



# Consistent protection of MySQL/MariaDB with Veeam

Pascal Di Marco  
Solutions Architect

# Contents

<b>Executive summary</b>	<b>2</b>
<b>Target audience</b>	<b>2</b>
<b>Introduction</b>	<b>2</b>
<b>Scripted database consistency at-a-glance</b>	<b>3</b>
In guest script overview	3
Editing and storing shell scripts	4
Guest account management	6
Job configuration	7
<b>Sample scripts</b>	<b>10</b>
Hot backup – Database online dump	11
Hot backup – Database freezing	12
Cold Backup – Database shutdown	14
<b>Recovery</b>	<b>17</b>
Guest recovery	17
Additional dump restoration	19
Veeam U-AIR database restoration	20
<b>About the Author</b>	<b>21</b>
<b>About Veeam Software</b>	<b>21</b>

## Executive summary

To answer the Availability challenge, modern tools should not only consider safely transferring data blocs from one container to another . Application consistency is also a key point that allows for the best possible combination of safety and Recovery Time Objectives (RTOs) to keep business-critical operations online.

This guide will outline how to achieve application consistency of MySQL and MariaDB databases with Veeam® Backup & Replication™, version 8 to version 9 .5.

## Target audience

Making today's virtual datacenter available requires more and more competencies beyond the backup tools themselves. Using modern and efficient software such as Veeam Backup & Replication necessitates a minimum amount of knowledge about the protected applications to ensure the best possible protection level. This publication is especially intended to be read by a technical audience, such as backup administrators or systems engineers.

## Introduction

As an agentless solution, Veeam Backup and Replication can be used to enforce application consistency, either via Microsoft Volume Shadow Copy Service (VSS) integration, or integration with third-party tools.

When these mechanisms are unavailable, either because the application is running on a Linux guest, or because of a non-integrated application, Veeam Backup & Replication, since version 8, offers the ability to leverage pre-freeze and post-thaw scripting, effectively allowing consistent backup or replication for any application.

## Scripted database consistency at-a-glance

### In guest script overview

The workflow on how scripts are executed during the backup process is illustrated below:

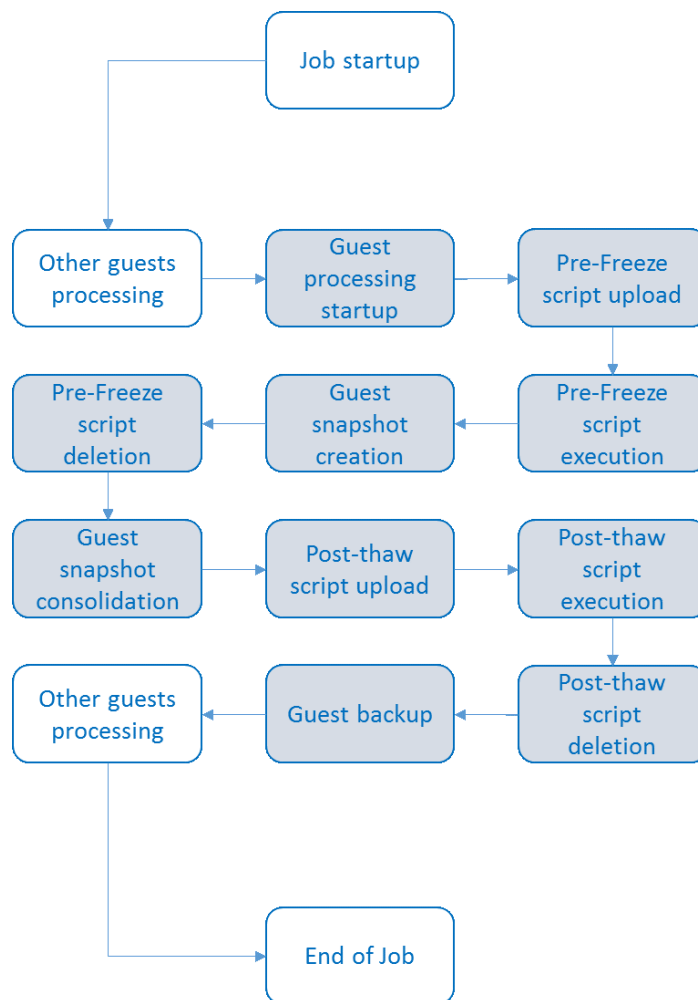


Figure 1. In guest scripting workflow

To protect a MySQL or MariaDB instance, which typically run on a Linux guest, a shell script is required, usually prefixed as `.sh`. These can be located anywhere on the Veeam Backup & Replication server.



## Editing and storing shell scripts

When storing a script on Windows, which is to be executed on a Linux guest, special attention should be paid to line endings.

**End of line** (EOL) characters are different on Windows and Linux, which means an editor should be used that supports **end of line** conversion (such as Atom, Sublime Text, Notepad++ or others).

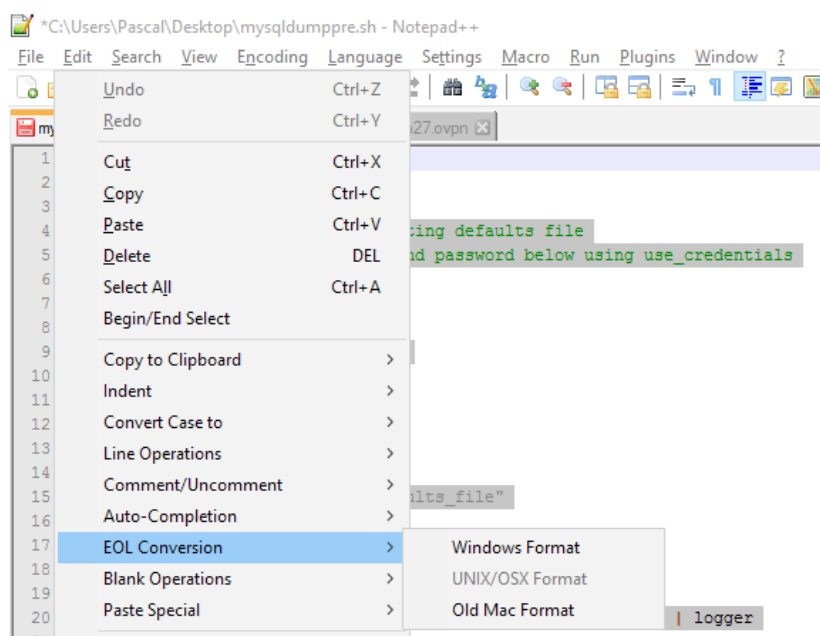


Figure 2. EOL conversion with notepad++

Scripts are stored on the Veeam Backup & Replication server and uploaded to the guest using **SSH port 22**.

