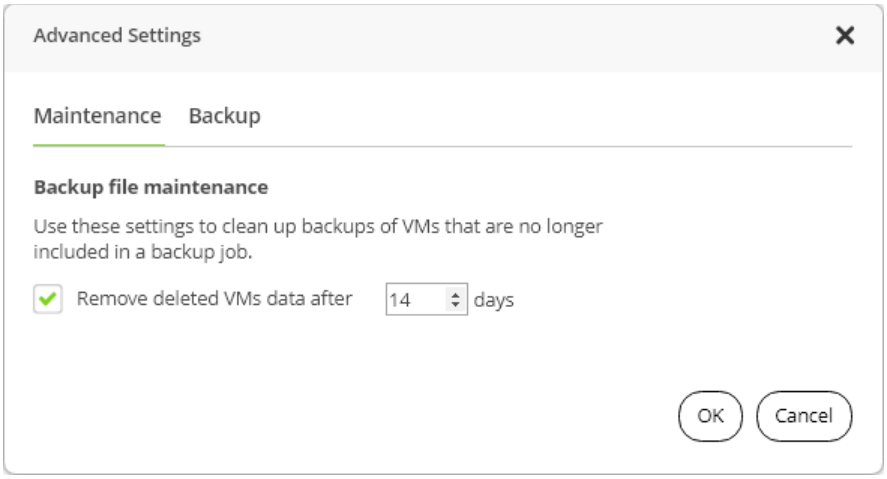

4. To create full backups regularly, select the **Run active full backup periodically** check box. You have the following options:

   - **Weekly on selected days**: with this option, AHV Backup Proxy creates an active full backup every week on selected days. Use the drop-down list to select the required week days.
   - **Monthly on**: with this option, AHV Backup Proxy creates an active full backup on selected days of selected months. Use drop-down lists to select required days and months. For example: every first Monday of all months. By default, all months are selected.

> **NOTE**
>
> Before scheduling periodic full backups, make sure that you have enough free space on the backup repository. As an alternative, you can create active full backups manually when needed. For more information, see Creating Active Full Backup.

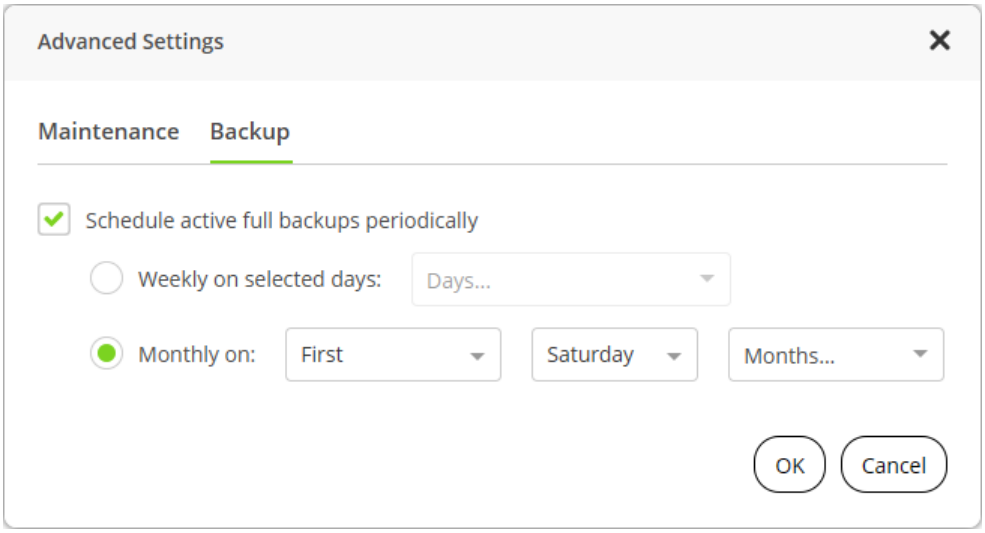

Step 4. Define Job Schedule

At the **Schedule** step of the wizard, define a schedule for the backup job.

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.
2. Define scheduling options for the job.
   - To run the job at specific time daily, on defined week days, select **Daily at this time**.
   - To run the job once a month on specific days, select **Monthly at this time**.
   - To run the job repeatedly throughout a day with a specific time interval, select **Periodically every.**
3. In the **Automatic Retry** section, define whether AHV Backup Proxy must attempt to run the job again if the job fails for some reason. During a job retry, AHV Backup Proxy processes failed VMs only.

   Specify the number of retry attempts and the period of time between the attempts.

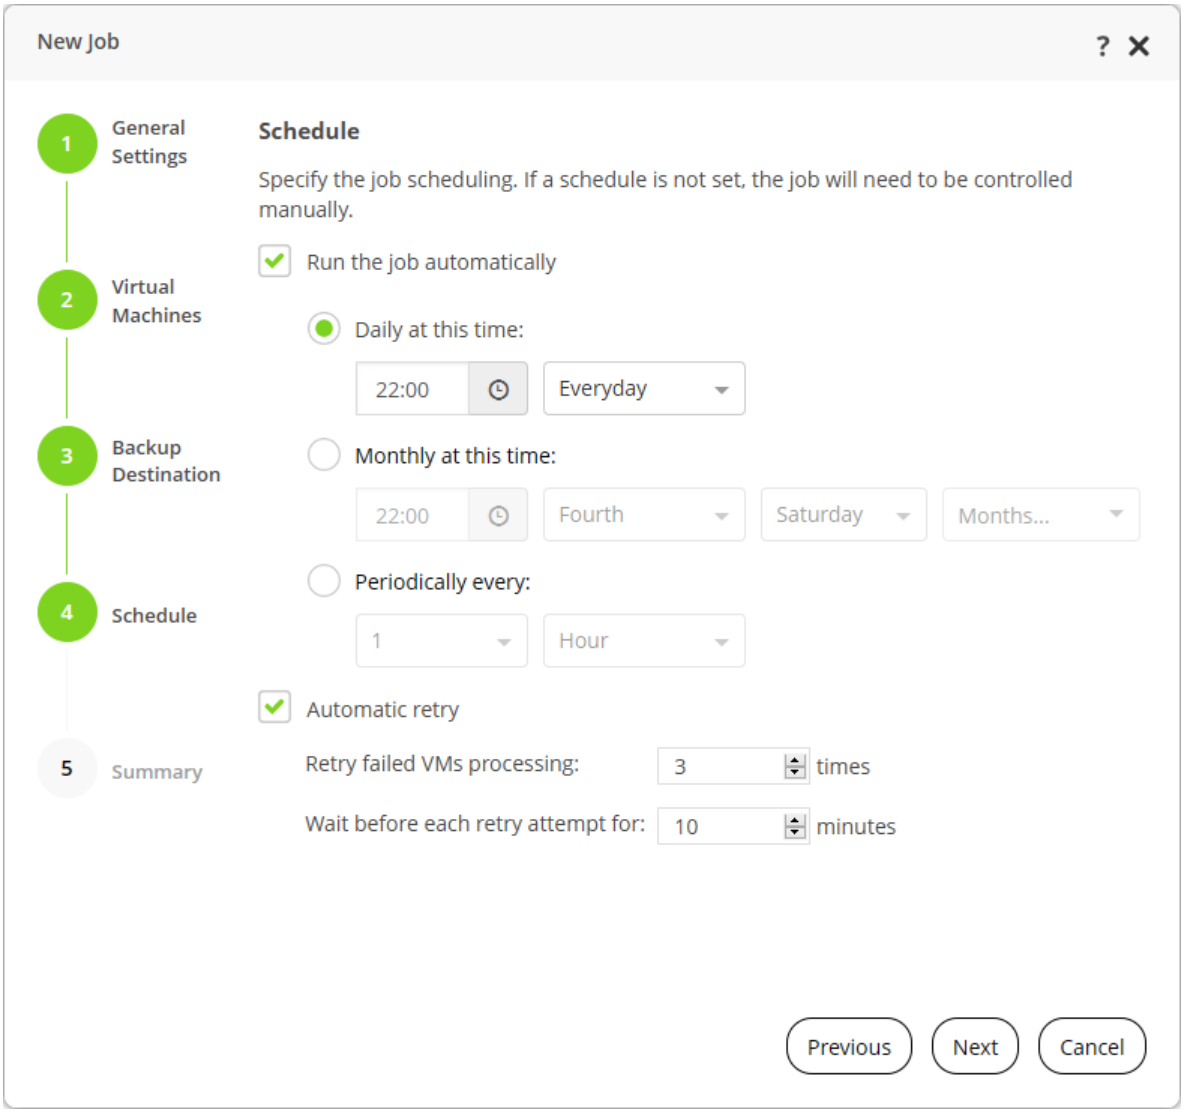

NOTE

Scheduling options in backup jobs of AHV Backup Proxy are similar to backup jobs in Veeam Backup & Replication. For more information, see Backup Job Scheduling and Creating Backup Jobs sections of the Veeam Backup & Replication User Guide.

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, finish configuration of the backup job.

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

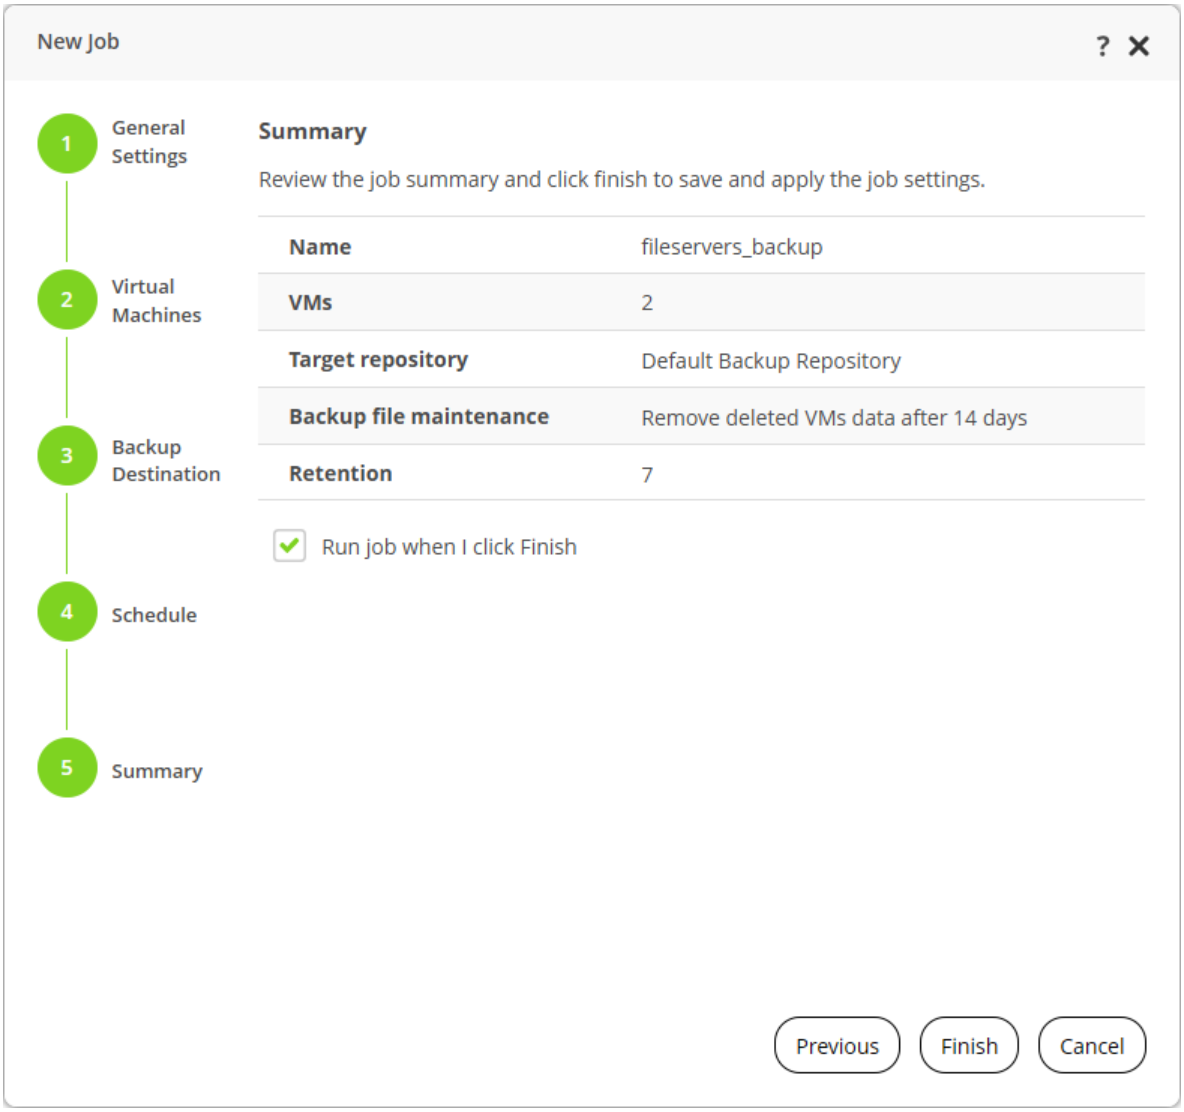

Creating Active Full Backup

When you create an active full backup, AHV Backup Proxy creates a new full backup file and adds it to the backup chain. The active full backup resets the backup chain. All subsequent incremental backups use the active full backup as a starting point. The previously used full backup remains on the backup repository until it is removed from the backup chain according to the retention policy.

To perform active full backup, do the following:

1. At the main panel, click the **Backup Job** tab.
2. Select the required backup job and click **Active Full** on the toolbar.
3. Click **Proceed** to start the creation of an active full backup.

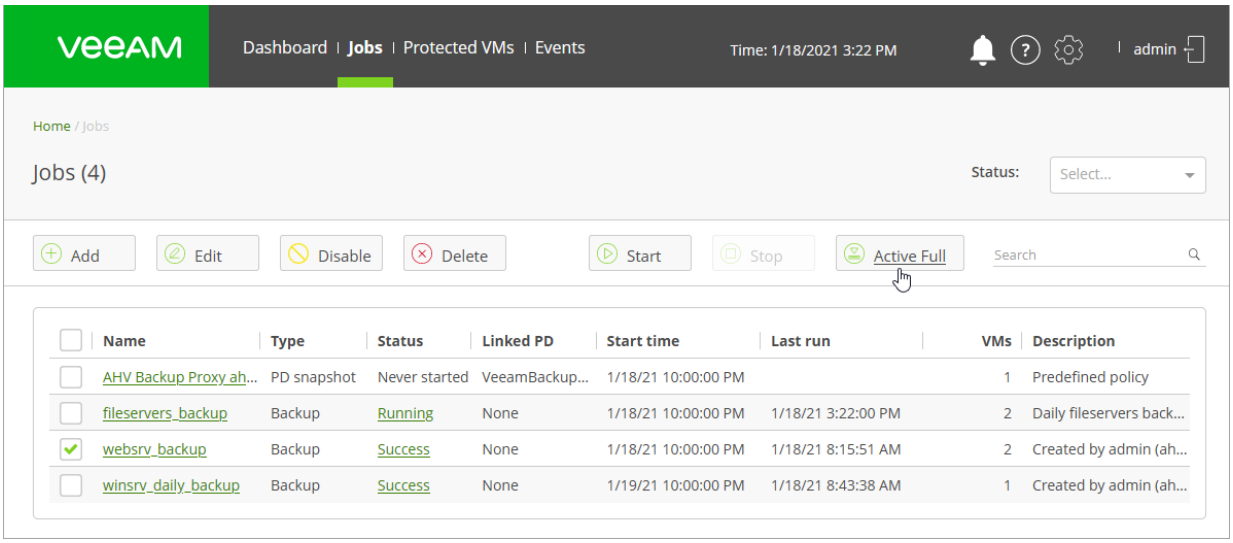

Snapshot-Only Jobs

To create snapshots of Nutanix AHV VMs, AHV Backup Proxy uses built-in capabilities of Nutanix AHV. You can use these snapshots to revert VMs to a state in time when the snapshot was taken. Keep in mind that snapshots are not fully-functional backups but only chunks of data which have been modified since the last snapshot was taken.

Snapshot files are not transferred to Veeam backup repositories, they are stored locally on the AHV cluster on which the source VM resides. For more information on Nutanix AHV snapshots, see Nutanix AHV Snapshot Types.

In AHV Backup Proxy, snapshot is a job-driven process. To create a snapshot, you must configure a snapshot job. Before you create a snapshot job, check prerequisites and limitations.

With AHV Backup Proxy, you can create two types of snapshot jobs:

- PD Snapshot Jobs: creates snapshots of a protection domain.
- Snapshot-Only Jobs: creates snapshots of selected VMs.

In This Section

- Nutanix AHV Snapshot Types
- Snapshot Chain and Retention Policy

Nutanix AHV Snapshot Types

AHV Backup Proxy uses built-in functionality of Nutanix AHV to create snapshots of VMs and volume groups. There are the following types of snapshots:

- **Backup snapshot**: A snapshot created by a backup job of AHV Backup Proxy. Backup snapshots are not visible in the consoles of AHV Backup Proxy and Veeam Backup & Replication. Backup snapshots are used for Veeam Changed Block Tracking (CBT) and Entire VM Restore.
- **VG snapshot**: A snapshot of a volume group which is created by a backup job of AHV Backup Proxy. VG snapshots are not visible in the consoles of AHV Backup Proxy and Veeam Backup & Replication. VG snapshots are used for Veeam Changed Block Tracking and Entire VM Restore.
- **Snapshot**: A snapshot of a VM which is created by a snapshot-only job of AHV Backup Proxy. To learn how to manage snapshots, see Managing Backups and Snapshots.
- **PD snapshot**: A snapshot for VMs and volume groups included in the protection domain (PD). AHV Backup Proxy creates PD snapshots in two cases:
  - When you run a PD snapshot job.
  - When you run a backup job and select a PD as a source object and at least one consistency group has 2 or more elements (VMs or/and VGs).
- **Prism Central snapshot**: Snapshots created manually in the Data Protection dashboard of Nutanix Prism. AHV Backup Proxy does not use Prism Central snapshots in any jobs.

For more information on protection domains, volume groups, consistency groups, see The Nutanix Bible.

Snapshot Chain and Retention Policy

How PD Snapshot Job Works

When you start a PD snapshot job, AHV Backup Proxy uses Nutanix REST API to send a command to create a PD snapshot to AHV cluster.

After every successful snapshot job session, AHV Backup Proxy creates a new restore point. If AHV Backup Proxy detects that the number of restore points in the snapshot chain exceeds the retention limit, the earliest restore point is removed from the chain. Retention policy is set at the **Schedule** step of the **New Job** wizard.

Every 60 seconds AHV Backup Proxy synchronizes with Veeam Backup & Replication and sends data on new or deleted snapshots.

How Snapshot Job Works

When you start a snapshot job, AHV Backup Proxy uses Nutanix REST API to send a command to create VM snapshots to AHV cluster. AHV Backup Proxy creates individual VM snapshots for each VM added to the snapshot job.

After every successful snapshot job session, AHV Backup Proxy creates a new restore point. If AHV Backup Proxy detects that the number of restore points in the snapshot chain exceeds the retention limit, the earliest restore point is removed from the chain. Retention policy is set at the **Schedule** step of the **New Job** wizard.

Every 60 seconds AHV Backup Proxy synchronizes with Veeam Backup & Replication and sends data on new or deleted snapshots.

Related Topics

- Nutanix AHV Snapshot Types
- Creating PD Snapshot Jobs
- Creating Snapshot Jobs

Before You Begin

Before you create a snapshot job, check the following prerequisites and limitations:

- Snapshot and PD snapshot jobs do not support instantly recovered VMs that have not completed migration to the Nutanix cluster storage. These jobs will fail untill VM data has migrated back to the Nutanix cluster.
- A PD snapshot job fails if 2 or more VMs have the same name in one protection domain. The issue occurs because Nutanix REST API does not allow to manage VMs via UUID when restoring from PD snapshots. In the case of a PD backup job, it failovers to a regular backup job with a warning. That means that snapshots will be created separately for each VM.

Creating PD Snapshot Jobs

In AHV Backup Proxy, you can create and schedule a protection domain (PD) snapshot job. The created PD snapshot is not transferred to the Veeam Backup & Replication repository. PD snapshots are stored in the AHV container. You can create PD snapshot jobs, even if the Veeam Backup & Replication server is not added the AHV Backup Proxy infrastructure.

> **NOTE**
>
> Note that you can add only one protection domain to a PD snapshot job.

To configure a PD snapshot job, launch the New Job wizard and configure the PD snapshot job settings:

1. Specify general settings.
2. Select a PD.
3. Configure snapshot retention settings.
4. Define the job schedule.
5. Finish working with the wizard.

Launch New Job Wizard

To launch the **New job** wizard, do the following:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. Click **Add at the toolbar to launch the New job** wizard.

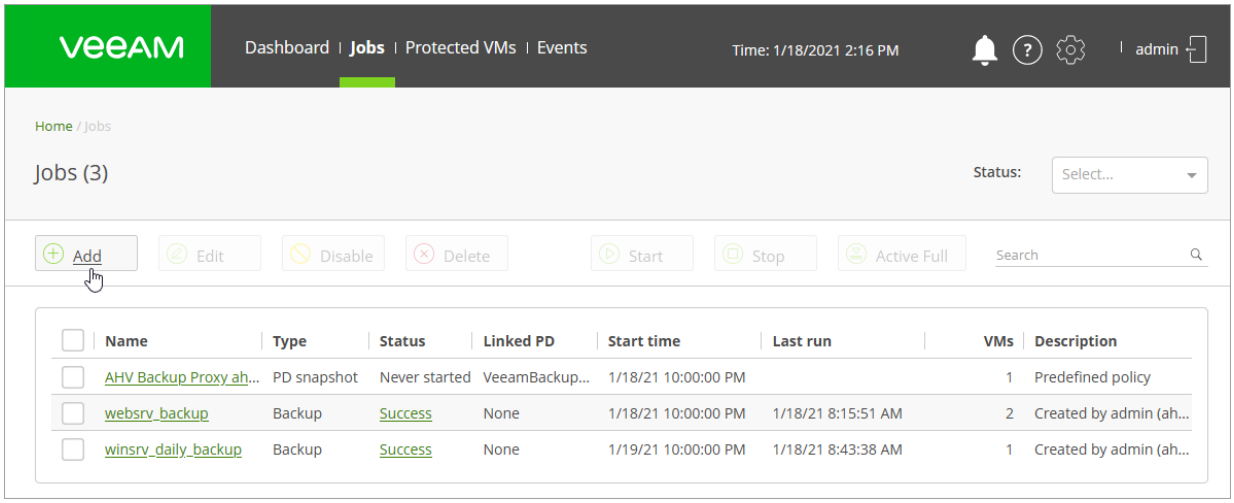

Step 1. Specify General Settings

At the **General Settings** step of the **New Job** wizard, specify a name, description for the job and select the job type.

1. In the **Name** field, enter a name for the PD snapshot 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. In the **Select the job type** field, select **PD snapshot job**.

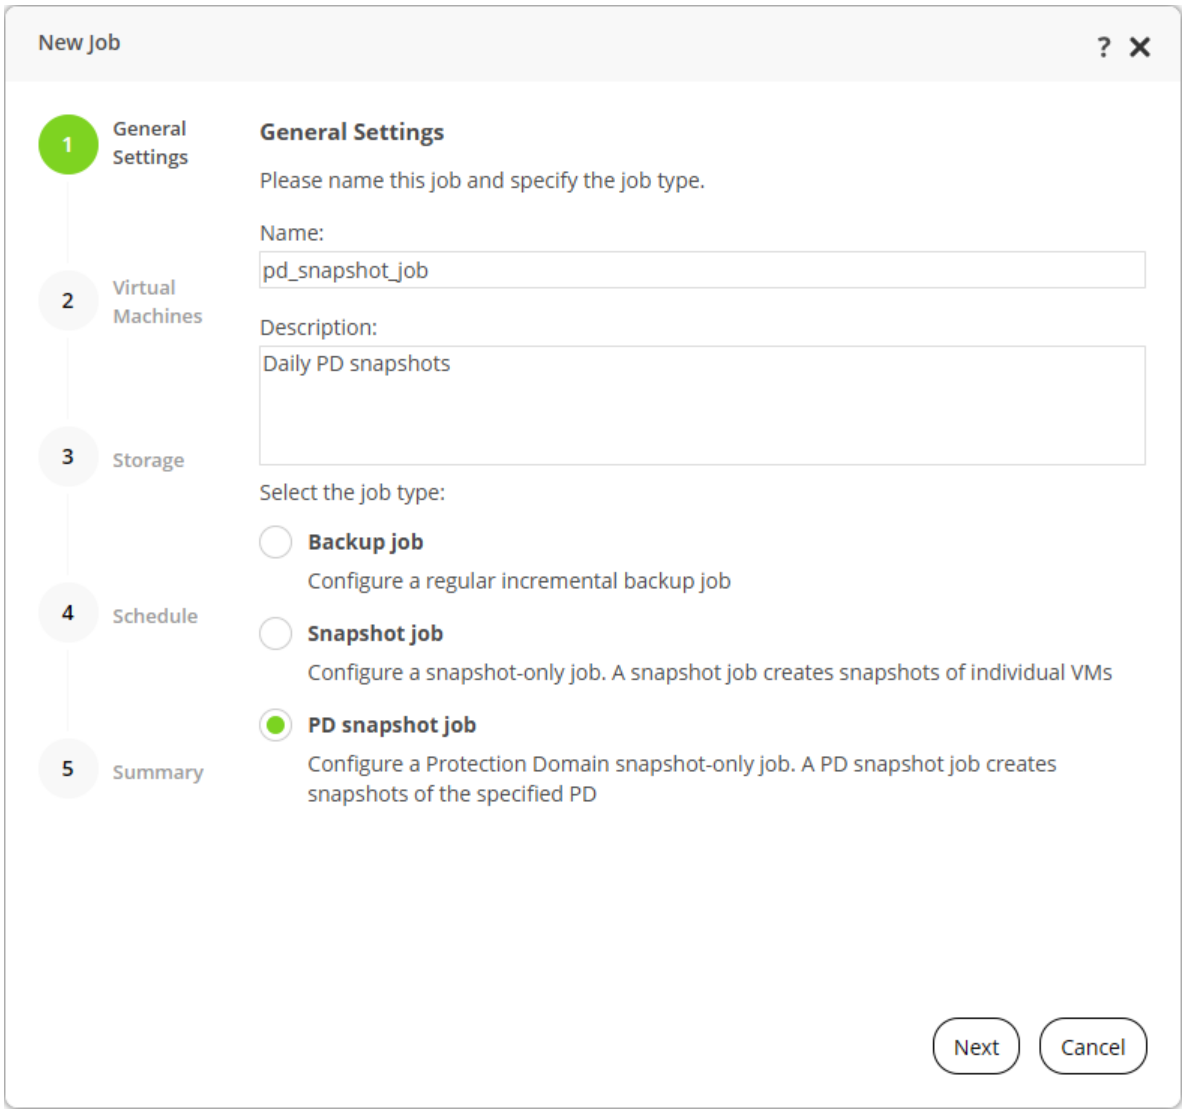

Step 2. Select PD

At the **Virtual Machines** step of the wizard, specify a protection domain that you want to add to the snapshot job.

1. On the toolbar, click **Add**.
2. In the **Add Objects** window, click **+** to expand a cluster which you want to protect.
3. Select a protection domain and click **Add**.

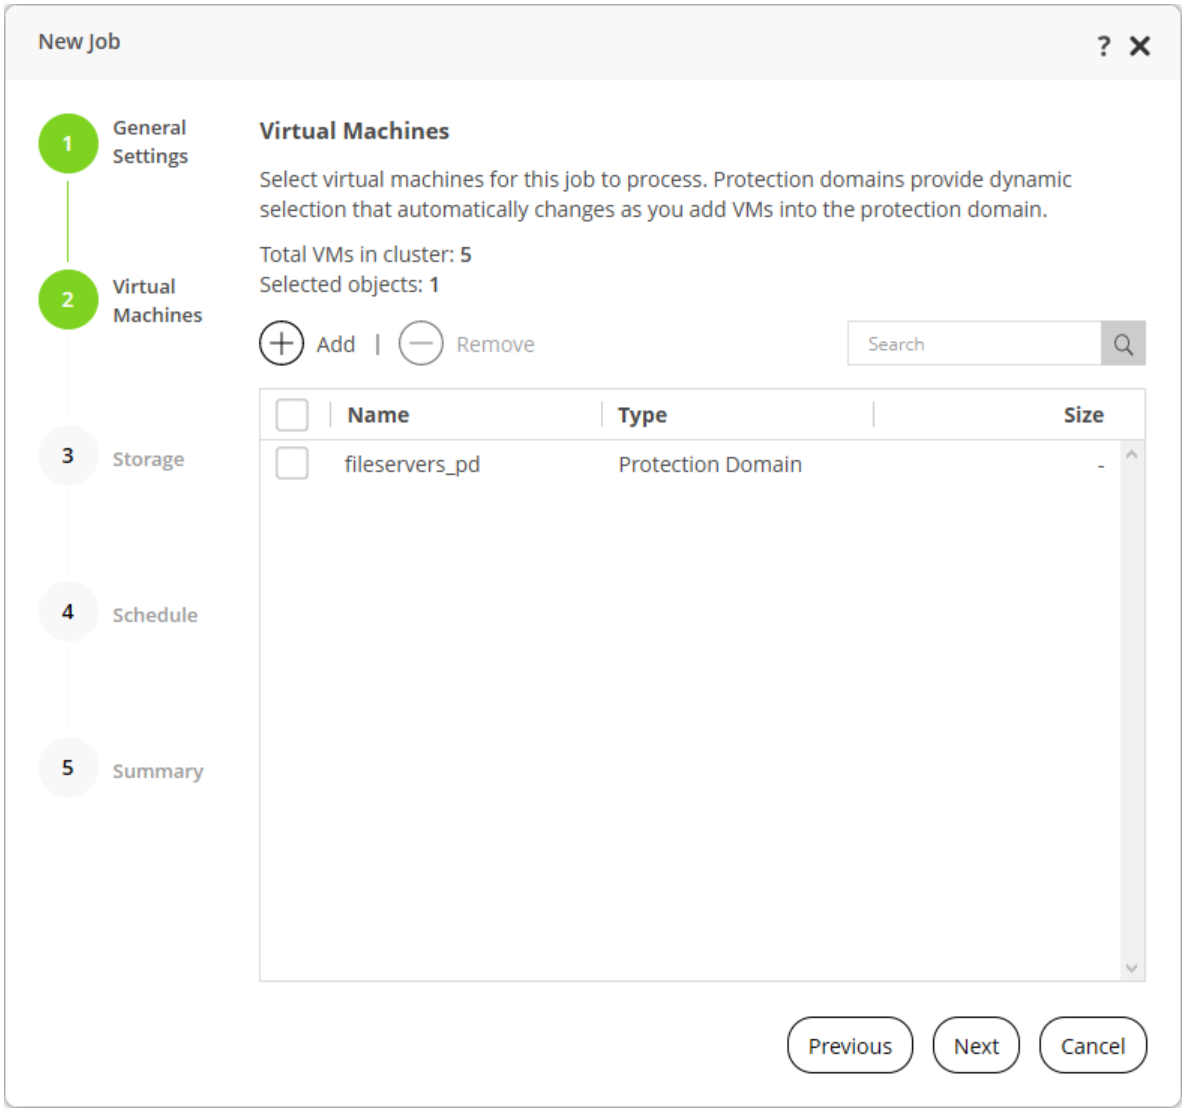

Step 3. Configure Retention

At the **Storage** step of the wizard, specify retention policy settings for restore points.

At the **Restore points to keep** field, specify the number of restore points. When the specified number is exceeded, the earliest restore point is removed from the snapshot chain. For more information, see Snapshot Chain and Retention Policy.

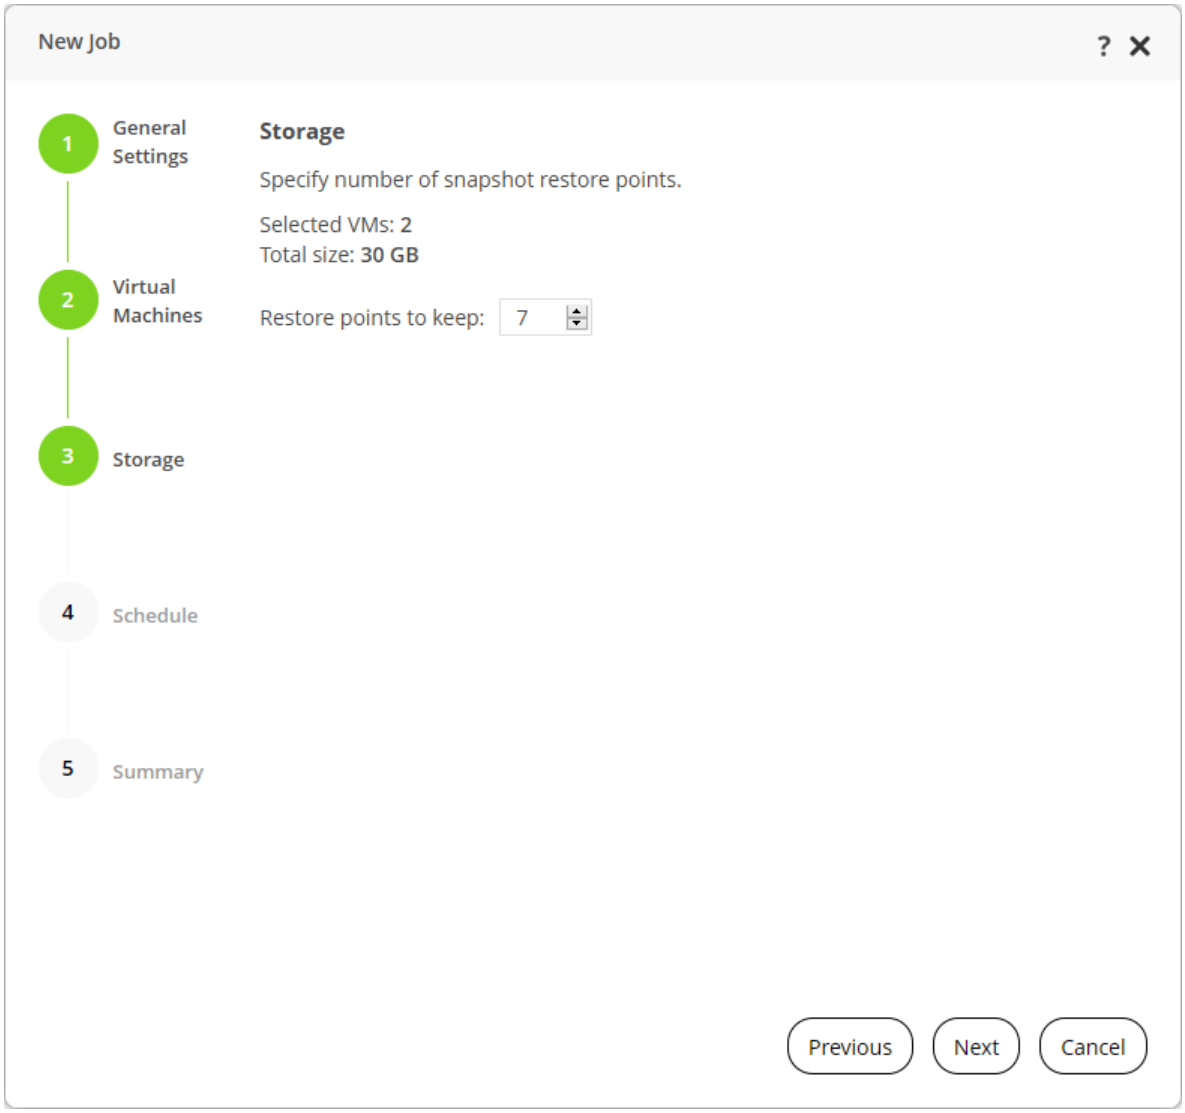

Step 4. Define Schedule

At the **Schedule** step, specify the job scheduling. If you don't set a schedule, you will have to launch the job manually.

1. Select the **Run the job automatically** check box.
2. Define the 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 below to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields below to configure the necessary schedule.
   - To run the job repeatedly throughout a day with a specific time interval, select **Periodically every**. In the field below, select the time period and necessary time unit: *Hours* or *Minutes*.
3. In the **Automatic Retry** section, define whether AHV Backup Proxy must attempt to run the job again if the job fails for some reason. During a job retry, AHV Backup Proxy processes failed VMs only.

   Specify the number of retry attempts and the period of time between the attempts.

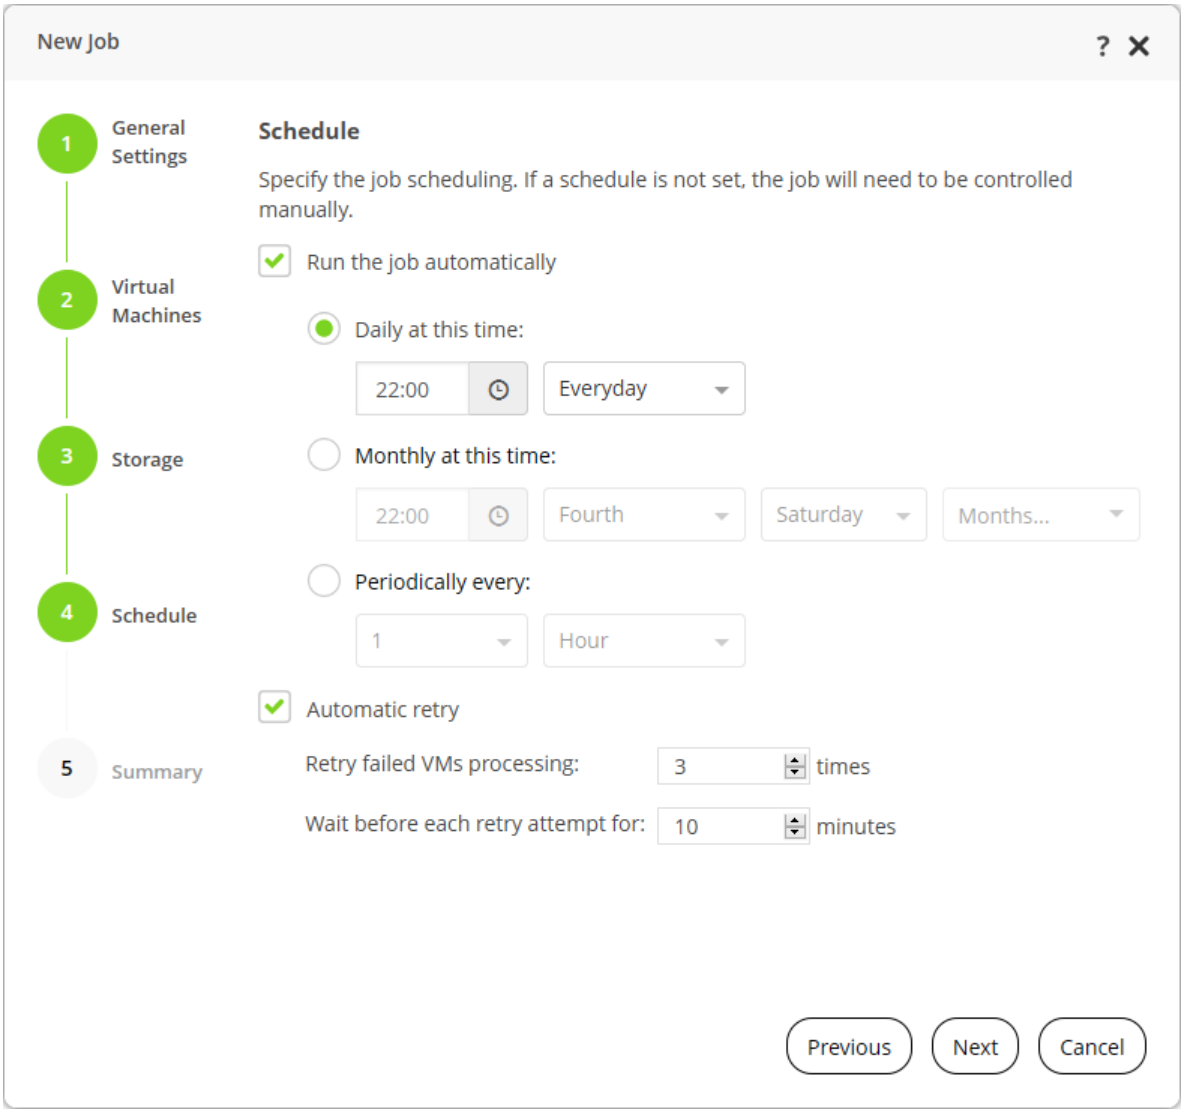

Step 5. Finish Working with Wizard

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

1. Review details of the PD snapshot job.
2. Select the **Run job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

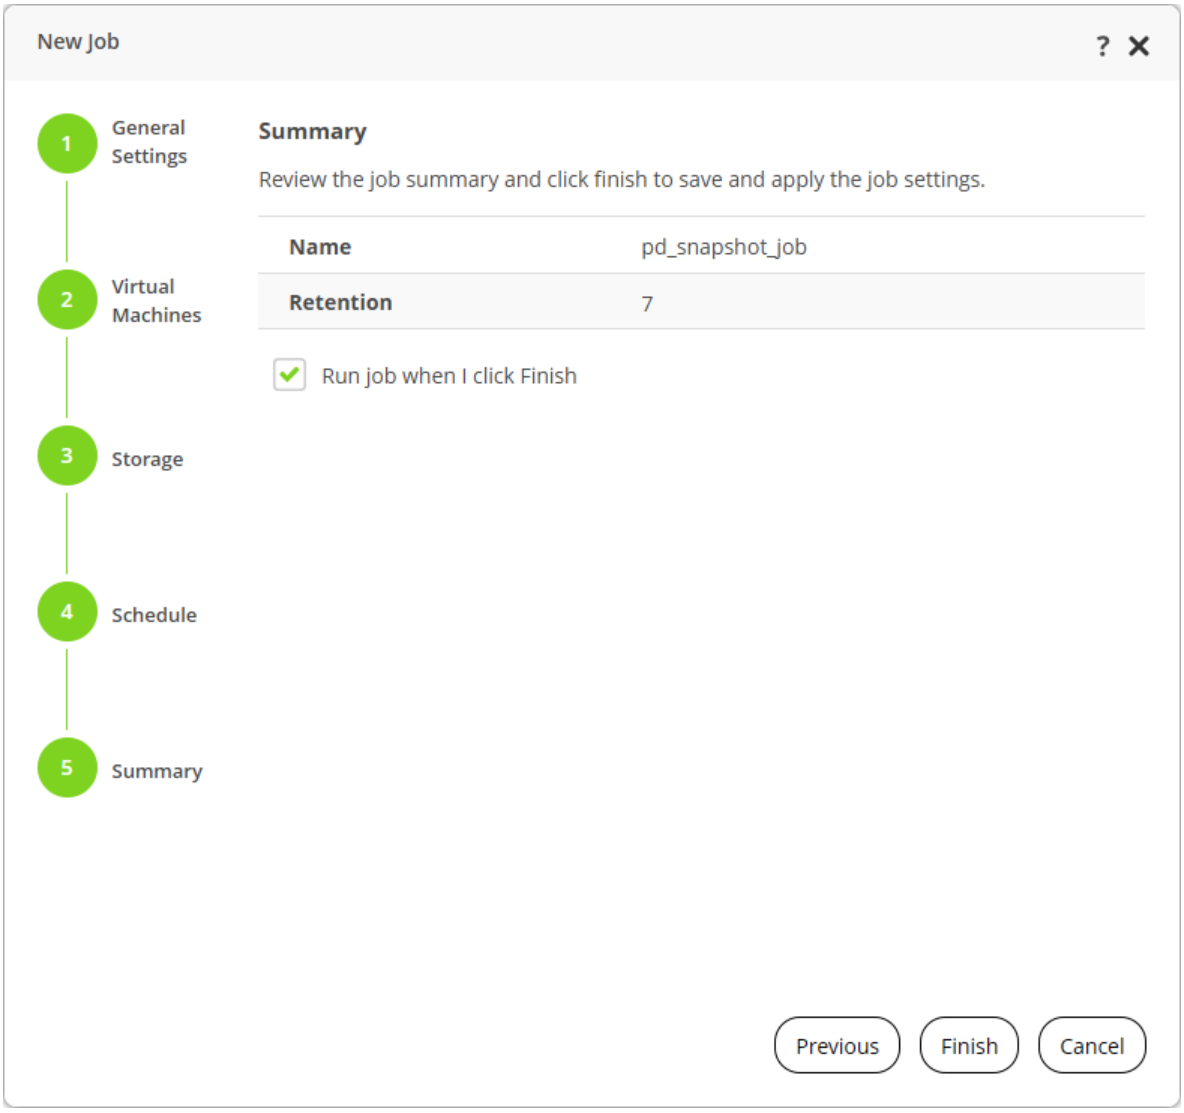

Creating Snapshot Jobs

In AHV Backup Proxy, you can create and schedule a snapshot job. AHV Backup Proxy creates snapshots of individual VMs selected in the job settings. Created snapshots are not transferred to the Veeam Backup & Replication repository. Snapshots are stored in the AHV container.

To create a snapshot job, launch the New Job wizard and configure the job settings:

1. Specify the job name and description.
2. Select VMs.
3. Configure snapshot retention settings.
4. Define the job schedule.
5. Finish working with wizard.

Launch New Job Wizard

To launch the **New job** wizard, do the following:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. Click **Add at the toolbar** to launch the **New job** wizard.

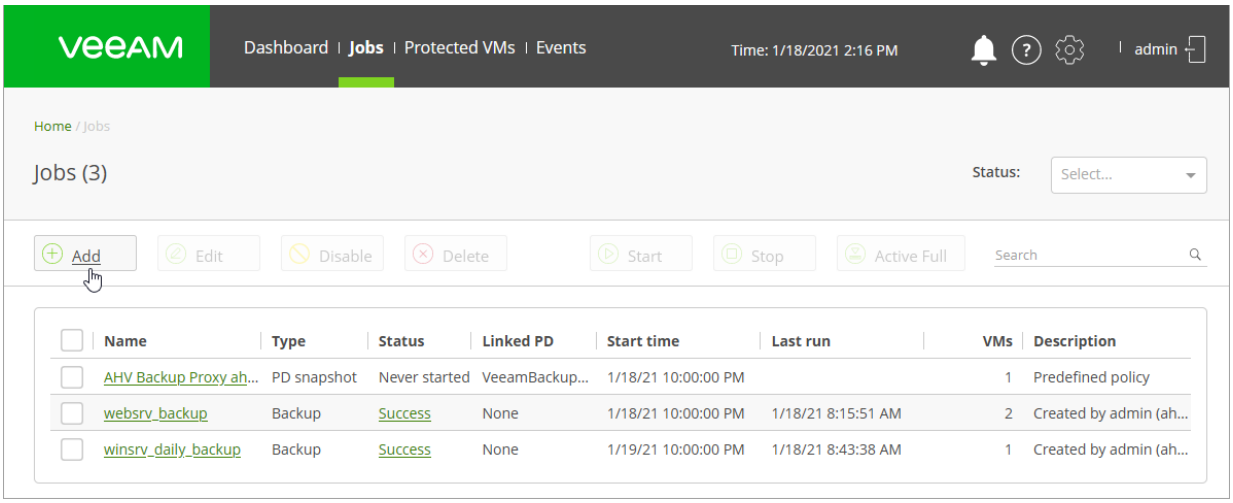

Step 1. Specify General Settings

At the **General Settings** step of the **New job** wizard, specify a name, description for the job and select the job type.

1. In the **Name** field, enter a name for the snapshot 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. In the **Select the job type** field, select **Snapshot job**.

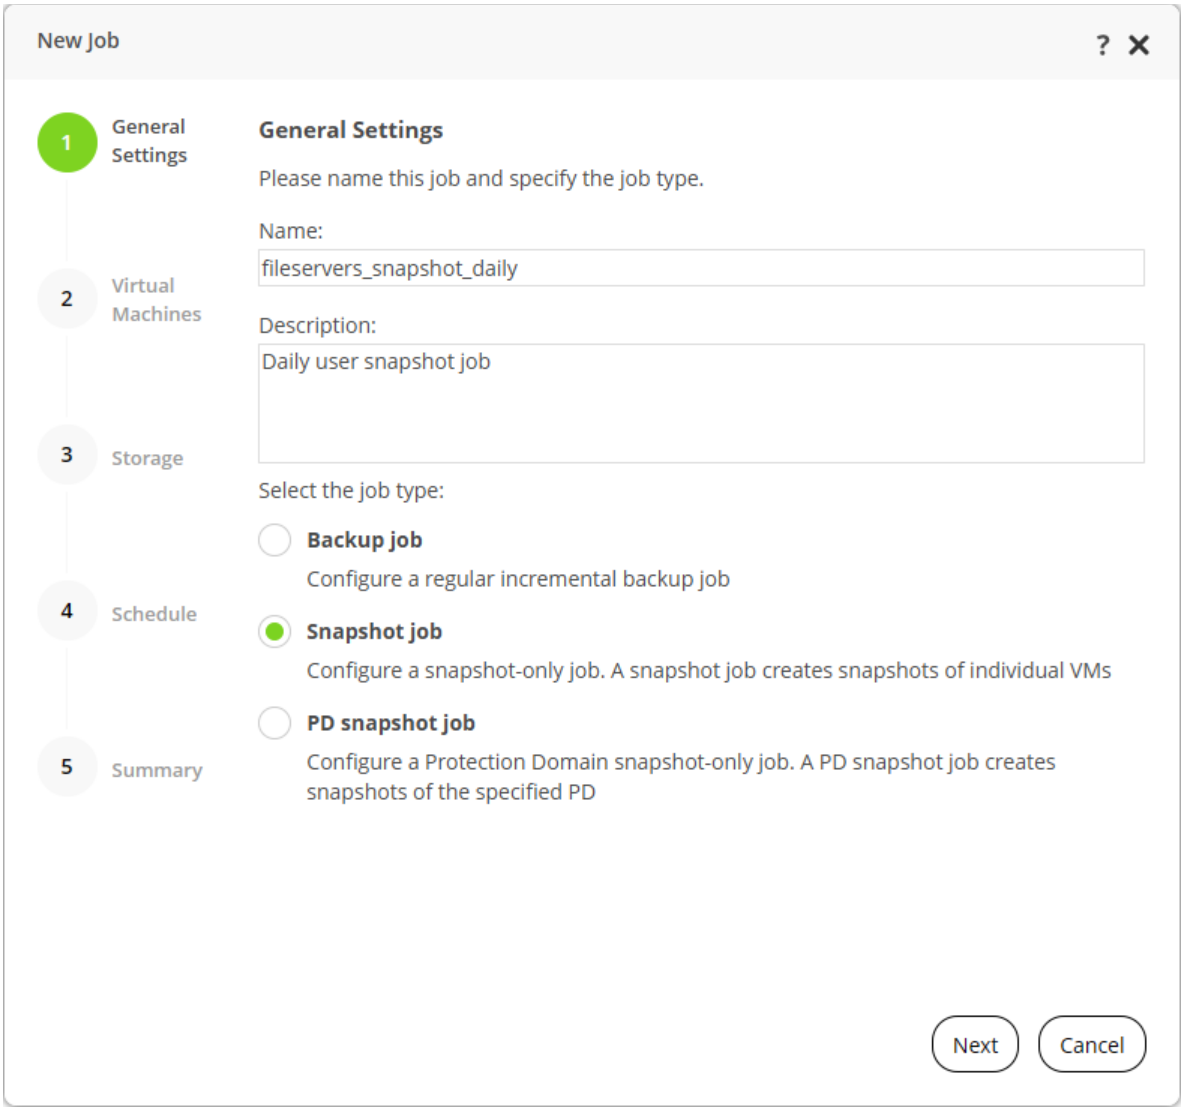

Step 2. Select VMs

At the **Assign VMs** step of the wizard, specify VMs or a protection domain that you want to add to the snapshot job.

> **NOTE**
>
> If you add a protection domain to a snapshot-only job, AHV Backup Proxy creates a snapshot for each VM not a PD snapshot. For more information, see Nutanix AHV Snapshot Types.

To learn how to add objects to a snapshot job, see:

- Adding VMs to snapshot job
- Adding Protection Domains to snapshot job

Adding VMs to snapshot job

To add VMs to a snapshot job:

1. At the **Virtual Machines** step of the wizard, click **Add on the toolbar**.
2. In the **Add Objects** window, click **+** on the left of the cluster name to expand the list of VMs.

3. From the list of VMs, select VMs that you want to back up and click **Add**.

   To quickly find the necessary VM, you can use the search field at the top of the **Select objects** field.

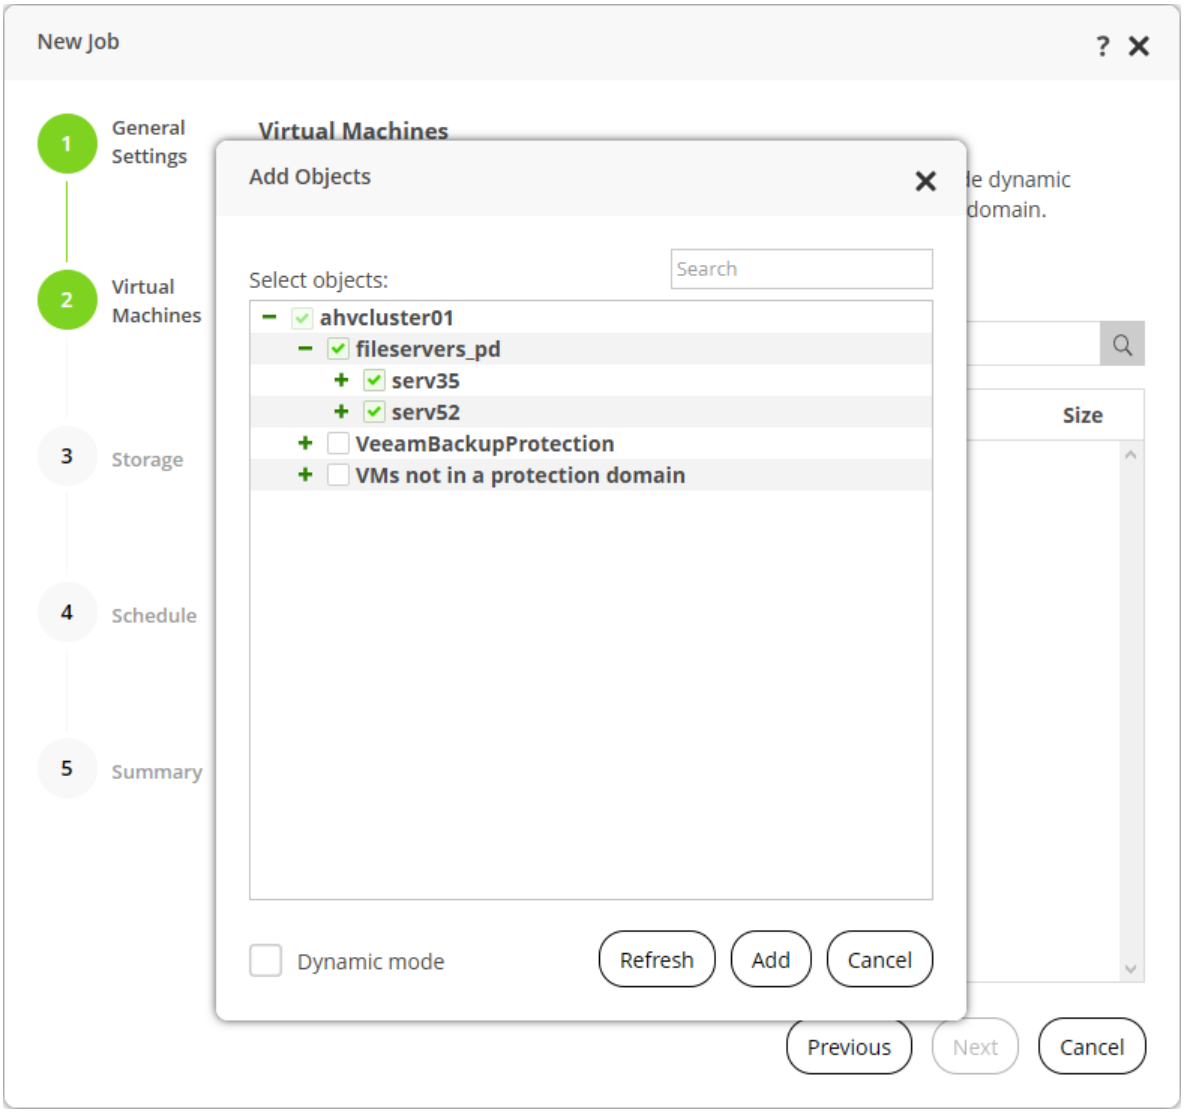

Adding Protection Domains to snapshot job

If you want to select a protection domain instead of individual VMs, do the following.

1. At the **Virtual Machines** step of the wizard, click **Add on the toolbar**.
2. In the **Add Objects** window, select the **Dynamic mode** check box.
3. Click **+** on the left of the cluster name to expand the list of protection domains.

4. From the list of protection domains, select one protection domain and click **Add**.

   To quickly find the necessary protection domain, you can use the search field at the top of the **Select objects** field.

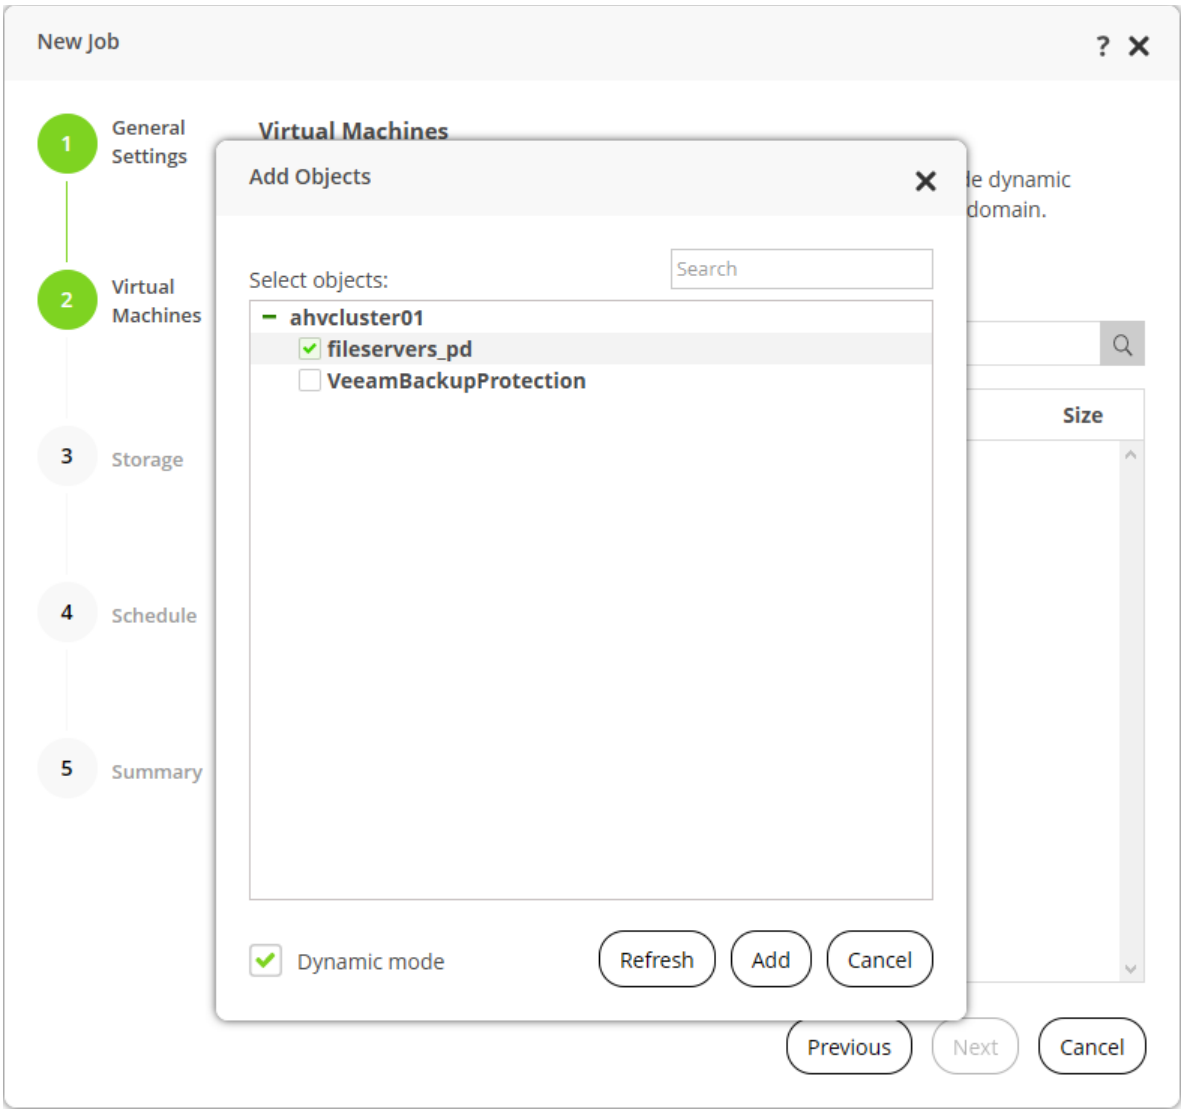

Step 3. Configure Retention

At the **Storage** step of the wizard, specify retention policy settings for restore points.

At the **Restore points to keep** field, specify the number of restore points. When the specified number is exceeded, the earliest restore point is removed from the snapshot chain. For more information, see Snapshot Chain and Retention Policy.

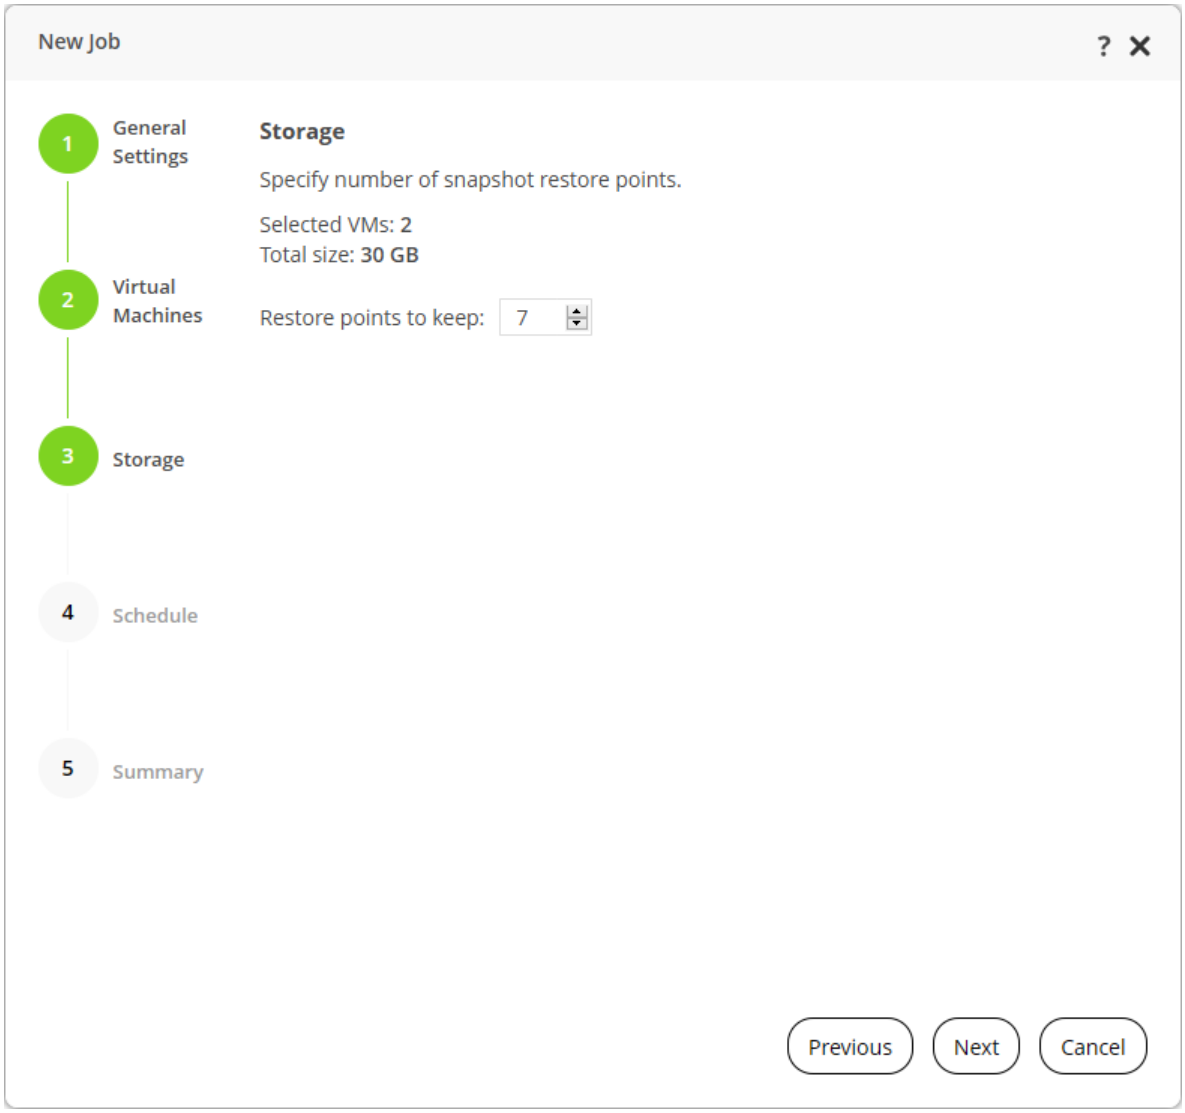

Step 4. Define Job Schedule

At the **Schedule** step, specify the job scheduling. If you don't set a schedule, you will have to launch the job automatically.

1. Select the **Run the job automatically** check box.
2. Define the 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 below to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields below to configure the necessary schedule.
   - To run the job repeatedly throughout a day with a specific time interval, select **Periodically every**. In the field below, select the time period and necessary time unit: *Hours* or *Minutes*.
3. In the **Automatic Retry** section, define whether AHV Backup Proxy must attempt to run the job again if the job fails for some reason. During a job retry, AHV Backup Proxy processes failed VMs only.

   Specify the number of retry attempts and the period of time between the attempts.

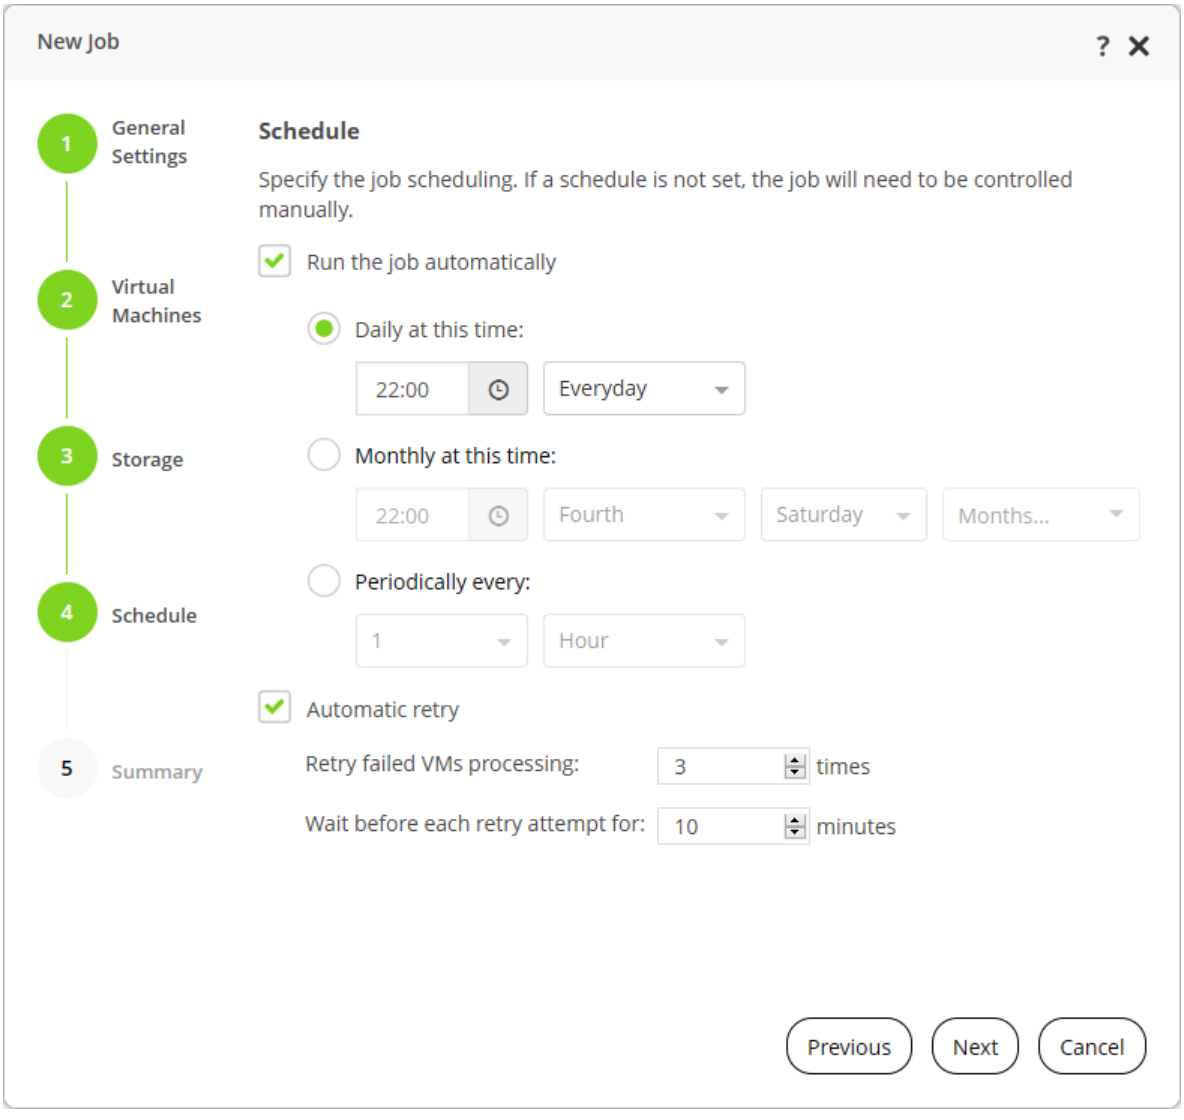

Step 5. Finish Working with Wizard

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

1. Review details of the snapshot job.
2. Select the **Run 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.

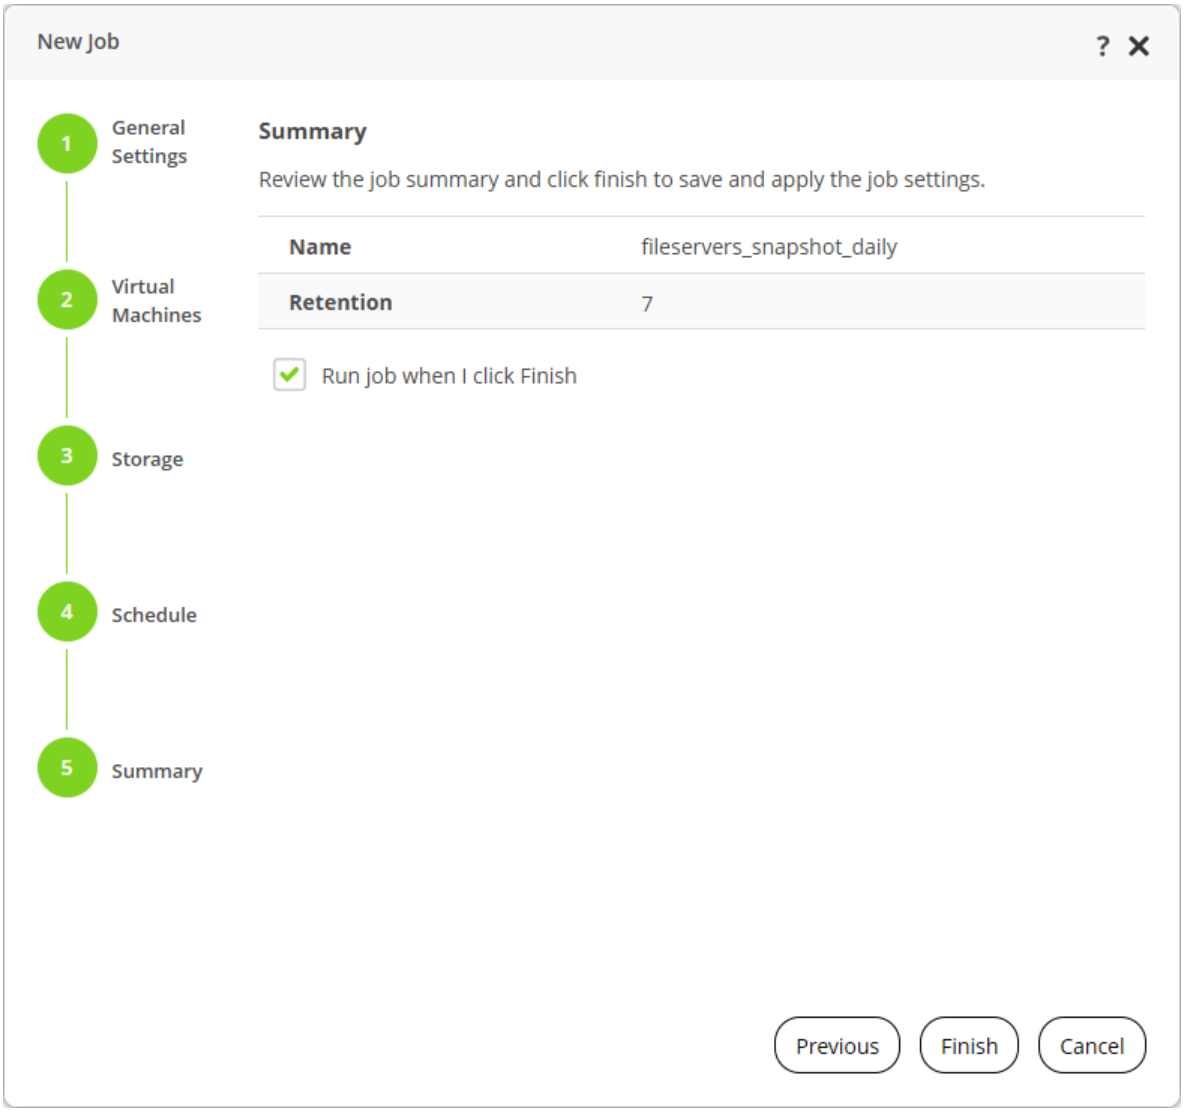

VeeamZIP

VeeamZIP functionality of Veeam Backup & Replication allows to quickly perform backup of one or several AHV VMs.

VeeamZIP is similar to a full VM backup. The VeeamZIP job always produces a full backup file (VBK) that acts as an independent restore point. You can store the backup file to a backup repository, to a local folder on the backup server or to a network share.

When you perform backup with VeeamZIP, you do not have to configure a backup job and schedule it. Instead, you can start the backup process for selected VMs immediately. This type of backup requires minimum settings — you should only select the backup destination.

Backup files produced with VeeamZIP jobs are displayed in the **Home** view, under the **Backups** > **Disk (VeeamZIP)** node. To restore VM data from VeeamZIP backups, you can right-click it in the **Home** view and select the necessary restore option. You can also double-click the necessary VeeamZIP backup file on the machine where Veeam Backup & Replication is installed.

To learn how to create VeeamZIP backups, see Creating VeeamZIP Backups.

> **IMPORTANT**
>
> Mind the following:
>
> - Veeam Backup & Replication does not enforce backup repository throttling rules during VeeamZIP jobs.
> - You cannot use a Veeam Cloud Connect repository as a target for VeeamZIP jobs.

Creating VeeamZIP Backups

To create a VeeamZIP backup, do the following:

1. In the Veeam Backup & Replication console, open the **Inventory** view and select Nutanix AHV at the inventory pane.
2. At the working area, select the required VM or VMs (press and hold **CTRL** to select several VMs) right-click and select VeeamZIP.

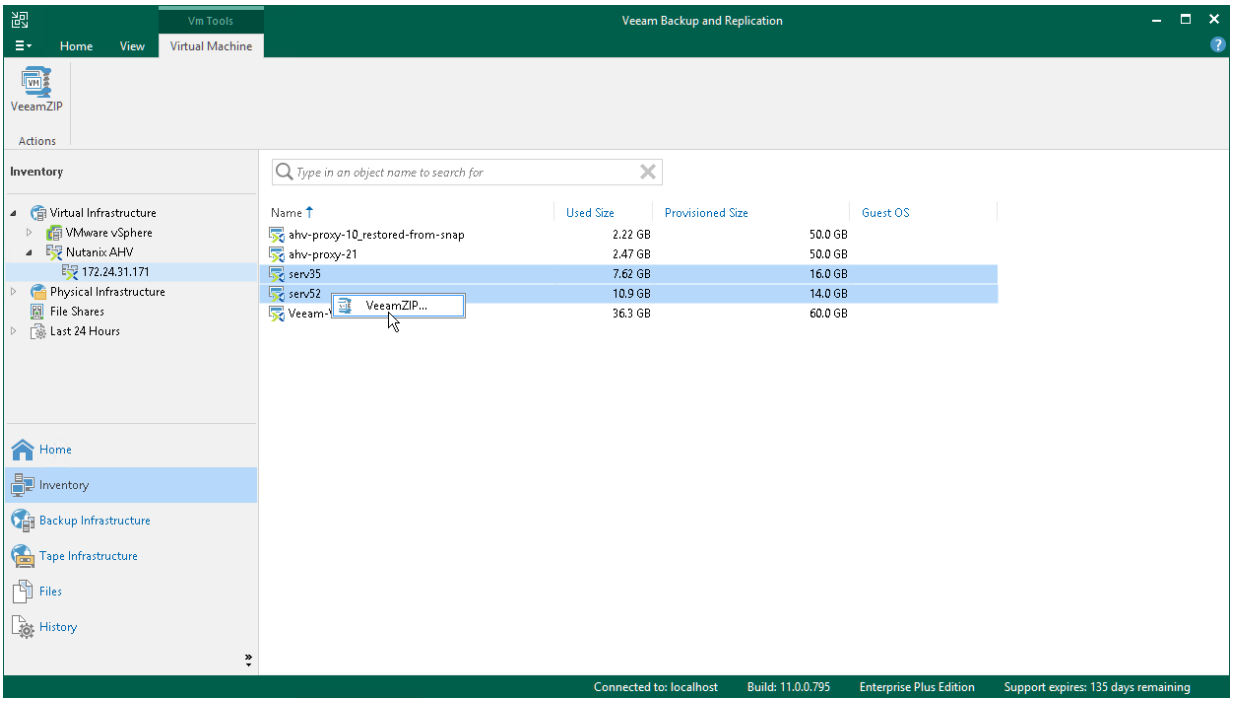

3. Select the repository where the VeeamZIP backup will be stored. You can select any repository added to Veeam Backup & Replication infrastructure or choose to store the VeeamZIP backup on a local or shared folder.

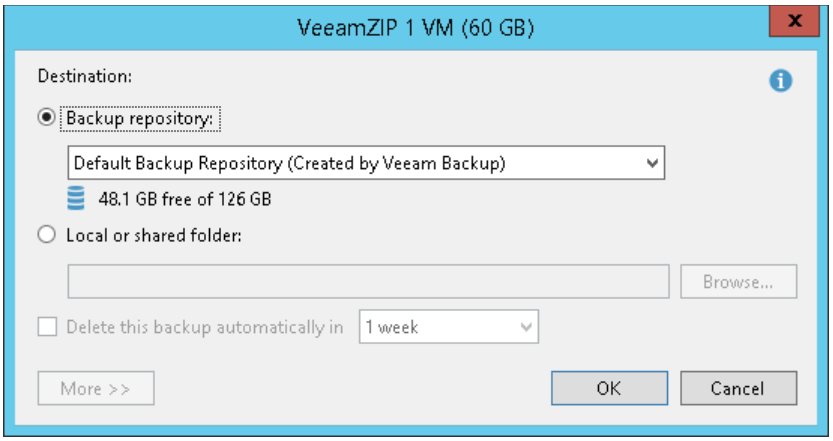

> **IMPORTANT**
>
> - VeeamZIP retention is not supported for AHV VMs.
> - In the **Destination** section of the VeeamZIP wizard, if you select the **Local or shared folder** option and specify an SMB share that requires authentication, VeeamZIP process fails. As a workaround, you can register the SMB share as a backup repository and, in the **VeeamZIP** wizard, select the **Backup repository** option, where you will be able to select the added SMB share repository.

4. Click **OK** and wait for the VeeamZIP job to finish.

   To see the VeeamZIP backup, open the **Home** view, and select **Disk (VeeamZIP)** on the inventory pane.

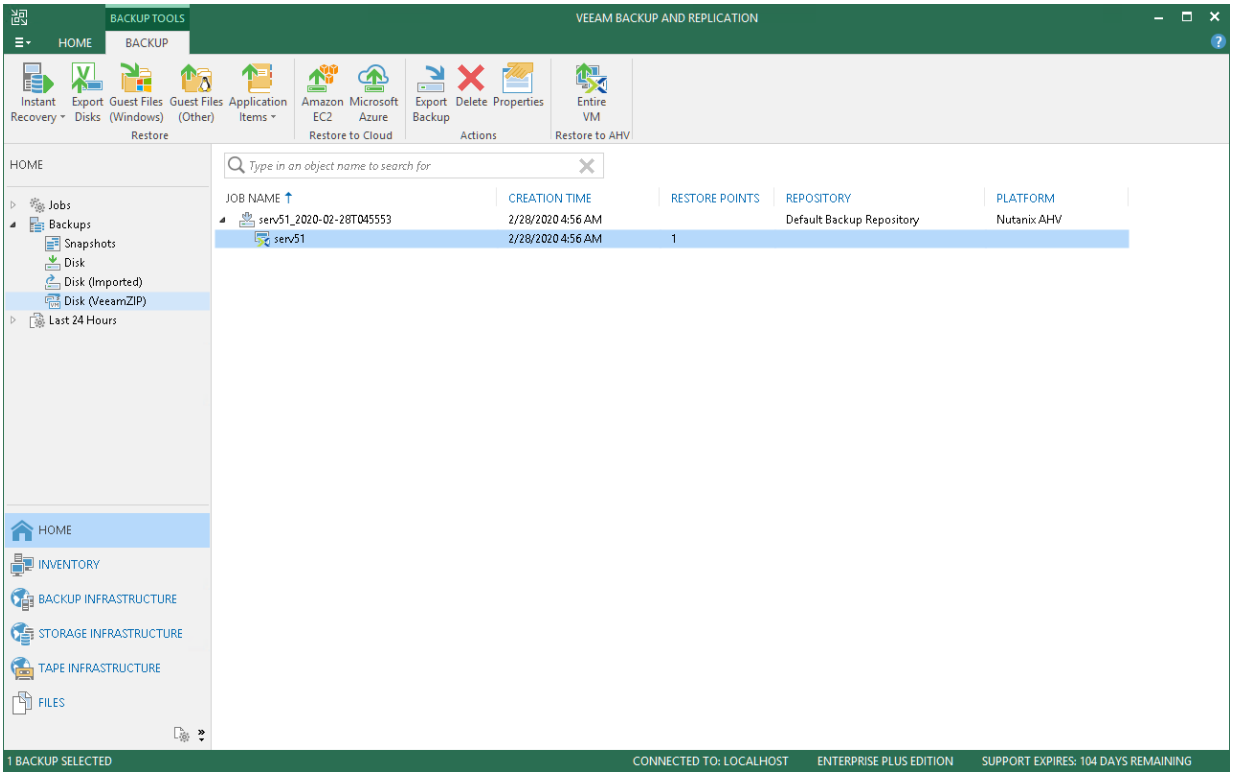

Backup Copy Jobs

Having just one backup does not provide the necessary level of safety. The primary backup may get destroyed together with production data, and you will have no backups from which you can restore data.

Backup experts advise that to build a successful data protection and disaster recovery plan, you must follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups.
- 2: You must use at least two different types of media to store the copies of your data (e.g. local disk and cloud).
- 1: You must keep at least one backup offsite (in the cloud or in a remote site).

Thus, you must have at least two backups and they must be in different locations. If a disaster takes out your production data and local backup, you can still recover from your offsite backup.

To help you adopt the 3-2-1 rule, Veeam Backup & Replication offers the backup copy functionality. Backup copy allows you to create several instances of the same backup data in different locations, whether onsite or offsite. Backup copies have the same format as those created by backup jobs and you can recover your data from them when you need it.

Backup copy is a job-driven process. Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes.

Creating Backup Copy Jobs

For instructions on how to create a backup copy job, see the Creating Backup Copy Jobs section of the Veeam Backup & Replication User Guide.

IMPORTANT

- The Immediate copy mode is not supported for backups of Nutanix AHV VMs.
- You cannot copy backups of Nutanix AHV VMs between HPE StoreOnce repositories. HPE StoreOnce Catalyst Copy is not supported for backups of Nutanix AHV VMs.

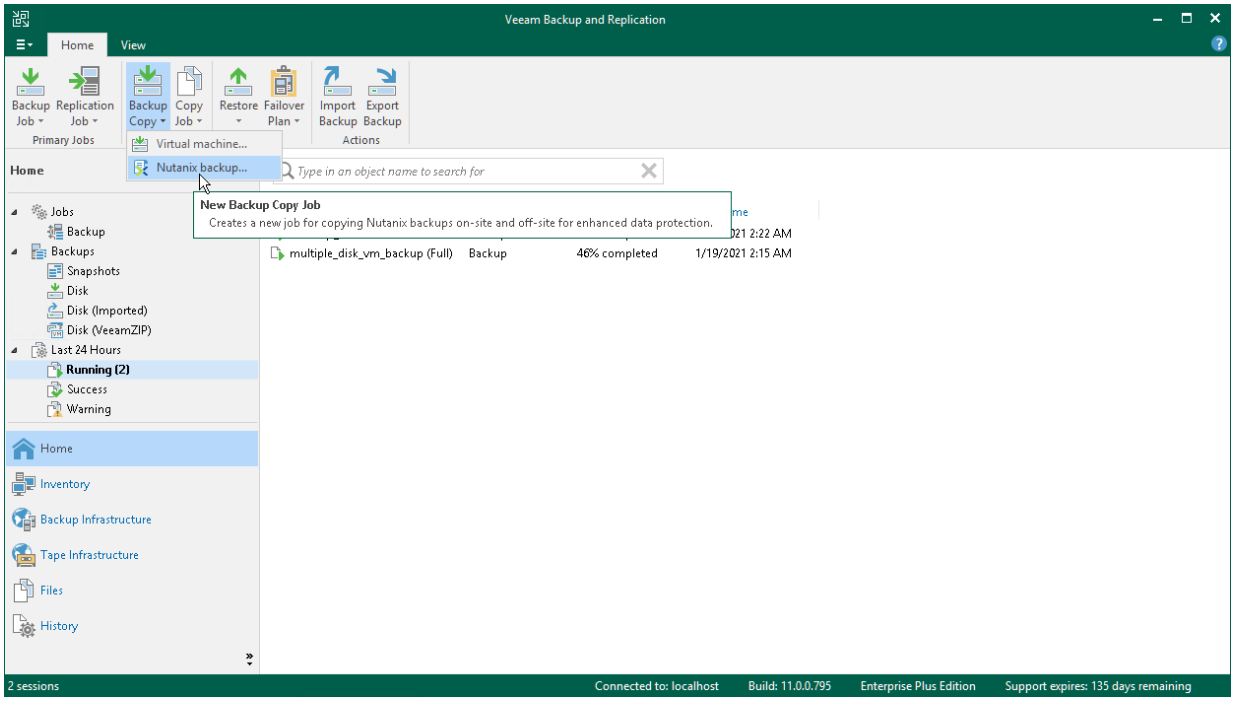

Backup to Tape Jobs

Veeam Backup & Replication allows to automate the process of copying AHV VM backups to tape devices and lets you specify scheduling, archiving and media automation options. For more information on how it works, see the How Machines Backup to Tape Works.

