



# **Backup, Replication and Data Recovery**



Getting Started Guide

April, 2018

This Page is intentionally left blank.

## Table of Contents

1. About Backup .....	4
1.1 Creating Backup Jobs .....	4
1.1.1 Performing Active Full Backup .....	20
2. Backup Copy .....	21
2.1 Creating Backup Copy Jobs .....	21
3. Replication .....	28
3.1 Creating Replication Jobs .....	28
4. Data Recovery .....	39
4.1 Instant VM Recovery .....	39
4.1.1 Performing Instant VM Recovery .....	39

## 1. About Backup

Veeam Backup & Replication is built for virtual environments. It operates at the virtualization layer and uses an image-based approach for VM backup.

Veeam Backup & Replication copies VM data from the source datastore at a block level. It retrieves VM data, compresses and deduplicates it, and stores in backup files on the backup repository in Veeam's proprietary format.

In Veeam Backup & Replication, backup is a job-driven process. To perform backup, you need to configure backup jobs. A backup job is a configuration unit of the backup activity. The backup job defines when, what, how and where to back up. One backup job can be used to process one or several VMs.

The first backup job session always produces a full backup of the VM image. Subsequent backup job sessions are incremental — Veeam Backup & Replication copies only those data blocks that have changed since the last backup job session.

### 1.1 Creating Backup Jobs

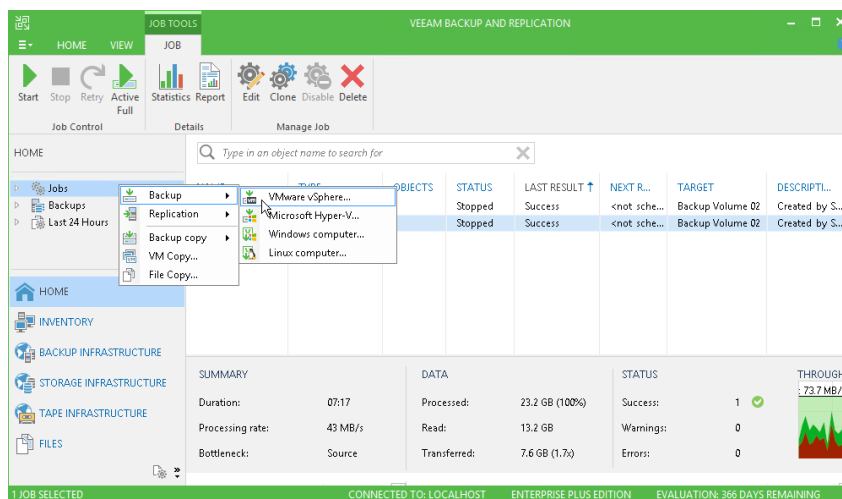
You can configure a backup job and start it immediately or save the job and run it later. Jobs can be started manually or scheduled to run automatically at specific time.

Before creating a backup job, check prerequisites. Then use the **New Backup Job** wizard to configure the backup job.

#### Step 1. Launch the New Backup Job Wizard

To launch the **New Backup Job** wizard, do one of the following:

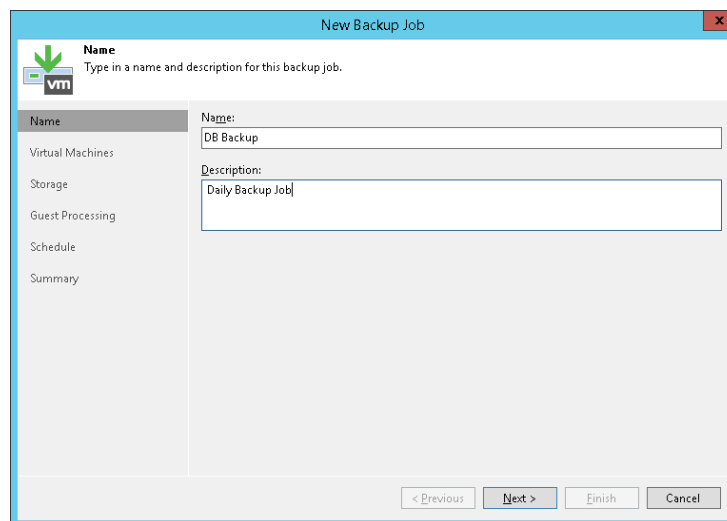
- On the **Home** tab, click **Backup Job** and select **VMware**.
- Open the **Home** view. In the inventory pane right-click **Jobs** and select **Backup > VMware**.
- Open the **Inventory** view. In the working area select the VM(s), click **Add to Backup** on the ribbon and select **New job** or right-click the VM(s) and select **Add to backup job > New job**. Veeam Backup & Replication will start the **New Backup Job** wizard and add the VM(s) to this job. You can add other VMs to the job later on, when you pass through the wizard steps.
- You can quickly add the VM(s) to an already existing job. To do this, open the **Inventory** view, in the working area select the VM(s) and click **Add to Backup > name of the job** on the ribbon or right-click the VM(s) and select **Add to backup job > name of the job**.



### Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the backup job.

1. In the **Name** field, enter a name for the backup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.



### Step 3. Select VMs to Back Up

At the **Virtual Machines** step of the wizard, select VMs and VM containers that you want to back up.

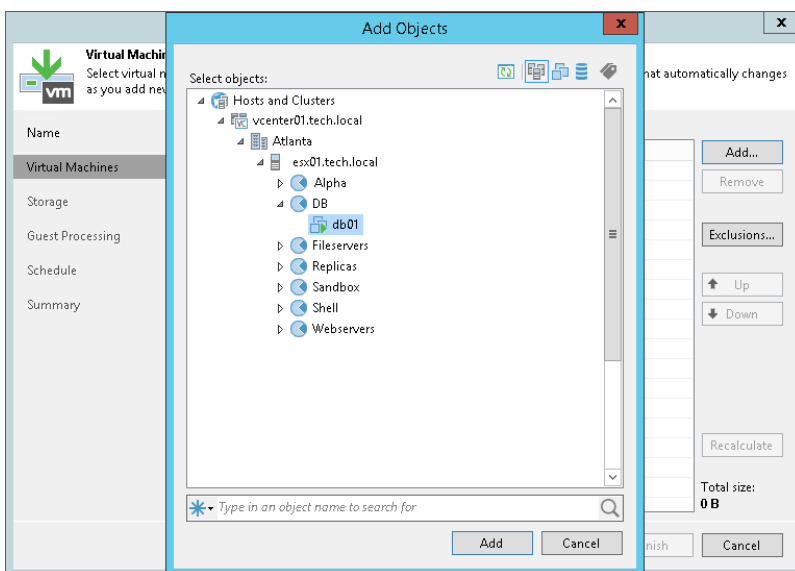
To add VMs to the job:

1. Click **Add**.

2. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you switch to the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
3. Select the VM or VM container in the list and click **Add**.

To quickly find the necessary object, use the search field at the bottom of the **Add Objects** window.

1. Click the button to the left of the search field and select the type of object to search for: *Everything*, *Folder*, *Cluster*, *Host*, *Resource pool*, *VirtualApp* or *Virtual machine*.
2. Enter the object name or a part of it in the search field.
3. Click the **Start search** button on the right or press **[ENTER]**.



### Step 4. Exclude Objects from Backup Job

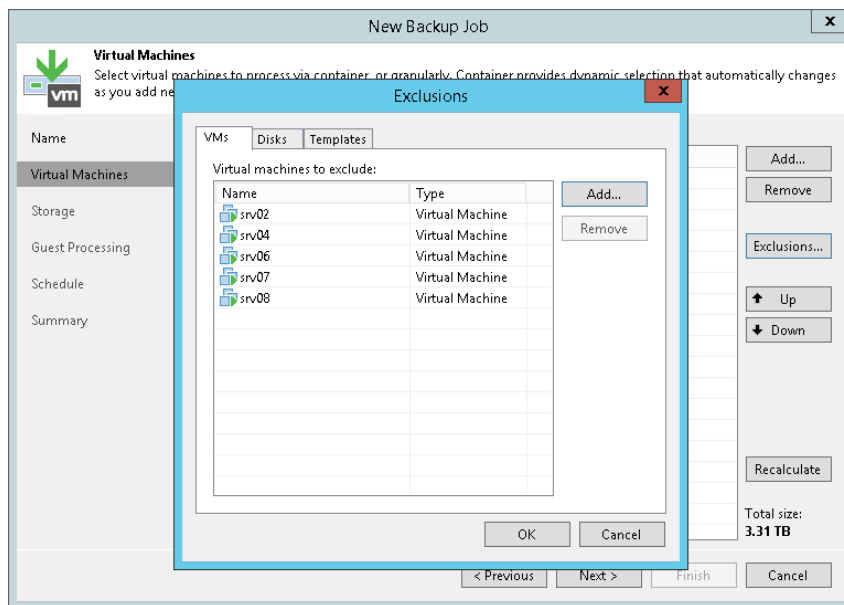
After you have added VMs and VM containers to the job, you can specify which objects you want to exclude from the backup. You can exclude the following types of objects:

- VMs from VM containers
- Specific VM disks
- VM templates

To exclude VMs from a VM container:

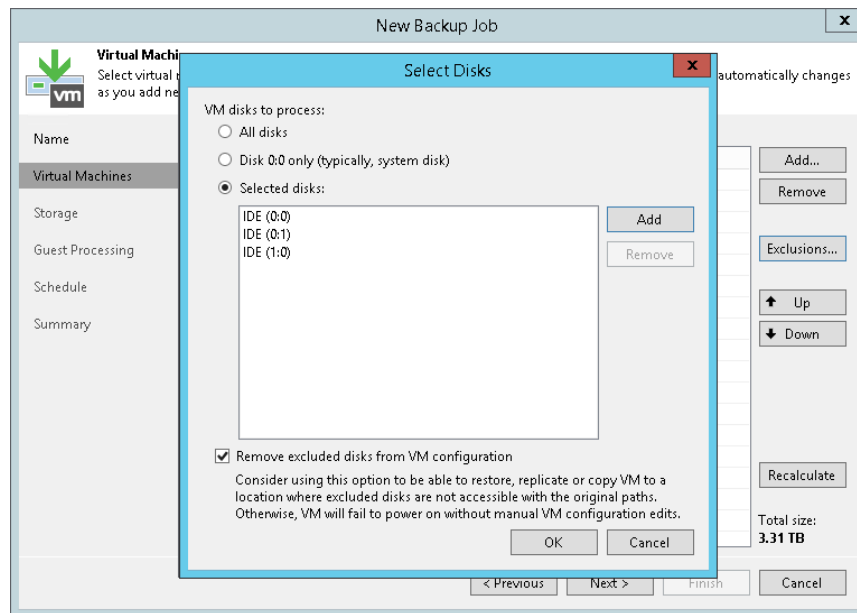
1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Click the **VMs** tab.
3. Click **Add**.

4. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
5. In the displayed tree, select the necessary object and click **Add**. Use the **Show full hierarchy** check box to display the hierarchy of all VMware Servers added to the backup infrastructure.
6. Click **OK**.



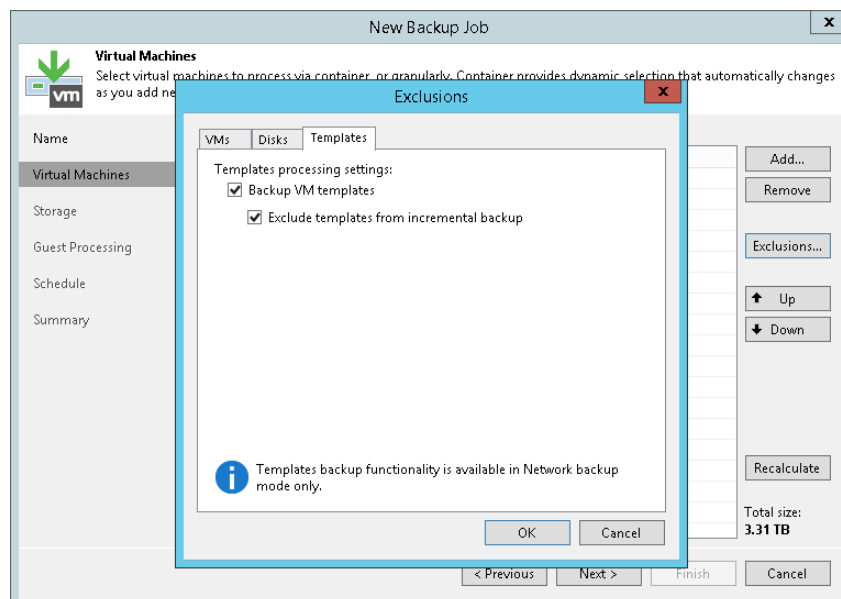
To exclude VM disks:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Click the **Disks** tab.
3. Select the VM in the list and click **Edit**. If you want to exclude disks of a VM added as part of the container, click **Add** to include the VM in the list as a standalone object.
4. Choose disks that you want to back up. You can choose to process all disks, 0:0 disks (typically, system disks) or add to the list custom IDE, SCSI or SATA disks.
5. Select the **Remove excluded disks from VM configuration** check box.