The Guest Interaction Proxy, as a Windows-only guest feature, will not be used for Linux guests.

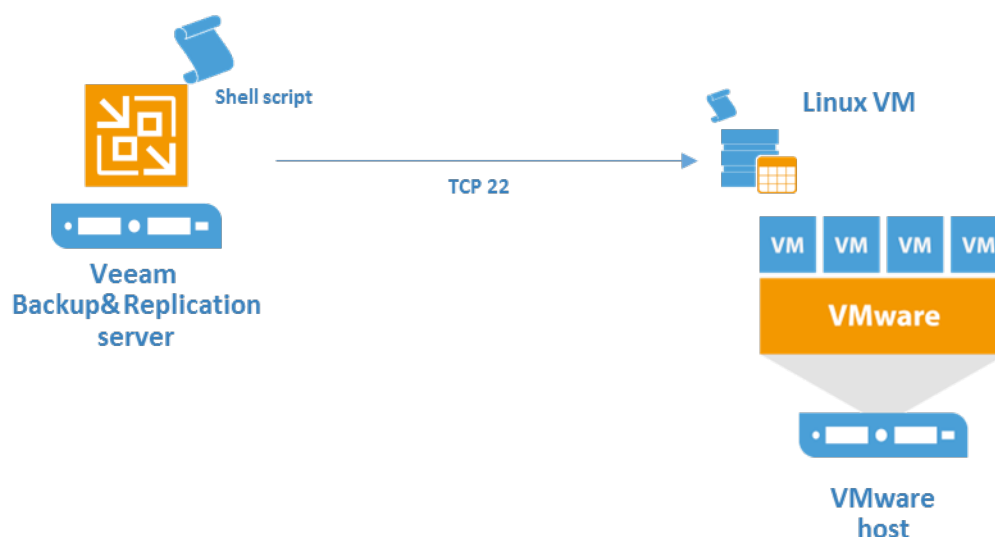


Figure 3. Shell script transfer over infrastructure

If the network connection between the Veeam Backup & Replication server and the Linux guest is unavailable, the script upload process and guest login will failover to the VMware **VIX** communication channel. Because both mechanisms require that VMware Tools be up and running on the guest, it is suggested the Veeam ONE™ **VM Configuration Assessment** report be used to detect whether VMware Tools are or are not installed the guest.

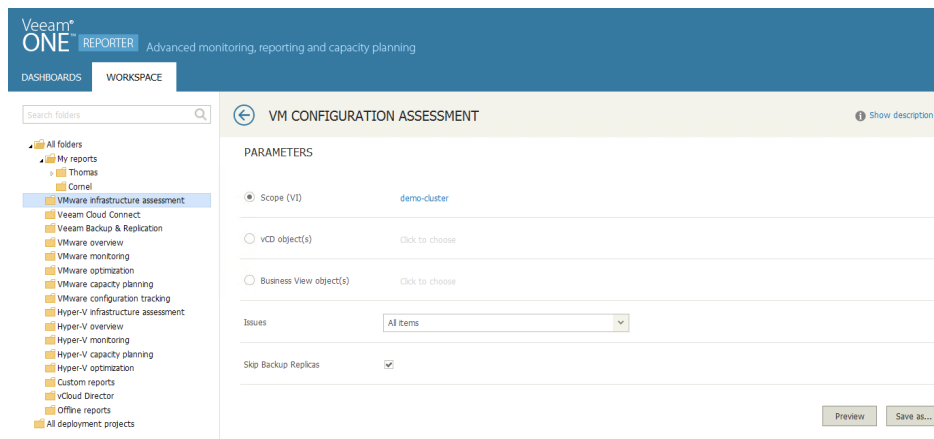


Figure 4. VM Configuration Assessment report location

Category	Potential Issue	VMs	Host	Datastore	Recommendation
Application-Aware Image Processing	VMware Tools Not Running	23			Make sure VMware Tools are up and running on these VMs
		cloudconnect2.democenter.int_s6aj0	esx1.democenter.int	Democenter\demo-netapp-nfs	
		VMware_Verification_Lab	esx1.democenter.int	Democenter\demo-netapp-nfs	
		win1	esx1.democenter.int	Democenter\demo-nimble-iscsi	
		Endpoint	esx1.democenter.int	vms-bronze\3par.vmfs.bronze.2	
		server2	esx1.democenter.int	vms-bronze\3par.vmfs.bronze.2	
		backup	esx1.democenter.int	vms-gold\3par.vmfs.gold.1	
		esxi	esx1.democenter.int	vms-gold\3par.vmfs.gold.1	
		vbr6_cert_5	esx1.democenter.int	vms-gold\vmx.vmfs.0	
		vbr6_cert	esx1.democenter.int	vms-gold\vmx.vmfs.2	
		hq-srv102	esx1.democenter.int	vms-silver\netapp1.nfs.1	
		VMware_Sandbox_Lab	esx2.democenter.int	Democenter\demo-netapp-nfs	
		win2	esx2.democenter.int	Democenter\demo-nimble-iscsi	
		gw3	esx2.democenter.int	vms-gold\3par.vmfs.gold.0	
		dr-hv1	esx2.democenter.int	vms-gold\vmx.vmfs.0	
		vbr6_cert_3	esx2.democenter.int	vms-gold\vmx.vmfs.1	
		vbr6_cert_4	esx2.democenter.int	vms-gold\vmx.vmfs.1	
		vbr6_cert_2	esx2.democenter.int	vms-gold\vmx.vmfs.2	
		Network Extension Appliance democenter(esx4)	esx4.democenter.int	Democenter\demo-cloudconnect	
		vtl	esx4.democenter.int	vms-bronze\3par.vmfs.bronze.2	
		hq-vbproxy2	esx4.democenter.int	vms-gold\vmx.vmfs.1	
		hq-vbproxy4	esx4.democenter.int	vms-gold\vmx.vmfs.1	
		sp-vbproxy1	esx4.democenter.int	vms-gold\vmx.vmfs.0	
		vbr6_cert_1	esx4.democenter.int	vms-gold\vmx.vmfs.0	

Figure 5. VM Configuration Assessment report sample

Depending of which method will be used to upload the scripts, the guest directory used may differ. If SSH is used, scripts will be temporarily stored as **/tmp/<UID>\_scriptname**. If VMware **VIX** is used, scripts will then be temporarily stored in **/tmp/vmware-root/scriptname**.