In the Veeam Backup & Replication console, you can copy your backups to a tape device. Before you can create backup to tape jobs, you must add the target tape server to the Veeam Backup & Replication infrastructure, create media pools, and configure media vaults. For instructions, see the following sections of the Veeam Backup & Replication User Guide:

- Adding Tape Servers
- Getting Started with Tapes
- Creating Backup to Tape Jobs

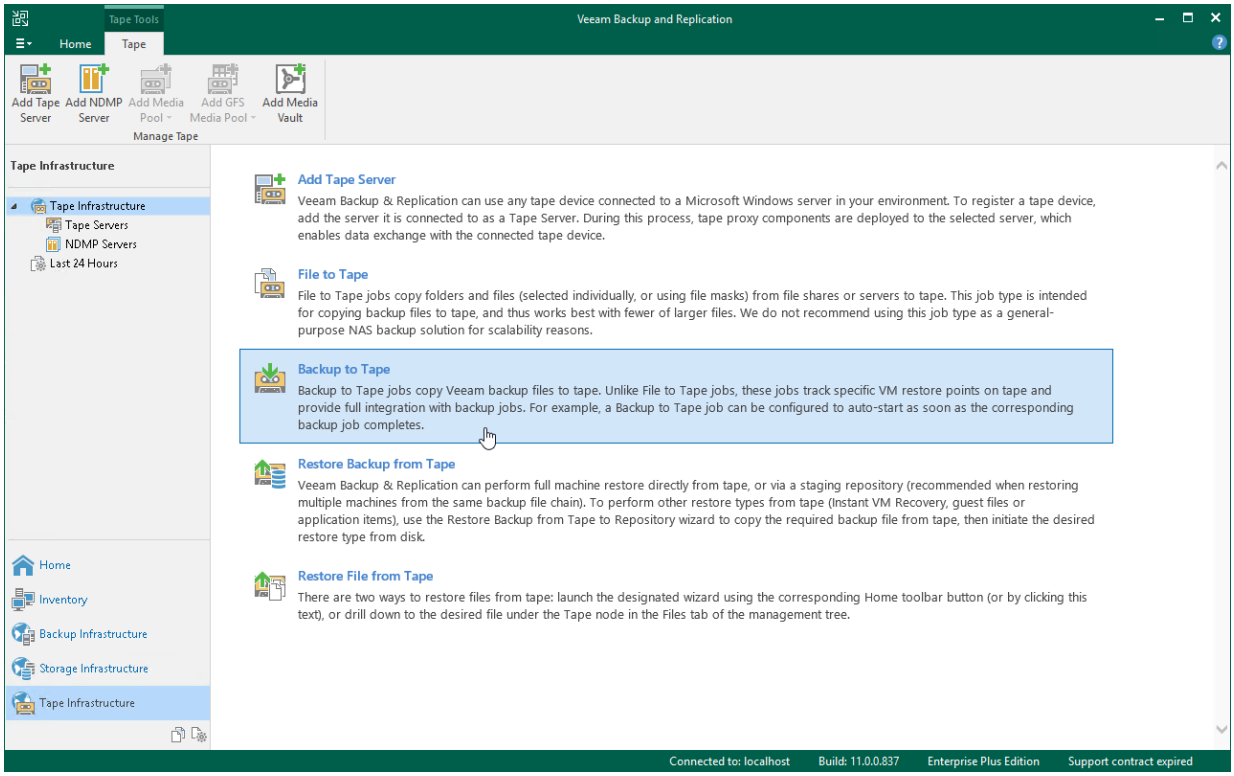

Managing Jobs

To view all jobs configured on AHV Backup Proxy, open the **Jobs** tab in the main menu of the web console. The list of available jobs is displayed in the working area. You can edit job properties, start and stop jobs, view job statistics and delete unnecessary jobs.

In this Section

- Starting Jobs Manually
- Stopping Jobs
- Editing Job Settings
- Deleting Jobs
- Disabling Jobs

Starting Jobs Manually

You can start a backup job manually. This can be necessary if you have not defined a schedule for the backup job, or if you want to create an additional restore point and do not want to change the schedule.

When you start a job manually, AHV Backup Proxy creates an incremental backup of VMs added to the job. If it is the first session of the job, AHV Backup Proxy creates a full backup.

Starting Jobs in AHV Backup Proxy Web Console

To start a job in the AHV Backup Proxy web console:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. In the list of jobs, select a job and click **Start** on the toolbar.

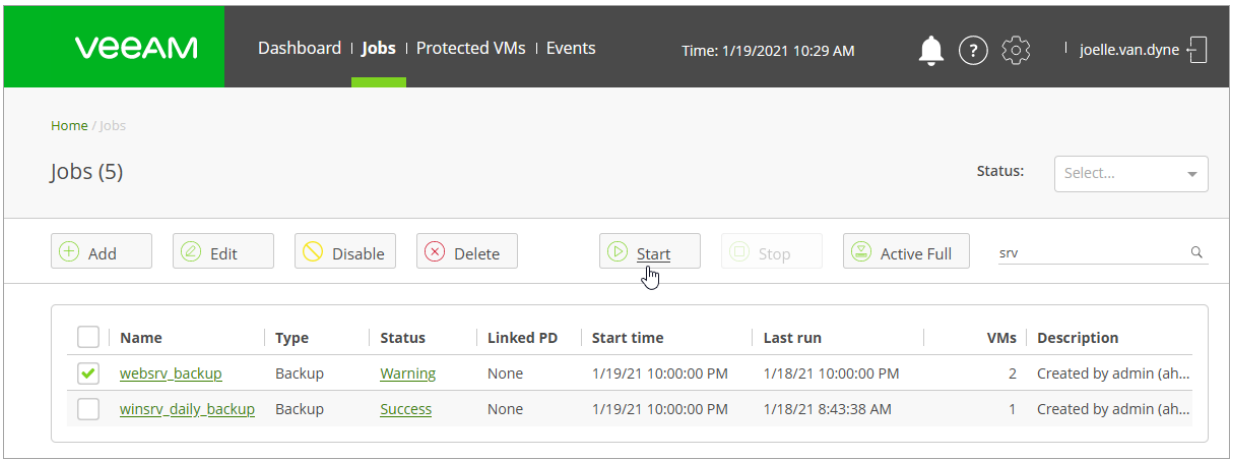

Starting Jobs in Veeam Backup & Replication Console

To start a job in the Veeam Backup & Replication console:

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

3. In the working area, select the job and click **Start** on the ribbon or right-click the job and select **Start**.

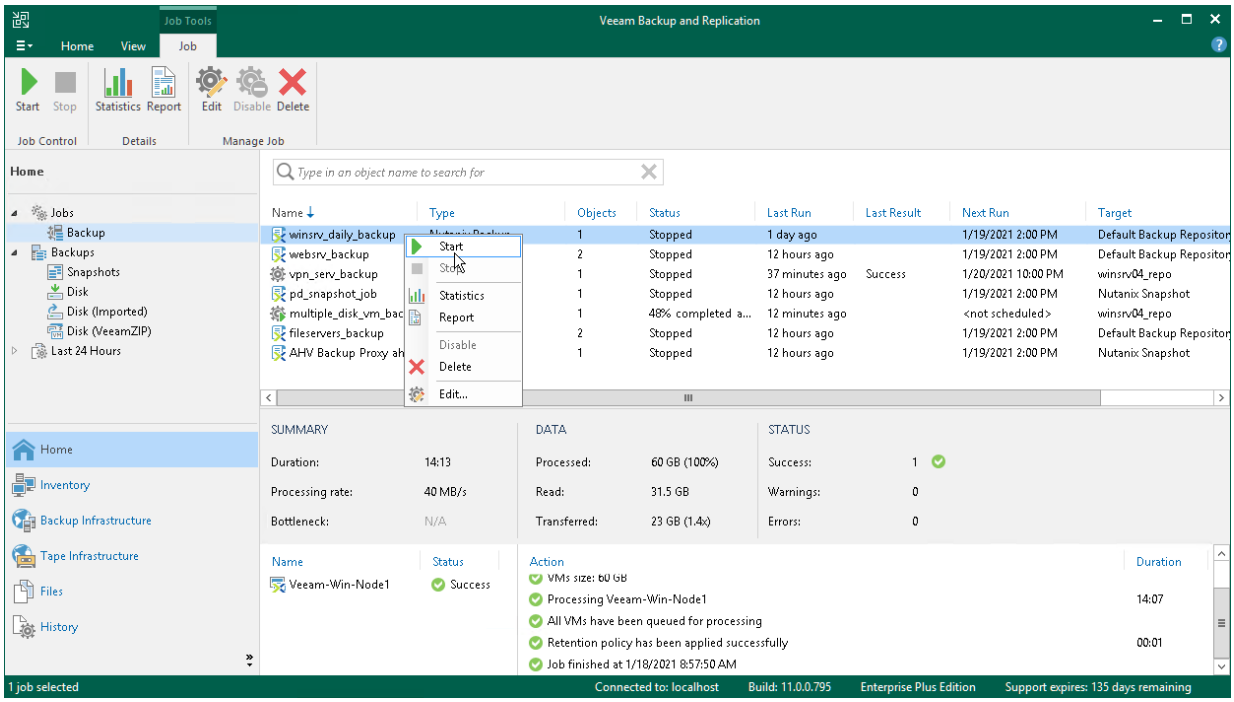

Stopping Jobs

You can stop a job, for example, if VM processing is about to take long, and you do not want the job to produce workload on the production environment during business hours.

When you stop a running job, AHV Backup Proxy will produce a new restore point only for those VMs that have already been processed by the time you stop the job.

Stopping Jobs in AHV Backup Proxy Console

To stop a job in the AHV Backup Proxy web console:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. In the list of jobs, select a running job that you want to stop and click **Stop** on the toolbar.
3. After you click **Stop**, AHV Backup Proxy will issue a warning that one or more VMs will have incomplete, unusable restore points.
4. In the displayed window, click **OK**.

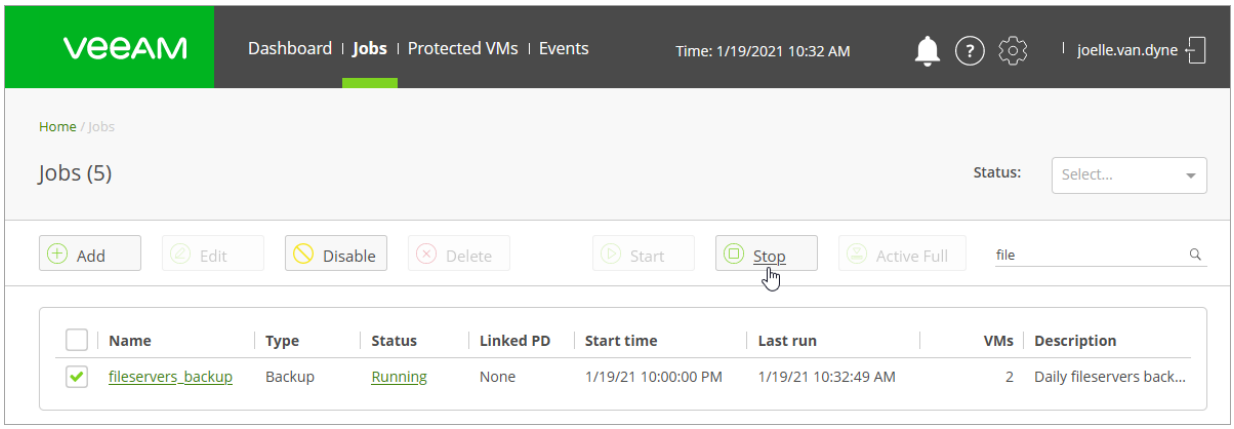

Stopping Jobs in Veeam Backup & Replication Console

To stop a job in the Veeam Backup & Replication console:

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

3. In the working area, select the job and click **Stop** on the ribbon or right-click the job and select **Stop**.

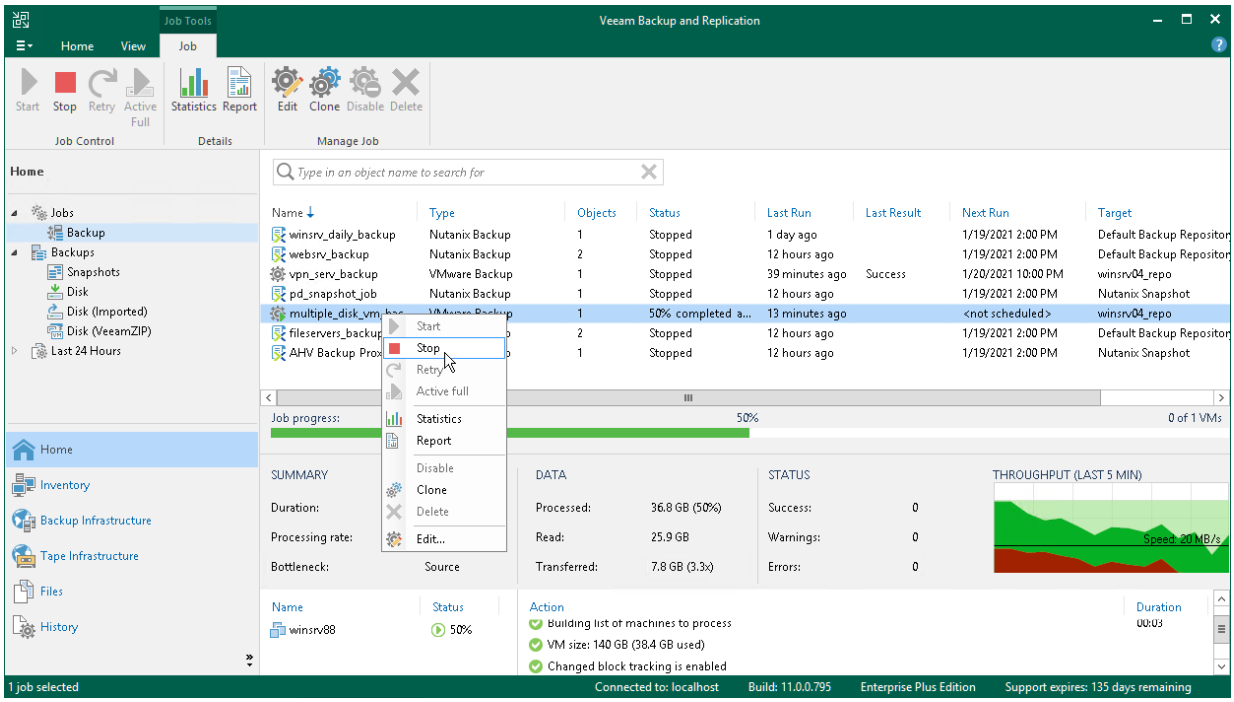

Editing Job Settings

You can edit created jobs at any moment. For example, you may want to change scheduling settings for a job or add some VMs to a job.

To edit job settings:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. In the list of jobs, select a job that you want to edit and click **Edit** on the toolbar.

   AHV Backup Proxy will open the **Edit Backup Job** wizard. Follow the same steps as you have followed when creating the job and change job settings as required.

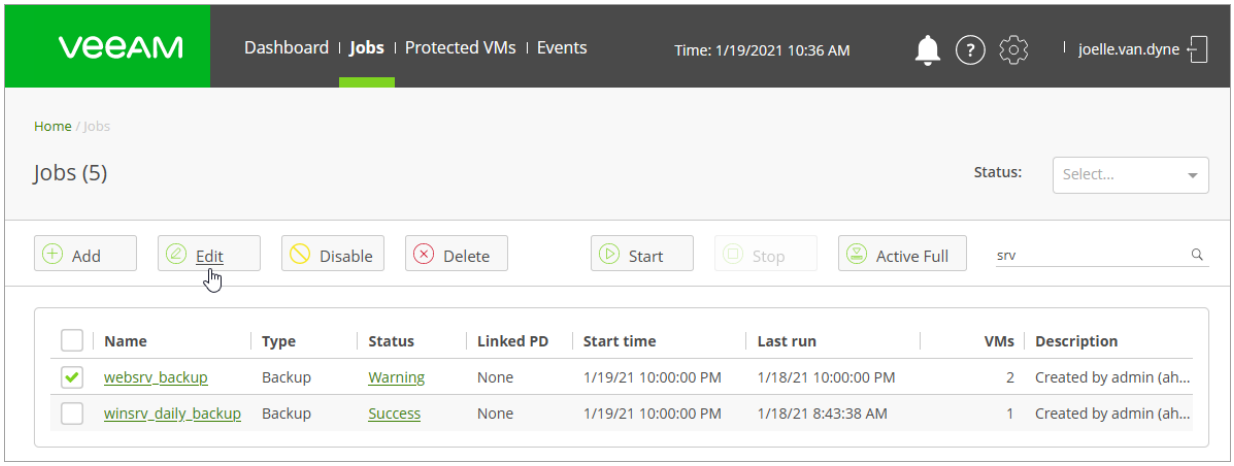

NOTE

In the Veeam Backup & Replication console, if you right-click the job and select **Edit**, Veeam Backup & Replication will open the job settings in the AHV Backup Proxy web console.

Deleting Jobs

You can permanently delete a job from AHV Backup Proxy and from the configuration database. When you delete a job, backup files created by the job are not deleted. Backup files remain in the repository and will be tagged as "imported".

If you want to delete backup files from the disk as well, see Managing Backups and Snapshots.

Deleting Jobs in AHV Backup Proxy Web Console

To delete a job:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. In the list of jobs, select a job and click **Delete** on the toolbar.

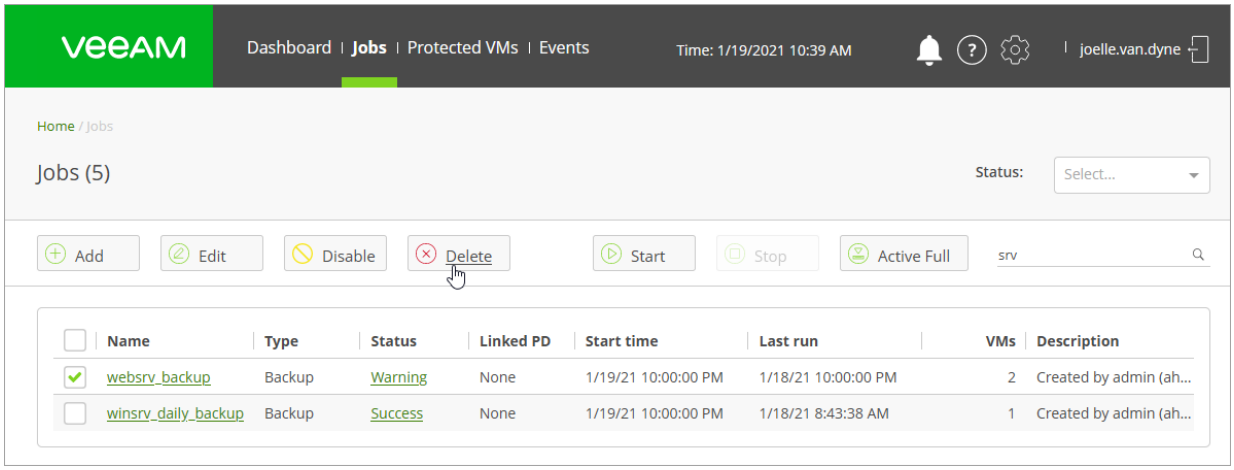

Deleting Jobs in Veeam Backup & Replication Console

To delete a job:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.

3. In the working area, select the job and click **Delete** on the ribbon or right-click the job and select **Delete**.

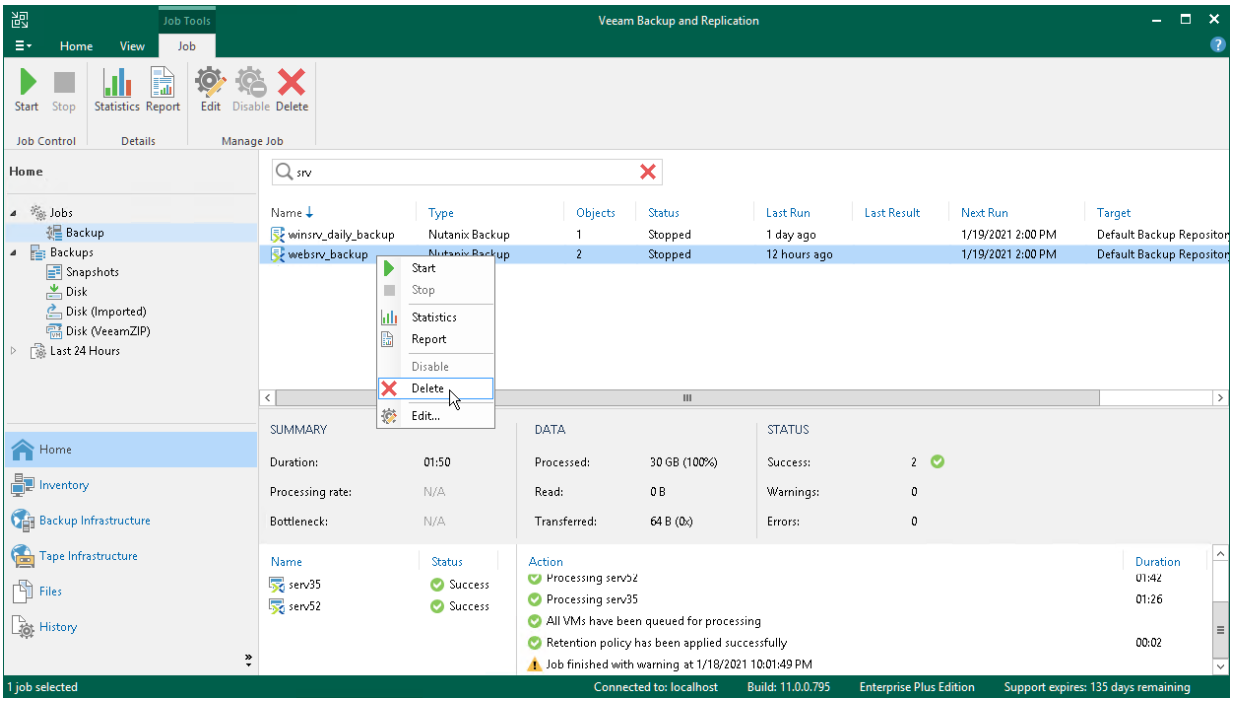

Disabling Jobs

You can temporary disable scheduled jobs. The disabled job is not deleted from AHV Backup Proxy, it is paused and is not run by the specified schedule. You can enable a disabled job at any time.

To disable a job, use one of the following guides.

- Disabling Jobs in AHV Backup Proxy Web Console
- Disabling Jobs in Veeam Backup & Replication Console

Disabling Jobs in AHV Backup Proxy Web Console

To disable a job:

1. At the main panel of the AHV Backup Proxy web console, click the **Jobs** tab.
2. In the list of jobs, select a job and click **Disable** on the toolbar.

To enable a disabled job, select it in the list and click **Enable** on the toolbar.

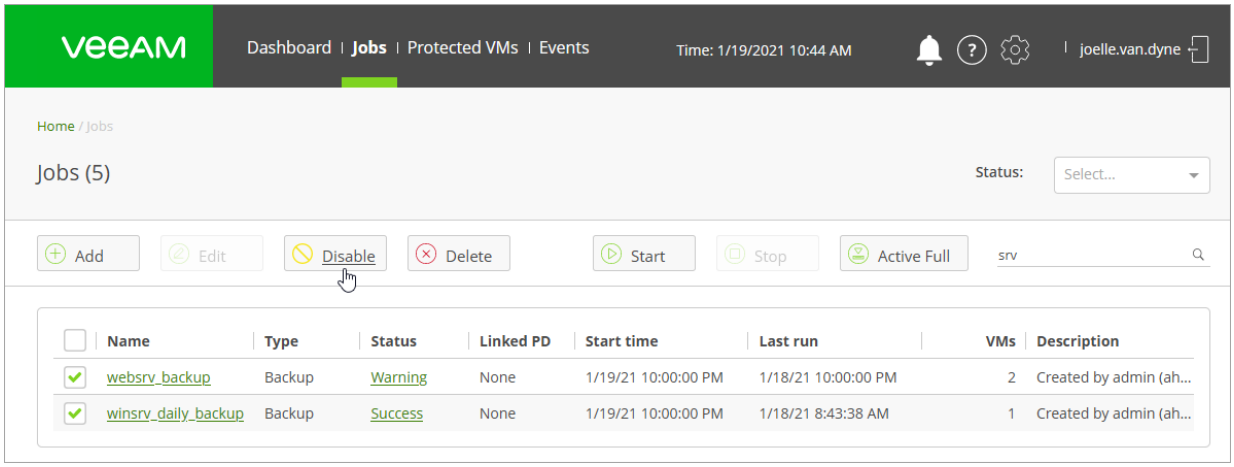

Disabling Jobs in Veeam Backup & Replication Console

To disable a job:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.

3. In the working area, select the job and click **Disable** on the ribbon or right-click the job and select **Disable**.

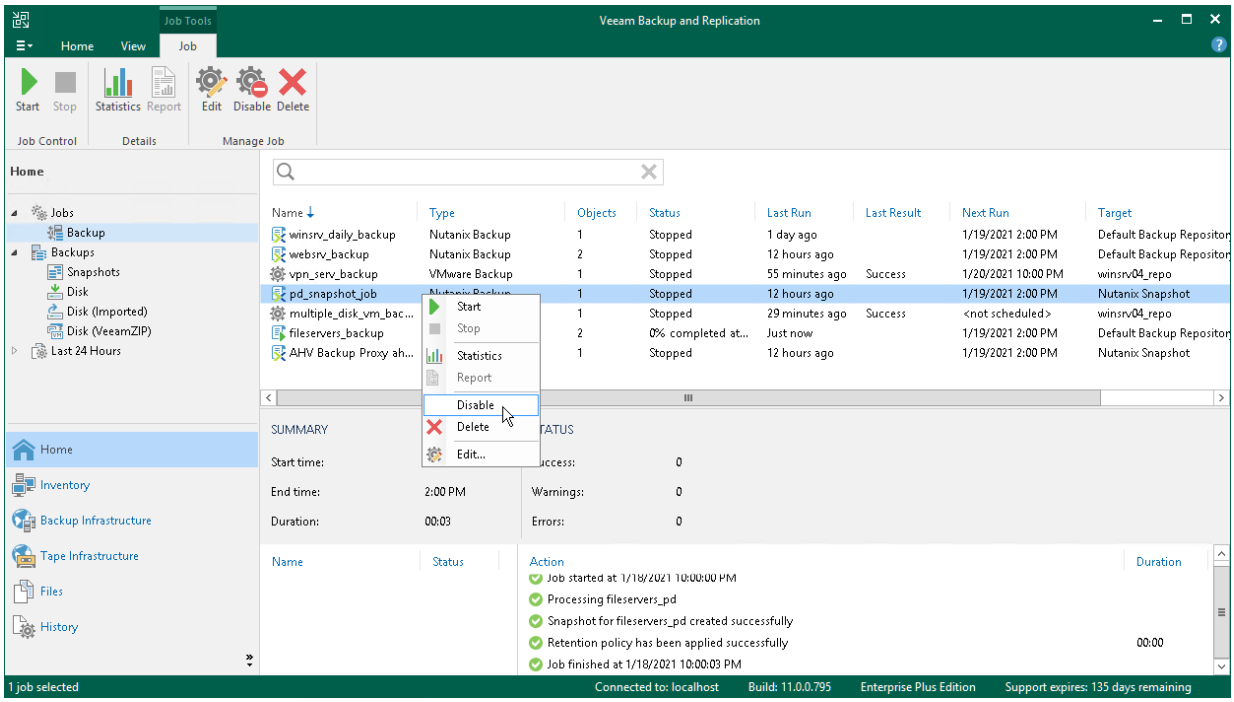

Managing Backups and Snapshots

You can perform the following operations with backup files and snapshots:

- Viewing Backup Properties
- Importing Backups
- Exporting Backups
- Deleting Backups
- Deleting Snapshots

Viewing Backup Properties

In the Veeam Backup & Replication console, you can view summary information about created backups. The summary information provides the following data:

- Repository and folder where the backup is stored
- Available restore points
- Date of restore points creation
- Compression and deduplication ratios
- Data size and backup file size

To view summary information for backup files:

1. In the Veeam Backup & Replication console, 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.

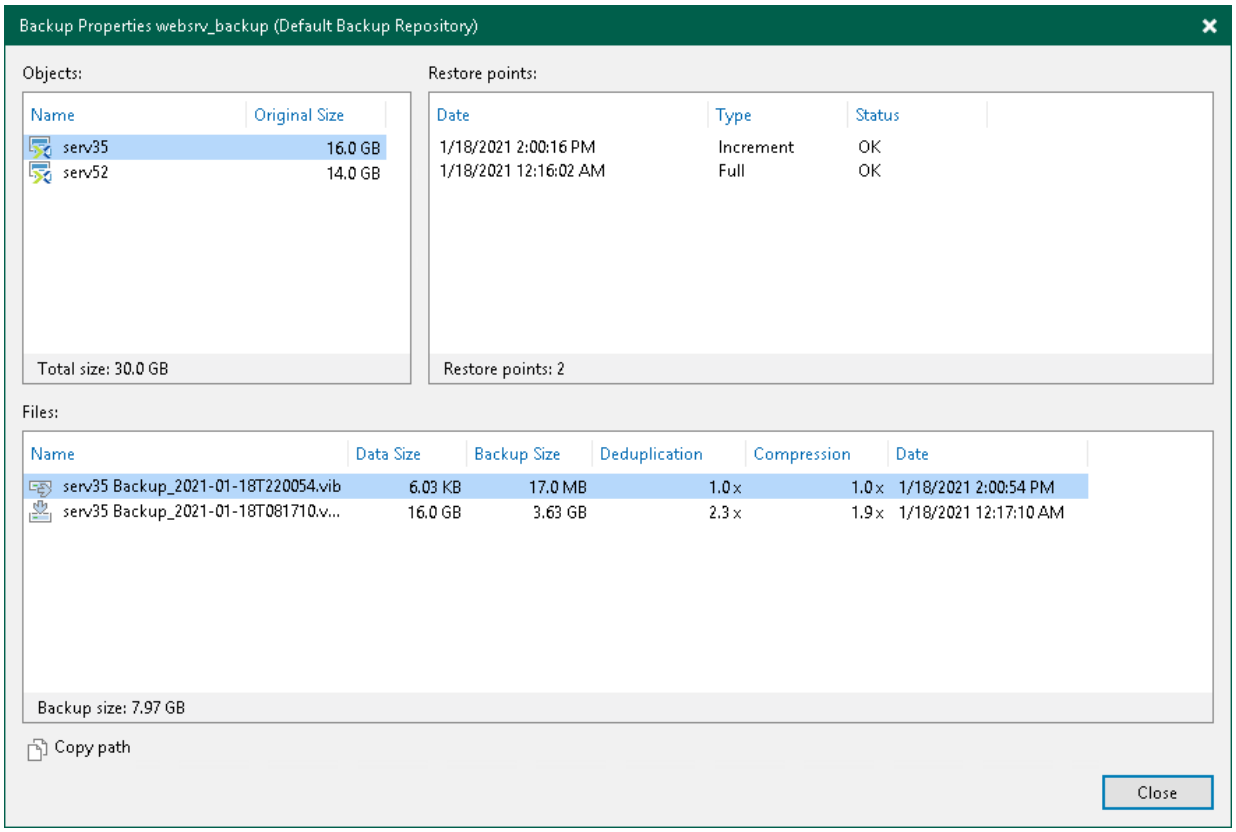

Importing Backups

If you have backups of Nutanix AHV VMs created by another instance of AHV Backup Proxy, you can import these backup files to the current instance. Once imported, you can use the backups to restore VMs and restore VM disks.

> **NOTE**
>
> You cannot import backups from unsupported backup repository types. This can affect imports from backup copy jobs. For the list of supported backup repository types, see Configure Backup Repository.

To import Nutanix AHV VM backups, do the following:

1. Transfer the backups to a required backup repository. Note that AHV Backup Proxy must have access to this repository.
2. In the Veeam Backup & Replication console, rescan the repository. For more information, see the Rescanning Backup Repositories section of the Veeam Backup & Replication User Guide.
3. At the main panel of the AHV Backup Proxy console, click the settings icon and select **Manage Veeam Server**.

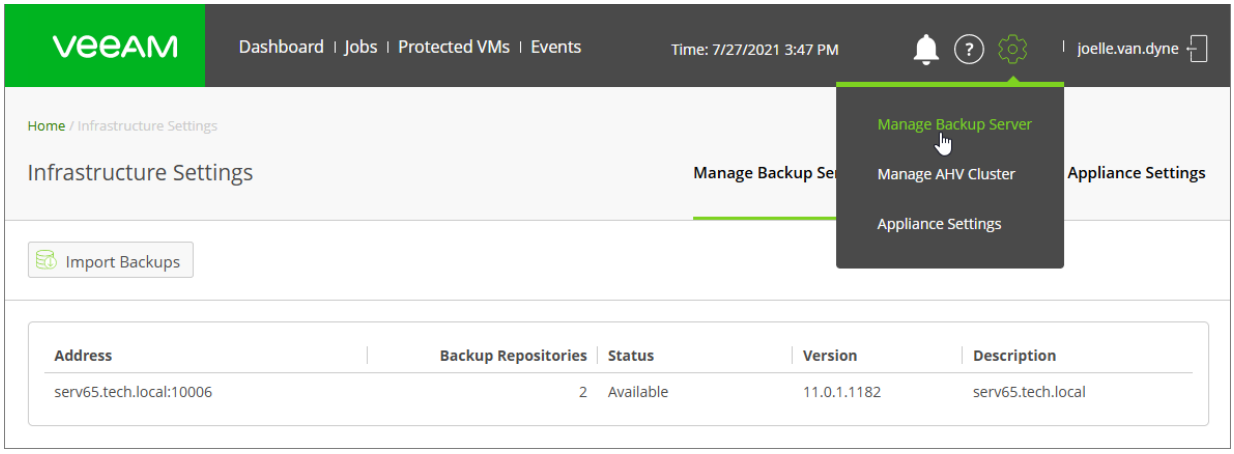

4. Select the required Veeam Backup & Replication server from the list of backup servers.
5. Click the **Import backups** button and click **Proceed** to confirm the action.

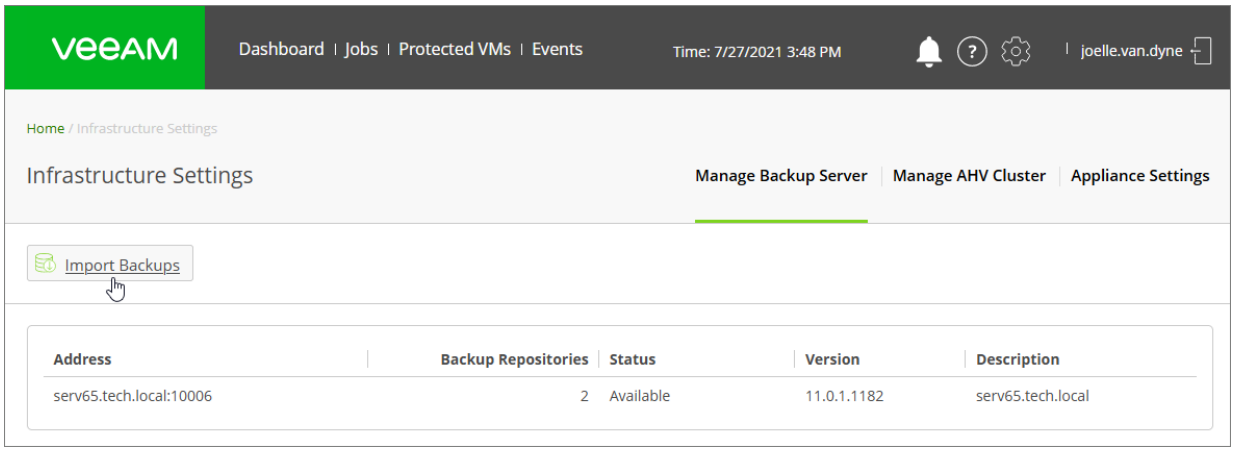

AHV Backup Proxy will scan the backup repositories and import all compatible backups of AHV VMs.

6. Once the backups are imported, click the **Protected VMs** at the main panel of the web console. You will see imported backups labeled with the icon.

   If you want to see the status of the import, click **Events** on the main menu.

> **TIP**
>
> When you import backups from a backup repository, AHV Backup Proxy rescans the repository and removes information about manually deleted backups from the database.

Exporting Backups

Exporting backups allows you to synthesize a complete and independent full backup file out of selected restore points that are located in your backup repositories. That means, you can transform any incremental backup chain (all dependent .vbk, .vib files) into a standalone .vbk file.

Export applies to Full and Incremental restore points located in backup repositories and object storage repositories.

For instructions on how to export backups, see the Performing Export section of the Veeam Backup & Replication User Guide.

Mind the following:

- The restore point that is being exported as a new full backup file is saved to the same repository, wherein the selected restore point resides.
- Once export is complete, the exported backup files will be attached under the **Backups > Disk (Imported)** node.
- If a restore point that is being exported resides on the tenant side, a new full backup file will also be exported to the same repository (on the tenant side) from which the source restore point is being taken.
- If you select a backup job consisting of multiple virtual machines, Veeam Backup & Replication will synthesize a separate full backup file per each machine.
- Export session results are saved to the configuration database and available for viewing. For more information, see the Viewing Session Statistics section of the Veeam Backup & Replication User Guide.

Deleting Backups

If you want to delete records about backups from the Veeam Backup & Replication configuration database and, additionally, delete backup files from the backup repository, you can use the **Delete from disk** operation. When you delete backup files from a disk, Veeam Backup & Replication deletes the whole chain from the backup repository. Thus, on the next run of the backup job, Veeam Backup & Replication will create full backups for VMs included in the job.

Mind the following:

- Do not delete backup files from the backup repository manually. Use the **Delete from disk** option instead. If you delete backup files manually, subsequent backup or replication job sessions will fail.
- Veeam Backup for Nutanix AHV supports backup jobs per-VM. You can perform the **Delete from disk** operation for separate VMs in the backup. If you delete backup files for one VM, on the next run of the backup job Veeam Backup & Replication will create a full backup for VMs whose backup files are deleted. For all other VMs, Veeam Backup & Replication will create increments.

  To learn more about per-VM backup files, see the Per-VM Backup Files section of the Veeam Backup & Replication User Guide.

To delete backup files from the backup repository, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane of the **Home** view, select **Backups**.
3. In the working area, select the backup or separate VM in the backup and click **Remove from** > **Disk** on the ribbon. You can also right-click the backup and select **Delete from disk**.

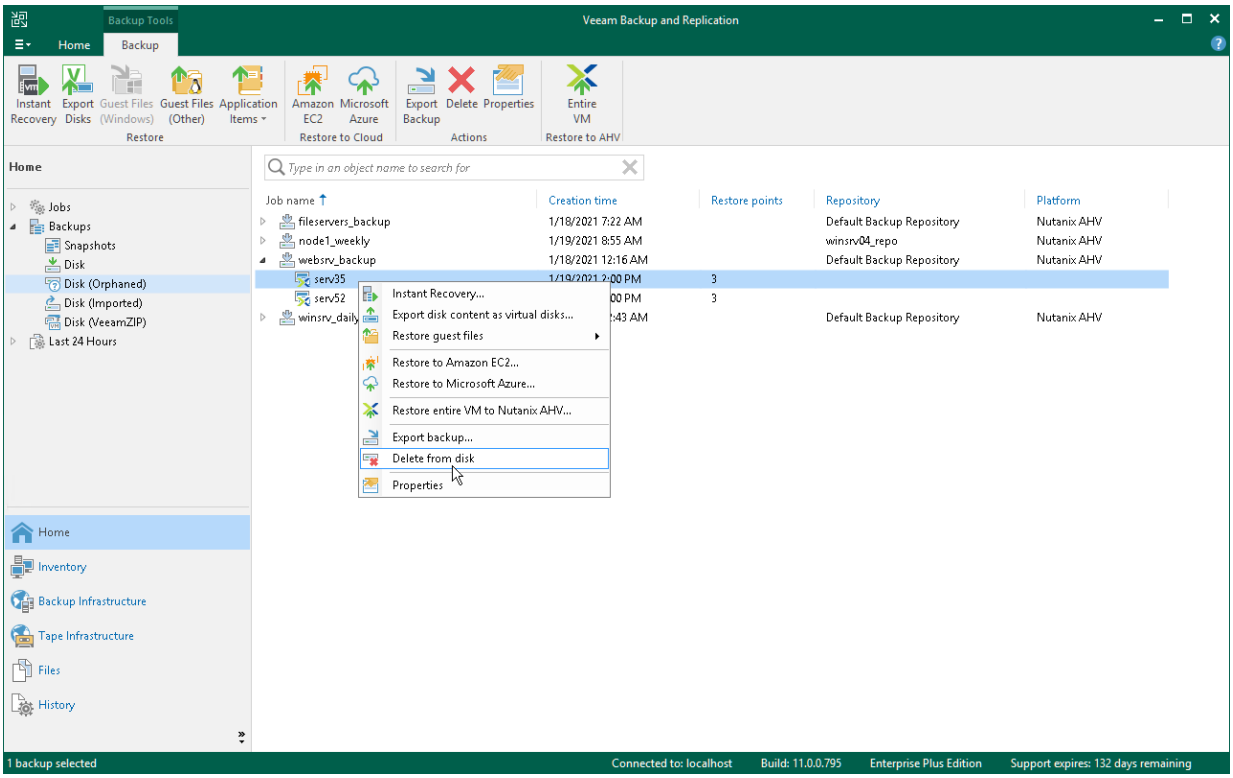

Deleting Snapshots

You can use the AHV Backup Proxy web console to remove snapshots from the AHV cluster. When you remove snapshots from the AHV cluster, AHV Backup Proxy automatically deletes records about these snapshots from the Veeam Backup & Replication configuration database.

> **TIP**
>
> You cannot remove AHV VM snapshots using the Veeam Backup & Replication console.

To remove a snapshot, do the following:

1. At the main panel of the AHV Backup Proxy web console, click the **Protected VMs** tab.
2. At the right of the working area, click the **View as tree** toggle to enable the tree view.
3. In the **Job name** list, expand a necessary snapshot job.
4. Select check boxes next to snapshots that you want to delete.
5. Click **Delete** on the toolbar.

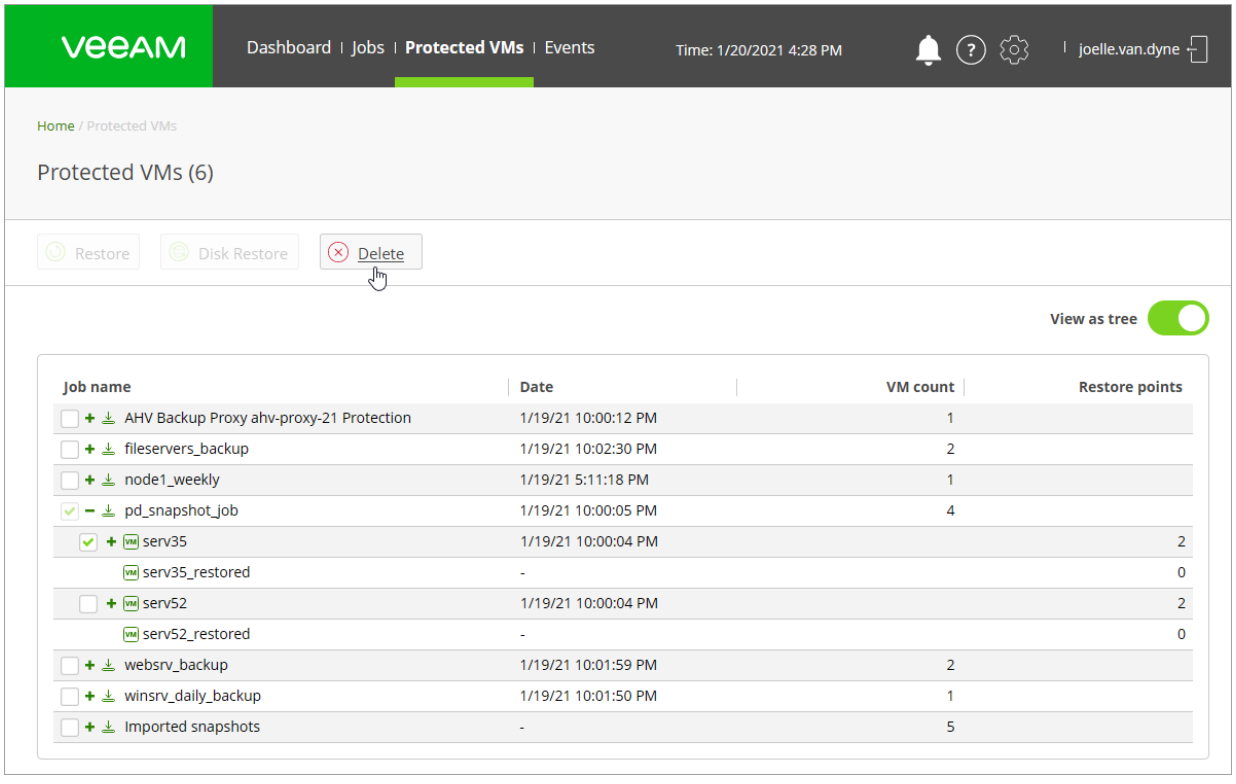

Data Recovery

Veeam Backup for Nutanix AHV offers the following recovery options for various disaster recovery scenarios:

- Entire VM Restore: recover a VM from a backup file or from a snapshot to an AHV cluster. You can restore from backups of VMware vSphere, Hyper-V, AHV VMs or from backups created by Veeam Agent for Windows or Veeam Agent for Linux.
- VM Disks Restore: recover a specific hard drive of a VM from backup or backup snapshot, and attach it to the original VM or to a new VM.
- Instant VM Recovery: instantly start a VM directly from a backup file.
- VM Disk Export: restore VM disks from backups created by AHV Backup Proxy and convert them to disks of the VMDK, VHD or VHDX format.
- Guest OS File Recovery: recover individual guest OS files from Windows, Linux, Mac and other guest OS file systems.
- Application Items Restore: restore application items (Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, and Microsoft SQL Server).
- Restore to Microsoft Azure: restore VMs from backups to Microsoft Azure.
- Restore to Amazon EC2: restore VMs from a backup to Amazon EC2 as EC2 instances.

Entire VM Restore

You can restore VMs to an AHV cluster from backup files and snapshots created by AHV Backup Proxy. You can perform the restore using both the AHV Backup Proxy web console and the Veeam Backup & Replication console. Note that the restore functionality differs depending on the console you use to perform the restore.

For instructions on how to perform full VM restore, see one of the following guides:

- Restore VMs using Veeam Backup & Replication Console: restore VMs of any platforms to the original or different AHV cluster.
- Restore VMs using AHV Proxy Console: restore AHV VMs to the original AHV cluster.

Restoring VMs Using Veeam Backup & Replication Console

In Veeam Backup & Replication console, you can restore VMs to Nutanix AHV clusters. You can restore from backups of VMware, Hyper-V, AHV VMs, from backups created by Veeam Agents and also from backups of EC2 instances and Azure VMs created by backup copy jobs.

Before you restore VMs, check the prerequisites and limitations. Then launch the **Full VM Restore** wizard to restore VMs.

1. Launch the Full VM Restore wizard.
2. Select VMs to restore.
3. Select the restore mode.
4. Select a target cluster.
5. Select a target storage container.
6. Specify the restored VM name.
7. Specify network mapping.
8. Specify a restore reason.
9. Review restore settings.

Before You Begin

Before you start the AHV VM restore process, mind the following:

- You can restore VMs to Nutanix AHV clusters from the following backups:
  - VM backups created with AHV Backup Proxy
  - VM snapshots created with AHV Backup Proxy
  - VMware vSphere VM backups created with Veeam Backup & Replication or Veeam Agent for Windows/Linux
  - Microsoft Hyper-V VM backups created with Veeam Backup & Replication or Veeam Agent for Windows/Linux
  - Backups of physical machines created with Veeam Agent for Windows or Veeam Agent for Linux
  - VeeamZIP backups
  - Backups of EC2 instances and Microsoft Azure VMs created by backup copy jobs
- SureBackup is not supported for backups created by AHV Backup Proxy.
- If you restore from a backup of a VMware, Hyper-V VM or from a backup created by Veeam Agent, restored VM may have network connection problems. To solve the problem, you must install Nutanix Guest Tools on the restored VM.
- When you restore from backups of VMware vSphere and Microsoft Hyper-V VM backups created by Veeam Backup & Replication or Veeam Agents, the restored VM has default hardware resources: 1 CPU core, 1 GB RAM. You can change default values in the settings of the restored VM.

- Within one restore session VM disks are restored sequentially.
- [For Archive Tier] If you want to restore a VM to AHV cluster, you must retrieve the required backup from the archive tier before starting the restore. For more information, see the Retrieving Backup Files section of the Veeam Backup & Replication User Guide.

Changing Default Disk Controller Interface

When you restore from backups of VMware vSphere and Microsoft Hyper-V VM backups created by Veeam Backup & Replication or Veeam Agents, all restored VM disks will be as SATA drives. To change the default disk controller interface, do the following:

1. On the Veeam Backup & Replication server, go to `%Program Files%\Veeam\Plugins\Nutanix AHV\Service` and open the `appsetting.json` with a text editor.
2. Change the default disk controller interface value for the `Any2AhvDefaultDisksBusType` parameter. You can use the following options: `"SCSI"`, `"PCI"`, `"IDE"`, `"SATA"`. Note that the parameter is not case-sensitive.

   ```
   "Any2AhvDefaultDisksBusType": "Sata"
   ```

3. Restart the *Veeam AHV Service* to implement the changes.

Step 1. Launch Restore Wizard

To launch the **Full VM Restore to Nutanix AHV** wizard:

1. In the Veeam Backup & Replication console, open the **Home** view and select **Backups** in the inventory pane.
2. Right-click a backup from which you want to restore and select **Restore entire VM to Nutanix AHV**.

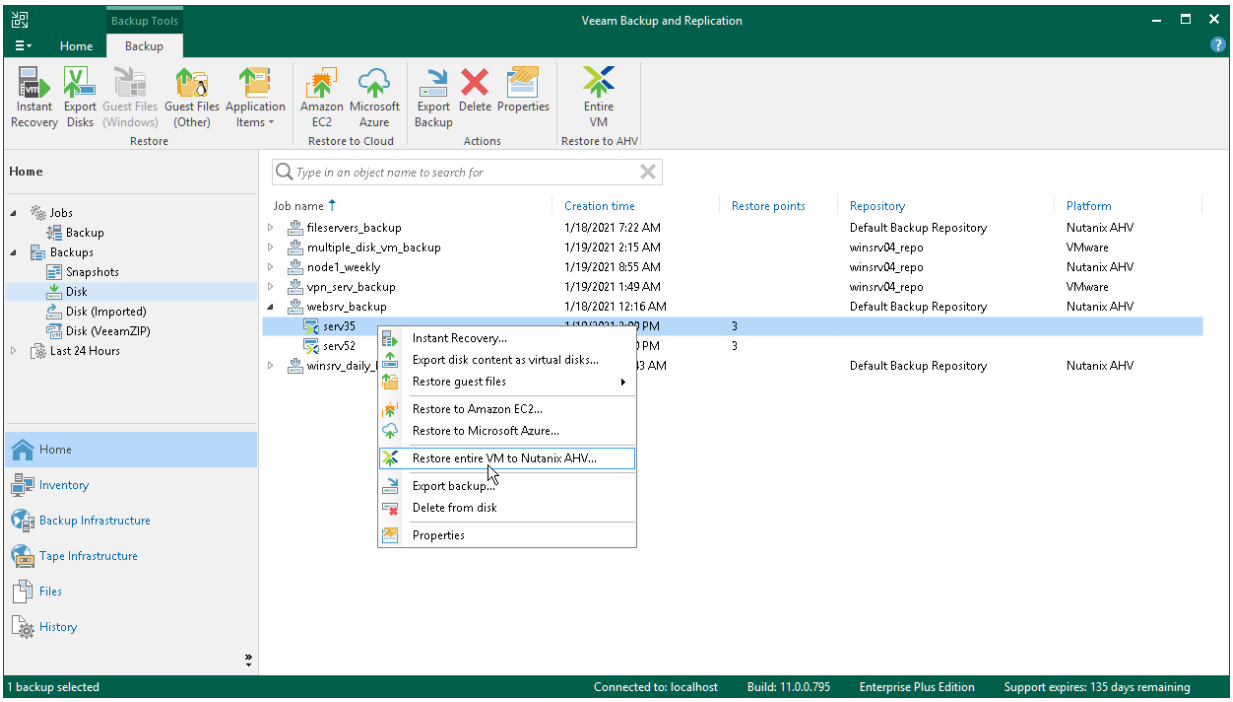

Alternatively, to launch the **Full VM Restore to Nutanix AHV** wizard, you can do the following:

1. On the **Home** tab, click **Restore** and select **Nutanix AHV**.

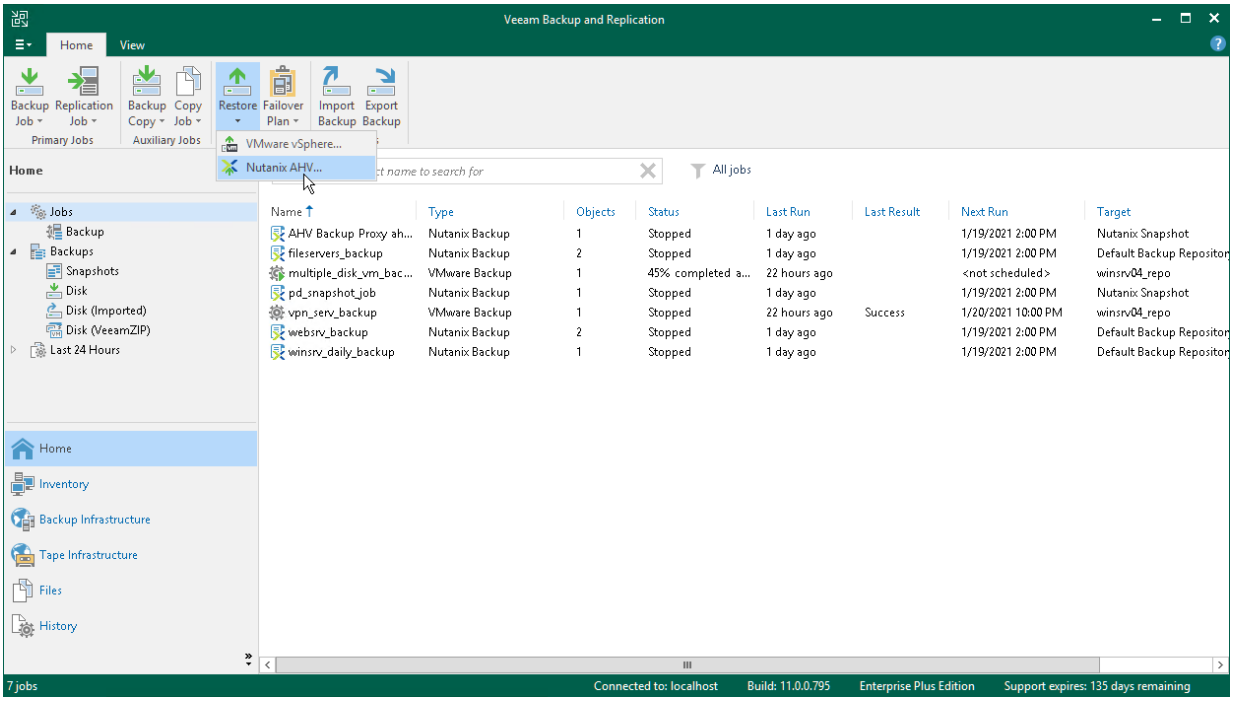

2. Select whether you want to restore from a backup file or a snapshot:
   - Select **Restore from AHV backup**, if you want to restore from backup files created by a backup job.
   - Select **Restore from AHV backup snapshot**, if you want to restore from a snapshot created by a backup job.
   - Select **Restore from AHV user snapshot**, if you want to restore from a snapshot created by a snapshot job.
   - Select **Restore from AHV PD snapshot**, if you want to restore from a PD snapshot created by a PD snapshot job.

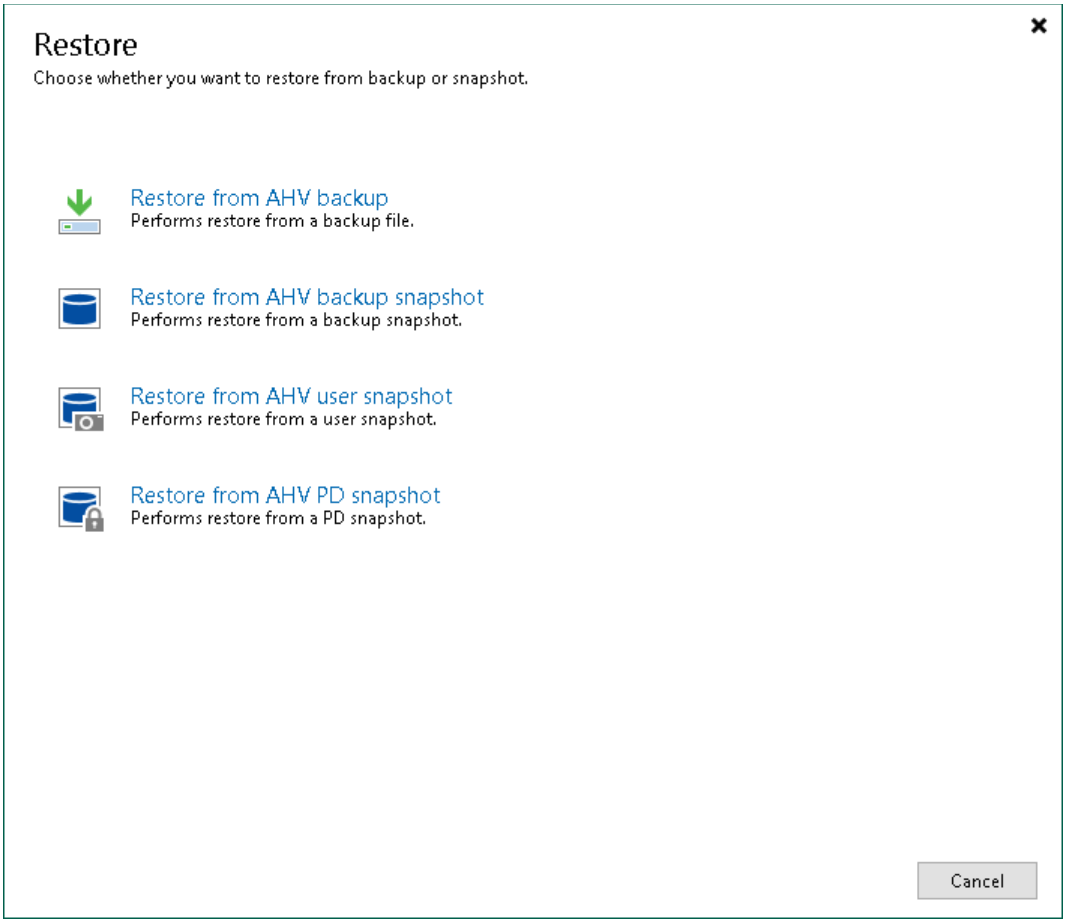

Step 2. Select VMs

At the **Virtual Machines** step of the **Full VM Restore to Nutanix AHV** wizard, select VMs that you want to restore.

To select VMs:

1. On the right of the wizard window, click **Add VM.**
2. **In the Backup Browser window, expand a backup job tree that contains backups of VMs you want to restore and select a VM.** Hold [Ctrl] to select multiple VMs.
3. By default, AHV Backup Proxy restores VMs from the latest backup. You can select to restore to an earlier restore point. Restore from a snapshot is significantly faster than restore from a backup. Note that you can restore from a snapshot only if a VM has at least one successful backup.

   To restore a VM to an earlier point in time:

   a. Select the VM in the list and click **Point**.

   b. In the **Restore Points** window, select a backup.

   c. Click **OK**.

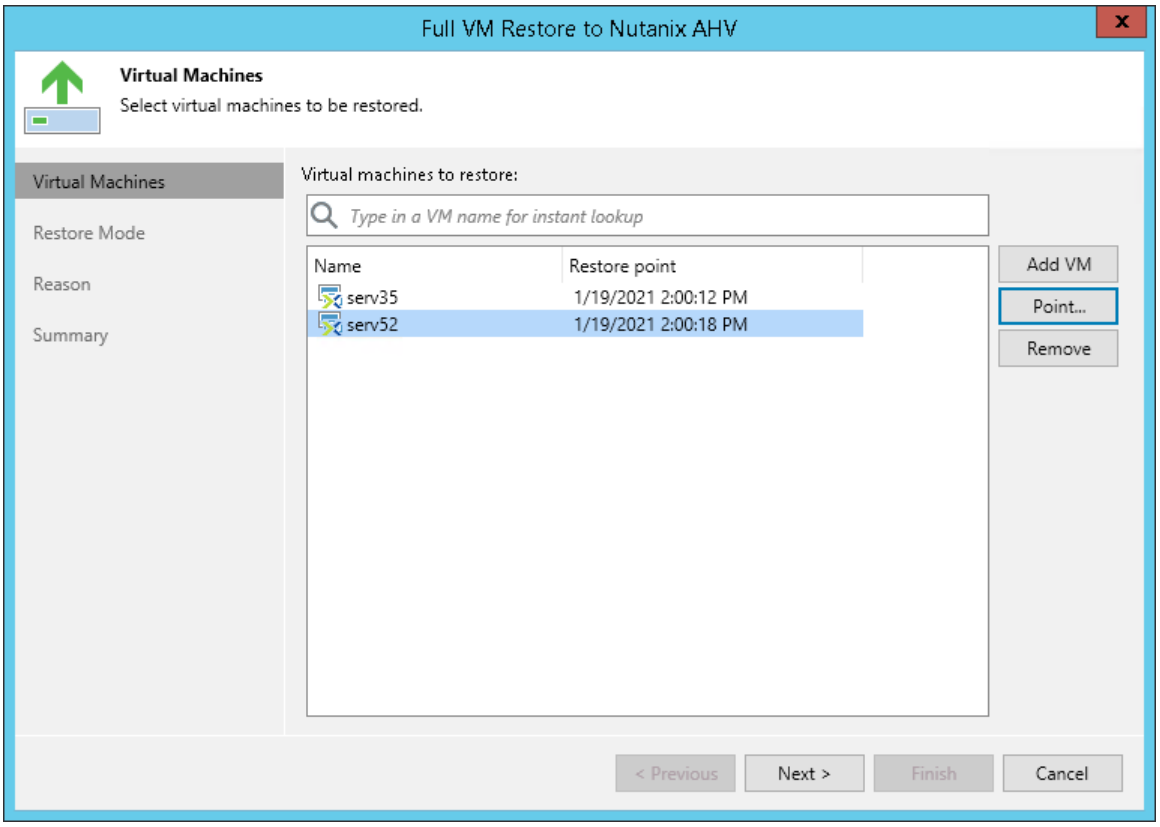

Step 3. Select Restore Mode

At the **Restore Mode** step of the **Full VM Restore to Nutanix AHV** wizard, specify how VMs will be restored.

- Select **Restore to original location** if you want to restore VMs with initial settings to their original location. If you select this option and click **Next**, you pass directly to the Reason step of the wizard.

  If a VM with the same name exists on the original AHV container, it will be replaced by the VM restored from the backup or snapshot.

- Select **Restore to a new location, or with different settings**, if you want to restore VMs to a different location and/or with different settings (such as VM name, VM location, and network settings). If you select this option, the **Full VM Restore to Nutanix AHV** wizard will include additional steps for customizing VM settings.

> **NOTE**
>
> - If you restore from a backup snapshot, you cannot restore the VM to a different cluster or different container.
> - If you restore from a user snapshot, you cannot restore the VM to a different cluster or different container.
> - If you restore from a PD snapshot, you cannot restore the VM to a different cluster or different container. Also, you cannot select a network for restored VMs.

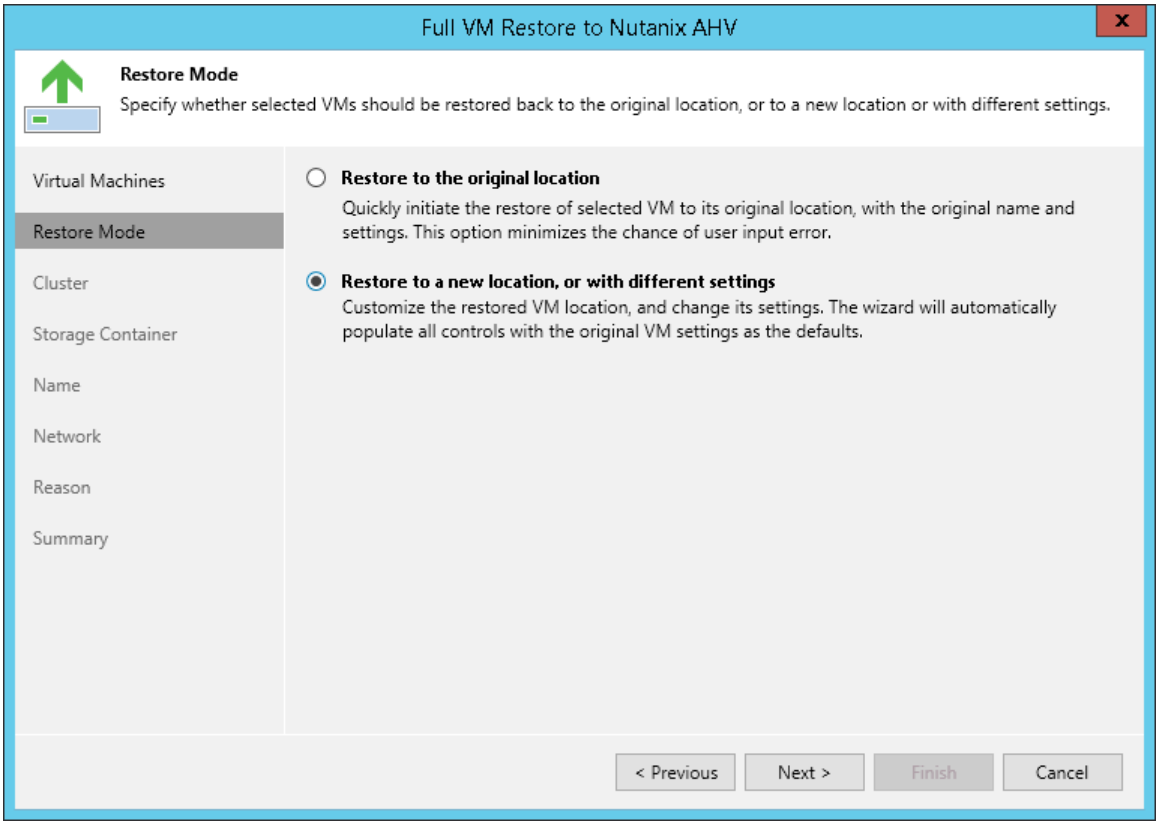

Step 4. Specify Target Cluster

The **Cluster** step of the wizard is available if you have selected to change the location and settings of the restored VM.

By default, VMs are restored to the original cluster. To change the default target cluster, do the following:

1. In the **VM location** field, select a VM and click **Cluster**.

   Hold [Ctrl] and select multiple VMs to select a cluster for multiple VMs.

2. In the **Select Cluster** window, select a cluster and click **OK**.

> **NOTE**
>
> The **Cluster** step of the **Full VM Restore to Nutanix AHV** wizard is not available if you restore a VM from a snapshot.

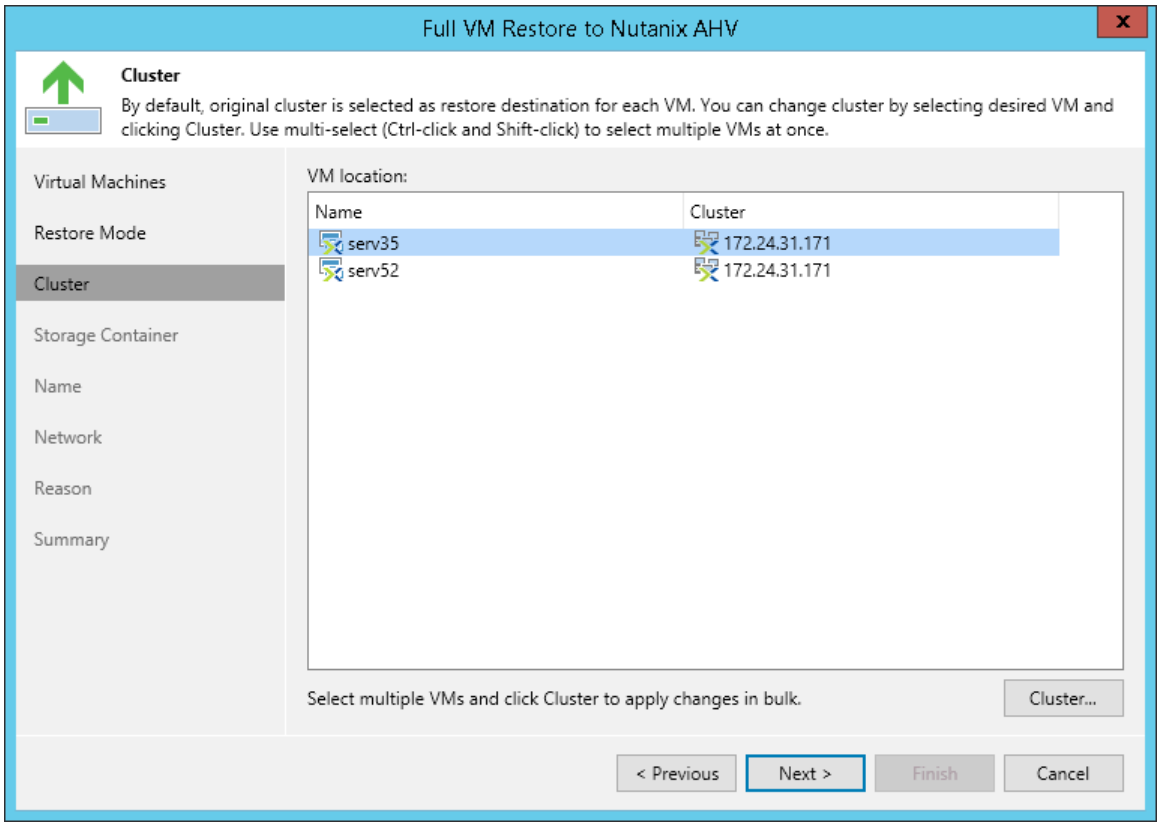

Step 5. Specify Storage Container

The **Storage Container** step of the wizard is available if you have selected to change the location and settings of the restored VM.

To change the default target container, do the following:

1. In the **Storage container** field, select a VM and click **Container**.

   Hold [Ctrl] and select multiple VMs to apply the setting for multiple VMs.

2. In the **Select Storage Container** window, select the container where you want to store the backup files and click **OK**.

> **NOTE**
>
> If you restore from a snapshot, you cannot select a different container. VM disks of the restored VM will be on the default container.

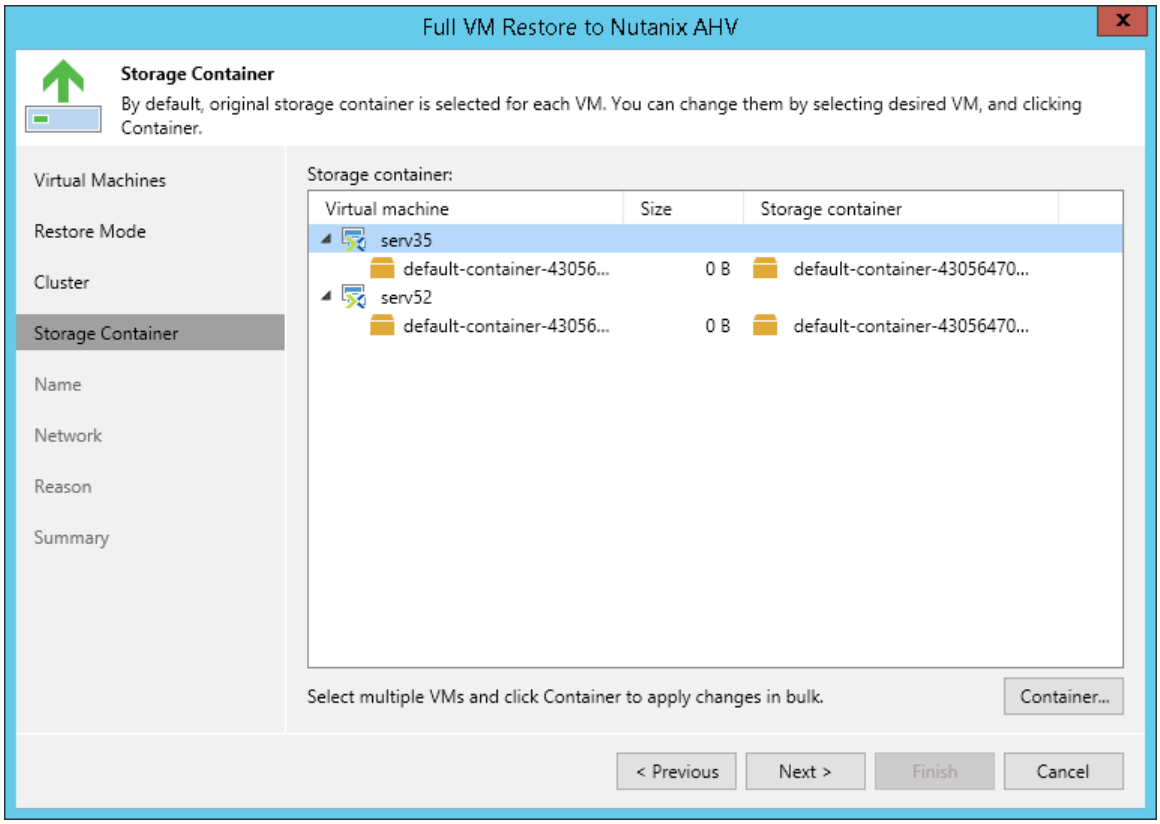

Step 6. Specify VM Name

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

By default, VMs are restored with original names. Note that if you restore VMs to the same location and don't set new names for VMs, original VMs will be overwritten by restored VMs.

To specify a new name for the restored VM:

1. In the **Virtual machines** field. Select a VM and click **Name**.

   Hold [Ctrl] to apply the setting for multiple VMs.

2. At the **Rename VM** window, specify a name for the restored VM:
   - In the **Set name to** field, specify a new name for the restored VM.
   - Check the **Add prefix** check box and specify a prefix that will be added to the name of the restored VM.
   - Check the **Add suffix** check box and specify a suffix that will be added to the name of the restored VM.

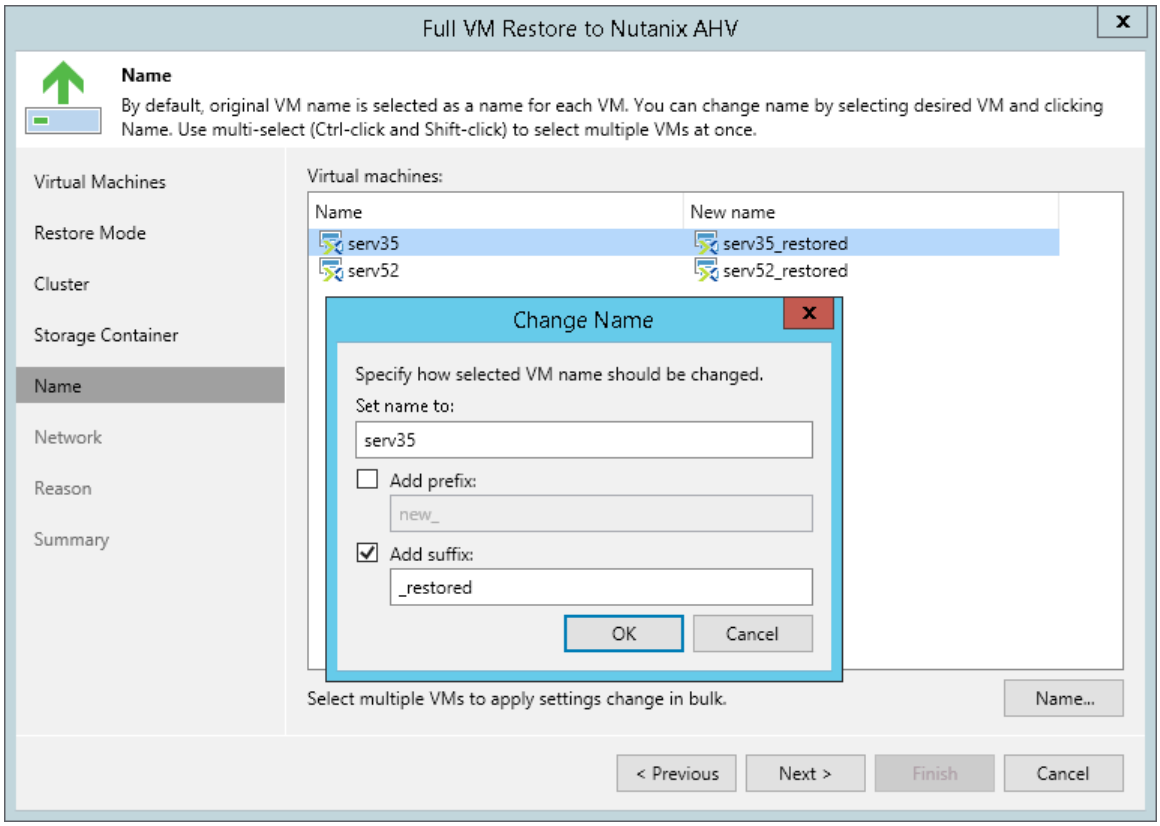

Step 7. Specify Network Settings

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. AHV Backup Proxy 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 and click **Network**. If you want to apply the setting to multiple VMs, hold [Ctrl] and select multiple VMs,

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

2. In the **Select Network** window, expand a cluster and choose a network to which the VM must have access upon restore. To quickly find a network, 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.

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

> **NOTE**
>
> If you restore from snapshots, you cannot map a new network. You can only choose to disconnect the original network.

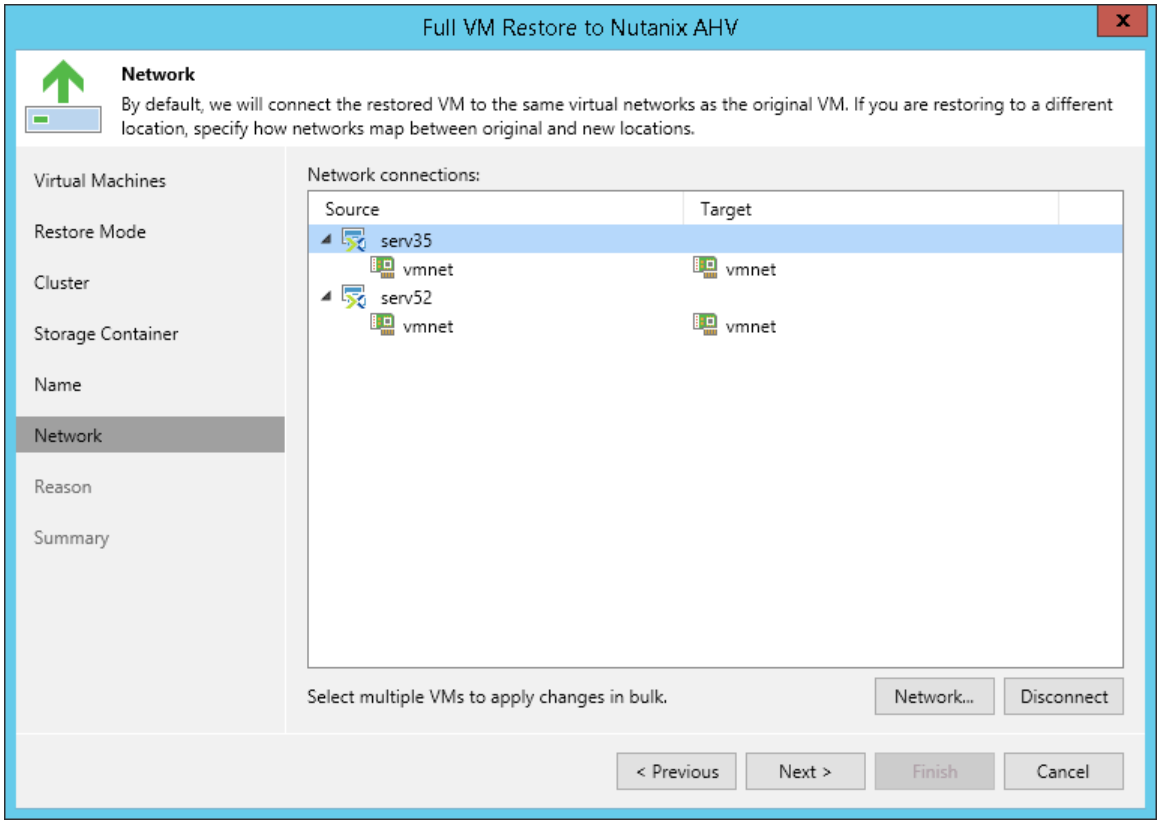

Step 8. Specify Restore Reason

[Optional] 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.

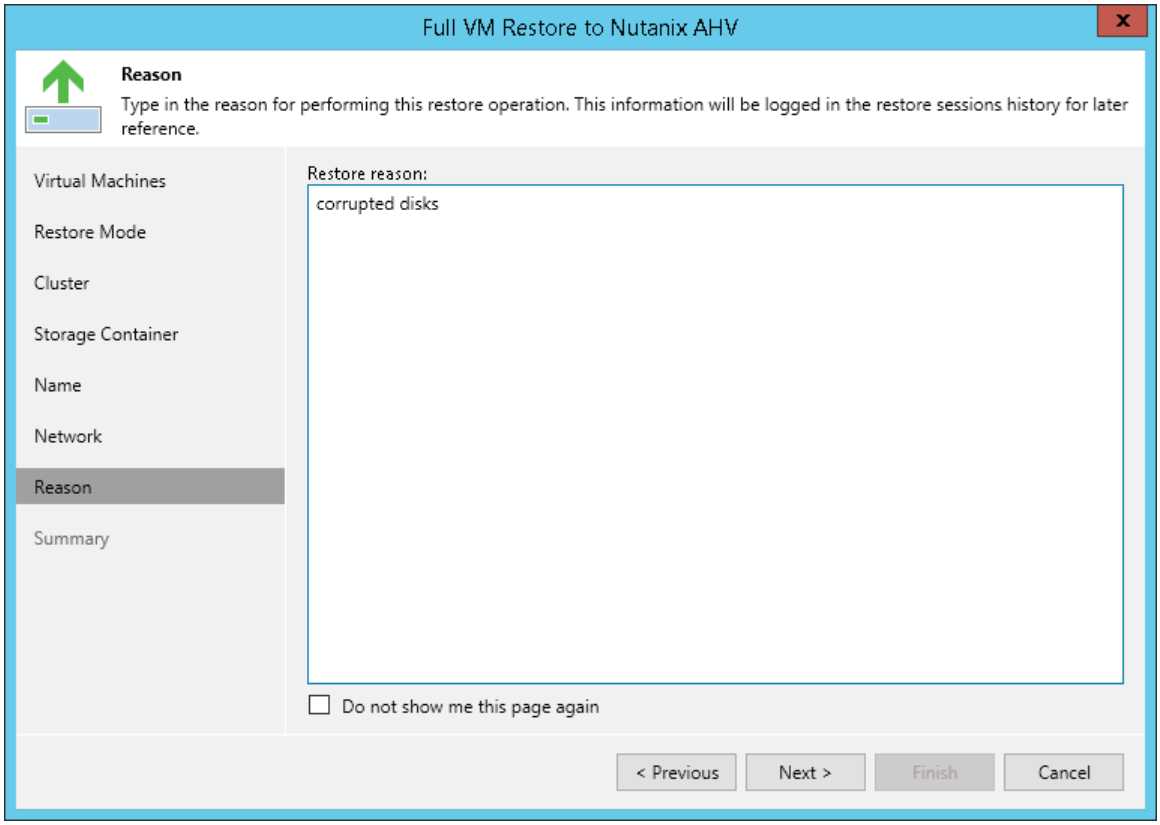

Step 9. Review Restore Settings

At the **Summary** step of the wizard, specify additional settings for VM restore:

1. If you want to start the restored VM on the target host, select the **Power on VM after restoring** check box.
2. Check the specified settings and click **Finish**. Veeam Backup & Replication will restore selected VMs in the specified destination.