To exclude VM templates:

1. At the **Virtual Machines** step of the wizard, select a VM container and click **Exclusions**.
2. Click the **Templates** tab.
3. Clear the **Backup VM templates** check box.
4. If you want to include VM templates into the full backup only, leave the **Backup VM templates** check box selected and select the **Exclude templates from incremental backup** check box.



### Step 5. Define VM Backup Order

You can define the order in which the backup job must process VMs. Setting VM order can be helpful, for example, if you add some mission-critical VMs to the job and want the job to

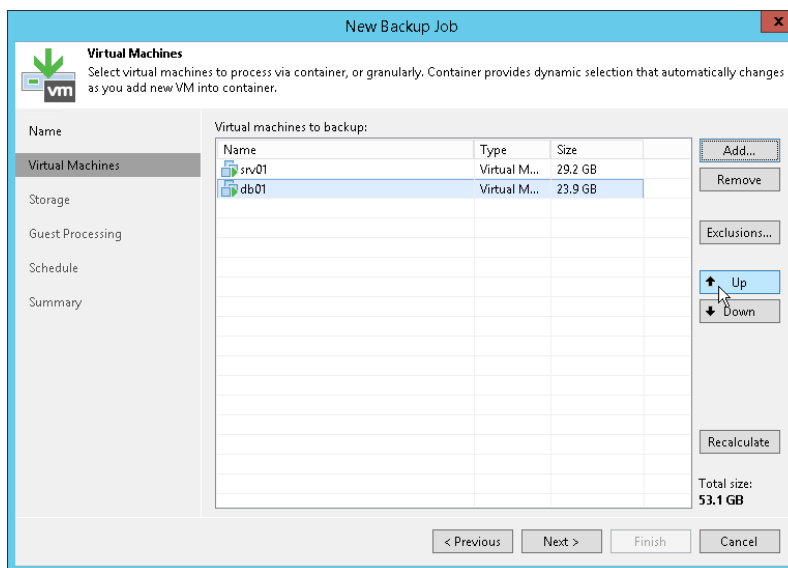


process them first. You can set these VMs first in list to ensure that their processing fits the backup window.

VMs inside a VM container are processed at random. To ensure that VMs are processed in the defined order, you must add them as standalone VMs, not as a part of the VM container.

To define the VM backup order:

1. At the **Virtual Machines** step of the wizard, select a VM or VM container.
2. Use the **Up** and **Down** buttons on the right to move the VM or VM container up or down in the list.



### Step 6. Specify Backup Storage Settings

At the **Storage** step of the wizard, select backup infrastructure components for the job — backup proxy and backup repository, and specify backup storage settings.

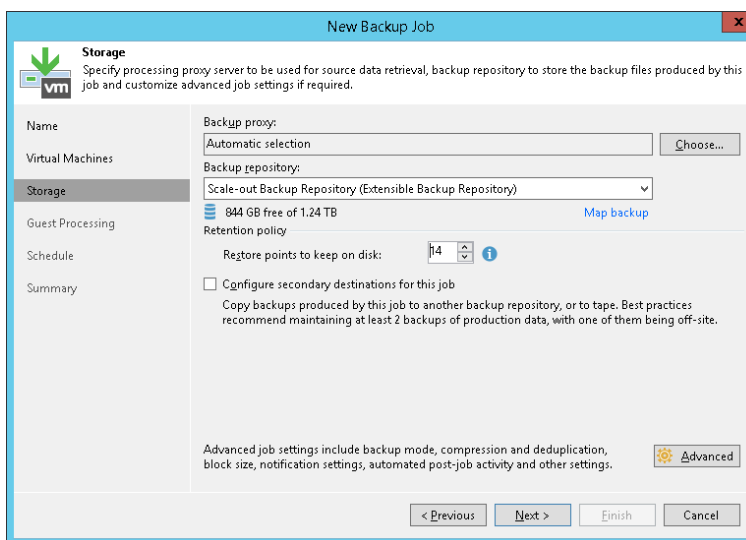
1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
  - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and automatically assign an optimal backup proxy to process VMs in the job.
  - If you choose **Use the selected backup proxy servers specified below**, you can explicitly select backup proxies that the job must use. It is recommended that you select at least two backup proxies to ensure that the backup job starts if one of the proxies fails or loses its connectivity to the source datastore.
2. From the **Backup repository** list, select a backup repository where the created backup files must be stored.
3. You can map the job to a specific backup stored on the backup repository. Backup job mapping can be helpful if you have moved backup files to a new backup

repository and want to point the job to existing backups on this new backup repository.

To map the job to a backup, click the **Map backup** link and select the backup on the backup repository. Backups can be easily identified by job names. To find the backup, you can also use the search field at the bottom of the window.

4. In the **Retention policy** section, specify the number of restore points that you want to store on the backup repository.
5. If you want to archive backup files created with the backup job to a secondary destination (backup repository or tape), select the **Configure secondary destination for this job** check box. With this option enabled, the **New Backup Job** wizard will include an additional step — **Secondary Target**. At the **Secondary Target** step of the wizard, you can link the backup job to the backup copy job or backup to tape backup job.

You can enable this option only if a backup copy job or backup to tape job is already configured on the backup server.



### sdaduStep 7. Specify Advanced Backup Settings

At the **Storage** step of the wizard, specify advanced settings for the backup job:

#### Backup Settings

To specify settings for a backup chain created with the backup job:

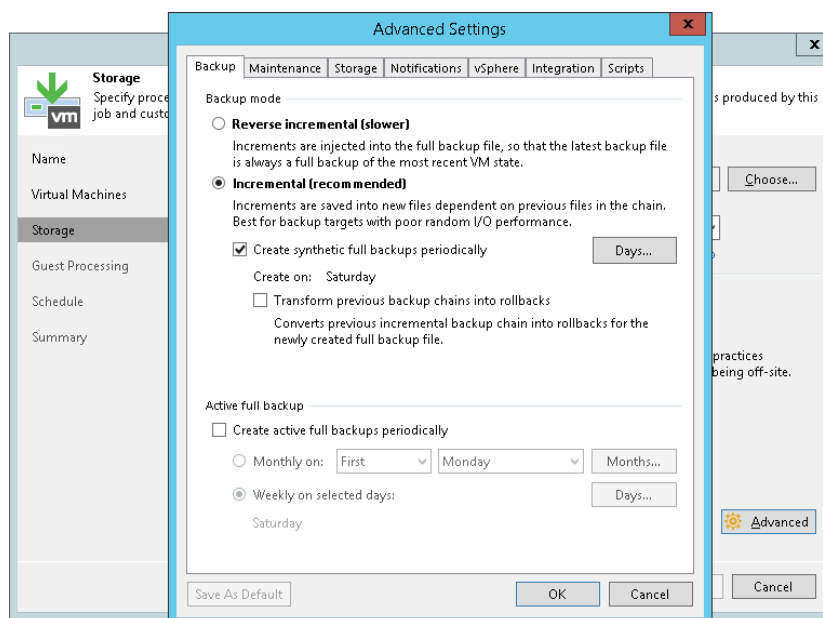
1. At the **Storage** step of the wizard, click **Advanced**.
2. On the **Backup** tab, select the backup method that you want to use to create the backup chain on the backup repository:
  - To create a reverse incremental backup chain, select **Reverse Incremental**.
  - To create an incremental backup chain, select **Incremental** and enable synthetic full and/or active full backups (see items 3-4).

- To create a forever forward incremental backup chain, select **Incremental** and do not enable synthetic full and/or active full backups (see items 3-4).
- 3. If you choose the incremental backup method, you can select to periodically create synthetic full backups and/or active full backups.
- To create a synthetic full backup, select the **Create synthetic full backups periodically** check box and click **Days** to schedule synthetic full backups on the necessary week days.

You can additionally choose to transform the previous full backup chain into the reverse incremental backup chain. To do this, select the **Transform previous full backup chains into rollbacks** check box.

- To create full backups regularly, select the **Create active full backups periodically** check box. Use the **Monthly on** or **Weekly on selected days** options to define scheduling settings.

Before scheduling periodic full backups, you must make sure that you have enough free space on the backup repository. As an alternative, you can create active full backups manually when needed.



### Maintenance Settings

You can instruct Veeam Backup & Replication to periodically perform maintenance operations — service actions that will help make sure that the backup chain remains valid and consistent.

To specify maintenance settings for the backup job:

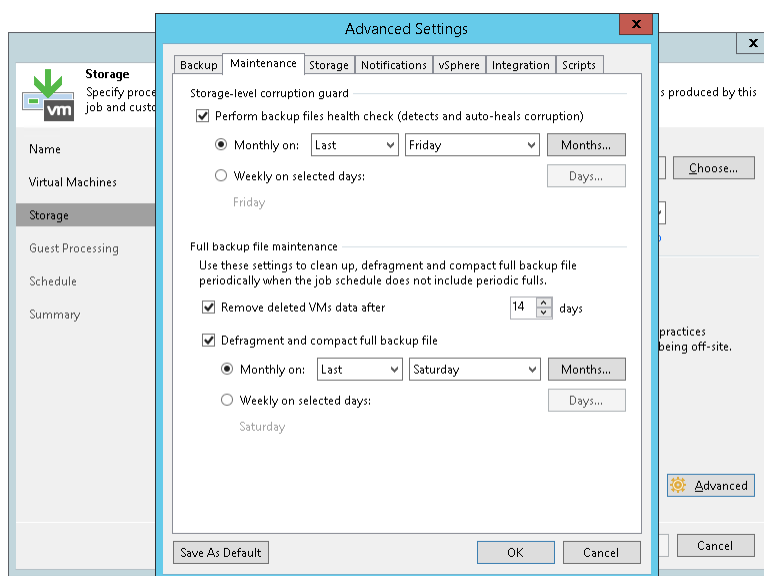
1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Maintenance** tab.

- To periodically perform a health check for the latest restore point in the backup chain, in the **Storage-level corruption guard** section select the **Perform backup files health check** check box and specify the time schedule for the health check.

Select the **Remove deleted VMs data after** check box and specify the number of days for which you want to keep backup data for deleted VMs. If a VM is no longer available (for example, it was deleted or excluded from the job), Veeam Backup & Replication will keep its data on the backup repository for the period that you have specified. When this period is over, data of the deleted VM will be removed from the backup repository.

By default, the deleted VM retention period is 14 days. Do not set the deleted VM retention period to 1 day or a similar short interval.

- To periodically compact a full backup, select the **Defragment and compact full backup file** check box and specify the schedule for the compact operation.



## Storage Settings

To specify storage settings for the backup job:

- At the **Storage** step of the wizard, click **Advanced**.
- Click the **Storage** tab.
- By default, Veeam Backup & Replication deduplicates VM data before storing it on the backup repository. Data deduplication provides a smaller size of the backup file but may reduce the backup job performance. To disable data deduplication, clear the **Enable inline data deduplication** check box.
- If you want to include data blocks of the hiberfil.sys file and pagefile.sys file to the backup, clear the **Exclude swap file blocks** check box.
- From the **Compression level list**, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.
- In the **Storage optimization** section, select what type of backup target you plan to use: *Local target (16 TB + backup files)*, *Local target*, *LAN target* or *WAN target*.

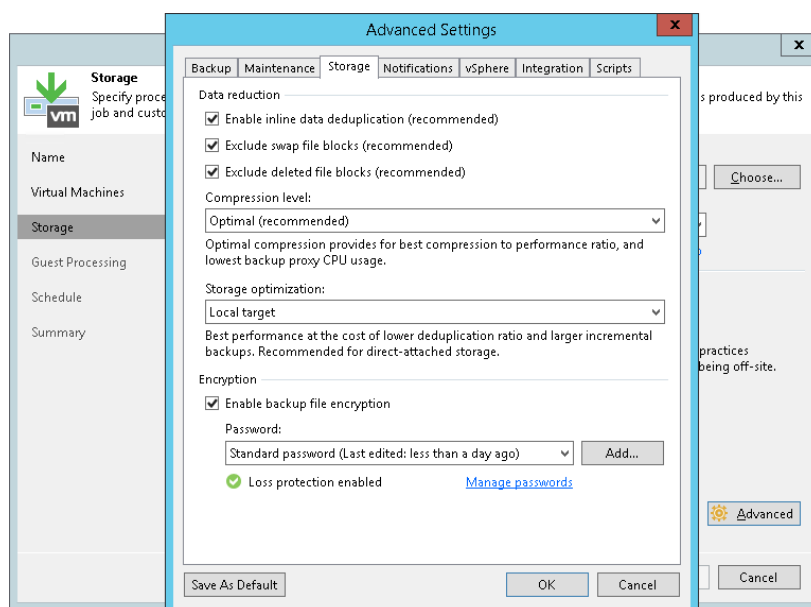
Depending on the chosen storage type, Veeam Backup & Replication will use data blocks of different size to optimize the size of backup files and job performance.