## Guest account management

Veeam Backup & Replication requires credentials to manage the Linux guest connection and script execution. To configure the credentials, select the **main menu** and open the **Credentials Management** dialog.

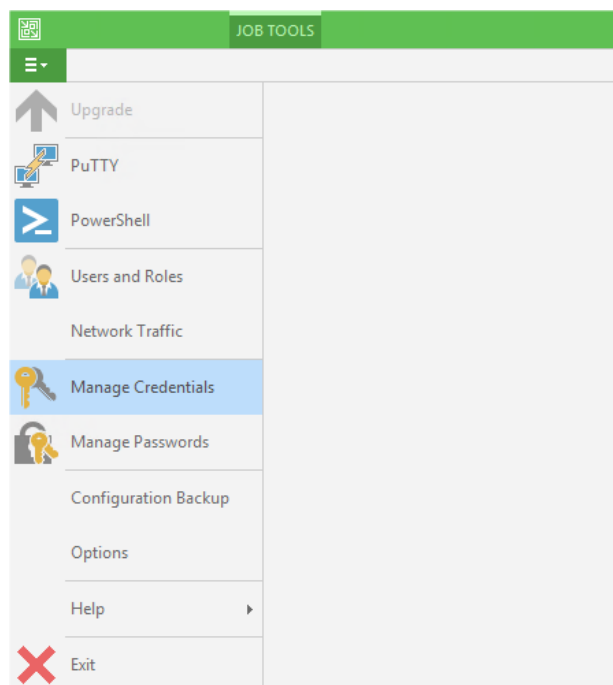


Figure 6. Guest credential management

Create and edit the user account. As a best practice, it is recommended to use the **description** field and clearly indicate the target guest for future management easiness.

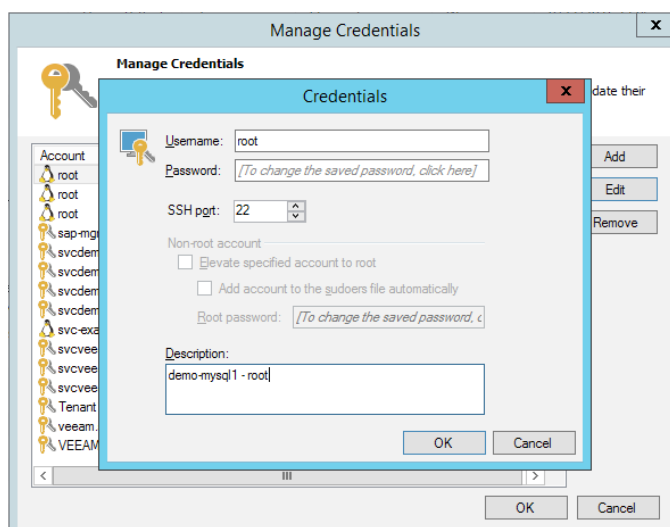


Figure 7. Guest credentials edition

For security reasons, when Veeam Backup & Replication backs up its own configuration for disaster recovery purposes into a **BCO** configuration file, passwords are not saved by default. This can be modified by enabling configuration backup file encryption: Open the **main menu**, select the **Configuration Backup** dialog and select the **encrypt configuration backup** option.




Figure 8. Encryption of a Veeam backup configuration file

The Veeam Documentation Center can be consulted for more on Veeam configuration backup here: [https://helpcenter.veeam.com/backup/vsphere/vbr\\_config.html](https://helpcenter.veeam.com/backup/vsphere/vbr_config.html).

## Job configuration

Veeam Backup & Replication gives you the ability to execute pre- and post-scripts on a per-guest basis. This means that the guests running the databases can be part of any backup or replication job.

In the job creation **Wizard**, during the **Guest Processing** step, select the **Enable application-aware image processing** option.

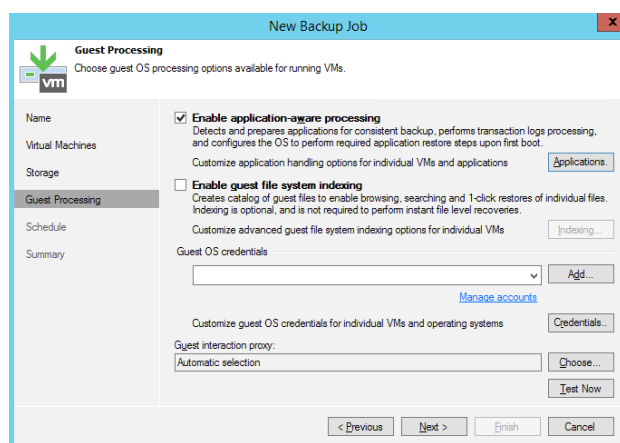


Figure 9. Enabling guest processing

Then, click on the Applications button to customize the application-aware image processing options. Select each **MySQL** and **MariaDB** server, which will use the scripts created and click **Edit**.

Depending on the job configuration, it might optionally be necessary to manually select which guests are eligible for application aware processing.

For example, with a job configured to back up a whole datastore, the default behavior will be to execute pre-freeze and post-thaw shell scripts on all Linux guests processed by the job, regardless of whether they are hosting the concerned application or not.

To avoid such a situation, it is possible to manually select which guests to apply the application-aware processing tasks to, while also making script selection granular for each and every guest.

In the **Application Aware Processing Option** window, click on the **Add** button, and select each necessary guest. It might probably be required to select **Show full hierarchy** beside the container-type selection area.