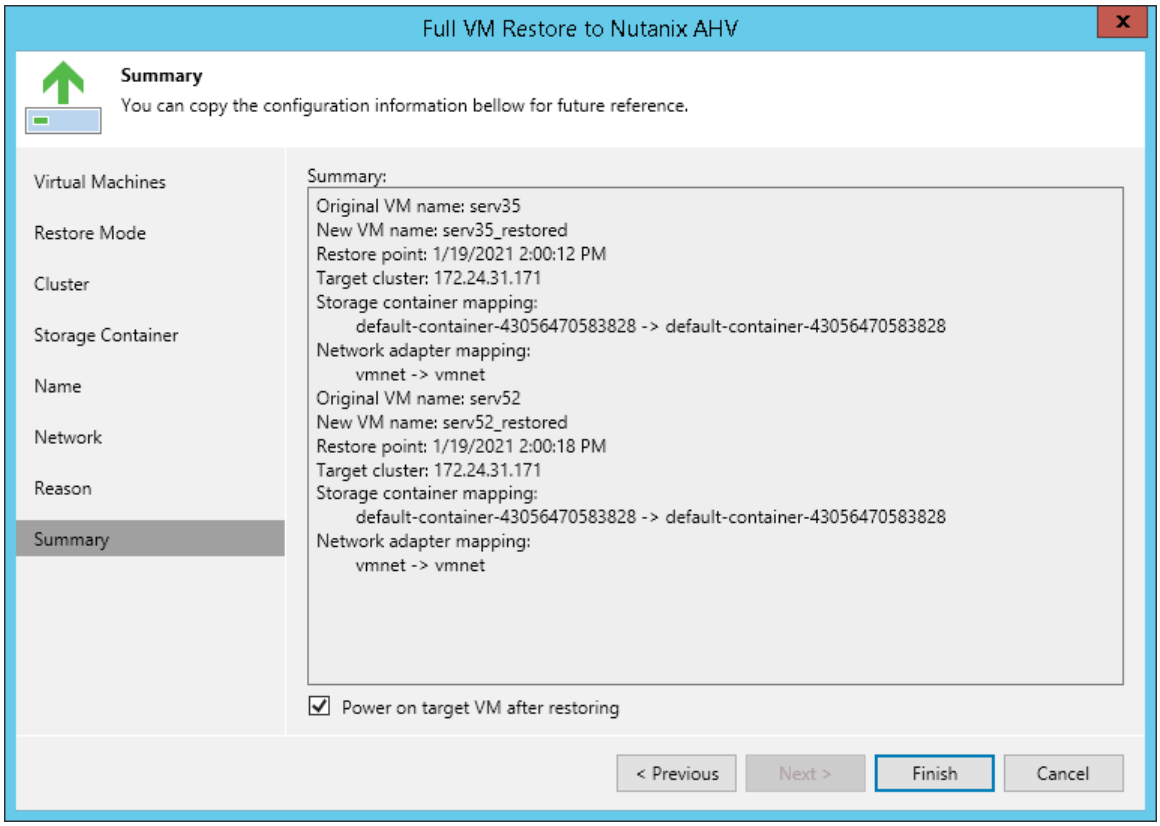

Restoring VMs Using AHV Backup Proxy Console

In the AHV Backup Proxy web UI, you can restore AHV VMs from backups and snapshots to a Nutanix AHV cluster.

During the restore process, AHV Backup Proxy retrieves VM disk data from the backup, copies it to the datastore where disks of the original VM were located, and registers a restored VM on the Nutanix AHV cluster node.

Before you restore VMs, check the prerequisites and limitations. Then launch the Full VM Restore wizard and follow its steps:

1. Select VMs to restore.
2. Select the restore mode.
3. Assign a name for the restored VM.
4. Select the container where the restored VM will be deployed.
5. Specify network settings.
6. Specify a reason for VM restore.
7. Finish working with the wizard.

Before You Begin

Before you start the AHV VM restore process, mind the following limitations:

- If a VM added to the restore session contains multiple disks, AHV Backup Proxy restores these disks sequentially, one disk at a time. Note that VMs added to the restore session are processed in parallel.
- If a source VM has a set affinity and if during restore this host is not available on the cluster (or original VM was backed up from different cluster), you must set the affinity policy manually before starting the VM.
- SureBackup is not supported for backups created by AHV Backup Proxy.
- When you restore VMs using the AHV Backup Proxy web console, you can restore only to the original AHV cluster. If you want to restore VMs to another cluster, perform the restore using the Veeam Backup & Replication console. For more information, see Restoring VMs using Veeam Backup & Replication Console. Alternatively, you can import the backups from another cluster and and restore from them.
- When you restore from a snapshot, you cannot change the target container.

  When you restore from a backup, you can select only one container for all VM disks. This means, you cannot select a different container for each VM disk.
- [For restore from user snapshots and PD snapshots] You cannot change network settings of the VM in the restore wizard. You can change these parameters after the restore is complete via Prism Element or Prism Central.

Launch Full VM Restore Wizard

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

1. At the main panel of the AHV Backup Proxy web console, open the **Protected VMs** tab.

2. Click **Restore** on the toolbar to launch the **Full VM Restore** wizard.

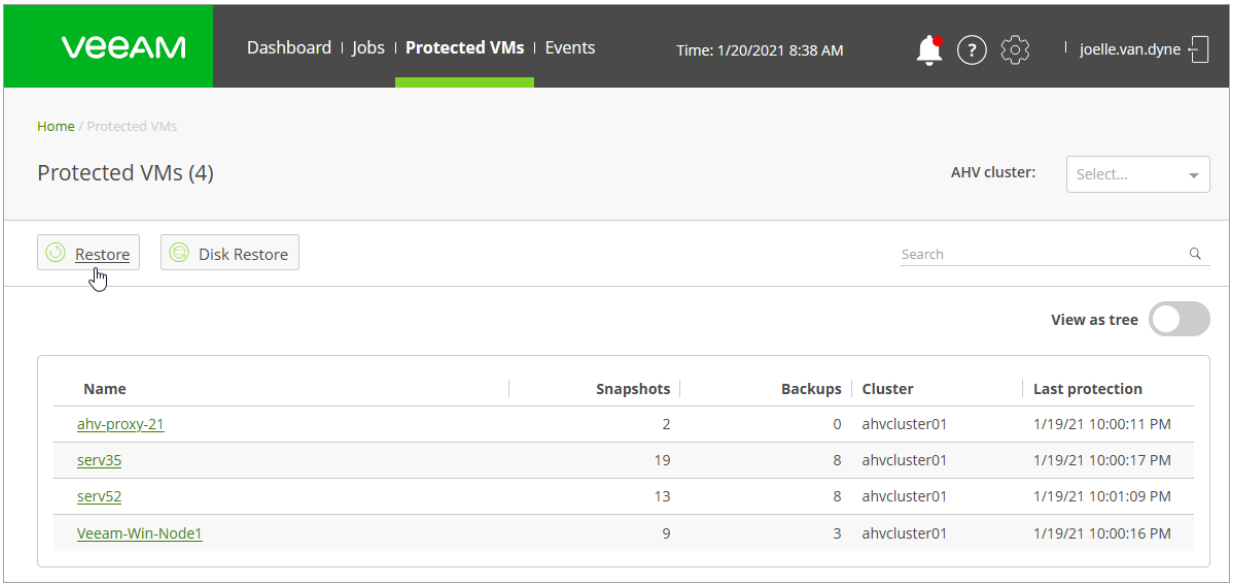

Step 1. Select VM to Restore

At the **Virtual Machines** step of the **Full VM Restore** wizard, select a VM that you want to restore.

1. Click **Add** on the toolbar.

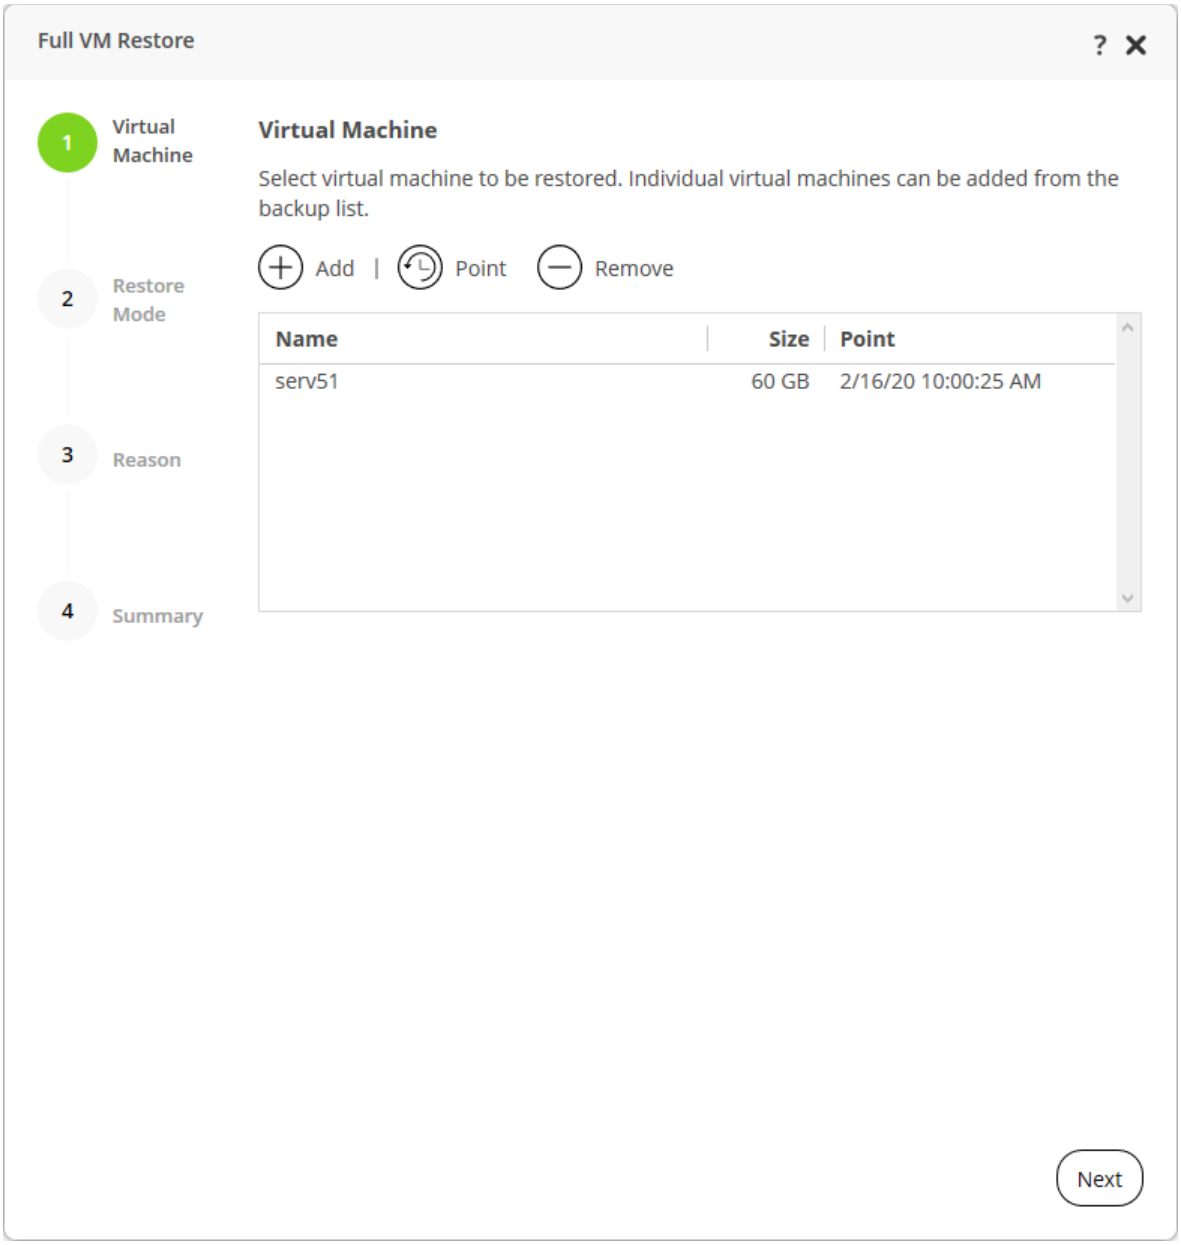

2. In the **Backup Browser** window, expand the required job and select a VM which you want to restore.

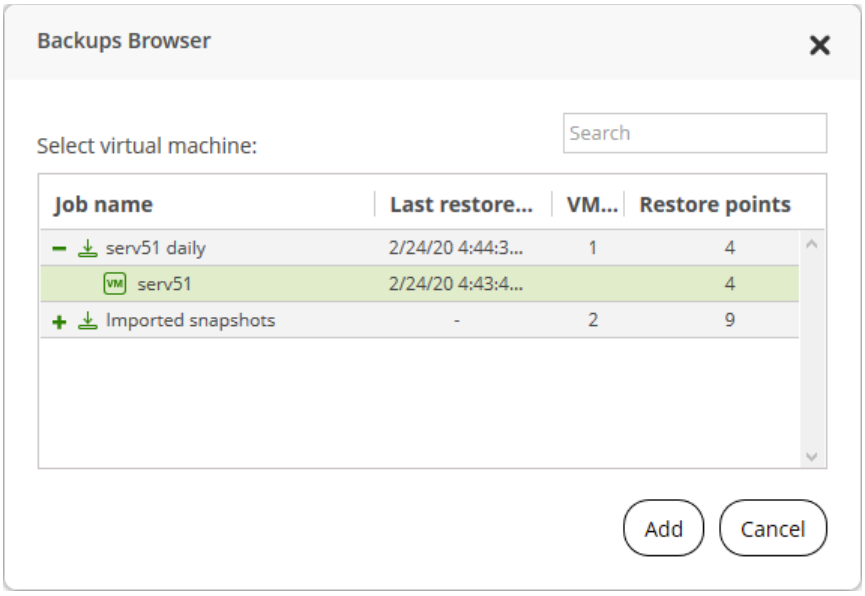

> **NOTE**
>
> In the AHV Backup Proxy web console, you can add only one VM to the restore session. If you want to restore multiple VMs in parallel, use the Veeam Backup & Replication console. For more information, see Restoring VMs Using Veeam Backup & Replication Console.

3. By default, AHV Backup Proxy restores VMs from the latest backup snapshot. You can leave the default setting or select to restore from an earlier backup, backup snapshot, PD snapshot, or user snapshot.

   Restore from a backup snapshot is significantly faster than restore from a backup. Note that you can restore from a backup snapshot only if a VM has at least one successful backup.

   To restore a VM to an earlier point in time:

   a. Click **Point on the toolbar**.

   b. In the **Restore Points** window, select a backup or a backup snapshot.

   c. Click **OK**.

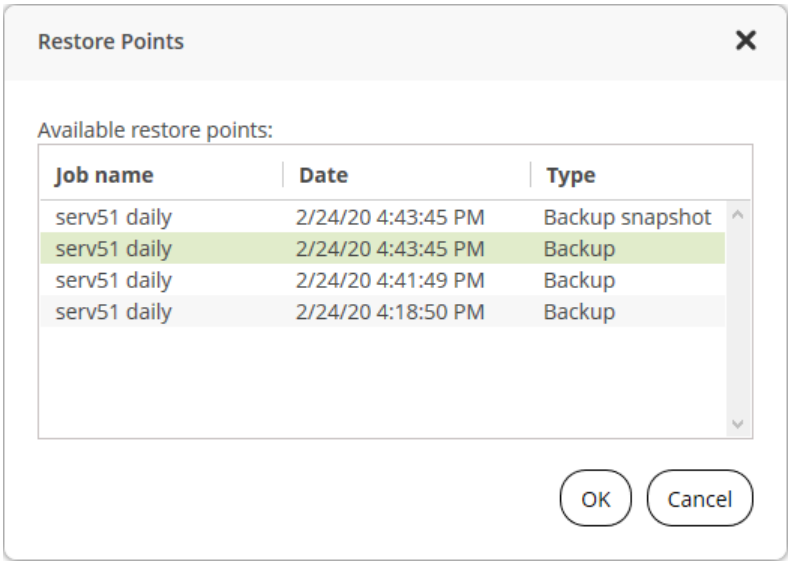

> **NOTE**
>
> AHV Backup Proxy stores the last backup snapshot for each job. For example, if a VM is protected by two jobs, the AHV Backup Proxy will store two backup snapshots.

Step 2. Select Restore Mode

At the **Restore Mode** step of the **Full VM Restore** wizard, specify how the VM must be restored:

- Select **Restore to original location** if you want to restore the VM with initial settings to its original location. If you select this option and click **Next**, you pass directly to the Reason step of the wizard.

  If a VM with the same name exists on the source AHV container, AHV Backup Proxy asks if you want to overwrite the original VM. If you select **Yes**, AHV Backup Proxy removes the VM and replaces it with the restored VM.

- 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 VM name, VM location, and network settings). If you select this option, the **Full VM Restore** wizard will include additional steps for customizing VM settings.

NOTE

- If you restore from a backup snapshot, you cannot restore the VM to a different cluster or different container.
- If you restore from a user snapshot, you cannot restore the VM to a different cluster or different container, and you cannot map a new network for the restored VM.
- If you restore from a PD snapshot, you cannot restore the VM to a different cluster or different container, and you cannot map a new network for the restored VM.

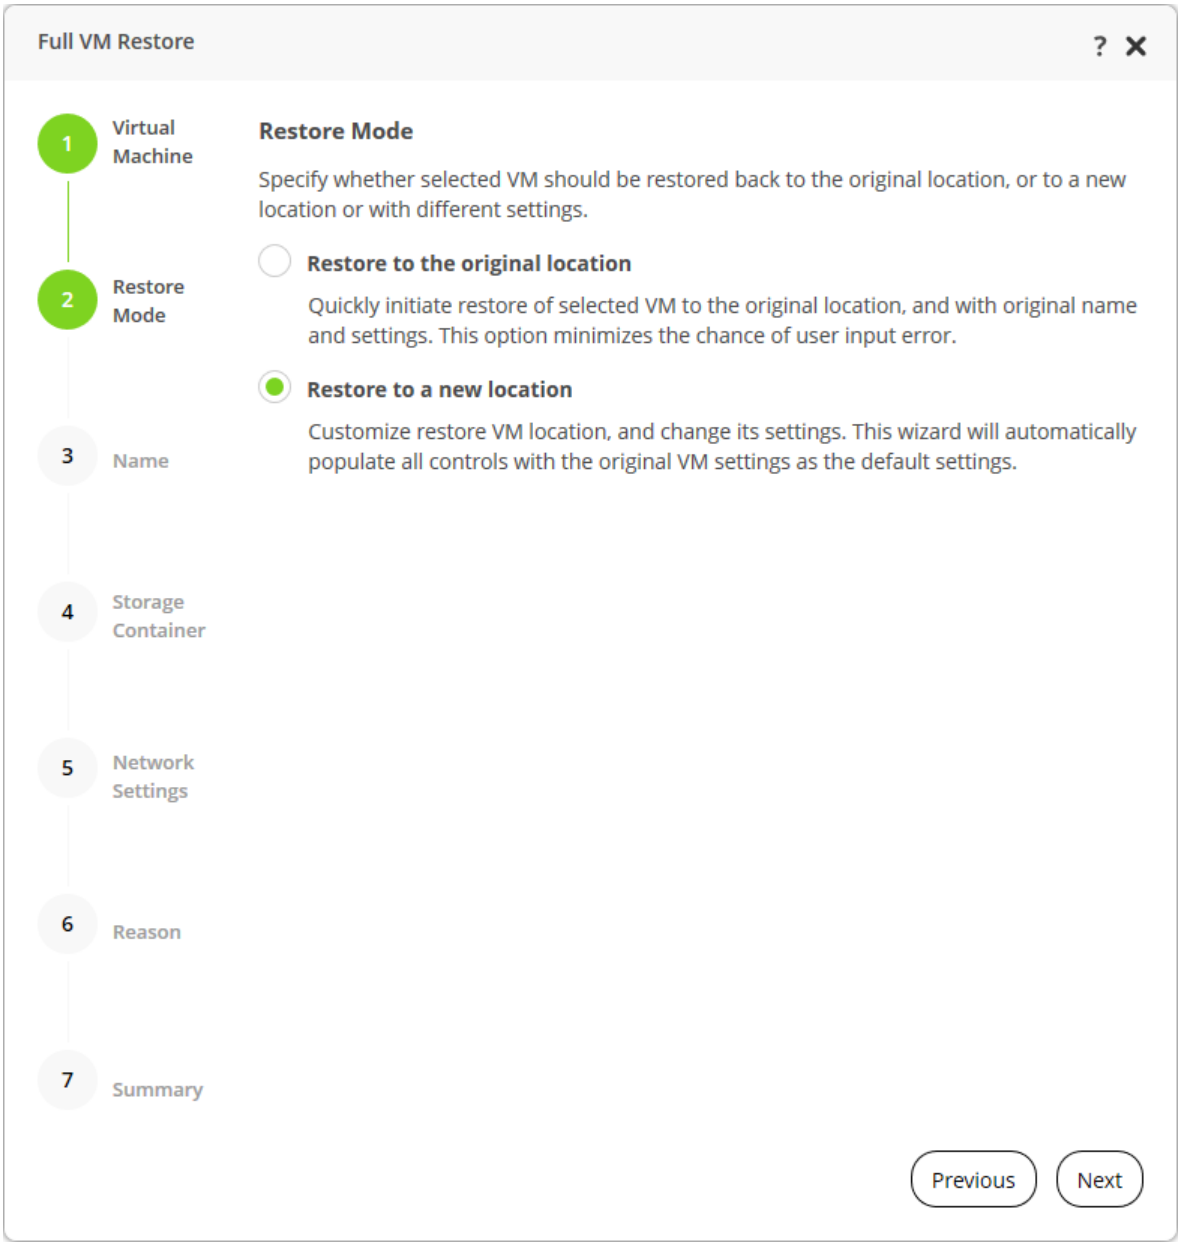

Step 3. Assign Name for Restored VM

The **VM name** step of the wizard is available only if you have selected the option to change the location and settings of the restored VM.

By default, VMs are restored with original names. To specify a new name for the restored VM:

1. Click the **Name** button.
2. In the **Rename VM** window, specify a name for the restored VM:
   - In the **Set name to** field, specify a new name for the restored VM.
   - Check the **Add prefix** check box and specify a prefix that will be added to the name of the restored VM.
   - Check the **Add suffix** check box and specify a suffix that will be added to the name of the restored VM.

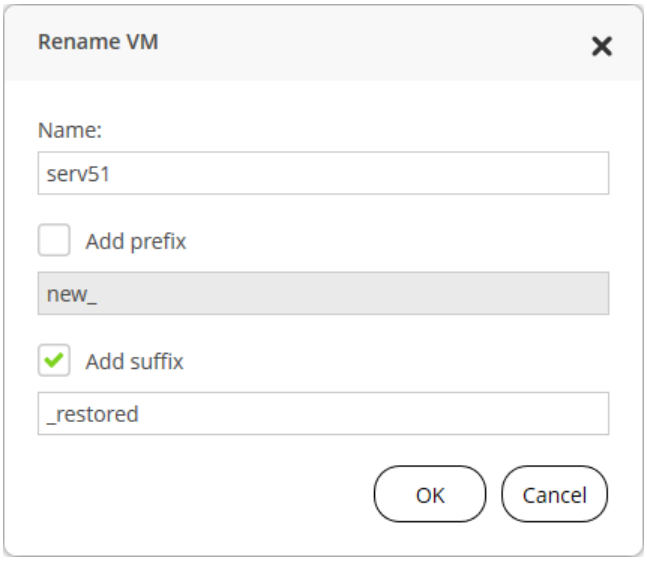

3. If you select the **Preserve virtual machine ID** check box and the source VM still exists, AHV Backup Proxy will replace the original VM by the VM from the backup.

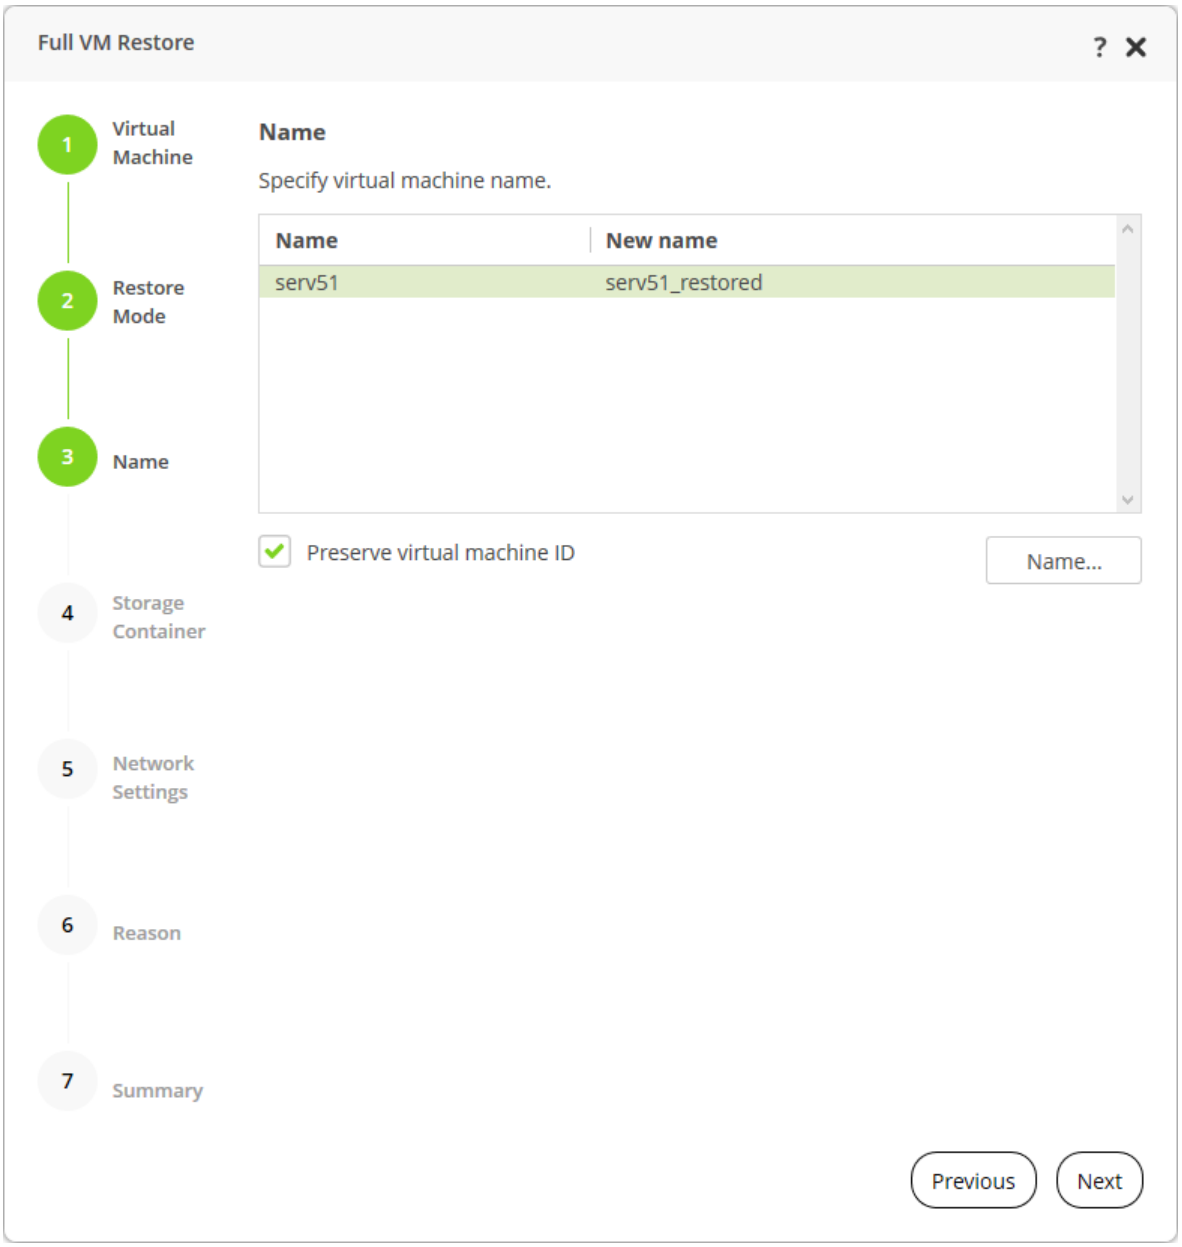

Step 4. Select Container

The **Storage Container** step of the wizard is available if you have selected to change the location and settings of the restored VM.

> **NOTE**
>
> If you restore from a snapshot, you cannot select a different container. VM disks of the restored VM will be on the default container.

To change the default target container, do the following:

1. In the **Disk location** field, expand the VM tree.
2. Click the **Container button**.

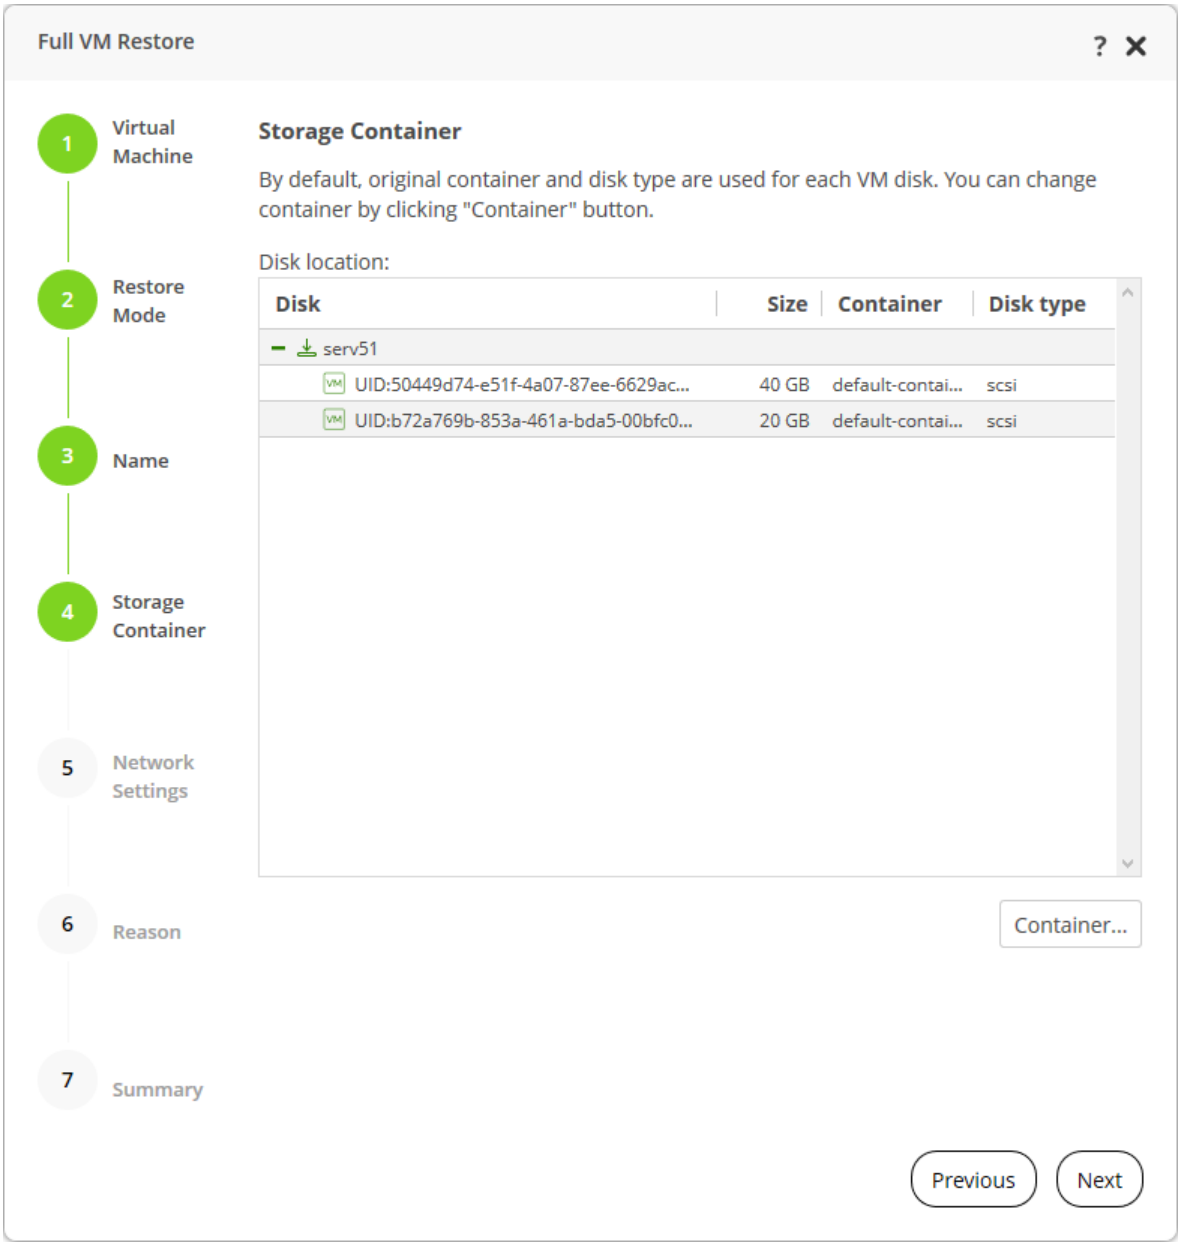

3. In the **Select Container** window:
   a. Expand the AHV cluster tree.

b. Select the target container where you want to restore the VM.

c. Click **OK**.

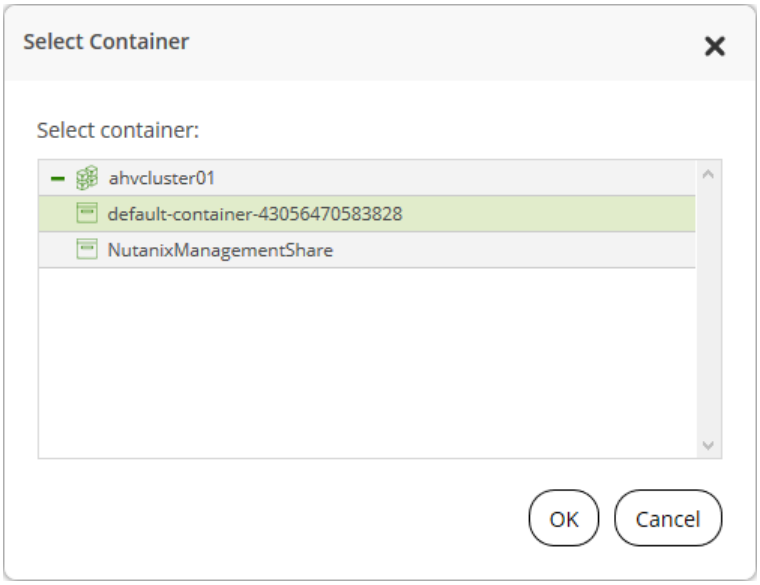

Step 5. Specify Network Settings

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 the source site networks to target site networks. AHV Backup Proxy 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. Click the **Network** button.

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

2. In the **Select Network** window, select a network to which the VM must have access after the restore.

   If you do not want to connect the restored VM to any virtual network, select the VM in the list and click **Disconnected**. In that case, you must connect the VM to network manually after the restore process.

NOTE

If you restore from a PD snapshot or a user snapshot, you cannot map a new network. You can only choose to disconnect the original network.

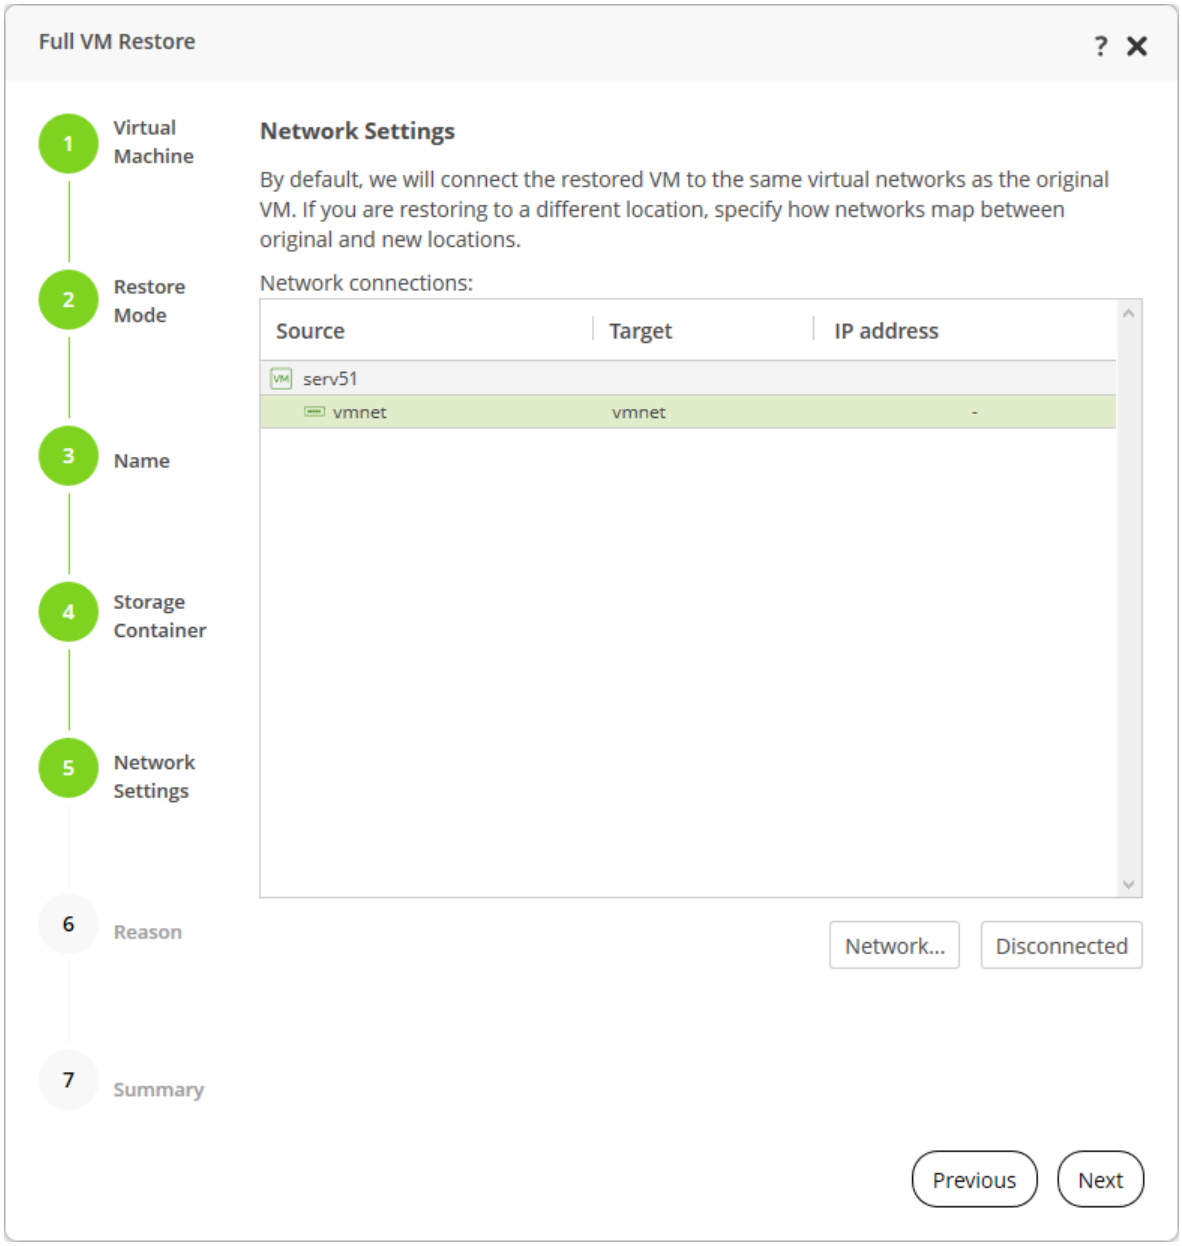

Step 6. Specify Reason for Restore

[Optional] 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.

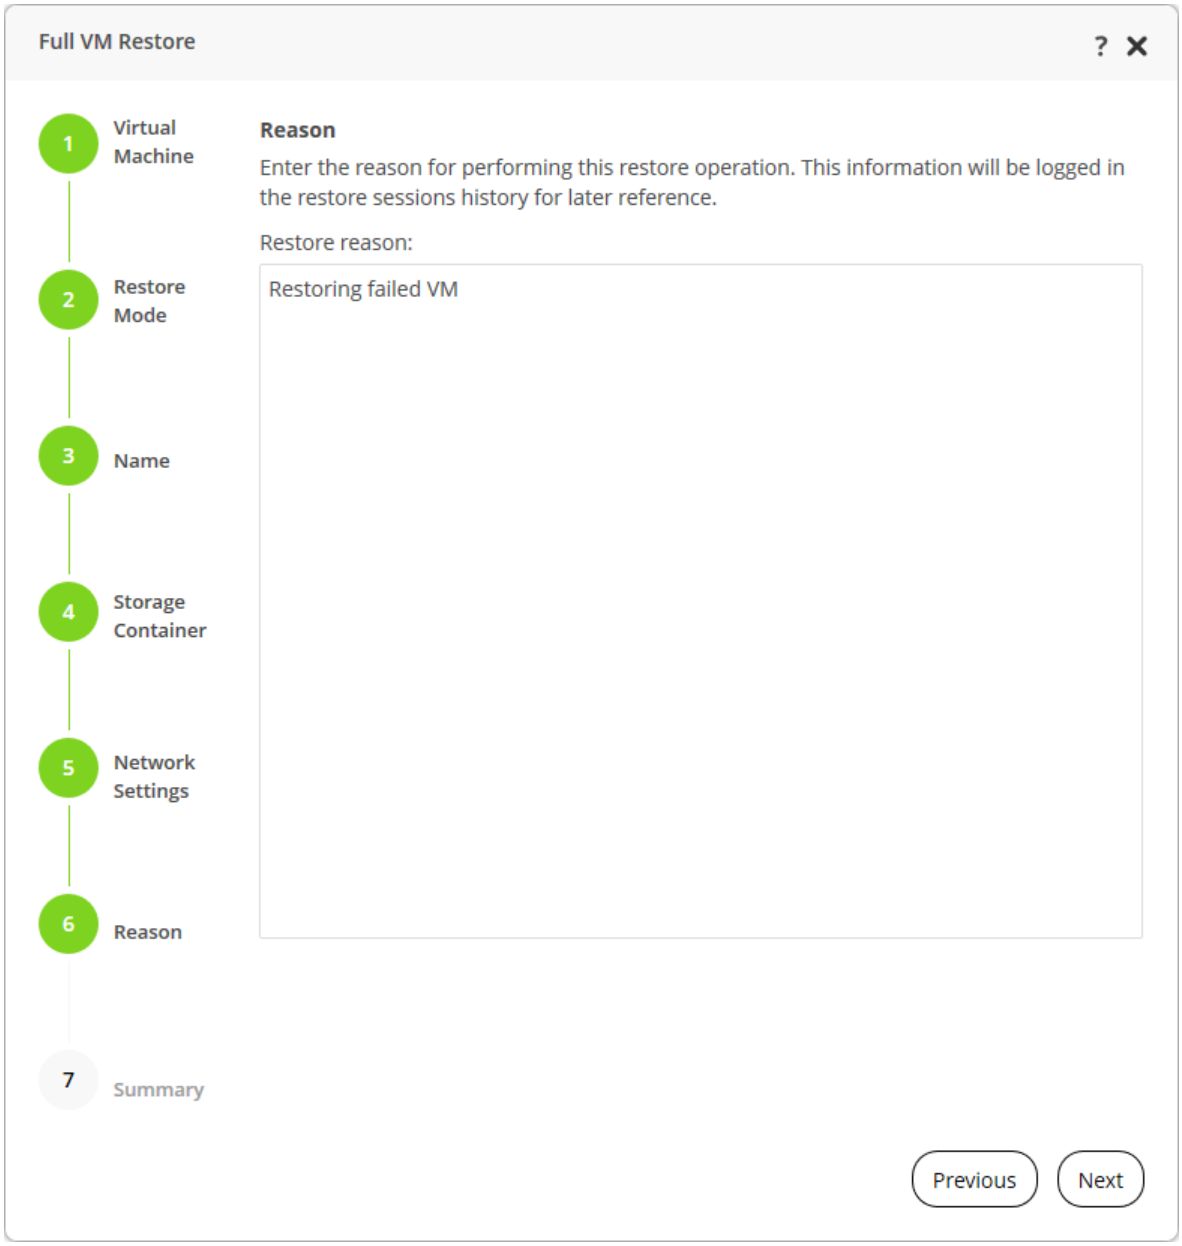

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard:

1. Review the restore settings.
2. If you want to turn on the VM after the restore, select the **Power on VM after restore** check box .
3. Click **Finish** to close the wizard and start the restore process.

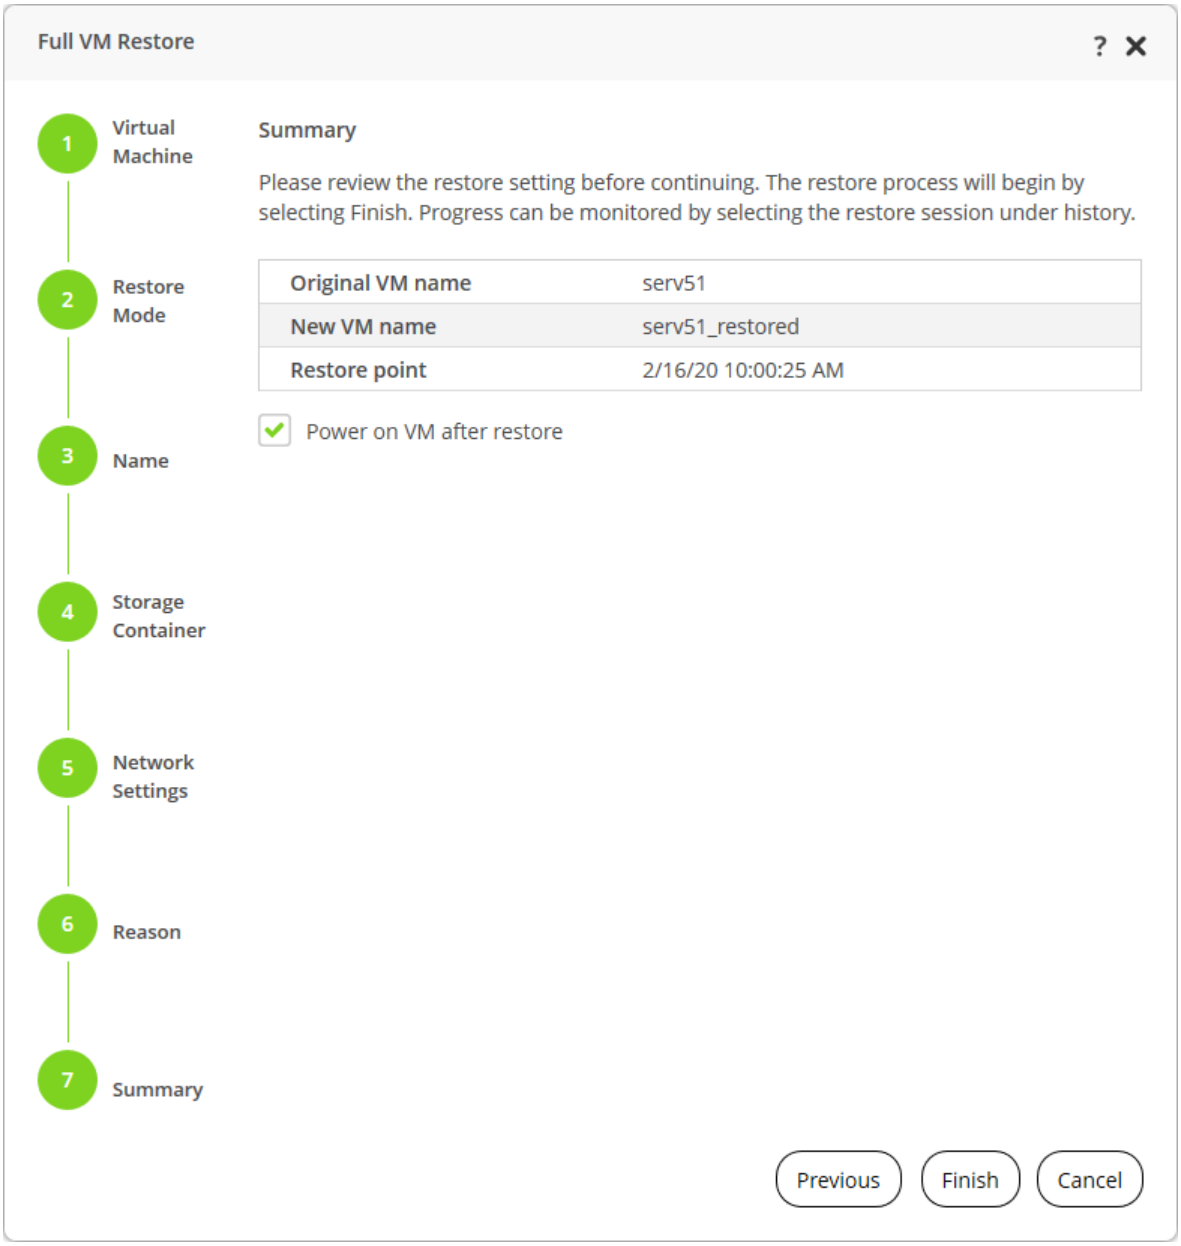

VM Disks Restore

If a VM virtual disk becomes corrupted for some reason, you can restore it from a backup. The restored virtual disk can be attached to the original VM to replace the corrupted disk or connected as a new disk to any other VM.

Before you restore disks, check the limitations. Then launch the Virtual Disk Restore wizard and follow its steps:

1. Select a virtual machine.
2. Select a restore point.
3. Configure mapping settings.
4. Specify a reason for the restore.
5. Finish working with wizard.

Before You Begin

Before you start the VM disk restore process, mind the following limitations:

- If a VM added to the restore session contains multiple disks, AHV Backup Proxy restores these disks sequentially, one disk at a time. Note that VMs added to the restore session are processed in parallel.
- You can perform VM disk restore only from backups and backup snapshots. For more information on snapshot types, see Nutanix AHV Snapshot Types.

> **IMPORTANT**
>
> When you finish configuring the restore settings, AHV Backup Proxy turns off the target VM and turns it on after the restore process completes.

Launch Virtual Disk Restore Wizard

To launch the **Virtual Disk Restore** wizard, do the following:

1. At the main panel of the AHV Backup Proxy web console, click the **Protected VMs** tab.
2. On the **Protected VMs** tab, click **Disk Restore** on the toolbar to launch the **Virtual Disk Restore** wizard.

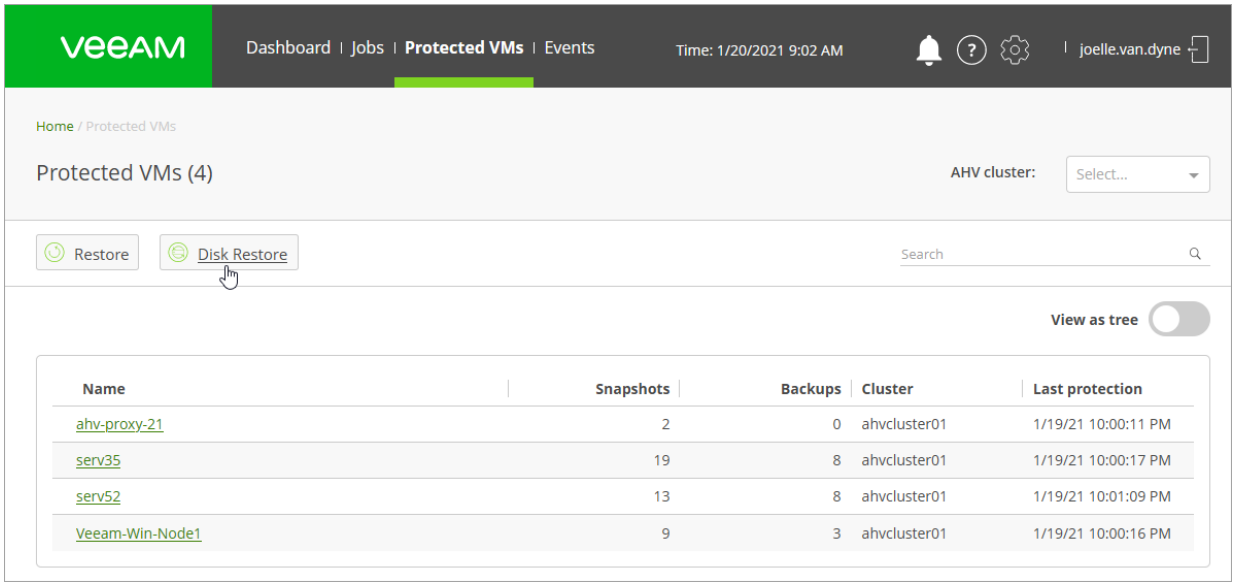

Step 1. Select Virtual Machine

At the **Virtual Machine** step of the wizard, select VMs that contain disks that you want to restore.

1. In the list of jobs, expand the job tree.
2. Select the VM which has disks that you want to restore and click **Next**.

   To quickly find a VM, you can use the search field.

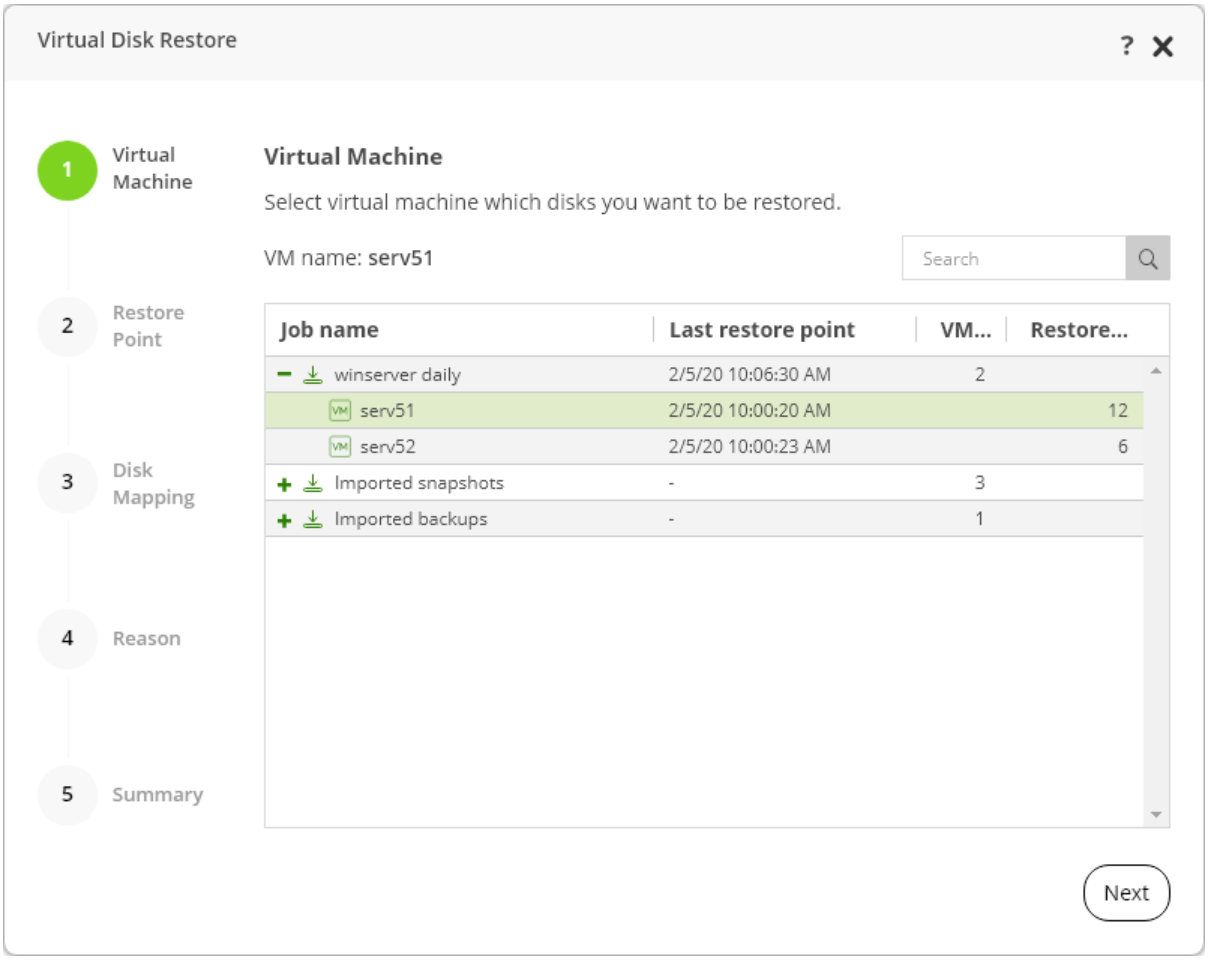

Step 2. Select Restore Point

At the **Restore Point** step of the wizard, you can select the required restore point for the VM.

By default, AHV Backup Proxy restores VM disks from the latest backup snapshot. If you want to restore from an earlier restore point, you have the following options.

- **Restore from a backup**

  To restore a VM to an earlier state, select the necessary backup restore point and click **Next**.

- **Restore from a backup snapshot**

  To restore from a snapshot, select the required snapshot from the list and click **Next**.

  Restore from a snapshot is significantly faster than restore from a backup. Note that you can restore from a snapshot only if the VM has at least one successful backup.

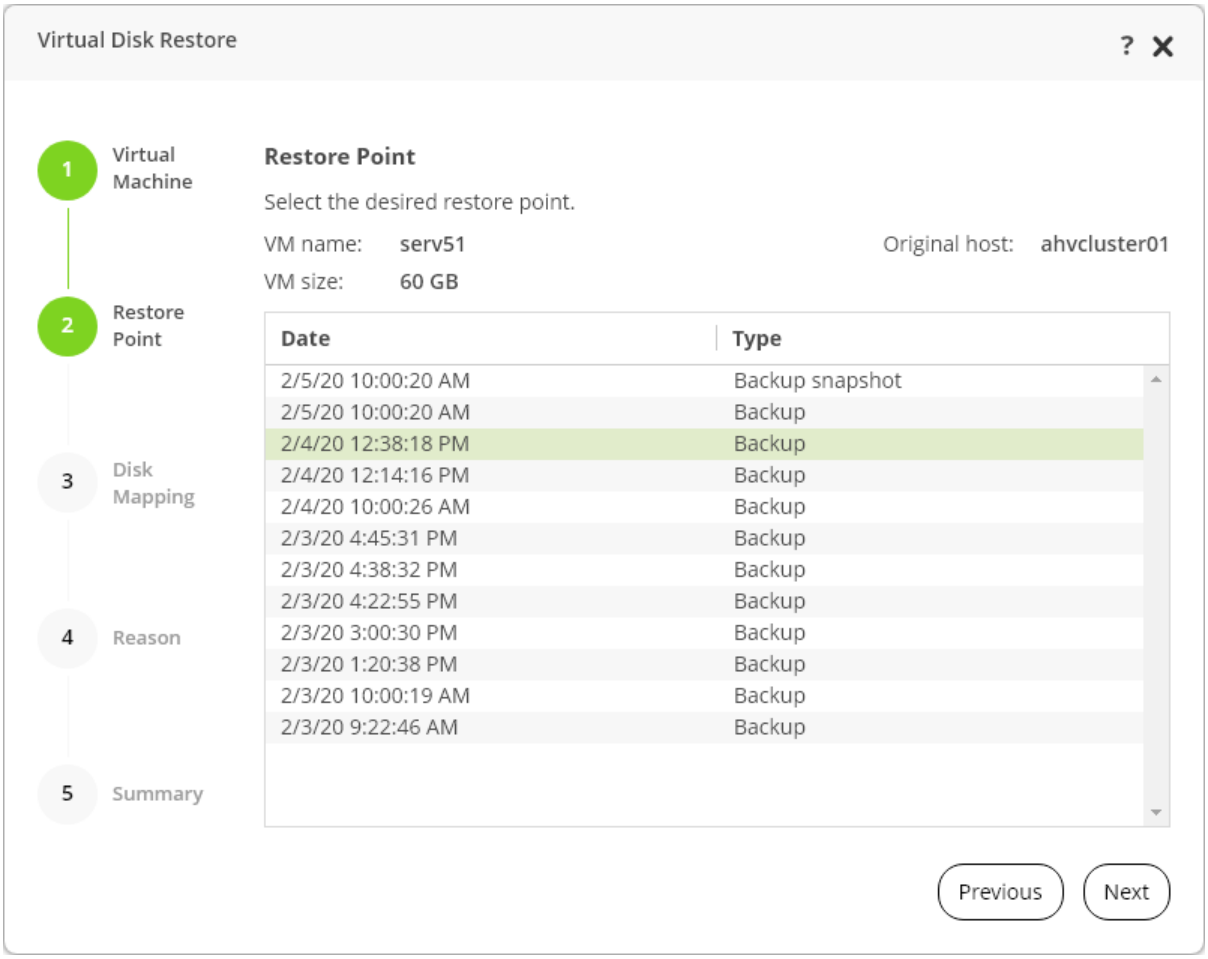

NOTE

Veeam Backup for Nutanix AHV stores the last snapshot of VM for each job. For example, if a VM is protected by two jobs, the appliance will store two snapshots.

Step 3. Configure Mapping Settings

At the **Disk Mapping** step of the wizard, configure VM disk mapping settings.

1. Click the **Browse** button and select a VM to which you want to attach the restored disk.

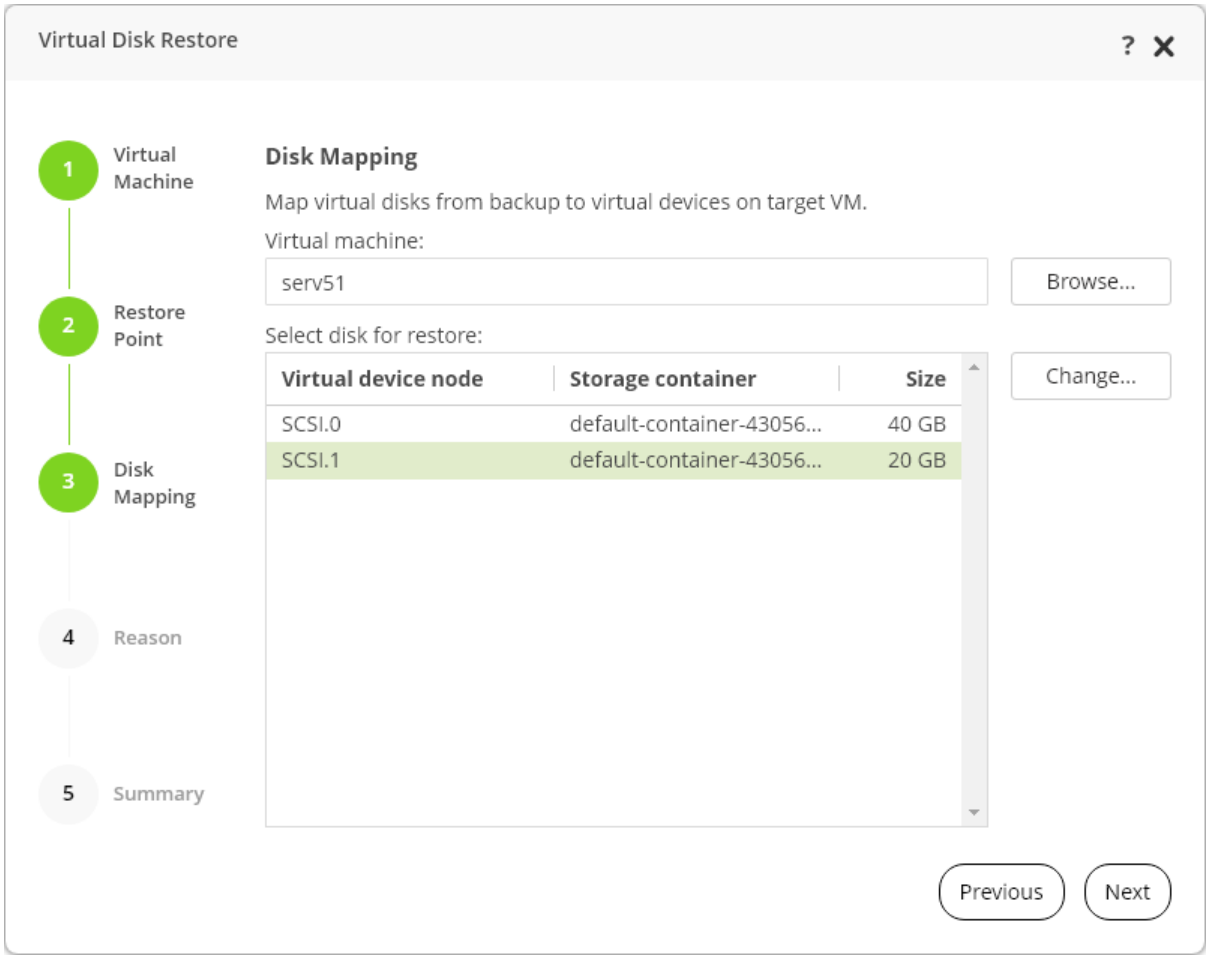

> **NOTE**
>
> You cannot select disks which were backed up within a Volume Group.

2. [Optional] If you want to change the target container, disk address or the bus type, do the following:
   a. Select the disk from the **Virtual device node** column and click the **Change** button.
   b. At the **Virtual Disk Properties** window, click **Choose** to select another container.
   c. From the **Bus Type** drop-down list, select the bus type for the restored disk. You can select from the following options: SCSI, PCI, IDE, SATA.
   d. From the **Disk Address** list, select the disk address.
      - If you select the default disk address, after successful restore process, the existing virtual disk will be replaced.
      - If you select a non-default disk address, the restored virtual disk will be added to the VM as a new disk.

e. Click **OK**.

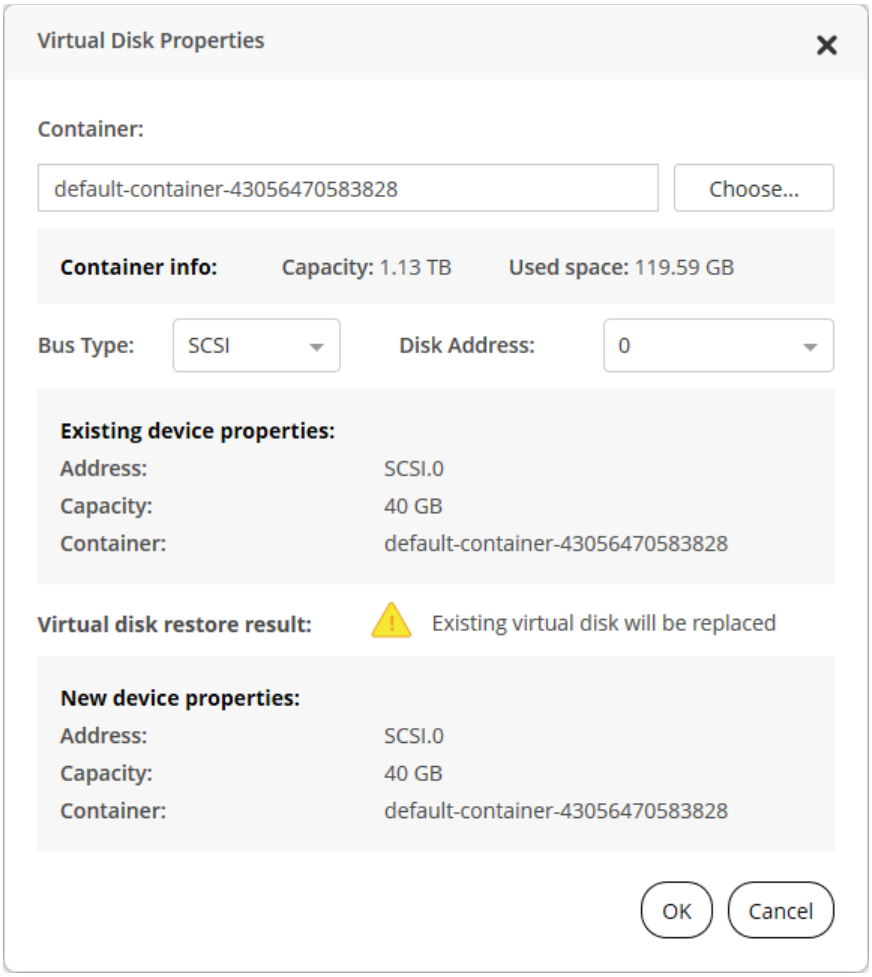
Virtual Disk Properties
Container:
default-container-43056470583828
Choose...
Container info:
Capacity: 1.13 TB
Used space: 119.59 GB
Bus Type:
SCSI
Disk Address:
0
Existing device properties:
Address:
SCSI.0
Capacity:
40 GB
Container:
default-container-43056470583828
Virtual disk restore result:
Existing virtual disk will be replaced
New device properties:
Address:
SCSI.0
Capacity:
40 GB
Container:
default-container-43056470583828
OK
Cancel

Step 4. Specify Reason for Restore

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

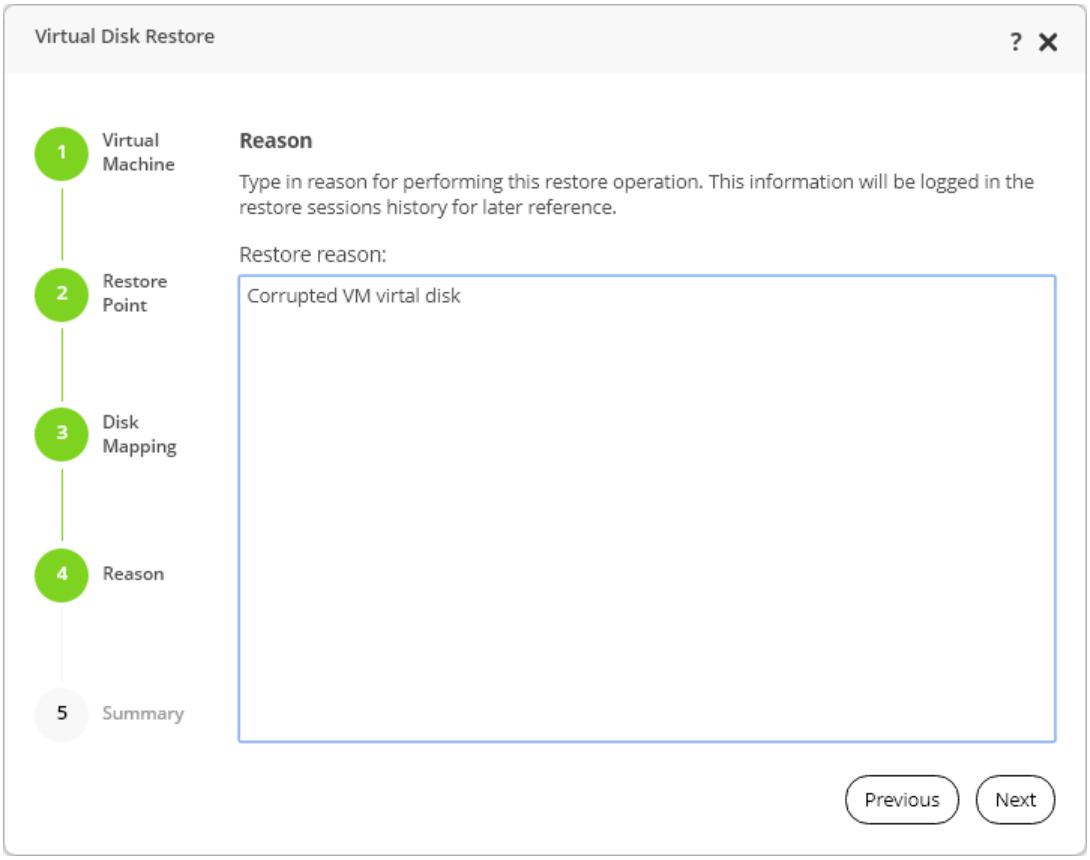

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured restore settings and click **Finish** to start the disk restore process.

After you click **Finish**, AHV Backup Proxy turns off the target VM and turns it on after the restore process completes.

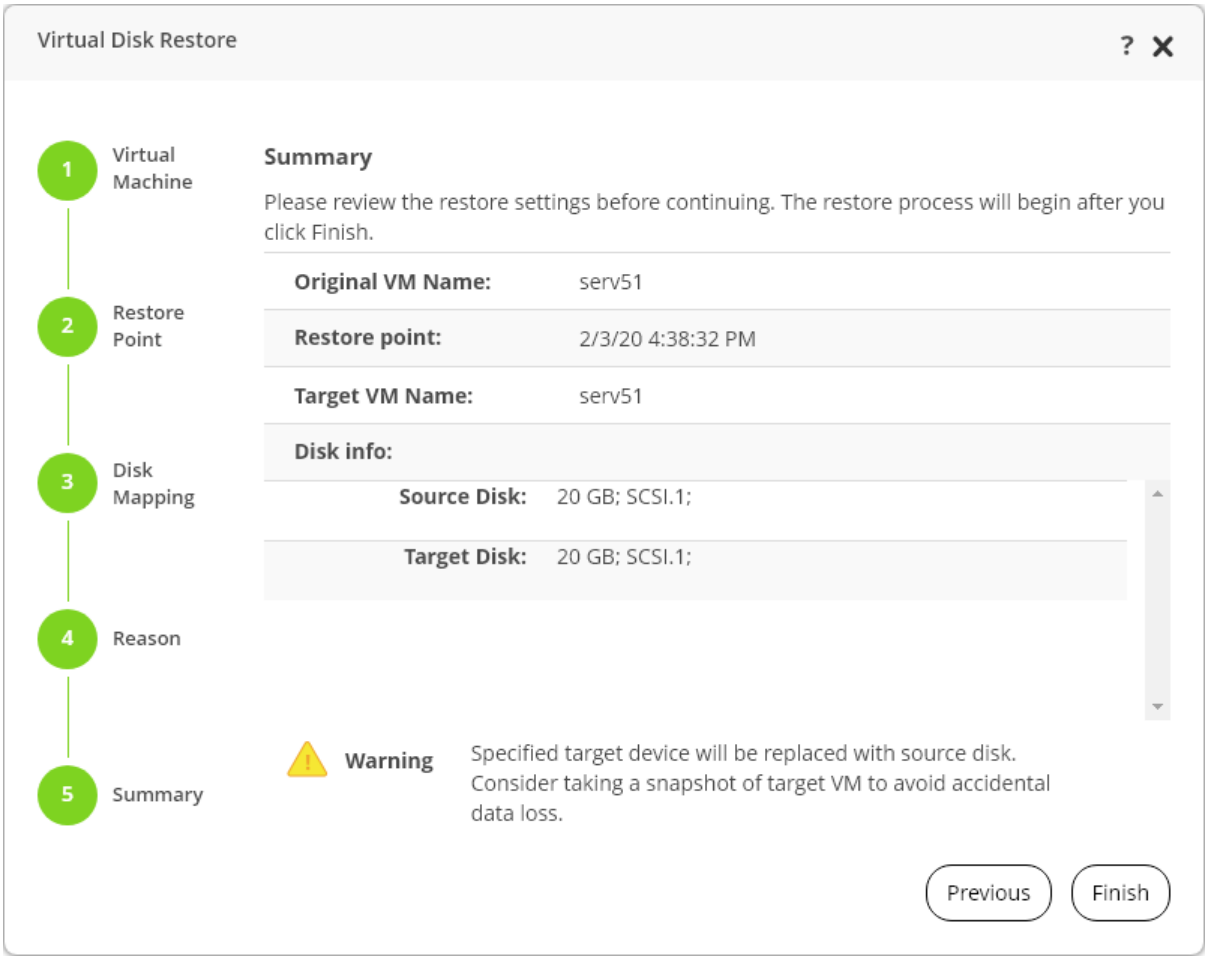

Instant Recovery

With the Instant Recovery feature of Veeam Backup & Replication, you can immediately restore different workloads as VMs to your production environment by running it directly from compressed and deduplicated backup files. Instant VM Recovery helps to improve recovery time objectives, minimize disruption and downtime of production VMs.

To perform instant recovery to Nutanix AHV, Microsoft Hyper-V or VMware vSphere, follow the steps of the **Instant Recovery** wizard. Depending on the target platform, see one of the following sections:

- Performing Instant Recovery of Workloads to Nutanix AHV
- Performing Instant Recovery of Workloads to VMware vSphere
- Performing Instant Recovery of Workloads to Hyper-V

Performing Instant Recovery of Workloads to Nutanix AHV

You can immediately restore virtual or physical machines into a Nutanix AHV cluster by running it directly from a compressed and deduplicated backup files. Instant Recovery feature helps improve recovery time objectives, minimize disruption and downtime of production VMs.

Supported Workloads

To recover machines to a Nutanix AHV cluster, you can use the following backups:

- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV
- Backups of Microsoft Hyper-V and VMware vSphere VMs created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux
- Backups of VMs created by vCloud Director
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure VMs created by Veeam Backup for Microsoft Azure
- Backups of Google Cloud Platform VMs instances created by Veeam Backup for Google Cloud Platform
- Backups of RHV VMs created by Veeam Backup for Red Hat Virtualization

How to Perform Instant Recovery to AHV

Before you start Instant Recovery, check prerequisites. Then use the **Instant Recovery to Nutanix AHV** wizard to configure the process.

Before You Begin

Before you perform Instant Recovery, consider the following:

- We highly recommend to use a separate host for a mount server to free up the VBR server RAM and avoid SQL issues due to lack of RAM.

  Minimum 512 MB of RAM on a mount server is required per each VM disk that you want instantly recover. For example, if you restore a VM with 4 disks, you need additional 2 GB of RAM on a mount server.
- Make sure that all NFS services except vPower NFS Services are stopped.
- ARM architecture VMs are not supported.
- Instant Recovery to Nutanix AHV is supported for Nutanix AOS 6.0 and later.
- Instant Recovery is not supported for backups made in 1.0 version of Veeam Aviability for Nutanix AHV. In this case, you can perform only entire VM restore.
- Instant Recovery from tapes, Archive Tier, Cloud Connect to Nutanix AHV is not supported.
- Instant Recovery is not supported for file-level backups created by the Kasten platform, Veeam Agent for Linux, Veeam Agent for Microsoft Windows, Veeam Agent for Unix, Veeam Agent for Mac.

- You cannot perform Instant Recovery to Nutanix AHV from *ObjectStorage (Imported) and ObjectStorage (VeeamZIP) nodes*.
- You must add the Nutanix AHV cluster (to which you want to restore machines) to the backup infrastructure. AHV Backup Proxy it is not required for Instant Recovery.
- All required drivers should be installed initially before the instant recovery process. You cannot add or modify drivers in the VM during the instant recovery launch.
- Veeam Backup for Nutanix AHV requires 64 MB of RAM for a VM via Instant Restore. If you have a VM with less than 64 MB of RAM, Veeam Backup for Nutanix AHV will increase this RAM to 64 MB during the restore process.
- You can restore a machine from a backup that has at least one successfully created restore point.
- If you restore a machine to the production network, make sure that the original machine is powered off to avoid conflicts.
- If a restored VM has an attached volume group, the disks from the volume group will not be restored.
- If you restore a Windows machine from a non-AHV backup, the SCSI drivers, NutanixGuestTools, they will be not installed on the published Windows VM. If you plan to perform Instant Recovery of such Windows VMs to an AHV cluster, install SCSI drivers and NutanixGuestTools on this machine before you back up the machine.
- When you instantly restore a VM from old backups, the restored VM has default virtual hardware settings: 2 CPU cores, 4 GB RAM and 1 network adapter. If you want to change the default settings:
  - Turn off the VM. Note that you must not switch off the instant recovery session before turning off the VM.
  - Set the required virtual resources (CPU/RAM/bootable disk) in the Nutanix AHV cluster web console.
  - If it is necessary, open CVM terminal via SSH and update boot configuration via acli:

```
admin@NTNX-20FM3F010153-B-CVM:{ip_arrd}:~$ acli
<acropolis> vm.update {vm_name} uefi_boot=true
```

Step 1. Launch Instant Recovery to Nutanix AHV Wizard

To launch the **Instant Recovery to Nutanix AHV** 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 to Nutanix AHV** — if you want to restore VMware vSphere VMs from a VM backup created with Veeam Backup & Replication.
  - **Agent** > **Entire machine restore** > **Instant Recovery to Nutanix AHV** — if you want to restore physical machines from a backup created with <%EP%> or <%VAL%>.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand a backup and select workloads that you want to restore. Then click **Instant Recovery** on the ribbon or right-click one of the workloads and select **Instant recovery**.

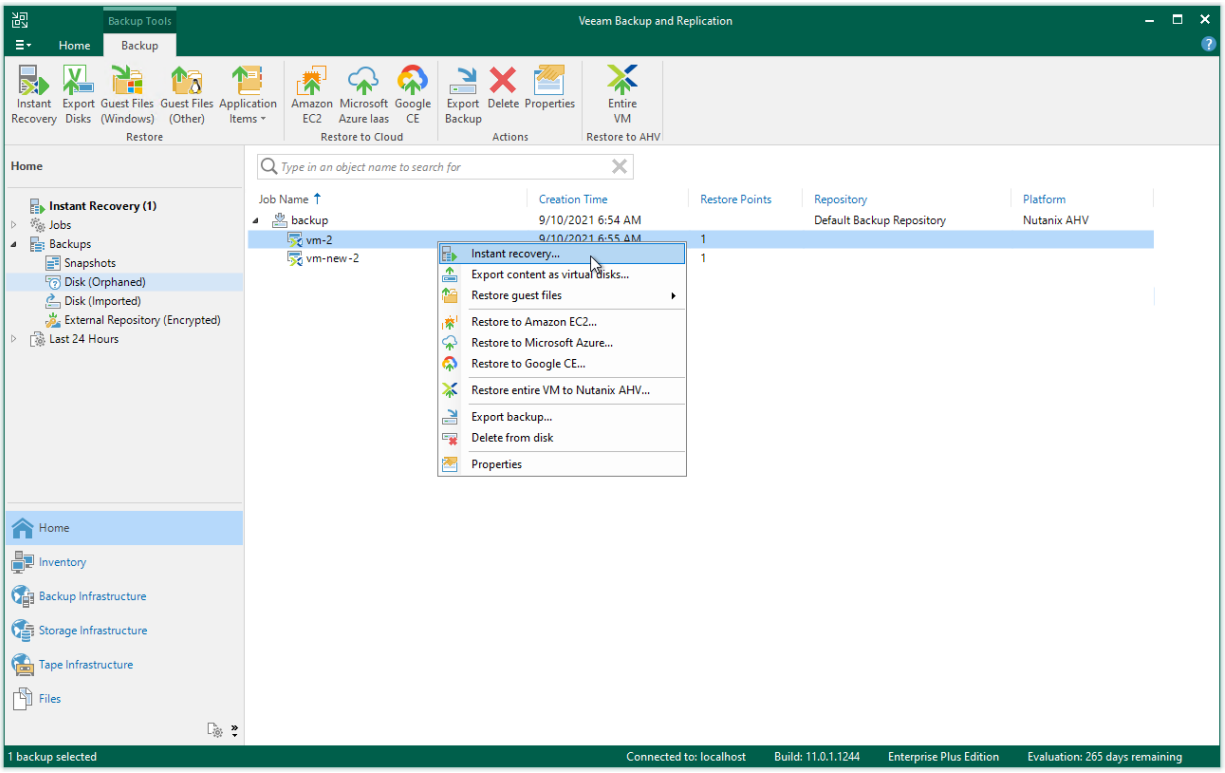

Step 2. Select VM and Restore Point

At the **Machines** step of the wizard, you can add more workloads that you want to recover. To select workloads:

1. Click **Add and select where to** browse for the necessary workloads:

   - **From infrastructure** – use this option to browse the virtual environment and select workloads or workload containers to restore. If you choose a workload container, Veeam Backup & Replication will expand it to a plain workload list.

     When you add a workload 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 workload, Veeam Backup & Replication will display a warning next to this workload.

   - **From backup** – browse existing backups and select workloads under backup jobs.

   In the **Backup Browser** window, select workloads you want to restore and click **Add**.

   To quickly find workloads within the created backups, you can use the search field:

   - Enter a workload name or a part of it in the search field. Veeam Backup & Replication will display matches.
   - If the workload is not in the list, click the **Show more** link to browse the backups.

2. Click **Point to** select a 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.

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

> **TIP**
>
> A name of a backup repository or object storage repository where a restore point resides is displayed in brackets right after a name of a backup job.

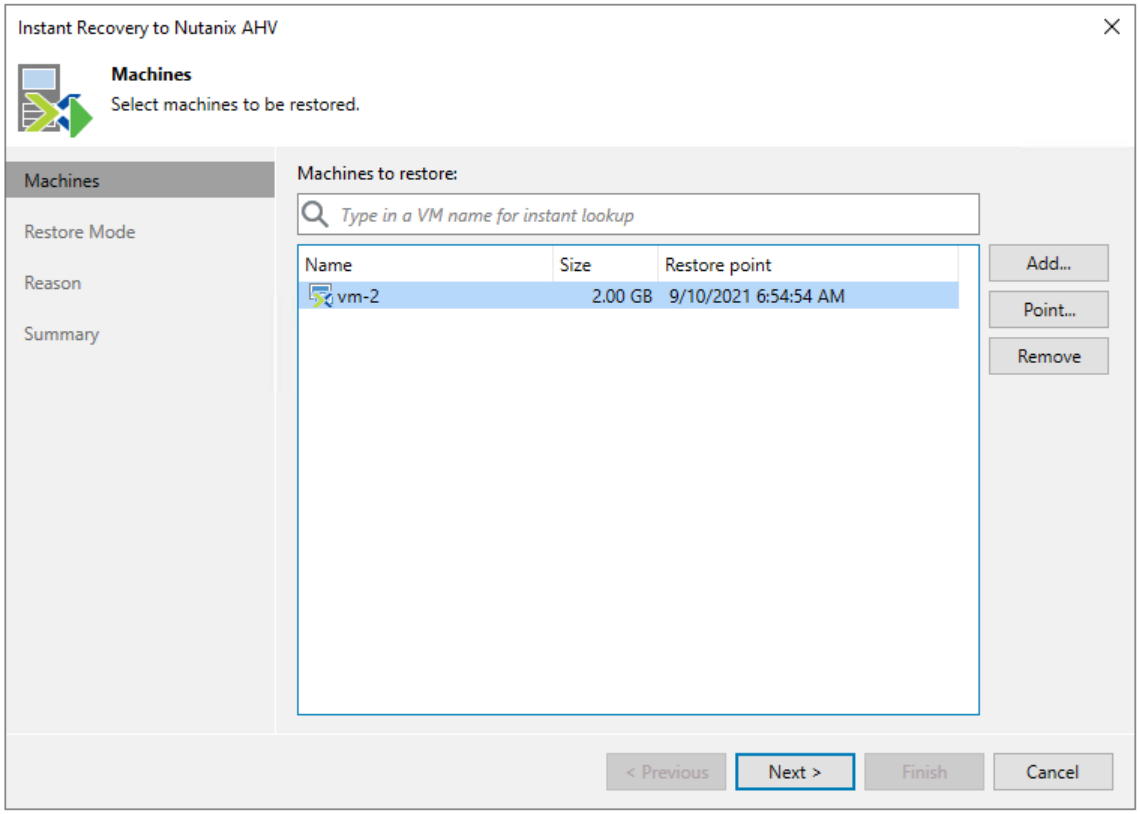

Step 3. Select Restore Mode

This step is available only if you restore Nutanix AHV VMs.