7. To encrypt the content of backup files, select the **Enable backup file encryption** check box. In the **Password** field, select a password that you want to use for encryption. If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password.

If the backup server is not connected to Veeam Backup Enterprise Manager, you will not be able to restore data from encrypted backups in case you lose the password. Veeam Backup & Replication will display a warning about it.

### NOTE:

If you enable encryption for an existing backup job, during the next job session Veeam Backup & Replication will create a full backup file. The created full backup file and subsequent incremental backup files in the backup chain will be encrypted with the specified password.

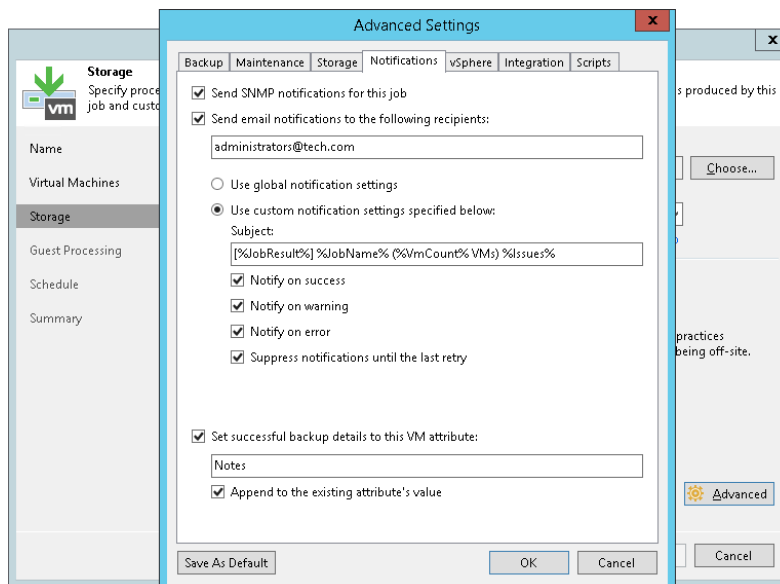


## Notification Settings

To specify notification settings for the backup job:

1. At the Storage step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.
4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications about the job completion status by email. In the field below, specify recipient's email address. You can enter several addresses separated by a semicolon.

5. You can choose to use global notification settings or specify custom notification settings.
  - To receive a typical notification for the job, select Use global notification settings. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server.
  - To configure a custom notification for the job, select Use custom notification settings specified below check box. You can specify the following notification settings:
    - a. 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 VMs in the job) and *%Issues%* (number of VMs in the job that have finished with the *Warning* or *Failed* status).
    - b. Select the **Notify on success**, **Notify on warning** and/or **Notify on error** check boxes to receive email notification if the job completes successfully, fails or completes with a warning.
    - c. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.
6. Select the **Set successful backup details to this VM attribute** check box to write information about successfully performed backup and backup results (backup date and time, backup server name and path to the backup file) to a VM attribute. In the field below, enter a name of the attribute. If the specified attribute does not exist, Veeam Backup & Replication will create it.



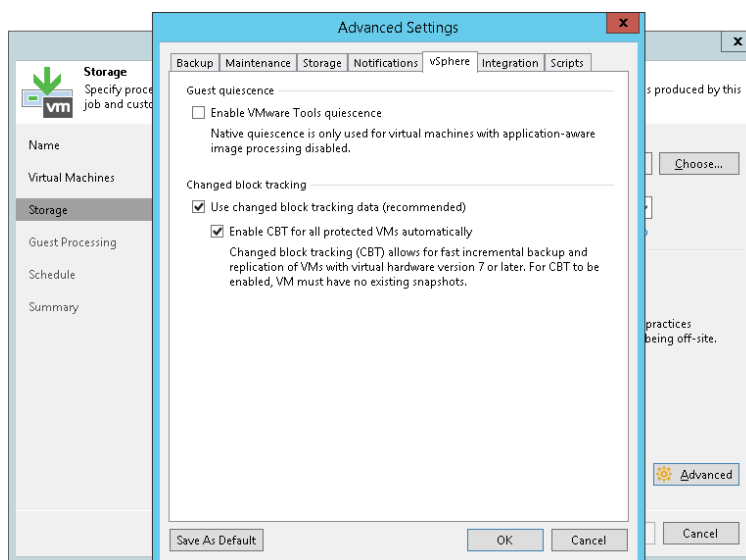
## vSphere Settings

To specify VMware vSphere settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **vSphere** tab.
3. Select the **Enable VMware tools quiescence** check box to freeze the file system of processed VMs during backup.
4. In the **Changed block tracking** section, specify if VMware vSphere CBT must be used for VM backup. By default, this option is enabled. If you want to force using CBT even if CBT is disabled at the level of the ESX(i) host, select the **Enable CBT for all processed VMs automatically** check box.

### IMPORTANT!

You can use CBT for VMs with virtual hardware version 7 or later.



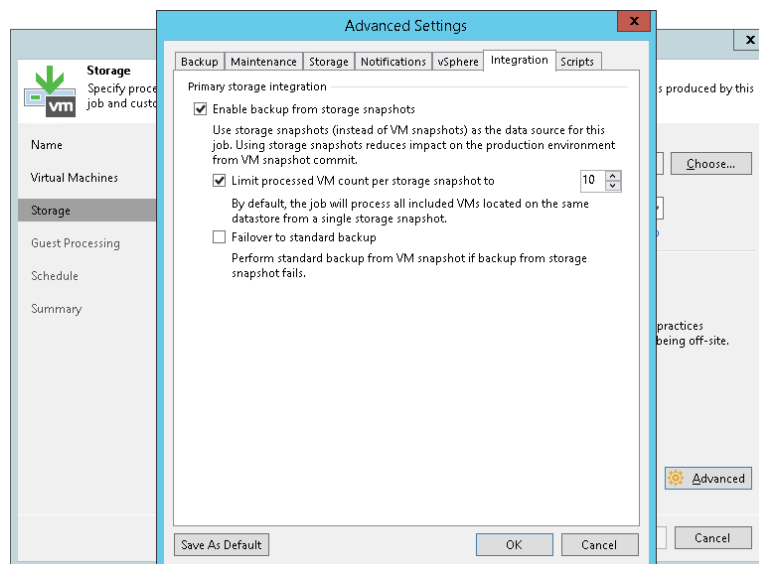
## Integration Settings

On the **Integration** tab, you can define whether you want to use the Backup from Storage Snapshots technology to create the backup.

To specify storage integration settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Integration** tab.
3. By default, the **Enable backup from storage snapshots** option is enabled. If you do not want to use Backup from Storage Snapshots, clear this check box.

4. If you add to the job many VMs whose disks are located on the same volume or LUN, select the **Limit processed VM count per storage snapshot to <N>** check box and specify the number of VMs for which one storage snapshot must be created.
5. To fail over to the regular VM processing mode and back up or replicate such VMs in the regular processing mode, select the **Failover to standard backup** check box.



## Script Settings

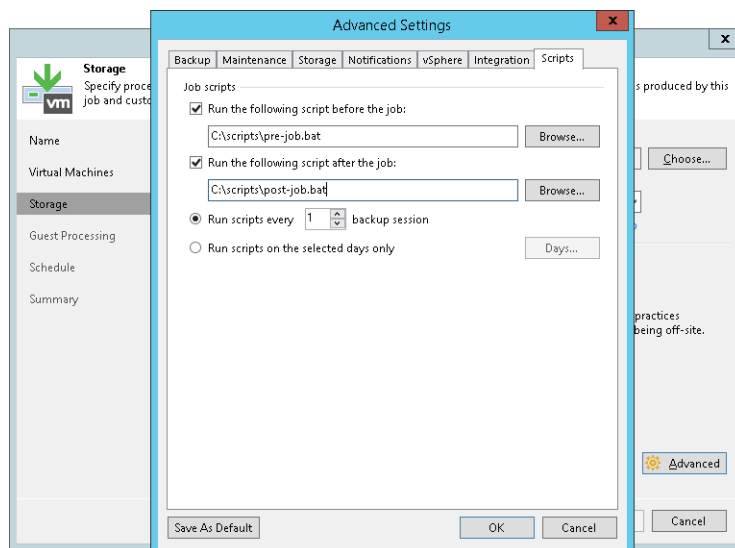
To specify script settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. If you want to execute custom scripts before and/or after the backup job, select the **Run the following script before the job** and **Run the following script after the job** check boxes and click **Browse** to choose executable file(s) from a local folder on the backup server. The scripts are executed on the backup server.

You can select to execute pre- and post-backup actions after a number of backup sessions or on specific week days.

- If you select the **Run scripts every <N> backup session** option, specify the number of the backup job sessions after which the script(s) must be executed.
- If you select the **Run scripts on the selected days only** option, click **Days** and specify week days on which the script(s) must be executed.






### Step 8. Specify Secondary Target

The **Secondary Target** step of the wizard is available if you have enabled the **Configure secondary destination for this job** option at the **Storage** step of the wizard.

At the **Secondary Target** step of the wizard, you can link the backup job to a backup to tape or backup copy job. As a result, the backup job will be added as a source to the backup to tape or backup copy job.

To link jobs:

1. Click **Add**.
2. From the jobs list, select a backup to tape or backup copy job that must be linked to the backup job. You can link several jobs to the backup job, for example, one backup to tape job and one backup copy job. To quickly find the job, use the search field at the bottom of the wizard.



## Step 9. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, you can enable the following settings for VM guest OS processing:

- Application-aware processing
- Transaction log handling for Microsoft SQL VMs
- Archived log handling for Oracle VM
- VM guest OS file exclusion
- Use of pre-freeze and post-thaw scripts
- VM guest OS file indexing

You must specify a user account that will be used to connect to the VM guest OS and deploy the runtime process:

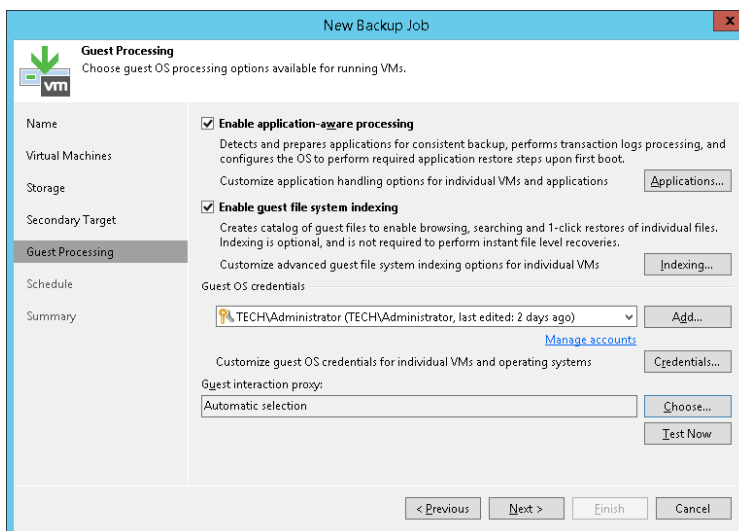
1. From the **Guest OS credentials** list, select a user account with local administrator privileges on the VM guest OS. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials.
2. By default, Veeam Backup & Replication uses the same credentials for all VMs in the job. If some VM requires a different user account, click **Credentials** and enter custom credentials for the VM.
3. If you have added Microsoft Windows VMs to the job, specify which guest interaction proxy Veeam Backup & Replication can use to deploy the runtime process on the VM guest OS. On the right of the **Guest interaction proxy** field, click **Choose**.
  - Leave **Automatic selection** to let Veeam Backup & Replication automatically select the guest interaction proxy.
  - Select **Use the selected guest interaction proxy servers only** to explicitly define which servers will perform the guest interaction proxy role. The list of servers contains Microsoft Windows servers added to the backup infrastructure.

To check if Veeam Backup & Replication can communicate with VMs added to the job and deploy the runtime process on their guest OSes, click **Test**

**Now**. Veeam Backup & Replication will use the specified credentials to connect to all VMs in the list.

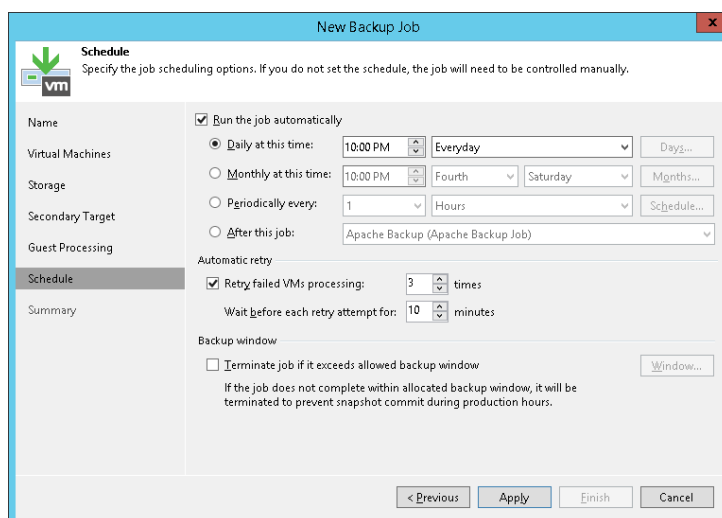
### NOTE:

The guest interaction proxy functionality is available in the Enterprise and Enterprise Plus Editions of Veeam Backup & Replication.