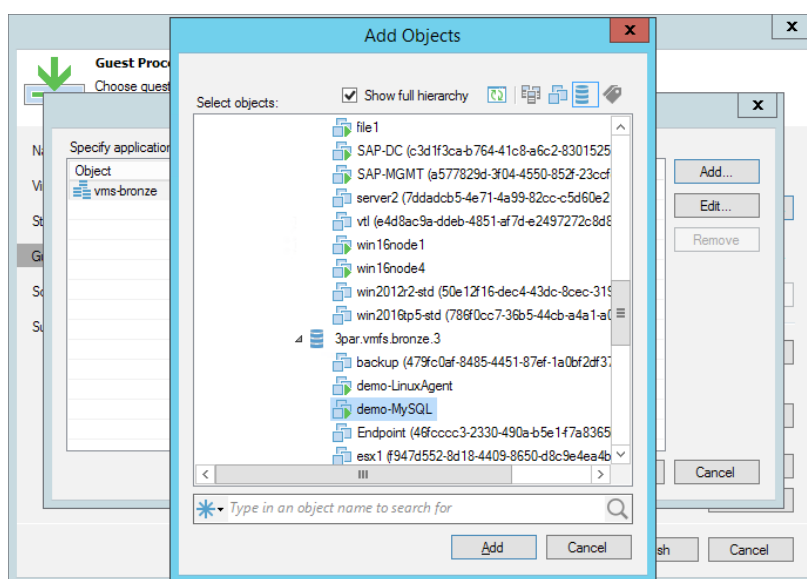


Figure 10. Adding granular object to the application-aware processing tasks

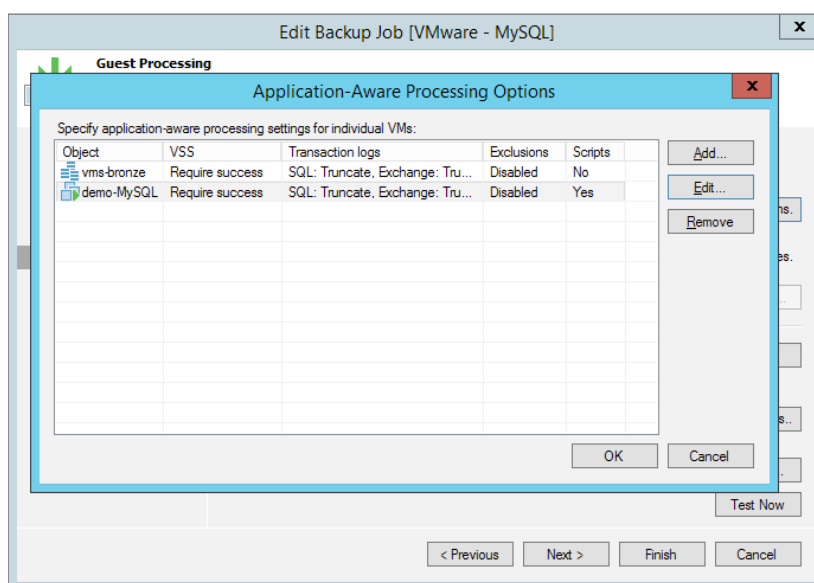


Figure 11. Editing individual guests

Once the **Processing settings** window is open, select the **Scripts** tab and browse for the pre-freeze and post-thaw in question. Select **Bash shell script files (\*.sh)** under the file types you wish to show.




Figure 12. Script file type filter

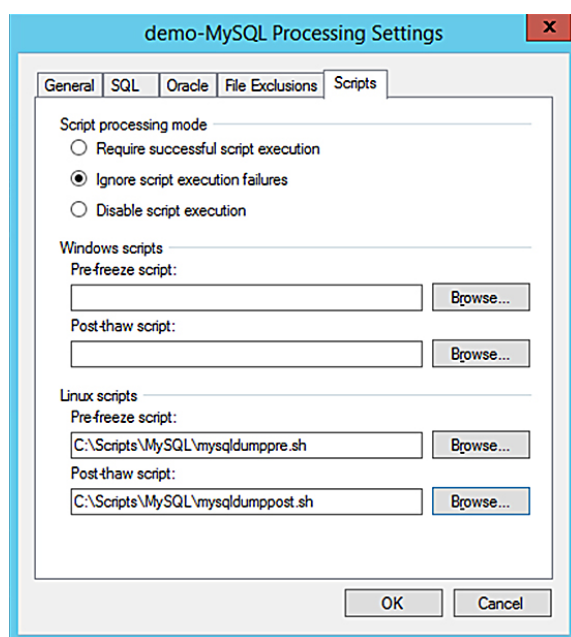


Figure 13. Scripts selection

## Sample scripts

MySQL and MariaDB can be backed up in many ways. The following three methods, however, are most commonly used.

Two of the following methods will use application commands to dump database contents or freeze writes:

### 1. Hot Backup – Database Dump

### 2. Hot Backup – Database Freeze

These commands must be executed by users with RELOAD privileges on the databases. In the following example, the user **"user"** with password **"Password"** has been granted proper privileges on the MySQL database.

The third option does not require MySQL user permissions, but does require permission to start and stop application services through **init.d** or **systemctl**:

### 3. Cold Backup – Database Shutdown

There are two authentication options available for the provided scripts:

Option 1: Hardcode the username and password in the script: **-u user -p Password**

Option 2: Use the **MySQL default configuration file** to grant users access to MySQL commands.

Usually the preferred method is to use the configuration file, because keeping the password within the same Linux installation is likely to be more secure. Depending on the MySQL install, the default file might be **/etc/my.cnf** or **/etc/mysql/debian.cnf**, or it may not even exist.

If the default file does not exist, it can be created based on default templates typically residing in **/usr/share/mysql/my-\*.cnf**.

In this configuration file, specify the user with required privileges on the database as follows:

```
vi /etc/my.cnf

[client]
user          = user
password      = Password
port          = 3306
socket        = /var/lib/mysql/mysql.sock
```

Depending on the option chosen, uncomment and modify **use\_credentials** or **default\_file** values in the pre-freeze scripts below.

## Hot backup – Database online dump

### Workflow

The pre-freeze script will dump all databases hosted on the guest to a single file under the **/tmp directory**. Before the VM snapshot creation, the **mysqldump** native command will dump a copy of the database while service will remain available.

The dump will be deleted by post-thaw script after the guest snapshot has been successful.

### Advantages

The database will stay online and read-write, while a separate and independent file will be generated.

### Disadvantages

This method is a bit more complex to set up than orchestrating a database shutdown.

The process of dumping the databases might be time consuming and affect backup windows. It also requires extra space.

To avoid over-filling issues such as a full root filesystem or inconsistent backups, it is strongly recommended to store the dump files on a dedicated and monitored filesystem.

This method will not provide the shortest possible RTO because the restore process can require that you recreate databases from the dump files.