At the **Restore Mode** step of the wizard, specify how VMs will be restored:

- Select **Restore to the original location** if you want to restore 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.
- Select **Restore to a new location, or with different settings** if you want to restore VMs to a different location and/or with different settings (such as VM location, network settings, format of restored virtual disks and so on). If this option is selected, the **Instant Recovery to Nutanix AHV** wizard will include additional steps for customizing VM settings.

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

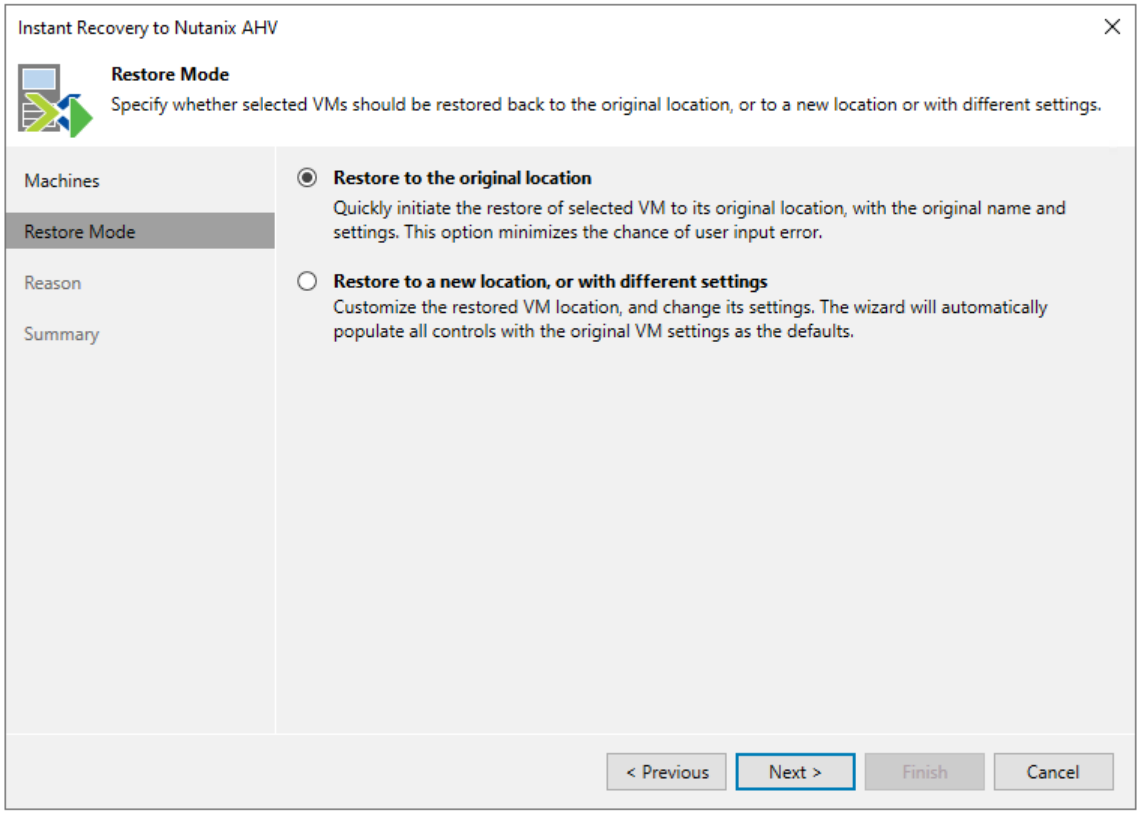

Step 4. Select Target Cluster

At the **Cluster** step of the wizard, specify a Nutanix AHV cluster where you want to locate the recovered VMs:

1. In the **VM location** list, select one or more VMs and click **Cluster**.
2. In the **Select Cluster** window, select a standalone host or cluster where the selected VMs must be registered.

   To quickly find the necessary object, use the search field at the bottom of the **Select Cluster** window:

   a. Click the button on the left of the field to select the necessary type of object that should be searched for.

   b. Enter an object name or a part of it in the search field.

   c. Click the **Start search** button on the right or press [Enter].

IMPORTANT

If a selected VM has an attached volume group, the disks from the volume group will not be restored.

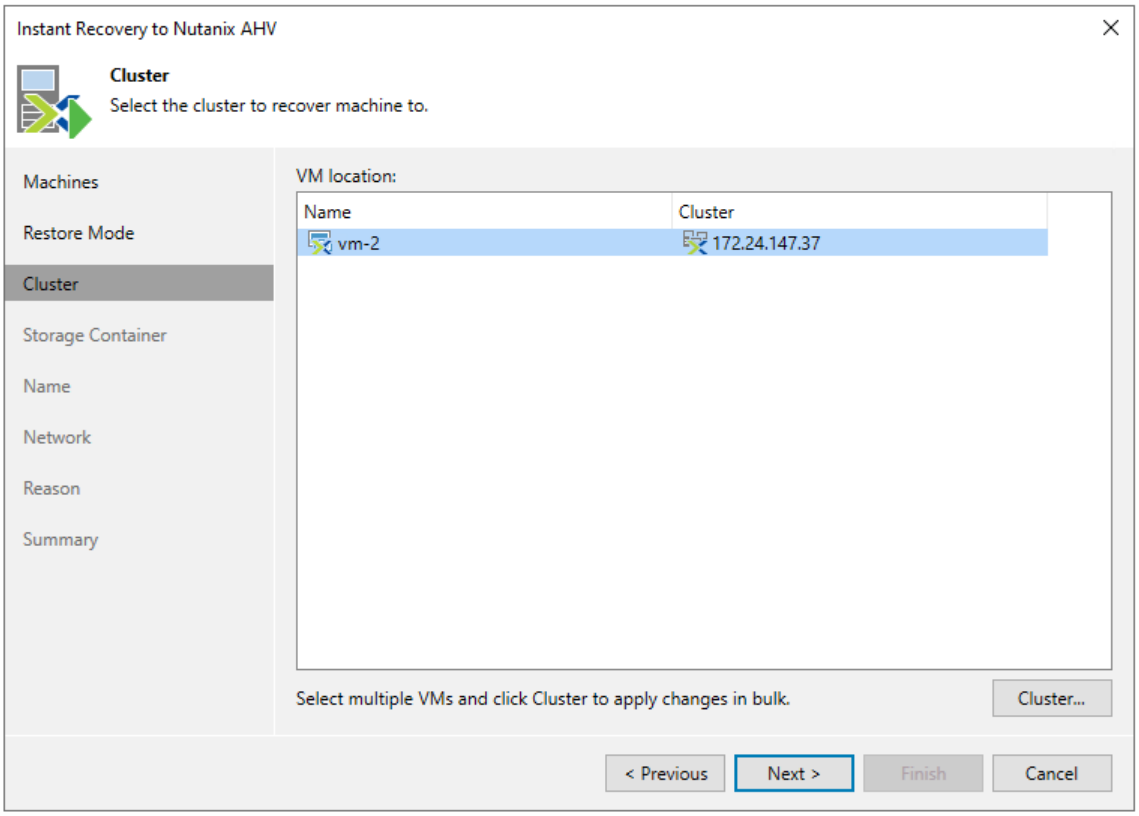

Step 5. Select Target Storage Container

At the **Storage Container** step of the wizard, specify a storage container where configuration files and disks of the recovered VMs will be stored. You can select a different container for each VM.

To specify a storage container:

1. Select one or more VMs in the list and click **Container**.
2. In the **Select Storage Container** window, select a container where VM files will be stored.

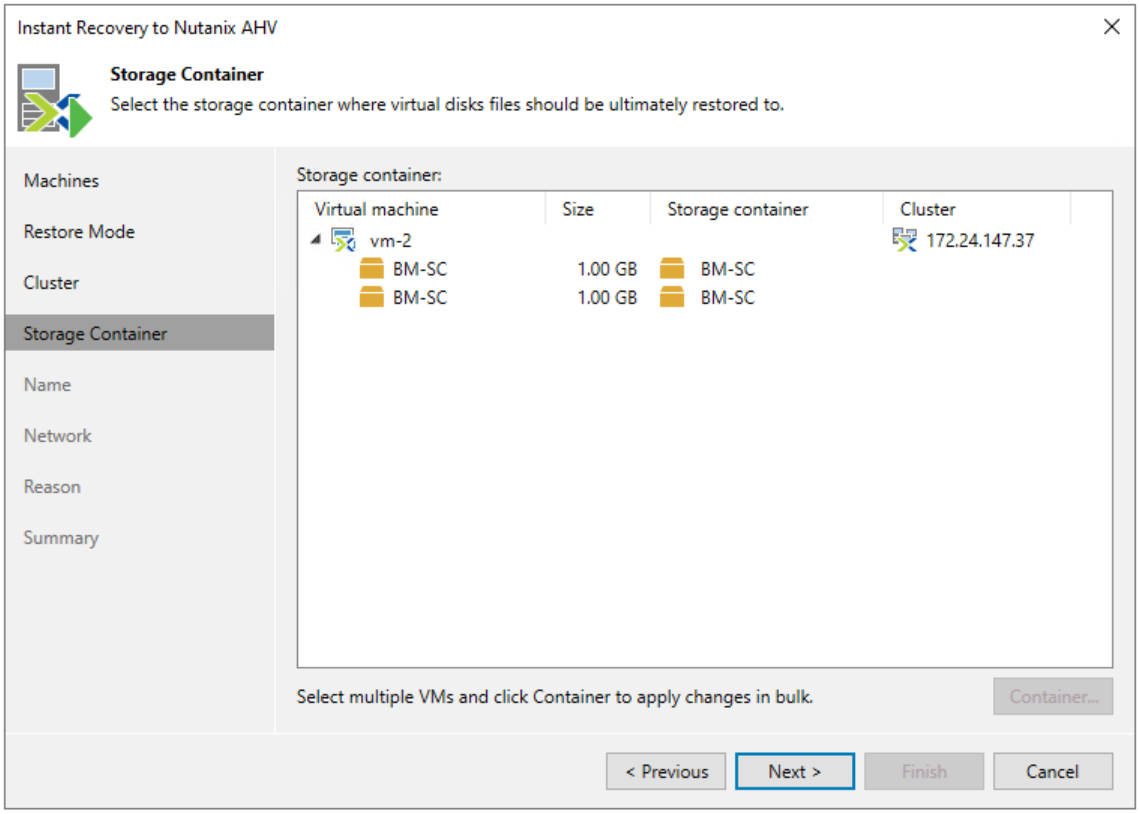

Step 6. Assign Name for Restored VM

At the **Name** step of the wizard, specify a name of restored VMs. We recommend to specify a new name for a VM to prevent conflicts if the original workload still resides in the production environment. The name change is not required if the original workload no longer exists, for example, it was permanently deleted.

To change a VM name:

1. In the **Virtual machines** list, select one or more VMs and click **Name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule by adding a prefix or suffix to the original VM name.

You can also change VM names directly in the list:

1. Select a VM in the list.
2. Double click the **New Name** field and enter the name to be assigned to the recovered VM.

> **IMPORTANT**
>
> If a VM with the specified name already exists in the target AHV cluster, the restored VM will overwrite the existing VM. Thus, the existing VM will be deleted.

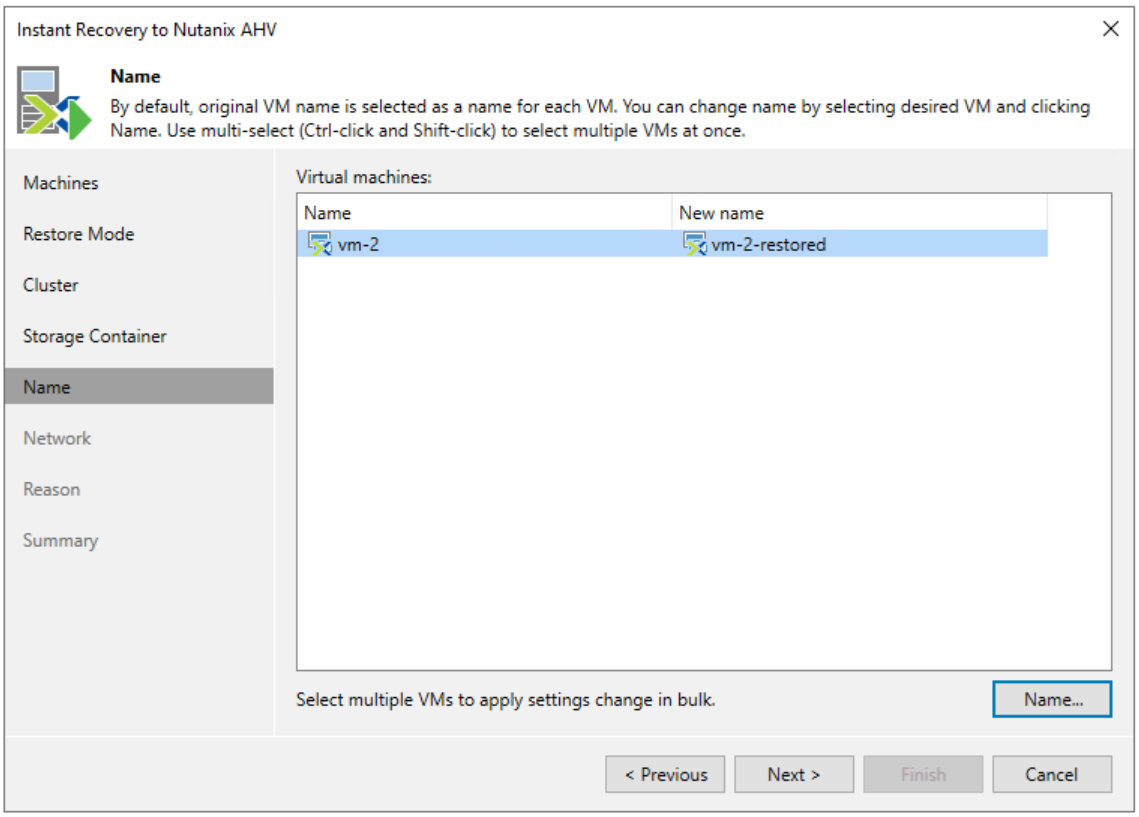

Step 7. Specify Network Mapping

At the **Network** step of the wizard specify network mapping settings. These settings allow you to map source networks, where the original VM resides, to target networks where the recovered VM will be registered. Veeam Backup & Replication will use the network mapping table to update configuration files of restored VMs on the fly, during the Instant VM Recovery process.

To specify network mapping settings:

1. In the **Network connections** list, select one or more 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 restore.

   To quickly find a network, enter a name or a part of it in the search field. Then click the **Start search** button or press [Enter].

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

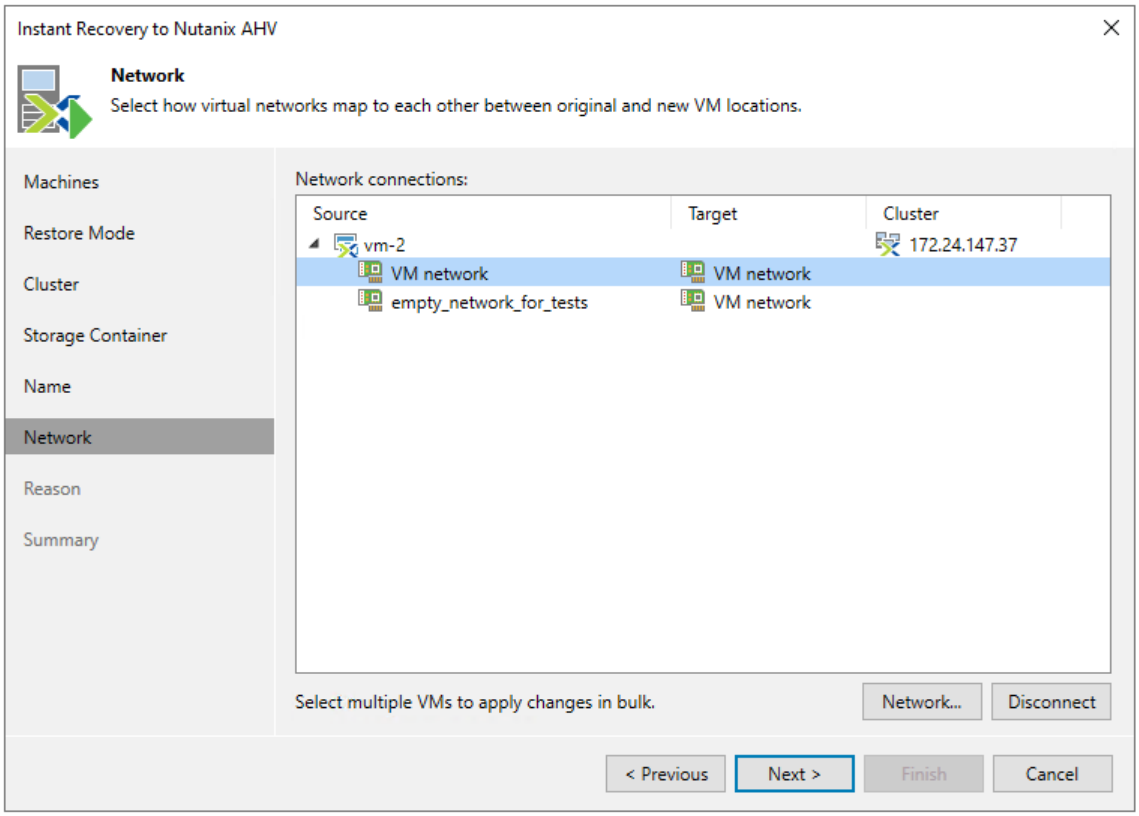

Step 8. Specify Restore Reason

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.

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.

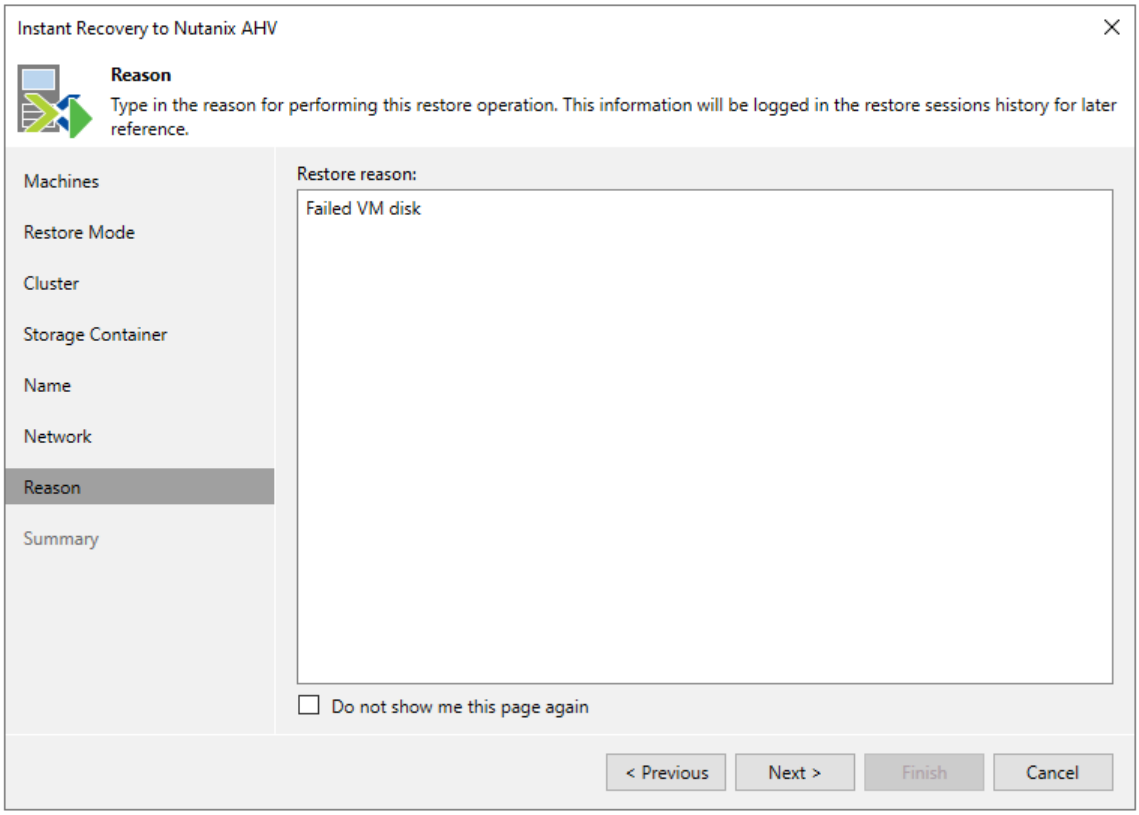

Step 9. Verify Instant Recovery Settings

At the **Summary** step of the wizard, check settings of the 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.

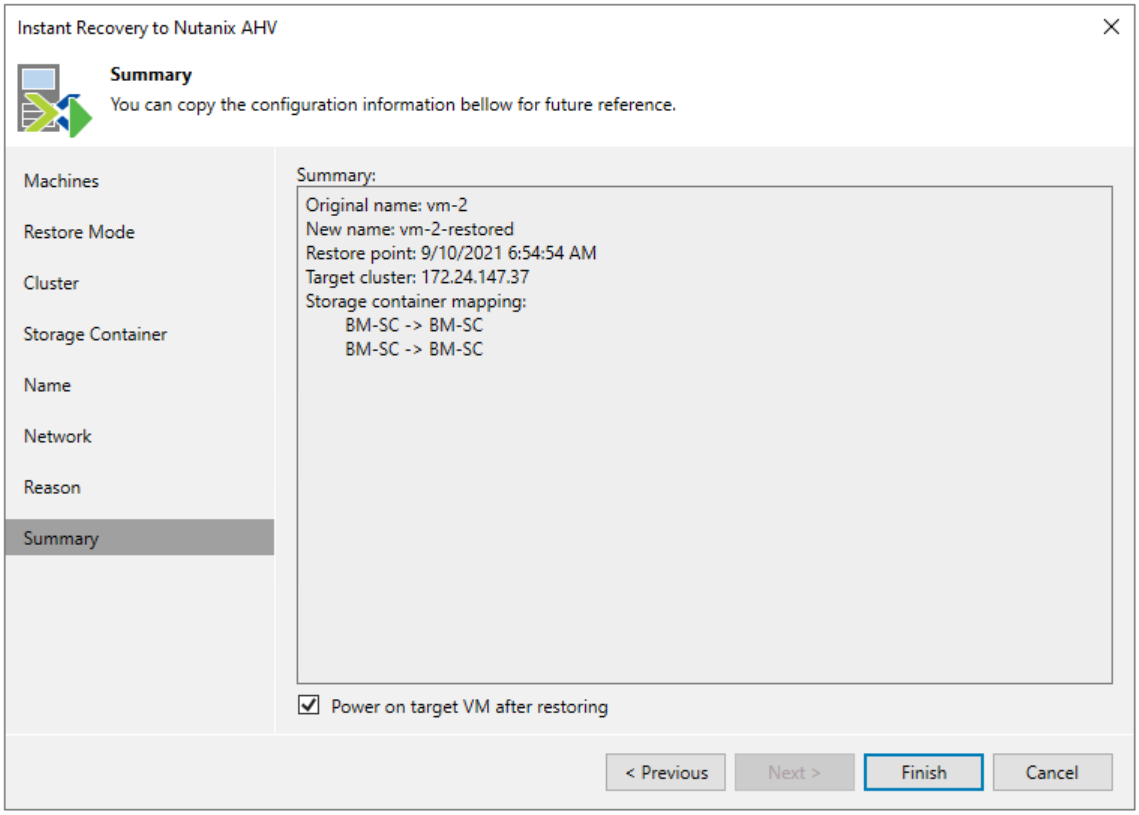

Step 10. Finalize Instant Recovery

After the VMs was recovered, you can finalize the process. Depending on the test results you can decide whether to migrate VMs to production environment or stop publishing.

Migrating Recovered VMs

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

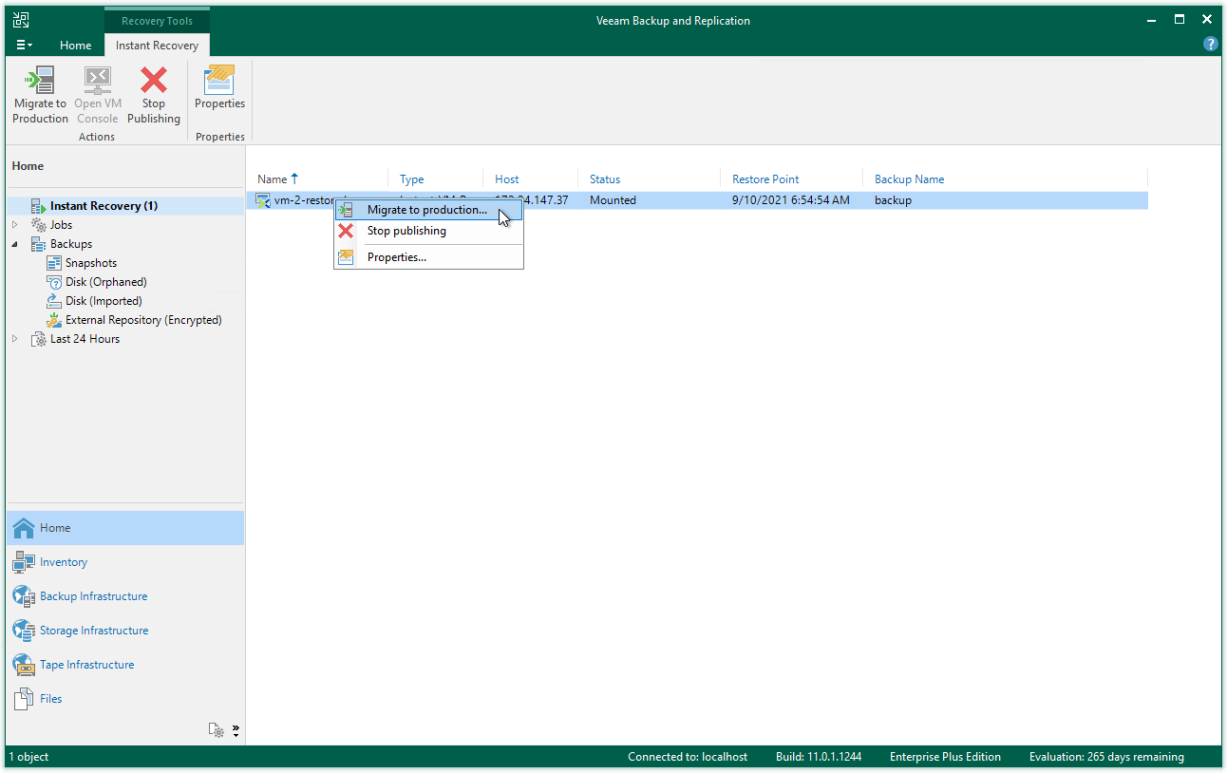

Stop Publishing Recovered VMs

If 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.

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

To remove a recovered VM, in the working area, right-click a VM and select **Stop publishing**

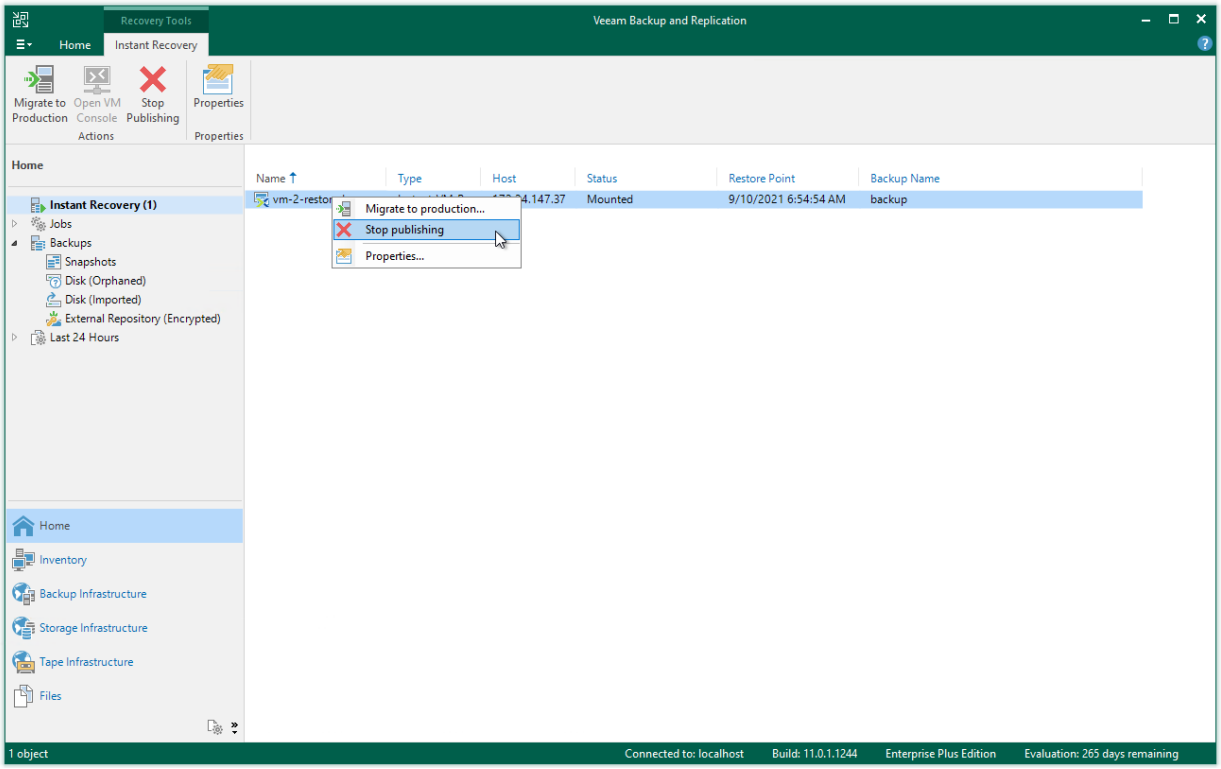

Performing Instant Recovery of Workloads to VMware vSphere

To perform instant recovery to VMware vSphere, follow the steps of the **Instant Recovery to VMware vSphere** wizard. For more information, see the Performing Instant Recovery of Workloads to VMware vSphere VMs section of the Veeam Backup & Replication User Guide.

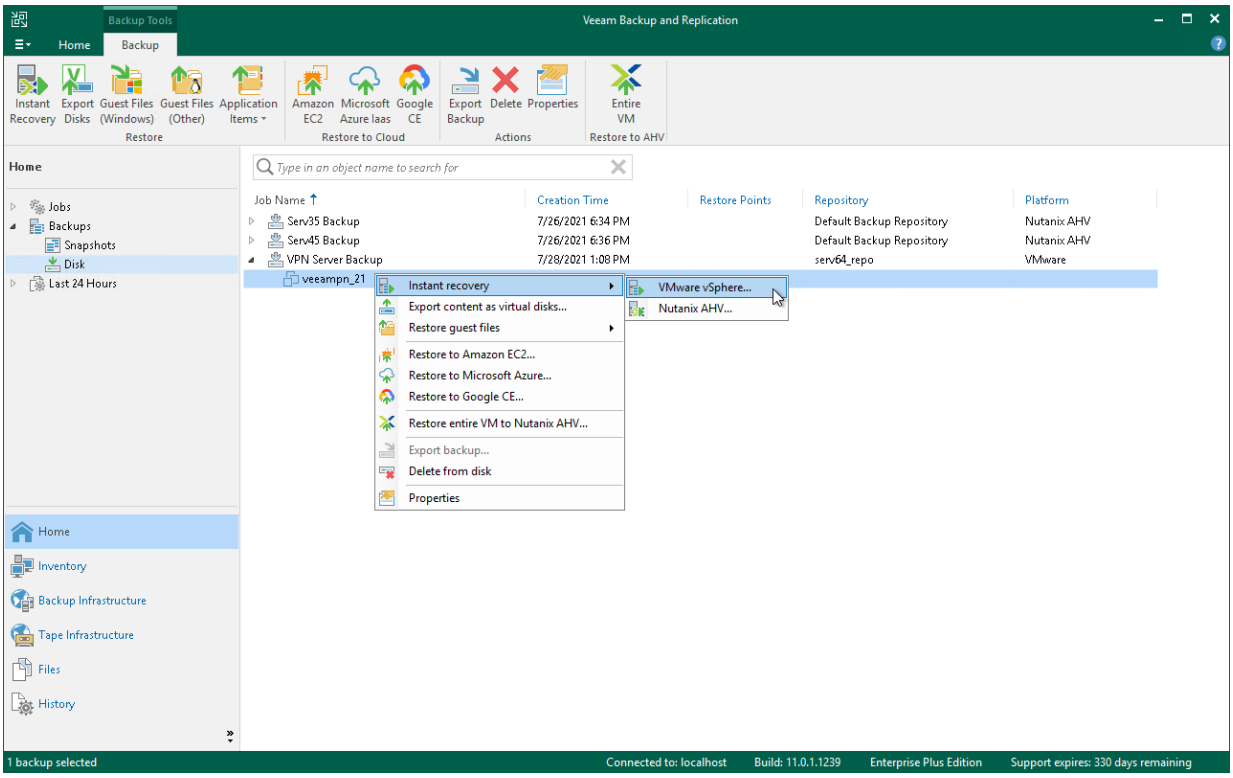

Performing Instant Recovery of Workloads to Hyper-V

To perform instant recovery to Microsoft Hyper-V environment, follow the steps of the **Instant Recovery to Microsoft Hyper-V** wizard. For more information, see the Performing Instant Recovery of Workloads to Hyper-V VMs section of the Veeam Backup & Replication User Guide.

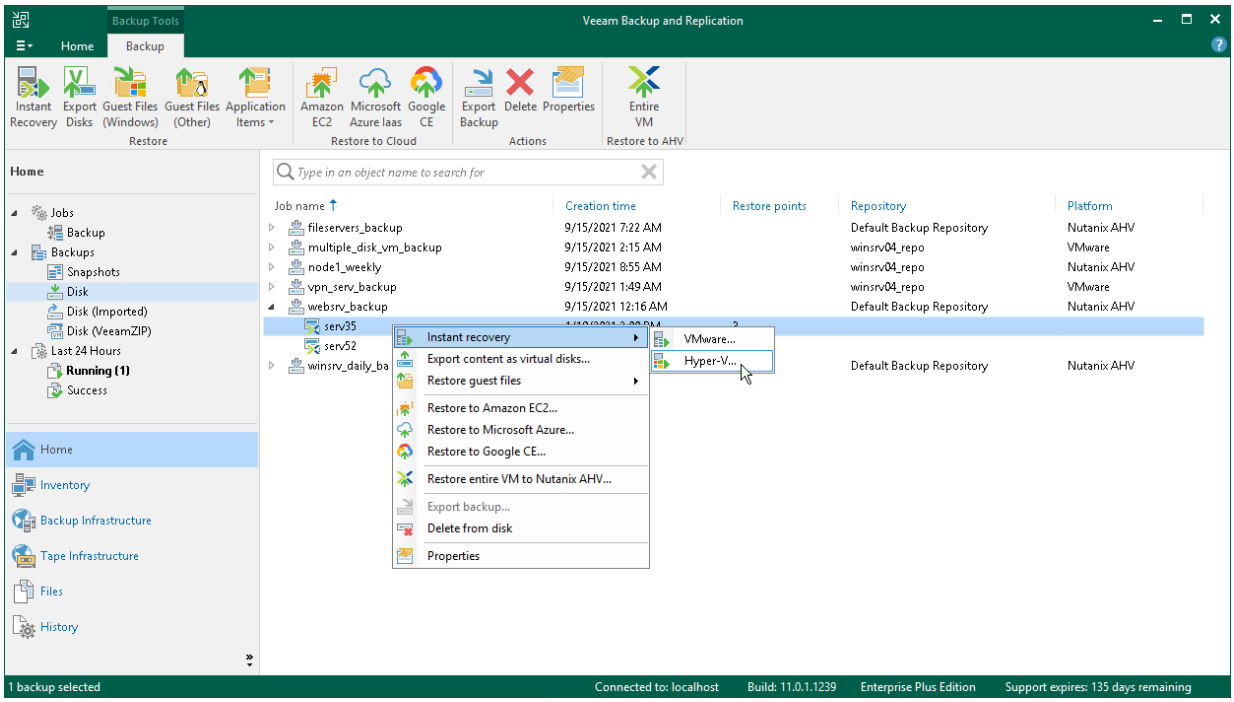

VM Guest OS Files Restore

You can use Veeam Backup & Replication to recover individual VM guest OS files and folders from VM backups and snapshots created by AHV Backup Proxy. VM guest OS file (or file-level) restore does not require you 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 snapshot to the necessary point in time. You can restore files and folders to original VM or save them to a local folder on the Veeam Backup & Replication server or to a network shared folder.

You can perform guest OS files restore for any VM guest OS file system. Veeam Backup & Replication offers different tools and methods for different file systems.

For guest OS files restore from FAT, NTFS, ReFS file systems and guest OS files restore from snapshots, Veeam Backup & Replication deploys a small temporary VM (helper appliance) on the AHV cluster. The helper appliance is a helper VM running a stripped down Linux kernel that has a minimal set of components. The appliance is around 50 MB. It requires 1024 MB RAM and takes around 10 seconds to boot. The helper appliance network settings are configured in the Appliance Settings step when you add a new AHV Nutanix AHV cluster.

To learn how Veeam Backup & Replication performs the guest OS files restore from backups and snapshots created by AHV Backup Proxy, see the following sections:

- File-Level Restore from Backups
- File-Level Restore from Backup Snapshot
- File-Level Restore from User Snapshots and PD Snapshots
- How to Restore VM Guest OS Files

File-Level Restore from Backups

You can use Veeam Backup & Replication to restore from the following file systems:

- **Restore from Linux, Unix and other file systems**: During restore from a backup, Veeam Backup & Replication mounts disks of the VM from the backup to a helper appliance (for Linux and other OSes), and launches the Veeam Backup Browser with the VM file system on the Veeam Backup & Replication server.
- **Restore from FAT, NTFS and ReFS**: During restore from a backup, Veeam Backup & Replication mounts disks of the VM from the backup to the backup server (for Microsoft Windows guest OS), and launches the Veeam Backup Browser with the VM file system on the Veeam Backup & Replication server. If you restore files to the original location, Veeam Backup & Replication creates an additional mount point on the mount server associated with the backup repository on which the backup file resides. The second mount lets you keep the VM traffic in one site and reduce load on the network.

File-Level Restore from Backup Snapshots

During restore from a backup snapshot, Veeam Backup & Replication mounts disks of the VM from the backup snapshot to the Veeam Backup & Replication server (for Microsoft Windows guest OS) or to a helper appliance (for Linux and Unix guest OSes), and launches the Veeam Backup Browser with the VM file system on the Veeam Backup & Replication server. In more detail, Veeam Backup & Replication does the following:

1. Deploys a helper appliance on the AHV cluster.
2. Uses Nutanix REST API to get information about a selected backup snapshot and files on that VM.
3. Creates a volume group with the VM files.

4. Attaches the VG to the helper appliance.
5. Powers on the helper appliance and mounts disks of a VM from the backup snapshot to the helper appliance (from Linux, Unix guest OS) or to the Veeam Backup & Replication server (for Windows guest OS).
6. Launches the Veeam Backup Browser with the VM file system.
7. When you finish restoring files and close Veeam Backup Browser, Veeam Backup & Replication detaches the volume group from the helper appliance, deletes the VG and unregisters the helper appliance from the AHV cluster.

File-Level Restore from User Snapshots and PD Snapshots

During restore from a user snapshot or a PD snapshot, Veeam Backup & Replication mounts disks of the VM from the snapshot to the Veeam Backup & Replication server (for Microsoft Windows guest OS) or to a helper appliance (for Linux and Unix guest OSes), and launches the Veeam Backup Browser with the VM file system on the Veeam Backup & Replication server. In more detail, Veeam Backup & Replication does the following:

1. Uses Nutanix REST API to get information about a selected snapshot.
2. Deploys a helper appliance on the AHV cluster.
3. Creates a temporary VM.
4. Creates a volume group (VG) using disks of this VM.
5. Attaches the volume groups to the helper appliance.
6. Deletes the temporary VM.
7. Powers on the helper appliance and mounts disks of a VM from the snapshot to the helper appliance (fro Linux, Unix guest OS) or to the Veeam Backup & Replication server (for Windows guest OS).
8. When you finish restoring files and close Veeam Backup Browser, Veeam Backup & Replication detaches the volume group from the helper appliance, deletes the VG and unregisters the helper appliance from the AHV cluster.

How to Restore VM Guest OS Files

To restore VM guest OS files, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, under **Backups** select **Disk** if you want to restore from a backup or **Snapshot** if you want to restore from a snapshot.
    3. In the working area, expand the necessary backup or snapshot node, right-click the VM whose VM guest OS files you want to restore and select **Restore guest files** > **Microsoft Windows/Linux and other OSes** and follow the steps of the **File Level Restore** wizard. For more information, see one of the following sections of the Veeam Backup & Replication User Guide:
    - Restoring VM Guest OS Files (FAT, NTFS or ReFS)

- Restoring VM Guest OS Files (Multi-OS)

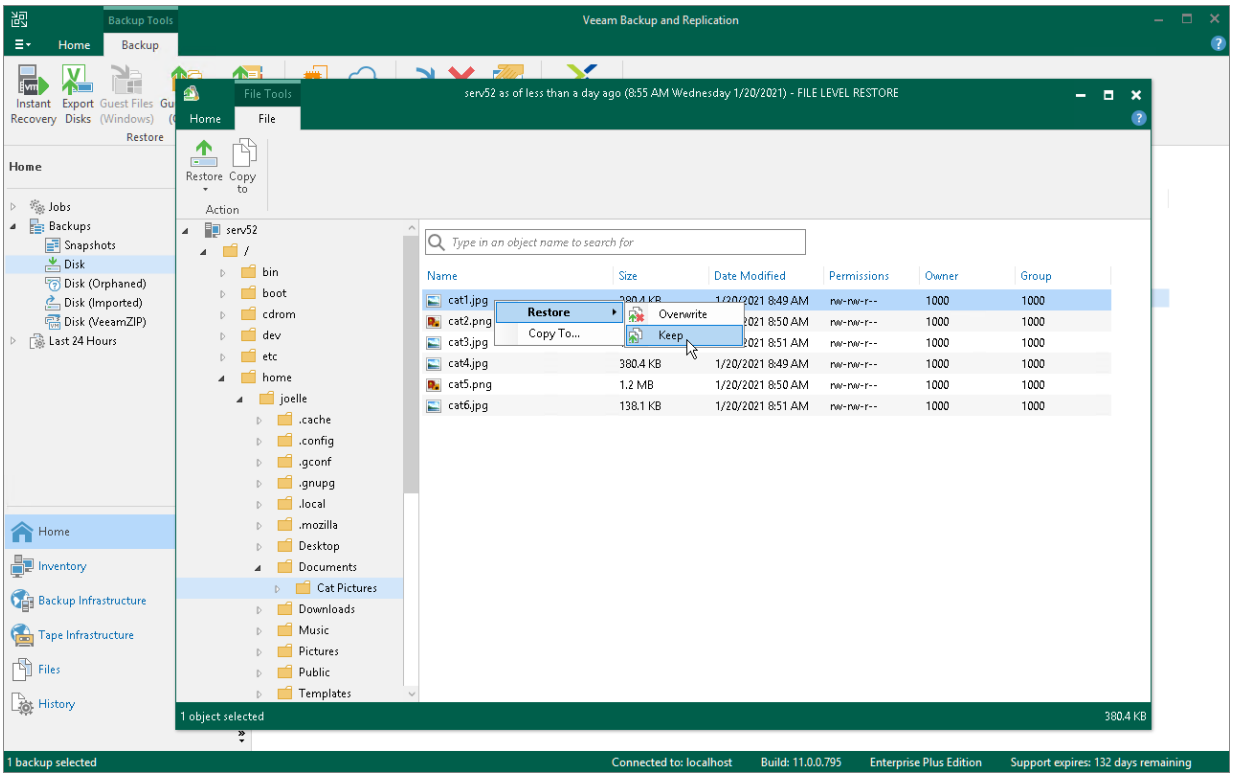

NOTE

You can also use Veeam Backup Enterprise Manager to restore guest OS files of AHV VMs. For more information, see the Restoring Guest OS Files section of the Veeam Backup Enterprise Manager User Guide.

Note that VM guest OS file indexing is not supported for Nutanix AHV VM backups.

Related Topics

Nutanix AHV Snapshot Types

Application Items Restore

Veeam Backup & Replication has auxiliary tools (Veeam Explorers) that let you restore items for the following applications:

- Microsoft Active Directory
- Microsoft Exchange
- Microsoft SharePoint
- Microsoft SQL Server
- Oracle Database

With the help of Veeam Explorers, you can restore Microsoft Active Directory objects and containers, Microsoft Exchange mailboxes, folders and messages, Microsoft SharePoint sites and lists, Microsoft SQL Server and Oracle databases.

> **IMPORTANT**
>
> The backup or the snapshot from which you want to restore must be application-consistent. AHV Backup Proxy creates application-consistent backups only if the VM meets requirements described in the Application-Consistent Snapshot Guidelines section of the Prism Web Console Guide.

To restore application items from backups or snapshots of Nutanix AHV VMs, perform the following steps:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select the **Backups** or **Snapshots** node.

   > **NOTE**
   >
   > Restore of application items from a snapshot requires a helper appliance. Make sure all required ports between the Veeam Backup & Replication server and the helper appliance are open. For more information, see Proxy Appliance (Multi-OS FLR) and Mount Server sections of the Veeam Backup & Replication User Guide.

3. In the working area, select the required machine in the backup or snapshot, click **Application Items** on the ribbon and select the application (Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server or Oracle Database).

   You can also right-click the machine or snapshot and select **Restore application items**.
4. Veeam Backup & Replication will open the application items wizard. You can use this wizard to automatically extract the application database from the backup and open it in Veeam Explorer.

5. In the **Backup Browser** window, click **Application Items** and select the necessary application from the list. Veeam Backup & Replication will start the Veeam Explorer for the selected application that will allow you to restore application items from the backup or snapshot.

   To be able to restore application items in the Veeam Explorer console, you must manually attach the database. For more information on working with Veeam Explorers, see Veeam Explorers User Guide.

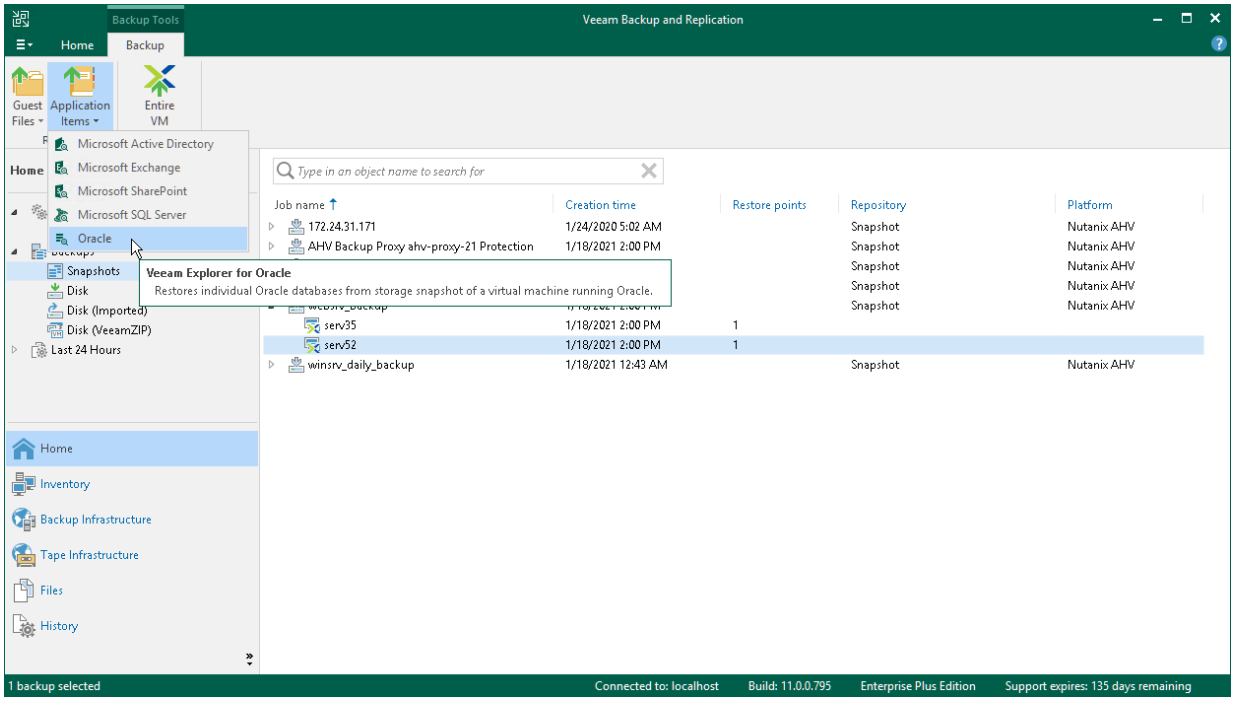

VM Disk Export

You can restore VM disks from backups created by AHV Backup Proxy and convert them to disks of the VMDK, VHD or VHDX format. You can store the exported disks locally on any server added to the backup infrastructure. You can also store the exported disks on a datastore connected to an ESXi host. After the disk export, you can mount the exported disk to any VM in the virtual infrastructure.

To export a VM disk:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups > Disk**.
3. In the working area, expand the necessary backup node, right-click the VM whose disk you want to restore and select **Export disk contents as virtual disks** and follow the steps of the **Export Disk** wizard. For more information, see the Exporting Disks section of the Veeam Backup & Replication User Guide.

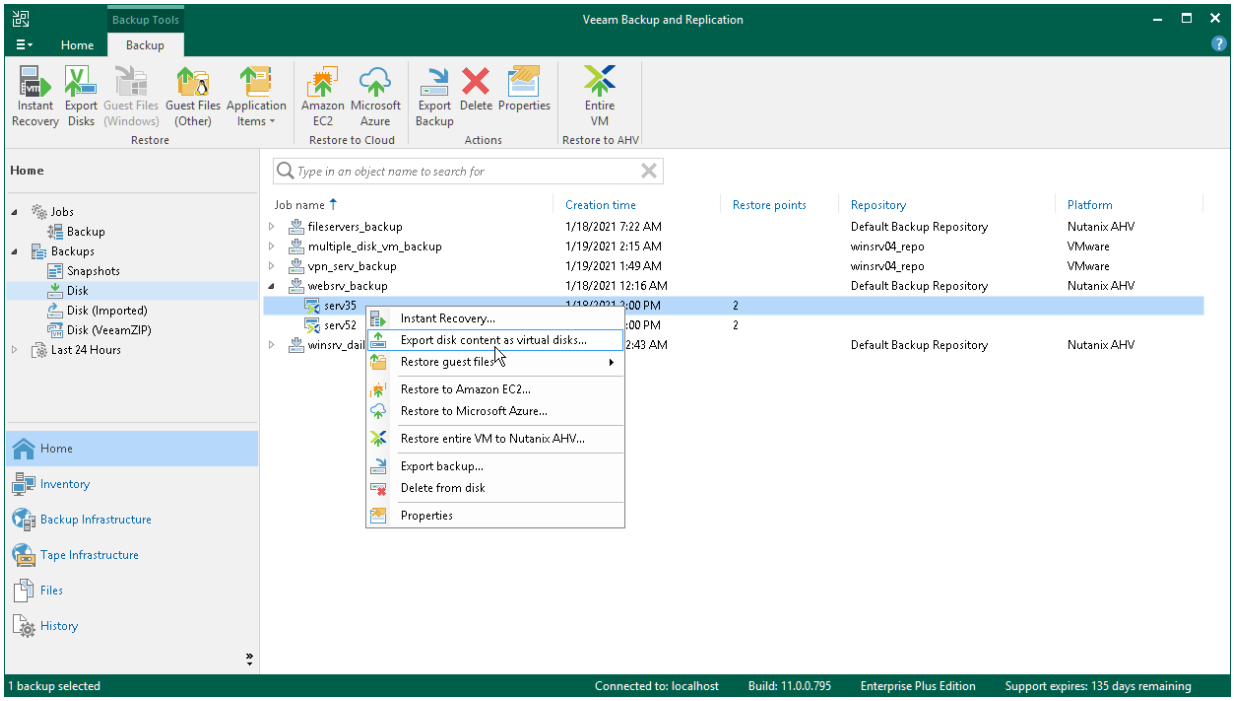

Restore to Amazon EC2

You can restore VMs from backup files created by AHV Backup Proxy to Amazon EC2 as EC2 instances.

To restore a VM to Amazon EC2, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups > Disk**.
3. In the working area, expand the necessary backup node, right-click the VM that you want to restore, select **Restore to Amazon EC2** and follow the steps of the **Restore to Amazon EC2** wizard. For more information, see the Restore to Amazon EC2 section of the Veeam Backup & Replication User Guide.

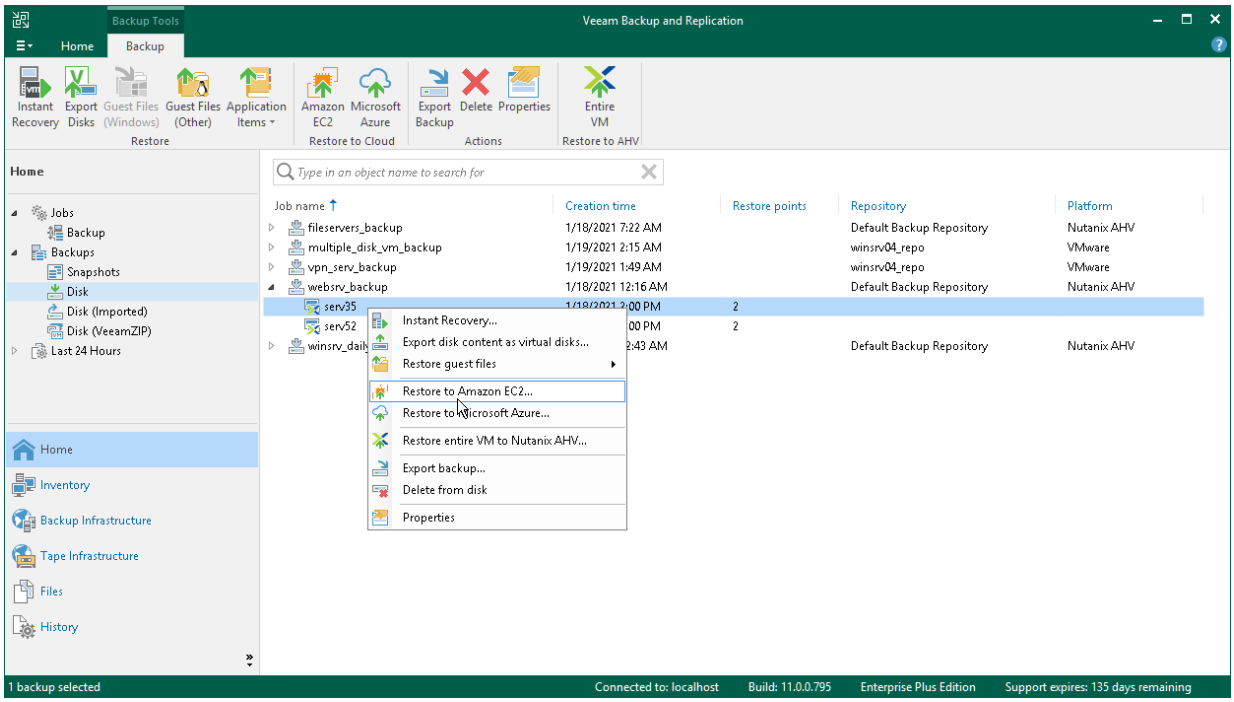

Restore to Microsoft Azure

You can restore VMs from backup files created by AHV Backup Proxy to Azure Blob storage as Microsoft Azure VMs.

To restore a VM to Microsoft Azure, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups > Disk**.
3. In the working area, expand the necessary backup node, right-click the VM that you want to restore, select **Restore to Microsoft Azure** and follow the steps of the **Restore to Azure** wizard. For more information, see the Restore to Microsoft Azure section of the Veeam Backup & Replication User Guide.

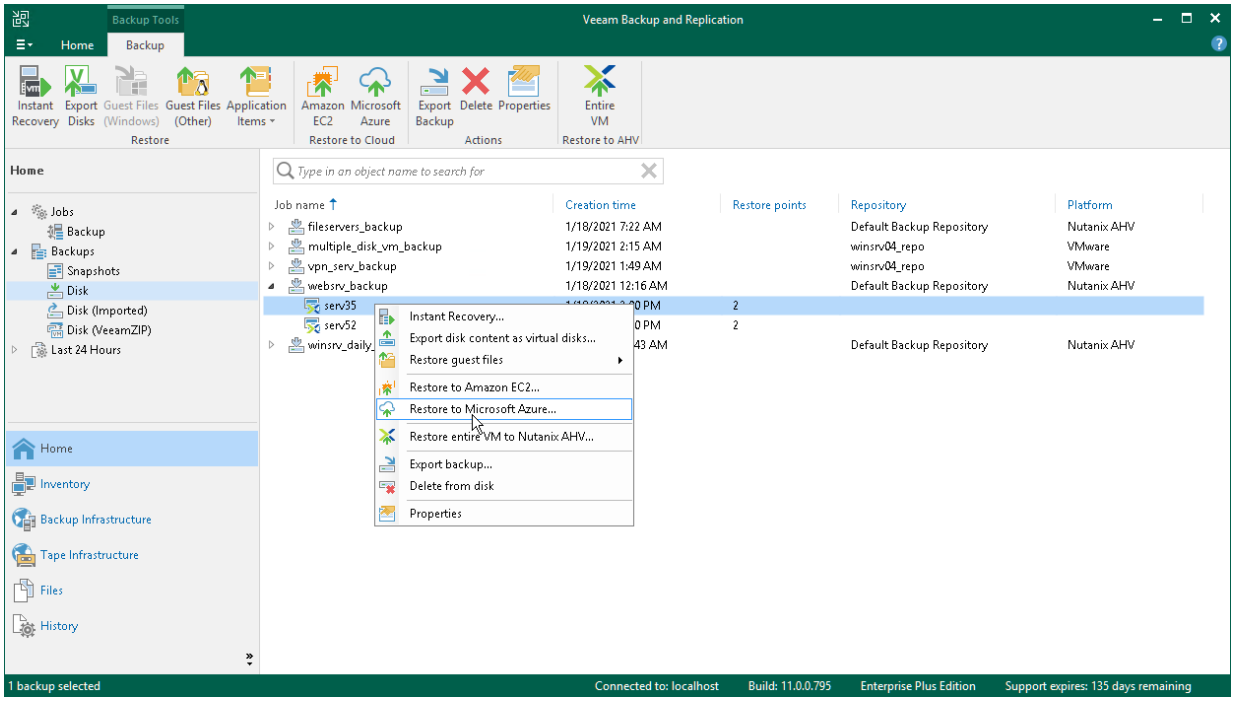

Logs and Support

If you have any questions or issues with Veeam Backup for Nutanix AHV, you can search for a resolution on Veeam R&D Forums or submit a support case in the Veeam Customer Support Portal.

When you submit a support case, we recommend to send all log files to ensure that comprehensive information is provided to Veeam Support Team.

To export logs of AHV Backup Proxy server, you can use either Veeam Backup & Replication or AHV Backup Proxy. Both options provide the same set of logs: logs of the AHV Backup Proxy service, Veeam Data Movers, VM Backup API and logs of the Veeam Updater component.

For instructions on how to export logs, see the following sections:

- Exporting Logs in AHV Backup Proxy Console
- Exporting Logs in Veeam Backup & Replication Console

Exporting Logs in AHV Backup Proxy

To export logs in the AHV Backup Proxy web console, do the following.

1. At the main panel, click the settings icon and select **Appliance Settings**.
2. At the **Summary** tab of the **Appliance Settings** view, click **Support Bundle**.

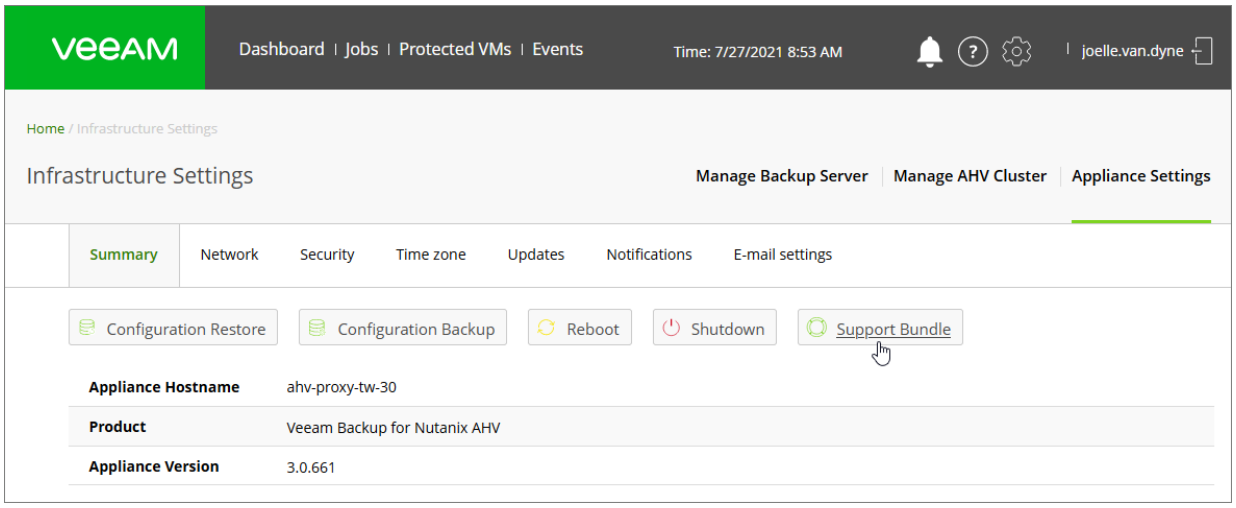

3. Select the necessary logs and click **Download**. You can select the following logs:
   - **Service logs**: logs of the main service of AHV Backup Proxy (Veeam Backup Agent Service).
   - **Agent logs**: logs of Veeam Data Movers that are responsible for data transfer between an AHV Backup Proxy and backup repositories.
   - **Library logs**: logs of VM Backup API.

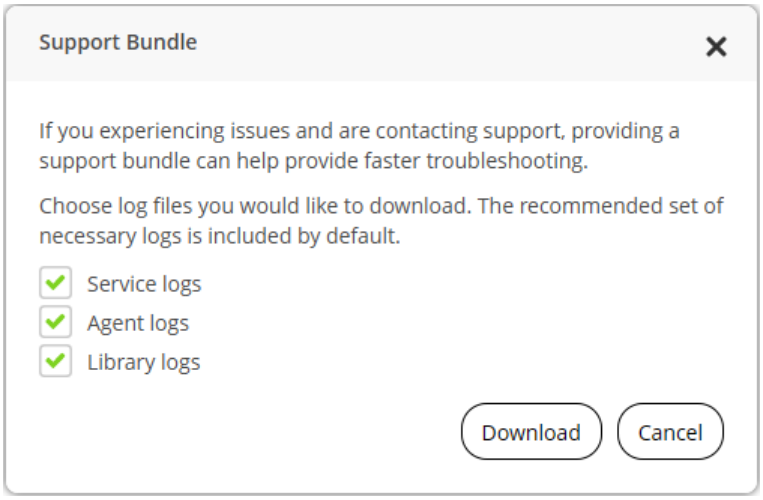

Exporting Logs in Veeam Backup & Replication

To export logs for AHV Backup Proxy VM, do the following:

1. From the main menu of the Veeam Backup & Replication console, select **Help > Support Information**.
2. At the **Scope** step of the **Export Logs** wizard, select the **Export all logs for selected components** option.
3. In the **Managed servers** list, select the AHV Backup Proxy VM.
4. Complete the wizard as described in the Exporting Logs section of the Veeam Backup & Replication User Guide.

> **NOTE**
>
> In the **Export Logs** wizard, you cannot export logs for separate AHV Backup Proxy backup jobs and backup files.

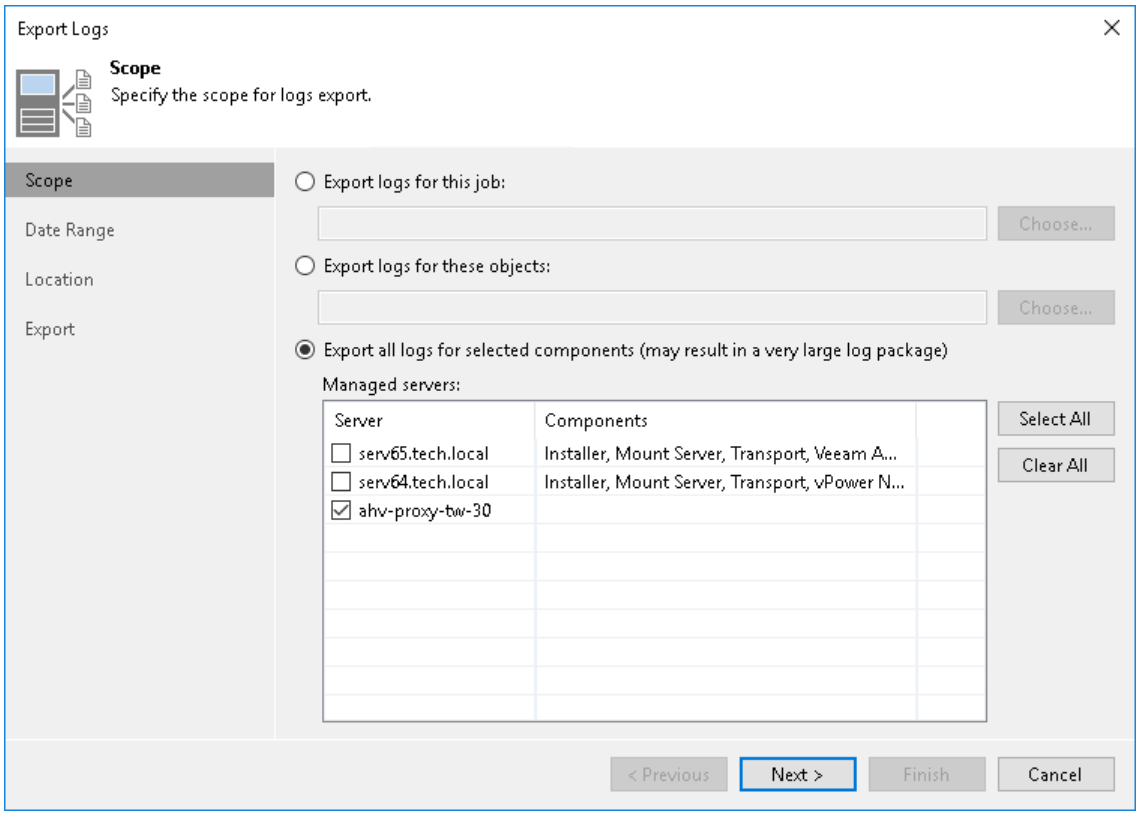

Appendices

See in this section:

- Appendix A. Manual Deployment
- Appendix B. Custom Scripts

Appendix A. Manual Deployment

Apart from main deployment scenario, you can also deploy AHV Backup Proxy manually.

For manual deployment, you must use the Nutanix Prism console to deploy AHV Backup Proxy on a Nutanix AHV cluster where user VMs that you plan to back up reside. One AHV Backup Proxy processes VMs within one Nutanix AHV cluster only. Thus, to back up and restore VMs on several Nutanix AHV clusters, you must deploy one AHV Backup Proxy on each cluster.

> **IMPORTANT**
>
> Before you start deploying AHV Backup Proxy, make sure that the virtual environment meets the requirements.
>
> - Platforms and virtual hardware must meet system requirements.
> - All registered server names must be resolvable into IPv4 addresses. IPV6 is not supported by the proxy appliance VM.

To deploy and configure AHV Backup Proxy on a Nutanix AHV cluster, do the following:

1. Upload the installation disk image of AHV Backup Proxy to the AHV cluster.
2. Create VM for AHV Backup Proxy.
3. Log in to the AHV Backup Proxy web console.
4. Install AHV Backup Proxy.
5. Add the AHV cluster to the AHV Backup Proxy infrastructure.
6. Add a Veeam Backup & Replication server.

Related Topics

Restoring Configuration Settings During Deploy

Step 1. Upload AHV Backup Proxy Virtual Disk Image

To deploy AHV Backup Proxy manually, you must upload the proxy installation virtual disk image to the target AHV cluster. You can find the installation image (`NutanixAHVPluginProxy_X.X.XX.vmdk`) in the `C:\Program Files\Veeam\Plugins\Nutanix AHV\Service\PlatformPlugins\AhvProxyManagement\Images` folder of the server where Nutanix AHV Plug-in is installed.

To deploy AHV Backup Proxy, do the following.

1. Log in to the Nutanix Prism web console (Prism Element console).

   > **NOTE**
   >
   > For more information on how to upload images on Nutanix Prism web console, see the Configuring Images section of the Prism Web Console Guide.
   >
   > If you use Nutanix Prism Central, image configuration instructions will be different. For more information, see the Adding an Image section of the Prism Central Guide.

2. In the main menu of Nutanix Prism web console, click the gear icon and select **Image Configuration**.
3. Click **Upload Image**.

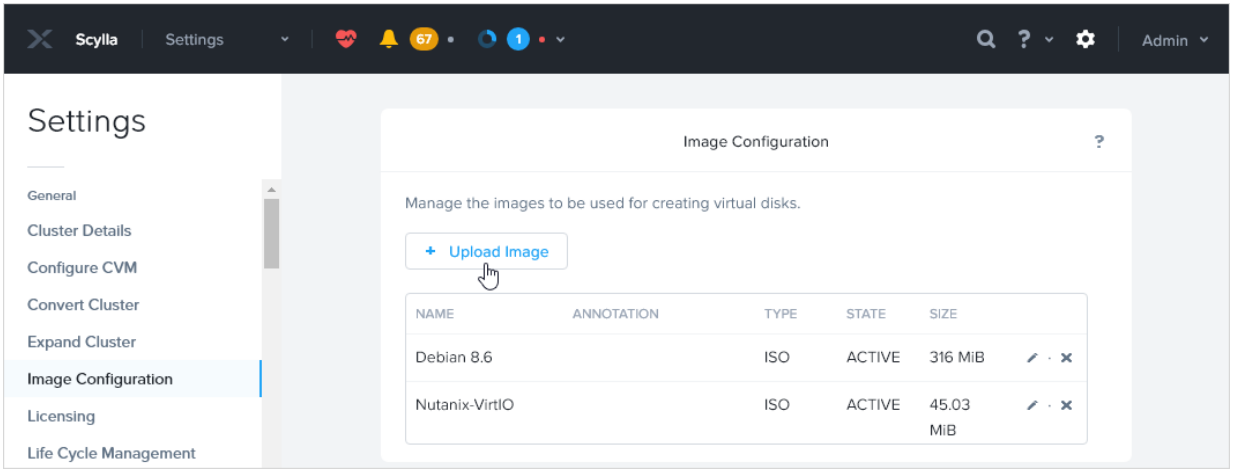

4. In the **Create Image** window, specify the image parameters:
   a. In the **Name** field, enter a name for the image.
   b. In the **Image Type** list, select **DISK**.
   c. In the **Storage Container** list, select the storage container where you want to store the image.
   d. Select the **Upload a file** option and click **Choose File**.

      You can copy the installation image (`NutanixAHVPluginProxy_X.X.XX.vmdk`) from the `C:\Program Files\Veeam\Plugins\Nutanix AHV\Service\PlatformPlugins\AhvProxyManagement\Images` folder of the server where Nutanix AHV Plug-in is installed.

   e. In the file browser, select the required image and click **Open**.

f. In the **Create Image** window, click **Save**.

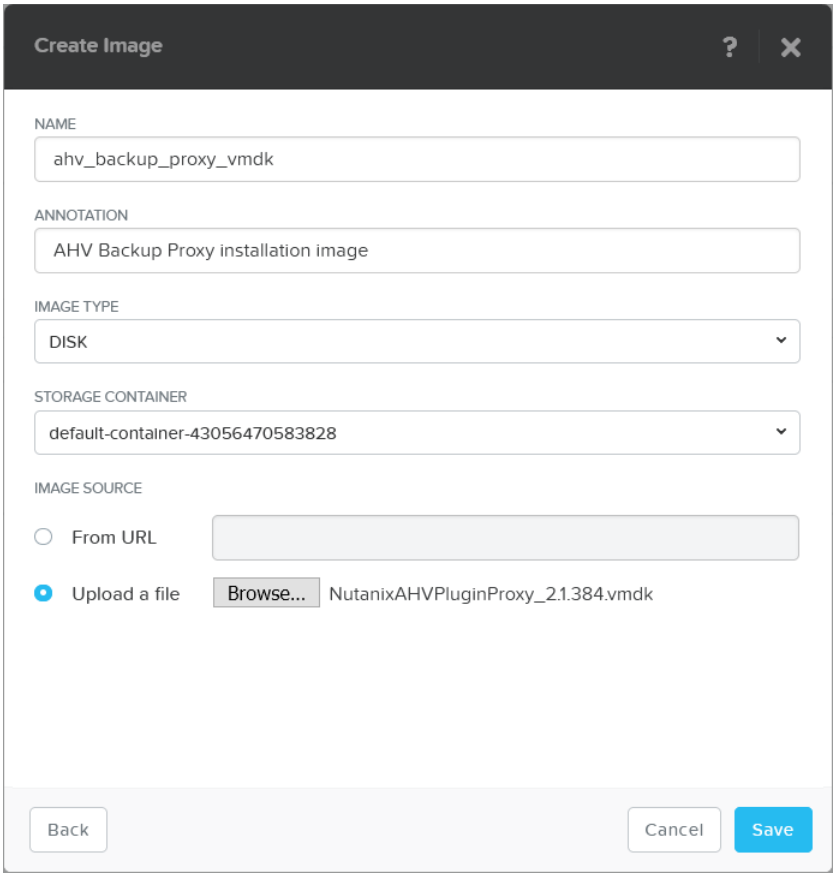

Step 2. Create VM for AHV Backup Proxy

Once you upload the AHV Backup Proxy image, you must create a VM where the AHV Backup Proxy will be deployed.

> **NOTE**
>
> If you use Nutanix Prism Central, the VM management console will be on a different location. For more information, see the Creating a VM section of the Prism Central Guide.

To create a VM for the AHV Backup Proxy, do the following:

1. On the main menu of Nutanix Prism web console, click **Home** and select **VM**.

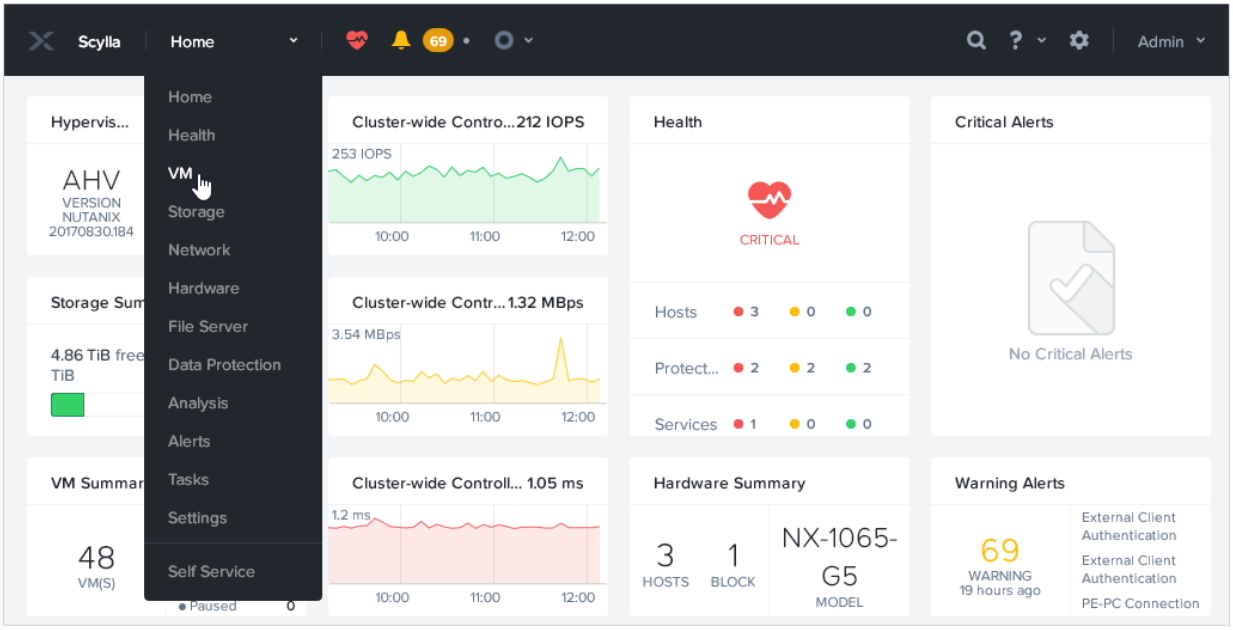

2. Click **Create VM**.

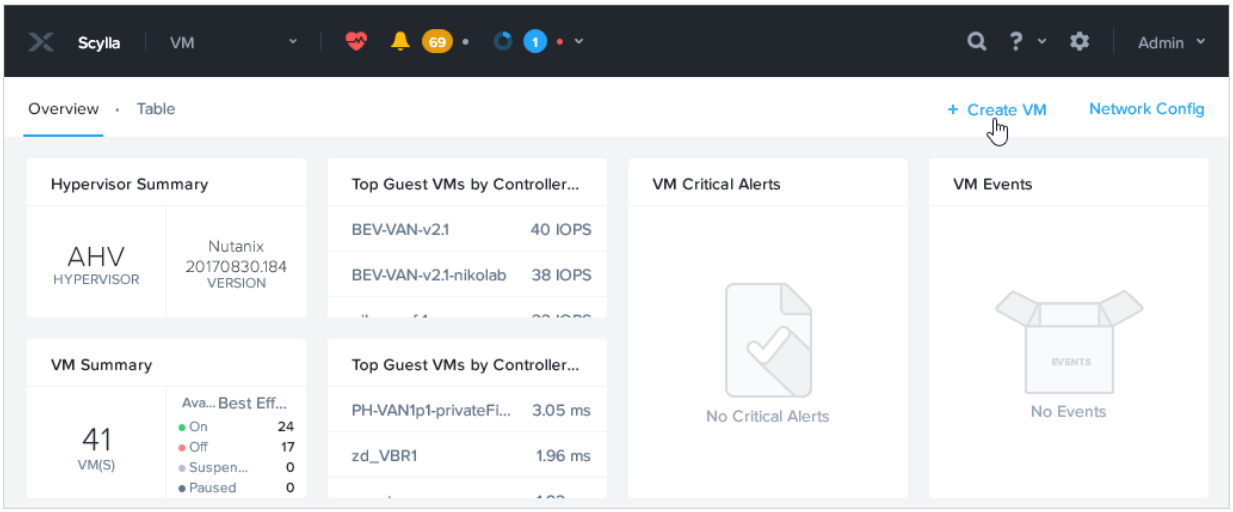

3. In the **Create VM** window, assign a VM name and specify the compute details. Note that the VM must meet the following requirements:

   - 2 CPUs, 2 cores per CPU plus 1 CPU core for each additional concurrent task (default concurrency level: 3 VMs at a time).

- 4 GB RAM plus 1 GB for each concurrent task.

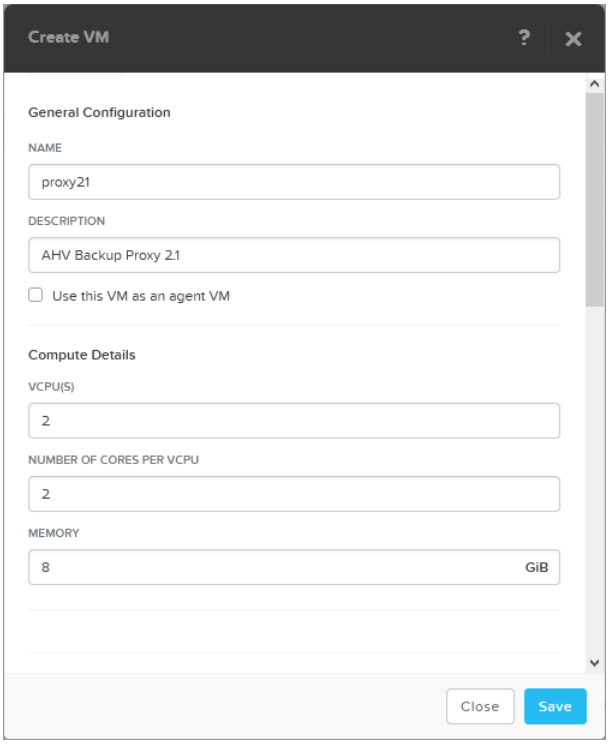

4. At the **Disks** section, click **Add New Disk** to add the AHV Backup Proxy installation disk.

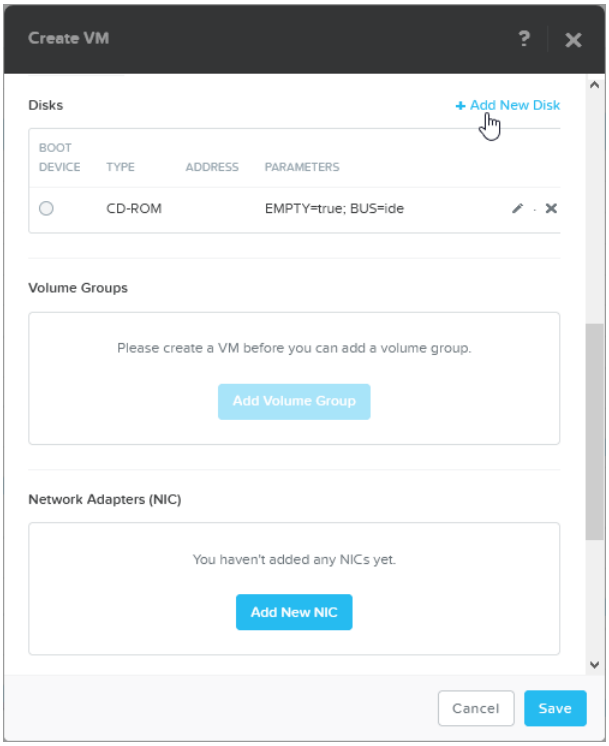

5. In the **Add Disk** window, configure the disk characteristics:
   a. In the **TYPE** drop-down list, select **DISK**.
   b. In the **OPERATION** drop-down list, select **Clone form Image Service**.
   c. In the **IMAGE** drop-down list, select the AHV Backup Proxy VMDK disk which you have uploaded in the previous step.

d. Click **Add**.

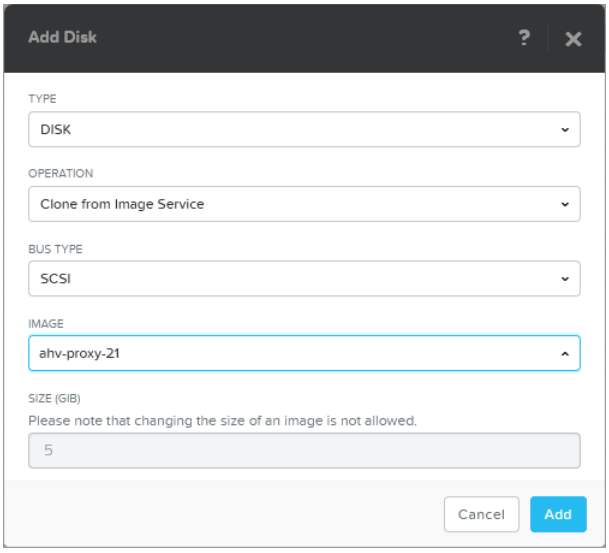

6. At the **Network Adapters (NIC)** section, click **Add New NIC**.

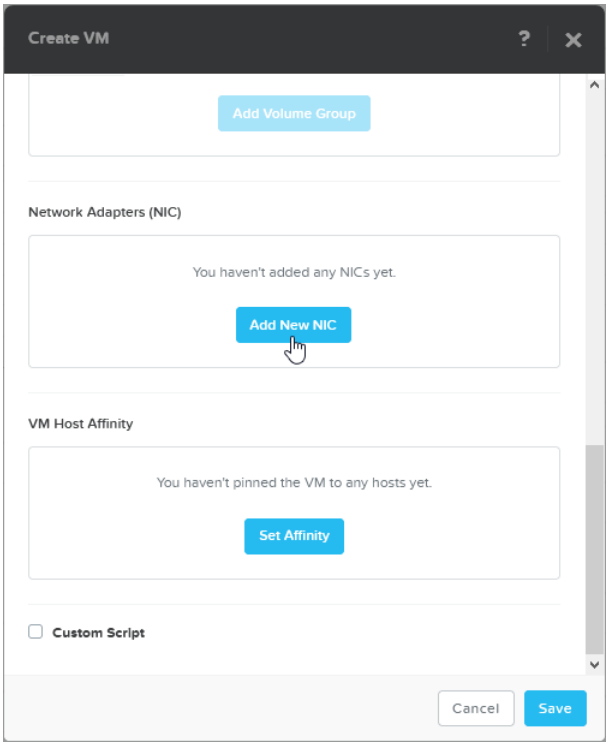

In the **Create NIC** window:

a. Select the required VLAN from the **VLAN Name** drop-down list.

b. Click **Add**.

> **IMPORTANT**
>
> Make sure the specified virtual network has access to your DHCP server, Veeam Backup & Replication Server and Nutanix Cluster.

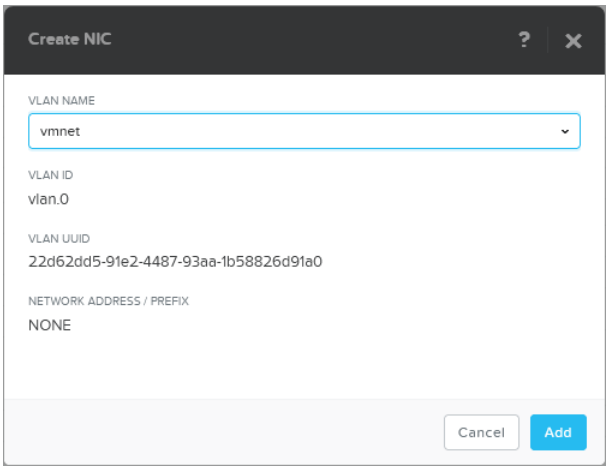

7. AHV Backup Proxy version 2.0 and later has cloud-init integration. For this reason, AHV Backup Proxy VM takes 4-5 minutes to start. To avoid this, do the following.

   a. At the bottom of the **Create VM** window, select the **Custom Script** check box.

   b. Select the **Type or paste script** check box and copy the following script:

```
#cloud-config
growpart:
 mode: auto
 devices: /
 ignore_growroot_disabled: true
write_files:
 - path: /etc/cloud/cloud.cfg.d/90_dpkg.cfg
   content: |
     [ None ]
```

8. Click **Save** to apply settings.

9. After saving the VM settings, you must resize the added image disk to **50GiB**. To resize the disk, perform the following:

   a. In the **VM** view of the Prism web console, click the newly created VM and click **Update**.

   b. Open the AHV Backup Proxy disk settings and at the **Size (GIB)** field type *50*.

c. Click **Update**.