### Step 10. Define the job schedule

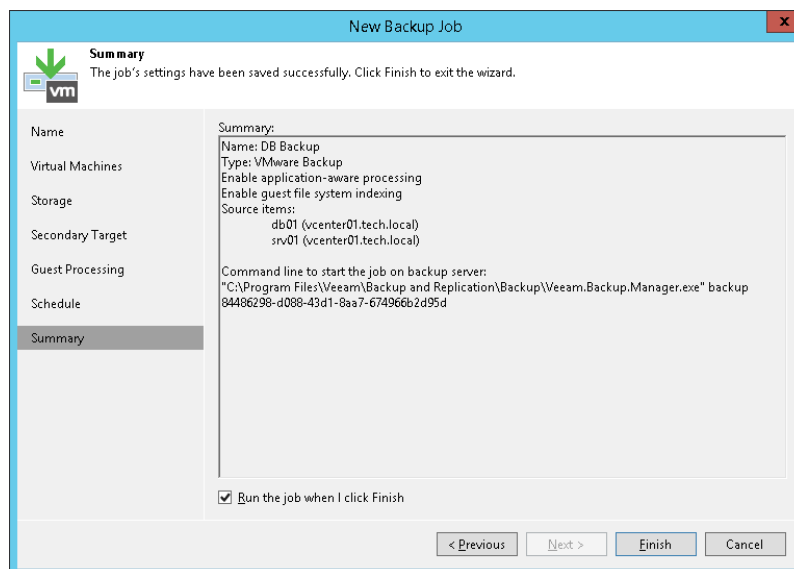
At the **Schedule** step of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.



### Step 11. Finish Working with the Wizard

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

1. Review details of the backup job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

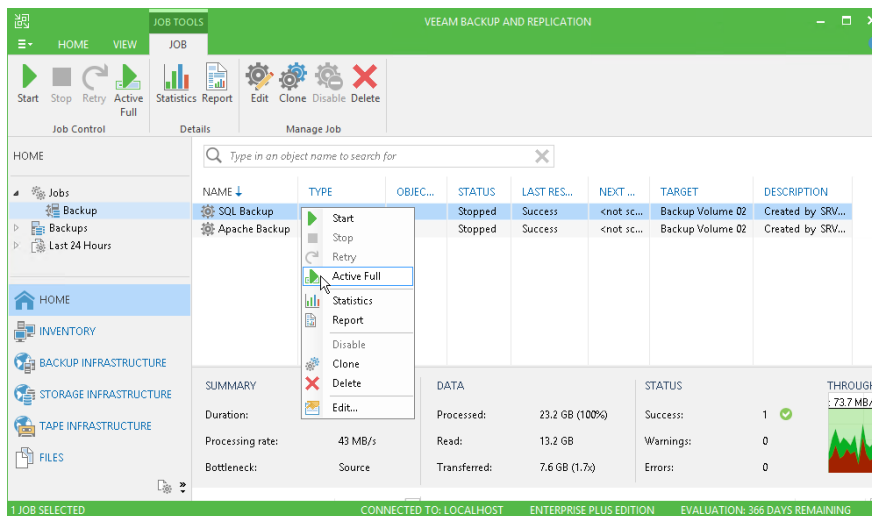


## 1.1.1 Performing Active Full Backup

You can create an ad-hoc full backup — active full backup, and add it to the backup chain on the backup repository. The active full backup resets the backup chain. All subsequent incremental backups use the active full backup as a starting point. The previously used full backup will remain on the backup repository until it is removed from the backup chain according to the retention policy.

To perform active full backup:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Active Full** on the ribbon or right-click the job and select **Active Full**.



## 2. Backup Copy

With backup copy, you can create several instances of the same backup file and copy them to secondary (target) backup repositories for long-term storage. Target backup repositories can be located in the same site as the source backup repository, or can be deployed offsite. The backup copy file has the same format as the primary backup, so you can restore necessary data directly from it in case of a disaster.

### 2.1 Creating Backup Copy Jobs

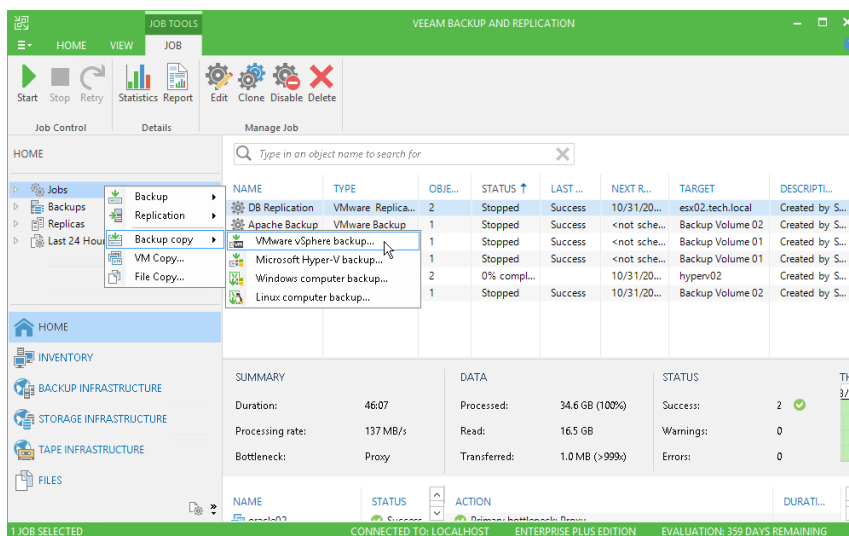
To copy a backup to a secondary location, you must configure a backup copy job. The backup copy job defines how, where and when to copy backups. One job can be used to process backups of one or more machines.

Use the **New Backup Copy Job** wizard to configure a backup copy job.

#### Step 1. Launch New Backup Copy Job Wizard

To run the **New Backup Copy Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Copy** and select the necessary platform: *VMware vSphere backup*, *Windows computer backup* or *Linux computer backup*.
- Open the **Home** view, in the inventory pane right-click **Jobs** and select the necessary option:
  - **Backup Copy** > *VMware vSphere backup* — if you want to create a copy of a VM backup.
  - **Backup Copy** > *Windows computer backup* — if you want to create a copy of a Veeam Agent backup created for Microsoft Windows machine(s).
  - **Backup Copy** > *Linux computer backup* — if you want to create a copy of a Veeam Agent backup created for Linux machine(s).
- Open the **Home** view, right-click anywhere in the working area and select the necessary option:
  - **Backup Copy** > *VMware vSphere backup* — if you want to create a copy of a VM backup.
  - **Backup Copy** > *Windows computer backup* — if you want to create a copy of a Veeam Agent backup created for Microsoft Windows machine(s).
  - **Backup Copy** > *Linux computer backup* — if you want to create a copy of a Veeam Agent backup created for Linux machine(s).



## Step 2. Specify Job Name and Description

At the **Job** step of the wizard, specify basic settings 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.

In the **Copy every** field, specify the time interval according to which the synchronization process must start.

**Job**

Backup copy job efficiently creates local and remote copies of your backups, making it easy to maintain multiple copies of your data. Type in a name and description for the job, and specify backup copy interval.

**Name:**  
DB Backup Copy Job

**Description:**  
Daily Backup Copy Job

**Copy every:**  
1 Day starting at 12:00 AM

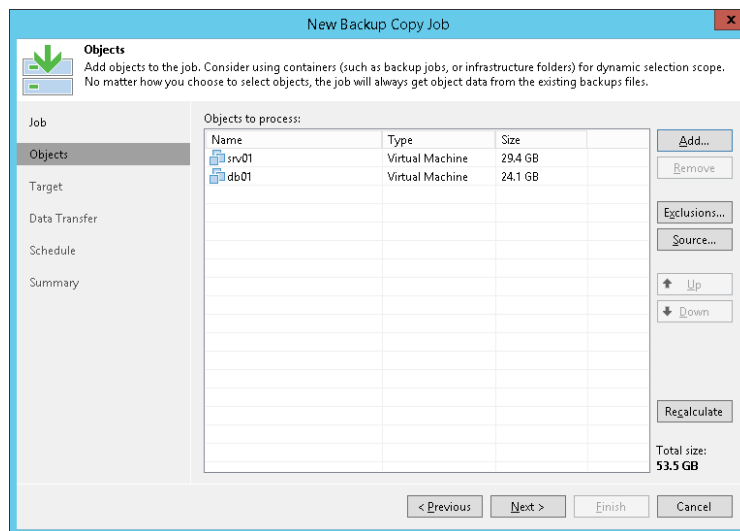
Controls how often backup copies are created. Backup Copy job creates a new backup file for each copy interval, and starts copying the most recent restore point of each processed object into this backup file immediately, or as soon as the new restore point appears in the source backup repository.

< Previous Next > Finish Cancel

### Step 3. Select Machines to Process

At the **Objects** step of the wizard, select machines whose restore points you want to copy to the target backup repository.

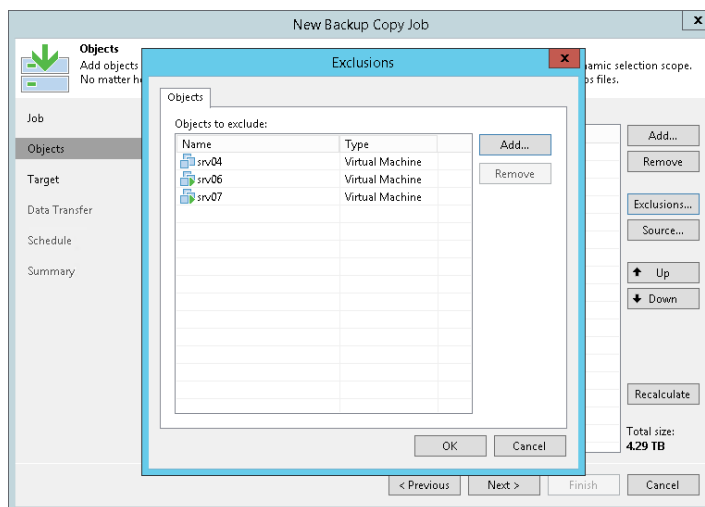
1. Click **Add**.
2. Select the machine(s) that you want to process with the job.



### Step 4. Exclude Objects from the Backup Copy Job

If you have added VM containers to the list of processed machines, you can specify which objects you want to exclude from the backup copy job.

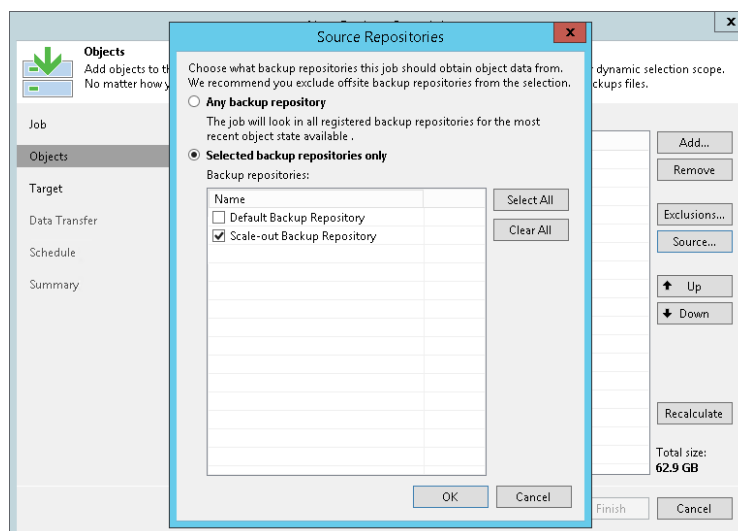
1. At the **Objects** step of the wizard, select a VM container added to the job and click **Exclusions**.
2. Click the **Objects** tab.
3. Click **Add**.
4. Use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
5. In the displayed tree, select the necessary object and click **Add**. Use the **Show full hierarchy** check box to display the hierarchy of all hosts added to Veeam Backup & Replication.
6. Click **OK**.



### Step 5. Select Source Backup Repositories

By default, Veeam Backup & Replication searches for restore points on all backup repositories configured in the backup infrastructure. However, you can select backup repositories in which Veeam Backup & Replication must search for restore points of selected machine(s).

1. At the **Objects** step of the wizard, click **Source**.
2. Choose backup repositories on which restore points must be searched for.