### Pre-freeze script

```
#!/bin/bash

# config:
# when running on debian we can use existing debian-sys-maint account using defaults file
# otherwise, specify username and password below using use_credentials

#use_credentials="-uroot -p"
defaults_file="/etc/my.cnf"
dump_file="/tmp/mysql_dump.sql"
database="--all-databases"

if [ -f $defaults_file ]
then
    opts="--defaults-file=$defaults_file"
elif [ -n $use_credentials ]
then
    opts="$opts $use_credentials"
else
    echo "$0 : error, no mysql authentication method set" | logger
    exit 1
fi

opts="$opts $database"

echo "$0 executing mysqldump" | logger
mysqldump $opts >$dump_file 2>/dev/null
if [ $? -ne 0 ]
then
    echo "$0 : mysqldump failed" | logger
    exit 2
else
    echo "$0 : mysqldump succeeded" | logger
    sync;sync
fi
```



## Post-thaw script

```
#!/bin/bash
dump_file="/tmp/mysql_dump.sql"
if [ -f $dump_file ]
then
    echo "$0 deleting mysql dump file $dump_file" | logger
    rm -f $dump_file > /dev/null 2>&1
    exit 0
else
    echo "$0 could not locate mysql dump file $dump_file" | logger
    exit 1
fi
```

## Hot backup – Database freezing

### Workflow

**MySQL** tables will be flushed to disk and in **Read Only** state during snapshot creation, and writable once the VM snapshot has been created.

### Advantages

The database will stay online, but **read only**. No additional storage is required.

This method provides a short RTO because no further action than booting the restored guest is required.

### Disadvantages

The database is partially unavailable during **VM** snapshot creation. A timeout has to be set to force the release of the **read-only** state, even if the snapshot has not completed within the allotted time period. In the sample provided, the timeout was set to 300 seconds, but might be modified depending the database size and activity.

## Pre-freeze script

```
#!/bin/bash

# config:
# when running on debian we can use existing debian-sys-maint account using defaults file
# otherwise, specify username and password below using use_credentials

#use_credentials="-uroot -p"
defaults_file="/etc/my.cnf"
timeout=300
lock_file=/tmp/mysql_tables_read_lock
###

if [ -f $defaults_file ]; then
    opts="--defaults-file=$defaults_file"
fi

if [ -n $use_credentials ]; then
    opts="$opts $use_credentials"
fi

sleep_time=$((timeout+10))

rm -f $lock_file
echo "$0 executing FLUSH TABLES WITH READ LOCK" | logger
mysql $opts -e "FLUSH TABLES WITH READ LOCK; system touch $lock_file; system nohup sleep $sleep_time; system echo\ lock released|logger; " > /dev/null &
mysql_pid=$!
echo "$0 child pid $mysql_pid" | logger
c=0

while [ ! -f $lock_file ]
do
    # check if mysql is running
    if ! ps -p $mysql_pid 1>/dev/null ; then
        echo "$0 mysql command has failed (bad credentials?)" | logger
        exit 1
    fi
    sleep 1
    c=$((c+1))
    if [ $c -gt $timeout ]; then
        echo "$0 timed out waiting for lock" | logger
        touch $lock_file
        kill $mysql_pid
    fi
done
echo $mysql_pid > $lock_file
exit 0
```

## Post-thaw script

```
#!/bin/bash

lock_file=/tmp/mysql_tables_read_lock
###

mysql_pid=$(cat $lock_file)
echo "$0 sending sigterm to $mysql_pid" | logger
pkill -9 -P $mysql_pid
rm -f $lock_file
exit 0
```

## Cold Backup – Database shutdown

### Workflow

The application service will be stopped during snapshot creation, and restarted once the VM snapshot has been created.

These scripts will shut down and restart the MySQL service using **init.d** or **systemctl** commands, depending on the database packages. The account calling these scripts should have enough permissions to start and stop MySQL processes.

### Advantages

This is easy to set up and requires no extra space.

This method will also provide a short RTO, since no further action, other than booting the restored guest, is required.

### Disadvantages

The databases will be totally unavailable while the guest snapshot is created.

### Pre-freeze script for MariaDB

```
#!/bin/bash
timeout=300
if [ -f /var/run/mariadb/mariadb.pid ]
then
    mysql_pid=$(cat /var/run/mariadb/mariadb.pid) >/dev/null 2>&1
else
    echo "$0 : MariaDB not started or bad MariaDB pid file location" | logger
    exit 1
fi
echo "$0 : Processing pre-freeze backup script" | logger
systemctl stop mariadb & > /dev/null 2>&1
c=0
while [ true ]
do
    if [ $c -gt $timeout ]
    then
        echo "$0 : timed out, MariaDB shutdown failed" | logger
        exit 2
    fi
    # check if MariaDB is running
    if [ -f /var/run/mariadb/mariadb.pid ]
    then
        echo "$0 : Waiting 5 more seconds for MariaDB shutdown" | logger
        sleep 5
        c=$((c+5))
    else
        echo "$0 : MariaDB stopped" | logger
        sync;sync
        break
    fi
done
```

## Post-thaw script for MariaDB

```
#!/bin/bash
timeout=300

echo "$0 : processing post-thaw backup script" | logger

if [ -f /var/run/mariadb/mariadb.pid ]
then
    MariaDB_pid=$(cat /var/run/mariadb/mariadb.pid) >/dev/null 2>&1
    echo "$0 : MariaDB already started with PID $MariaDB_pid " | logger
    exit 1
fi

systemctl start mariadb & > /dev/null 2>&1

c=0
while [ true ]
do
    if [ $c -gt $timeout ]
    then
        echo "$0 : timed out, MariaDB startup failed" | logger
        exit 2
    fi
    # check if MariaDB is running
    if [ -f /var/run/mariadb/mariadb.pid ]
    then
        MariaDB_pid=$(cat /var/run/mariadb/mariadb.pid) >/dev/null 2>&1
        echo "$0 : MariaDB started with pid $MariaDB_pid " | logger
        break
    else
        echo "$0 : Waiting 5 more seconds for MariaDB startup"
        sleep 5
        c=$((c+5))
    fi
done
```

## Pre-freeze script for MySQL

```
#!/bin/bash
timeout=300

if [ -f /var/run/mysqld/mysqld.pid ]
then
    mysql_pid=$(cat /var/run/mysqld/mysqld.pid) >/dev/null 2>&1
else
    echo "$0 : Mysql not started or bad mysql pid file location" | logger
    exit 1
fi

echo "$0 : Processing pre-freeze backup script" | logger
/etc/init.d/mysqld stop mysql & > /dev/null 2>&1

c=0
while [ true ]
do
    if [ $c -gt $timeout ]
    then
        echo "$0 : timed out, mysql shutdown failed" | logger
        exit 2
    fi
done
```

```
# check if mysql is running
if [ -f /var/run/mysqld/mysqld.pid ]
then
    echo "$0 : Waiting 5 more seconds for mysql shutdown" | logger
    sleep 5
    c=$((c+5))
else
    echo "$0 : Mysql stopped" | logger
    sync;sync
    break
fi
done
```