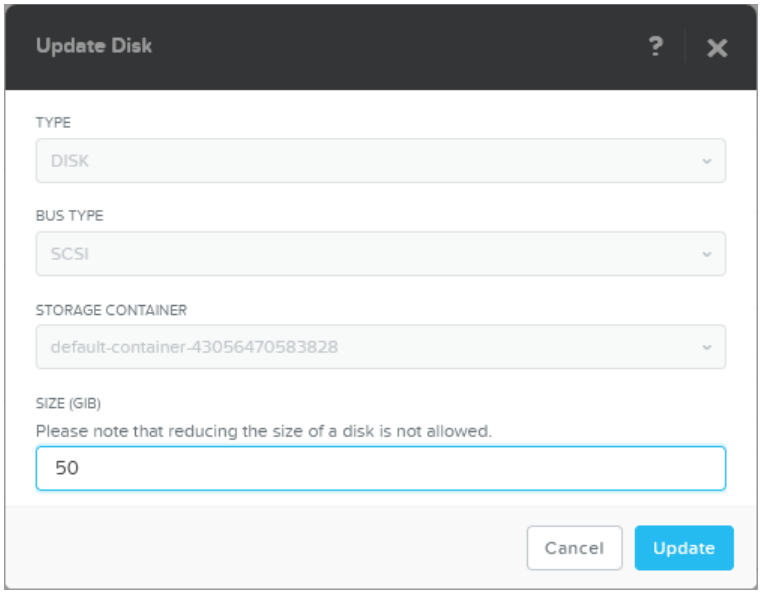

d. In the **Update VM** window, select the added disk as a boot device.

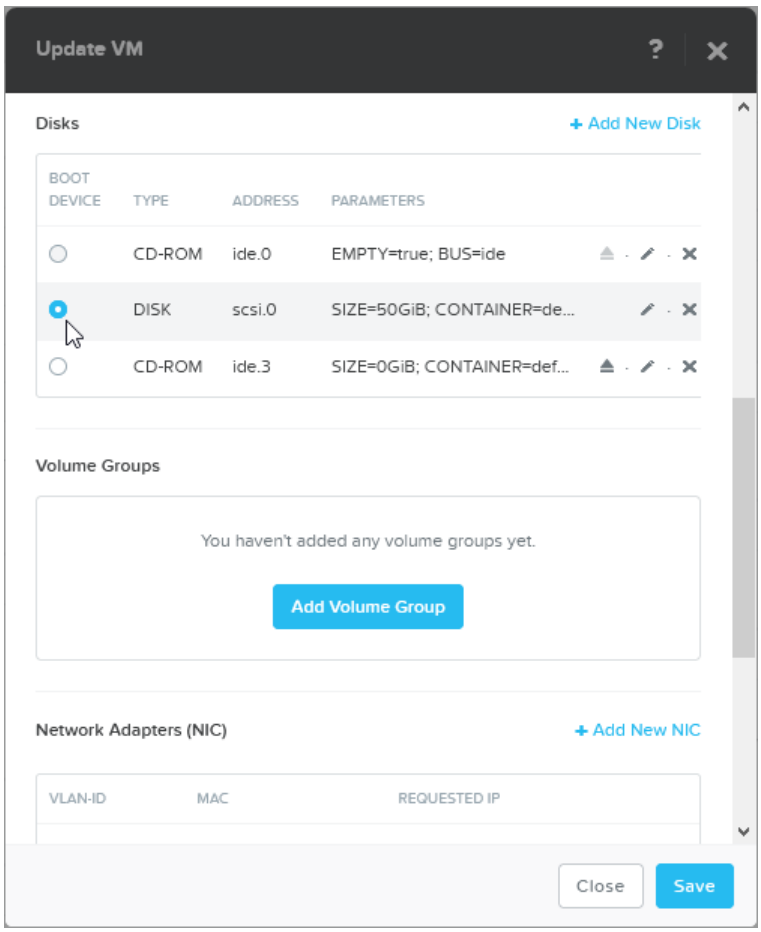

e. Click **Save** to apply settings.

> **IMPORTANT**
>
> You must resize the disk only at the end of configuration. If you resize the disk before saving initial settings of the VM, the disk resize will not be applied correctly.

Step 3. Log in to AHV Backup Proxy Web Console

To log in to the web console of AHV Backup Proxy, do the following.

1. From the list of virtual machines, select the VM for the proxy appliance.
2. On the main menu of Nutanix Prism web console, click **Home** and select **VM**.

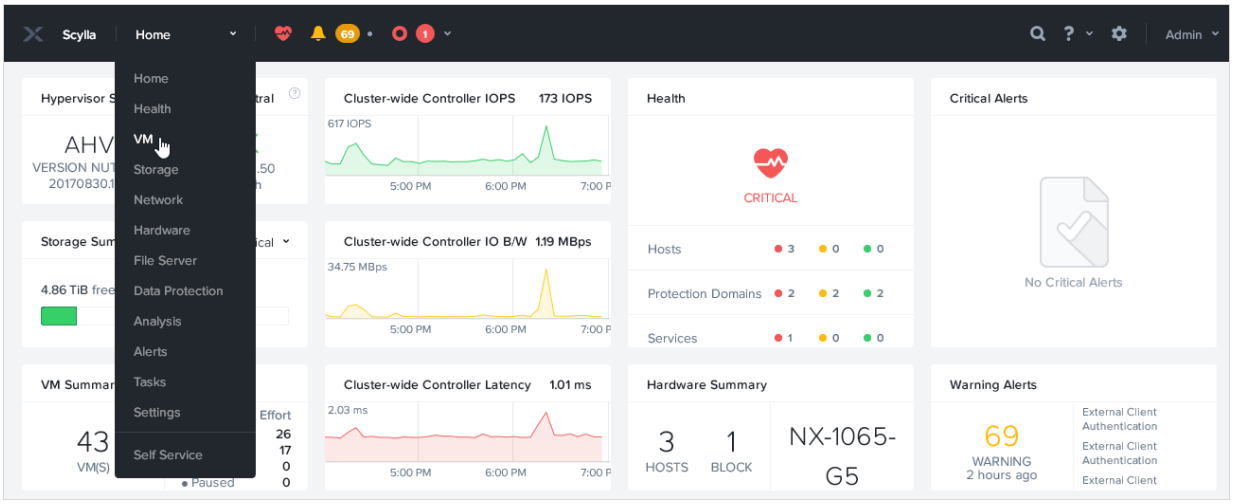

3. Switch to the **Table** view.
4. Click the VM with the image disk of Veeam Backup for Nutanix AHV and click **Power on** on the toolbar.

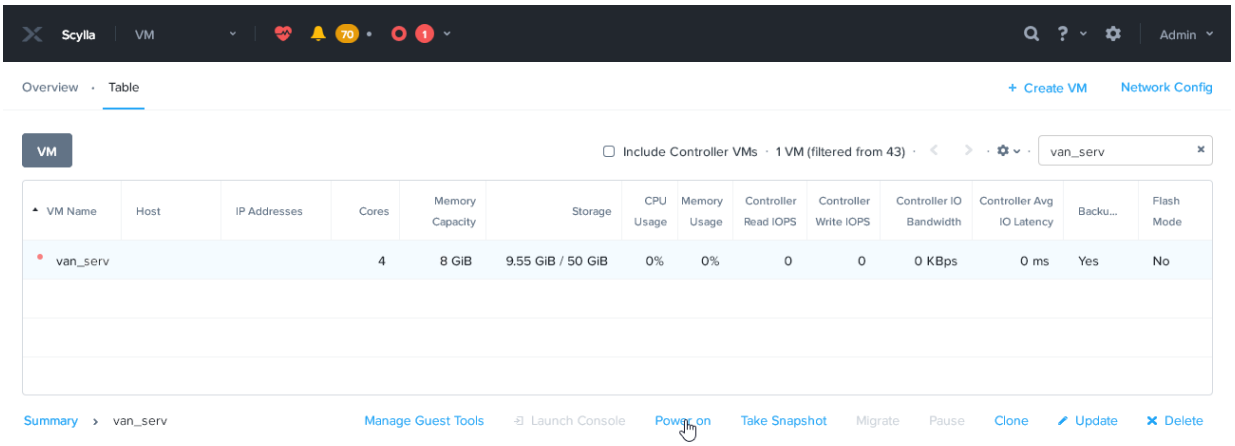

5. After the VM is powered on, click **Launch Console** on the toolbar.

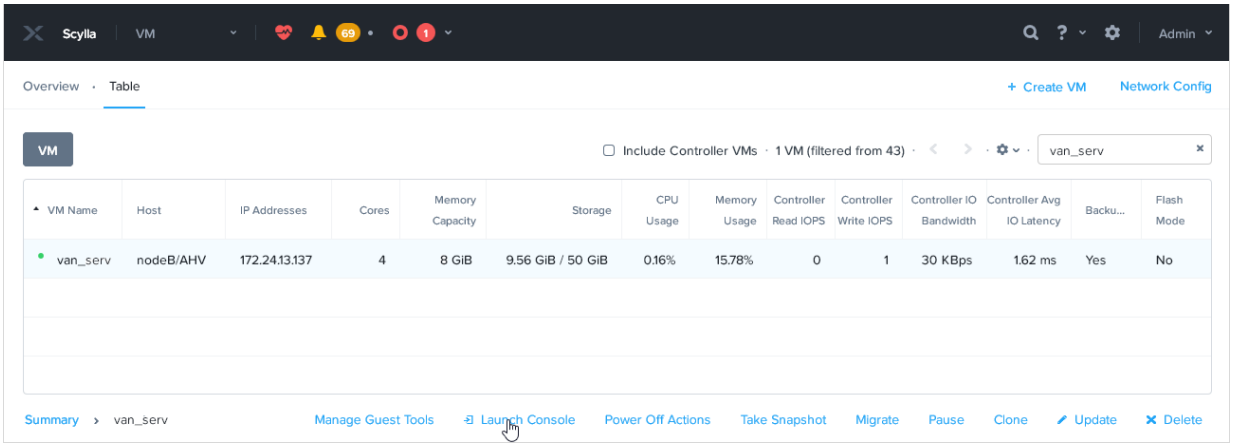

6. On the VM console, you will see the URL address. Use a web browser to go to the displayed address.

   If the URL address of the web console is not displayed, wait 10 seconds and press **ENTER**.

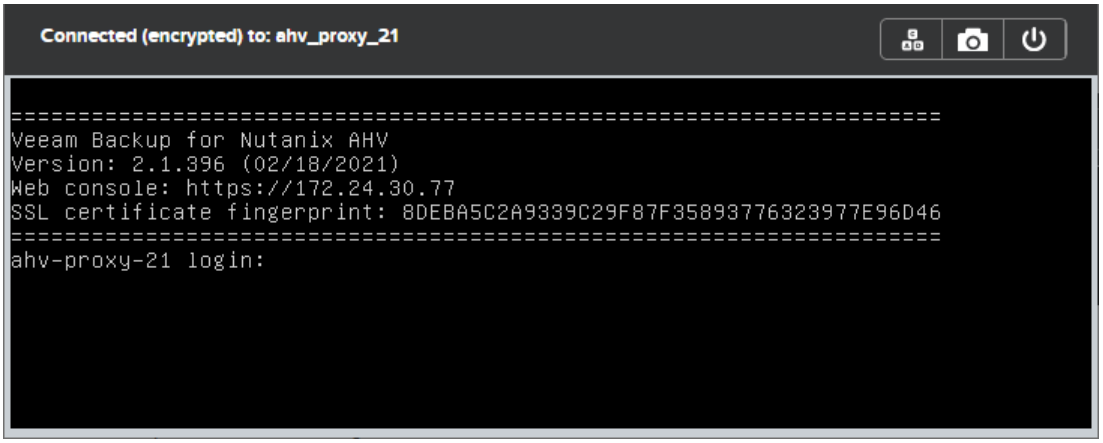

   If your network doesn't have a DHCP server, you must set a static IP address on the VM. To set a static IP address, do the following.

   a. Log in to the VM console (login: *veeam;* password: *veeam)*.

   b. Open the interface settings with a text editor.

```
sudo nano /etc/network/interfaces
```

   c. In the opened file, find the following text.

```
auto eth0
        iface eth0 inet dhcp
```

   Assign a static IP address and specify the netmask and gateway.

```
auto eth0
        iface eth0 inet static
        address <IP_address>
        netmask <netmask>
        gateway <gateway>
```

d. Save the changes and restart the network and the Veeam service.

```
sudo service networking restart
sudo service VeeamBackupAgent restart
```

e. To log in to the web console of AHV Backup Proxy, open https://<static_IP_address>:8100 in a web browser.

g. At the welcome screen, enter default credentials (login: *veeam*; password: *veeam*) to log in to the web console.

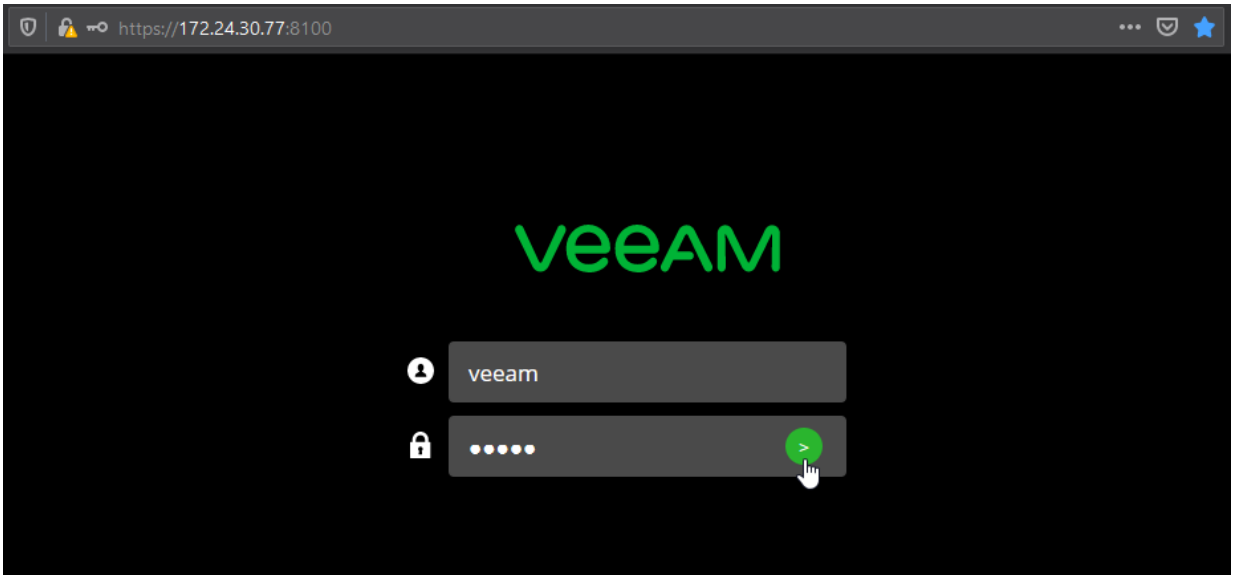

Step 4. Install AHV Backup Proxy

After you log in to the web console of AHV Backup Proxy, the installation wizard will be launched. Go through steps of the wizard:

1. At the **Veeam Backup for Nutanix AHV** wizard, click **Install** to start the installation wizard.

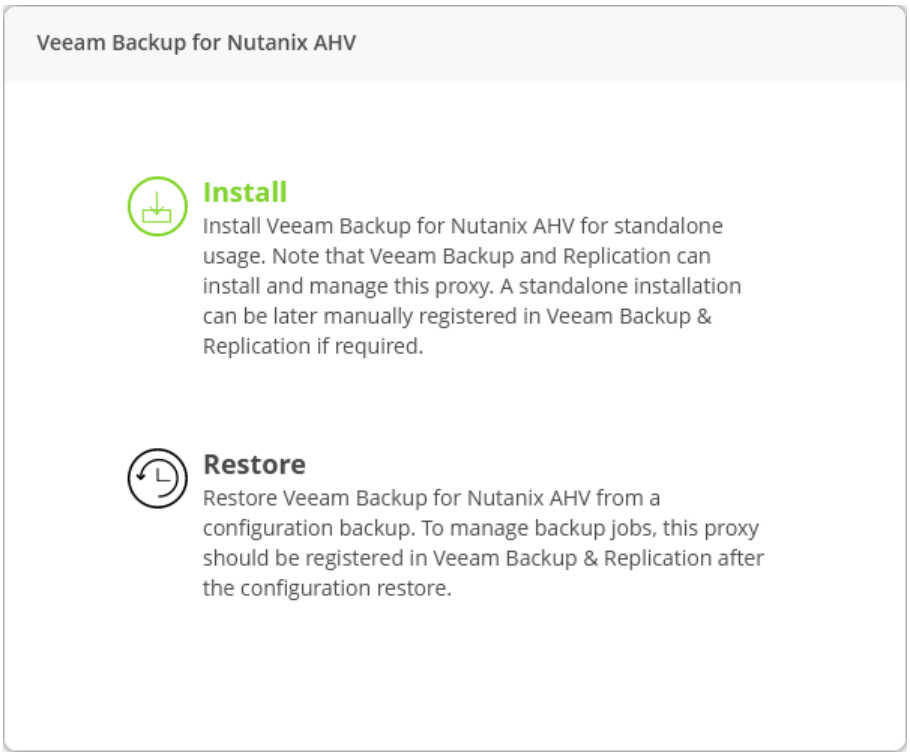

> **TIP**
>
> If you have a configuration backup of AHV Backup Proxy, you can deploy a new AHV Backup Proxy with settings from the backup. For instructions, see Restoring Configuration Settings During Deploy.

2. At the **End User License Agreement** step of the wizard, you must accept the license agreement for Veeam and 3rd party components that Veeam incorporates. If you do not accept the license agreement, you will not be able to pass to next step of the setup wizard.
   a. Read the license agreement.

      To view the license agreement for 3rd party components, click **View**.

   b. Select the **I accept the terms of the Veeam license agreement** check box.

c. Select the **I accept the terms of the third party components license agreements** check box.

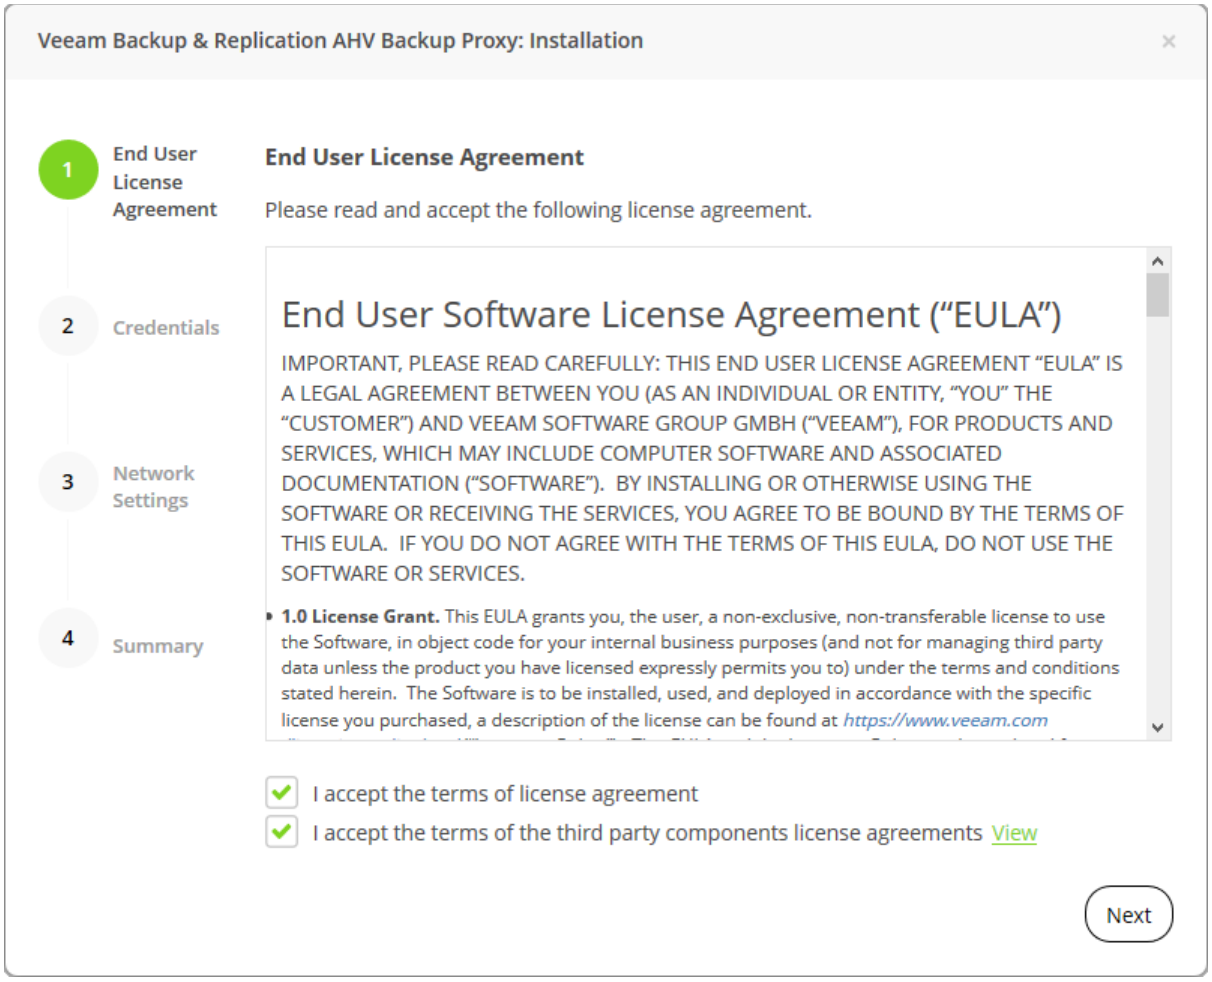

3. At the **Credentials** step of the wizard, change the default password for the admin user.

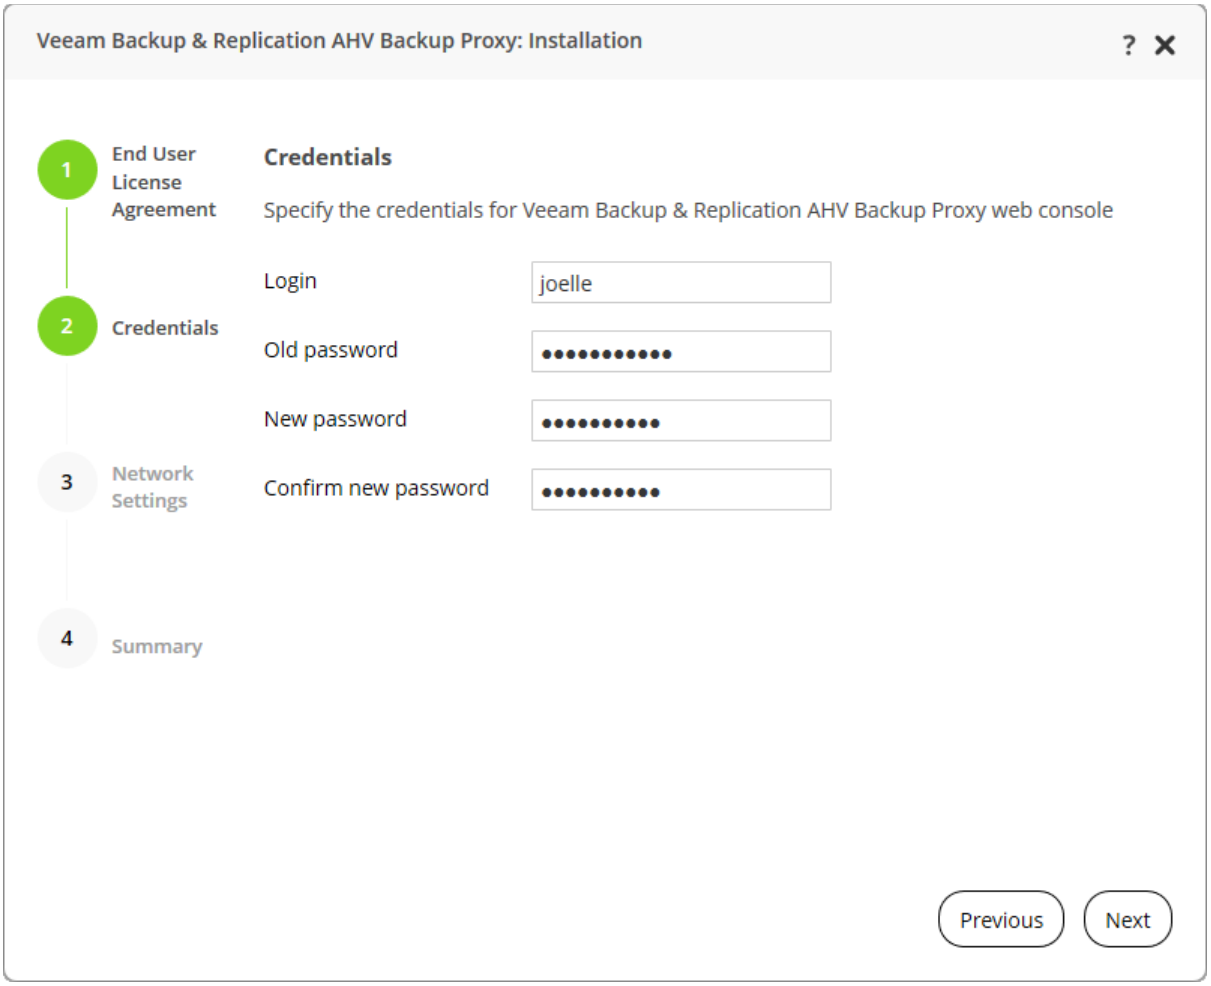

4. At the **Network Settings** step of the wizard, specify the AHV Backup Proxy hostname and network settings or select the **Obtain an IP address automatically** check box.

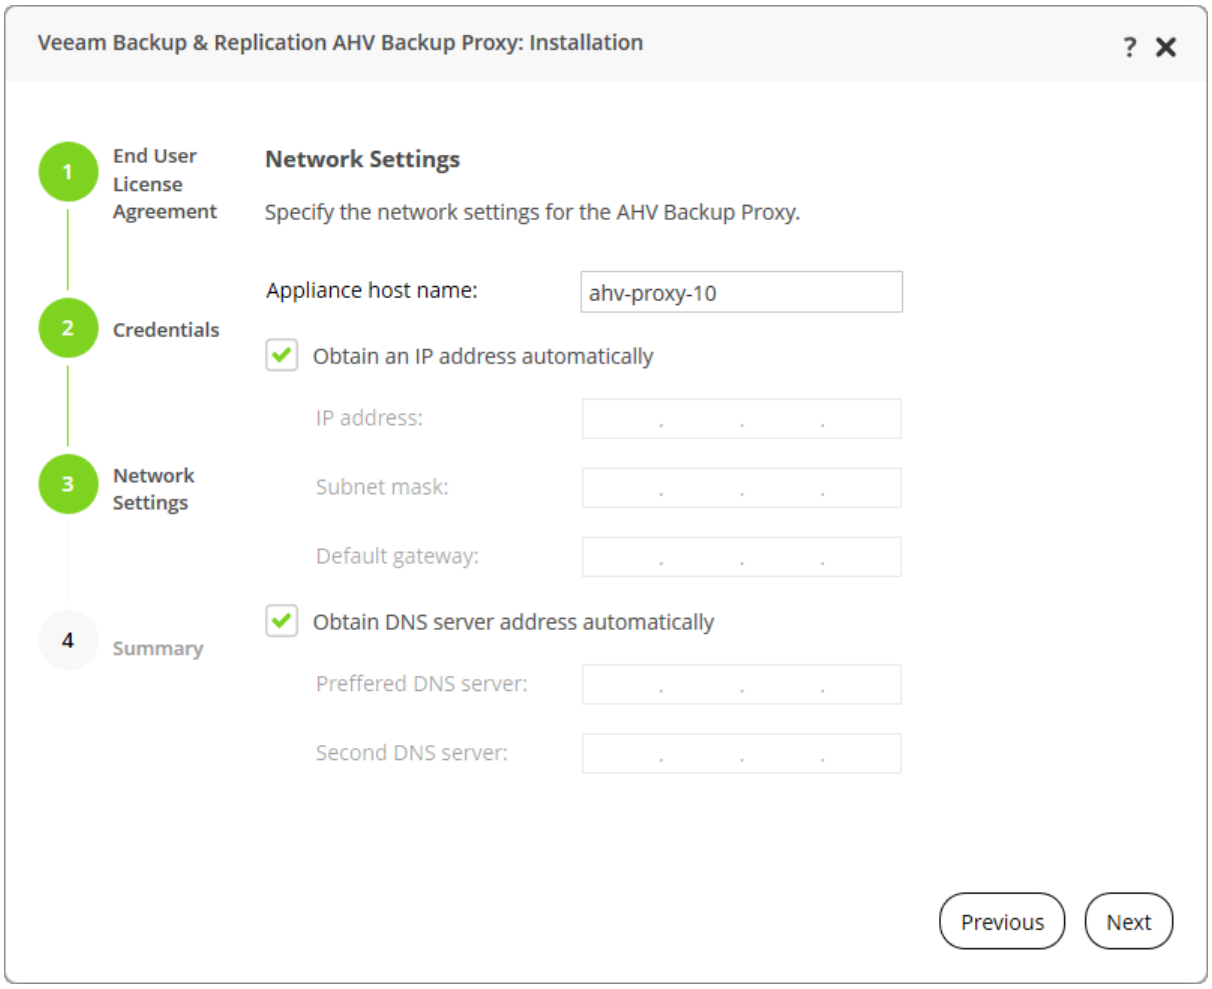

5. At the **Summary** step of the wizard, review the appliance settings and click **Finish**.

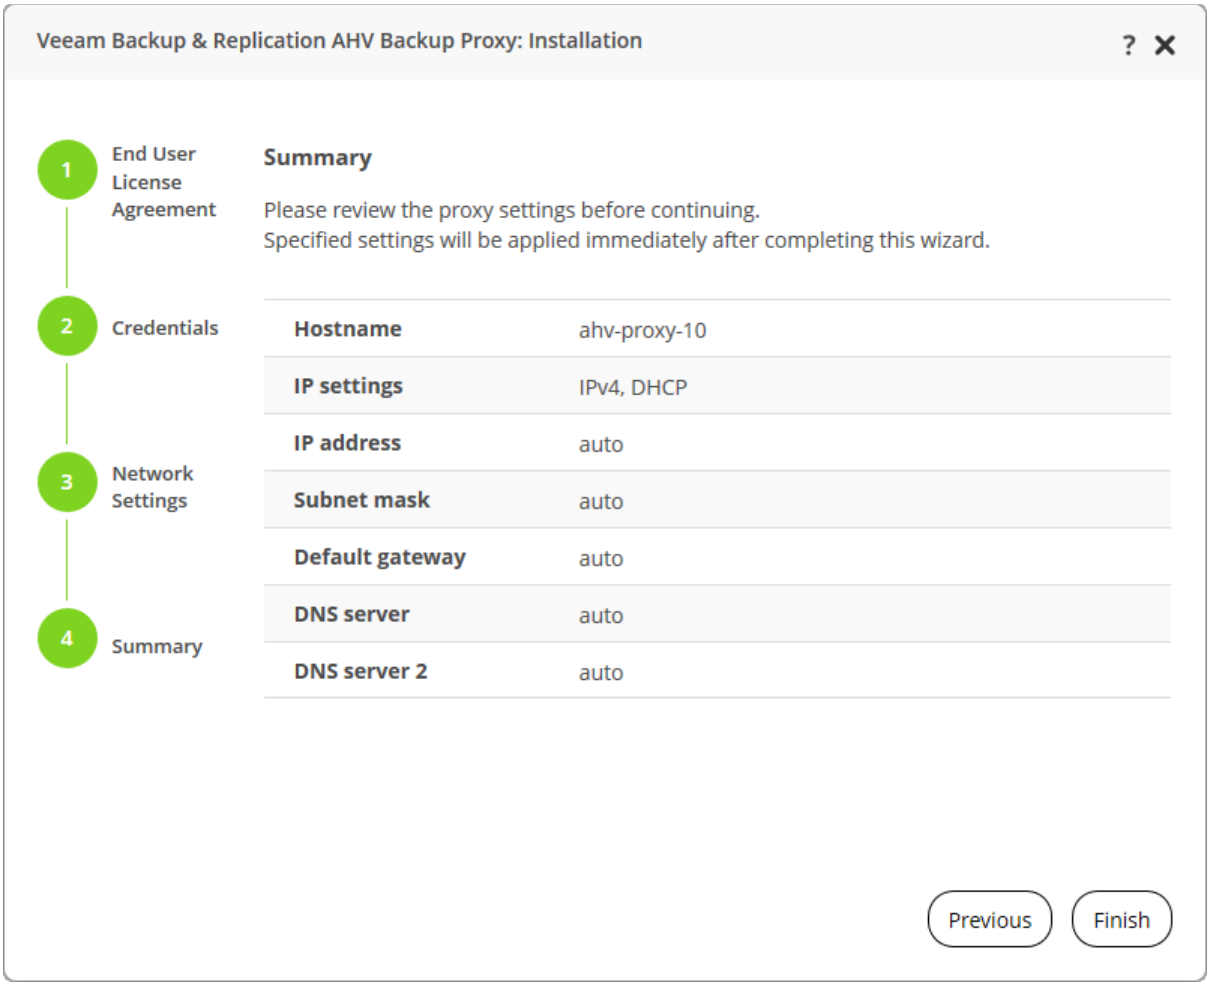

6. After finishing the wizard, you will be redirected to the welcome screen of AHV Backup Proxy. Specify new credentials and log in to the web console of AHV Backup Proxy.

Step 5. Add AHV Cluster to AHV Backup Proxy Infrastructure

After you install AHV Backup Proxy, you must add the Nutanix AHV cluster to the AHV Backup Proxy infrastructure.

Limitations

- You can add only one Nutanix AHV cluster per AHV Backup Proxy.
- You cannot use Prism Central IP addresses to define the AHV cluster in the AHV Backup Proxy settings.
- [For Nutanix Community Edition 2017.07.20] When you add a cluster to the AHV Backup Proxy infrastructure, use the built-in admin account for authentication.

Procedure

To add a Nutanix cluster to the AHV Backup Proxy infrastructure, do the following.

1. Log in to the AHV Backup Proxy web console.
2. At the main panel, click the settings icon and select **Manage AHV Cluster**.

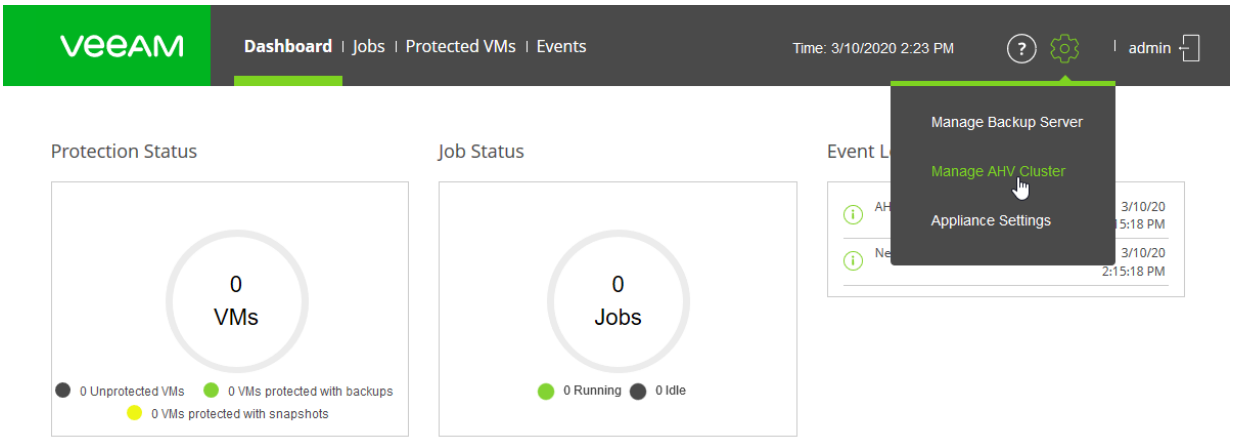

3. Click the **Add button on the toolbar**.

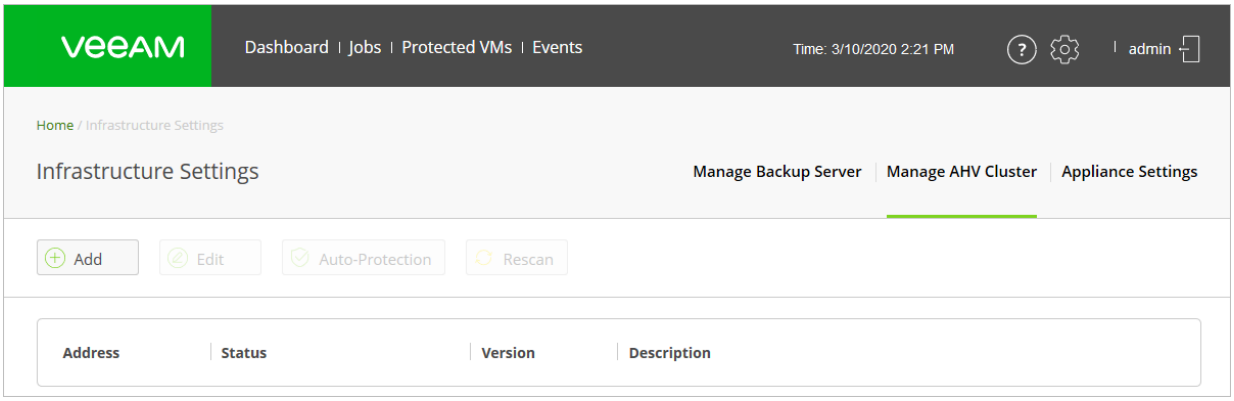

4. Specify the Nutanix AHV cluster settings:
   a. In the **DNS name or IP address** field, specify the IP address or FQDN of the AHV server.

b. In the **Port** field, specify a port which will be used by AHV Backup Proxy to communicate with the AHV cluster (default port: *9440*).

c. [Optional] In the **Description** field, specify a description of the AHV cluster for further reference.

d. In the **Username** and **Password** fields, specify credentials to connect to the AHV cluster.

The account which you want to use to connect to the AHV cluster must have the *Cluster Admin* role permissions.

e. Click **OK**.

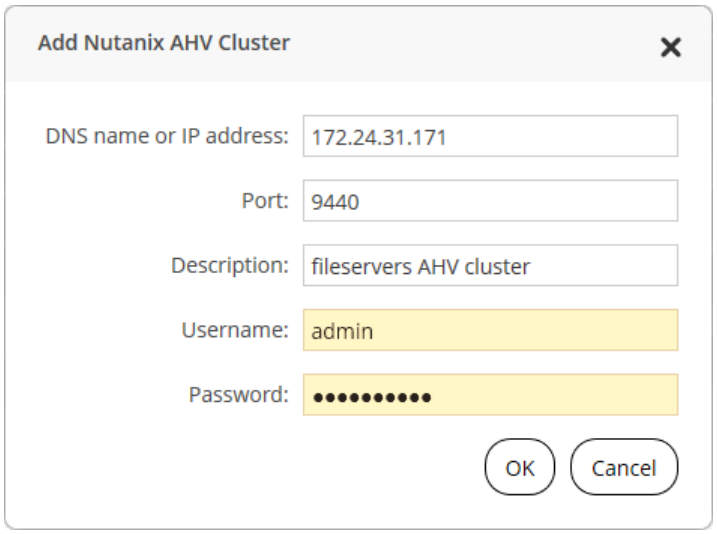

f. AHV Backup Proxy checks the authenticity of the TLS certificate on the AHV cluster. If you trust the legitimacy of the certificate, click **Accept the Risk and Continue** to connect to the AHV cluster.

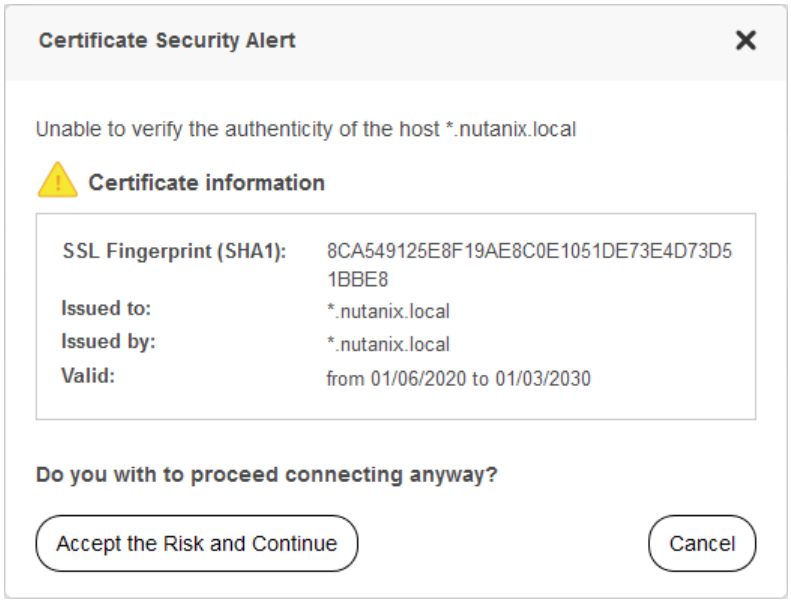

Step 6. Add Veeam Backup & Replication Server

After you add an AHV Cluster, you must configure the connection between AHV Backup Proxy and Veeam Backup & Replication. To do this, you must add the Veeam Backup & Replication server to the AHV Backup Proxy infrastructure:

1. At the main panel of the web console, click the settings icon and select **Manage Backup Server**.
2. Click the **Add button on the toolbar**.
3. In the **Add Veeam Server** window, specify connection settings for the Veeam Backup & Replication server:
   a. In the **DNS Name or IP** field, specify the FQDN or IP address of the Veeam Backup & Replication server.
   b. In the **Port** field, specify a port which will be used by the backup server to communicate with the AHV Backup Proxy server. Default port: 10006.
   c. [Optional] In the **Description** field, specify a description for the Veeam Backup & Replication server.
   d. In the **User Name** and **Password** fields, specify credentials to connect to the Veeam Backup & Replication server.
   e. Click **Add**.

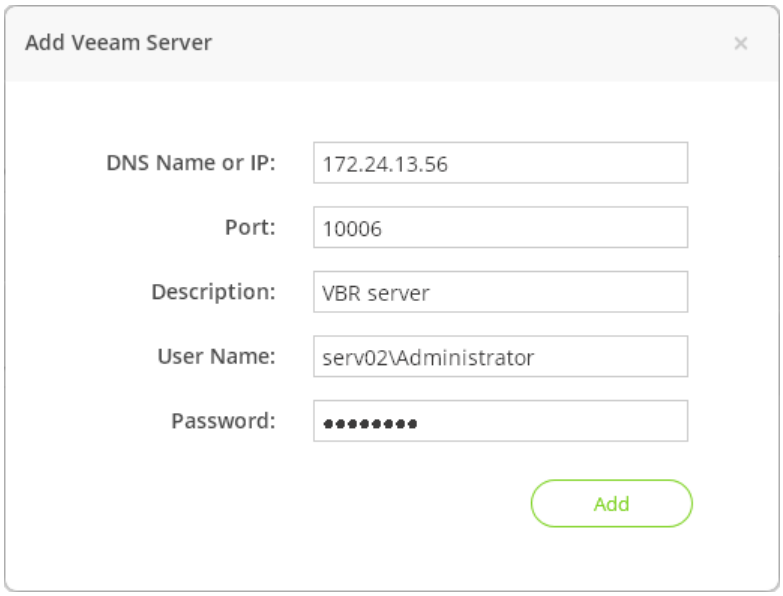

Backup files are stored in the backup repository, in the folder with the same name as the user name used in the Veeam Backup & Replication server settings. If you change the user name for the Veeam Backup & Replication server, on the next run of a backup job a new backup chain will be created. Old backup chain will remain in the folder created for the previous user.

If you want to continue the old backup chain, do the following:

1. In the backup repository server, transfer old backup files to the folder with new backup files.
2. In the Veeam Backup & Replication console, rescan the backup repository. For more information, see the Rescanning Backup Repositories section of the Veeam Backup & Replication User Guide.

NOTE

You can add only one Veeam Backup & Replication server to the AHV Backup Proxy infrastructure.

Step 7. Restoring Configuration Settings During Deploy

To restore configuration settings, do the following.

1. Create a new VM for AHV Backup Proxy. For more information, see Create VM for AHV Backup Proxy.
2. Log in to AHV Backup Proxy web console. For more information, see Log in to AHV Backup Proxy Web Console.
3. In the **Install-Restore** wizard, click **Restore**.

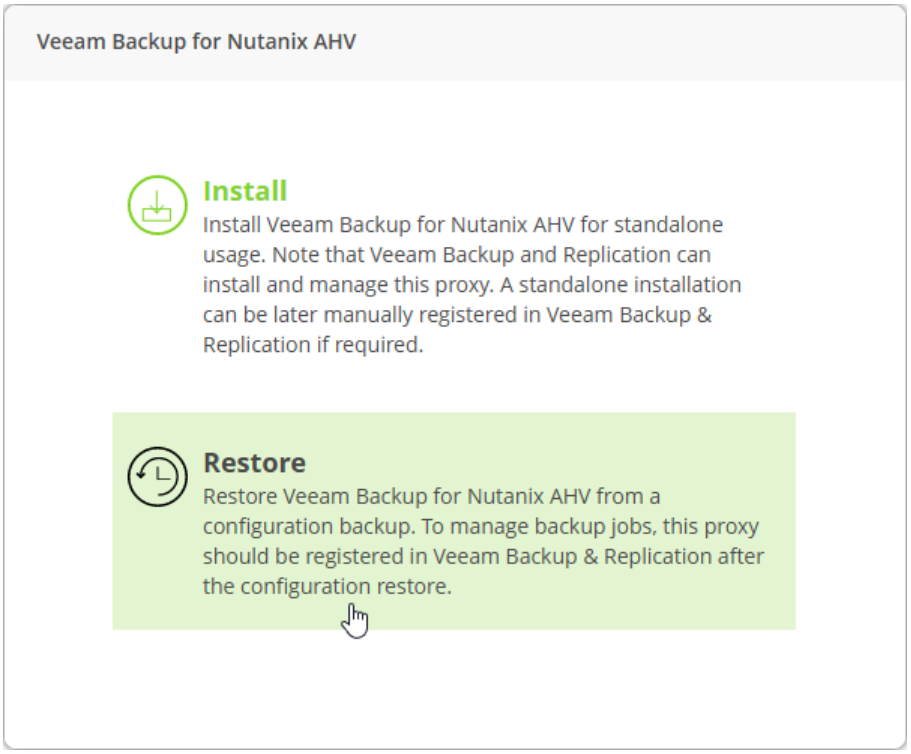

4. At the **End User License Agreement** step of the wizard, accept the terms of license agreement and click **Next**.

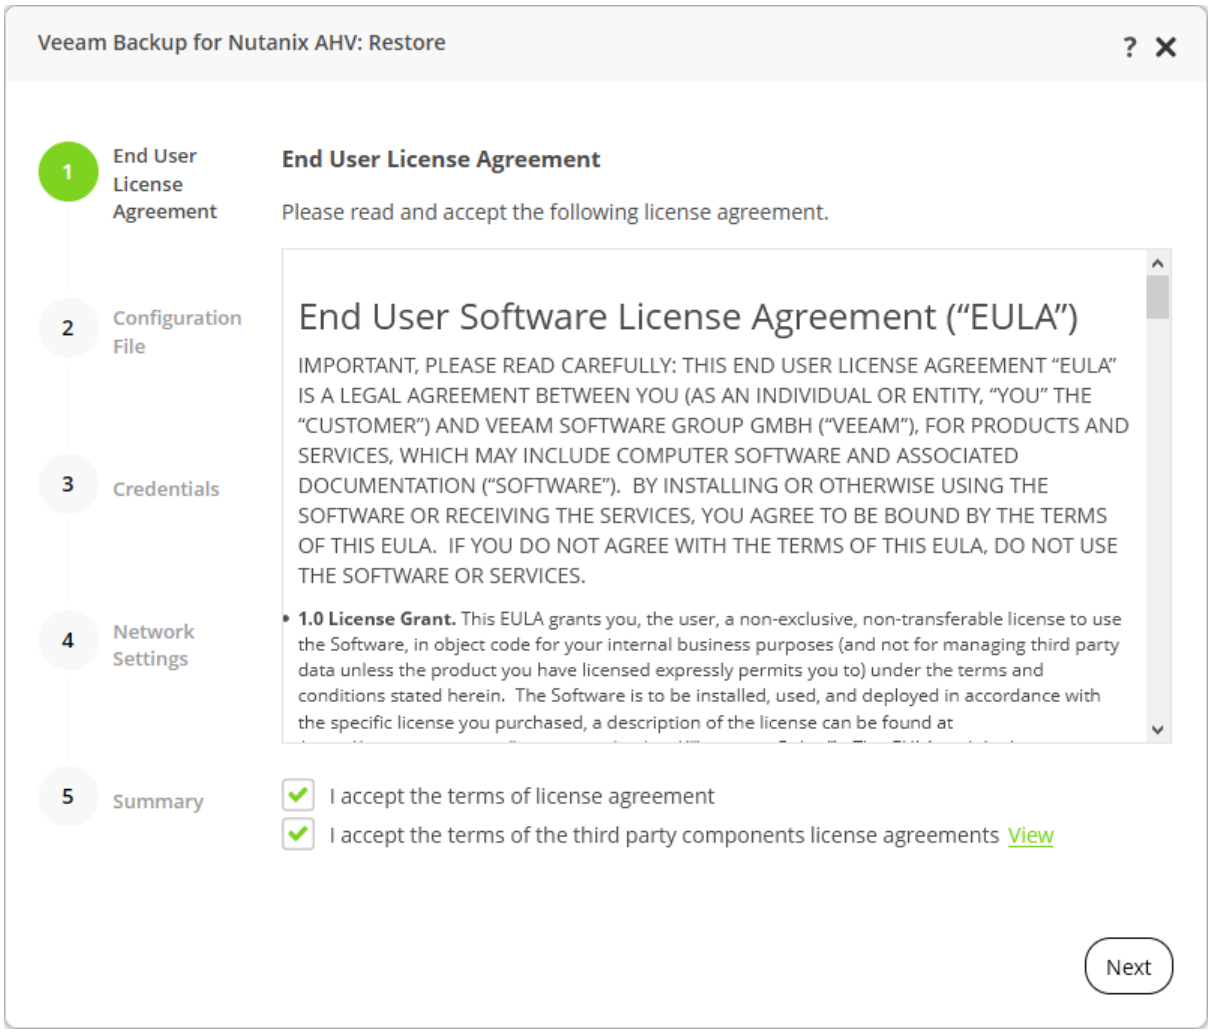

5. At the **Configuration File** step of the wizard, select the configuration file from internal storage. The default naming scheme for the configuration backup file is `config_AhvProxyName-DD_MM_YYYY_HH_MM_SS.zip`.

   Before uploading the configuration backup file, the wizard will request you to enter the password that you have specified when backing up the configuration.

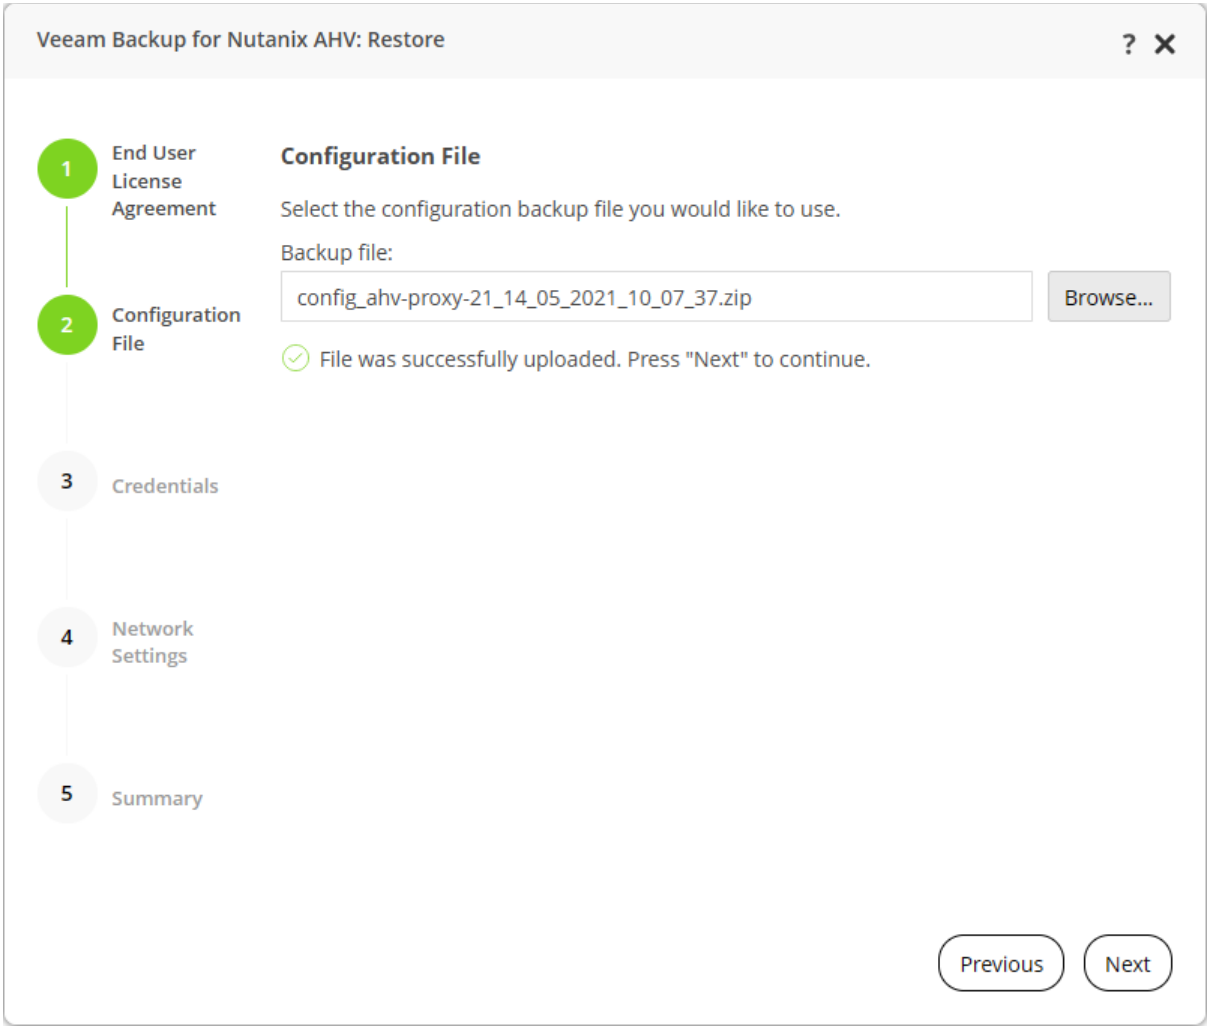

6. At the **Credentials** step of the wizard, specify credentials that will be used to log in to the web console. You can use existing credentials or you can assign new credentials.

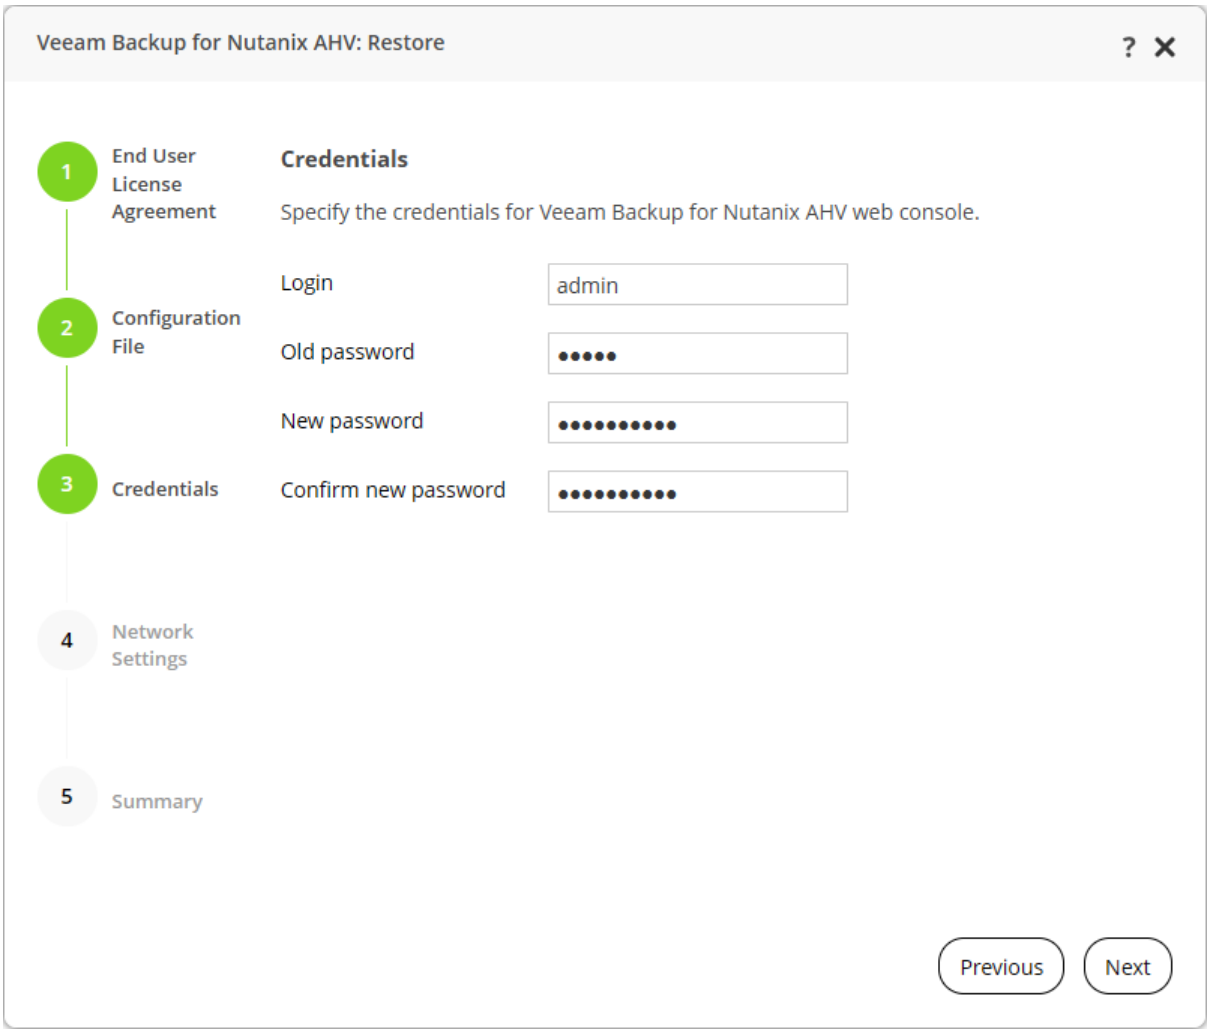

7. At the **Network Settings** step of the wizard, specify hostname and network settings for AHV Backup Proxy. You can use the existing settings from the configuration backup or you can configure new settings.

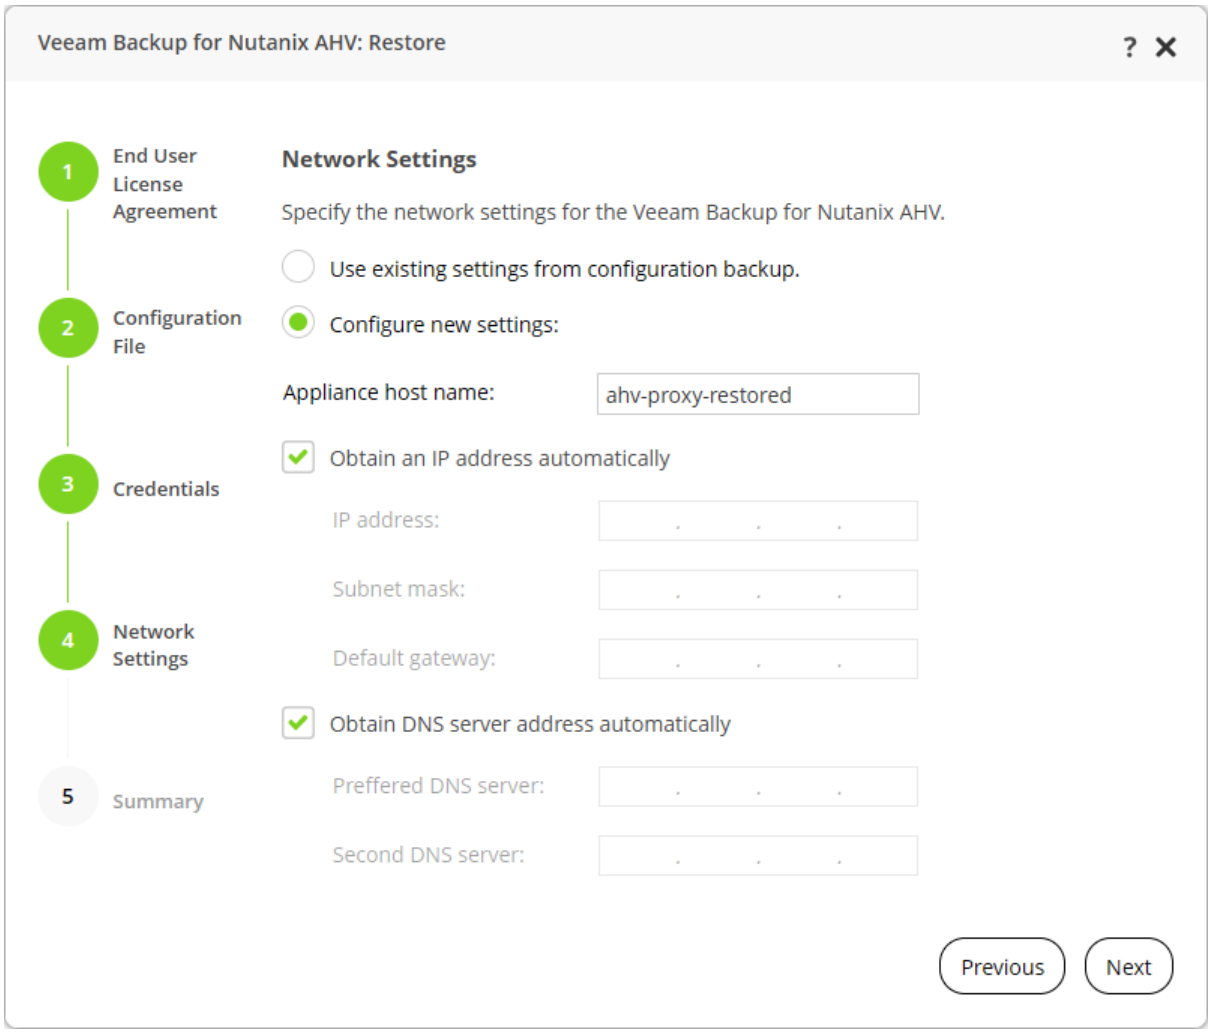

8. At the **Summary** step of the wizard, click **Finish**.

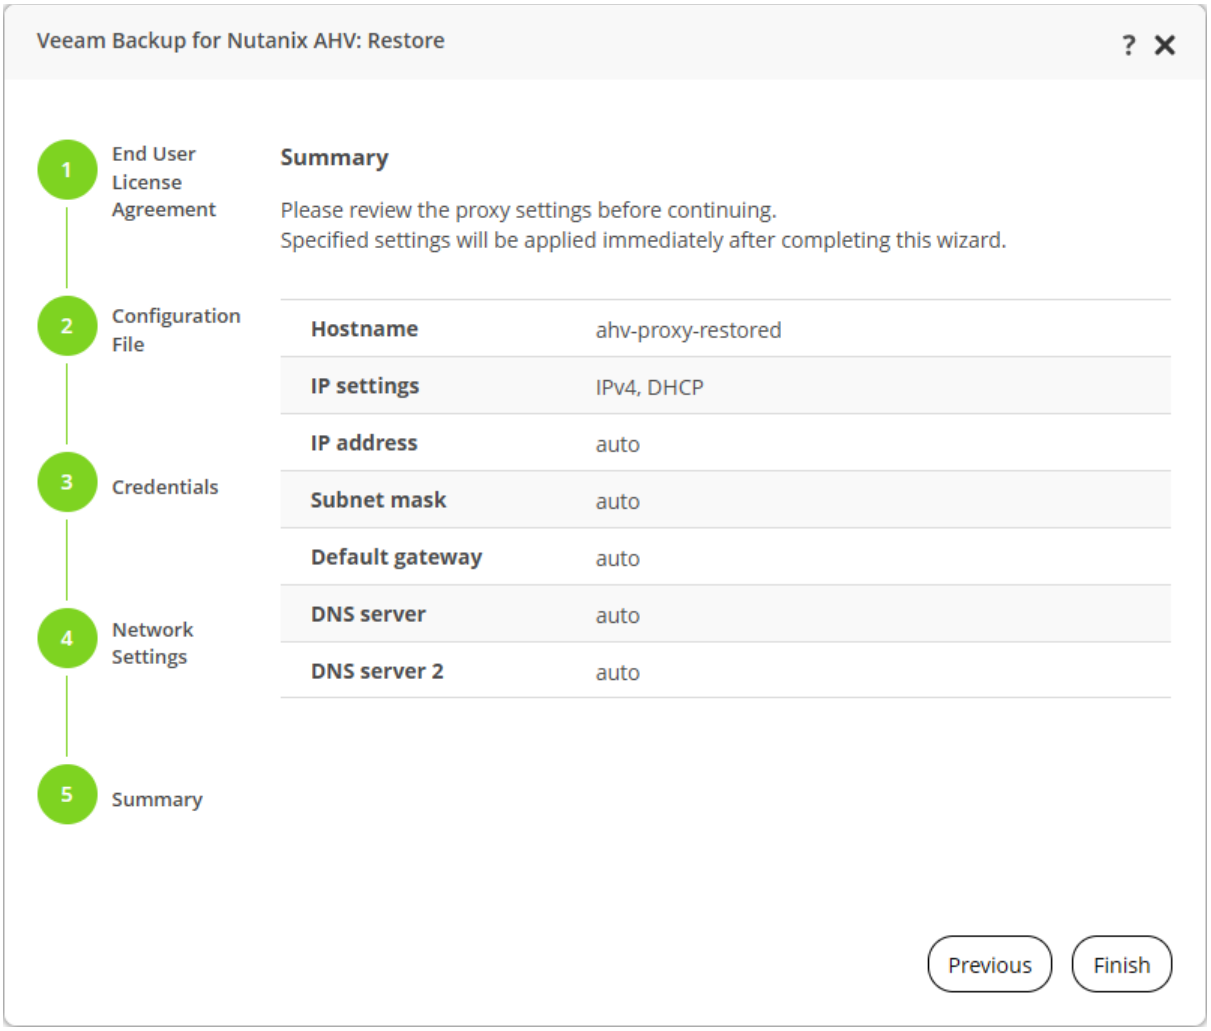

After finishing configuration, the wizard will redirect you to AHV Backup Proxy web console with restored settings.

Related Topic

Backing Up Configuration Settings

Appendix B. Custom Scripts

In AHV Backup Proxy, you can instruct backup jobs and PD snapshot jobs to run custom scripts. You can use the custom scripts to coordinate application activity for backup processes and create protection domain snapshots that will be used as the source of the backup.

AHV Backup Proxy runs the custom scripts at the beginning of the job to which the script is added. The custom scripts are available only for the Dynamic mode of backup jobs and PD snapshot jobs.

Roadmap

To configure Veeam AHV Backup Proxy to use custom scripts to create PD snapshots, go through the following steps:

1. Prepare the custom script.
2. Enable the custom scripts functionality in the AHV Backup Proxy.
3. Upload the script to the AHV Backup Proxy VM.
4. Add the custom script to a backup job.
5. Check Logs of custom scripts and backup jobs.

Step 1. Preparing Custom Script

First of all, you must create a script that will create PD snapshots and/or perform pre-freeze and post-thaw operations.

You can use the sample custom script as an example. To find the sample script, go to the `/opt/VeeamBackupAgent/CustomScripts/BackupJobs` directory in the AHV Backup Proxy VM.

> **IMPORTANT**
>
> The custom script content is not considered to be a part of the Veeam product line. Troubleshooting of custom scripts is not supported as per Veeam Support Policy.

Mind the following:

1. The sample script is written in Python3, but you can use any language. The used interpreter is specified on the shebang line.
2. The sample script contains implemented classes and methods that are used to create a PD snapshot.
3. You can find the correct use examples of classes and methods in the body of the main method.
4. The sample script contains not implemented methods that you can implement yourself:
   - `doSomeChecks(pdVms)`
   - `freezeVms(self, pdVmsList)`
   - `unFreezeVms(self, pdVms)`
5. The sample script does not check validity of the Nutanix AHV cluster certificate.
6. With the custom script, you cannot specify parameters required for connecting to Nutanix AHV cluster. These parameters must be stored separately. For more information, see Defining AHV Cluster Administrator Credentials.
7. If you add the custom script to a backup or PD snapshot job, PD snapshots will be created even for PDs that do not have any Consistency Groups that contain 2 or more entities.

   If you don't add the custom script to a backup or PD snapshot job, PD snapshots will be created only if the PD has at least one Consistency Group that contains 2 or more entities (for example, 2 VMs or one VM and one Volume Group).

Custom scripts support the following arguments:

Argument	Description
--pdName	The name of the protection domain for which you want to create a snapshot.
--jobName	The name of the backup job for which the script is prepared. This argument is used to create backup job logs in the right directory.
--logLevel	A level of logging.

Step 2. Enabling Custom Scripts Functionality

To enable the functionality of custom scripts on AHV Backup Proxy, do the following:

1. Log into to the Nutanix Prism console.
2. Go to the VMs dashboard, where you can find the AHV Backup Proxy VM.
3. Select the AHV Backup Proxy VM and click Launch Console.
4. Open the `appsettings.json` file.

```
sudo nano /opt/VeeamBackupAgent/appsettings.json
```

5. Find the **3rdPartyJobScripts** parameter and change the value of the **Enable** argument to *true*.

```
"3rdPartyJobScripts": {
   "Enable": true,
   "ExecutionTimeLimitSec": 1200,
   "FailScriptAfterTimeout": true,
   "FailScriptExecutionAfterPolicyCancel": true
 }
```

6. Configure other parameters if required. For the list of available parameters, see the table below.

Argument	Description
Enable	Possible values: • **true**: backup job uses the custom script specified in the backup job settings. • **false**: backup job does not use custom scripts.
ExecutionTimeLimitSec	Possible values: a positive integer (0, MAX_INT)
FailScriptAfterTimeout	Possible values: • **true**: if the script execution exceeds the timeout period, the script execution is interrupted with a timeout error. • **false**: if the script execution exceeds the timeout period, the backup job waits for the script execution indefinitely. When the script execution is finished, the job ends with a warning.

Argument	Description
FailScriptExecutionAfterPolicyCancel	Possible values: • **true**: if the backup job session is cancelled while the script is running, execution of the script will be interrupted. The backup job will fail with an error. • **false**: if the backup job session is cancelled while the script is running, the script execution will be continued in the background. The backup job will fail with an error.

7. Restart the **VeeamBackupAgent** service.

```
sudo service VeeamBackupAgent restart;
```

Defining AHV Cluster Administrator Credentials

To use the script, you must specify credentials that will be used to connect to the AHV cluster and its IP address.

1. Log into the AHV Backup Proxy VM console.
2. Open the `nutanix.conf` file.

```
sudo nano /opt/VeeamBackupAgent/CustomScripts/nutanix.conf
```

3. Specify the cluster IP address and credentials.

```
nutanixClusterIp = "172.22.22.01"
nutanixLogin = "joelle.van.dyne"
nutanixPass = "Password"
```

NOTE

If you perform backup of the AHV Backup Proxy configuration settings, credentials specified in the custom script will be stored in the plain text format. To avoid this, you can relocate the `nutanix.conf file` and change its default path in the script.

Step 3. Uploading Custom Script

To upload a custom script and set required permissions, do the following.

1. Upload your scripts to the AHV Backup Proxy VM. To upload scripts, you can use WinSCP or any other tool. The default directory for scripts is: `/opt/VeeamBackupAgent/CustomScripts/BackupJobs/`.

   Note that SSH is disabled by default in the AHV Backup Proxy. To enable SSH, run the following command in the VM console.

   ```
   sudo service ssh start
   ```

2. Set the file execution permissions.

   ```
   sudo chmod +x /opt/VeeamBackupAgent/CustomScripts/BackupJobs/custom_script
   .sh
   ```

Step 4. Adding Custom Script to Job

In the AHV Backup Proxy web console, you can add a custom script to a backup job or a PD snapshot job. The example in this section shows how to add the script to a backup job. Same instructions can be used for PD snapshot job.

To add custom scripts to a backup job, do the following:

1. Log into the AHV Backup Proxy web console. For instructions, see Accessing Web Console.
2. Open the **New Job** wizard to create a new backup job.
3. At the **Virtual Machines** step of the wizard, click **Add** and select the **Dynamic mode** check box.

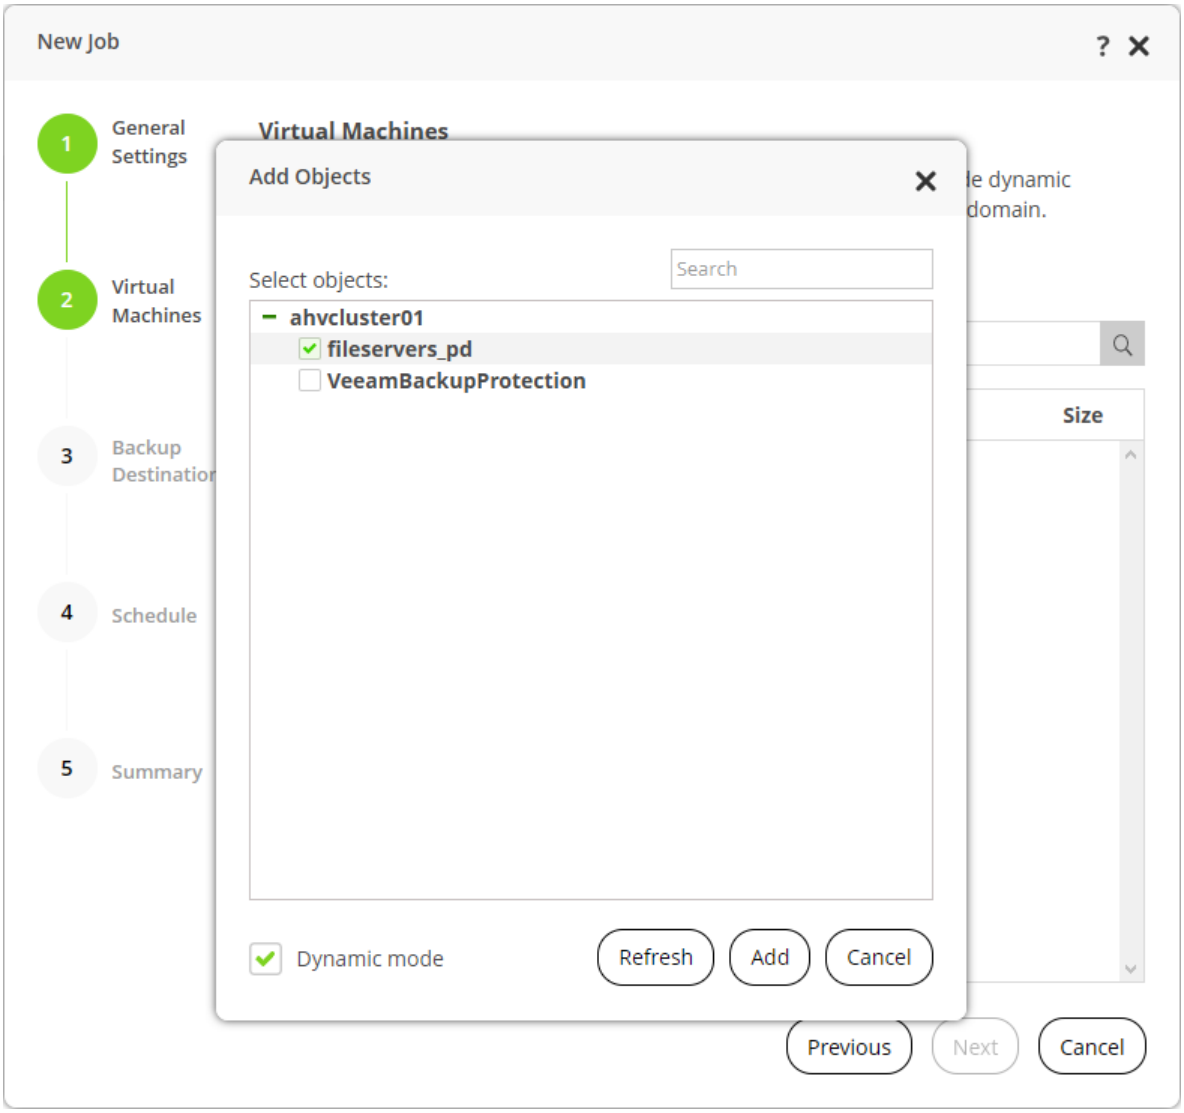

4. At the **Backup Destination** step of the wizard, click **Advanced** and select the script:
   a. At the **Advanced Settings** window, select the **Script** tab and select the **Run the following script to create PD snapshot for this job** check box.

b. In the field specify the path where the script is located. If the script is in the default directory, leave this field unchanged.

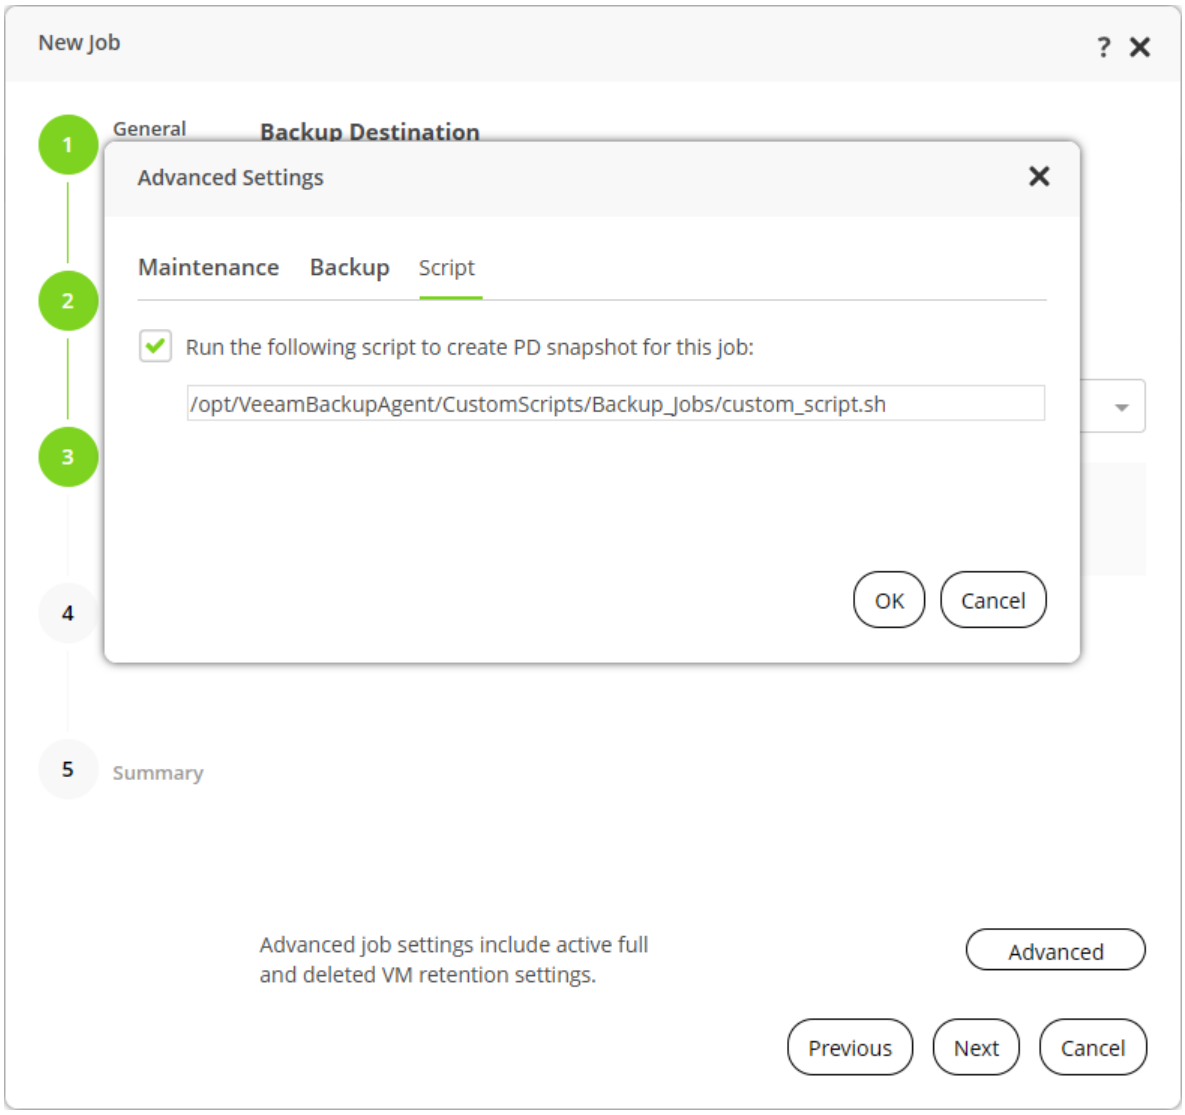

5. Go through other steps of the wizard to finish configuring the job and close the **New Job** wizard.
6. Run the created backup job. Within this job, PD snapshots will be created by the specified scripts.

Step 5. Checking Custom Script Logs

Logs that contain custom script execution results are stored in the `/var/log/nxbackupagent/policy_name/CustomScript` directory in the AHV Backup Proxy VM.

Logs of backup jobs are stored in the following directory: `/var/log/nxbackupagent/policy_name/`.

Also, you can see the logs in the backup job session in the AHV Backup Proxy console.

Defining Error Messages in Custom Script

In a custom script, you can define messages that must be placed in logs upon successful or failed execution of the script.

- You can use the following template to define a message if the custom script execution was successful.

```
{
"oob_schedule_id": 44509,
"errorMessage": "",
"status": "Success"
}
```

  Where 44509 is the identifier of created PD snapshot.

- You can use the following template to define a message if the script execution failed.

```
{
"oob_schedule_id": "",
"errorMessage": "Unable to get vm_id list for Protection Domain='VeeamBack
upProtection': Authentication failed.",
"status": "Failed"
}
```

  In this case the backup job will fail with an error.

- You can use the following template to define a message if the custom script execution finished with a warning.

```
{
"oob_schedule_id": 44509,
"errorMessage": "Pre-freeze took more time than expected",
"status": "Warning"
}
```

  Where 44509 is the identifier of created PD snapshot.

Veeam Backup & Replication

Veeam ONE

Feature Comparison

Contents

LICENSING DETAILS **3**

LICENSING CONSIDERATIONS **4**

VEEAM UNIVERSAL LICENSE (VUL) ALONG WITH A SOCKET LICENSE **4**

LICENSES FOR VEEAM CLOUD AND MANAGED SERVICE PROVIDERS (VCSP) **4**

VEEAM BACKUP & REPLICATION™ **5**

BACKUP 5

ENTERPRISE APP BACKUP PLUGINS 7

STORING BACKUPS 7

REPLICATION 9

RESTORE 9

INSTANT RECOVERY 9

FILE & ITEM LEVEL RECOVERY 10

DATALABS 11

SELF-SERVICE 12

MANAGEMENT 13

VEEAM ONE™ **14**

MONITORING AND REPORTING 14

MONITORING AND ALERTING FOR VIRTUAL INFRASTRUCTURE 15

REPORT CUSTOMIZATION AND AUTOMATION 16

GENERAL 17

INTELLIGENT AUTOMATION 17

CAPACITY PLANNING 18

CHARGEBACK AND BILLING 18

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam® Backup & Replication™ delivers powerful data protection across ALL cloud, virtual and physical workloads in a simple, flexible and reliable solution. Protect your data no matter the location or workload type. The fully featured, portable Veeam Universal License (VUL) license is the best option for any environment, whether it be on-premises, or in hybrid- or multi-cloud deployments.

Licensing details

Veeam has previously offered socket-based licenses and have now pivoted to the feature-complete and portable VUL for full functionality across workloads and environments. This document will outline the features that are available by license type.

VUL is required for many features, so it is always recommended that you migrate your license to VUL for full capability. A Veeam Universal License, VUL, is always required for NAS, all physical agents, Nutanix AHV VMs, cloud-native workloads and much more.

Workloads	Veeam Universal License	Per-socket
Environments	On premises, cloud and hybrid	On premises only
Virtual Machines: VMware vSphere, Microsoft Hyper-V & Nutanix AHV	✔	VMware & Hyper-V
ENHANCED Physical Servers, Workstations and Endpoints: Windows, Linux, Mac, Oracle Solaris and IBM AIX	✔	–
ENHANCED Enterprise Apps for Microsoft, Oracle, SAP HANA and more	✔	–
ENHANCED Cloud-native backup and recovery for AWS, Microsoft Azure and Google Cloud Platform	✔	–
ENHANCED NAS File Shares	✔	–
More to come!	✔	–

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Licensing considerations

Veeam Universal License (VUL) along with a Socket License

If you still use a socket license, we recommend you talk to your sales representative about migrating it to VUL for the ultimate capability and flexibility. However, if you must continue using sockets, please be aware of the following limitations:

- When combining VUL with a legacy socket-based license in a merged and centrally managed environment, **the socket edition and support level always takes precedence and dictates the feature edition level of the entire environment**.
  - Example: Combing a Standard legacy socket license with VUL will restrict certain platform capabilities like the Scale-out Backup Repository.
- If a feature isn't available with a socket license, like cloud machine protection for example, it will be available when a VUL is merged with socket licenses.
- In the presence of a socket license, **all VMware vSphere and Microsoft Hyper-V VMs must consume a socket license**. You **cannot** use VUL to protect some of your VMware and Hyper-V VMs and Sockets to protect others. Nutanix AHV and Red Hat Virtualization (RHV) requires the VUL license. Migrating socket licenses to VUL gives the freedom to protect all workloads regardless of hypervisor, OS, workload or environment.
- VMs on hypervisor hosts that are licensed with a socket license can now be protected with any Veeam product without needing an additional VUL license, so long as the protection is managed by the same backup server. For example, protecting a Windows failover cluster based on vSphere VMs with Veeam Agent *for Microsoft Windows* does not additionally consume a VUL license for the agent.