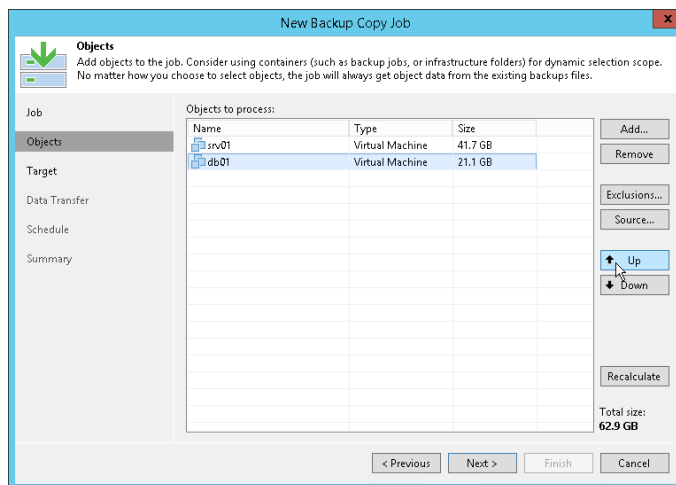
### Step 6. Define Machines Processing Order

To define machine processing order:

1. At the **Objects** step of the wizard, select a virtual or physical machine, or VM container added to the job.



2. Use the **Up** and **Down** buttons on the right to move the virtual or physical machine, or VM container up or down in the list.



### Step 7. Define the Backup Copy Target

At the **Target** step of the wizard, define the target backup repository for the backup copy job and retention policy settings.

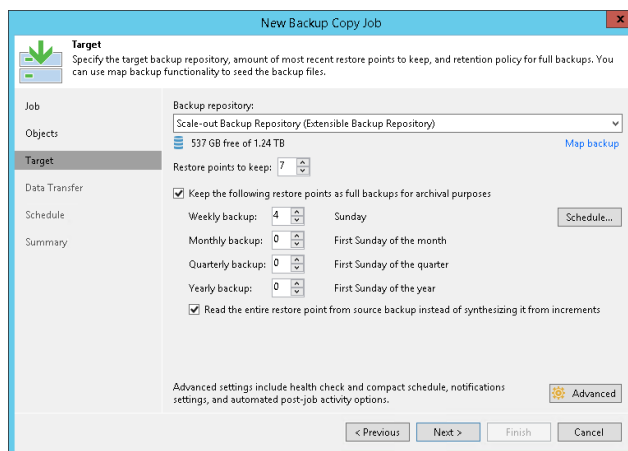
1. From the **Backup repository** list, select a backup repository in the target site where copied backups must be stored.
2. In the **Restore points to keep** field, specify the number of restore points that must be retained on the target backup repository. When this number is exceeded, Veeam Backup & Replication will remove the earliest restore point from the backup chain.

The maximum number of restore points for the backup copy job is 999.

3. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep the following restore points as full backups for archival purposes** check box. In the fields below, specify the number of daily, weekly, monthly, quarterly and yearly full intervals for which backups must be retained. Use the **Schedule** button to define the schedule by which GFS full backups must be created.
4. If you do not want Veeam Backup & Replication to synthesize archive backup files on the target backup repository, select the **Read the entire restore point from source backup instead of synthesizing it from increments** check box.

### IMPORTANT!

You cannot enable GFS retention settings if you use a backup repository with rotated drives as the target backup repository.

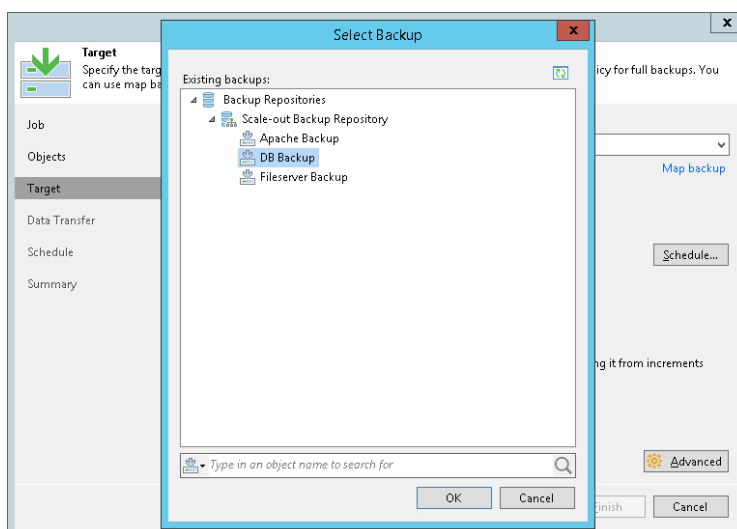


### Step 8. Map a Backup File

If you plan to copy backups over WAN or slow connections, you can use backup mapping.

To map a backup copy job to the backup:

1. Click the **Map backup** link.
2. 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.

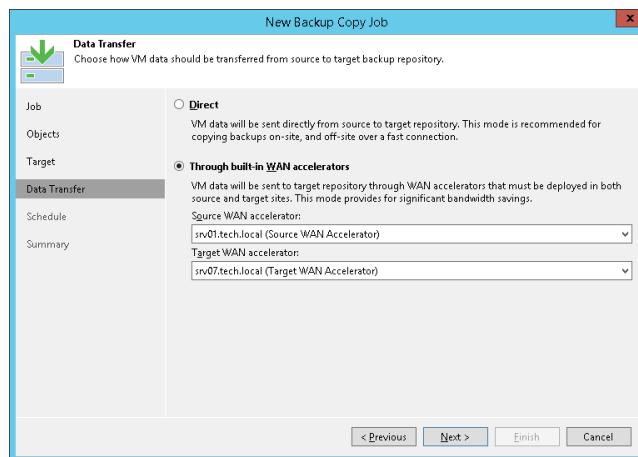


### Step 9. Specify Data Path Settings

To use WAN acceleration for the backup copy job:

1. Select the **Through built-in WAN accelerators** option.
2. From the **Source WAN accelerator** list, select a WAN accelerator configured in the source site.
3. From the **Target WAN accelerator** list, select a WAN accelerator configured in the target site.

Be extremely careful when assigning WAN accelerators to the backup copy job. If you make a mistake and assign the WAN accelerator in the target site to be used as the source WAN accelerator, data will go in the backward direction and workload on the WAN will increase.

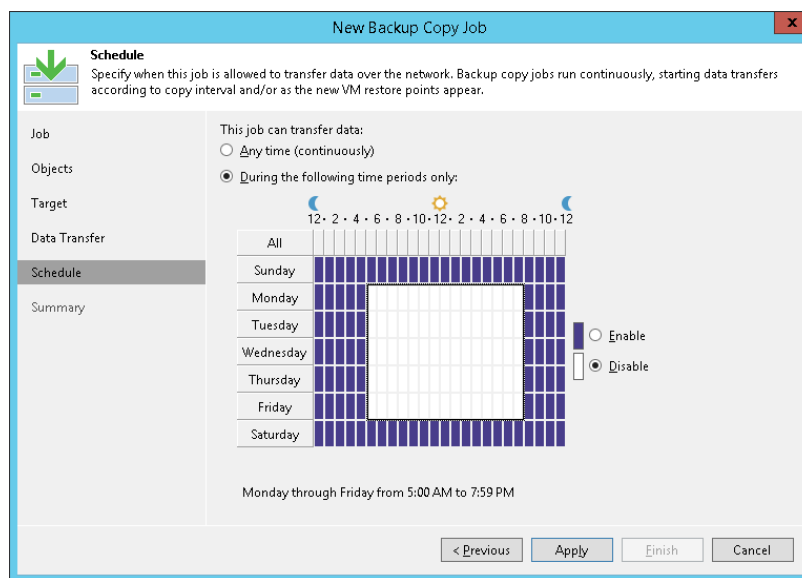


### Step 10. Define the Backup Copy Window

At the **Schedule** step of the wizard, define the time span in which the backup copy job must not transport data between source and target backup repositories.

To define a backup window for the backup copy job:

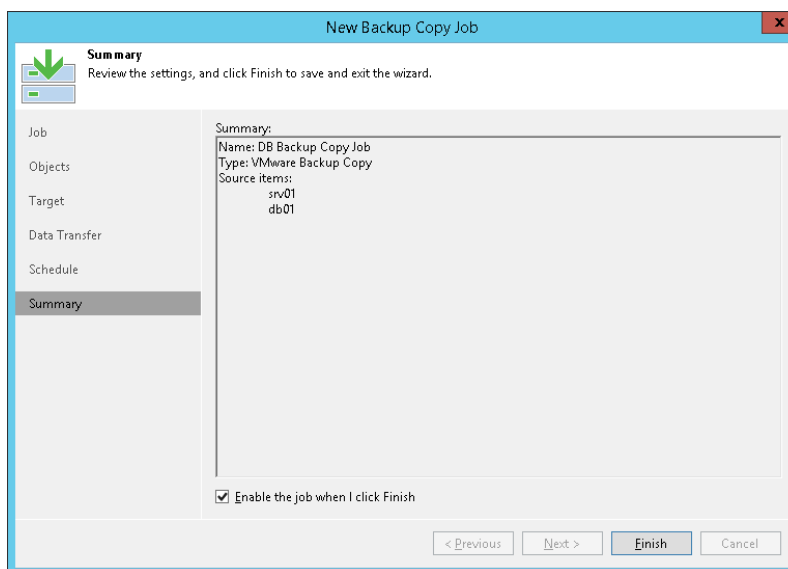
1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enable** and **Disable** options to mark the selected area as allowed or prohibited for the backup copy job.



## Step 11. Finish Working with Wizard

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

1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.



## 3. Replication

Replication provides the best recovery time objective (RTO) values, as you actually have a copy of your VM in a ready-to-start state. That is why replication is commonly recommended for the most critical VMs that need minimum RTOs.

### 3.1 Creating Replication Jobs

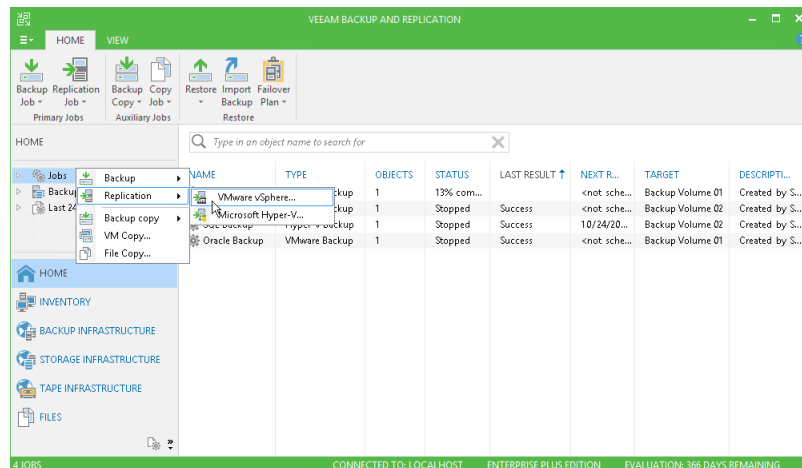
To create VM replicas, you must configure a replication job. Use the **New Replication Job** wizard to configure a replication job.

#### Step 1. Launch New Replication Job Wizard

To run the **New Replication Job** wizard, do one of the following:

- On the **Home** tab, click **Replication Job** and select **VMware**.
- Open the **Home** view, in the inventory pane right-click **Jobs** and select **Replication > VMware vSphere**.

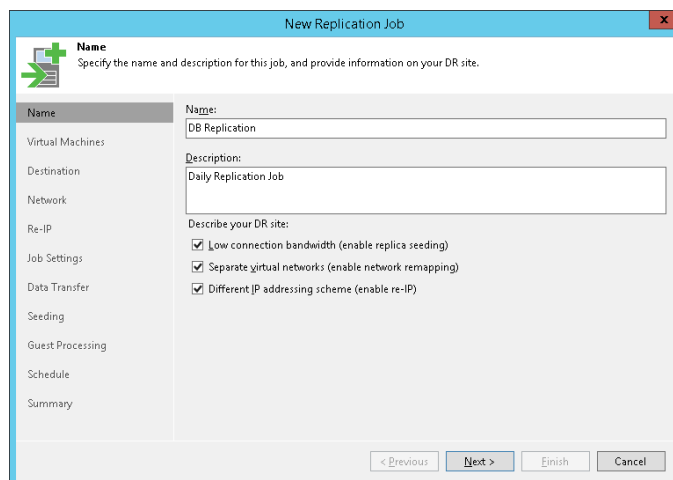
- Open the **Inventory** view. In the working area, select the VM(s), click **Add to Replication** on the ribbon and select **New job** or right-click the VM(s) and select **Add to replication job > New job**.
- You can quickly include the VM(s) to already existing jobs. To do this, open the **Inventory** view. In the working area, select the VM(s) and click **Add to Replication > name of the job** on the ribbon or right-click VMs and select **Add to replication job > name of the job**.



### Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify the job name and description and define advanced settings for the replication job.