## Post-thaw script for MySQL

```
#!/bin/bash
timeout=300

echo "$0 : processing post-thaw backup script" | logger
if [ -f /var/run/mysqld/mysqld.pid ]
then
    mysql_pid=$(cat /var/run/mysqld/mysqld.pid) >/dev/null 2>&1
    echo "$0 : Mysql already started with PID $mysql_pid" | logger
    exit 1
fi

/etc/init.d/mysqld start mysql & > /dev/null 2>&1

c=0
while [ true ]
do
    if [ $c -gt $timeout ]
    then
        echo "$0 : timed out, mysql startup failed" | logger
        exit 2
    fi
    # check if mysql is running
    if [ -f /var/run/mysqld/mysqld.pid ]
    then
        mysql_pid=$(cat /var/run/mysqld/mysqld.pid) >/dev/null 2>&1
        echo "$0 : MySQL started with pid $mysql_pid" | logger
        break
    else
        echo "$0 : Waiting 5 more seconds for mysql startup"
        sleep 5
        c=$((c+5))
    fi
done
```

## Recovery

Depending on the backup method used and the type of outage that occurs, either at the guest or application level, the recovery process may differ from a single file transfer to a full guest recovery.

### Guest recovery

Because the cold backup and freeze method will leave the database consistent and able to startup without additional operation, restoring the VM from the backup files is the only operation to perform.

Assuming that the Veeam Backup Repository can sustain an intensive I/O pattern (especially random read), the guest recovery might benefit from the Veeam **Instant VM Recovery** feature, which allows you to boot up the guest directly from the Veeam Backup Repository within minutes, and to live migrate it later on. Using deduplication appliances, very slow spinning drives, or low-end NAS appliances might make the Instant VM Recovery feature slower than a full guest restore.

To access the Instant VM Recovery wizard, select the **Jobs** menu on the left pane of the **Backup and Replication** main view, then launch the Restore wizard on the top ribbon. Select the proper **guest** and restore **point** to restore.

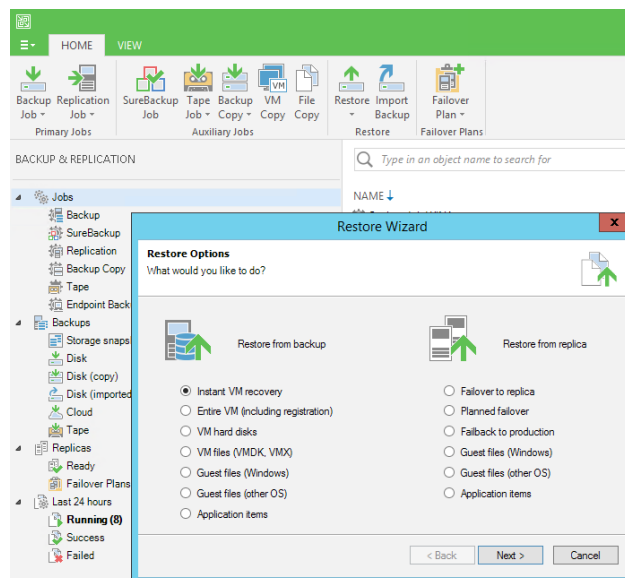


Figure 14. Launching the Instant VM Recovery wizard

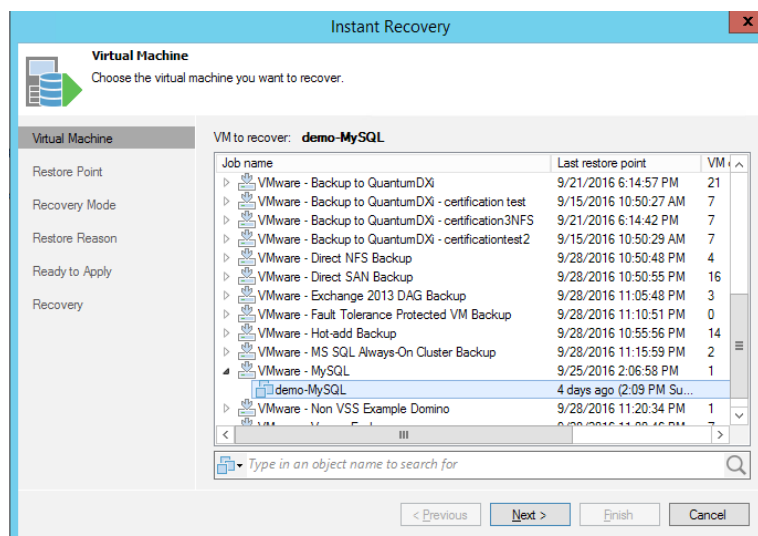


Figure 15. Selecting the guest

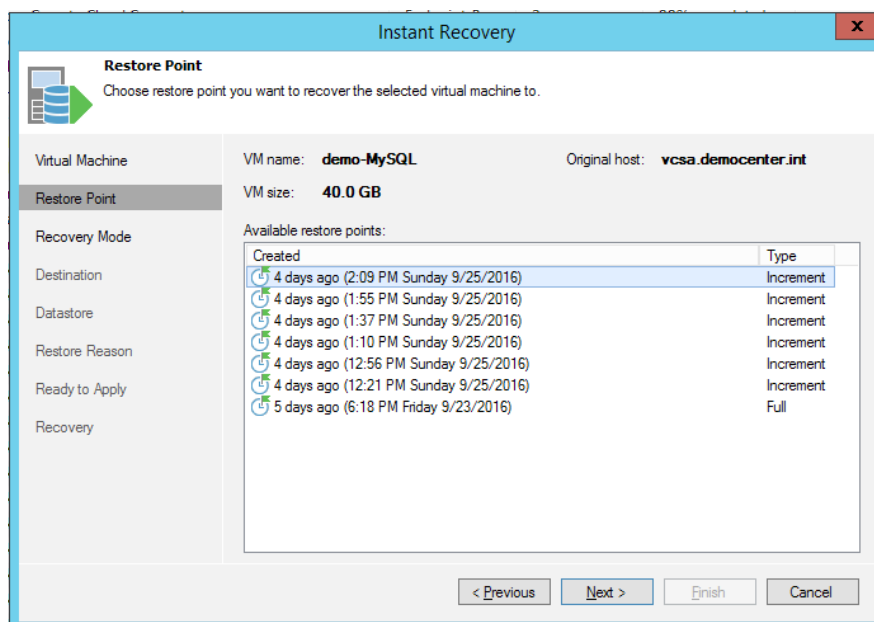


Figure 16. Selecting the restore point

Instant VM Recovery allows you to instantly boot the guest at its original location, or at any user-defined environment.