Licenses for Veeam cloud and managed service providers (VCSP)

VCSP Rental license functions in a similar fashion to the Subscription license, but uses a point-based system with point consumption depending on capability levels. For more information, please refer to your VCSP licensing documentation or speak with your Service Provider.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam Backup & Replication™

✔ Fully supported ✔ Partially supported

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
BACKUP						
Support for VMware vSphere and Microsoft Hyper-V hypervisors	✔	✔	✔	✔	✔	Complete data protection with full spectrum recovery for the market- leading hypervisors. Protect VMs, data, applications and more with the ability to instantly recover complete VMs as well as perform granular recovery of key application items.
Support for Nutanix AHV hypervisor	✔				✔	Complete data protection with full spectrum recovery. Protect VMs, data, applications and more with the ability to instantly recover complete VMs as well as perform granular recovery of key application items.
Support for Red Hat Virtualization hypervisor	✔				✔	Data protection and granular recovery for Red Hat Virtualization hypervisor VMs, data and applications.
Application-aware, image-based backups	✔	✔	✔	✔	✔	Create application-consistent, VMware, Hyper-V, and Nutanix AHV image-level VM backups with advanced, application-aware processing (transaction log backup for Oracle and SQL Server requires Enterprise or higher edition).
AWS-native backup and recovery					✔	Create AWS-native snapshots of Amazon EC2, Amazon RDS and Amazon EFS, as well as Amazon VPC backup. Automatically tier backups to Amazon S3 , S3 Glacier and S3 Glacier Deep Archive. Restore options include full- and file-level recovery in AWS, recovery on-premises, Instant VM Recovery® and recovery to Microsoft Azure and Google Cloud Platform.
Microsoft Azure-native backup and recovery					✔	Create Azure-native snapshots of Azure VMs, as well as Azure SQL Database backup. Automatically tier backups to Azure Blob storage tiers and Azure Archive Storage. Restore options include full- and file-level recovery in Azure, recovery on-premises, Instant VM Recovery® and recovery to AWS and Google Cloud Platform.
Google-native backup and recovery					✔	Create Google-native snapshots of Google Cloud Platform Compute Engine instances. Automatically tier backups to Google Cloud Storage, including Archive Storage. Restore options include full- and file-level recovery in Google Cloud, recovery on-premises, Instant VM Recovery® and recovery to AWS and Microsoft Azure.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
NAS Backup	✓				✓	Protect and recover NAS file shares (SMB and NFS) effectively and at scale. Back up directly to Veeam repositories for short and long-term storage objectives. The first 500GB in each share is free with additional capacity licensed with VUL at 500GB per license. Long-term file version archival support for NAS data is not available with Community edition.
Agents for Microsoft Windows, Linux and MacOS	✓				✓	Agent-based backup for Windows, Linux and MacOS servers and workstations with automated deployment and management directly within the backup console.
Agents for AIX and Solaris					✓	Agent-based backup for AIX and Solaris servers and workstations to Veeam backup repositories.
Backup Copy jobs	✓	✓	✓	✓	✓	Copy any backups to your DR storage location of choice automatically, including validation and remediation to ensure copies are available and reliable.
End-to-end encryption	✓	✓	✓	✓	✓	Secure the backup data and network transfers with end-to-end AES 256-bit encryption without any negative impact on built-in compression and WAN acceleration data reduction ratios. All editions include at-source (during backup), in-flight (network traffic) and at-rest (tape) encryption. VUL, Enterprise Plus and Enterprise editions includes lost password protection.
ROBO Guest Interaction Proxy			✓	✓	✓	Reduce the load on the central backup server, improve scalability and simplify networking for environments with a large number of remote office/branch offices (ROBOs) when performing application- aware processing and guest file system indexing.
Proxy affinity			✓	✓	✓	Avoid cross-node network traffic across multiple all-in-one Veeam appliances by associating backup proxies and backup repositories located on the same appliance together.
Storage latency control			✓	✓	✓	Set the maximum acceptable I/O latency level for production datastores to ensure that backup and replication activities do not impact storage availability to production workloads. Enterprise edition includes a global latency setting. VUL and Enterprise Plus edition enables setting customization on a per-datastore basis.
Built-in WAN Acceleration			✓	✓	✓	Get replicas and backup copy jobs off site up to 50x faster and save bandwidth. Enterprise edition supports built-in WAN Acceleration to Veeam Cloud Connect targets only. VUL and Enterprise Plus edition supports built-in WAN Acceleration to any target.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
Backup from Storage Snapshots				✔	✔	Create image-based backups and replicas as often as necessary with little to no impact on production from the world's leading storage providers for backups from primary and secondary storage snapshots. This includes a growing list of storage alliance and ecosystem partners.
Primary Snapshot Orchestration	✔	✔	✔	✔	✔	To support customers in implementation of better RTO and RPO windows, additional storage snapshots, including the option for Oracle and SQL Server transaction log backup, can be created on the primary storage arrays for a wide range of storage systems. Orchestration for secondary storage snapshots is included for several storage vendors in VUL and Enterprise Plus edition.
ENTERPRISE APP BACKUP PLUGINS						
Veeam Plug-in for Oracle RMAN, SAP HANA and SAP on Oracle				✔	✔	Stream SAP HANA and Oracle RMAN backups into Veeam repositories and recover with ease, leveraging native SAP and Oracle backup and restore capabilities, including simultaneous backup copies to multiple backup targets for Oracle Database Enterprise Edition Supported platforms for databases are Microsoft Windows, Linux, Solaris and AIX.
STORING BACKUPS						
Hardened Repository	✔	✔	✔	✔	✔	Keep your backups safe in hardened, malware- and hacker-proof repositories, with immutable backups preventing encryption and deletion by ransomware and malicious actors.
Veeam Cloud Connect Backup		✔	✔	✔	✔	Get your backups off site with fully integrated, fast and secure backup, and restore to the cloud with one of our Veeam-powered service providers. This built-in capability makes it easy to establish a connection with a Cloud Repository, target the repository for Backup Jobs, and even backup directly to the cloud.
VMware Cloud Director Backup	✔	✔	✔	✔	✔	Back up vApp and VM metadata and attributes and restore vApps and VMs directly to vCloud with full support for fast provisioned VMs. All editions include integrated visibility of vCloud Director infrastructure, backup via VeeamZIP (including backup of vApp and VM metadata and attributes) and direct restore to vCloud. Enterprise edition supports scheduled incremental backup jobs of vCloud VMs, as well as customer portal integrations via RESTful API. VUL and Enterprise Plus supports self-service, tenant managed backup and restore via enterprise manager, as well as native vCloud Director authentication.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
Scale-out Backup Repository™ (SOBR) with object storage support			✓	✓	✓	Provide an abstraction layer across individual storage devices to create a single virtual pool of backup storage for your backups, with Performance, Capacity and Archive Tiers automating management for the entire backup lifecycle. With Enterprise edition, users can create two Scale-out Backup Repositories with three active extents and unlimited inactive extent (placed in maintenance mode). VUL and Enterprise Plus edition have no limitations on the number of repositories or extents.
SOBR Capacity Tier			✓	✓	✓	Leverage native object storage integration with AWS, Microsoft Azure, IBM Cloud, Google Cloud Storage and most S3-compatible object storage offerings for on-premises and in the cloud without "cloud tax" and vendor lock-in associated with secondary storage appliances. Protect your backups against ransomware and malicious actors with S3 Object Lock based immutability.
SOBR Archive Tier				✓	✓	Reduce the costs of long-term data archival and retention by up to 20 times and replace manual tape management with new native support for Amazon S3 Glacier (including Glacier Deep Archive) and Microsoft Azure Archive Storage through end-to-end backup life cycle management.
Per-VM backup files			✓	✓	✓	Store each VM in a separate backup file to enhance backup performance when backing up to enterprise-grade storage and deduplicating storage appliances. This is accomplished by leveraging parallel VM processing to enable multiple, concurrent write streams.
De-duplicating storage appliance integration			✓	✓	✓	Get faster backup performance by backing up to deduplicating storage integrations (Dell EMC Data Domain Boost, HPE StoreOnce Catalyst, Quantum DXi and ExaGrid Accelerated Data Mover).
Native tape support	✓	✓	✓	✓	✓	Back up files to any LTO compatible tape devices. All editions support File to Tape jobs. Enterprise and above editions add Backup to Tape jobs for tighter integration with primary backup jobs, dedicated media pool type to simplify GFS retention and global media pools to simplify the usage of multiple tape devices. VUL and Enterprise Plus also provide NDMP support for NAS device protection.
Veeam External Repository					✓	Connect to object storage repositories managed by cloud-native Veeam backups, offering to perform restores from those backups, or copy them out of the cloud to achieve the 3-2-1 rule compliance.
File versions archiving					✓	Archive all or only specific file versions for a longer period of time with file share backup jobs.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
REPLICATION						
Veeam CDP (Continuous Data Protection)[1]				✔	✔	Eliminate downtime and minimize data loss for Tier-1 VMware workloads with built-in CDP and achieve immediate recoveries to a latest state or desired point in time providing the best recovery point objectives (RPOs).
Image-based VM replication	✔	✔	✔	✔	✔	Replicate VMs periodically on site for High Availability or off site for DR.
Assisted failover and failback	✔	✔	✔	✔	✔	Replica rollback and assisted failover and failback.
Replication from a backup	✔	✔	✔	✔	✔	Create replicas directly from backup files without impacting production.
Planned failover	✔	✔	✔	✔	✔	Facilitate data-center migrations with zero data loss.
Veeam Cloud Connect Replication		✔	✔	✔	✔	Ensure Availability of your mission-critical applications with fully integrated, fast and secure cloud-based DR through a Veeam- powered Disaster Recovery as a Service (DRaaS) provider. In the event of disaster or other use case, replicas stored in the repository can be used for partial- or full-site failover.
1-Click failover orchestration			✔	✔	✔	Built-in failover plan orchestration enabling easy, 1-click-site failover to minimize unplanned downtime.
VMware Cloud Director Replication				✔	✔	Dedicated replication job type enables service providers to perform vApp replication within and across vCloud Director instances. The replication job processes vApp VMs and metadata (such as networking or VM start order) to create a ready-to-use replica vApp in the target vCloud Director that can be leveraged instantly in case of a disaster.
RESTORE						
Entire VM restore	✔	✔	✔	✔	✔	Restore selected VMs to the original or a new location.
VM disks and files restore	✔	✔	✔	✔	✔	Restore individual virtual disks and other VM files (such as VMX) from both backups and Storage Snapshots.
Direct Restore to AWS, Azure, and Google Cloud Platform	✔	✔	✔	✔	✔	Restore or migrate on-premises, Windows based or Linux-based VMs, physical servers and endpoints directly to AWS, Microsoft Azure, and Google Cloud Platform.
INSTANT RECOVERY						
Instant VM Recovery	✔	✔	✔	✔	✔	Accelerate DR and lower RTOs with Veeam's patented instant recovery engine. Restore ANY backup instantly to a VMware, Microsoft Hyper-V or Nutanix AHV VM by running it directly from backup.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
Instant disk recovery	✔	✔	✔	✔	✔	Reduce your instant recovery footprint by instantly restoring only the required disks of large vSphere VMs, as opposed to restoring the entire VM (for example only OS disk, or only data disks).
Instant database recovery			✔	✔	✔	Instantly recover any Microsoft SQL Server and Oracle database from backup to the latest state or to an earlier point-in-time to any production database server or cluster (physical or virtual) in minutes regardless of its size.
Instant file share recovery	✔	✔	✔	✔	✔	Publish SMB file shares from backup to the latest state, or to an earlier point-in- time, on the selected mounted server, and enable your users to instantly access their data in this temporary SMB file share while you're getting the problem fixed or data restored.
FILE & ITEM LEVEL RECOVERY						
File-Level Recovery	✔	✔	✔	✔	✔	Recover files from 19 common file systems used by Windows, Linux, BSD, Mac OS, Novell, Solaris and AIX directly from image-level backup.
Veeam Explorer™ for Storage Snapshots	✔	✔	✔	✔	✔	Restore individual VMs, guest files and application items from storage snapshots.
Veeam Explorer for Microsoft Active Directory	✔	✔	✔	✔	✔	Search and restore for all Active Directory (AD) object types, such as users, groups, computer accounts and contacts, including user and computer password recovery. All editions support single user and computer account restore (including the ability to restore passwords) as well as export into LDIFDE format. Enterprise and above editions include multi-select restore and container restore, Group Policy Objects (GPO) restore, AD-integrated DNS records restore and Configuration Partition objects restore.
Veeam Explorer for Microsoft Exchange	✔	✔	✔	✔	✔	Get instant visibility into Microsoft Exchange backups for item recovery of individual Exchange items (emails, appointments, notes, contacts, etc .), online archive mailboxes and hard-deleted items. Leverage comprehensive e-Discovery functionality, including query result size estimation and detailed export reports. All editions, including Community Edition, support Exchange mailbox item restore via save, send and PST export. VUL, Enterprise Plus and Enterprise editions support restoring back to the original mailbox.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
Veeam Explorer for Microsoft SQL Server	✓	✓	✓	✓	✓	Restore individual SQL databases with ease, without needing an extensive SQL background or having to search for database and transaction-log files. All editions support local, point-in-time SQL database file restores. VUL, Enterprise Plus and Enterprise editions include agentless, transaction log backup and replay, transaction-level recovery of databases and SQL objects (tables, stored procedures, views, etc .) back to the original or new SQL server.
Veeam Explorer for Microsoft SharePoint	✓	✓	✓	✓	✓	Get instant visibility into SharePoint backups with advanced search-and-browse capabilities for quick recovery of individual SharePoint items and entire sites. All editions support SharePoint item restore via save, send and export. VUL, Enterprise Plus and Enterprise editions also support entire-site restore, as well as restore to the original location.
Guest file system indexing	✓	✓	✓	✓	✓	A catalog of guest files enabling effortless search of individual files to find and restore a file without knowing its precise location, or the time when it was deleted. All editions provide a catalog of guest files in backups currently on disk. VUL, Enterprise Plus and Enterprise editions also provide a catalog of archived backup files and integrate with one click to enable restores directly from search results.
Veeam Explorer for Oracle			✓	✓	✓	Restore individual Oracle databases with ease without needing an extensive Oracle background or having to search for database and transaction-log files. Includes agentless, transaction-log backup, archived log management and transaction- level recovery of databases back to the original or new Oracle Server, as well as the ability to publish and export Oracle databases. VUL and Enterprise Plus edition also provides a simple user interface for database restores in emergency situations from native RMAN backups.
DATALABS						
Secure Restore	✓	✓	✓	✓	✓	Increase security and reduce interruptions caused by viruses and recover back to a safe, uncompromised restore point through optional antivirus scan directly in backup files prior to recovery; agnostic support for Windows Defender, ESET, Kaspersky and Symantec Protection Engine, as well as the ability to extend to other third-party antivirus tools.
SureBackup®[2]			✓	✓	✓	Automatically test and verify recoverability of every backup by running the VM directly from a backup file in a fenced-off copy of your production environment, including support for custom application test scripts.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
SureReplica[1]			✓	✓	✓	Automatically test and verify recoverability of every replica by performing a failover to a fenced-off copy of your production environment, including support for custom application test scripts.
Staged Restore			✓	✓	✓	Clean up sensitive data from the selected restore point in a fenced-off copy of your production environment before restoring the modified VM state to production.
On-Demand Sandbox™			✓	✓	✓	Run one or more VMs directly from a backup a fenced-off copy of your production environment to troubleshoot issues, test upgrades, train staff on a working copy of the production environment, without impacting business operations.
On-Demand Sandbox from Storage Snapshots[1]				✓	✓	Use storage snapshots to quickly create fenced-off copies of your production environment operating at full I/O performance.
Data Integration API			✓	✓	✓	Put your backups to work and reuse your data by enabling third-party applications and scripts to instantly access the content of any Veeam backup for data mining, security analysis, compliance checks, and other data reuse scenarios.
SELF-SERVICE						
1-Click File and VM recovery portal for help desk operators			✓	✓	✓	Restore guest files and VMs with a single click through a web UI.
Microsoft Exchange item recovery portal for help desk operators			✓	✓	✓	Restore missing mailbox items back to the original mailbox with a single click through a web UI.
Microsoft SQL Server database recovery portal for SQL DBAs			✓	✓	✓	Restore individual databases back to the original server or a new SQL server with a single click through a web UI.
Oracle database recovery portal for Oracle DBAs			✓	✓	✓	Restore Oracle databases to the original server or a new Oracle server with a single click through a web UI.
Self-service file restore portal				✓	✓	File-restore portal with automatic VM detection and automatic delegation based on the Local Administrator group membership.
Recovery delegation				✓	✓	Enable full self-service for all web UI recovery features by delegating recoveries for individual VMs and groups of VMs to specific users or groups such as local IT staff, application owners, department members, etc.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup & Replication	Community*	Standard Socket	Enterprise Socket	Enterprise PLUS Socket	VUL	Description
Native VMware Cloud Director plug-in				✔	✔	This capability allows service providers to extend the vCloud Director tenant UI to include Veeam Backup & Replication functionality, enabling tenants to manage their own backups and restores without leaving the convenience of the vCloud Director web console. This integration is based on Veeam's existing vCloud Director Self-Service Backup & Restore Portal.
Role-Based Access Control (RBAC)1				✔	✔	Allow users to securely manage and monitor backups and restores of their own VMs through the self-service backup and restore portal for VMware vSphere with access delegation rules based on vSphere role, VM privilege or VM tags.
MANAGEMENT						
Built-in management for Veeam Agents	✔				✔	The ability to organize, automate, and deploy Windows, Linux and Mac agents with protection groups based on Active Directory containers or CSV files and configure custom automatic update and deployment schedules.
Veeam Backup Enterprise Manager — Centralized Management Web UI		✔	✔	✔	✔	Get a web-based, consolidated view of your distributed deployment in a single pane of glass, without having to login to individual backup servers, including federation of multiple backup servers, centralized reporting and consolidated alerting. Paid editions include monitoring and reporting across multiple backup servers, as well as the ability to start and stop jobs. VUL, Enterprise Plus and Enterprise editions also include full job management functionality and the ability to perform restores.
PowerShell module	✔	✔	✔	✔	✔	Use PowerShell Cmdlets to create custom workflows and automate Veeam tasks
RESTful API				✔	✔	Integrate into existing systems and execute and automate Veeam tasks using the RESTful API framework. Service Providers using rental licenses have API access across all editions.

[1] VMware only

[2] VMware and Hyper-V Only

* Community Edition is limited to 10 licenses

Learn more
www.veeam.com

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam ONE™

✔ Fully supported ✔ Partially supported

Veeam ONE	Community	Socket-based	VUL	Notes/ description
MONITORING AND REPORTING				
Performance monitoring and alerting for backup activity	✔	✔	✔	Performance monitoring and alerting for backup activity displays Veeam Backup & Replication™ jobs, including their latest state and performance statistics. Veeam ONE™ has no limitations on the number of Veeam Backup & Replication servers that can be monitored. Note that Veeam ONE Community Edition is limited to one server.
Email notification and reports	✔	✔	✔	You can receive alarm notifications through a customizable email. Email customizations are not available in Veeam ONE Community Edition, and emails will not include all alarm details, such as time and affected object.
Veeam backup general reporting	✔	✔	✔	Veeam ONE features several reports, including overviews of Veeam Backup & Replication configurations, backup infrastructure audits, compliance reporting, protected virtual machines (VMs), capacity planning, tape reports and more.
Veeam Backup *for AWS*, Veeam Backup *for Microsoft Azure*, and Veeam Backup *for Google Cloud Platform*- native data protection support	✔	✔	✔	Monitoring and alerting: Veeam ONE users can now monitor the backup and snapshot job statuses of cloud-based workloads through predefined alarms and notifications. Backup status reporting: A predefined protected VMs report now includes VMs that are hosted in AWS, Microsoft Azure and Google Cloud Platform public clouds. This helps users in verifying the data protection status of each discovered VM, based on the agreed service level agreement (SLA).
Veeam Backup *for Nutanix AHV*	✔	✔	✔	Backup status reporting: A predefined protected VMs report now includes VMs that are hosted on Nutanix AHV. This helps users in verifying the data protection status for each discovered VM based on the agreed service level agreement (SLA).

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam ONE	Community	Socket-based	VUL	Notes/ description
Veeam Agent monitoring and reporting	✔	✔	✔	Support for Veeam backup agents includes several alarms and reports that help you analyze trends and determine how policies and backup jobs have evolved, as well as determine the backup status of agents, including which agents have no backup copies and how many are protected or unprotected. This is given with the backup state along with the total number of successes, warnings and failures.
Veeam Cloud Connect reporting	✔	✔	✔	Reports include information about the Veeam Cloud Connect infrastructure, like user quota usage, capacity planning for cloud repositories and configuration data for cloud gateways and repositories.
Performance monitoring and alerting for Veeam Cloud Connect	✔	✔	✔	Performance monitoring and alerting for Veeam Cloud Connect displays jobs for Veeam Cloud Connect backup servers only, including their latest state and performance statistics. Note: To monitor and report on managed clients, Veeam ONE needs to be deployed at the tenant's site. Veeam ONE Community Edition is limited to one week of performance.
Heatmaps	✔	✔	✔	Heatmaps help plan resource usage and easily optimize backup activities with real-time visibility into backup repositories and proxies. Veeam ONE Community Edition is limited to three proxy servers and three repositories.
MONITORING AND ALERTING FOR VIRTUAL INFRASTRUCTURE				
Performance monitoring and alerting for virtual infrastructures	✔	✔	✔	Display virtual infrastructure hosts, datastores and statistics. In Veeam ONE Community Edition, performance graphs only show one week of performance and emails include limited alarm details.
Application-level monitoring and management	✔	✔	✔	Ensure the availability of mission-critical applications, monitor select groups of applications and create additional alarms to monitor VM service states. Make sure processes aren't exceeding available resources, identify if critical applications, like anti-virus software, have been disabled, or indicate if unwanted changes have taken place. With Veeam ONE Community Edition, you can monitor an unlimited number of applications on up to three VMs.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam ONE	Community	Socket-based	VUL	Notes/ description
VMware vSphere and Microsoft Hyper-V reporting	✓	✓	✓	Individual report packs for vSphere and Hyper-V provide general configuration overviews for vCenter servers, datastores, clusters hosts and VMs, as well as SCVMMs, volumes, clusters, hosts and VMs, including information on the most common alarms raised for virtual infrastructure objects Optimization report packs also allow users to evaluate the efficiency of resource usage and to optimize VM resource provisioning to help achieve improved output and increased ROI for virtualized environments. With Veeam ONE Community Edition, optimization reports are limited to show only three VMs and historical performance reports are limited to the last 24 hours of data.
Multi-tenant reporting and monitoring	✓	✓	✓	Delegate reporting to system administrators over a limited scope of assigned VMs (VMware vSphere only).
Infrastructure assessment: Veeam Backup & Replication	✓	✓	✓	Assessment reporting is designed to verify that your Veeam Backup & Replication installation meets all known best practices.
Virtual infrastructure assessment	✓	✓	✓	Assessment reporting is designed to verify that your VMware vSphere and/or Microsoft Hyper-V infrastructure meets all known best practices.
VMware Cloud Director (VCD) monitoring and alerting	✓	✓	✓	Companies want to stay ahead of the issues impacting their virtual datacenters and ultimately their customers. To help, Veeam provides capabilities for monitoring teams of service providers to track the availability and configuration of their virtual data centers, including organizations and vApps.
VMware Cloud Director (VCD) reporting	✓	✓	✓	Businesses want to ensure that they are using the ideal configuration and optimizing their resources to get the most out of their virtual infrastructure when building and maintaining virtual data centers to provide a reliable service to their customers.
REPORT CUSTOMIZATION AND AUTOMATION				
Raw data analysis		✓	✓	Create custom Microsoft Excel reports using templates you define.
Customizable, predefined report templates	✓	✓	✓	Customize and save predefined report templates to create reports that fit your specific information and formatting needs. Veeam ONE Community Edition is limited to saving one customized report template.
Automated report generation and delivery		✓	✓	Schedule reports to be automatically generated and delivered to your inbox, dashboards, web portals or archives.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam ONE	Community	Socket-based	VUL	Notes/ description
GENERAL				
Business view	✔	✔	✔	Categorize and group clusters, hosts, VMs and datastores by criteria such as business unit, department, purpose and SLA. Leverage VMware Tags to assist with role- based access control and eliminate human error, including tagging sensitive data by location to meet compliance. Easily monitor apps, regardless of geographic location and categorize business view objects based on multiple parameters to gather insight on how infrastructure investments are being leveraged. You also can assign administrators to manage specific parameters.
Alarm customization and modelling		✔	✔	Ensure complete monitoring coverage by creating custom alarms for events unique to your virtual environment. Model alarms against past performance data to understand potential alarm frequency and avoid inadvertent alert storms and missed events.
INTELLIGENT AUTOMATION				
Veeam Intelligent Diagnostics	✔	✔	✔	Proactive identification and alerting of known issues within your Veeam environment, providing self-support resolutions via hotfixes or reconfigurations. For Veeam ONE Community Edition, all of the alarms will continue to work but the details for each alarm will not be shown.
Remediation actions	✔	✔	✔	Predefined and automated resolution of common backup and virtual infrastructure-related issues with self-healing remediation actions help to reduce IT management costs . With Veeam ONECommunity Edition, remediation actions are available, but require manual approval.
Predefined performance dashboards for backup and virtual infrastructures	✔	✔	✔	Get at-a-glance views of backup jobs and the backup infrastructure, as well as view the VMs, hosts and clusters that are the highest consumers of CPU, memory, network and datastore resources. Also, view datastore utilization metrics, including metrics relating to VM snapshots and guest disk space usage.
Customized reporting dashboards	✔	✔	✔	Create your own at-a-glance views of your virtual environment. Veeam ONE Community Edition allows for one customized reporting dashboard to be created.

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam ONE	Community	Socket-based	VUL	Notes/ description
CAPACITY PLANNING				
Capacity planning: VMware vSphere and Microsoft Hyper-V		✔	✔	Individual report packs for VMware vSphere and Microsoft Hyper-V are designed to forecast when the available virtual infrastructure resources will reach their minimum levels. The packs provide recommendations on resource allocation and load balancing to optimize performance and resource utilization in your respective virtual environment and to avoid possible performance bottlenecks.
Deployment project reports		✔	✔	Leverage fully automated "what-if" modeling scenarios to safely add or remove VMs or hosts from your infrastructure and to accurately calculate the required resources needed for each project.
CHARGEBACK AND BILLING				
Chargeback analysis reports		✔	✔	Chargeback reports help you understand the costs associated with individual virtual infrastructure users, as well as help you perform chargeback by analyzing VM configuration and resource usage, including CPU, memory and storage. Four reports are included: Host Configuration Chargeback, Host Resource Usage Chargeback, VM Configuration Chargeback and VM Performance Chargeback.
Veeam backup: Billing	✔	✔	✔	This report provides chargeback capabilities for Veeam Backup & Replication and helps assess storage management costs for backup and replication operations, providing data that can be used in financial and billing reports.

Learn more
www.veeam.com

© 2021 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

ExaGrid
Technical Specifications

Application Support – All ExaGrid Appliances

Backup Applications, Utilities, and Methods

- Acronis Backup & Recovery
- Arcserve Backup, D2D, UDP
- BridgeHead
- Carbonite Amanda
- Commvault
- Dell EMC NetWorker
- HYCU
- IBM TSM (Spectrum Protect)
- Idera SQL Safe
- Micro Focus Data Protector
- Quest LiteSpeed, NetVault
- Quest vRanger
- Red Gate SQL Backup
- StorageCraft ShadowProtect
- Veeam Backup & Replication
- Veritas Backup Exec
- Veritas NetBackup
- Veritas System Recovery
- Zerto Offsite Backup
- POSIX-compliant TAR files
- Oracle RMAN
- SQL Dumps

Site Configurations

- Scalable from 6TB to 2.69PB full backup (24TB to 6.14PB raw) with 32 servers in a single system
- Hub-and-spoke architecture
- Mix and match different size/age appliance models
- ExaGrid sites (one hub and up to 15 spokes)
- Standard IP WAN connection between sites
- WAN encryption for secure data transfer
- Bandwidth throttling for WAN efficiency

Appliance Configurations

ExaGrid Model	EX6	EX10	EX18	EX27	EX36	EX52	EX84
Chassis	2U	2U	2U	2U	2U	3U	4U
Full Backup	6TB	10TB	18TB	27TB	36TB	52TB	84TB
Disks	6	8	12	12	12	16	24
Disk Size	4TB	4TB	4TB	6TB	8TB	8TB	8TB
Memory	8 GiB	16 GiB	32 GiB	32 GiB	64 GiB	64 GiB	128 GiB
CPUs / Cores	1 CPU 4 cores		1 CPU 4 hyper-threaded cores		1 CPU 6 hyper-threaded cores	1 CPU 8 hyper-threaded cores	1 CPU 16 hyper-threaded cores
NIC Ports	Four 1GbE; two optional 10GbE or 25GbE					Two 1GbE; two optional 10GbE, 25GbE or 40GbE	
All Servers: Intel or AMD processors / Enterprise SAS drives / FIPS 140-2 Self-Encrypting SAS drives for all SEC systems / RAID6 with hot spare / Redundant fans / Redundant power supplies / ExaGrid software / 19-inch rack-mount chassis							

Regulatory Approvals

- UL/CSA 60950-1
- IEC 60950-1
- EN 60950-1
- FCC Class A
- CB Certificate
- ICES-003
- EN 55024
- EN 55022/32
- CISPR 22/32
- VCCI
- EN 61000-3-2/3
- RoHS Compliant
- WEEE Declaration EU Directive
- IEC 62368-1
- EN 62368-1
- EN55035

Environmental Requirements

	Power	Temperature	Humidity	Cooling (Operating)
All Appliances	Frequency: 50-60 Hz / Voltage: 100-240 VAC Two NEMA 5-15R 3-prong receptacles/appliance Country kits available for international orders EX6: 141W (1.18A @ 120VAC) EX10: 172W (1.43A @ 120VAC) EX18: 207W (1.73 @ 120VAC) EX27: 235W (1.96A @ 120VAC) EX36: 288W (2.44A @ 120VAC) EX52: 361W (3.01A @ 120VAC) EX84: 433W (3.61A @ 120VAC)	10° to 35° C (50° to 95° F)	20% to 80% non-condensing	EX6: 481 BTU/hr EX10: 587 BTU/hr* EX18: 706 BTU/hr* EX27: 801 BTU/hr* EX36: 982 BTU/hr* EX52: 1231 BTU/hr* EX84: 1477 BTU/hr* ** including -SEC*

United States 350 Campus Drive / Marlborough, MA 01752 / (800) 868-6985
United Kingdom 200 Brook Drive / Green Park, Reading, Berkshire RG2 6UB / +44 (0) 1189 497 051
Singapore 1 Raffles Place, #20-61 / One Raffles Place Tower 2 / 048616 / +65 6808 5574

exagrid.com

ExaGrid reserves the right to change specifications or other product information without notice. ExaGrid and the ExaGrid logo are trademarks of ExaGrid Systems, Inc. All other trademarks are the property of their respective holders. ©2020 ExaGrid Systems, Inc. All rights reserved.

Administrator's Guide

For ExaGrid Software Version 6.0 and Later

PN: 210-0165-24

Copyright

No part of this document may be reproduced or transmitted in any form or by any means without the express written permission of ExaGrid Systems, Inc.

© 2003-2020 ExaGrid Systems, Inc. All rights reserved. Printed in USA.

This document, the associated software, and the associated online documentation are the property of ExaGrid Systems, Inc. or its licensors, as applicable, and are loaned to the user under the terms of the ExaGrid Systems, Inc. End User License Agreement. Unauthorized copying or use of the software or any associated materials is contrary to the property rights of ExaGrid Systems and is a violation of state and federal law.

Trademarks

ExaGrid is a registered trademark of ExaGrid Systems, Inc. ExaGrid Systems (Logo), InstantDR, InfiniteFiler, GRIDdisk and Intelligent Disk-based Data Protection are all trademarks of ExaGrid Systems, Inc. All third-party trademarks are the property of their respective owners.

License Agreement

This document, the associated software and the associated online documentation are the property of ExaGrid Systems, Inc. or its licensors, as applicable. The use of these materials and the software is strictly limited to those users who have signed the ExaGrid Systems, Inc. End User License Agreement.

For any software acquired directly or indirectly on behalf of a unit or agency of the United States Government, whether for civilian agencies or for units of the Department of Defense, the software is a commercial item or commercial computer software (and documentation), and pursuant to Section 12.212 of the Federal Acquisition Regulations (FAR) or DoD FAR Supplement Section 227.7202, is provided under restricted rights as enumerated in the End User License Agreement.

Preface

This guide provides an overview of the ExaGrid System and getting started instructions for the novice ExaGrid Administrator.

Detailed descriptions of the ExaGrid System and advanced administrator procedures can be found in the ExaGrid System's online help.

Related Documentation

Title	Part Number
Assembling and Initializing Your ExaGrid System	210-0164-xx
ExaGrid Release Notes	210-0174-xx

Contacting ExaGrid Technical Support

If you have technical questions about this product that are not answered in this document contact ExaGrid Support:

- Email: support@exagrid.com
- Phone: 800-868-6985 and at the prompt, press 2.

Glossary

Understanding these terms will help you better understand the ExaGrid System.

Disaster Recovery – In multi-site systems, in the event of a catastrophic loss of an ExaGrid Site, the ability to recover backup data that has been replicated to one or two surviving ExaGrid Sites.

ExaGrid Server - The hardware on which the ExaGrid Software runs. There are several ExaGrid Server models, each with a different capacity.

ExaGrid Site - A logical collection of ExaGrid Servers. An ExaGrid System can be made up of one or more ExaGrid Sites.

File Collection - A set of files associated with a backup application's backup to an ExaGrid Share. ExaGrid reports will report one or more file collections associated with a backup job.

Hub – ExaGrid Systems use a hub-and-spoke topology. In multi-site ExaGrid Systems, one site is designated the hub. When logged into the hub, the hub and all spoke sites are displayed in the navigation tree and system-level reports roll up information from all spokes. When logged into a spoke, only the spoke site and its hub are displayed in the navigation tree. system and site-level reports roll up information for the spoke site only.

InstantDR™ Share - A copy of a share at a second ExaGrid Site. When you create an ExaGrid Share you have the option of also creating an InstantDR Share on an ExaGrid Server in another ExaGrid Site. InstantDR Shares are used in disaster recovery operations, remote tape copy operations, and in testing data recovery from another ExaGrid Site.

Landing Zone Tier – The portion of an ExaGrid Server's disk space dedicated to incoming backup data. The size of the Landing Zone Tier varies by ExaGrid Server model. The Landing Zone Tier is used to store the most recent version of your backup(s). They are stored in their entirety. Since most restores are done from the most recent backup, the Landing Zone Tier ensures rapid restores.

Replica - A copy of a share that has been replicated to another ExaGrid Site.

Replication - A share setting that instructs the ExaGrid System to maintain copies of a share on one or more ExaGrid Sites.

Retention Tier – The portion of an ExaGrid Server that stores and manages deduplicated "Versions" of backup files.

Share - A mount point/folder on an ExaGrid Server to which backup applications write.

Share Migration - To move a share and its contents from one ExaGrid Server to another.

Spoke – ExaGrid Systems use a hub-and-spoke topology. In multi-site ExaGrid Systems, one site is designated the hub. When logged into the hub, the hub and all spoke sites are displayed in the navigation tree and system-level reports roll up information from all spokes. When logged into a spoke, only the spoke site and its hub are displayed in the navigation tree. system and site-level reports roll up information for the spoke site only.

Versions – The deduplicated copies of a backup job.

Contents

1. Introduction 1

Shares 1

ExaGrid Hub Sites and Spoke Sites 2

Data Flow 3

The ExaGrid User Interface 4

Logging In 7

2. Networking 9

Network Settings 10

Bonding 12

Ports Used Between ExaGrid Sites 12

3. TLS Certificates 13

4. Managing and Securing Access 19

Retention Time-Lock Overview 20

User Roles 24

Creating or Modifying ExaGrid Local Users 25

Two-Factor Authentication 27

Managing Active Directory 29

Share Access Policies 32

Managing SSH Keys 33
Support Accounts 34
Root Access Management 35

5. Using Server Message Block (SMB) Signing 37
38

6. ExaGrid Shares 39
Primary Shares 39
InstantDR™ Shares 42
Share Protection 43
Share Migration 43
Share Access Policies 44
Creating Shares 47

7. Share Replication 51
Creating Replicas and InstantDR Shares 52
Hub Site To Spoke Site Replication Example 53
Spoke Site to Hub Site Replication Example 55
Deleting Shares and Replicas 55
Replication Scheduler 56
Throttling Replication 56
Encrypting Replication 58

8. System Record Keeping 61
ExaGrid Site Worksheet 62
ExaGrid Server Worksheets 63
ExaGrid Share Worksheet 65

9. The IPMI Port 67
Configuring The IPMI Port 68

Resetting the IPMI Network Port User Password 69

10. Encryption Capable ExaGrid Servers 73

Encryption at Rest 74

Boot Protection 79

11. Recovering From Malicious Software Attacks, or a Loss of ExaGrid Servers 83

Before an Event Occurs 83

Recovering From Malicious Software Attacks 84

Recovering From The Loss of ExaGrid Servers - Disaster Recovery 88

12. ExaGrid Reports 91

Summary Reports 92

Detailed Reports 96

13. Monitoring Your ExaGrid System 97

SMTP Email Notifications 97

SNMP Traps 98

Status Monitor 98

ExaGrid MIB and SNMP Objects 99

System Logging 105

14. ExaGrid System Maintenance 109

Shutting Down and Restarting an ExaGrid Server 109

Software Upgrade 110

Adding Additional ExaGrid Servers 110

Hub Site and Spoke Site Systems 112

Reconfiguring a Site's Topology 112

Adding a Spoke Site 113

Defining NTP Servers 114

15. Troubleshooting 117

Monitoring Status of Ongoing Operations 117

Replacing Disks 118

Disk Safety Halt Protection 119

Disk Numbering 120

Replacing Power Supplies 120

Replacing Fans 121

Network Debugging Tools 123

Notifications 124

Consumption Issues 125

Replication 127

Performance Tuning 128

Introduction

Your ExaGrid System is a tiered backup storage device to which your backup environment sends backup data.

ExaGrid uses deduplication technology that allows multiple versions of your backups to be stored in a very small storage footprint.

This chapter provides a high-level overview of the ExaGrid System's features and capabilities. Consult the ExaGrid on line help for detailed descriptions and step-by-step instructions on how to use the ExaGrid System.

Caution – Do not install third party software on your ExaGrid Server. Installing third party software can result in a loss of data and void your ExaGrid Warranty.

ExaGrid Servers differ in capacity, physical size and network connectors. All ExaGrid Servers run the same ExaGrid software and can be commingled in an ExaGrid Site.

Shares

ExaGrid Shares are mount points to which your backup application backs up data.

There are two basic types of ExaGrid Shares:

- **Primary Shares:** Mount points created on ExaGrid Servers to which your backup application can write.
- **InstantDR Shares:** Replicas of Primary Shares.

For details see "ExaGrid Shares" on page 39

ExaGrid Hub Sites and Spoke Sites

The ExaGrid System uses a hub and spoke topology in which there is a single ExaGrid Site known as the hub, connected to one or more ExaGrid Sites known as spokes.

Note – ExaGrid Hub Sites can replicate shares to one or two spoke sites. ExaGrid Spoke Sites can replicate shares to their Hub Site; the share replicated to the Hub Site can then be replicated to another Spoke Site. ExaGrid Spoke sites cannot replicate shares between each other.

When logged into an ExaGrid Server on a hub site you can:

- See the hub site and all spoke sites displayed in the ExaGrid user interface's navigation tree. The hub site is always listed as the top-most site in the navigation tree.
- View reports for the hub site and all spoke sites
- View events for the hub site and all spoke sites
- Perform hub and spoke site related management tasks

When logged into an ExaGrid Server on a spoke site you can:

- See the spoke site and the hub site displayed in the ExaGrid user interface's navigation tree. Other spoke sites in the system are not displayed.
- View reports for the spoke site
- View events for the spoke site
- View shares on the hub to which the spoke shares replicate
- Perform spoke site related management tasks

Data Flow

The following diagram details the data flow descriptions presented in this section:

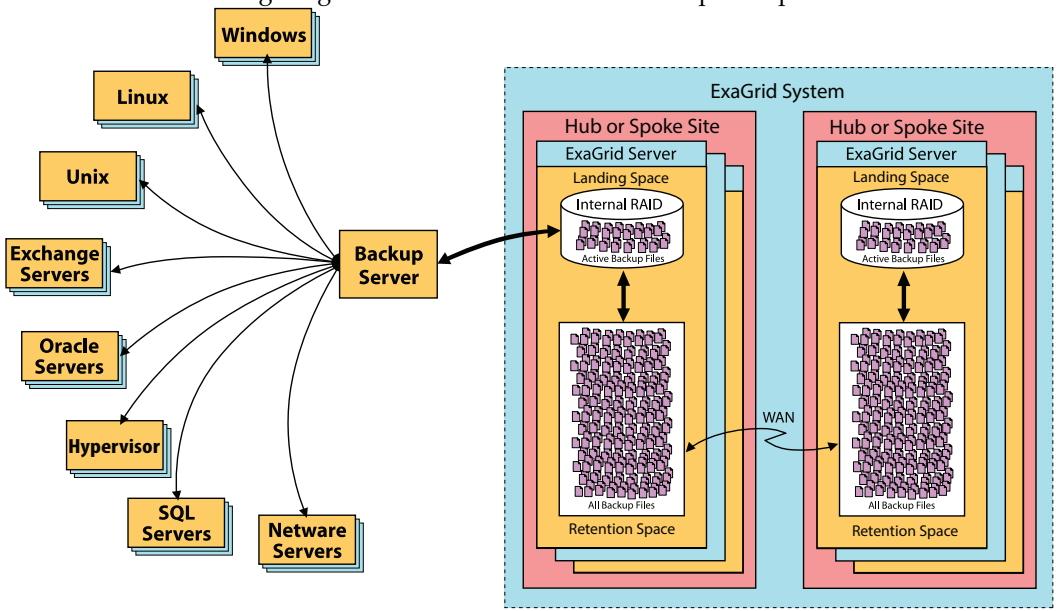

Data flows into an ExaGrid Share from a backup application:

- Backup application writes to an ExaGrid share.
- When the backup is complete, the ExaGrid deduplication engine will deduplicate the backed up data.

Note – As more and more backups land, are retained, and deduplicated, deduplication ratios will improve.

- A full copy of the most recent backup is maintained in the Landing Zone Tier.
- The deduplicated versions are kept in the Retention Tier.
- In multi-site ExaGrid Systems, you can specify whether or not to maintain a copy of a share on a second or optionally third ExaGrid Site.

When a backup application requests a restore:

- Since most restores are done from the most recent backup, the ExaGrid system keeps the most recent backup in its entirety. This ensures extremely fast restores.
- Restores from an older backup are done by re-assembling the backup from backup versions.

The ExaGrid User Interface

ExaGrid presents a unified, browser-based user interface for all ExaGrid Sites and ExaGrid Servers in an ExaGrid System. A single login gives you reports on, and administrative access to, all ExaGrid Sites and ExaGrid Servers in your ExaGrid System.

The following screen capture details the main components of the ExaGrid user interface when ExaGrid Site is selected:.

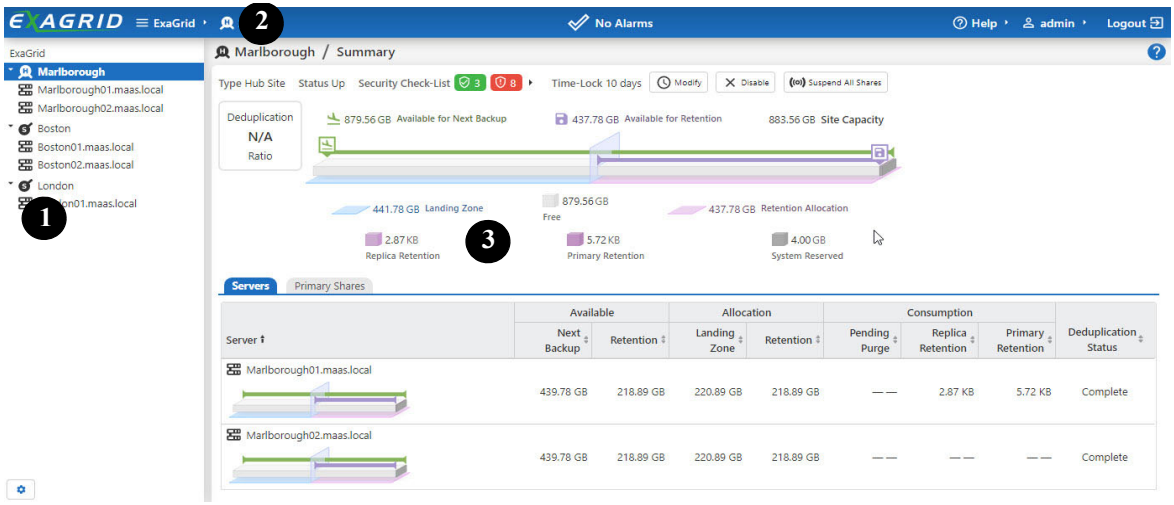

1. **Navigation Tree:** A tree view of the ExaGrid Sites and ExaGrid Servers in the ExaGrid System. The hub site is always displayed as the top-most site.
2. **Main Menu:** Drop down menus that support functions for the currently selected item in the navigation tree.
3. **Report/Management Area for Sites**:
   - Type: Identifies the site as either a hub or spoke
   - Security checklists: These lists allow you to quickly determine whether the security settings in your ExaGrid Site meet ExaGrid's recommended best practices. Click on a list item to display the page where the security setting can be configured. For details on ExaGrid's security and access best practices. See "Managing and Securing Access" on page 19.
   - Time Lock: Displays the site's Retention Time-Lock setting in days and allows you to make Retention Time-Lock change requests to the site's Security Officer. See "Retention Time-Lock Overview" on page 20.

- Suspend All Shares: If a ransomware or other malicious software attack is detected, you can suspend all shares in the ExaGrid Site and if necessary, begin recovering from the event. See “Recovering From Malicious Software Attacks” on page 84.
- High-level deduplication and consumption reports for the site, its servers, and its shares

The following screen capture details the main components of the ExaGrid user interface when ExaGrid Server is selected

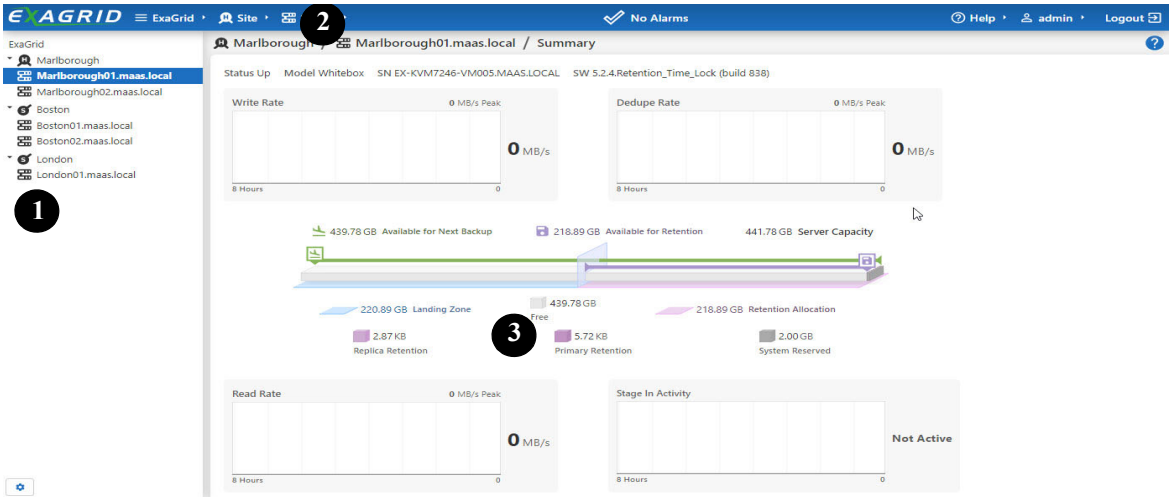

1. **Navigation tree:** A tree view of the ExaGrid Sites and ExaGrid Servers in the ExaGrid System. The hub site is always displayed as the top-most site.
2. **Main menu:** Drop down menus that support functions for the currently selected item in the navigation tree.
3. **Report/Management Area**: Provides:
   - Graphs that detail Write Rates, Deduplication Rates, Read Rates, and Stage in Activity.
   - A graphical representation of the capacity and consumption of the selected ExaGrid Server.

ExaGrid Menus

ExaGrid Menus vary by ExaGrid, Site, and Server.

To display a menu, in the top-most menu bar:

- Click **ExaGrid** to display menu items for ExaGrid.

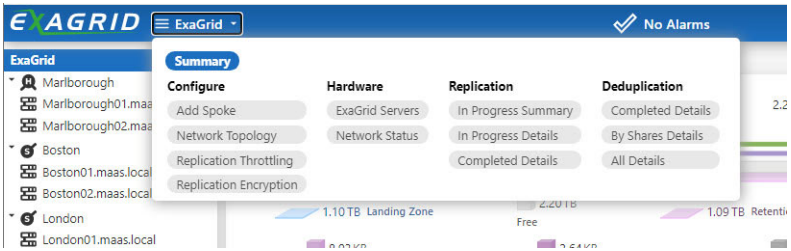

- Click **Site** to display menu items for the selected ExaGrid Site.

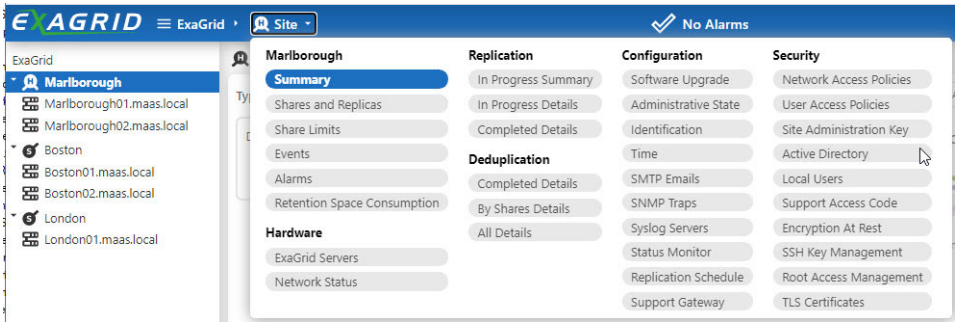

- Click **Server** to display menu items for the selected ExaGrid Server.

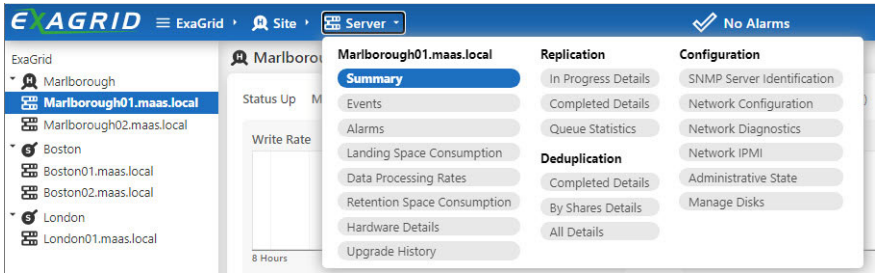

Logging In

You can access the ExaGrid user interface through a web browser on your network. To log in to the ExaGrid user interface:

1. Open a browser and enter the URL or IP address of an ExaGrid Server. You can specify any ExaGrid Server. For example:
   `https://ExaGridServerName` or `https://192.168.55.3`

   To take advantage of all of your ExaGrid System's capabilities, ExaGrid recommends that you log in using Hyper Text Transfer Secure (**HTTPS**) Protocol or employ TLS certificates.
2. Press **Return**. The ExaGrid Login screen is displayed.
3. Enter a user name and password for either an:
   - ExaGrid local user
   - Active Directory user or group that has been mapped into an ExaGrid role. For details, see " Managing Active Directory" on page 29.
4. If two-factor authentication is enabled for your user, when prompted, provide the access code sent to your cell phone. See "Two-Factor Authentication" on page 27.
5. Click **OK** and the ExaGrid user interface is displayed.

If you log in to an ExaGrid Hub Site, your user's role will determine what you can manage and view in the ExaGrid Hub sit and its spokes. If you log in to an ExaGrid Spoke Site, your user's role will determine what you can manage on the spoke site, and what you can view on the spoke site and its associated hub site.

For details on local users and their roles see "User Roles" on page 24.

Networking

Depending upon the model, ExaGrid Servers come with two, four or six network interface connectors (NICs). By default, all NICs are 1 Gb, but 10 Gb and 40 Gb NIC options are available.

Network ports are located in the back of your ExaGrid Server. The actual location varies by ExaGrid Model. NIC ports are used for backup and administration and are labeled, the IPMI port is a standard 1GbE port without a label.

The following is an example of the back of an ExaGrid Server - note the labeled NIC ports and the unlabeled IPMI port.

Network Configuration Rules

- At least one network interface must allow ExaGrid traffic from All Sites.This rule applies to single-site systems as well as multi-site systems.
- All interfaces configured for Any Site traffic must use an address that is reachable from all such addresses on all ExaGrid servers in the ExaGrid System - i.e. all hub Sties and all Spoke sites. An error will be returned if an interface set to All Sites cannot logically reach an interface in another site that is set to All Sites. This is a logical evaluation only based on IP networking rules - there is no testing done of actual physical connectivity.
- At least one network interface must allow Backup traffic.

- For multi-site systems, when changing ExaGrid traffic settings for ExaGrid Servers in one ExaGrid Site, wait for the settings to take effect before changing the settings in another ExaGrid Site.
- If you rename an ExaGrid Server:
  - If the server belongs to an Active Directory domain, its name cannot be changed without first leaving the domain. After the name is changed, the server can be added back into the Active Directory domain.
  - Do not rename another server in the site until servers in the site see the newly changed name.

Note – Redefining an ExaGrid Server's IP addresses or name will in most cases require you to also update your backup job configurations.

Network Settings

Note – If you make changes to the Network Settings page, messages will be displayed that evaluate your choices.

The following list details the network settings allowed on the Network Settings page:

- **Name:** The name of each interface for this ExaGrid Server. If the ExaGrid Server is a member of an Active Directory Domain, then its fully qualified domain name (FQDN) must be used.

  An ExaGrid Server, depending on the model, will have two or more physical interfaces. Interfaces are typically used directly, but can be bonded together into virtual interfaces.

- **Type:** The speed of the interface. Either 1Gb or optional 10Gb or 40Gb (depending on ExaGrid model).
- **Link:** Indicates the links connection status: A check mark indicates the connection status is OK. An X indicates the connection status is disconnected. A question mark indicates the connection status is unknown. A triangle indicates that the link status is degraded.
- **VLAN**: Allows tagging an IP address with the specified VLAN ID. Use the **+** link under **Traffic** to add an IP address and/or VLAN ID. When adding a VLAN ID to IP addresses used for ExaGrid traffic (Local Site or All Sites), ensure that all other ExaGrid Servers in the Local Site or All Sites are accessible over the VLAN with that ID.
- **Enabled:** A check mark indicates the interface is enabled.

- **IP MTU:** Maximum size of IP packets sent on the interface. In general, the maximum Ethernet packet size allowed on networking equipment connected to this interface must be larger than this number - refer to your networking equipment documentation
- **DHCP:** Whether DHCP is enabled.
- **IP Address:** The IP address assigned to the interface.
- **Mask:** The network mask applied to the interface.
- **Gateway:** The network gateway used by interface.
- **Traffic, ExaGrid:** Default setting is ***All Sites*** - All ExaGrid Servers in this site and others may use this interface to replicate and balance data and communicate between each other. All addresses configured for Any Site traffic must be reachable from all such addresses on all ExaGrid servers in the ExaGrid System. An error will be returned if an interface set to All Sites cannot reach an interface in another site that is set to All Sites. In most cases All Sites will be the setting used for the interface. The other options are used only when fine tuning your ExaGrid Server is required:

  Local Sites only ExaGrid Servers in site may use this interface to balance data and communicate with other ExaGrid Servers in this site. You can use this setting to add additional bandwidth in a site without having to provision route-able network ports to your switch or when some interfaces are on a *site.local* private network.

  None: Disallows any ExaGrid traffic from flowing through this interface.

- **Traffic, Backup:** Default is disabled (unchecked). This setting allows backup traffic to flow through its associated interface.

Caution – If a backup job is pointed at an interface whose Backup Traffic setting is disabled, the backup job will fail.

- **+** Allows you to apply multiple address definitions to a single network interface.
- **ExaGrid Server Name:** The name of the ExaGrid Server. Unqualified names may be up to 55 character long. Fully qualified names may be up to 63 characters long.
- **Default Gateway:** Static or DHCP
- **DNS Servers:** Name of DNS server. You may specify one or two DNS server names
- **DNS Search Path:** A space-separated list of domains. When looking up the address for unqualified host names on this server. Each domain will be appended to the name until one resolves.

Bonding

Network interfaces (NICs) may be bonded. Doing so provides a fall-back should one NIC fail. A bond requires at least two NICs. Bonded NICs must be of like type. For example, you cannot bond a 1GbE NIC with a 10 GbE NIC.

Caution – Bonding NICs does not automatically increase network bandwidth between the ExaGrid Server and the backup application(s). Work with your ExaGrid Customer Support engineer on the proper switch port configurations and backup configurations required to produce the concurrent backup streams and proper network paths that result in increased network bandwidth.

Ports Used Between ExaGrid Sites

To successfully join and replicate between ExaGrid Sites, the following protocol/ports must be opened in any firewall between the sites:

Port	Protocol	Purpose
22	TCP	SSH (including rsync)
88	TCP & UDP	krb5kdc (Kerberos)
123	UDP	NTP
464	TCP & UDP	kadmind (Kerberos)
749	TCP	kadmind (Kerberos)
5433	TCP	PostgreSQL
5434	TCP	PostgreSQL
5435	TCP	PostgreSQL
7201	TCP	Topology Service
8000	TCP	Management Service
28700 – 28799	TCP	ExaGrid Services
32768 – 61000	TCP	Replication Service (Ephemeral Ports)

TLS Certificates

TLS certificates eliminate the browser security warnings associated with pages in the ExaGrid browser-based user interface that require an HTTPS connection.

TLS certificates used with an ExaGrid Server must:

- Have a common host name and a DNS name that matches the DNS name in the browser. IPs are not allowed.
- The browser must know the Certification Authority that signed the certificate.

TLS certificates are applied to ExaGrid Servers from an ExaGrid Site's Security menu's TLS Certificates page. Only users with the Admin role or Support role can apply certificates.

From the TLS Certificates page you can:

- Keep the ExaGrid default certificate.
- Upload a Personal Information Exchange File.
- Configure an SCEP Server to automatically update and serve up certificates.
- Revert to the ExaGrid default certificate. This option is only available if TLS certificates are in use.

Caution – When reverting back to the ExaGrid default option, you MUST ensure that HSTS is disabled.

ExaGrid recommends that each ExaGrid Server in an ExaGrid Site use the same TLS settings.

When a certificate approaches its expiration date, the ExaGrid will generate an alarm to remind you to apply an updated certificate.

Applying TLS Certificates

Note – When you change TLS settings, the ExaGrid web server will restart. During that brief interval, connection to the ExaGrid will be lost.

TLS settings are applied, one at a time, to each ExaGrid Server in an ExaGrid Site.

To apply a TLS certificate to an ExaGrid Server:

1. From the ExaGrid left-hand navigation frame, select the ExaGrid Site in which the ExaGrid Server resides.
2. In the top, main menu select **Site** and all menu options for the site are displayed.
3. Under **Security**, select **TLS Certificates** and the TLS Certificates page is displayed

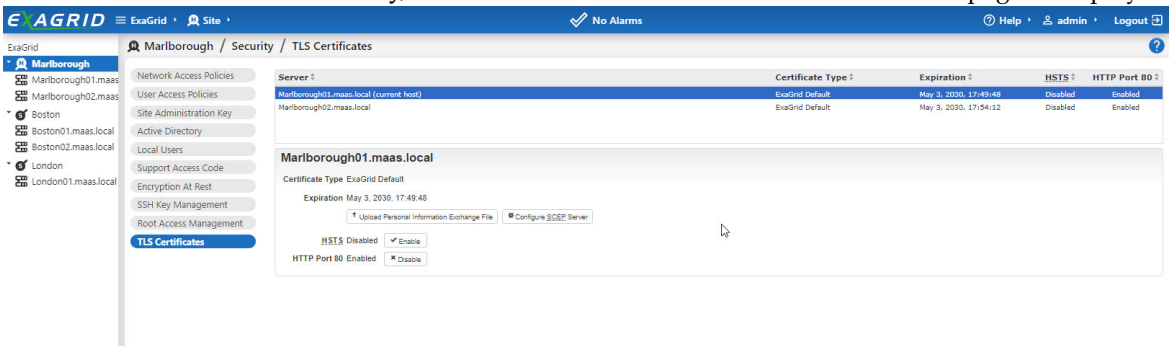

4. From the list of ExaGrid Servers, select the ExaGrid Server to which the certificate will be applied.
5. Select the certificate to be applied:

a. If you select **Upload a Personal Information Exchange File**, the Personal Information Exchange File dialog box is displayed:

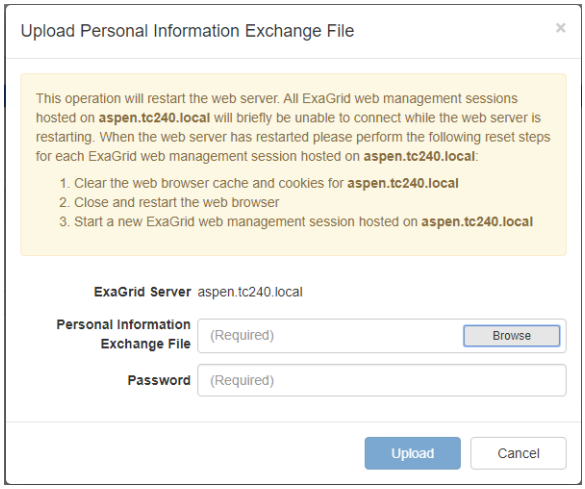

- Enter the path, or browse to the Personal Information Exchange File.
- Enter the Password associated with the Personal Information Exchange File.
- Click **Upload** and the Personal Information Exchange File will be uploaded.
- Once you have confirmed that the certificate is good and the browser warnings in the ExaGrid GUI pages that require HTTPS are gone, enable/disable Port 80 and HSTS according to your corporate best practices.

Note – If you enable HSTS and the applied certificate is bad, contact ExaGrid Support to resolve issues you may encounter.

- Click **Apply** and the selected certificated will be used on the selected ExaGrid Server.

b. If you select **Configure an SCEP Server**, the Configure SCEP Server dialog box is displayed:

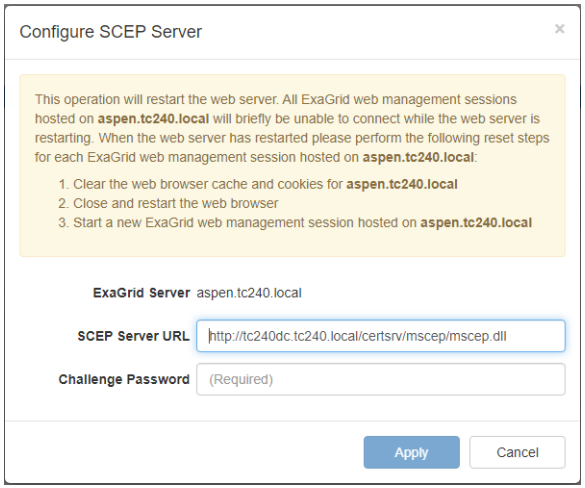

- Enter the **SCEP Server URL.**
- Enter the **Challenge Password** associated with the SCEP server.
- Once you have confirmed that the certificate is good and the browser warnings in the ExaGrid GUI pages that require HTTPS are gone, enable/disable Port 80 and HSTS according to your corporate best practices.

Note – If you enable HSTS and the applied certificate is bad, contact ExaGrid Support to resolve issues you may encounter.

- Click **Apply** and the SCEP Server URL will be applied.

c. If you have an existing certificate you can select **Revert to the ExaGrid default certificate** and the Confirm Revert to ExaGrid Default Certificate dialog box is displayed:

Caution – When reverting back to the ExaGrid default option, you MUST first Disable HSTS.

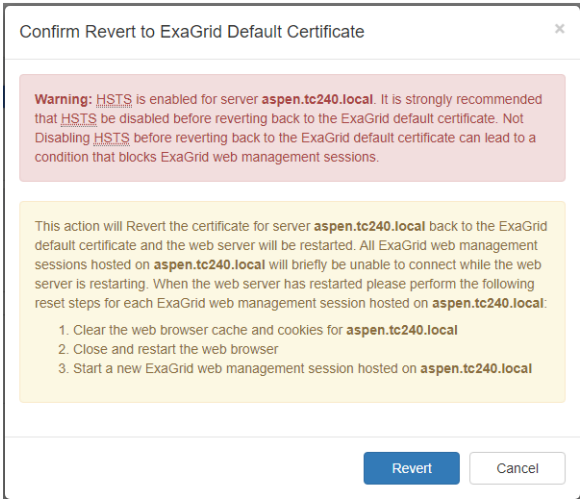

- Click **Apply** and the certificate settings will revert to the ExaGrid default certificate.
- Enable/disable Port 80 according to your corporate best practices.

Note – If HSTS is enabled and you apply the default ExaGrid certificate, contact ExaGrid Support to resolve issues you may encounter.

Once the TLS setting have been applied, the ExaGrid Server's web server will restart. During that brief interval, connection to the ExaGrid Server will be lost.

When the web server has restarted, you must:

1. Clear your web browser's cache and cookies.
2. Log out of the ExaGrid Server.
3. Close your web browsers.
4. Open a new browser session and log back into the ExaGrid Server.

Managing and Securing Access

Access to the ExaGrid user interface, its features, and backed up data is controlled in several ways:

- The ExaGrid resides in your data center and your access protocols limit physical access to the ExaGrid and its data.
- Time-locks are set on deduplicated and retained data. This protects against ransomware and other malicious activity. For details see “Retention Time-Lock Overview” on page 20.
- Two-Factor Authentication can be enabled on a user’s account. For details see “Two-Factor Authentication” on page 27.
- Users can be assigned roles that limit their access to a specific set of ExaGrid functions. For details see “User Roles” on page 24.
- ExaGrid Servers can be members of an Active Directory Domain. ExaGrid roles can be mapped to users or groups in their server’s Active Directory Domain. For details see “ Managing Active Directory” on page 29.
- User Access Policies define the backup users who have access to a share. For details see “User Access Policies” on page 45.
- Network Access Policies that define network access to a share. For details see “Network Access Policies” on page 44.
- Hardware options that employ Encryption at Rest and Boot Protection. For details see “Encryption Capable ExaGrid Servers” on page 73.

At the top of an ExaGrid Site’s Summary page are drop down lists that detail whether the site’s security settings meet ExaGrid’s recommended best practices. Use the links in the lists to set the site’s security settings:

Caution – ExaGrid highly recommends reviewing these security settings on a regular schedule, including ExaGrid’s Retention Time-Lock settings. Doing so ensures that the maximum security settings are applied to your backed up data.

Retention Time-Lock Overview

ExaGrid's Tiered Backup Storage architecture has:

- A high-speed, disk cache that makes up ExaGrid's front-end Landing Zone Tier where the latest backup data resides.
- A separate Retention Tier that contains all deduplicated and retained data.

This proven tiered architecture lends itself to a unique, time-lock ransomware solution where:

> Data is written directly to the network facing ExaGrid Landing Zone. The backed up data is then kept in a "non-network facing" long-term retention tier. As data is put into the Retention Tier, it is deduplicated and stored in a series of objects and metadata. As with other object storage systems, the ExaGrid objects and metadata cannot be modified. This ensures that only new objects are created and only old objects are deleted when their defined retention date is reached.

Because the Retention Tier is not network facing and not accessible to ransomware and other malicious activity, a time-lock can be set on the ExaGrid Retention Tier that blocks deletes for up to 30 days.

This combination of a non-network facing Retention Tier, a definable period of time where delete operations are not performed, and objects that never change, protect ExaGrid's retained data from ransomware and other malicious activity.

If the data in the ExaGrid Landing Zone is deleted or encrypted, the data in the Retention Tier is not affected. As soon as an attack is identified, with a single click you can suspend (lock down) all shares in an ExaGrid Site and begin restoring backup data from the ExaGrid Retention Tier to your primary storage.

Key features of ExaGrid's Time-Lock Retention solution include:

- A Retention Time-Lock policy that can delay delete commands from 0 to 30 days.
- Required security officer approval for changes to time-lock settings.
- Compromised data backed up to the ExaGrid Landing Zone Tier is deduplicated and sent to the Repository Tier but does not delete or change previous backups in the Repository Tier.
- Two-factor Authentication protects against unauthorized access and changes to Retention Time-Lock setting. For details see "Two-Factor Authentication" on page 27.
- Optionally, separate passwords for Administrators on Primary and Secondary (DR) ExaGrid Sites.
- The Retention Tier usually only needs 2% extra capacity per week for the Time-Lock. For example, setting the Time-Lock to 14 days typically will require only 4% of additional Retention Tier capacity.

For details on how to recover from a ransomware or malicious software attack, see "Recovering From Malicious Software Attacks, or a Loss of ExaGrid Servers" on page 83.

Modifying Retention Time-Locks

By default, ExaGrid's Retention Time-Locks are set to 10 days. For most customers the default setting will provide more than enough time to detect and then recover from a ransomware or malicious software attack.

Retention Time-Locks can be modified by any user assigned the Admin role. Retention Time-Locks can be manually set from 0 (not recommended) to 30 days.

Changes to Retention Time-Locks require the approval of an ExaGrid User assigned the Security Officer role. Current time-locks stay in effect until the Security Officer approves a change. Before modifying Retention Time-Locks, ensure a user with the Security Officer role exists in the site and they are aware that a modification request is coming.

To modify Retention Time-Locks on shares in an ExaGrid Site:

1. Log into the ExaGrid Site as a user who is assigned the Admin role, and the site's summary page is displayed:

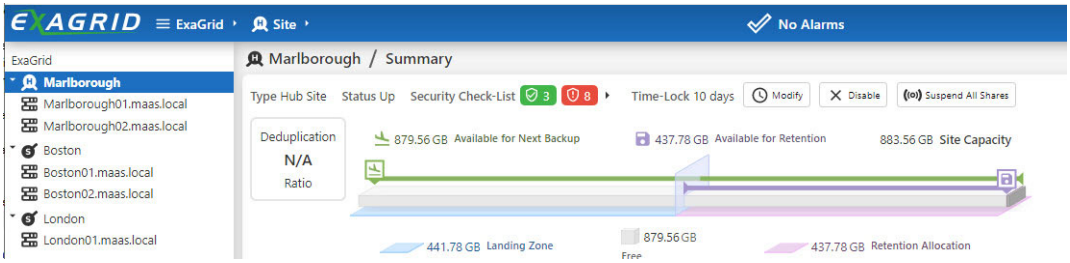

2. At the top of the page, next to the current time-lock setting, click **Modify** and the Modify Time-Lock dialog box is displayed:

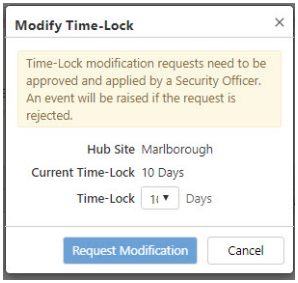

3. From the drop down menu select the new time-lock period.

4. Click **Request Modification** and the following message will appear at the top of the Security Officer's default login page:

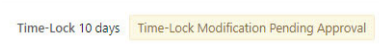

Approving and Rejecting Retention Time-Lock Modifications

It is the responsibility an ExaGrid Site's Local or Active Directory user assigned the Security Officer role to accept or reject modifications to a site's Retention Time-Locks.

To approve or reject a Retention Time-Lock modification request:

1. Log into the ExaGrid Site as a user who is assigned the Security Officer role, and the site's summary page is displayed:

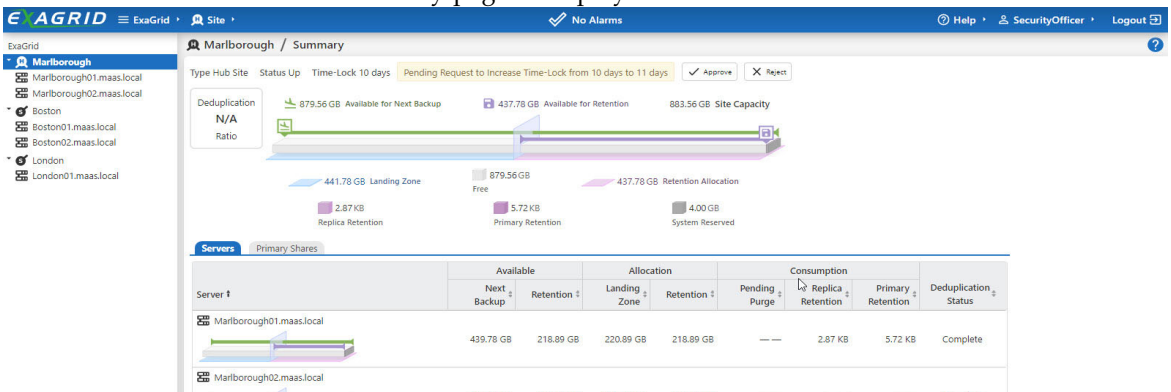

2. At the top of the page, next to message that details the Time-Lock change request, click either **Approve** or **Reject**.

3. If this is a request to increase the Time-Lock period: clicking Accept or Reject displays a simple confirmation dialog box:

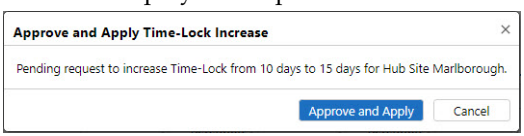

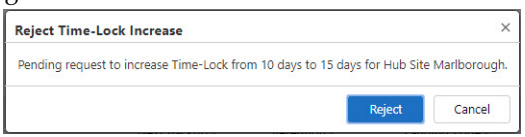

4. Click either **Approve and Apply** or **Reject** and the site's Retention Time-Lock is set accordingly.

5. If this is a request to decrease the Time-Lock period, if you click **Accept**, a dialog box explaining the risks involved in decreasing the Time-Lock period is displayed:

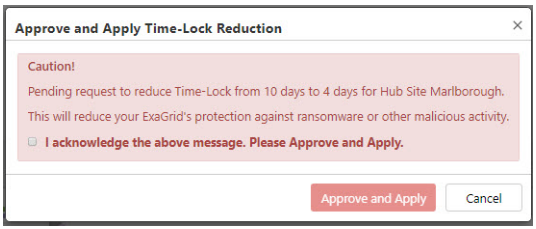

6. Because of the increased risks involved, you must click the **I Acknowledge...** check box before you click **Approve and Apply**.
7. Click **Approve and Apply** and the Time-Lock setting for the site is changed.
8. If this is a request to decrease the Time-Lock period, if you click **Reject**, a simple confirmation dialog box is displayed:

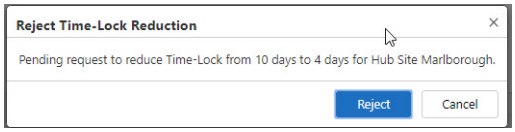

9. Click **Reject** and the Retention Time-Lock period will not be changed.

User Roles

Roles are assigned during the ExaGrid user creation process and determine what ExaGrid features and capabilities a user can access. See “Creating or Modifying ExaGrid Local Users” on page 25 for details.

Each ExaGrid System comes with one predefined ExaGrid user named Admin that is assigned the Admin role. The Admin role grants complete access to the ExaGrid user interface and can perform all management functions including creating new users.

Note – The Admin user role **CANNOT** approve or deny Retention Time-Lock change requests. Approving and denying such requests is the sole responsibility of the Security Officer role. See “Retention Time-Lock Overview” on page 20.

Domain users, and domain groups can be to mapped an ExaGrid role. See “ Managing Active Directory” on page 29 for details.

The following list defines the roles that can be assigned or mapped:

- Admin: Grants full access to all ExaGrid features and can perform all management functions including creation of new users and power management. **As a best practice, only a limited number of users should be assigned this role.**
- Operator: Grants access to typical day-to-day operations such as share creation (but cannot delete shares) without granting operating-system-level access to the entire ExaGrid System. **As a best practice, most users should be assigned this role.** Not assigning this role to at least one user will raise an alarm in the ExaGrid Site’s security check list.
- Security Officer: Grants approvals or denials to changes to Retention Time-Lock settings. **As a best practice, only one user should be assigned this role.** Not assigning this role to at least one user will raise an alarm in the ExaGrid Site’s security check list. See “Retention Time-Lock Overview” on page 20.
- Viewer: Grants access to ExaGrid reports and status information. Users with the Viewer role cannot perform any day-to-day or system-level operations.
- UPS Operator: Grants scripts access to ExaGrid’s power management features.
- Backup Only: Grants (backup applications) access to ExaGrid Shares for backup operations. This role does not give access to ExaGrid management functions.

Caution – Attempting to access ExaGrid CIFS shares using the **guest** account from certain versions of Windows will not succeed if your Windows username (domain) matches an ExaGrid user. To access an ExaGrid CIFS share in this case, use your Windows username rather than “guest”.

User Role Idle-Timeout Settings

To help prevent unauthorized use of an ExaGrid user account, idle-timeouts are set on user roles. Roles with the greatest access have the shortest idle-timeout periods:

- Security officer, Admin, and Support role time out after fifteen minutes
- Operator role times out after one hour
- Viewer role times out after twelve hours

These timeout settings are not configurable.

Creating or Modifying ExaGrid Local Users

To create or modify a new ExaGrid local user:

1. Ensure that you are logged in to the ExaGrid as a user assigned the Admin role.
2. In the navigation tree, click the ExaGrid Site.
3. From the top, main menu, click **Site** and all menu options for the site are displayed.
4. Under **Security**, select **Local Users** and the Manage Local Users page is displayed.
5. If you are modifying an existing user, select the user then click **Modify** and the Modify Local User dialog box is displayed.
6. If you are creating a new user, click **New** and the Create Local User dialog box is displayed:

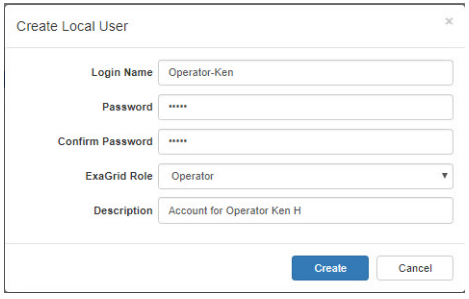

7. In the spaces provided enter a user name and password. See “Valid ExaGrid Local User Name Character Set” on page 26 and “Valid ExaGrid Local User Password Character Set” on page 26.
8. From the drop down menu select the required role.

9. In the space provided, enter a description of the user.
10. Click **Create** or **Apply**, and the local user is created or modified.

Valid Local User Names

When creating a new local user for a site, there are a number of names reserved for internal Linux accounts.

If you try to create a local user with one of these reserved names you will encounter an error stating that the account already exists. For example, the user name **operator** with a **lowercase "o"** is a reserved Linux account name and cannot be used for an ExaGrid Local User.

Valid ExaGrid Local User Name Character Set

Local user names must:

- Not contain more than 32 characters
- Begin with an upper or lower case letter A-Z or a-z

For the rest of the name, you can use any of the following:

- Upper/lower case A-Z
- Numeric 0-9
- . - _ (period, hyphen, underscore)

Valid ExaGrid Local User Password Character Set

Local user passwords cannot contain spaces but can contain any printable, 7-bit ASCII character.

Two-Factor Authentication

As a best practice, local users and active directory users assigned the following roles should set ExaGrid's two-factor authentication on their account.

- Admin
- Operator
- Security Officer
- Viewer

The following roles cannot use two-factor authentication:

- The Backup Only role is meant for backup applications writing to ExaGrid Shares.
- The UPS Operator role is meant for power control scripts.

Setting Two-Factor Authentication

Before two-factor authentication can be set, an ExaGrid user must first install a third party two-factor authentication application on the device (cell phone, lap top, etc.) that will be used to authenticate their log in.

To set Two-Factor Authentication:

1. In the upper right-hand corner of the ExaGrid user interface, find the user account name you are currently logged in to. For example, in the following screen capture we are using the user named **admin.**
2. Click on the user account name.

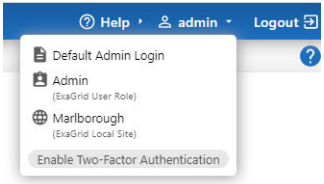

3. From the drop down menu, select **Enable Two-Factor Authentication** and the Two-Factor Authentication dialog box opens:

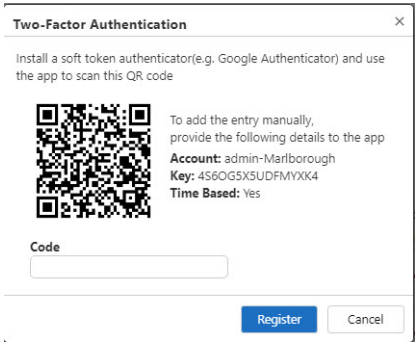

4. If supported by your authentication application, scan the QR Code. For other applications you can manually enter the **Key** provided in the dialog box.
5. Your authentication application will generate a code on the device you use for authentication.
6. Enter the code in the Two-Factor Authentication dialog box's **Code** text box.
7. Click **Register**, and a list of recovery codes is displayed. These codes can be used (once each) to recover from the loss of the soft token.
8. Click **Download** and the list of recovery codes can be downloaded for safe keeping.
9. Click **Close** and two-factor authentication is enabled on the user account.

Going forward, use the two-factor authentication application on your device of choice to obtain a code. This code must be entered into the ExaGrid Challenge Code screen. For example:

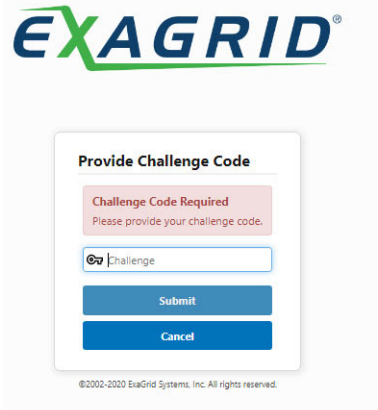

Enter the code in the space provided then click **Submit**.

Managing Active Directory

Active Directory membership is defined at the ExaGrid Site level. All ExaGrid Servers in a multi-server ExaGrid Site must belong to the same Active Directory Domain. All ExaGrid Sites in a multi-site ExaGrid System do not have to belong to the same Active Directory Domain.

Note – Active Directory trusts can be used to give a user access to an ExaGrid site in a different domain.

Once the ExaGrid Servers in an ExaGrid Site are added to a domain, users and groups in that domain can be mapped into ExaGrid roles. When a user or group is mapped into a role, they will have access to all ExaGrid features associated with that role.

Active directory users and groups with mapped roles can be added to a share's User Access Policy. See "Share Access Policies" on page 32 for details.

ExaGrid User Interface Access Using Active Directory Credentials

Microsoft supports a wide range of characters in Active Directory usernames and passwords. However ExaGrid's user interface login is currently limited to:

Active Directory usernames and passwords that contain any printable, 7-bit ASCII character including spaces. See"Mapping an Active Directory User or Group to an ExaGrid Role" on page 31.

Adding ExaGrid Servers to An Active Directory Domain

To add ExaGrid Servers to an Active Directory Domain:

Note – You must be logged in through Hyper Text Transfer Secure (**HTTPS**) Protocol otherwise the Manage Active Directory page will be view only.

---

1. Log into a user account that has the Admin role.

2. In the left-hand navigation tree, select the ExaGrid Site that contains the ExaGrid Servers to be added to the Active Directory Domain.
3. From the top, main menu select **Site** and all menu options for the site are displayed.
4. Under **Security**, select **Active Directory** and the Server Domain Membership tab for Active Directory is displayed:

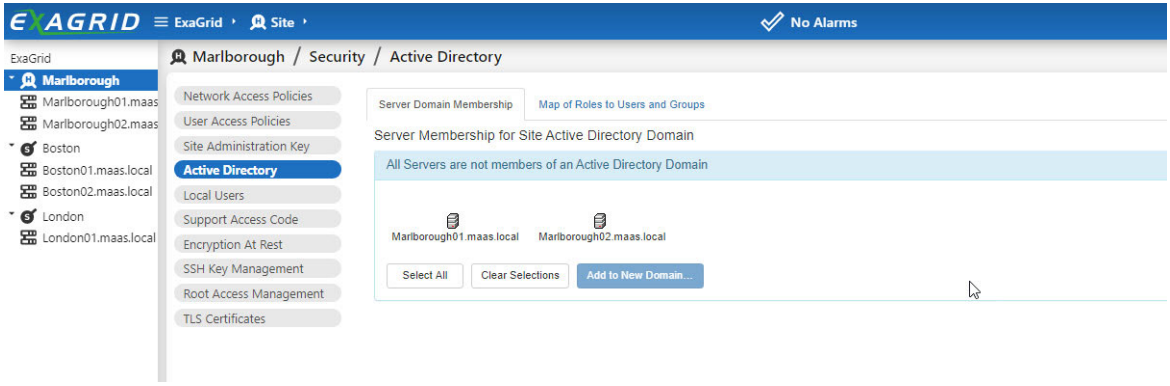

5. Click **Select All.**
6. Click **Add to New Domain** and the Add to Domain dialog box is displayed:

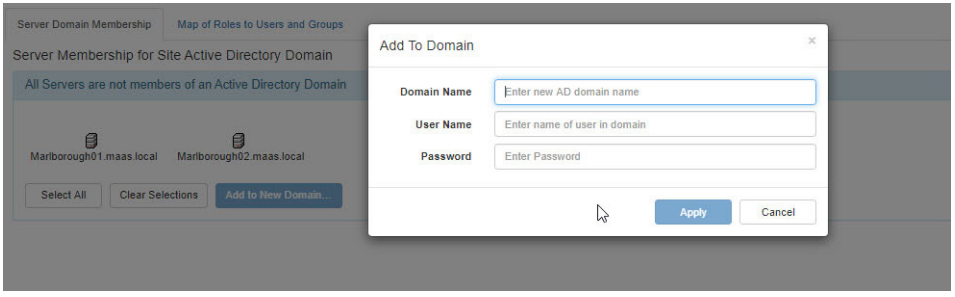

7. In the spaces provided enter the Domain Name, the domain User Name, and the domain Password for the specified user. The User and Password must have sufficient permissions to "Join a computer to a domain" or the operation will fail.
8. Click **Apply** and the selected ExaGrid Servers in the ExaGrid Site are added to the specified domain.