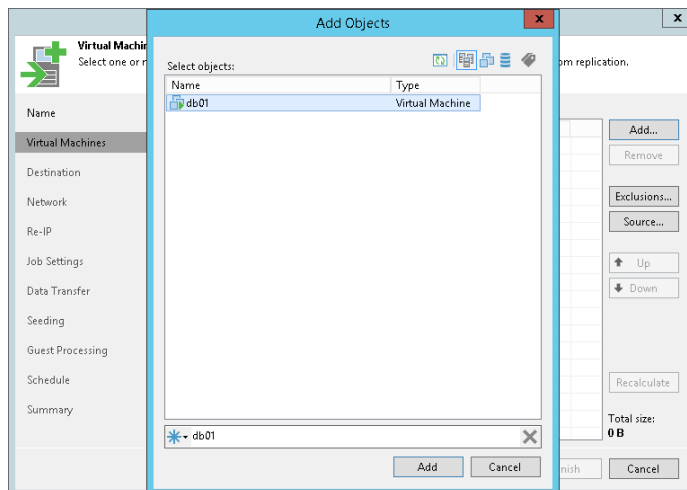
1. In the **Name** field, enter a name for the replication job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created a job, date and time when the job was created.
3. If you plan to replicate VMs to a DR site, you can use a number of advanced settings for the job:
  - Select the **Low connection bandwidth** check box to enable the Seeding step in the wizard.
  - Select the **Separate virtual networks** check box to enable the Network step in the wizard.
  - Select the **Different IP addressing scheme** check box to enable the Re-IP step in the wizard. Re-IP possibilities can be used to automate reconfiguration of replica IP addresses for Microsoft Windows VMs if IP schemes in the DR and production sites do not match.



### Step 3. Select VMs to Replicate

At the **Virtual Machines** step of the wizard, select VMs and VM containers that you want to replicate.

1. Click **Add**.
2. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**.
3. Select the object and click **Add**.

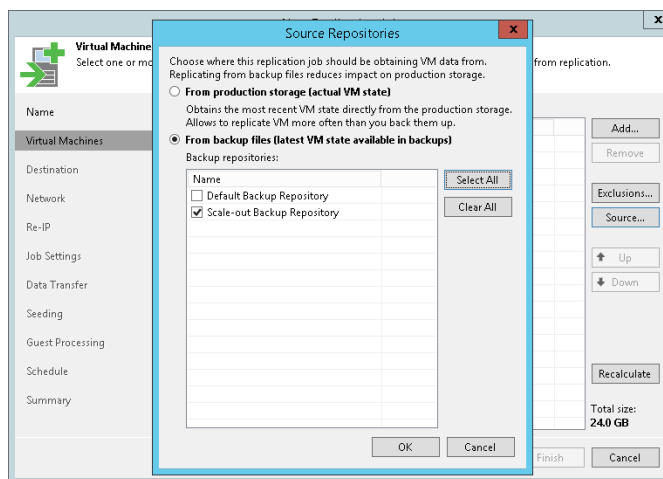


### Step 4. Specify Data Source

You can select a data source from which VM data must be read.

1. At the **Virtual Machines** step of the wizard, click **Source** on the right of the VMs list.
2. In the displayed window, select one of the following options:
  - **From production storage.** In this case, Veeam Backup & Replication will retrieve VM data from datastores connected to the source ESX(i) host.

- **From backup files.** In this case, Veeam Backup & Replication will read VM data from the backup chain already existing on the backup repository. This option can be used in the replica from backup scenario.



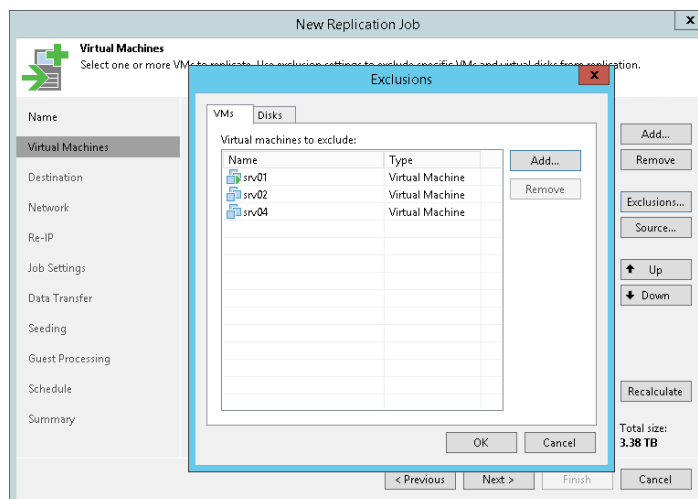
### Step 5. Exclude Objects from Replication Job

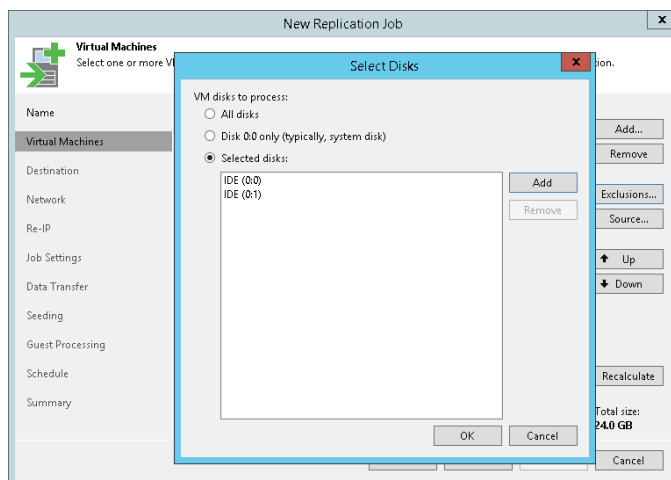
After you have added VMs and VM containers to the job, you can specify which objects you want to exclude from replicas. You can exclude the following types of objects:

- VMs from VM containers
- Specific VM disks

**NOTE:** To make the replication process faster and reduce the size of created replicas, Veeam Backup & Replication automatically excludes the following objects from replication:

- VM log files
- VM templates from VM containers

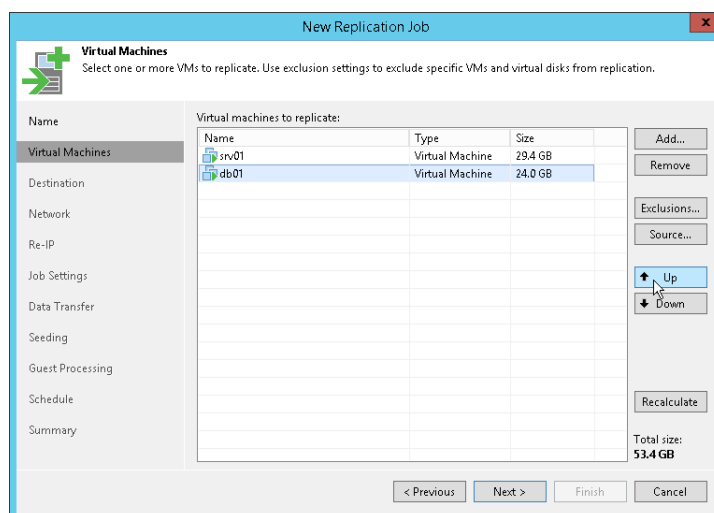




### Step 6. Define VM Replication Order

To define VM replication order:

1. At the **Virtual Machines** step of the wizard, select a VM or VM container.
2. Use the **Up** and **Down** buttons on the right to move the VM or VM container up or down in the list.



### Step 7. Specify Replica Destination

At the **Destination** step of the wizard, select a destination for the VM replica(s).

1. Click **Choose** next to the **Host or cluster** field and select an ESX(i) host or cluster where VM replica(s) must be registered.
2. If all or majority of VM replicas must belong to the same resource pool, click **Choose** next to the **Resource pool** field and select the target resource pool.

If you want to place VM replicas to different resource pools:

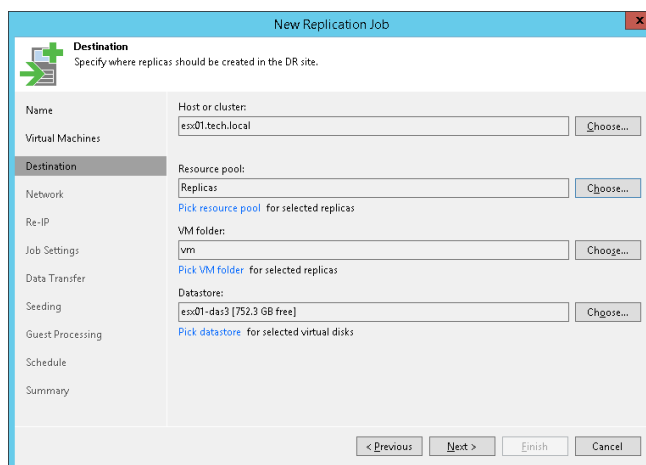
- a. Click the **Pick resource pool for selected replicas** link.



- b. In the **Choose Resource Pool** window, click **Add VM** on the right and select the VMs and click **Add**.
  - c. Select the added VM in the **Replica VM resource pool** list and click **Resource Pool** at the bottom of the window.
  - d. From the list of available resource pools, choose the target resource pool for the VM.
3. If all or majority of VM replicas must be placed in the same folder, click **Choose** next to the **VM folder** field and choose the target folder.

The **VM folder** section is disabled if you selected a standalone ESX(i) host as a target for VM replicas.

4. If files for all or majority of VM replicas must be stored on the same datastore, click **Choose** next to the **Datastore** field and select the target datastore.

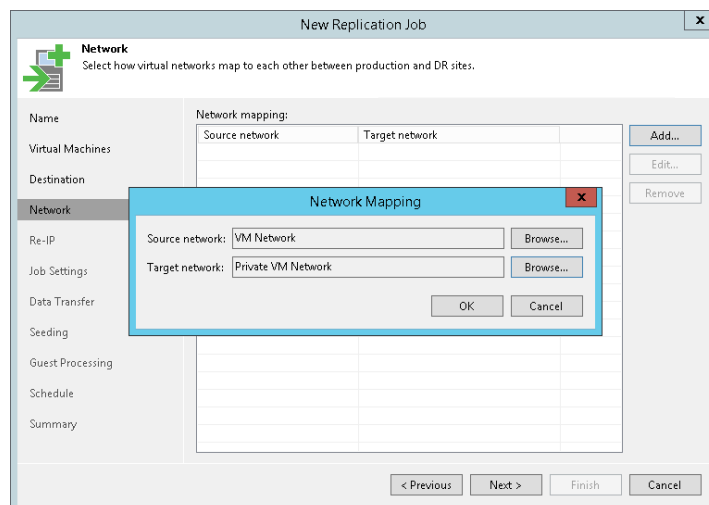


## Step 8. Create Network Map Table

The **Network** step of the wizard is available if you have selected the **Separate virtual networks** option at the Name step of the wizard. You can use this step to configure network mapping settings for the VM replica(s).

To configure a network mapping table:

1. Click **Add**.
2. Click **Browse** next to the **Source network** field and select the production network to which VMs added to the job are connected.
3. Click **Browse** next to the **Target network** field and select the network in the DR site to which VM replicas must be connected.
4. Repeat steps 2-3 for all networks to which VM replicas must be connected.



### Step 9. Configure Re-IP Rules

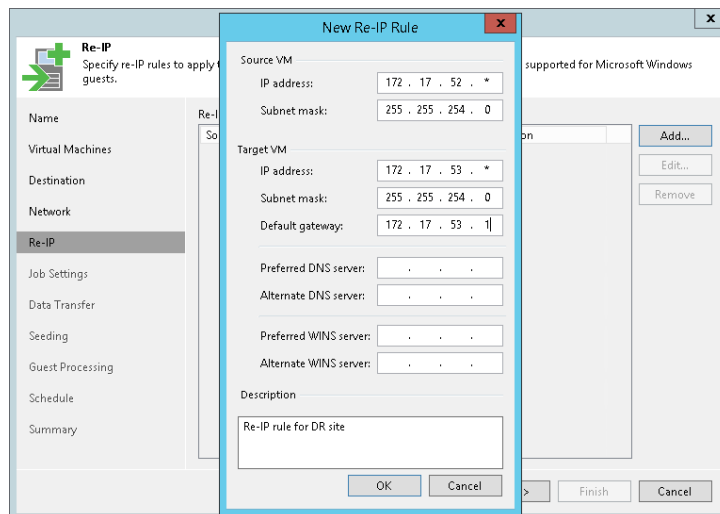
The **Re-IP** step of the wizard is available if you have selected the **Different IP addressing scheme** option at the Name step of the wizard. You can use this step to configure re-IP rules for Microsoft Windows VMs.

To configure a re-IP rule:

1. Click **Add**.
2. In the **Source VM** section, describe an IP numbering scheme adopted in the source site. To facilitate configuration, Veeam Backup & Replication detects an IP address and subnet mask for the backup server and pre-populates values in the **Source VM** section.
3. In the **Target VM** section, describe an IP numbering scheme adopted in the DR site. Specify an IP address, subnet mask and default gateway that will be used for VM replicas. If necessary, define the DNS and WINS server addresses.
4. In the **Description** field, specify a brief outline of the rule or any related comments.