Booting the guest in a separate environment might be useful for testing purposes to modify the name of the guest or if the original location is unavailable.

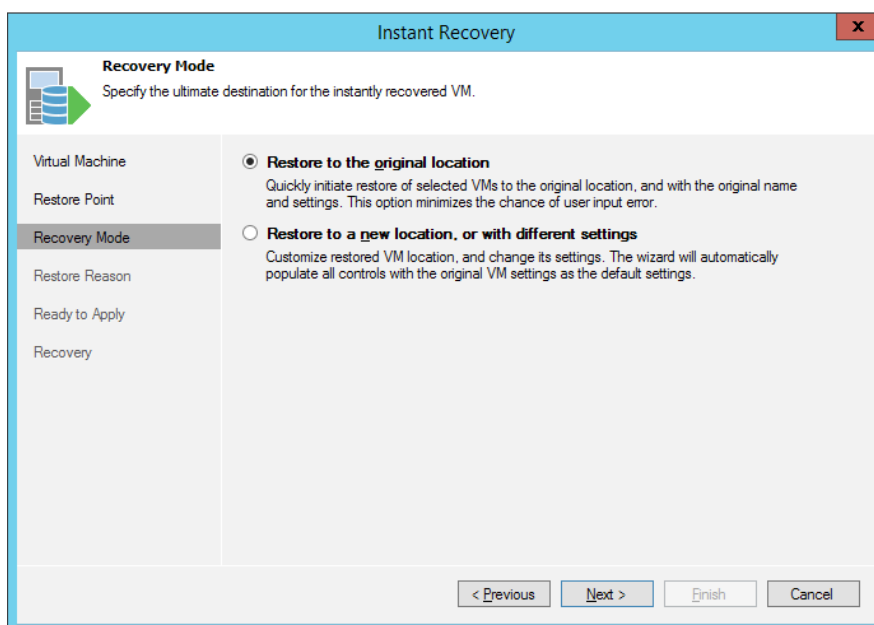


Figure 17. Restore location selection

Refer to Veeam User Guide for more about Veeam Instant VM Recovery here:

[https://helpcenter.veeam.com/backup/vsphere/instant\\_recovery.html](https://helpcenter.veeam.com/backup/vsphere/instant_recovery.html).

Once the Instant VM Recovery process ends, the guest is booted from Veeam V-Power NFS data store and left running. This situation must be considered temporary because the guest should be migrated to production storage as soon as possible using:

- Live VMware storage vMotion (if licensed) with no downtime
- Planned VMware cold vMotion with downtime
- Planned Veeam Quick Migration with limited downtime

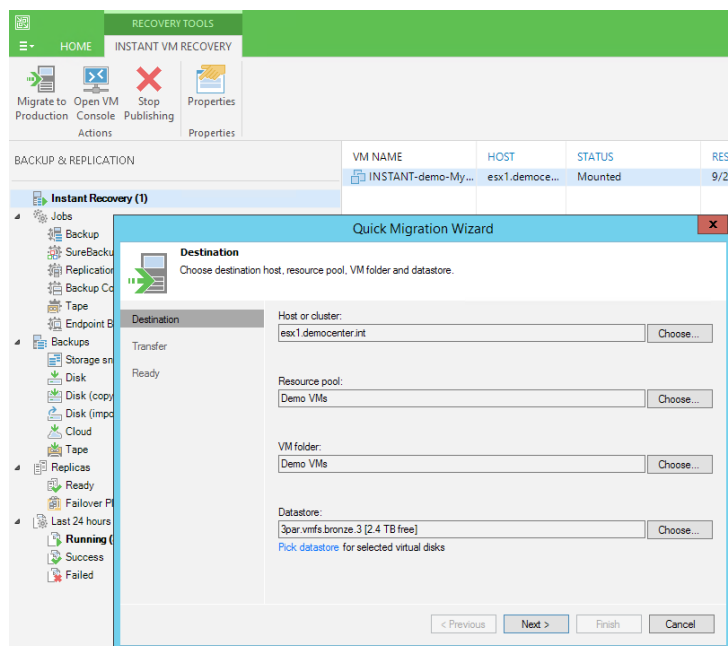


Figure 18. The Veeam quick migration wizard

Refer to the Veeam User Guide for more about Veeam quick migration here:

[https://helpcenter.veeam.com/backups/vsphere/quick\\_migration.html](https://helpcenter.veeam.com/backups/vsphere/quick_migration.html)

## Additional dump restoration

Assuming that:

1. The issue is not limited to a database outage
2. The entire VM has to be recovered from the Veeam Backup file
3. The database dump method has been used

Then, the additional operation of injecting the dump file into the database is necessary, using file redirection.

For example, if a whole set of databases has been dumped into a single file, the following command line should be used:

```
#> mysql -uuser -ppassword < /tmp/mysql-dump.sql
```

If the guest is hosting more than one database, each of which being dumped in a separate file, it is necessary to recreate the database prior to injecting the dump file:

```
#> mysql -uuser -ppassword
mysql> create database database_1;
Query OK, 1 row affected (0.02 sec)
# mysql -uuser -ppassword database_1 < /tmp/mysql_dump_database_1.sql
```



## Veeam U-AIR database restoration

To recover a database item, whether it is a granular or a full database restoration, the **Veeam U-AIR®** wizard can be used in conjunction with any relevant database management tool such as **MySQL Workbench**.

Veeam U-AIR is a unique Veeam feature that relies on V-Power NFS and Virtual Lab features. It allows you to boot the backed-up guest directly from the Veeam Backup files and into an isolated network.

Connectivity between the **recovery guest** and the production network is provided by an automatically deployed virtual appliance in charge of NAT.