Mapping an Active Directory User or Group to an ExaGrid Role

Note – If a suitable Domain User or Group containing Users that need access to the ExaGrid does not exist, it must first be created using standard Active Directory management tools.

To map an Active Directory User or Group to an ExaGrid role:

1. Log into a user account that has the Admin role.
2. In the left-hand navigation tree, select the ExaGrid Site that contains the ExaGrid role to be mapped into an Active Directory Group or user.
3. From the top, main menu click **Site** and all menu options for the site are displayed.
4. Under **Security**, select **Active Directory** and the Server Domain Membership tab for the Manage Active Directory is displayed.
5. Click the **Map of Roles to Users and Groups** tab and a table is displayed that lists how roles are mapped into users or groups is displayed:

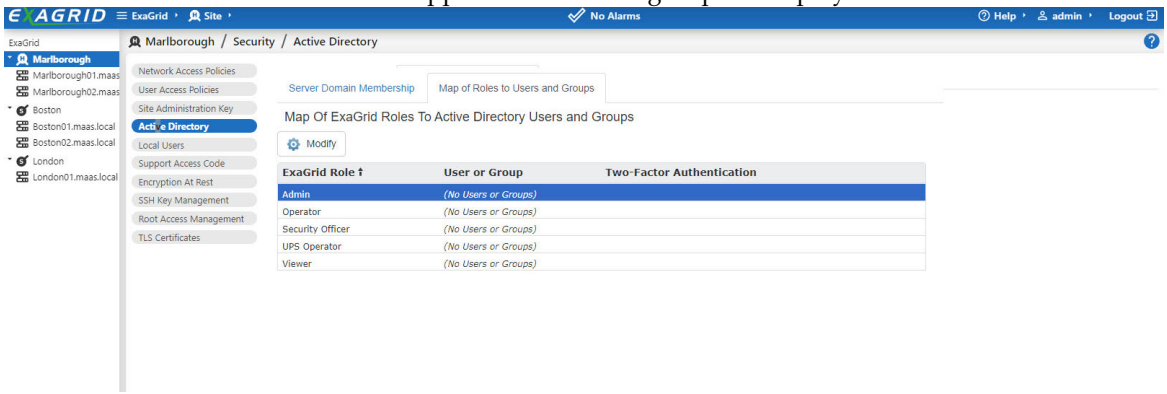

6. Select the role to be mapped into a user or group.

7. Click **Modify** and the Map Users and Groups to Role dialog box is displayed:

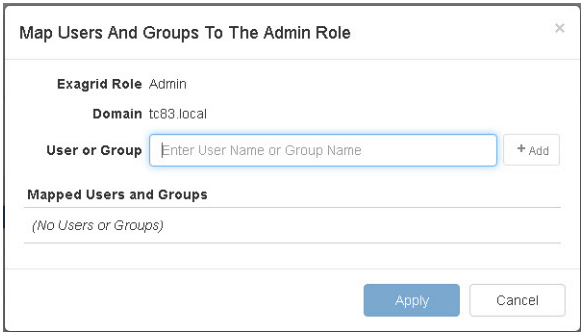

8. In the space provided, enter the user or group name to map to the selected role.
9. Click **Apply** and the user or group is mapped to the ExaGrid role. The specified user or any user in the group can access the ExaGrid with this role using their Domain credentials.

Share Access Policies

Policies that define share access are created at the ExaGrid Site level. Depending on the share type, you must create at least one:

- User Access Policy that defines the backup application users who have access to a share.
- Network Access Policy that defines network access to a share.

Note – Veeam Shares do not require a Network Access Policy.

A single User and Network Access Policy can be used by multiple shares.

For details see “Share Access Policies” on page 44.

Managing SSH Keys

SSH keys are used by:

- ExaGrid Support logging in as the **root** user.
- Veeam agents logging in as a backup user to access ExaGrid-Veeam data mover shares.

To manage SSH Keys:

1. Log in to the ExaGrid system as a user assigned the Admin role.
2. From the navigation tree, select an ExaGrid site.
3. In the to main menu, select **Site**, and all menu options for the site are displayed.
4. Under **Security**, select **SSH Key Management** and the SSH Key Management page is displayed:

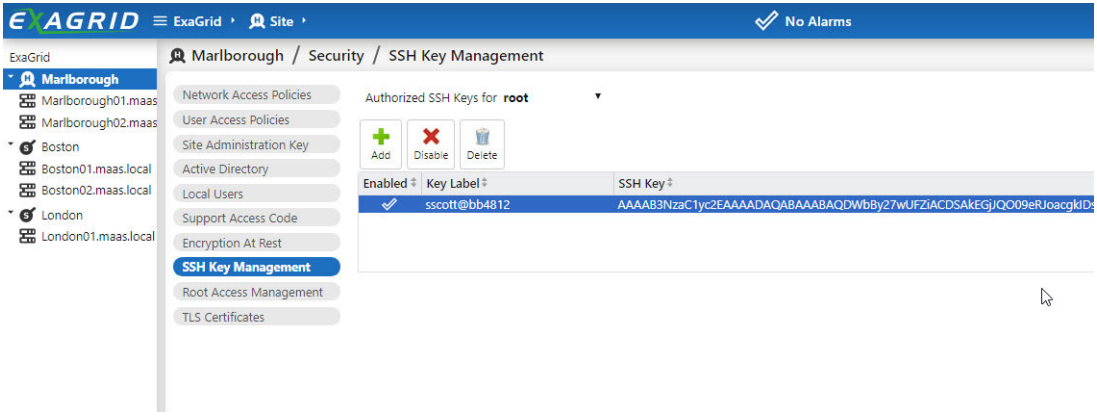

5. From the **Authorize SSH Keys For:** drop down menu, select a user.
6. Click **Add** and the Add SSH Key dialog box for the selected user is displayed:.

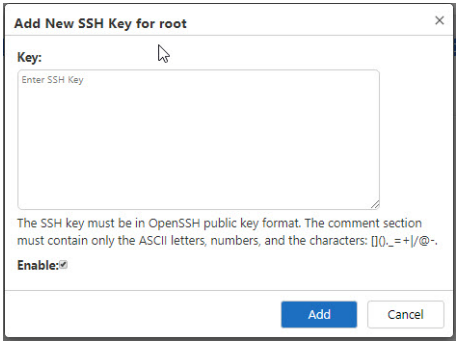

7. In the space provided, enter or paste the SSH Key.
8. Click **Add** and the key is added for the user.

Support Accounts

The ExaGrid support account is a role-based account created at installation with a user name of **support** and a password set to a unique code called the Support Access Code.

Note – The ExaGrid Support role cannot be mapped to an Active Directory Group.

The **support** user password can be changed from either the:

- Local User page - When you manually change the **support** user's password on the Local User's page the **root** user's password will also be changed. See "Creating or Modifying ExaGrid Local Users" on page 25.
- Support Access Code page - When you generate a new access code, the generated code becomes the new password for both the **support** user AND for the **root** user. This is typically used to give ExaGrid Support representatives temporary root access to ExaGrid servers. Once the support representative has completed their tasks, a new password can be generated and they will no longer have access to root.

To generate a new support access code:

1. Log in to the ExaGrid system as a user assigned the Admin role.
2. From the navigation tree, select an ExaGrid site.
3. In the top, main menu, click **Site**, and all menu options for the site are displayed.
4. Under **Security**, select **Support Access Code** and the Support Access Code page is displayed

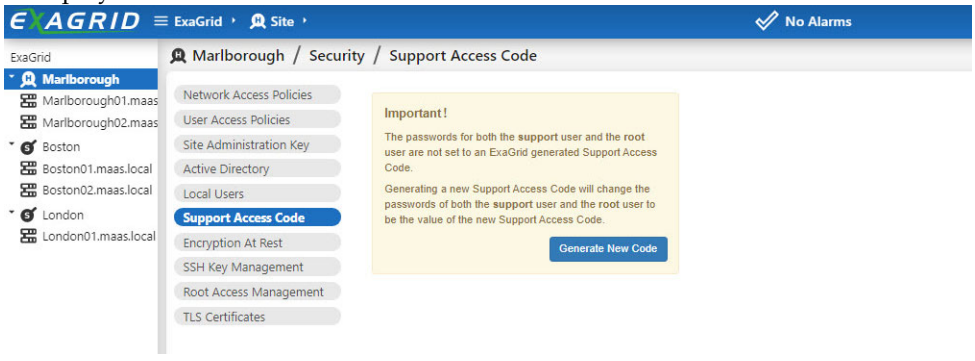

5. Click **Generate New Code** and a new **support** user password and **root** user password will be created:

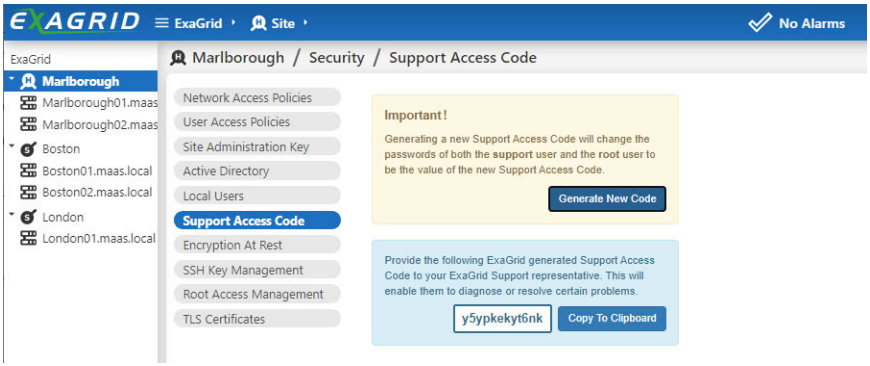

6. You can copy the new code to the clipboard and paste it in to a message or file as needed.

Root Access Management

You can use a combination of passwords and SSH keys to control access to the **root** user account for ExaGrid Servers in a site. Three options are provided for managing root access:

- A password or SSH key
- A password for one day that reverts to **SSH Key only** on time out
- Authorized SSH keys only

To manage **root** user access:

1. Log in to the ExaGrid system as a user assigned the Admin role.
2. From the navigation tree, select an ExaGrid site.
3. In the top, main menu, click **Site**, and all menu options for the site are displayed.
4. Under **Security**, select **Root Access Management** and the Root Access Management page is displayed.

5. Click **Configure** and the Configure Root User Login dialog box is displayed:

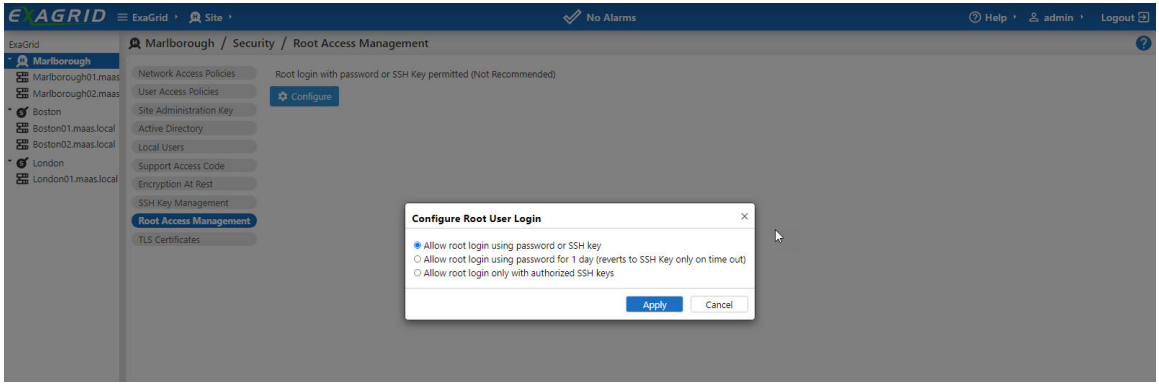

6. Make your selection as needed.
7. Click **Apply** and **root** user access will be configured.

Using Server Message Block (SMB) Signing

SMB Signing is a feature that allows communications to be digitally signed at the packet level. Digitally signing the packets enables the recipient of the packets to confirm their point of origination and their authenticity. This security mechanism in the SMB protocol helps avoid issues like tampering of packets and "man in the middle" attacks. SMB signing does not encrypt data.

In addition, SMB signing augments the ExaGrid Share Network Access Policy and User Access Policy to control who and what has access to data on ExaGrid Shares.

To use SMB signing with your ExaGrid System you must enable SMB signing on your Microsoft® Windows® backup servers through either:

- Enabling the Local Security Policy for the Microsoft network client: Digitally signed communications (always).
- Setting up a security policy that includes the enabled Microsoft network client: Digitally signed communications (always) on a Windows domain through the use of Group Policy Objects (GPO), and applying it to the backup server.

To ensure these changes are fully in place, restart your backup server.

Caution – Backups must not be running while you are enabling SMB signing.

Caution – Once SMB signing is enabled, you must specify credentials on your backup application and on the ExaGrid System. Not doing so will cause your backups to fail.

Determine which Windows account (NOT the local system account) or domain account credentials you will specify in your backup application and on the ExaGrid System.

If the ExaGrid Site is not part of an Active Directory Domain, enable and define local security credentials on your ExaGrid System:

1. In the left-hand navigation tree, click the ExaGrid Site name.
2. From the top, main menu, click **Site**, and all menu options for the site are displayed.
3. Under **Security**, select **Local Users.**
4. Create a local user with the **Backup Only Role.**

   The local user credentials you specify on the ExaGrid System must match an existing local or domain user account. Domain accounts must be entered without the domain information.
5. Create a **User Access Policy** specifying the local user (when the ExaGrid Site is not added to an Active Directory Domain) or the Active Directory Domain user(s) or group(s).
6. Assign the User Access Policy in the ExaGrid share configuration.

ExaGrid Shares

ExaGrid Shares are mount points to which your backup application backs up data.

There are two basic types of ExaGrid Shares:

- **Primary Shares:** Mount points created on ExaGrid Servers to which your backup application can write.
- **InstantDR Shares:** Replicas of Primary Shares.

Primary Shares

There are three types of primary shares:

1. Backup Application Specific Shares
2. Universal Backup Shares
3. Utility Shares

Backup Application Specific Shares

Application specific shares are tailored to specific backup applications. ExaGrid provides detailed step-by-step instructions on how to use your backup application with these shares:

1. In the ExaGrid top, main menu, click **Help**.
2. From the drop down menu, select the **Online Library**. The Online library is displayed.
3. Expand the option that applies to your backup application.
4. Scroll down until you see the guide that applies to your application.

5. Right mouse click on the link to the guide and save or print it as needed.

To create an application specific share, see “Creating Shares” on page 47.

Universal Backup Shares

ExaGrid Universal Backup Shares are designed to accommodate those backup applications that do not have an ExaGrid Share designed specifically for their use.

Universal Backup Shares are deduplicated and can be replicated and migrated like application specific ExaGrid Shares.

ExaGrid Systems has qualified a number of backup applications for use with the Universal backup Share. Best practices for the qualified applications can be found in your ExaGrid System’s Online Library.

In addition, ExaGrid is continually qualifying backup applications for the Universal Backup Share. The ExaGrid web site (www.exagrid.com) will have the most up-to-date list of applications qualified for use with the Universal Backup Share.

In most cases, if your backup application has not yet been qualified by ExaGrid for use with the Universal Backup Share; the application will work with the Universal Backup Share provided you follow these simple guidelines:

Note – ExaGrid Support may have limited expertise in applications not qualified by ExaGrid Systems

---

1. Create an ExaGrid Universal Backup Share. See “Creating Shares” on page 47.
2. Write backups to the ExaGrid Universal Backup Share as recommended by the application’s best practices for writing to disk,
3. In addition to following your backup application’s best practices, if supported by your backup application, ensure that:
    - Compression is **OFF**
    - Encryption is **OFF**
    - Multiplexing is **OFF**
    - Deduplication is **OFF**
    - Concurrent write operations should not exceed 20. ExaGrid recommends to initially limit concurrent writes to four or five. Over time, increase concurrent writes until you reach a number that maximizes the performance of your backups. Each backup application, due to its particular pattern it writes will experience different levels of performance.
4. Run a series of test backups and restores to monitor deduplication ratios and performance and make adjustments as needed.When testing your backups look for:

- Errors or warnings in the ExaGrid Event Viewer
- Unusual deduplication ratios
- Unusual deduplication rates
- A sudden increase in backup durations when the retention period for your first backups is reached.
- Restores from backups at the other edge of your retention period that take an excessive amount of time.

Note – Your ExaGrid System does not employ real-time reporting. When monitoring backup results it can take up to fifteen (15) minutes for a report to be updated.

The Universal Backup Share is not be suitable for all backup applications. This includes backup applications that:

- Create file and folder names that do NOT use UTF-8 encoding.
- Archive data into many small, individual files
- Backup data into many small, individual files
- Overwrite large amounts of backup data sent previously to the ExaGrid Universal Backup Share. Overwriting a small amount of "metadata" – e.g. control files, catalogs, etc. – should perform acceptably.
- Treat the ExaGrid Universal Share as primary storage

Utility Shares

Utility Shares are designed to accommodate any data type. Data can be copied manually, through a script, or from an application to a Utility Share.

Caution – Utility Shares are not protected by ExaGrid's Retention Time-Lock settings and can be vulnerable to ransomware and other malicious activity. See "Retention Time-Lock Overview" on page 20

When you create a Utility Share you must allocate a specific amount of storage space to the Utility Share. The space allocated is taken proportionately from the host ExaGrid Server's Landing Zone Tier and Retention Tier.

- Data in Utility Shares is not deduplicated and can not be replicated to another ExaGrid Site.
- You can only create one Utility Share per ExaGrid Server.
- When you create a Utility Share, the amount of available Landing Zone Tier and Retention Tier reported in Summary Reports will be automatically adjusted.

- No more than twenty percent (20%) of an ExaGrid Server's space can be allocated to a Utility Share.
- The maximum allowable size of a Utility Share is 13 TB

Caution – When allocating space, you must consider the impact of the space allocation on your backup Landing Zone Tier as well as your Retention Tier. It is possible to allocate space for a Utility Share such that it interferes with your ExaGrid Server's ability to accept and retain backups from other backup applications.

To create a Utility Share, see "Creating Shares" on page 47

InstantDR™ Shares

InstantDR Shares are replicas of primary shares that are reside on one or more ExaGrid Site. InstantDR shares can be used to:

- Recover replicated data after a disaster scenario. In multi-site ExaGrid Systems, backup data in ExaGrid Shares that are set to replicate to other sites, can be recovered.
- Test recovery of replicated data. You can use InstantDR shares to prove that the data replicated to other ExaGrid Sites is recoverable.
- Perform remote tape copy operations. Instant DR shares can also be used to make a tape copy of data that has been replicated to other ExaGrid Sites. For details, see the ExaGrid guide that describes how to use your backup application. A copy of your backup application guide can be found in the Online Library.
- Can be used with Veritas' OpenStorage Technology to achieve:
  - Seamless integration with the ExaGrid disk backup system
  - Unified control and tracking of all copies of data from the Veritas application's interface -- including copies replicated to other ExaGrid Sites and tape copies.
  - Simplified disaster recovery, as Backup Exec is already aware of all copies of data and individually tracks retention periods for each copy
  - Fast recovery times due to the combination of Veritas' and ExaGrid InstantDR capability
  - Short backup window, by load balancing across Veritas media servers into the ExaGrid System
  - Flexibility to retain different amounts of backup data at a local site versus an off-site disaster recovery location, by better utilizing the unique ExaGrid architecture.

To see how InstantDR shares are added to replicas, see "Creating Replicas and InstantDR Shares" on page 52.

Share Protection

Shares, and the data they contain, are protected in several ways:

- The ExaGrid System and the data backed up to it are in your data center.
- ExaGrid Servers can be members of an Active Directory Domain. ExaGrid Roles can be assigned to groups in their server's Active Directory Domain, thereby limiting access to an ExaGrid Server and the ExaGrid Shares it hosts. For details see " Managing Active Directory" on page 29
- User Access Policies that define the backup application users who have access to a share. See "User Access Policies" on page 45.
- Network Access Policies that define network access to a share. For details see "Network Access Policies" on page 44
- Multi-site configurations provide an additional level of data security. Copies of backup files can be kept on a second site. In case of a catastrophic event where one site has been damaged or destroyed, copies of backups on the second site can be used to rebuild your backups.
- ExaGrid provides a hardware option that employs Encryption at Rest and Boot Protection. See "Encryption Capable ExaGrid Servers" on page 73 for details.

Share Migration

Over time, your backup requirements will change. Because of these changes, it might be necessary to expand your ExaGrid System and/or add and move shares from one ExaGrid Server to another. To help you balance your backup load, the ExaGrid System allows you to move primary shares from one ExaGrid Server to another. This process is known as Migration.

If it becomes necessary to migrate an ExaGrid Share, contact your ExaGrid Support Representative and they will perform the migration.

Once a share has been migrated, it is treated as a new share. If the share is set to replicate to other sites, its replication history will be restarted and its replication reports will begin anew.

Note – Once a share has been migrated, backup jobs writing to the share need to be reconfigured to write to the share's new location. However, Veritas backup jobs using the ExaGrid OpenStorage (OST) plugin do not need to be reconfigured.

See the "*Using .... With an ExaGrid System*" guide for the specific backup application for more details.

Share Access Policies

Policies that define share access are created at the ExaGrid Site level. Depending on the share type, you must create at least one:

- User Access Policy that defines the backup application users who have access to a share.
- Network Access Policy that defines network access to a share.

Note – Veeam Shares do not require a Network Access Policy.

A single User and Network Access Policy can be used by multiple shares.

Network Access Policies

Network Access Policies provide the necessary networking information (IP address or name) for a backup application's server to access a share. A single Network Access Policy can contain multiple addresses and be used by multiple shares.

Modifying a Network Access Policy will change access to all the shares that use the policy.

To create a Network Access Policy:

1. In the left-hand navigation tree, select the ExaGrid Site on which the you want to create the Network Access Policy.
2. In the top, main menu, click **Site** and all menu options for the site are displayed.

3. Under **Security,** select **Network Access Policy** and the Network Access Policy dialog box is displayed:

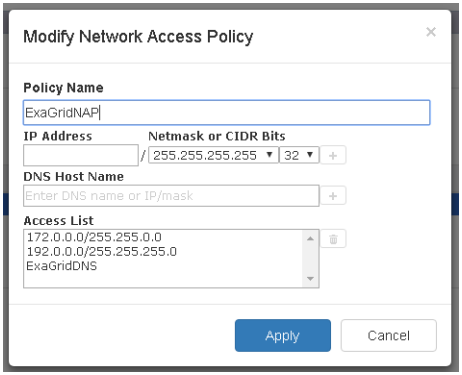

4. In the space provided, enter the Network Access Policy name.
5. In the space provided, enter the IP address of the backup server(s) that will access and backup to this share. You can also use the backup server's hostname (as known by the DNS server(s) you added to the ExaGrid network configuration) as described below.
6. Select a subnet mask for the IP address.
7. Click the **+** plus and the IP/subnet mask pair are entered into the Access list.
8. Or, in the space provided, enter the DNS hostname of the backup server.
9. If a backup server uses multiple IP addresses, or if more than one backup server will write to the share, repeat this process for all servers involved.
10. Click **Apply** and the Network Access Policy is created.

User Access Policies

User access policies allow you to control which users can access a share. A single User Access Policy can be used by multiple ExaGrid Shares.

Modifying a User Access Policy changes access for all shares that use it.

If the servers in the ExaGrid Site are members of an Active Directory Domain, share access control can be expanded to include users and groups who are also members of the Active Directory Domain. See " Managing Active Directory" on page 29.

Local ExaGrid users and their credentials can be added to the ExaGrid configuration and used to control access to CIFS shares (via SMB) or Veeam Data Mover shares. Access to shares using other share transports (NFS, OpenStorage Technology, etc.) can be controlled only by Network Access Policies. See "Network Access Policies" on page 44.

To create a User Access Policy:

1. In the navigation tree, select the ExaGrid Site on which the you want to create the User Access Policy.
2. In the top, main menu, click **Site** and all menu options for the site are displayed.
3. Under **Security**, select **User Access Policy** and the User Access Policy dialog box is displayed:

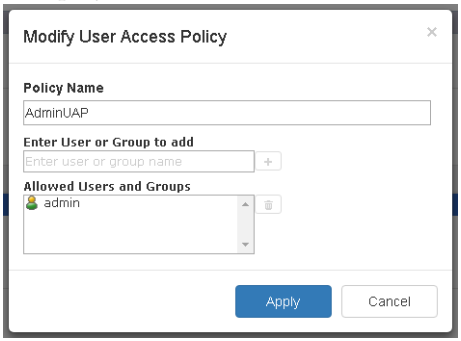

4. In the space provided, enter a name for the User Access Policy.
5. In the space provided, enter an ExaGrid local user name, an Active Directory user name, or an Active Directory group name.
6. Click the **+** plus and the name or group is added to the Allowed User and Groups list.
7. Repeat the process for all users and groups that need access from this User Access Policy.

Creating Shares

Caution – Never use external commands or utilities to create or delete ExaGrid shares. Always create and delete ExaGrid shares from the ExaGrid user interface.

Before you begin:

- Ensure that there is a Network Access Policy and a User Access Policy on the ExaGrid Site that will host the share that are applicable to your backup application server. See "Share Access Policies" on page 44. for details.
- When the Site is not joined to an Active Directory Domain, and you are creating a Veeam share to take advantage of the ExaGrid-Veeam Accelerated Data Mover transport, you must first create a local ExaGrid Veeam user. See the *Using Veeam Backup and Replication Software with an ExaGrid System* guide located in the ExaGrid Online Library for details.

To create a share:

1. In the navigation tree, select the ExaGrid Site on which the you want to create the share.
2. In the top, main menu, click **Site** and all menu options for the site are displayed.
3. Under the site's name, select **Shares and Replicas** and the Shares and Replicas page is displayed:

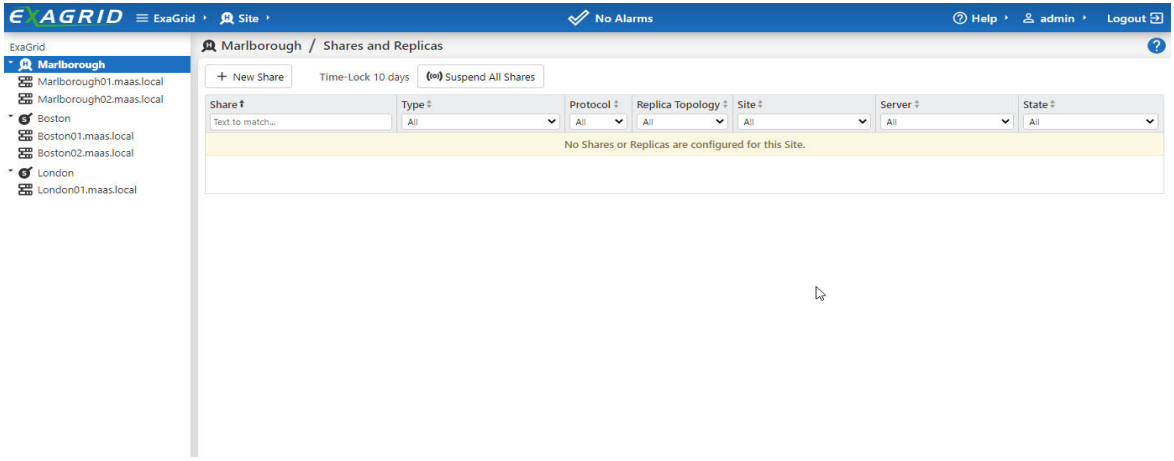

4. Click **+ New Share**, and the Create New Share dialog box is displayed:

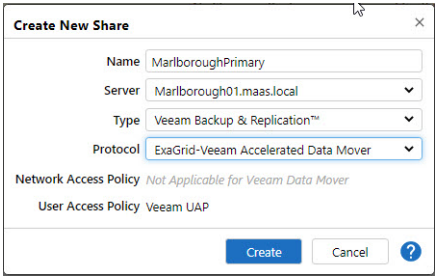

5. In the space provided, enter a name for the share. You must use unique share names on each ExaGrid Server. In addition, ExaGrid recommends using unique share names across your ExaGrid System. Doing so will:
   - Make ExaGrid reports easier to interpret.
   - Avoid share name conflicts in disaster recovery scenarios.
   - Avoid share name conflicts if you decide to migrate a share to another ExaGrid Server.
6. From the drop down menu, select the share type.
7. Select the protocol the application backup server will use to write to the ExaGrid Share: CIFS, NFS, OST or ExaGrid-Veeam Accelerated Data Mover.
8. If applicable to the share type and more than one Network Access Policy is available, from the drop down menu, select the Network Access Policy that will allow your backup server to write to the share.
9. From the drop down menu, select the User Access Policy that defines which users and groups can access the share. Optionally, (not recommended), you can select Guest User. Selecting Guest User will allow anyone access to the share.
10. If creating a Utility share, select the amount of space assigned to the share.

11. Click **Create** and the new ExaGrid Share is created:

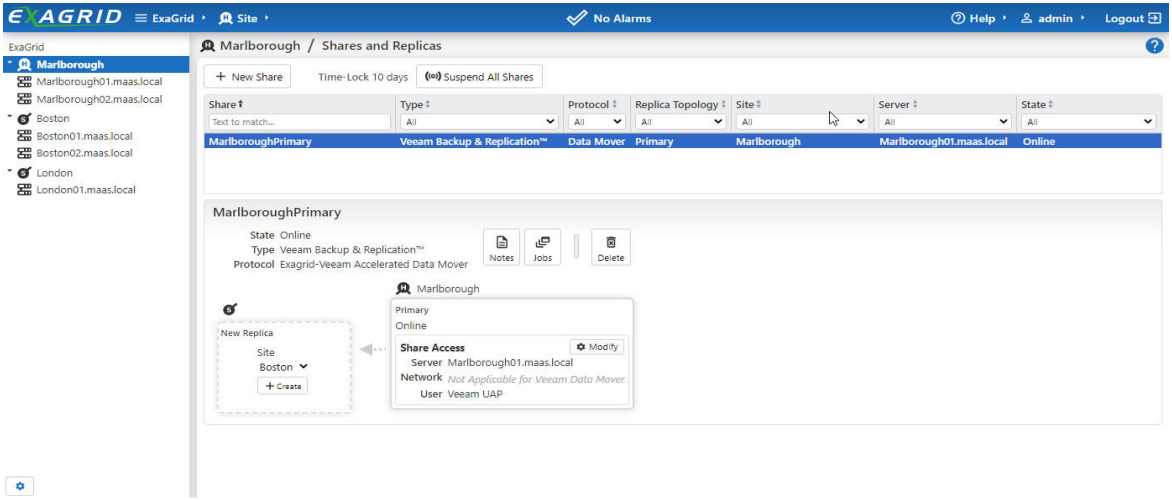

12. The newly created share will appear in the list of shares and because the new share is automatically selected, its details are displayed at the bottom of the page.

13. From the share's details section you can:

- Read detailed **Notes** that pertain to this specific share type
- View ExaGrid activity **Jobs** associated with the share
- **Delete** the Share

14. In addition, if this share is part of a multi-site ExaGrid System, you can replicate the share's data to other ExaGrid Sites and create an InstantDR share on the other ExaGrid Sites. See "Share Replication" on page 51.

Share Replication

Once an ExaGrid Share has been created (See “Creating Shares” on page 47), if it resides in a multi-site ExaGrid System, you can replicate that share to one or two ExaGrid Sites.

A copy of a share that has been replicated is known as a *replica*.

Replicating copies to other ExaGrid Sites provides an additional layer of protection for your data.

ExaGrid Spoke Sites can replicate shares to their ExaGrid Hub Site. If there is another Spoke Site, the shares replicated to the Hub Site can then be replicated to the other Spoke Site. Spoke Sites cannot directly replicate shares between each other.

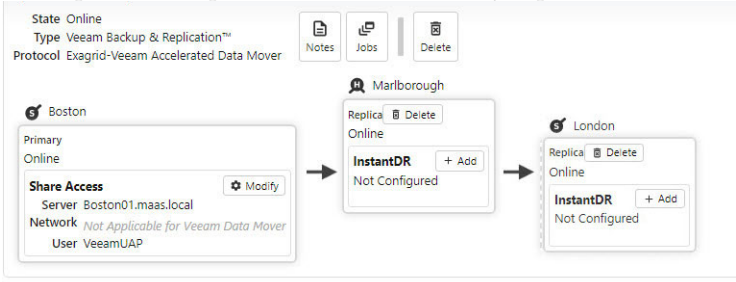

ExaGrid Hub Sites can replicate shares directly to one or two Spoke Sites:

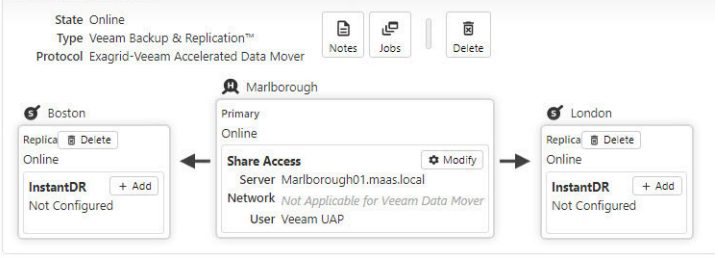

For Hub Sites, consider replicating all shares to the same spoke sites. Doing so will make a disaster recovery easier and faster.

If you create a share and then some time later replicate that share to another ExaGrid Site, the "age" of the share's replication queue is based on when the share was created and not when it was marked to be replicated.

Creating Replicas and InstantDR Shares

Note – Network and User Access Policies are kept separately on each ExaGrid Site. If a site does not have a defined User Access Policy, you must first create one before attempting to add an InstantDR Share.

Regardless of whether you are replicating data from a Hub or a Spoke, the basic process for creating replicas and InstantDR shares is the same:

1. Create an ExaGrid Share as described in "Creating Shares" on page 47
2. At the bottom of the ExaGrid Shares and Replicas page, select the **New Replica** box.

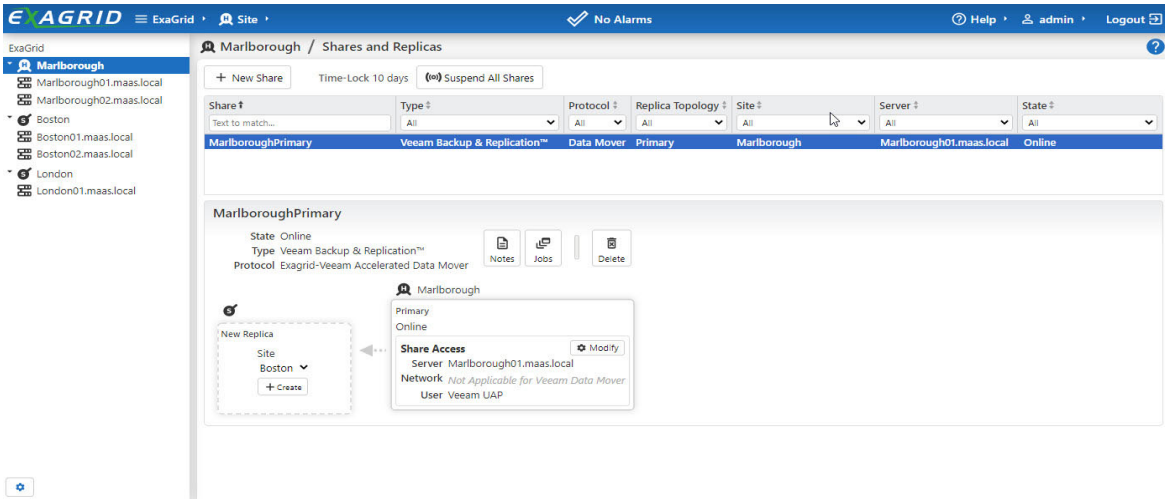

3. If you are creating a replica from a Hub that has multiple spokes, from the **Site** drop down menu, select the site on which the replica will be created.
4. Click **+Create** and the replica is created.
5. Once the replica has completed initialization, you can add an InstantDR share to the replica, by clicking **+Add** and the Add InstantDR Share dialog box opens.
6. From the **Server** drop down menu, select the ExaGrid Server on which the InstantDR Share will be created.

7. Depending on the share's Protocol, a User Access Policy (UAP), and/or a Network Access Policy (NAP) is required. For example:
   - CIFS- NAP is required, UAP is optional
   - NFS - NAP is required, UAP is not allowed
   - ExaGrid-Veeam Accelerated Data Mover - NAP is not allowed, UAP is required
   - OST - NAP is required, UAP is optional

   If only one UAP or NAP exists, it will be automatically selected if needed. For details See "Share Access Policies" on page 44.
8. Click **Add** and the InstantDR Share is created.

Hub Site To Spoke Site Replication Example

Once a share is created on a Hub Site it can be replicated to a spoke site.

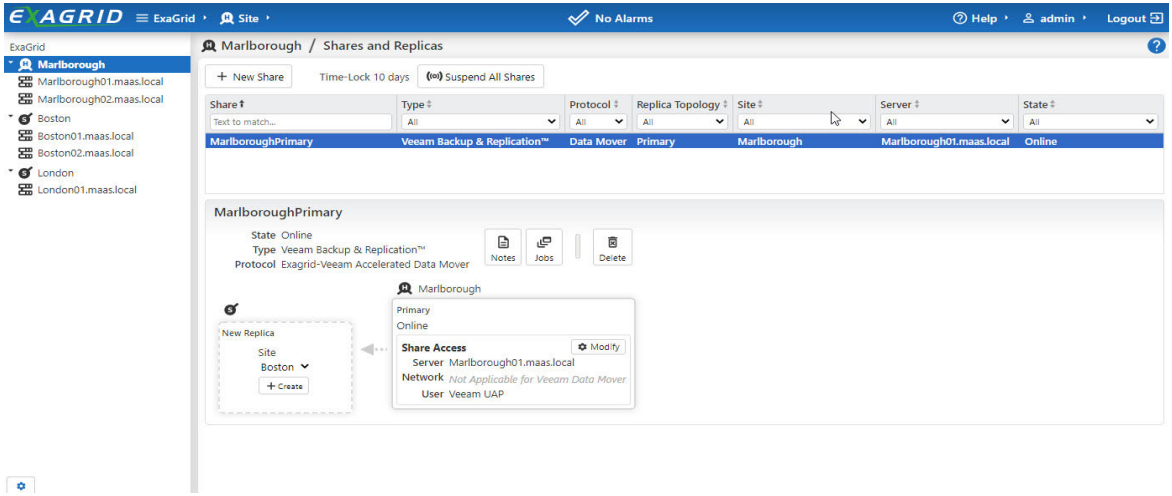

To replicate a share from a Hub Site to a Spoke Site

1. From the Site drop down menu in the New Replica box, select the site on which the replica will reside:

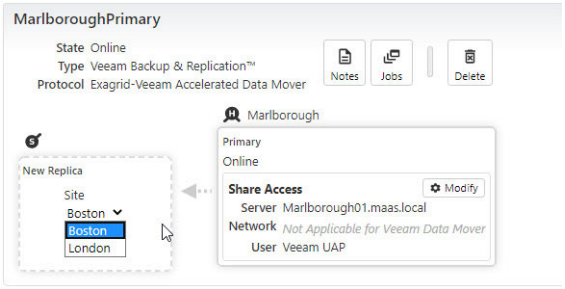

2. In the New Replica box click **Create** and the new replica is created on the selected site:

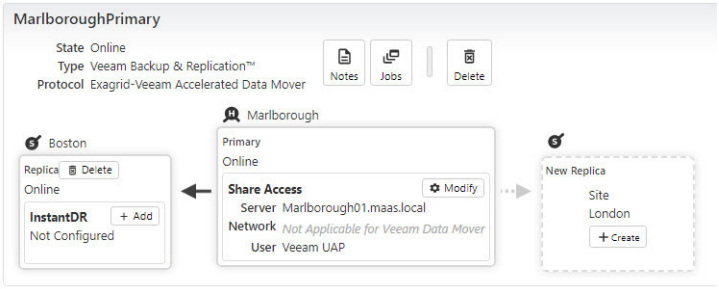

Once the share has been replicated, the Hub's share can be replicated to other ExaGrid Sites as needed.

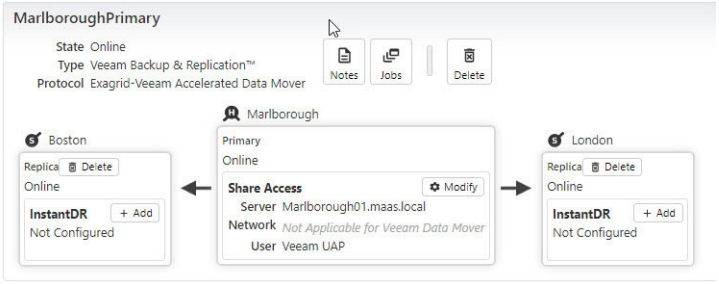

Spoke Site to Hub Site Replication Example

Once a share is created on a Spoke Site, it can be replicated to a Hub Site.

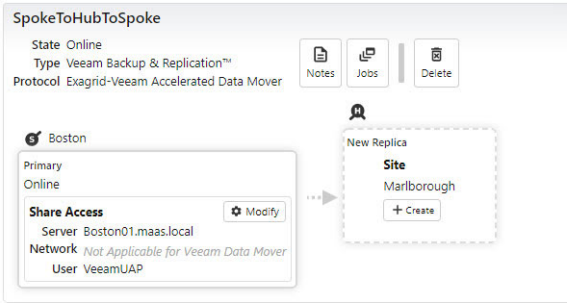

To replicate a share from a Spoke Site to a Hub Site, in the New Replica box click **Create**, and the replica is created on the Hub Site.

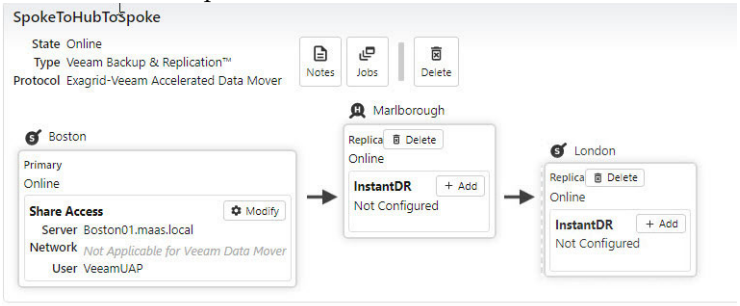

Once the Spoke Site's share has been replicated to the Hub Site, if there are other ExaGrid Sites in the system, it can be replicated to a second site by clicking **Create** in the New Replica box.

3.

Deleting Shares and Replicas

When deleting shares and replicas:

- Only a user assigned/mapped to the Admin role can delete shares and replicas
- Deleting a share will also delete its replicas

- If a Spoke Site's share is replicated to the Hub Site and then to another Spoke Site, you cannot delete the replica on the Hub Site without first deleting the replica on the Spoke Site.
- If a replica has an InstantDR share, you must first delete the InstantDR share before you can delete the replica.

Replication Scheduler

The ExaGrid user interface allows you to define when to allow a site's replication to occur. This is useful in those environments where, for specific days of the week or time of day a network must be reserved for other traffic.

Note – For additional control over replication bandwidth usage, see "Throttling Replication" on page 56.

To display the replication scheduler:

1. In the navigation tree, select the ExaGrid Site.
2. In the top, main menu, click **Site** and all menu options for the site are displayed.
3. Under **Configuration**, select **Replication Schedule** and the ExaGrid Site's replication schedule is displayed.

Follow the instructions on the page to schedule replication for the selected ExaGrid Site.

Throttling Replication

The Replication Bandwidth Throttling Management page is used to create and assign bandwidth throttling policies between an ExaGrid Hub Site its associated Spoke Sites.

Replication Bandwidth Throttling is set at the ExaGrid System level in the left-hand navigation tree.

You can create a policy for each ExaGrid Hub and ExaGrid Spoke pair in your ExaGrid System. Within a single policy, separate policies are applied to traffic going:

- To the spoke, from the hub

- From the spoke, to the hub

You can create more policies than you have hub and spoke pairs. However, only one policy can be applied to a hub and spoke pair.

Changes to a policy or the application of a new policy goes into effect within two minutes.

To create and apply a Replication Bandwidth Throttling policy:

1. In the ExaGrid navigation tree, select the top-level ExaGrid System icon.
2. From the top main menu, click **ExaGrid** and all of the menu options for the ExaGrid are displayed.
3. Under **Configure**, select **Replication Throttling** and the Replication Bandwidth Throttling Management page is displayed.
4. Click **Policy Manager...** and the Replication Throttling Policy Manager dialog box is displayed:

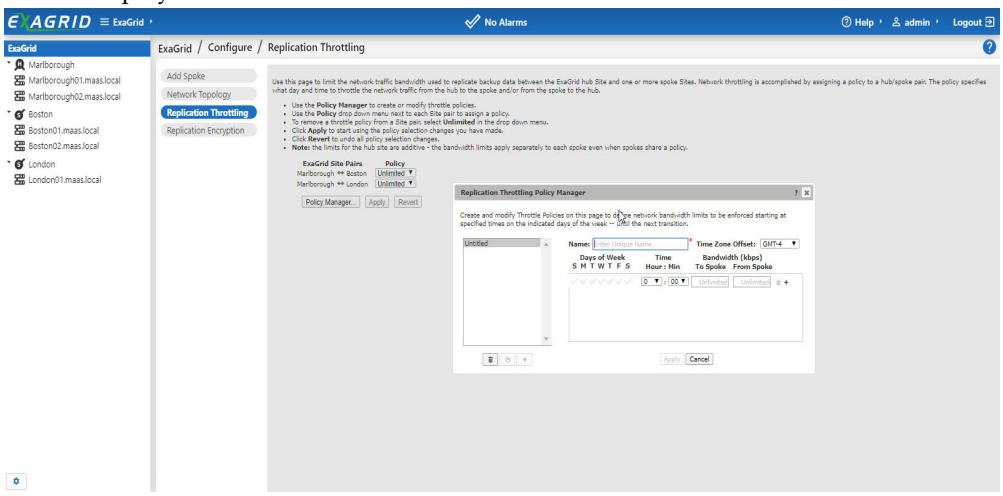

5. Enter a name. The policy name can be up to 23 characters long. It can only consist of the characters A-Z, a-z, 0-9, hyphen, and underscore. It can also have leading dots and a trailing dollar sign.
6. Select a time zone off set from GMT. Shifts for Daylight Savings Time must be done manually.
7. Define the days of the week and the time ranges (24 hour clock) when throttling will take place.
8. Set the **To Spoke** and **From Spoke** bandwidth in kbps (kilobits per second). There is a minimum bandwidth rate required by network communications between ExaGrid Sites for management and status reporting. **Unlimited** means bandwidth throttling is not be applied in that direction.

9. Use the **+** plus to add additional transitions as needed. A single policy can contain a maximum of sixteen transitions.
10. Click **Apply** and the new policy appears in the list of polices in the left-hand side of the dialog box.
11. Close the dialog box.

Tip – The fastest way to create a new policy is to select an existing policy and then click the Clone button. This will clone the selected policy. You can then edit the cloned policy according to your needs.

To apply a policy to a hub and spoke pair:

1. From the drop down menu next to a hub spoke pair, select a policy.
2. Click **Apply** and the policy is applied to the hub and spoke pair.

To remove a policy from a hub and spoke pair:

1. From the Replication Bandwidth Throttling Management page, select **Unlimited** from the drop down menu next to a hub spoke pair.
2. Click **Apply** and the hub and spoke pair will not have an assigned policy.

Encrypting Replication

The Replication encryption page is used to configure the encryption of replicated backup data as it travels across the network between an ExaGrid Hub Site and an ExaGrid Spoke Site and vice versa. When enabled, standard AES-256 encryption is used to protect the deduplicated backup data sent between sites. This encryption method conforms to the FIPS 197 publication. If the remote site consists of secure ExaGrid models, the unencrypted deduplicated backup data is stored on FIPS-compliant encrypted disk drives; if the remote site consists of non-secure ExaGrid models, the unencrypted deduplicated backup data is stored on non-encrypted disk drives.

Note – Replication to a Cloud ExaGrid Site is always encrypted.

---

1. In the ExaGrid navigation tree, select the top-level ExaGrid System icon.
2. From the top main menu, click **ExaGrid** and all of the menu options for the ExaGrid are displayed.

3. Under **Configure**, select **Replication Encryption** and the Replication Encryption page is displayed.

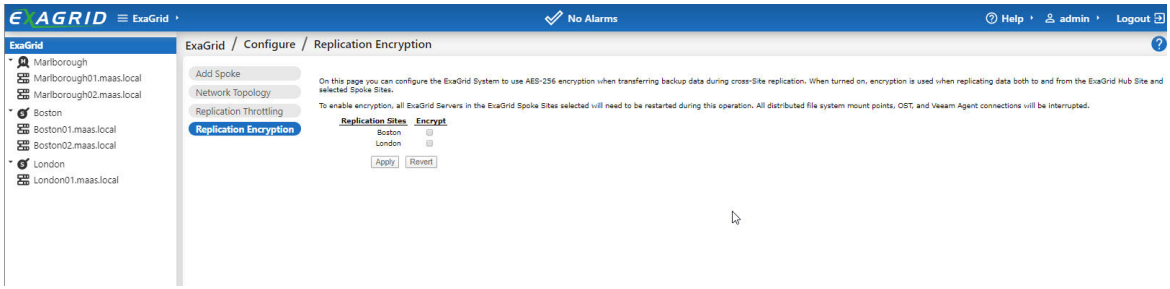

4. Select the replication site(s) to be encrypted.
5. Click **Apply** and the data will be encrypted.

System Record Keeping

In the event of a disaster where an ExaGrid Server, or an entire ExaGrid Site is lost, having kept records of your ExaGrid System will minimize the amount of time and effort required to restore your data.

Use the worksheets in this chapter to record information associated with your ExaGrid System.

Gathering and recording this information a head of time will make the disaster recovery process easier.

System setting can change over time, so it is important to keep these records up to date.

In the event of a disaster, verify this information immediately prior to starting the disaster recovery process.

Caution – Names are case sensitive, enter them exactly as they where assigned to the lost ExaGrid resources.

ExaGrid Site Worksheet

In the event of a disaster, to recover a lost ExaGrid Site, you need to provide information about the ExaGrid Site involved in the event.

Note – The ExaGrid Hub Site is always displayed in the top-most position in the ExaGrid navigation tree.

Use the following worksheet to record the required information

ExaGrid Site Worksheet	
Site's Physical Location	
Site's Name	
Was this site the Hub?	**___YES or NO**
Hub Site's IP Address	
Hub Site's Administrative Key (keys expire after 24 hours)	

To obtain the Hub Site's IP address and Administrative Key see "ExaGrid Site Administrative Key" on page 62

ExaGrid Site Administrative Key

To find an ExaGrid Site's IP and Administrative Key information:

1. In the ExaGrid navigation tree, select the ExaGrid Site.
2. In the top, main menu click Site, and all menu options for the site are displayed.
3. **Under Security,** select **Site Administration Key** and the Site Administration key page is displayed.

ExaGrid Server Worksheets

In the event of a disaster, to recover lost ExaGrid Servers, you will have to provide network information about the ExaGrid Servers involved in the event.

To find network configuration information for an ExaGrid Server:

1. Select an ExaGrid Server.
2. In the top, main menu click Server and all menu options for the server are displayed.
3. Under Configure, select **Network Configuration** and the ExaGrid Server Network Configuration page is displayed.

Note – If the network information can not be displayed, contact your network administrator.

Use the following worksheet to record the required information:

ExaGrid Server Name:_____

ExaGrid Site Name:_____

ExaGrid Server Network Configuration

NIC Name	Speed	Enabled	MTU	DHCP or Static	IP plus VLAN	Masks	Gateway	ExaGrid Traffic	Backup Traffic

ExaGrid Server Network Interface Bonds

NIC Name	Speed	Unbonded or Bond Membership	Bond Mode

ExaGrid Share Worksheet

Keep detailed records for each share. Whether online or paper, keep the records in a secure place. The following is an example of a share record.

ExaGrid Share Worksheet	
Share name:	
Date created:	Created by:
Date modified:	Modified by
Date deleted:	Deleted by:
Replicated to a second ExaGrid Site	Yes____ No____
Replicated to a third ExaGrid Site	Yes____ No____
Backup Server Type	
Type of data being backed up	
Copy to tape	Yes____ No____

The IPMI Port

The IPMI port is located in the back of your ExaGrid Server. The actual location varies by ExaGrid Model. NIC ports used for backup and administration are labeled, the IPMI port is a standard 1GbE port without a label. For example:

Warning – The following procedure exposes a system interface, that if improperly used can cause a loss of data and void your system warranty.

Configuring The IPMI Port

Before you can use an IPMI network port, it must first be configured from the ExaGrid user interface

1. From the ExaGrid navigation tree, select the ExaGrid Server whose IPMI network port you want to configure.
2. From the top, main menu, click **Server**, and all menu options for the server are displayed.
3. Under **Configure**, select **Network IPMI** and the IPMI network Configuration page is displayed:

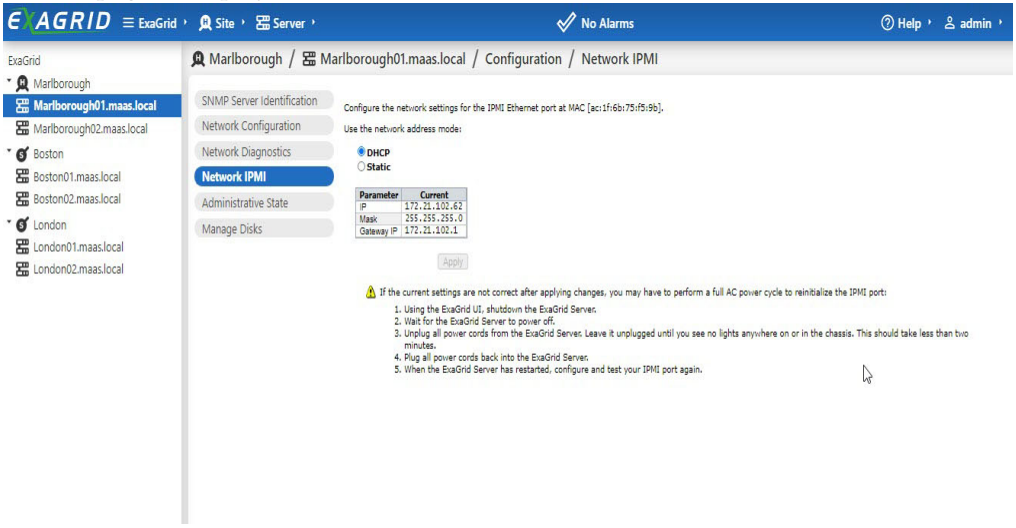

4. In the spaces provided enter:

   a. An IP address
   b. A subnet mask
   c. Gateway IP
5. Click **Apply** and the IPMI network port will be configured.

Resetting the IPMI Network Port User Password

ExaGrid Servers come with a default IPMI password.

If you are configuring the password for the first time, obtain the default user name and password from your ExaGrid support engineer.

ExaGrid recommends that you change the default password.

Once the password is reset, record the new password and keep it in a secure location.

Warning – ExaGrid **CANNOT** retrieve lost passwords. If a password is lost, your data is at risk.

To change the password:

1. From a web browser session, enter the IPMI network port address and the login page is displayed:

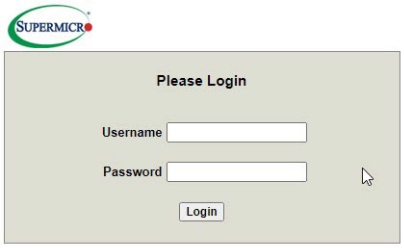

2. Enter the existing IPMI username and password.

3. Click **Login** and the System Information page is displayed:

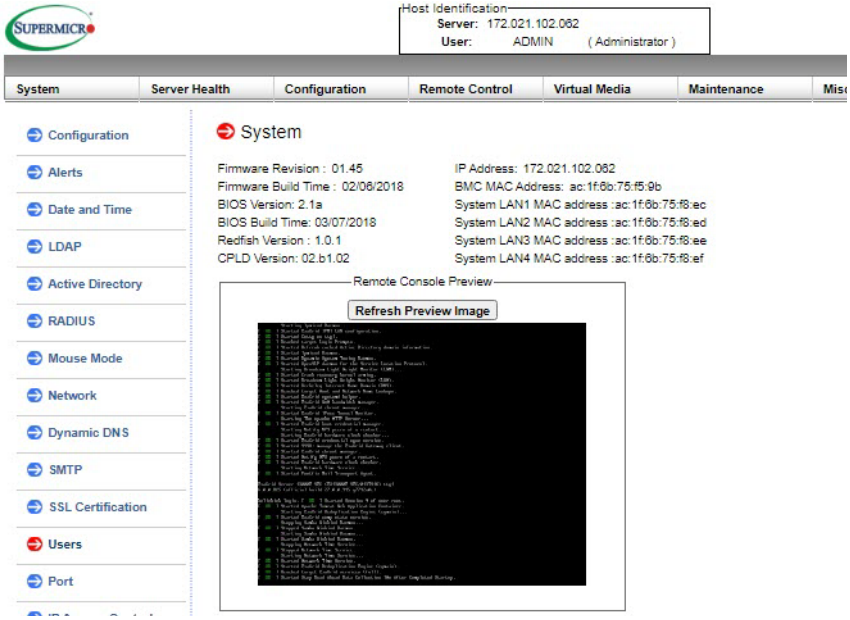

4. In the main menu, click **Configuration** and from the drop down menu, select **Users** and the user list is displayed:

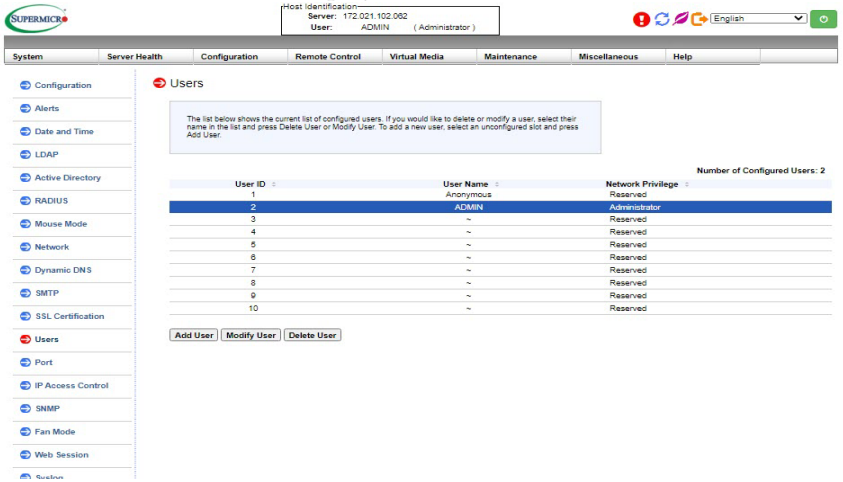

5. Select the **ADMIN** user.

6. Click **Modify User** and the Modify User page is displayed:

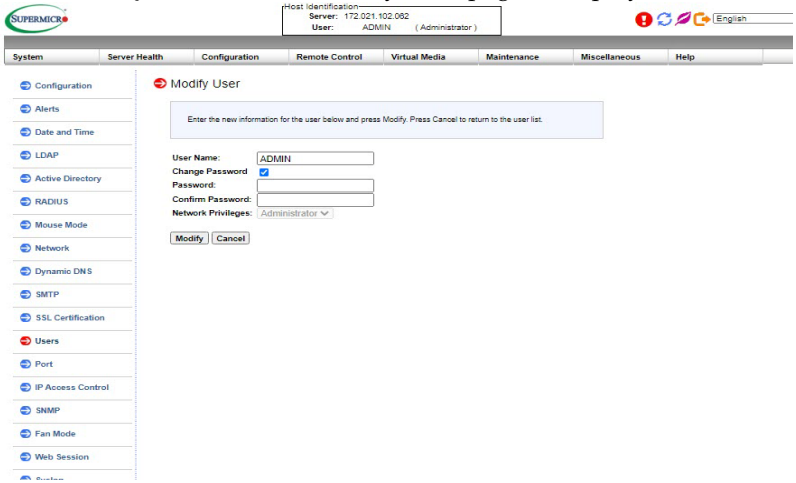

7. Select (Check) **Change Password.**
8. In the spaces provided, enter the new password.
9. Click **Modify**, and the password is changed.

Caution – Record the new password and store it in a safe location. ExaGrid **CANNOT** retrieve lost passwords. Lost passwords can result in data loss.

Encryption Capable ExaGrid Servers

Encryption Capable ExaGrid Servers run the same ExaGrid software as other ExaGrid Servers, with the added features of Encryption at Rest and Boot Protection.

You can not mix Encryption Capable and non-Encryption Capable ExaGrid Servers in an ExaGrid Site. You may replicate a Share(s) from an Encryption Capable Site to a non-Encryption Capable Site or vice versa, however it is typical to have both sites be the same - either all Encrypted Sites, or no Encrypted Sites.

Encryption Capable ExaGrid Server's can be identified by the -SEC suffix appended to their name. For example: EX63E-SEC.

Encryption at Rest and Boot Protection are enabled from the ExaGrid user interface at the ExaGrid Site level. This ensures that all ExaGrid Servers in a site use the same settings and password.

Warning – ExaGrid **CANNOT** retrieve lost passwords. If a password is lost your data is at risk.

Encryption at Rest

Warning – ExaGrid **CANNOT** retrieve lost passwords. If a password is lost your data is at risk.

When Encryption at Rest is enabled, if a disk is removed from an Encryption Capable ExaGrid Server, its data is encrypted and cannot be read. However, if the entire Encryption Capable ExaGrid Server is removed it can be restarted and its data can be read.

Once Encryption at Rest has been enabled and the password established, it cannot be disabled.

Warning – You will be required to provide your Encryption at Rest password if an Encryption Capable ExaGrid Server is replaced as a result of a disaster or if your ExaGrid Server requires a major overhaul i.e. chassis swap, new RAID card, etc. In such cases, if you can not provide your password, **all of your data will be lost**.

Encryption at Rest Passwords

Warning – ExaGrid **CANNOT** retrieve lost passwords. If a password is lost your data is at risk.

Encryption at Rest passwords:

- Are case sensitive
- Must be between 8 and 32 characters
- Must **not** contain the <SPACE> character
- Must contain at least
  - One number
  - One lowercase letter
  - One uppercase letter
  - One non-alphanumeric character from the set: @#^*-_=+

Note – If you can not enable or change your Encryption at Rest password, check for a failed drive or a drive logging errors in your system. Follow normal support procedures to replace the problem drive. Encryption at Rest password management will function normally once the problem drive is replaced.

Enabling Encryption at Rest

When an ExaGrid Encryption Capable Server is first initialized it appears in the ExaGrid navigation tree in an "alarmed state" until you explicitly enable or opt out of the security Encryption at Rest and/or Boot protection.

If you mouse over the ExaGrid Site name, details about the alarm are displayed.

To enable Encryption at Rest:

1. In the left-hand navigation tree, select an ExaGrid Site
2. From the top, main menu, click **Site**, and all menu options for the site are displayed.
3. Under **Security**, select **Encryption at Rest** and the Encryption at Rest Management page is displayed.
4. Next to an ExaGrid Server's name click Configure and the Enable Encryption at Rest dialog box opens::

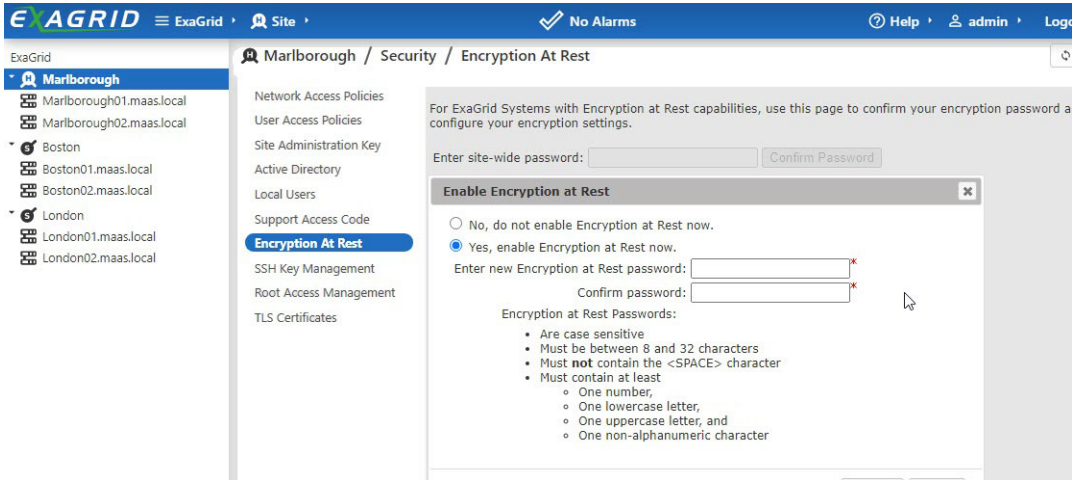

5. Select **Yes, enable encryption at rest**.
6. In the space provided, enter the Encryption at Rest password.

Note – Selecting **No, do not enable Encryption at Rest now** will clear the Encryption at Rest alarm and does not apply a password for encryption at rest. You can revisit this page at another time and enable Encryption at rest.

Warning – ExaGrid CANNOT retrieve lost passwords. If a password is lost your data is at risk.

7. Click **Apply** and your selection will be applied

Adding a New Encryption Capable Server to a Site

Note – You can not mix Encryption Capable and non-Encryption Capable ExaGrid Servers in an ExaGrid Site.

To add a new Encryption Capable server to an Encryption Capable ExaGrid Site:

1. Initialize the ExaGrid Server and add it to the site as detailed in the ExaGrid *Assembling and Initializing your ExaGrid Server* guide.
2. When the server has been initialized and added to the site, in the left-hand navigation tree, select the ExaGrid Site.
3. From the top, main menu, click **Site**, and all menu options for the site are displayed.
4. Under **Security**, select **Encryption at Rest** and the Encryption at Rest Management page is displayed:
5. In the space provided, enter the site's Encryption at Rest password.

Warning – ExaGrid **CANNOT** retrieve lost passwords. If a password is lost your data is at risk.

6. Click **Confirm Password** and the ExaGrid Servers in the site are listed. The new ExaGrid Server will be listed as **Not Configured**.

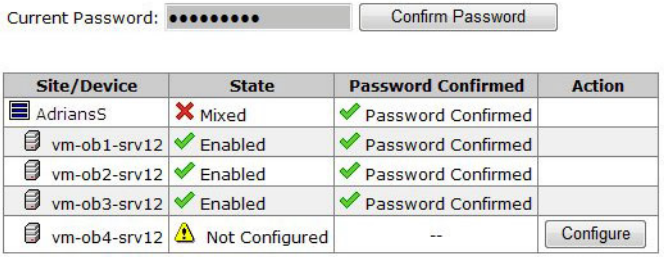

Click the **Configure** button and the new ExaGrid Server will be configured with the same password and boot protection setting as the other ExaGrid Servers in the site.

Server's Encryption Password Does Not Match Site's Password

In the rare event that an Encryption Capable Server's password does not match that of the other servers in a site, the ExaGrid navigation tree will display an alarm.

To address the alarm:

Navigate to the Encryption at Rest Management page:

1. In the left-hand navigation tree, select the ExaGrid Site.
2. From the top, main menu, click Site, and all menu options for the site are displayed.
3. Under **Security**, select **Encryption at Rest** and the Encryption at Rest Management page is displayed.
4. In the space provided, enter the site's Encryption at Rest password.

Warning – ExaGrid cannot retrieve lost passwords. If a password is lost your data is at risk.

---

5. Click **Confirm Password** and the ExaGrid Servers in the site are listed. The new ExaGrid Server will be listed as having an **Incorrect Password**.

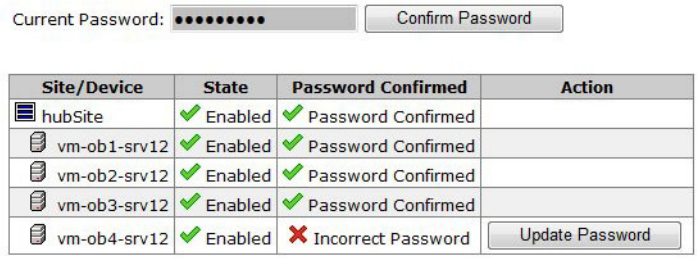

6. Click the **Update Password** button and the update password dialog box is displayed:

7. In the space provided, enter the current Encryption at Rest password for the server.

Warning – ExaGrid cannot retrieve lost passwords. If a password is lost your data is at risk.

---

8. Click **Update**, and if you entered the correct password, the server will use the same password as the other servers in the site.

Boot Protection

Warning – ExaGrid cannot retrieve lost passwords. If boot protection is enabled and the password is lost, at the next reboot **your data will be inaccessible until the Encryption at Rest password is provided**.

To enable Boot Protection, you must contact ExaGrid Support.

Boot Protection uses the Encryption at Rest password.

An Encryption Capable ExaGrid server encrypts the data on its disks. If a disk is taken, its data is encrypted and cannot be read. Enabling boot protection will protect data even if the entire server is taken. When a server with Boot Protection is restarted, data cannot be accessed until the Boot Protection encryption password is provided.

ExaGrid recommends using the IPMI network port at the back of your Encryption Enabled ExaGrid Server. Doing so will facilitate entering your encryption password when your ExaGrid Server is rebooted.

See "Using the IPMI Network Port to Enter Boot Protection Password" on page 80. for details on how to setup and use your ExaGrid Server's IPMI network port.

Software upgrades, system maintenance, power events can all require or cause system reboots. Each rebooted system will require a console connection to enter the boot protection password.

Tip – For system maintenance procedures that require reboots, deactivate Boot Protection for the duration of the procedure. Once complete, reactivate Boot Protection.

Enabling Boot Protection

Call ExaGrid Support. Your ExaGrid Customer Support representative will use a remote login, enable boot protection on your Encryption Enabled ExaGrid Servers.

Once Boot Protection is enabled:

1. In the left-hand navigation tree, select the ExaGrid Site
2. From the top, main menu, click **Site**, and all menu options for the site are displayed.

3. Under Security, select **Encryption at Rest** and the Encryption at Rest Management page is displayed.
4. In the space provided, enter the Encryption Password.
5. Click **Confirm Password** and the Encryption Management page will display the encryption details for the selected ExaGrid Site.
6. Click **Configure** and the Encryption at Rest Settings dialog box is displayed:

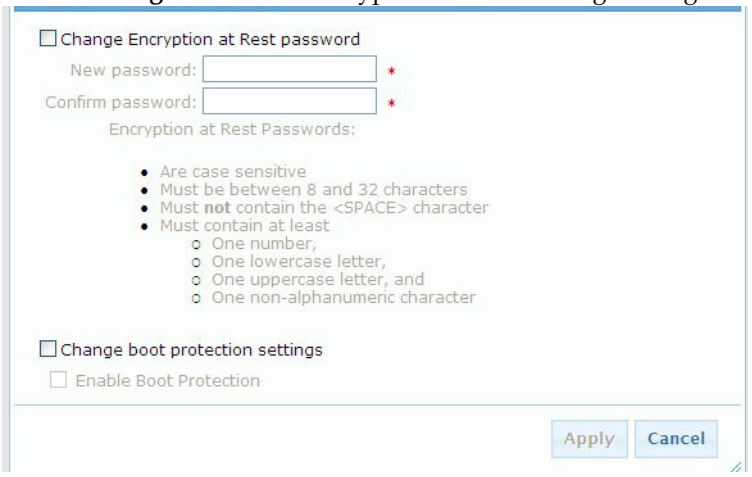

7. From this dialog box you can change the current Encryption at Rest password and/or enable/disable Boot Protection.
8. Make your selections as needed.
9. Click **Appl**y and you will be prompted to ensure that you want to enable Boot Protection.
10. If you enabled Boot Protection, the Encryption at Rest Management details page is updated with the Boot Protection listed as **On**.

Using the IPMI Network Port to Enter Boot Protection Password

Using the IPMI network port will facilitate entering your encryption password when boot protection is enabled.

Warning – The following procedure exposes a system interface, that if improperly used can cause a loss of data and void your system warranty.

To reboot an encryption enabled server:

Note – If you have not yet set up you IPMI port, username, and password, see "Configuring The IPMI Port" on page 68.

1. Enter the IPMI network port address from a secure (HTTPS) web browser session and the login page is displayed:

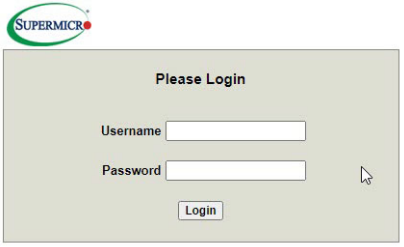

2. Log into the page with the login username and password and the System Information page is displayed:

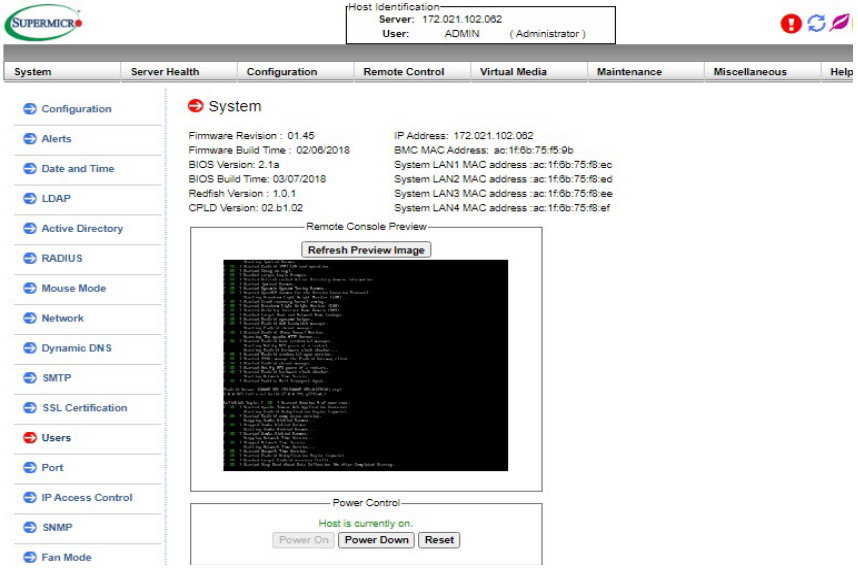

3. In the main menu, click the **Remote Control** tab and the Remote Control page is displayed:

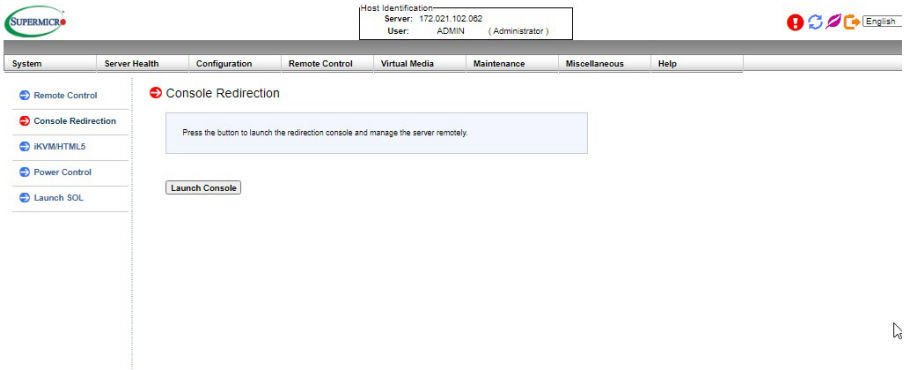

4. Click **Launch Remote Console** and a command window is displayed:

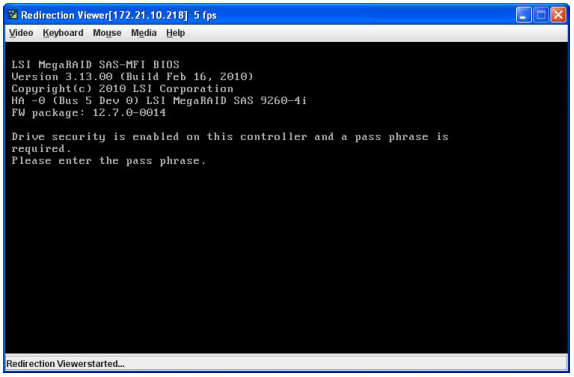

Note – The first time you launch the remote console, a Java application is downloaded. Install the application and then the command window will open.

---

5. At the **Please enter the pass phrase:** prompt, enter the Encryption password you assigned to ExaGrid Servers in this ExaGrid Site.
6. Press **Enter** and the ExaGrid Server will reboot.

Repeat this process for each ExaGrid Server in your ExaGrid Site as needed.

Recovering From Malicious Software Attacks, or a Loss of ExaGrid Servers

ExaGrid provides separate processes for recovering from:

- Ransomware and malicious software attacks. See “Recovering From Malicious Software Attacks” on page 84.
- The loss of one or more ExaGrid Servers. See “Recovering From The Loss of ExaGrid Servers - Disaster Recovery” on page 88.

In either of these scenarios, your ExaGrid Support team stands ready to assist and walk you through the recovery process.

Before an Event Occurs

ExaGrid recommends that you have a recovery plan in place that addresses your corporate Recovery Point Objectives and Recovery Time Objectives. The plan should include:

- Reviewing ExaGrid security settings on a regular schedule, including ExaGrid’s Retention Time-Lock settings, to ensure that the maximum security settings are applied to your backed up data. See “Managing and Securing Access” on page 19.
- Prioritize which servers will come back on line first.
- How and when to backup to new ExaGrid Shares.
- Prioritize which ExaGrid Servers and/or ExaGrid Shares will be restored first.

Recovering From Malicious Software Attacks

Recovering from a malicious software attack consists of:

- “Suspending All Shares” on page 84
- “Recovering From a Point In Time” on page 85
- “Post Recovery” on page 88

Suspending All Shares

Suspending all shares is the first step in recovering from a ransomware or malicious software attack.

Note – Utility Shares are not suspended and cannot be recovered to a point in time. See “Utility Shares” on page 41.

As soon as malicious activity is detected, on each ExaGrid Site impacted by the malicious activity, Suspend All Shares:

1. Log in to the ExaGrid Site as a user assigned the Admin role.

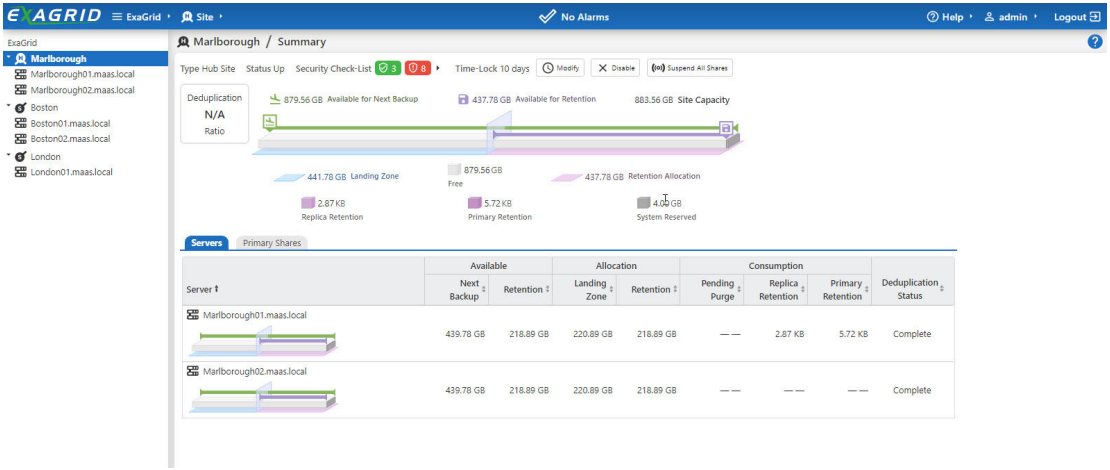

2. At the top of the ExaGrid Site's Summary page click **Suspend All Shares**, and the Suspend All Shares dialog box is displayed:

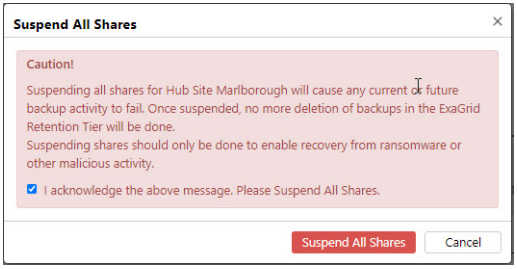

3. Read the CAUTION text in the dialog box and understand that once shares are suspended, current and future backups for this site will fail.
4. If you understand the CAUTION text and want to proceed, check (click) **I acknowledge the above message**.
5. Click **Suspend All Shares** and the dialog box closes and all primary shares in the current ExaGrid Site are suspended.

Note – Shares may also be suspended from an ExaGrid Site's Shares and Replicas page.

---

6. Once shares are suspended, use your backup application to access each suspended share and determine if the latest backups have been deleted or damaged. This process will vary by backup application. For details, consult the ExaGrid application specific guide in the ExaGrid Online Library.
7. If a suspended share is determined to be undamaged, click **Return to Online** and it will be brought back online and can be put back in to production.
8. After damaged shares have been identified, use the following Point In Time Recovery process to find and recover undamaged backups.

Recovering From a Point In Time

Suspended ExaGrid Shares are used to recover data from a specific point in time. Data is recovered from backups that reside within the window defined by the ExaGrid Site's Retention Time-Lock settings. See "Retention Time-Lock Overview" on page 20.

To recover data from a point in time:

1. Login to the ExaGrid Site as a user assigned the Admin role.
2. From the top, main menu, click **Site** and all menu options for the site are displayed.

3. Under the site's name click **Shares and Replicas** and the Shares and Replicas page is displayed:

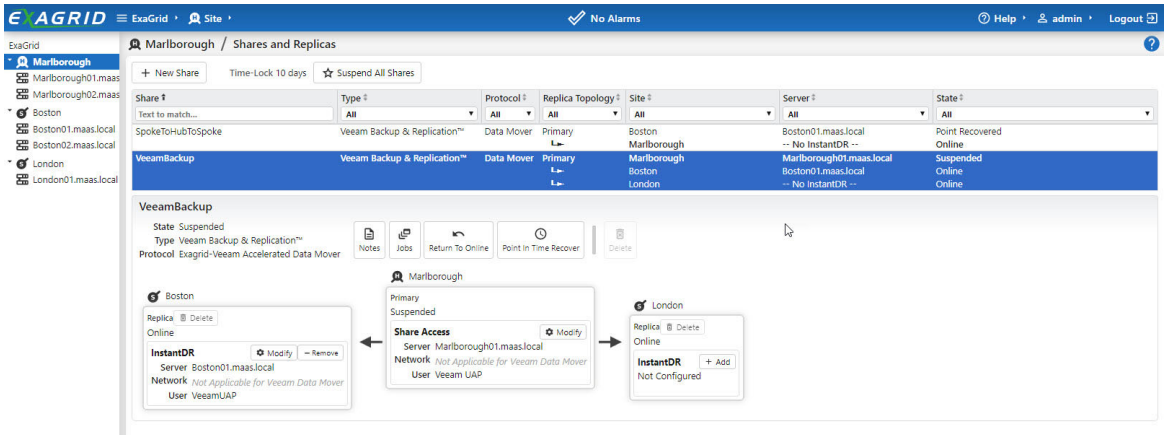

4. Select a suspended ExaGrid Share.
5. Click **Point In Time Recovery** and the Point In Time Recover dialog box is displayed:

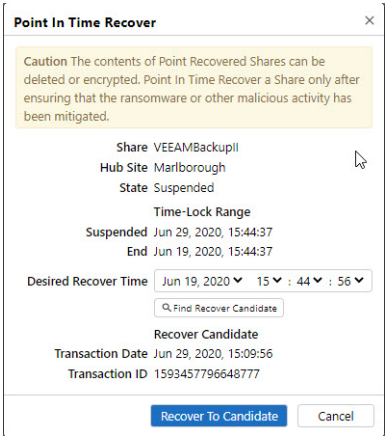

6. Next to **Desired Recover Time**, select a date and time of the last known good, completed backup.
7. Click **Find Recover Candidate** and a recovery candidate closest to the desired date and time is selected.
8. If the Desired Recover Time found does not meet with your approval, select another date and time and click **Find Recover Candidate.**

9. If the Desired Recover Time meets with your approval, click **Recover To Candidate** and a Recover Point for the Desired Recover Time is created:

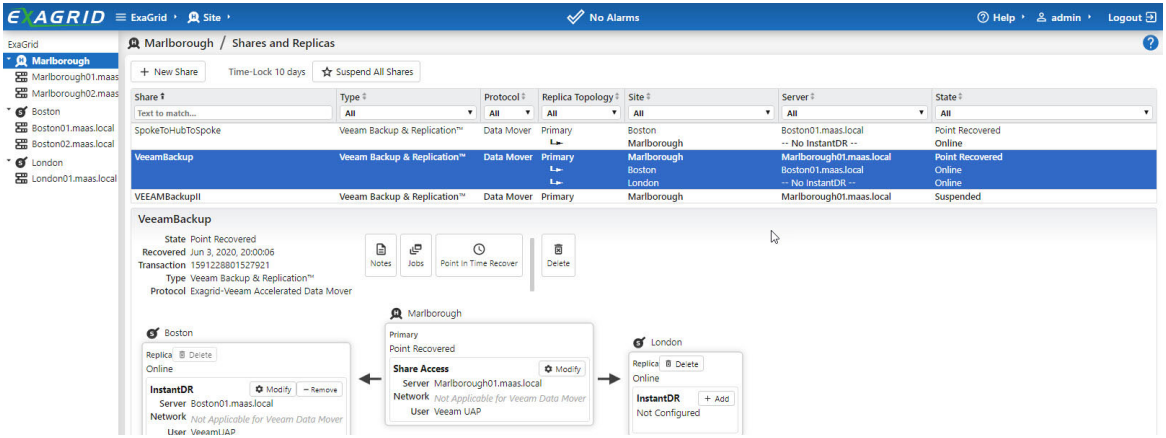

Note – The data from the specified time is recovered to the existing share, overwriting any existing content and the recovered share becomes read-write. No subsequent modifications to the recovered share content will be deduplicated or put into the retention tier.

10. Once the recover point is created, use your backup application to access the recovered share and determine whether it has been damaged. This process will vary by backup application. For details, consult the ExaGrid application specific guide in the ExaGrid Online Library.

11. If the Desired Recover Time is damaged, from the selected share's Shares and Replicas page, click **Point In Time Recovery** again and the Point In Time Dialog box is displayed:

12. Select a new **Desired Recover Time**. Repeat the process until undamaged backups are found.
13. When you have found a Desired Recover Time that is undamaged, use your backup application's standard procedures to recover the data.

Note – The data from the specified time is recovered to the existing share, overwriting any existing content and the recovered share becomes read-write. No subsequent modifications to the recovered share content will be deduplicated or put into the retention tier.