**NOTE:** You can use the asterisk character (\*) to specify a range of IP addresses, for example, 172.16.17.\*.

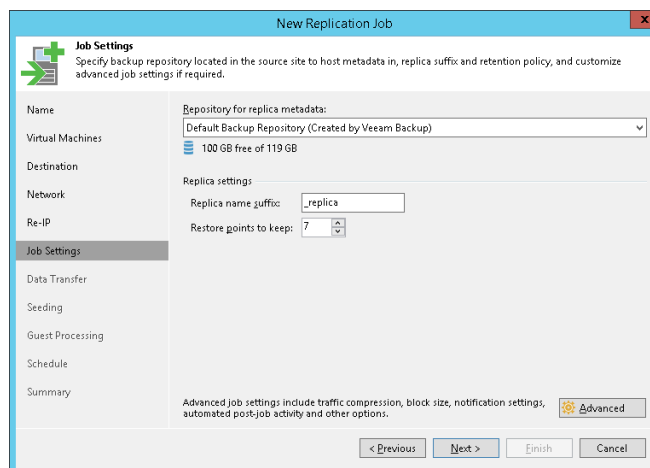
Do not use 0 to specify a range of IP addresses. In Veeam Backup & Replication, value 172.16.17.0 means a regular IP address 172.16.17.0, not an IP address range.



### Step 10. Specify Replication Job Settings

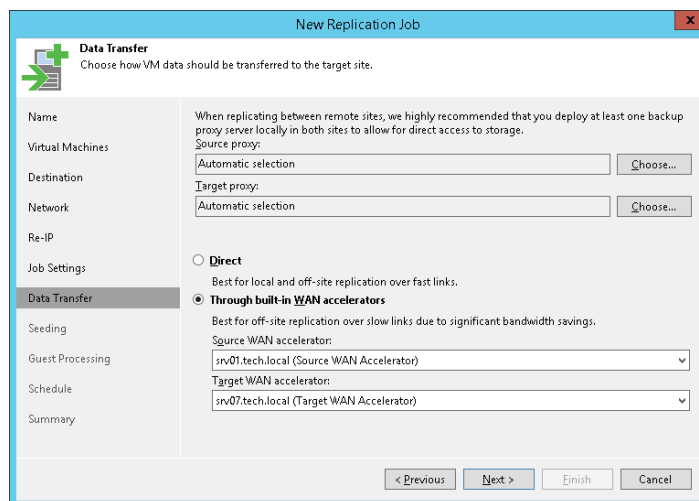
At the Job Settings step of the wizard, define replication job settings.

1. From the Repository for replica metadata list, select a backup repository that is located in the source site.
2. In the Replica name suffix field, enter a suffix for the name of VM replicas.
3. In the Restore points to keep field, specify the number of restore points that must be maintained by the replication job. If this number is exceeded, the earliest restore point will be removed.



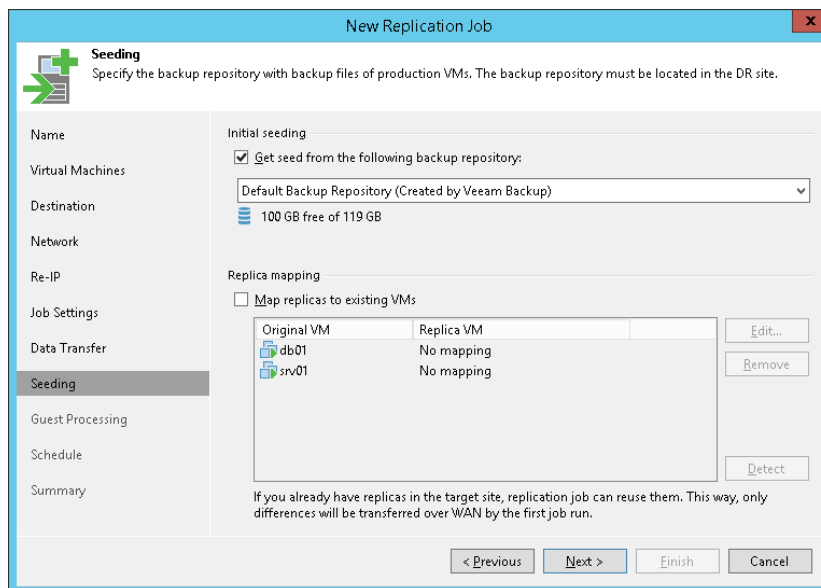
### Step 11. Specify Data Transfer Settings

At the Data Transfer step of the wizard, select backup infrastructure components that must be used for the replication process and choose a path for VM data transfer.



### Step 12. Define Seeding and Mapping Settings

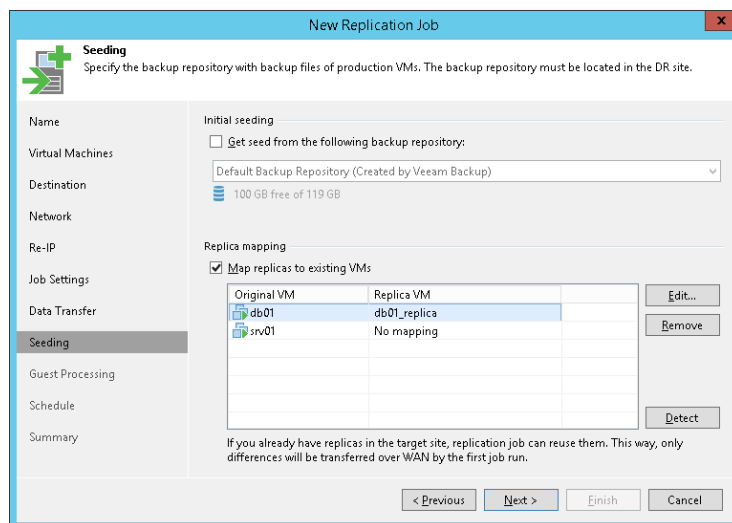
The **Seeding** step is available if you have selected the **Low connection bandwidth** option at the Name step of the wizard. You can use this step to configure replica seeding and mapping for the replication job.



### Configuring Replica Mapping

To set up replica mapping:

1. Select the **Map replicas to existing VMs** check box.
2. Click **Detect** to detect existing VM replicas. If any matches are found, Veeam Backup & Replication will populate the mapping table.



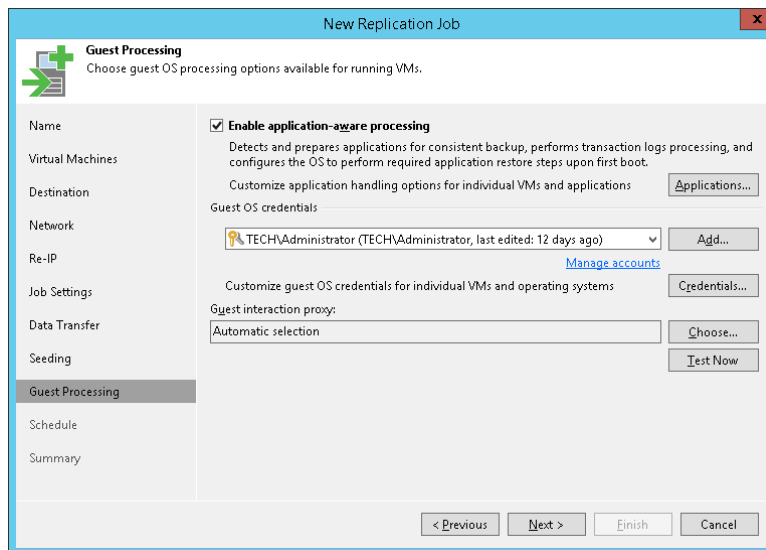
### Step 13. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, you can enable the following settings for VM guest OS processing:

You must specify a user account that will be used to connect to the VM guest OS and deploy the runtime process:

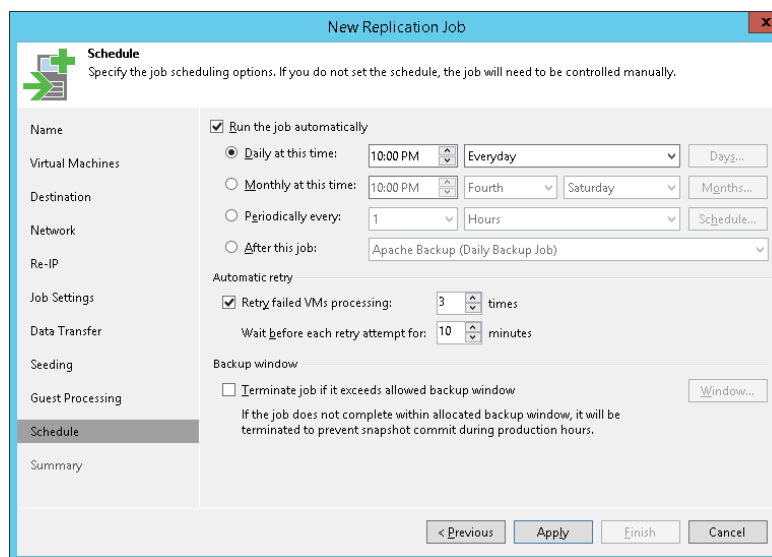
1. From the **Guest OS credentials** list, select a user account with local administrator privileges on the VM guest OS. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials.
2. By default, Veeam Backup & Replication uses the same credentials for all VMs in the job. If some VM requires a different user account, click **Credentials** and enter custom credentials for the VM.
3. If you have added Microsoft Windows VMs to the job, specify which guest interaction proxy Veeam Backup & Replication can use to deploy the runtime process on the VM guest OS. On the right of the **Guest interaction proxy** field, click **Choose**.
  - Leave **Automatic selection** to let Veeam Backup & Replication automatically select the guest interaction proxy.
  - Select **Use the selected guest interaction proxy servers only** to explicitly define which servers will perform the guest interaction proxy role. The list of servers contains Microsoft Windows servers added to the backup infrastructure.

To check if Veeam Backup & Replication can communicate with VMs added to the job and deploy the runtime process on their guest OSes, click **Test Now**. Veeam Backup & Replication will use the specified credentials to connect to all VMs in the list.



### Step 14. Define the Job Schedule

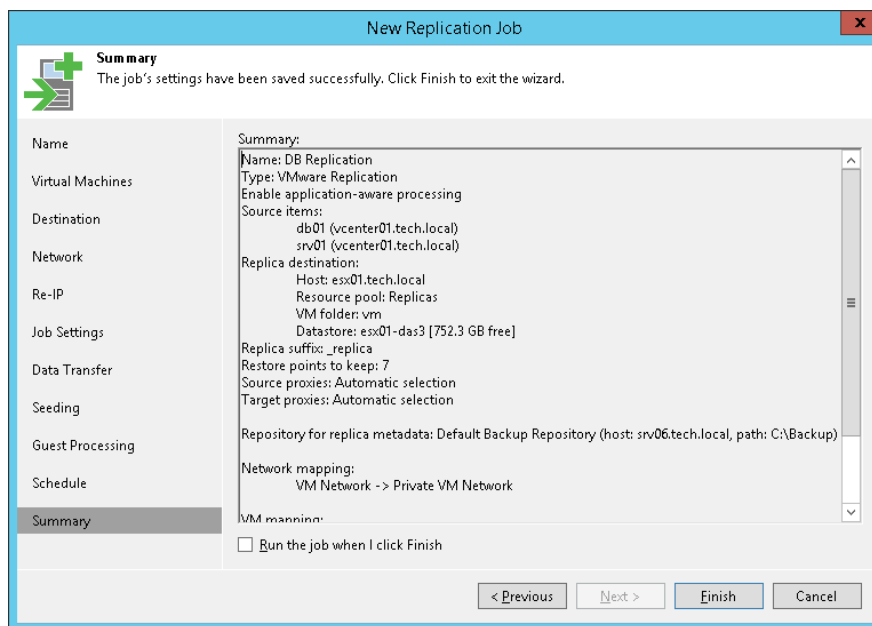
At the **Schedule** step of the wizard, select to run the replication job manually or schedule the job to run on a regular basis.



### Step 15. Finish Working with Wizard

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

1. Review details of the replication job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.



## 4. Data Recovery

Veeam Backup & Replication offers a number of recovery options for various disaster recovery scenarios:

- **Instant VM Recovery** enables you to instantly start a VM directly from a backup file
- **Entire VM recovery** enables you to recover a VM from a backup file to its original or another location
- **VM files restore** enables you to recover separate VM files (virtual disks, configuration files and so on)
- **Virtual disks restore** enables you to recover a specific hard drive of a VM from the backup file, and attach it to the original VM or to a new VM

### 4.1 Instant VM Recovery

#### 4.1.1 Performing Instant VM Recovery

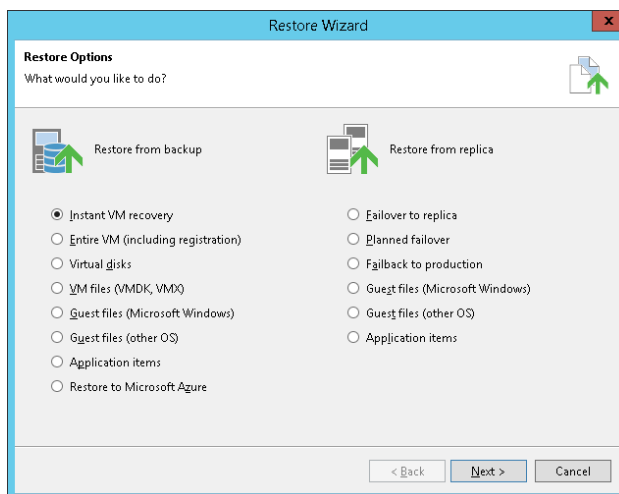
With Instant VM Recovery, you can immediately start a VM from a backup file stored on the backup repository. Instant VM Recovery accelerates the restore process, allows you to improve RTOs and decrease downtime of production VMs.

#### Step 1. Launch Instant VM Recovery Wizard

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

- On the **Home** tab, click **Restore** and select **VMware vSphere backup**. In the **Restore from backup** section, select **Instant VM recovery**.

- Open the **Home** view, in the inventory pane select **Backups**. In the working area, expand the necessary backup, select the VM you want to restore and click **Instant VM Recovery** on the ribbon.
- Open the **Home** view, in the inventory pane select **Backups**. In the working area, expand the necessary backup, right-click the VM you want to restore and select **Instant VM recovery**.



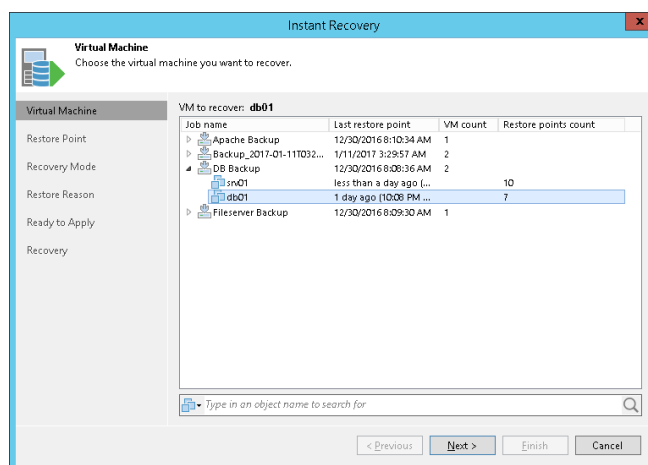
### Step 2. Select VMs

At the Virtual Machine step of the wizard, select the VM that you want to recover:

1. In the VM to recover list, expand the backup job.
2. Select the VM.

To quickly find a VM, you can use the search field at the bottom of the window.

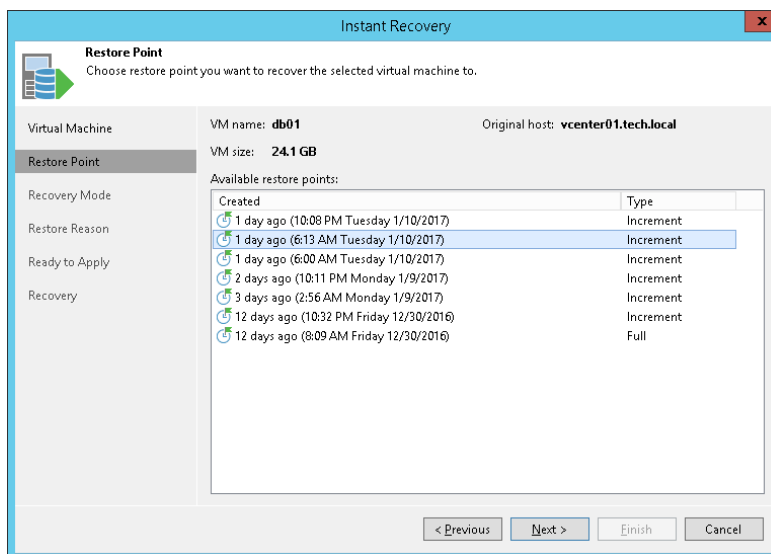
1. Enter a VM name or a part of it in the search field.
2. Click the Start search button on the right or press [ENTER].



### Step 3. Select Restore Point

At the Restore Point step of the wizard, select the restore point for the VM.





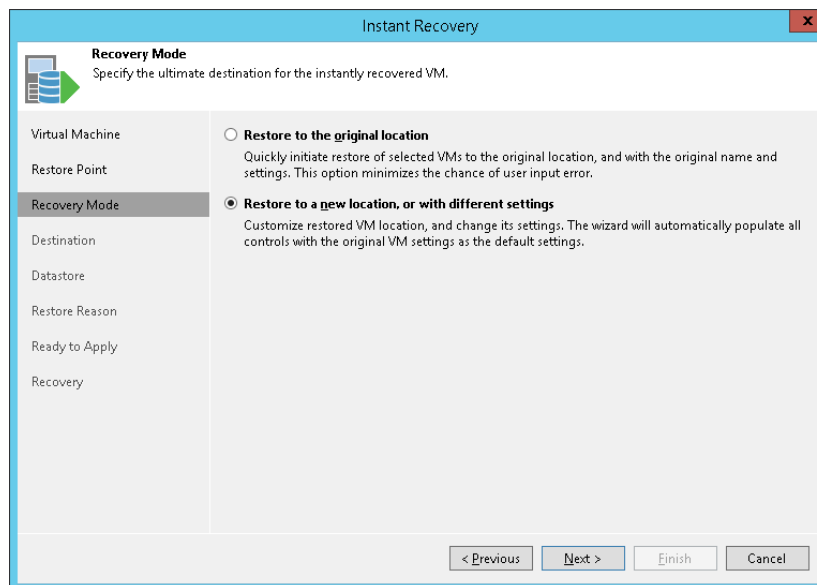
## Step 4. Select Recovery Mode

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

- Select **Restore to the original location** if you want to restore the VM with its initial settings and to its original location.
- 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 location, network settings, format of restored virtual disks and so on).

### 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.

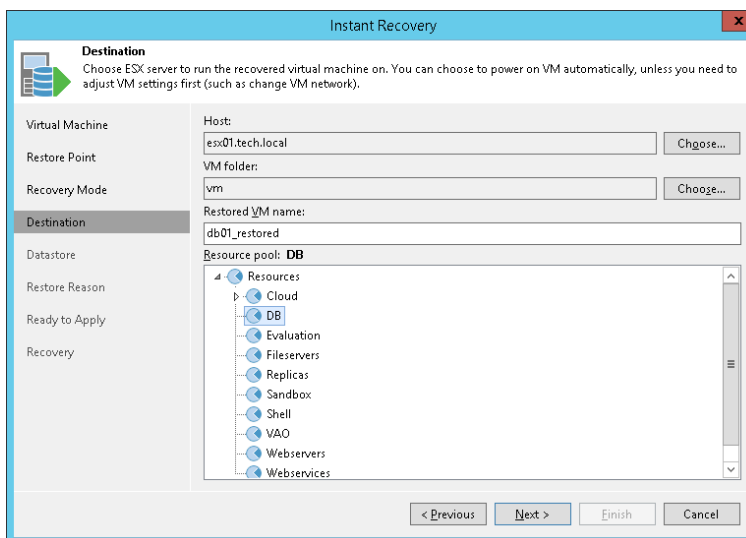


## Step 5. Select Destination for Restored VM

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

Select a destination for the restored VM:

1. In the **Host** field, specify a host on which the VM must run.
2. In the **VM folder** field, specify a folder to which the restored VM must be placed.
3. In the **Restored VM name** field, enter a name under which the VM must be restored and registered. By default, the restored VM has the name of the original VM.
4. In the **Resource pool** list, select a resource pool to which the VM must be placed.

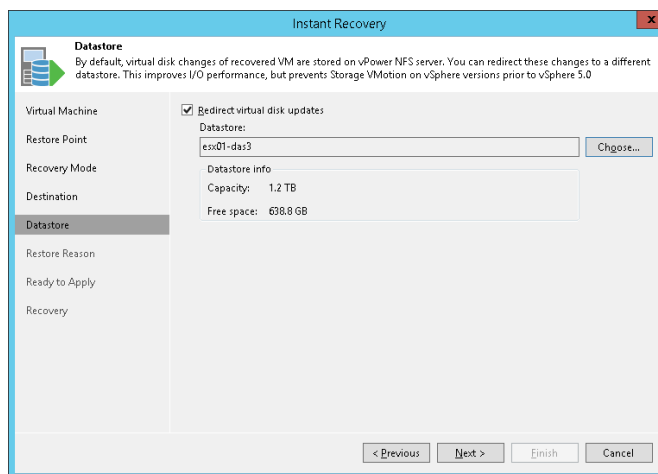


## Step 6. Select Destination for Virtual Disk Updates

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

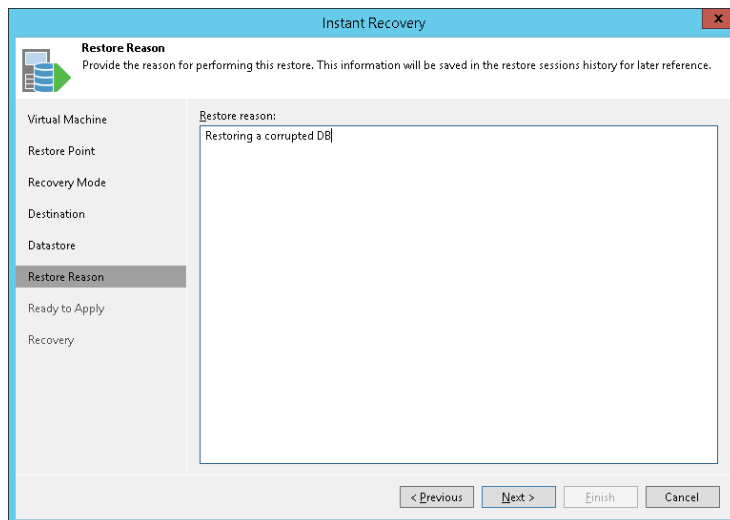
To redirect redo logs:

1. Select the **Redirect virtual disk updates** check box.
2. Choose the datastore.



### Step 7. Specify Restore Reason

At the Restore Reason step of the wizard, enter a reason for performing Instant VM Recovery for the VM.

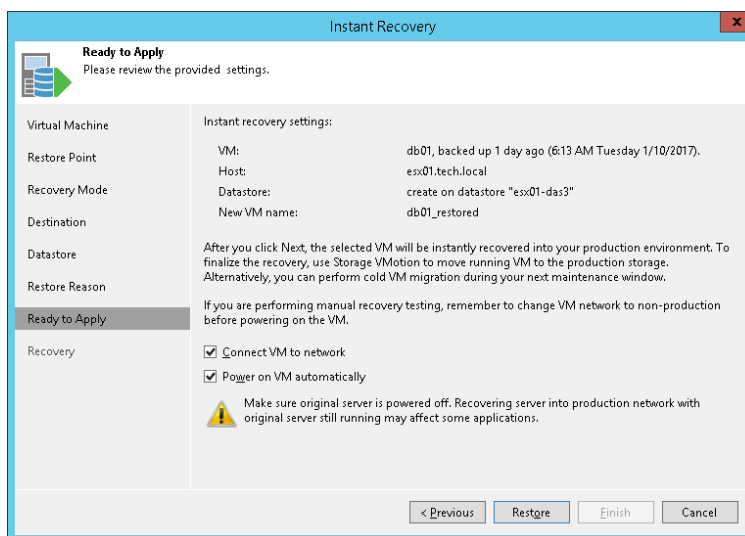


### Step 8. Verify Instant VM Recovery Settings

At the **Ready to Apply** step of the wizard, specify additional settings for Instant VM Recovery:

1. If you are recovering a production VM that has failed and want to restore it with initial network settings, select the **Connect VM to network** check box. If you are recovering a VM for testing disaster recovery while the initial VM is still running, leave this check box not selected.
2. To start a VM immediately after recovery, select the **Power on VM automatically** check box. If you are recovering the VM to the production network, make sure that the initial VM is powered off to avoid conflicts.

3. Check settings you have specified for Instant VM Recovery and click **Next**. Veeam Backup & Replication will recover the VM on the selected ESX(i) host.



### Step 9. Finalize Instant VM Recovery

All VMs restored with Instant VM Recovery are displayed in the **Home** view, under the **Instant Recovery** node.

To check the progress of Instant VM Recovery and view session details:

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

Alternatively, you can open the **History** view, select the **Instant Recovery** node under **Restore** in the inventory pane and double-click the instant VM restore session.

After the VM has been successfully recovered, you can finalize Instant VM Recovery in one of two ways:

- Migrate the recovered VM to the production environment
- Unpublish the recovered VM

### Migrating Recovered VM

To migrate a recovered VM to the production environment:

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

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

### Unpublish Recovered VM

To unpublish a recovered VM:

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