---

14. Repeat this process for each suspended share in the ExaGrid Site.

Post Recovery

Once your environment is clear of all threats and your data has been restored, to resume your backups:

- Create new shares on the ExaGrid Site(s). This step can be done in parallel to restoring your data to suspended shares.
- Begin backing up data to the new ExaGrid Shares.

For additional help and guidance, contact your ExaGrid support engineer.

Recovering From The Loss of ExaGrid Servers - Disaster Recovery

If a catastrophic event occurs where ExaGrid Servers or an entire ExaGrid Site is no longer available, the shares from the lost servers or site can be recovered if:

- The loss was part of a multi-site system
- Shares on the lost site were configured to replicate to other ExaGrid Sites
- One of the sites to which shares were replicated is still available

To perform a disaster recovery, in your surviving ExaGrid System's online library, consult:

- The *ExaGrid Disaster Recovery* Guide. This guide provides step-by-step instructions on how to recover your lost ExaGrid System.
- Contact ExaGrid Customer Support before initiating a disaster recovery operation.

- Once the ExaGrid System has been recovered, consult the ExaGrid application guide for your backup application. This guide provides step-by-step instructions on how to recover your data from your backup application.

ExaGrid Reports

Because your ExaGrid System employs a unified user interface, from a single log in, you can view reports for all ExaGrid Sites and ExaGrid Servers in your ExaGrid System.

Note – The ExaGrid System is not a real time reporting system. Most reports are updated every fifteen minutes. If a report displays data that appears to be out of date, you may be viewing data between update cycles

Report data is rolled up according to what is selected in the navigation tree. For example, when you select:

- The ExaGrid icon, you will see reports on the entire ExaGrid System. See "ExaGrid Summary Report" on page 93
- An ExaGrid Site name, you will see reports on all ExaGrid Servers in the selected ExaGrid Site. See "ExaGrid Site Summary Report" on page 94
- An ExaGrid Server, you will see reports specific to the selected ExaGrid Server. See "ExaGrid Server Summary Report" on page 95.

ExaGrid reports allow you to monitor how closely you are adhering to the initial sizing of your ExaGrid System:

- The amount of data being backed up
- The frequency at which data was backed up
- Expected data growth
- How long data would be retained.

Reported changes in consumption can indicate that you are out growing your ExaGrid System.

Summary Reports

A summary report is the default report displayed when you select an item in the ExaGrid navigation tree.

The purpose of a summary report is to, at a very high-level, answer basic questions about your backups:

- Will my next backup fit on my ExaGrid System's Landing Zone Tier?
- Will the next deduplicated backup fit into my ExaGrid System's Retention Tier?
- How well is my data deduplicating?

Summary reports provide rolled up, high-level details on:

- The entire ExaGrid
- An ExaGrid Site
- An ExaGrid Server

ExaGrid Summary Report

The ExaGrid System-level Summary report is the default landing page when you click on the ExaGrid name in the ExaGrid navigation tree.

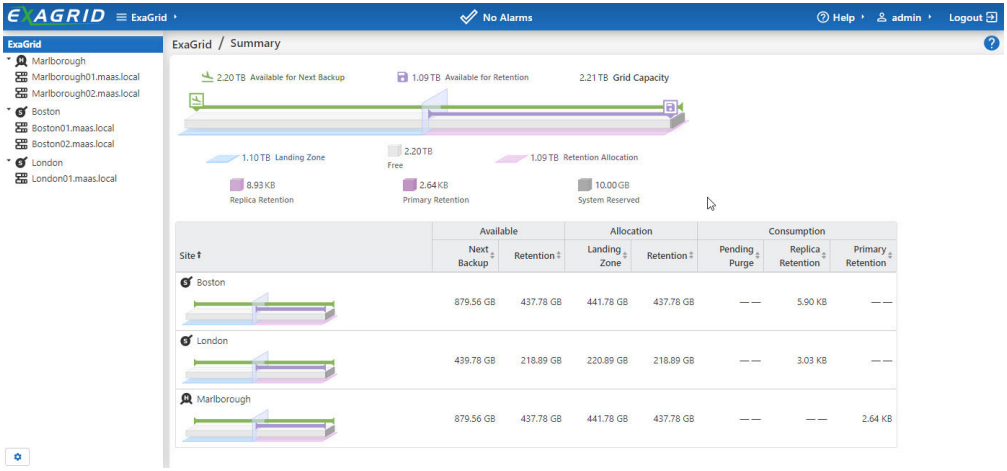

The ExaGrid System-level Summary Report is divided in to two sections:

1. The first and top-most section provides a system-level graphical representation of how space is allocated and consumed.
2. The second section provides detailed information on each ExaGrid Site in the system:
   - Available - How much space is available to the next backup and how much space is available for retention.
   - Allocation - How much space is allocated to the Landing Zone Tier and how much space is allocated to the Retention Tier.
   - Consumption - How much space is being consumed by data waiting to be purged, how much space is being consumed by data replicated to an ExaGrid Server, and how much space is being consumed by primary shares that reside on an ExaGrid Server.

ExaGrid Site Summary Report

The ExaGrid Site Summary report is the default landing page when you log in to an ExaGrid Site.

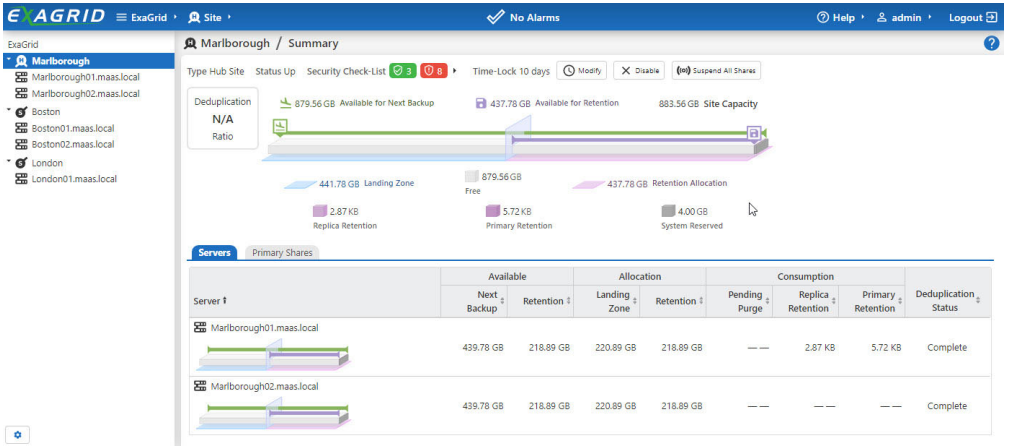

The ExaGrid Site-level Summary Report is divided in to three sections:

1. The first and top-most section allows you to quickly:

- View Security checklists and status: These lists allow you to quickly determine whether the security settings in your ExaGrid Site meet ExaGrid's recommended best practices. Click on a list item to display the page where the security setting can be configured. For details on ExaGrid's security and access best practices. See "Managing and Securing Access" on page 19.
- View settings for Time-Locks: Displays the site's Retention Time-Lock setting in days and allows you to make Retention Time-Lock change requests to the site's Security Officer. See "Retention Time-Lock Overview" on page 20.
- Suspend All Shares: If a ransomware or other malicious software attack is detected, you can suspend all shares in the ExaGrid Site and if necessary, begin recovering from the event. See "Recovering From Malicious Software Attacks" on page 84.

2. The second section provides a site-level graphical representation of how well data is deduplicating, how space has been allocated, and how much space is available.
3. The third section is divided into two tabs:

   The **Servers** tab provides detailed information on each ExaGrid Server in the site:

   - Available - How much space is available to the next backup and how much space is available for retention.
   - Allocation - How much space is allocated to the Landing Zone Tier and how much space is allocated to the Retention Tier.

- Consumption - How much space is being consumed by data waiting to be purged, how much space is being consumed by data replicated to an ExaGrid Server, and how much space is being consumed by primary shares that reside on an ExaGrid Server.

The **Shares** tab provides high-level information about the primary shares that reside on the ExaGrid Site.

ExaGrid Server Summary Report

The ExaGrid Server Summary report is the default report displayed when you select an ExaGrid Sever from the left-hand navigation.

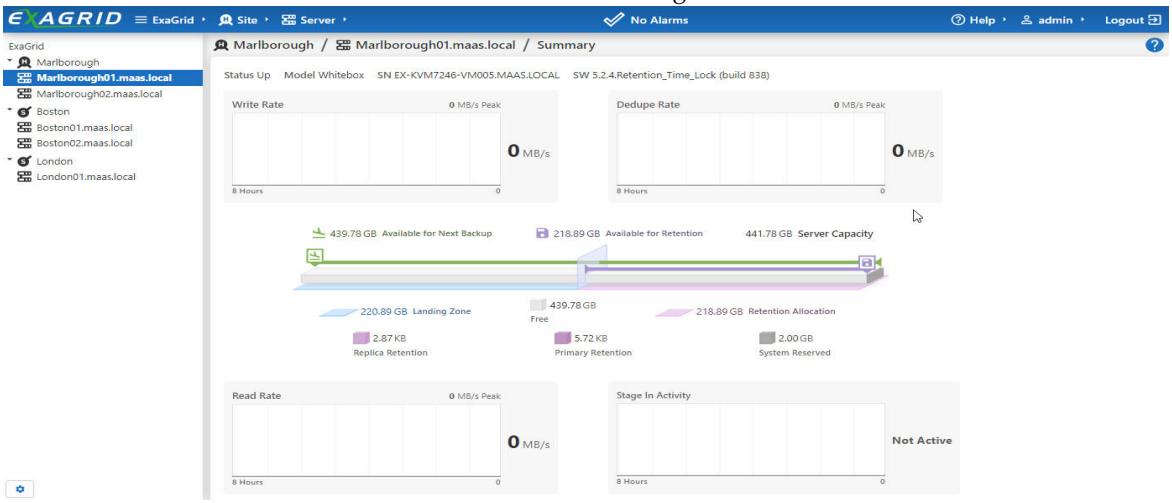

The ExaGrid Server-level Summary Report is divided in to four sections:

1. The first and top-most section allows you to view the ExaGrid Server's:
   - Status
   - Model number
   - Serial number
   - The version of ExaGrid Software currently installed
2. The second section of the report provides a graphical representation of the:
   - Write Rate - How fast data is being written to the ExaGrid Server
   - Dedupe Rate - How fast data is deduplicating on the ExaGrid Server.
3. The third section of the report provides a graphical representation of how space has been allocated and how much space is available to the ExaGrid Server.
4. The fourth section provides a graphical representation of the:

- Read Rate - How fast data is being read from the ExaGrid Server
- Stage-in activity on the ExaGrid Server.

Detailed Reports

ExaGrid provides detailed reports on ExaGrid hardware, software, data, deduplication, and replication.

Which reports, and how much detail is reported varies by what's selected in the left-hand navigation tree.

For example: If an ExaGrid Server is selected, the Replication In Progress report will only show the replication details for that ExaGrid Server. If an ExaGrid Site is selected, a Replication In Progress report will show replication details for all ExaGrid Servers in the ExaGrid Site.

The *ExaGrid Online Help* describes how to use each report and how to interpret results.

Monitoring Your ExaGrid System

As detailed in “ExaGrid Reports” on page 91 the ExaGrid system allows you to manually log in and view the status of your ExaGrid System.

In addition, the ExaGrid System also provides several automated tools that monitor and report status.

SMTP Email Notifications

You can configure the ExaGrid System to automatically send you or others in your organization an email when the ExaGrid System generates an event.

You can also configure the ExaGrid System to send a daily status report. The status report consists of a summary of the ExaGrid System’s Landing Zone Tier, Retention Tier, and optionally detailed backup information. The daily status report can be sent at a specific hour of the day, to several individuals in your organization, and can be sent in HTML, or text format. Setting the format to NONE will disable sending the status report to the associated email address.

To display the SMTP Email configuration page:

1. From the left-hand navigation tree, select an ExaGrid Site.
2. From the top, main menu, click **Site**, and all menu options for the site are displayed.
3. Under **Configuration**, select **SMTP Email,** and the SMTP page is displayed.

The online help has detailed instructions on how to setup SMTP Email alerting.

SNMP Traps

You can also configure your ExaGrid System to send alerts about events in the ExaGrid System from SNMP traps.

To display the SNMP Trap configuration page:

1. From the left-hand navigation tree, select an ExaGrid Site.
2. From the top, main menu, click **Site**, and all menu options for the site are displayed.
3. Under **Configuration**, select **SNMP Traps,** and the SNMP Traps page is displayed.

The online help has detailed instructions on how to setup SNMP Trap alerting.

Configure your MIB browser to use SNMP V2.

If you plan to use an SNMP management application, such as Hewlett-Packard's OpenView, you can download the ExaGrid Management Information Base (MIB):

1. In the ExaGrid top, main menu click **Help**.
2. From the drop down menu select **Online Library**. The ExaGrid Online Library is displayed.
3. Expand the **ExaGrid Documentation** option.
4. Scroll down and right-click on **ExaGrid Grid SNMP Management Information Base (MIB)** and save the MIB file to your required location.

Status Monitor

Your ExaGrid System has a "phone home" capability. The Status Monitor allows you to configure your ExaGrid System to send system information back to ExaGrid Support.

Reports are ExaGrid Site specific. Customers with multiple ExaGrid Sites can set up status monitoring for each of their ExaGrid Sites.

ExaGrid Support will automatically run tools against the report. Potential issues can be identified long before they become a problem.

You can send reports to ExaGrid from FTP or email.

To configure the Status Monitor for an ExaGrid Site:

1. From the left-hand navigation tree, select an ExaGrid Site.
2. From the top, main menu select **Site**, and all menu options for the site are displayed.
3. Under **Configuration**, select **Status Monitor.** The Status Monitor page is displayed.

ExaGrid MIB and SNMP Objects

The ExaGrid MIB is made up of industry standard MIB and ExaGrid specific MIB objects.

A copy of the ExaGrid MIB file can be found in the ExaGrid Online Library. To save a copy of the ExaGrid MIB file:

1. From the top, main menu, click **Help**.
2. From the drop down menu, click **Online Library**. The online library is displayed.
3. expand the ExaGrid Documentation option.
4. Scroll down the library entries and right click on the **ExaGrid MIB** file.
5. From the pop up menus, click **Save as...** and save the MIB file with a name and to a location as needed

---

Note – You must enable SNMP on each ExaGrid Server you want to monitor.

To enable SNMP on an ExaGrid Server:

1. From the left-hand navigation tree, select an ExaGrid Server
2. From the top, main menu click **Server** and all menu options for the server are displayed.
3. Under **Configuration**, select **SNMP Identity Configuration** and the ExaGrid SNMP Server Identification page is displayed:

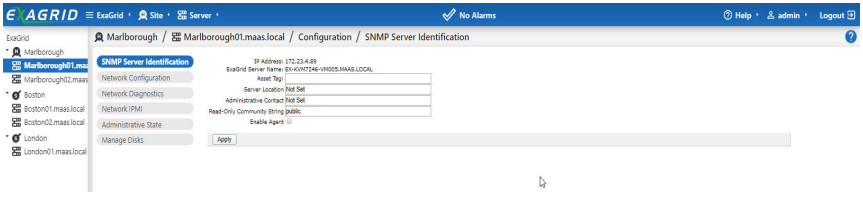

4. On the Identification page configure the following SNMP fields:

Field	Description
Asset Tag	A name or text string that allows you to easily identify the server. This tag appears in several ExaGrid reports
Server location	The MIB-2 system group sysLocation
Administrative Contact	The MIB-2 system group sysContact
Read-Only Community String	The default is set to Public. This is a password and should be set in accordance with your corporate security requirements.
Enable Agent	Must be checked to enable SNMP reporting for the ExaGrid Server.

5. Click **Apply** and the SNMP settings are applied to the selected ExaGrid Server.

Standard MIB Objects supported by ExaGrid

Caution – Built as a dedicated appliances, ExaGrid Servers often run at very high consumption levels for CPU, memory, and storage. Applying alerts to the ExaGrid SNMP object values as if they are general-purpose servers may produce false alerts/alarms. ExaGrid Servers are sized appropriately for the backup and retention specified as part of the ExaGrid sales/sizing process, and since recent backups are cached in the ExaGrid Landing Zone Tier, disk consumption SNMP object values may appear much larger than for general-purpose servers.

The ExaGrid MIB supports the following “standard” MIB objects:

- SNMPv2-MIB sysDescr as defined in RFC 3418.
- Interface MIB module defined in RFC 1213. This includes 64-bit counters useful for network utilization calculations - see below.

- Host Resources MIB defined in RFC 2790
  - The hrStorage table is implemented. This table includes memory and storage objects. See **Caution** above.

Note – This table is an array of hStorageEntry and is numerically indexed - as is typical for a collection of SNMP MIB objects. The indexes may change between releases. ExaGrid monitoring using SNMP must first find the appropriate hrStorageEntry index via hrStorageDescr and use values appropriately.

  - The Host Resources hrProcessorLoad MIB object is implemented for all processors. This MIB object contains: The average, over the last minute, of the percentage of time that this processor was not idle. See **Caution** above.

Example hrStorage MIB Objects

hrStorageDescr	hrStorageAllocationUnits	hrStorageSize hrStorageUsed
Physical memory	INTEGER: 1024	Varies by ExaGrid model
/home1	INTEGER: 16384	Combined Landing Zone Tier and Retention Tier in 16K blocks.
/home1/shares/<utilName>	INTEGER: 4096	Size of utility share in 4K blocks.

ExaGrid Specific MIB Objects

ExaGrid MIB objects are made up of WholeGigaBytes and FractionalGigabytes.

A 32 bit gauge does not hold values large enough to represent the bytes counts found in an ExaGrid System. To accommodate the 32 bit limitation each data item has been broken into two 32 bit counters, one for the “whole” part and one for the “fractional” part.

The “whole” part contains a scaled truncated value representing GigaBytes. The fractional part contains the remainder in bytes.

For example, to determine the byte count for a Landing Zone Tier objects, multiply the egLandingSpaceConfiguredWholeGigabytes value by 1,000,000,000 (10^9) and then add the egLandingSpaceConfiguredFractionalGigabytes.

ExaGrid Object Tree

The ExaGrid MIB specifies the following object tree:

exagridServerData (branch)

- exagridLandingSpace (branch)
  - egLandingSpaceConfiguredWholeGigabytes (32 bit Gauge)
  - egLandingSpaceConfiguredFractionalGigabytes (32 bit Gauge)
  - egLandingSpaceAvailableWholeGigabytes (32 bit Gauge)
  - egLandingSpaceAvailableFractionalGigabytes (32 bit Gauge)
- exagridRetentionSpace (branch)
  - egRetentionSpaceConfiguredWholeGigabytes (32 bit Gauge)
  - egRetentionSpaceConfiguredFractionalGigabytes (32 bit Gauge)
  - egRetentionSpaceAvailableWholeGigabytes (32 bit Gauge)
  - egRetentionSpaceAvailableFractionalGigabytes (32 bit Gauge)
- exagridDeduplicationRatio (branch)
  - egBackupDataAvailableWholeGigabytes (32 bit Gauge)
  - egBackupDataAvailableFractionalGigabytes (32 bit Gauge)
  - egBackupDataSpaceConsumedWholeGigabytes (32 bit Gauge)
  - egBackupDataSpaceConsumedFractionalGigabytes (32 bit Gauge)
- exagridPendingDeduplication (branch)
  - egPendingDeduplicationWholeGigabytes (32 bit Gauge)
  - egPendingDeduplicationFractionalGigabytes (32 bit Gauge)
  - egPendingDeduplicationAge (timeticks)
- exagridPendingReplication (branch)
  - egPendingReplicationWholeGigabytes (32 bit Gauge)
  - egPendingReplicationFractionalGigabytes (32 bit Gauge)
  - egPendingReplicationAge (timeticks)
- exagridServerStatus (branch)
  - egServerAlarmState (integer)

1. No alarms
2. One or more minor alarms outstanding (warning/yellow)
3. One or more major alarms outstanding (error/red)

- exagridRates (branch)
  - egServerReadRate (32 bit Gauge) - read from ExaGrid Server

- egServerWriteRateMegabytesSec (32 bit Gauge) - written to ExaGrid Server
- egServerDedupRateMegabytesSec (32 big Gauge) - rate of deduplication

How to Calculate

- Retention Capacity

  The amount of unused Retention Tier available is found by querying the egRetentionSpaceAvailableWholeGigabytes and egRetentionSpaceAvailableFractionalGigabytes MIB objects, scaling the whole portion and then adding the fractional portion.

- Deduplication Ratios

  Query all 4 objects in the exagridDeduplicationRatio group, scale the returned values to 64 bits and then convert available/consumed into a ratio.

- CPU utilization

  Average for the most recent minute CPU utilization is directly available as the hrProcessorLoad MIB object.

- Memory utilization

  Point in time memory utilization is directly available as the hrStorageUsed column in the appropriate row in the hrStorage table. The row index can change with each SNMP agent restart, so the hrStorageType object must be queried to determine the correct index to use.

- Network Utilization

  Several MIB objects supplied by the ifXTable in the ifMIB can assist in determining Network Utilization. Of particular interest are: ifHCInOctets and ifHCOutOctets. ifHCInOctets is defined as: The total number of octets received on the interface, including framing characters. This object is a 64-bit version of ifInOctets"

Note – Cisco has published an article titled *How to calculate bandwidth utilization using SNMP*. It is available at:

http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a008009496e.shtml

This article describes how these counters can be used over time to determine network utilization. With full duplex media, two formulas are required, one for input and one for output. The ifInOctets and ifOutOctets high capacity counters should be polled keeping track of the time between polls.

The Formulas become:
input utilization = change in ifInOctets * 8 * 100 / pollSeconds * ifSpeed
output utilization = change in ifOutOctets * 8 * 100 / pollSeconds * ifSpeed

System Logging

The ExaGrid user interface allows you to specify one or more external syslog servers to which ExaGrid system logging information from each ExaGrid server will be sent.

Once syslogs are received by a syslog server, they can be displayed, filtered and searched by the syslog server software.

Note – When System Logging is turned on, only events logged after it is turned on are sent to the specified syslog server(s).

The ExaGrid System compresses log files and retains them on the ExaGrid Server(s) for 6 months.

The ExaGrid System uses its event types to create syslogs. The following table shows how each event type maps into Syslog IDs:

Event Type	Syslog ID
Error	Error
Warning	Warning
Info	Informational
Audit	Informational

At a high-level, the events sent to the syslog include, but are not limited to:

- Management of ExaGrid Sites and ExaGrid Servers from the ExaGrid browser-based user interface. Including such things as: log in/log out, share creation/deleting, users/roles, server and site configuration, and server and site issues.
- SSH connections from outside of an ExaGrid System
- SMB share mounts
- NFS share mounts
- OST share mounts (or similarly-scoped event)
- Firewall events (packets being blocked)

Where applicable, events include the following information:

- The remote host by IP address and name
- The username of the login account involved
- The resource in question: ExaGrid Site, ExaGrid Server, ExaGrid Share, etc.

The following example shows ExaGrid syslog information being displayed by a third-party application:

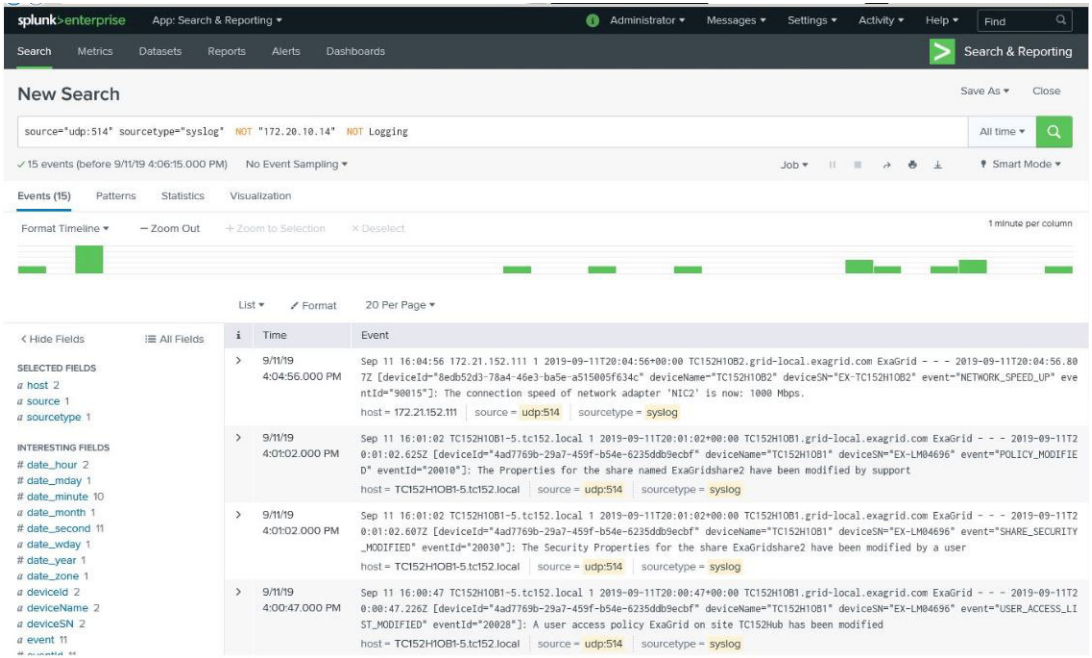

As can be seen in this example, the event, device name, host, etc. fields are included in the syslog.

Accessing the Syslog Management Page

To access the ExaGrid Site's Syslog Management page:

1. In the left-hand navigation tree, select the ExaGrid Site.
2. In the top, main menu, click **Site**, and all menu options for the site are displayed.
3. Under **Configuration**, select **Syslog Servers** and the Syslog Management page is displayed.
4. Under **Define Syslog Servers**, use the controls to:
   - **Delete** a syslog server from the list of syslog servers
   - **Edit** the details of an existing syslog server
   - **Add** a syslog server
   - **Refresh** the list of syslog severs

Adding Syslog Servers to an ExaGrid Site

To add a syslog server to an ExaGrid Site:

1. In the left-hand navigation tree, select the ExaGrid Site.
2. In the top, main menu, click **Site**, and all menu options for the site are displayed.
3. Under **Configuration**, select **Syslog Servers** and the Syslog Management page is displayed.

Note – The syslogs for all ExaGrid Servers in the selected ExaGrid Site will be sent to the specified syslog server.

---

4. Under **Define Syslog Servers**, click the **+** (Add) control and the Syslog Server Configuration Editor is displayed:

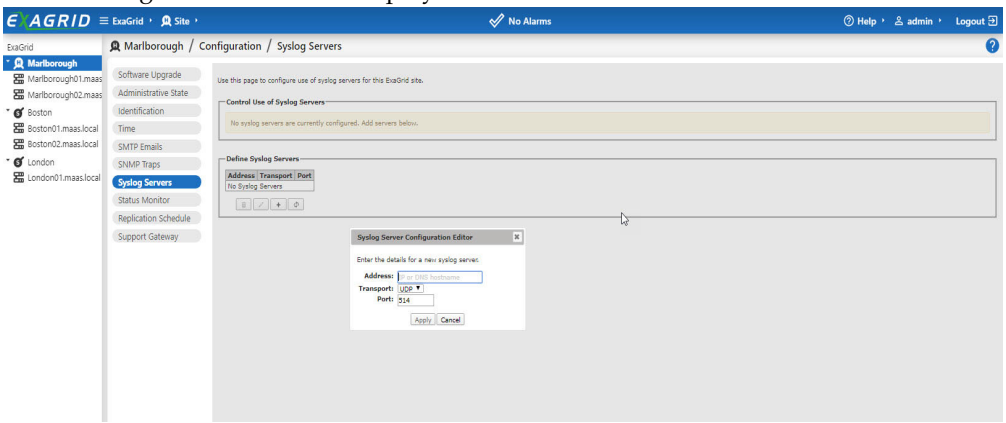

5. In the space provided enter the IP address of the syslog server.
6. From the **Transport** drop down menu select either **UDP** or **TCP**.
7. Review the default port number. The default port number is determined by the selected Transport: 514 for UPD and 601 for TCP.
8. Click **Apply** and the new syslog server is added to the list of syslog servers receiving syslogs from the selected ExaGrid Server.
9. Repeat the process as needed.

ExaGrid System Maintenance

Management and maintenance tasks for your ExaGrid System include:

1. Shutting down and restarting ExaGrid Servers
2. Adding additional ExaGrid Servers
3. Software upgrade

Shutting Down and Restarting an ExaGrid Server

Caution – NEVER shut down or restart an ExaGrid Server by cutting power or pressing the power or restart buttons on the front panel. ExaGrid Servers must be shutdown or restarted from the ExaGrid user interface.

To shut down an ExaGrid Server:

1. In the left-hand navigation tree, select the ExaGrid Server to be shut down.
2. From the top, main menu, click **Server** and all menu options for the server are displayed.
3. Under **Configuration**, click **Administrative State**. The Administrative State page is displayed.
4. Click **Shutdown**.

To restart an ExaGrid Server:

1. In the left-hand navigation tree, select the ExaGrid Server to be shut down.

2. From the top, main menu, click **Server** and all menu options for the server are displayed.
3. Under **Configuration**, click **Administrative State**. The Administrative State page is displayed.
4. Click **Restart**.

Software Upgrade

Software upgrades are available through your ExaGrid Customer Support Representative. Upgrades can be downloaded from the ExaGrid Support web site, or a copy can be sent to you.

To view the Software Upgrade page:

1. In the left-hand navigation tree, select an ExaGrid Site.
2. In the top, main menu, click Site, and all menu options for the site are displayed.
3. Under **Configuration**, select **Software Upgrade**. The Software Upgrade page is displayed.

From the software upgrade page you can:

- Determine the current version and status of ExaGrid Software across the selected ExaGrid site.
- Review the list of upgrade packages that have been uploaded and currently available to be installed.
- Enter a path, or browse to, an upgrade package to be uploaded to this ExaGrid site.
- Test/Retest an uploaded upgrade package.
- Apply an ExaGrid upgrade package to the selected site.
- Remove upgrade packages once they have been applied.

Adding Additional ExaGrid Servers

As your backup requirements grow, you may find it necessary to add additional ExaGrid Servers to your ExaGrid System.

Note – Step-by-step instructions on how to install and initialize a new ExaGrid Server can be found in the *Assembling and Initializing Your ExaGrid System* guide.

When you add an ExaGrid Server, you not only increase the amount of storage space in your ExaGrid System, you also increase the system's ability to take in and process backup data. More ExaGrid Servers mean more memory, CPUs and ports over which backup data can flow.

When a new ExaGrid Server is added, the ExaGrid System will automatically redistribute retained backup versions to balance the storage load across the ExaGrid Servers.

Note – Adding a new server will not automatically move (migrate) shares to a new server. Share migration is a manual process.

To use the additional ports/storage on the new ExaGrid Server(s):

1. Evaluate existing shares and the jobs writing to them to identify which shares, if moved (migrated) to the new ExaGrid Server, would improve your backup throughput.
2. Migrate the share(s) to the new ExaGrid Server(s).
3. Redirect backup job(s) to the new share location.

The following steps use ExaGrid's global deduplication capability to utilize the additional ports/storage n the new ExaGrid Server(s):

1. Create new shares on the new ExaGrid Server(s).
2. Add the new shares to the backup application
3. Evaluate existing shares and the jobs writing to them to identify which jobs can be changed to point to the new shares. Or,
4. Evaluate existing jobs to identify which client(s) in the job(s) can be moved to new jobs using the new shares.

ExaGrid 5.0 and later has global deduplication. Migrated moved jobs/clients will be deduplicated against their previous share/location.

Note – Consult the ExaGrid System's on line help for instructions on how to migrate a share.

Hub Site and Spoke Site Systems

ExaGrid Systems employ a hub and spoke architecture. An ExaGrid System has a single hub site and can have multiple spoke sites. A single site ExaGrid System is usually initialized as a hub site.

In multi-site systems, the hub site is displayed in the top-most position of the user interface's navigation tree.

When initializing ExaGrid Servers for a multi-site system, once both the hub and spoke sites are initialized, you can add the spoke site(s) to the parent hub site's system.

In most cases, the hub and its spokes will be identified (but not joined) during the initial ExaGrid Server initialization process. However, stand alone single-site systems, are usually defined as hubs sites. Such sites can be reconfigured as a spoke site and then added to a hub site's system. See "Reconfiguring a Site's Topology" on page 30. for details.

Reconfiguring a Site's Topology

Stand alone single site ExaGrid Systems are typically configured as hub sites. Before a stand alone single-site ExaGrid System can join another ExaGrid System, it must become a spoke site.

To reconfigure a stand alone single-site ExaGrid System to a spoke site:

1. Log into an ExaGrid Server on the spoke site to be reconfigured.
2. In the left-hand navigation tree, click **ExaGrid**.
3. In the top, main menu, Click **ExaGrid**, and all menu options for the ExaGrid are displayed.
4. Under **Configure**, click **Network Topology**. The Change Topology page is displayed.
5. Click **Make Spoke**. The site will be re-designated as a spoke.

Once reconfigured, the site can be added to another ExaGrid System as a spoke. For details see "Adding a Spoke Site" on page 113.

Adding a Spoke Site

To add a (physical or cloud) spoke site to an ExaGrid Hub:

1. Log into the spoke site to be added as a user assigned the Admin role.
2. In the left-hand navigation tree, select the spoke site.
3. In the top, main menu, click **Site**, and all menu options for the site are displayed.
4. Under **Security**, select **Site Administration Key**, and the Site Administration key page is displayed:

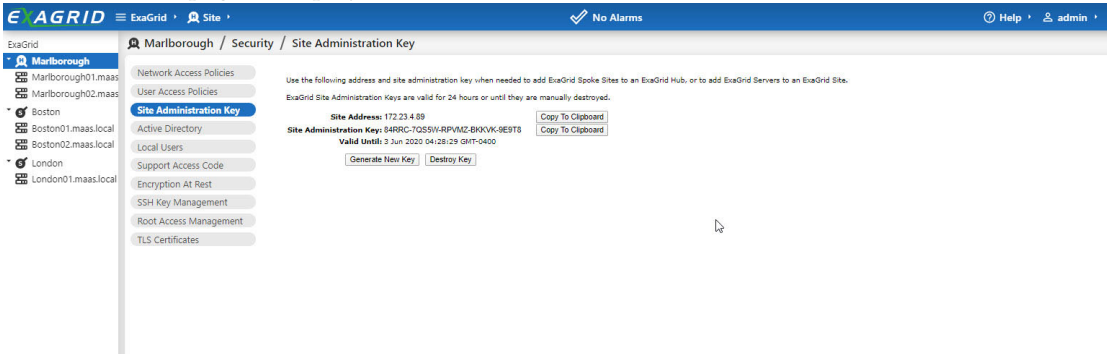

5. In a separate browser window or tab, log into the ExaGrid Hub site to which the spoke will be added. You must log in as a user who is assigned the Admin role.
6. In the left-hand navigation tree select the ExaGrid.
7. In the top, main menu click **ExaGrid** and the menu items for ExaGrid are displayed.

8. Under **Configure**, select **Add Spoke**, and the Add a New Spoke page is displayed:

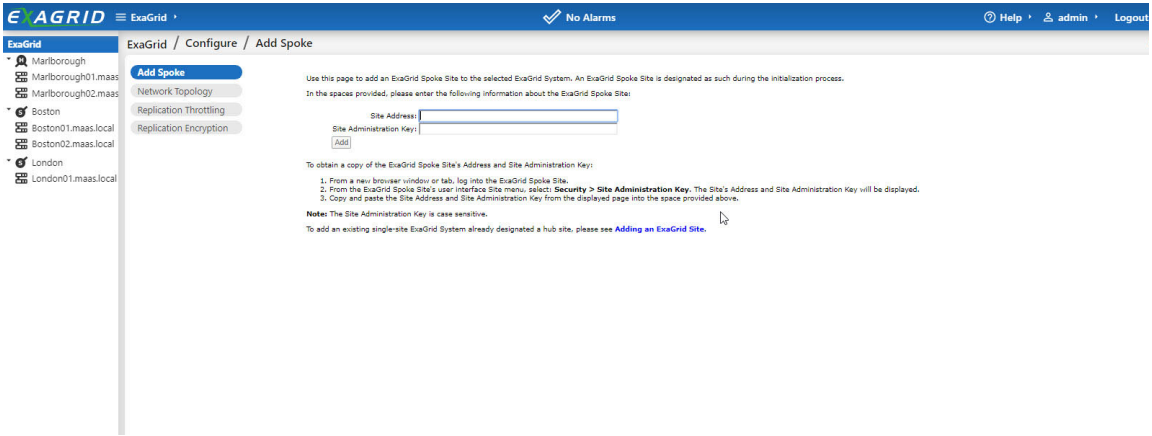

9. Use the Copy to Clipboard buttons in the Spoke's browser window to copy and paste the Site Address and Site Administration Key from the spoke site's browser window or tab in to the spaces provided in hub site's Add a New Spoke page.
10. In the hub site's Add a New Spoke page, click **Add**, and the spoke site will be added to the hub site.
11. Create Network and User Access Policies for the new Spoke by selecting the new Spoke and using the **Security** menus.

Note – Network and User Access Policies are kept separately on each ExaGrid Site.

Defining NTP Servers

The ExaGrid Hub Site is set at the time of manufacture to use the public NTP server pool to synchronize time. ExaGrid Spoke Sites are configured by default to use the ExaGrid Hub Site as their NTP server(s).

You can make changes to the ExaGrid Hub Site's factory defaults as needed.

Caution – ExaGrid Spoke sites that use an NTP server that differs from their ExaGrid Hub Site can experience unexpected results.

To display an ExaGrid Site's NTP Settings:

1. In the left-hand navigation tree, select the ExaGrid Site.

2. In the top, main menu, click **Site**, and all menu options for the site are displayed.
3. Under **Configuration**, select **Time** and the NTP settings page is displayed:

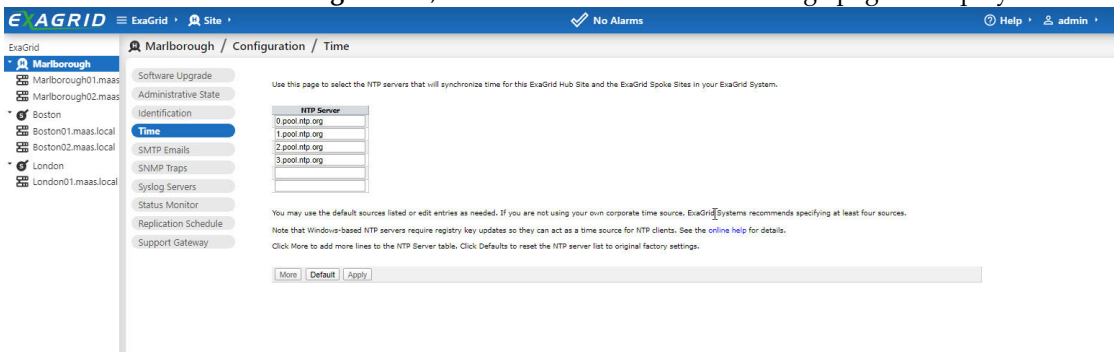

NTP Settings page for an ExaGrid Hub Site

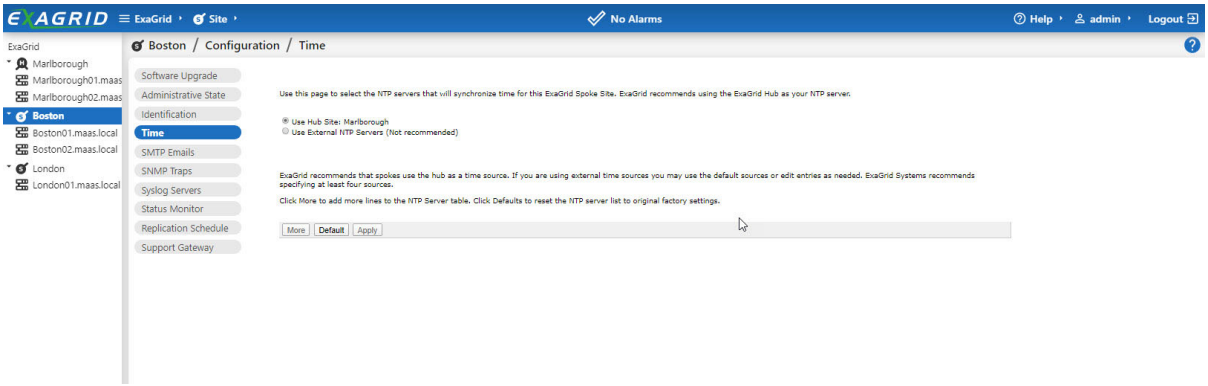

NTP Settings page for an ExaGrid Spoke Site

Windows® NTP Servers

If you use a Microsoft® Windows® server as an NTP server, that server must be running Windows 2008R1 or later.

To enable NTP Service on Windows:

1. On the Windows NTP server, from the **Start** menu, start the Run application:
2. In the space provided enter **regedit** and click **OK**. The Registry Editor is displayed.
3. Locate and then expand the following registry entry:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config\
```

4. In the right pane, right-click **AnnounceFlags**, then click **Modify**.
5. In the Edit DWORD Value dialog box, under Value data, enter **5**, and then click **OK**.
6. Locate and then expand the following registry subkey:

   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\
   TimeProviders\NtpServer\
   ```

7. In the right pane, right-click **Enabled**, and then click **Modify**.
8. In the Edit DWORD Value dialog box, enter **1** under Value data, and then click **OK**.
9. Exit the Registry Editor.
10. At the command prompt, run the following command to restart the Windows:

    ```
    net stop w32time && net start w32time
    ```

Troubleshooting

In the unlikely event that your ExaGrid System experiences a unexpected issue, use this section to help troubleshoot the issue.

Monitoring Status of Ongoing Operations

The ExaGrid System provides a window in to the background operations of your ExaGrid System. The Status Monitor is useful to determine the status of an operation. For example:

- During a reboot of an ExaGrid Server
- An ExaGrid Share migration, depending on how much data is being moved can take some time to complete.
- Initial replication of data.
- RAID rebuild

In addition, in the unlikely event that your ExaGrid user interface becomes unresponsive, you can use the status monitor to determine the status of your ExaGrid System.

To display the status of a reboot operation, refresh your browser until the Monitor Status page is displayed.

To check on the progress of non-reboot operations:

1. In a browser's address space, enter the IP address or DNS name of the ExaGrid of interest.
2. Follow the address with the port of :9080

3. Press return and the ExaGrid Status Monitor is displayed:

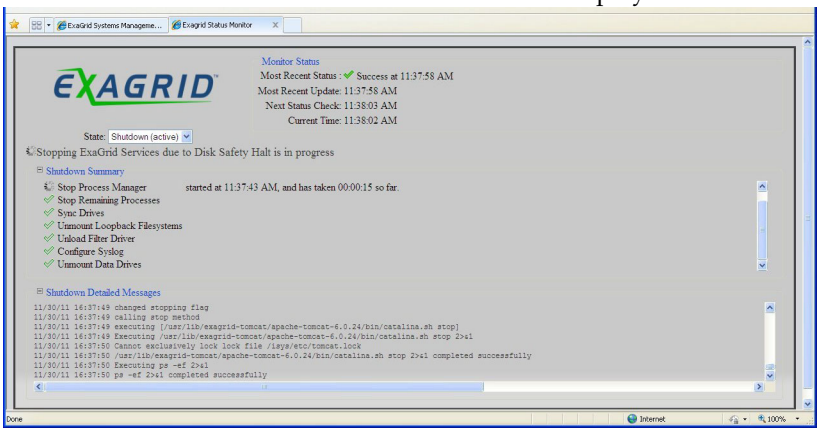

Replacing Disks

Each ExaGrid Server is RAID 6 compliant with a hot spare. This means that an ExaGrid Server can lose up to three disks and continue functioning. However, to ensure optimum performance, failed disks should be replaced as soon as possible. Replace the disk in accordance with the maintenance plan that came with your ExaGrid Server.

If an ExaGrid Server disk fails, an alert will be generated indicating that the failure has occurred and which disk has failed. The alert will be sent to all users who normally receive email alerts from your ExaGrid System. The alert message will also appear in the Events report in the ExaGrid System's user interface.

ExaGrid Server disk failure alerts contain the disk number that failed. The Disk Management page in the ExaGrid user interface will identify the physical location of the failed disk. See "Disk Numbering" on page 120. In addition, ExaGrid Customer Support engineers have disk layout maps for all ExaGrid models and can help identify the location of a failing disk.

Removing Disks

Before physically removing a disk from an ExaGrid System you must first, from the ExaGrid user interface, prepare the disk for removal:

1. In the left-hand navigation tree select the ExaGrid Server that contains the disk to be removed.
2. From the top, main menu click **Server**, and all menu options for the server are displayed.
3. Under **Configuration**, select **Manage Disks** and the Disk Management page is displayed.
4. Follow the instructions on the page to perform a software remove and get status of other disks in the ExaGrid Server.

Disk Safety Halt Protection

ExaGrid Servers use RAID 6 with a hot spare. Disk Safety Halt is initiated if an ExaGrid Server's RAID is degraded because:

- Two drives are in a failed state and there is no spare. The ExaGrid Server will automatically go into Disk Safety Halt mode. Once the failed drives are replaced, the ExaGrid Server's RAID will rebuild and will resume operation.

or

- Two drives fail - the second drive failed before the spare drive completes its rebuild. The ExaGrid Server will automatically go into Disk Safety Halt mode and when the spare drive completes rebuilding, the ExaGrid Server resumes operation.

While in Disk Safety Halt mode, an ExaGrid Server will not accept backups.

Although the primary ExaGrid user interface is not viewable during a disk safety halt, you can use ExaGrid's Status Monitor to view the status of the system's RAID and which disks are involved.

Disk Numbering

To view an ExaGrid Server's disk layout and status.

1. In the left-hand navigation tree select the ExaGrid Server that contains the disk to be removed.
2. From the top, main menu click **Server**, and all menu options for the server are displayed.
3. Under **Configuration**, select **Manage Disks** and the Disk Management page is displayed.

Some disk slots can be un-populated with blank covers.

Replacing Power Supplies

Depending upon the model, an ExaGrid Server chassis uses two or three power supplies.

In the event of a power supply failure, the ExaGrid Server can continue functioning. However, to maintain the safety margin found in having redundant power supplies, the failed power supply should be replaced as soon as possible. Replace the power supply in accordance with the maintenance plan that came with your ExaGrid Server.

If an ExaGrid Server power supply fails, a "90530" alert will be generated indicating that the failure has occurred. The alert will be sent to all users who normally receive email alerts from your ExaGrid Server. The alert message will also appear in the Events report in the ExaGrid System's user interface. In addition, the failed power supply's LED status indicator will change from green to amber.

ExaGrid Server power supply failure alerts use the following format:

```
Event ID: 90530
Severity/Type: Warning
Grid Name: Grid Name
Site Name: Your_Site
SiteGrid Name: name
Device Name: Server Name
Description: Power Supply Fault
Detail: Power supplies are indicating a fault condition
```

Replacing Fans

Depending upon the model, an ExaGrid Server chassis can use up to six customer replaceable fans.

In the event of a fan failure an event will be generated detailing which fan has failed. In such cases, contact ExaGrid Support and depending on your support maintenance contract, you will be sent a replacement fan, along with step-by-step instructions on how to replace the fan.

To replace a fan:

1. From the event message that identified the failed fan, record the fan's name and location.
2. Ensure backup jobs are not writing to the ExaGrid Server.

Caution – NEVER shut down or restart an ExaGrid Server by cutting power or pressing the power or restart buttons on the front panel. ExaGrid Servers must be shutdown or restarted from the ExaGrid user interface.

---

3. Shut down an ExaGrid Server:

1. In the left-hand navigation tree, select the ExaGrid Server to be shut down.
2. From the top, main menu, click **Server** and all menu options for the server are displayed.
3. Under **Configuration**, click **Administrative State**. The Administrative State page is displayed.
4. Click **Shutdown**.
5. Remove the ExaGrid Server from the rack.

Caution – At least two people are needed to lift the ExaGrid Server from the rack.

---

6. Remove the top cover.
7. Before continuing, properly connect an anti-static strap to yourself.

8. Use the following fan map to locate the fan identified in the fan failure event message:

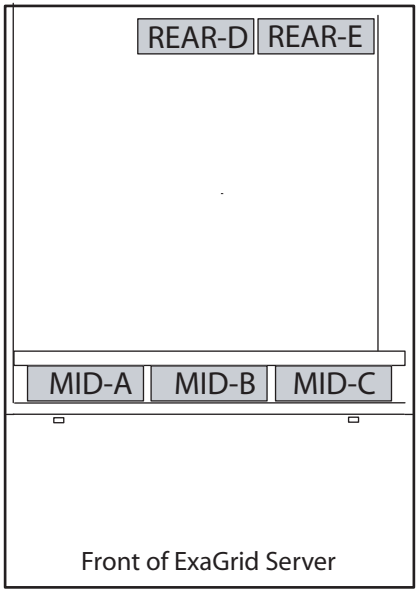

Note – The number of fans vary by ExaGrid Server model, however the naming convention and fan location is consistent across all models.

9. To remove the defective fan, press its side release latch and lift the fan out:

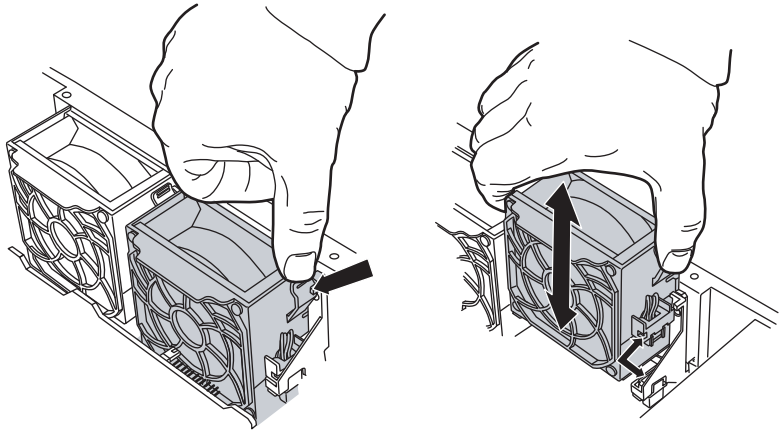

10. To insert the replacement fan, align the fan into the fan rails and gently push the fan in place until it is completely seated into its electrical connector.

11. Replace the top cover.
12. Re-rack the ExaGrid Server.
13. Reconnect cables.
14. Power on the ExaGrid Server.

Network Debugging Tools

ExaGrid Systems are appliances. Access to an ExaGrid System is through the ExaGrid System's web-based user interface.

In the unlikely event you need to debug your ExaGrid network settings, the ExaGrid System's user interface provides the following tools:

1. Ping
2. Traceroute
3. Network interface statistics
4. Network performance test

To use these tools:

1. In the navigation tree select either ExaGrid, an ExaGrid site or an ExaGrid Server.
2. In the top, main menu, click either ExaGrid, an Site or Server and all of the options for the selected component are displayed.
3. Under **Configure**, select **Network**. **Diagnostics,** and the network Diagnostics page is displayed.

Notifications

If you encounter issues with receiving email alerts or SNMP traps from your ExaGrid System:

- If your email server is relay enabled, a missing **Send from** password will cause the system to send all messages as relays. Relayed messages are often times treated as spam by spam filters and filtered out. If you must configure messages as relays, ensure that the recipient's spam filters allow relayed messages from this site.
- See "Network Debugging Tools" on page 123. to determine proper network connectivity for each ExaGrid Server.
- Double check the DNS definition for each ExaGrid Server:

1. In the left-hand navigation tree, select an ExaGrid Server.
2. In the top, main menu, click **Server**, and all menu options for the server are displayed.
3. Under **Configuration** select **Network Configuration** and the Network Configuration page for the selected server is displayed.
4. Ensure the DNS IP address is correct. If it is not, enter the correct IP address and click **OK**.

- For SNMP traps, ensure that you are using the latest ExaGrid MIB:

1. In the ExaGrid user interface's Main menu click **Help**.
2. From the drop down menu select **Online Library**. The ExaGrid Online Library is displayed.
3. Expand the ExaGrid Documentation option.
4. Scroll down and right-click on **ExaGrid Grid SNMP Management Information Base (MIB)**.
5. Save the MIB file to your required location.

Consumption Issues

In the unlikely event that your ExaGrid System begins experiencing rapid consumption of storage space:

Ensure the size of your backups and the length of time they are retained does not exceed your system's rating. When you purchased your ExaGrid System, your ExaGrid System Engineer sized your system according to your current and future backup requirements. If those requirements have changed, it is possible that you are now be exceeding your system's rated capabilities. Review your current backup sizes and retention periods.

Ensure your backup deduplication ratios are getting consistently better over time. Review the **Deduplication** reports. Each time a job runs, its deduplication ratios should improve.

The site-level Deduplication Summary report will give you a general overview of deduplication in your ExaGrid System. Use the Deduplication reports to gather more detailed information about your ExaGrid System's deduplication.

Landing Zone Tier Consumption

The ExaGrid System allows you to write your backup directly to disk, at the speed of disk, without any preprocessing. To achieve this, the ExaGrid System provides a dedicated Landing Zone Tier for landing backups. The size of the Landing Zone Tier varies by ExaGrid Server model. The larger the ExaGrid Server, the larger the Landing Zone Tier.

The Landing Zone Tier section of a report allows you to determine whether or not there is enough space available for the next backup to land. In most cases the space available will be 100% of the system's rating; even if a backup job was completed. This is because ExaGrid Servers can dynamically allocate space from their Retention Tier on an as needed basis. As long as your backup job size does not exceed your ExaGrid System's rating, your next backup will fit.

Note – Do not confuse Landing Zone Tier space with the total ExaGrid System space. Landing Zone Tier space size is determined by the ExaGrid Server model number. It is the maximum backup size for which the ExaGrid Server has been rated. Your ExaGrid Server has much more storage space than what is reported in the Landing Zone Tier section of this report.

When you purchased your ExaGrid System it was sized according to the amount of data to be retained and how long it was to be retained.

If the amount of data backed up to your ExaGrid System exceeds the amount for which it was originally sized, the Landing Zone Tier space will fill and begin to borrow space from the Retention Tier. When Retention Tier space is used as Landing Zone Tier space, the issue cannot be corrected until:

- The amount of data being backed up is reduced.
- Shares are migrated to a different ExaGrid Server.
- Additional ExaGrid Servers are added.

Landing Zone Tier used By Retention Tier

Under certain conditions, the Retention Tier can borrow space from the Landing Zone Tier. In such situations, the landing Zone Tier space available can be less than 100%.

When the Retention Tier space borrows space from the Landing Zone Tier space it is usually an indication that there has been a change to the amount of data being backed up and/or how long data is being retained. This can be an indicator that you have begun to out grow your ExaGrid System.

When Landing Zone Tier space is used as Retention Tier space, the issue cannot be corrected until one or more of the following things happen:

- The length of time backup jobs retention data.
- The amount of data being backed up is reduced.
- Shares are migrated to a different ExaGrid Server.
- Additional ExaGrid Servers are added.

When Landing Zone Tier is consumed by retained data, to determine if your next backup job will fit on the ExaGrid Server, compare the amount of Landing Zone Tier space available with the amount of space required to accommodate the next backup.

Retention Tier Consumption

Retention Tier is dedicated to retaining versions of your backups. The amount of Retention Tier space varies by ExaGrid Server model.

As backups land on your ExaGrid System and deduplicate, versions are put into the Retention Tier. Depending on how your system was initially sized, once 70 or 80 percent of the Retention Tier is consumed, consumption of the Retention Tier will cease because; as previous backups are deduplicated and put into the Retention Tier, your backup application, based on its retention schedule, will begin aging out earlier backups.

The amount of Retention Tier space consumed is determined by:

- Size of backups
- The number of backups retained
- How long backups are retained
- How well each backup deduplicates

If your backup requirements change and you begin to outgrow the ExaGrid Server, the Retention Tier space can begin using space from the Landing Zone Tier

In such cases you should evaluate the amount of data you are retaining and whether or not you need to retain less data or add capacity to your ExaGrid System.

Replication

In the unlikely event that your ExaGrid System experiences issues replicating deduplicated data to another site, use the following suggestions to troubleshoot the issue:

- See the ExaGrid "Network Debugging Tools" on page 123 to ensure that the necessary network bandwidth is available.
- Review the ExaGrid Site's replication schedule and ensure it provides a wide enough time window to replicate the ExaGrid Site's data.

To view an ExaGrid Site's replication schedule:

1. In the left-hand navigation tree, select an ExaGrid Site.
2. In the top, main menu, click **Site** and all menu options for the site are displayed.
3. Under **Configuration**, select **Replication Schedule** and the ExaGrid Site's replication schedule is displayed.

- Use the ExaGrid Server-level Replication Transfer Rate History report to identify daily replication trends. Look for changes in the replication rate. A change can indicate other network traffic is competing with the transfer of your deduplicated data.
- Use the Replication Transfer History report to identify when a backup job's data was available to be replicated.
- Use the Replication History report to identify how much data each job contributes to overall replication.

Performance Tuning

Your ExaGrid Server will accept data as fast as your backup server or media server can write it. Most backup application vendors provide performance tuning documentation. If you experience slower than expected writes to your ExaGrid System, ExaGrid recommends that you consult the backup vendor's documentation to get the most out of your backup application.

In general, ExaGrid has found that the following types of things, if addressed, will maximize backing up to an ExaGrid System:

- Backup stream concurrency is very important. Each ExaGrid server can handle 20 high-speed concurrent backup streams - this number is for guidance, it is not enforced. A higher number of lower-speed concurrent backup streams will still improve overall backup performance.
- Balancing the backup streams across ExaGrid servers maximizes backup performance. For example, if you have 5 ExaGrid servers, and your backup infrastructure is capable of creating 100 concurrent backup streams, configure the ExaGrid devices in your backup application(s) and your jobs to have each ExaGrid server receive 20 concurrent backup streams.
- Using your backup application(s) documentation and support organization, ensure backup servers and media servers have enough capabilities or proper settings to achieve a high number of concurrent backup streams.
- If your backup server or media server is equipped with a network card that has off loading capabilities, off load as much processing as possible to the network card. This will free up the server's CPU to process and write the backup faster.
- Review and adjust your backup server's or media server's network buffer size. Optimal settings will vary according to backup application and server type.
- NetBackup users should review their Shared memory settings (number and size of data buffers). Consult the NetBackup's Tuning guide that came with your NetBackup backup application.
- NetBackup users should review their client communications buffer size. When NetBackup clients are installed, buffers are set from 16K to 64K. ExaGrid has found that 128K or larger can enhance performance.
- NetBackup users should review their server and client communications buffer size and ensure that they are both set to the same value.

Administrator's Guide

Part Number: 210-0165-24

© 2003-2020 ExaGrid Systems

Printed in the U.S.A. All rights reserved.

No part of this document may be reproduced
or transmitted in any form or by any means without
prior written permission of ExaGrid Systems

.All rights reserved. Printed in U.S.A

DETAILED PRODUCT DESCRIPTION

ExaGrid Tiered Backup Storage

EXAGRID®

Tiered Backup Storage

Table of Contents

Executive Summary . . . 2
ExaGrid Basic Concept . . . 2
ExaGrid Technology Overview . . . 3
ExaGrid Appliance Models . . . 4
Scalable Architecture – Scale-Out . . . 5
Designed for Performance . . . 7
Data Deduplication . . . 7
Data Backup Performance . . . 7
Data Restore Performance . . . 8
Support for Heterogeneous Backup Applications and Utilities . . . 9
Advanced Backup Features . . . 10
Veritas OST . . . 10
Veeam Accelerated Data Mover . . . 11
Veeam SOBR . . . 12
Commvault . . . 13
Oracle RMAN Channels . . . 14
Reliability and Redundancy . . . 15
ExaGrid Retention Time-Lock for Ransomware Recovery . . . 16
RAID6 Internal Storage with Consistency Checking . . . 19
Flash-Backed RAID Cache . . . 19
Backup Data Checksums with Automatic Repair . . . 19
Deduplicated Metadata Transactional Consistency . . . 19
Internal Database Backups and Self-Describing Metadata . . . 19
Logging Filesystem . . . 19
Data Security . . . 20
Active Directory Support . . . 20
Periodic Assessments Using a Network Vulnerability Scanner . . . 20
Offsite Data Protection for Disaster Recovery . . . 21
Total Cost of Ownership . . . 24
Cost Effectiveness Up Front . . . 24
Cost Effectiveness as Data Grows . . . 25
Summary . . . 26
About ExaGrid . . . 28

Executive Summary

ExaGrid provides tiered backup storage with a unique disk-cache Landing Zone, long-term retention repository, and scale-out architecture. ExaGrid's Landing Zone provides for the fastest backups, restores, and instant VM recoveries. The disk-cache Landing Zone tier allows for the fastest backups and restores. The retention repository tier offers the lowest cost for long-term retention. ExaGrid's scale-out architecture includes full appliances and ensures a fixed-length backup window as data grows, eliminating expensive forklift upgrades and product obsolescence.

ExaGrid's Tiered Backup Storage makes backup better with:

- Fastest backups resulting in the shortest backup window
- Fastest restores
- Instant VM recoveries in seconds to minutes
- Fastest offsite tape copies
- A fixed length backup window as data grows, eliminating expensive future forklift upgrades
- Industry-leading deduplication ratios for storing long-term retention data for low-cost retention storage
- Deduplication efficiency across the entire system with global deduplication
- Fast, reliable, and up-to-date offsite disaster recovery using Adaptive Deduplication
- Linear scalability — scale-out architecture
- Low cost up front and over time as data grows
- No forklift upgrades or product obsolescence
- Retention Time-Lock to ensure recovery from a ransomware attack

ExaGrid Basic Concept

ExaGrid appliances work seamlessly with over 25 industry-leading backup applications and database utilities by presenting themselves as standard NAS shares (CIFS or NFS), a Veeam Data Mover target, or Veritas NetBackup OST target. ExaGrid appliances are easily integrated into existing backup environments, as illustrated below.

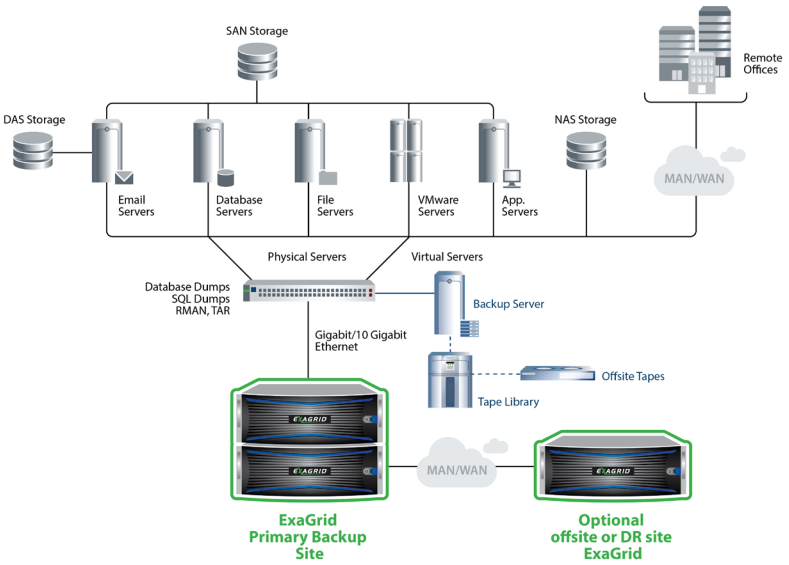

ExaGrid sits behind the existing backup server and replaces low-cost disk, inline deduplication appliances or tape backup storage, both onsite or offsite.

ExaGrid Technology Overview

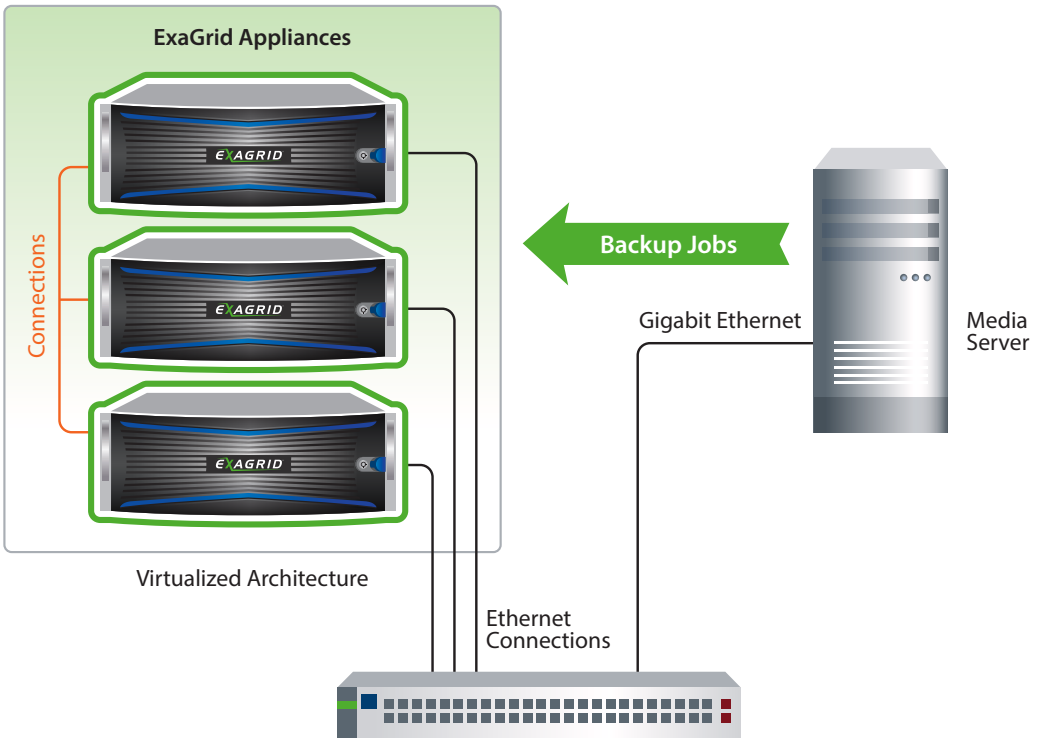

ExaGrid Appliances Connect to Form a Single Scale-Out System

ExaGrid appliances are comprised of processors, memory, RAID6, and a hot spare drive per appliance, using enterprise-class SAS drives and ExaGrid software. See the ***ExaGrid Technical Specifications*** data sheet for more details.

Each appliance plugs in and is virtualized into a shared system with a single user interface, global deduplication, and automatic load balancing. The media server is connected to the same network and sees the appliances as one or more NAS shares, Veeam Data Mover targets, or Veritas NetBackup OST targets. Since each appliance includes the appropriate amount of processor, memory, Landing Zone disk, deduplicated repository disk, and bandwidth for the rated data size, performance increases as more appliances are added to the system.

ExaGrid Appliance Models

The product line's multiple appliance models can be combined into a single scale-out system with a configuration of up to 32 appliances in a single system for a total of 6.14PB raw capacity, 5.37PB of usable capacity, and supports a full backup of up to 2.69PB, total backup logical data capacity of 43PB.

ExaGrid Model	Raw Capacity (TB)	Usable Capacity (TB)	Capacity for Weekly Full (TB)	Capacity for 16 Weekly Fulls (TB)	Max Backup Thruput (TB/hr)
ExaGrid Appliances					
EX6	24	12	6	96	1.74
EX10	32	20	10	160	2.40
EX18	48	36	18	288	3.60
EX27	72	54	27	432	6.09
EX36	96	72	36	576	7.78
EX52	128	104	52	832	10.87
EX84	192	168	84	1344	15.25
ExaGrid Appliances with Encryption					
EX10-SEC	32	20	10	160	2.40
EX18-SEC	48	36	18	288	3.60
EX27-SEC	72	54	27	432	6.09
EX36-SEC	96	72	36	576	7.78
EX52-SEC	128	104	52	832	10.87
EX84-SEC	192	168	84	1344	15.25
	TB / PB	**TB / PB**	**TB / PB**	**PB**	**TB/hr**
Example System Configurations					
EX168-G	384TB	336TB	168TB	2.68	30.5
EX504-G	1.15PB	1PB	504TB	8.06	91.5
EX1008-G	2.3PB	2.01PB	1PB	16.12	183
EX1344-G	3.07PB	2.68PB	1.34PB	21.5	244
EX1680-G	3.84PB	3.36PB	1.68PB	26.88	305
EX2100-G	4.8PB	4.2PB	2.1PB	33.6	381.25
EX2688-G	6.14PB	5.37PB	2.69PB	43	488

Scales to 32 appliances in a single scale-out system
(2.69PB full backup @ 488TB/hr. ingest rate)

Scalable Architecture – Scale-Out

The typical business or organization is seeing data growth of 30% a year which doubles data every 2.5 years. This data growth presents challenges to IT to ensure their backup system can scale easily to support that growth. With scale-up disk backup solutions that have a front-end server/controller architecture and add just disk shelves as data grows, you start with a short backup window, but as data grows, the backup window expands because only disk is added but no additional deduplication processing resources. Eventually, the backup window expands to a point where you must replace the front-end server with a more powerful server via a costly "forklift upgrade."

In contrast, ExaGrid's scalable approach with global deduplication adds full servers—including memory, processor, and bandwidth as well as disk. The figure below shows the differences between how the two different architectures cope with data growth over time.

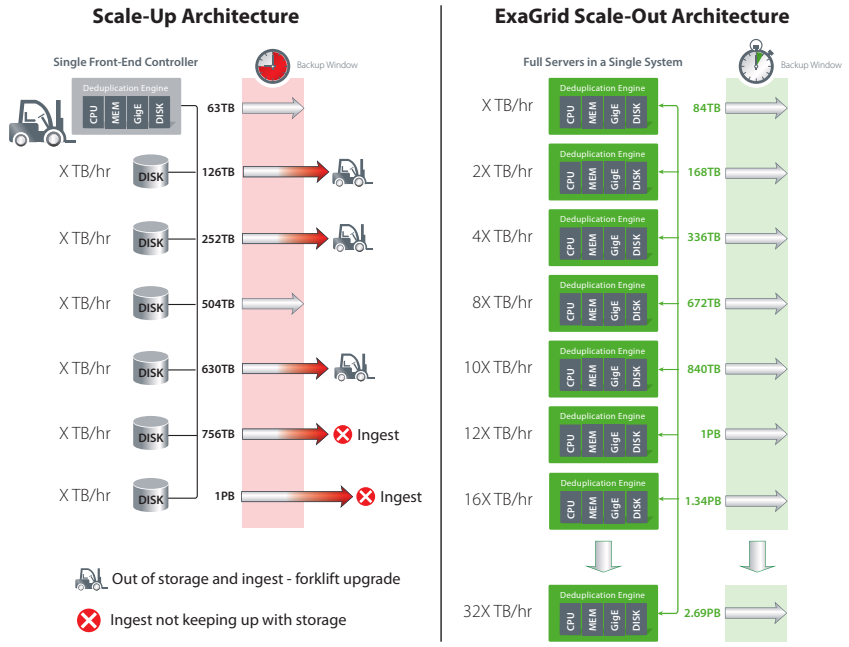

Scale-out vs Scale-up Backup Storage

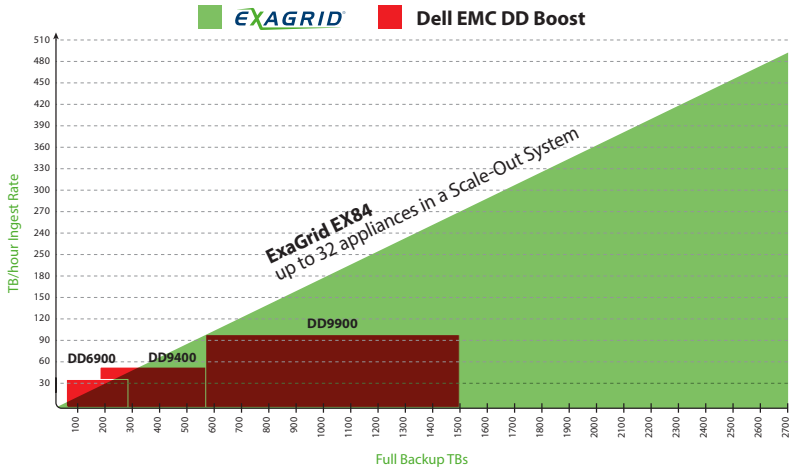

This unique approach provides the following benefits:

- **No Expansion of Backup Windows as Data Grows** – By adding full appliances, ExaGrid maintains consistently fast backup performance and a fixed length backup window as data increases.
- **Seamless Scalability with No Forklift Upgrades** – ExaGrid's scale-out approach avoids forklift upgrades by adding modular capacity with full servers to the existing system in full backup increments of 6, 10, 18, 27, 36, 52, or 84TB with up to 32 appliances in a single system that scales to 2.69PB full backup.
- **No Obsolescence of Previous Model Systems** – All ExaGrid appliances can be seamlessly added to any existing ExaGrid deployment and can be mixed and matched with older appliances of any size, thereby eliminating obsolescence. Adding a new ExaGrid appliance is as simple as plugging it in and pointing backup jobs at the appropriate target. If any model is no longer offered, the models still work in a scale-out system with older and newer models, maintenance and support is always offered and maintenance and support will not go up more than 3% per year allowing organizations to get their full lifetime use out of all appliances in the system.

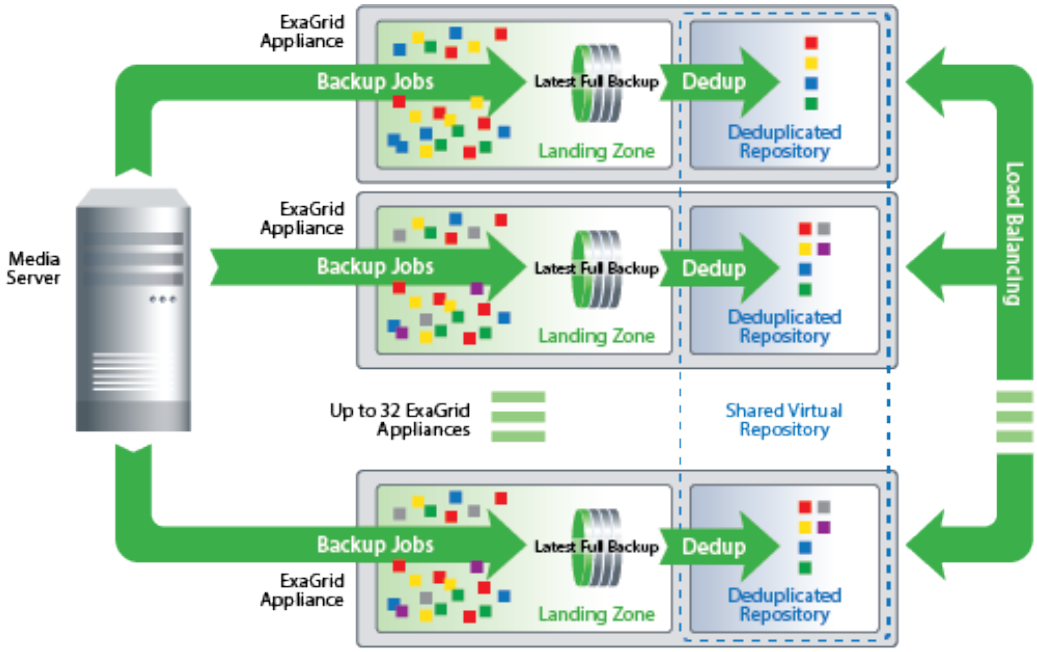

ExaGrid appliances include scalable computing software, which allows them to virtualize and share data storage capacity with one another (automatic load balancing across all repositories). This scalable system (shown above) can expand as data grows by adding appliances, providing up to 6.14PB raw capacity, 5.37PB usable capacity, and allowing full backups of up to 2.69PB in a single scale-out system. Multiple systems can be used at a single location, and up to 16 separate locations can be managed through a single user interface.

Designed for Performance

Data Deduplication

ExaGrid's innovative approach to data deduplication minimizes the amount of data to be stored by using zone-level data deduplication across all received backups.

ExaGrid's zone-level technology stores only the changed data at a granular level from backup to backup instead of storing full copies. ExaGrid uses zone stamps and similarity detection.

This unique approach reduces the disk space required by an average of 20:1 and from 10:1 up to 50:1 depending on data type, retention and backup rotation delivering unparalleled performance for the fastest backups and restores.

ExaGrid is the only vendor that allows backup applications such as Veeam and Commvault to keep their deduplication turned on and then ExaGrid further deduplicates the data for increased storage efficiency.

Data Backup Performance

The ExaGrid system performs deduplication in parallel with backups by using idle system cycles. Termed "adaptive deduplication," this approach keeps deduplication out of the backup path, resulting in faster backups and a shorter backup window. Adaptive Deduplication gives full system resources to backups for the shortest backup window and uses idle cycles to perform deduplication and offsite replication, allowing data to be replicated offsite for a strong offsite recovery point (RPO) for update to data for disaster recoveries.

ExaGrid supports data deduplication globally across all targets and all appliances in a scale-out system. Global deduplication ensures that all data is deduplicated regardless of the number of targets and/or the number of appliances in a system. In addition, global deduplication allows organizations the flexibility to redirect backup jobs to any target, on any appliance, at any time while maintaining data deduplication globally across the entire system.

Once the backup job is complete and off the network, the data is protected and immediately available to be restored. When deduplication and replication are complete, the data is ready for restore on the ExaGrid appliances at the disaster recovery site. ExaGrid uses Adaptive Deduplication that avoids the performance penalty of inline deduplication. ExaGrid deduplicates in parallel with backups coming in but performs no inline processing which allows for at least a 3X performance advantage over inline deduplication appliances. As shown in the diagram to the right, backup data is written directly from the backup server to ExaGrid's disk-cache Landing Zone at the highest possible rate with no inline processing to interfere, resulting in the shortest possible backup window.

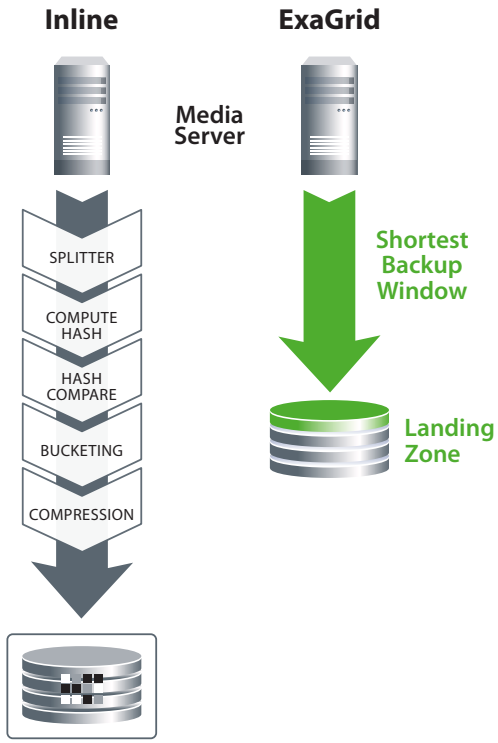

Data Restore Performance

Because ExaGrid's appliance allows each full backup to first land on the disk-cache Landing Zone before deduplication, the system maintains the most recent backup in its full, undeduplicated form for fast restores, Instant VM recoveries in seconds to minutes, and fast offsite tape copies. Since over 90% of restores and 100% of instant VM recoveries and tape copies are done from the most recent backup, this approach avoids the overhead incurred from "rehydrating" data during critical restores. As a result, restore, recovery, and copy times from an ExaGrid system are an order of magnitude faster than solutions that only store deduplicated data.

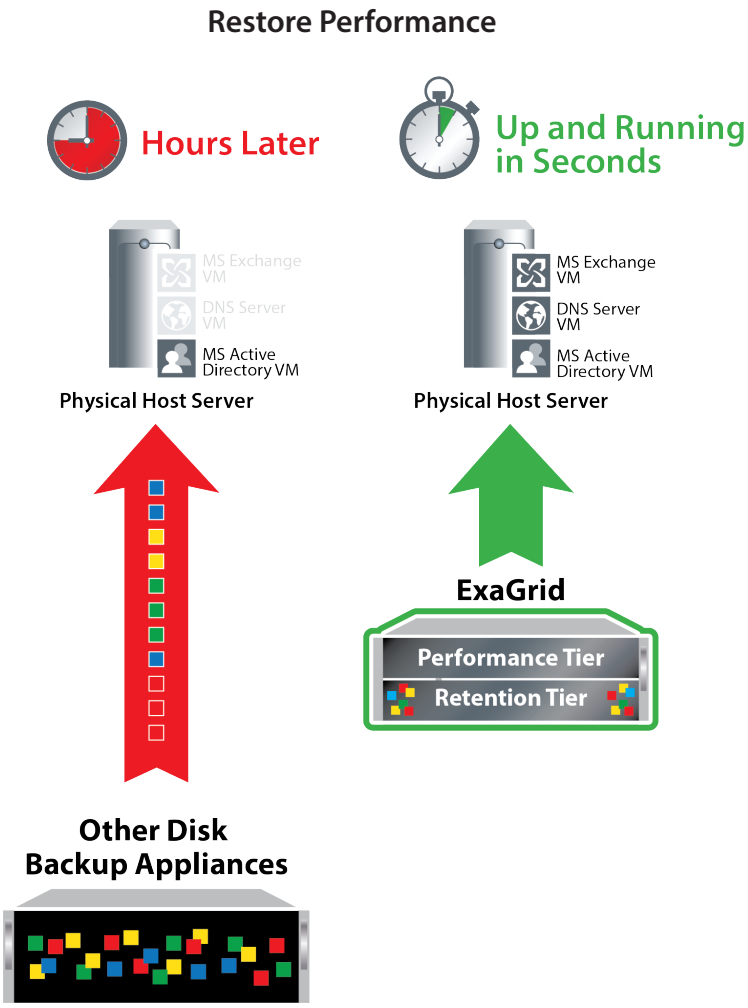

Support for Heterogeneous Backup Applications and Utilities

Customer environments are made up of many backup approaches, backup applications, and utilities and different disk-based backup approaches support these in different ways. Customers may have any number of backups occurring in their environment, including traditional backup applications such as Veeam, Veritas NetBackup, Commvault, IBM Spectrum Protect (TSM), Dell EMC NetWorker, HYCU, Zerto, and many others as well as direct-to-disk SQL dumps, Oracle RMAN direct backups, and specific UNIX utilities such as UNIX TAR. ExaGrid supports over 25 backup applications and utilities.

ExaGrid is able to support backup data from multiple sources, including a variety of backup applications and database dump utilities. Performing deduplication in the backup software limits the ability to have all data from all sources stored and deduplicated in a single target device. Unless 100% of your backup data passes through that particular backup application, a purpose-built disk-based backup appliance such as ExaGrid's is the best choice to meet the requirements of your entire environment.

In contrast, backup application software solutions that have incorporated deduplication by definition only support their own backup application, with its own backup server software and its own backup client agents. These solutions are not able to support backup data from other backup applications or utilities. For example, if you have a physical environment, for backup applications that employ data deduplication but want to use a separate utility for VMware (ex: Veeam) and also do direct database dumps, only the data running through the physical system's backup application will be deduplicated. Also, deduplication in backup software ranges from 2:1 to 10:1 and therefore uses a lot more disk and bandwidth as retention grows versus target-side deduplication appliances that employ far more aggressive deduplication algorithms, since with appliances, resources are dedicated to the task. Dedicated appliances achieve ratios on average of 20:1, using one-quarter to one-third the disk and bandwidth.

ExaGrid allows for Veeam and Commvault to leave deduplication turned on and ExaGrid will further deduplicate the data for increase storage efficiency.

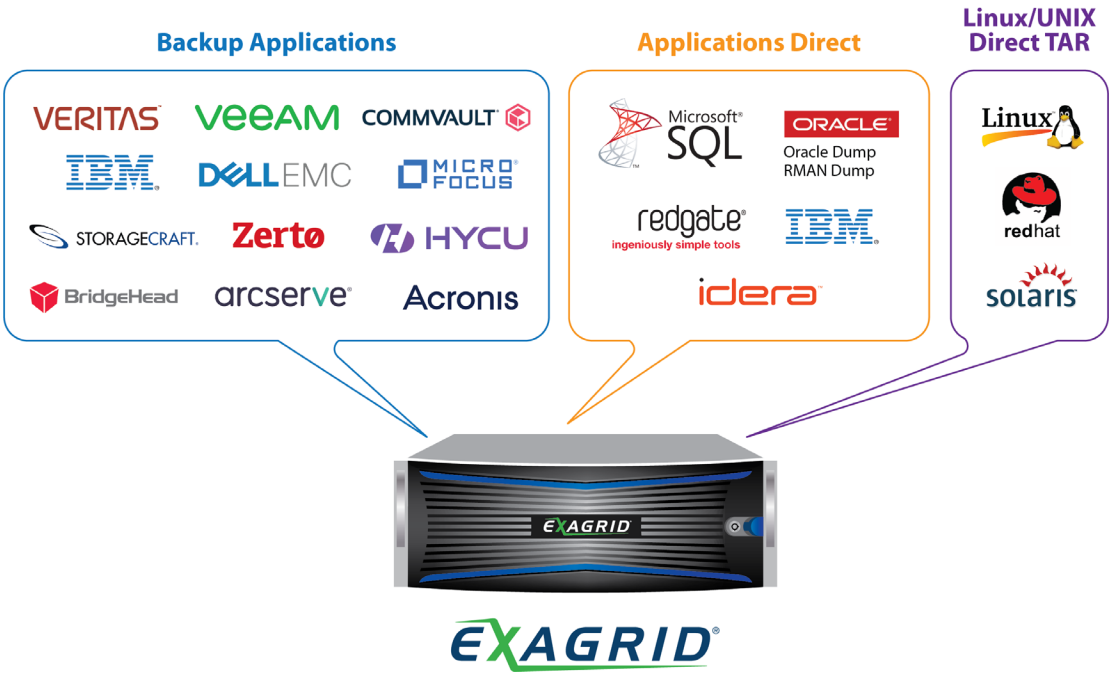

Advanced Backup Features

Veritas Open Storage Technology (OST)

An important consideration when looking at disk-based backup solutions is how well a particular solution supports advanced backup application features such as Backup Exec GRT (Granular Restore) and Veritas' OST (Open Storage Technology) for Backup Exec and NetBackup. Some solutions do not integrate well with these features; poorly implemented GRT solutions, for example, may take hours to restore an individual e-mail or may not work at all.

Veritas' Open Storage Technology is another popular feature that allows for more integrated offsite data protection, and it is important to check whether these features are supported if you are using Veritas NetBackup or Backup Exec, OST allows for faster performance, better management, and unbalanced onsite and offsite retention.

ExaGrid supports and is certified for:

- NetBackup and Backup Exec OST
- NetBackup AIR
- NetBackup appliance compatibility
- NetBackup Accelerator

ExaGrid is the only solution in the market for Veritas NetBackup Accelerator that can reconstitute a full backup in its Landing Zone so that a complete backup is ready for restore in its already hydrated form for faster restores and VM boots.