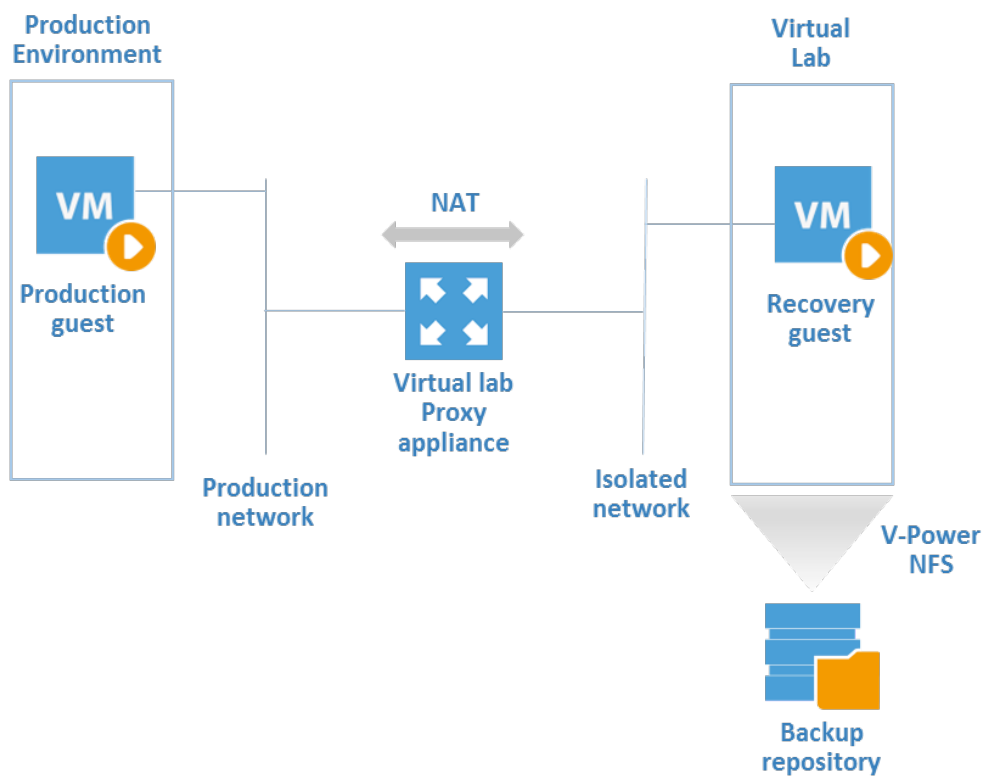


Figure 19. Veeam U-AIR

Refer to Veeam documentation for U-AIR use and requirements here:

<https://helpcenter.veeam.com/backup/70/uair/index.html>.

## About the Author



Pascal Di Marco is a Veeam Solutions Architect in charge of Southern EMEA. At Veeam since 2014, he is based in France. His former IT experience is mainly storage and virtualization consulting, especially for DELL/EMC and VMware products, which led him to focus on High Availability and get involved in providing Veeam users with the best possible resiliency.

## About Veeam Software

[Veeam](#)® recognizes the new challenges companies across the globe face in enabling the Always-On Business™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the Always-On Enterprise™ by helping organizations meet recovery time and point objectives (RTPO™) of < 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. [Veeam Availability Suite](#)™, which includes [Veeam Backup & Replication](#)™, leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 51,000 ProPartners and more than 267,500 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit <http://www.veeam.com>.

AVAILABILITY for the Always-On Enterprise™



# Veeam makes the Fortune 500 Available.

# 24.7.365

To enable **Digital Transformation**, enterprises trust Veeam  
to ensure Availability of all data and applications.



# Veeam Agent for Linux

---

Version 6

User Guide

October, 2024

© 2024 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

<b>CONTACTING VEEAM SOFTWARE .....</b>	<b>8</b>
<b>ABOUT THIS DOCUMENT .....</b>	<b>9</b>
<b>OVERVIEW .....</b>	<b>10</b>
Solution Architecture .....	11
Standalone and Managed Operation Modes .....	12
Data Backup .....	14
Backup Types .....	15
How Backup Works .....	18
Backup Job .....	21
Backup Repository .....	24
Backup Chain .....	26
Data Compression .....	32
Data Encryption .....	34
Backup of Database Systems .....	43
Backup to Object Storage .....	50
Data Restore .....	57
Volume-Level Restore .....	58
File-Level Restore .....	59
Veeam Recovery Media .....	60
Veeam Recovery Media Versions .....	61
Drivers in Veeam Recovery Media .....	62
Integration with Veeam Backup & Replication .....	63
Backup to Veeam Cloud Connect Repository .....	65
Managing Veeam Agent in Veeam Backup & Replication .....	66
<b>PLANNING AND PREPARATION .....</b>	<b>67</b>
System Requirements .....	68
System Requirements for Nosnap Veeam Agent for Linux .....	76
System Requirements for Nosnap Veeam Agent for Linux on Power .....	82
RAM Requirements for Backup of Large Number of Files .....	86
Permissions .....	89
Ports .....	91
Live Patching Support .....	96
<b>INSTALLATION AND CONFIGURATION .....</b>	<b>97</b>
Before You Begin .....	98
Installing Veeam Agent for Linux .....	99
Connecting to Veeam Software Repository .....	100
Installing Veeam Agent for Linux with Kernel Module .....	101



Installing Nosnap Veeam Agent for Linux .....	104
Installing Nosnap Veeam Agent for Linux on Power .....	105
Installing Veeam Agent for Linux in Offline Mode .....	106
Installing Veeam Agent for Linux with Kernel Module in Offline Mode .....	107
Installing Nosnap Veeam Agent for Linux in Offline Mode .....	114
Installing Nosnap Veeam Agent for Linux on Power in Offline Mode .....	118
Configuring UEFI Secure Boot .....	120
Upgrading Veeam Agent for Linux .....	125
Upgrading Veeam Agent for Linux with Kernel Module .....	126
Upgrading Nosnap Veeam Agent for Linux .....	130
Upgrading Nosnap Veeam Agent for Linux on Power .....	131
Granting Permissions to Users .....	132
Performing Initial Setup .....	133
Step 1. Accept License Agreements .....	134
Step 2. Create Custom Veeam Recovery Media .....	135
Step 3. Install Product License .....	138
Configuring Advanced Settings.....	140
Managing Veeam Agent Operation Mode .....	142
Viewing Operation Mode .....	143
Resetting to Standalone Operation Mode .....	144
Connecting to Veeam Backup & Replication.....	145
Synchronizing with Veeam Backup Server .....	146
Exporting Logs to Veeam Backup Server .....	147
Uninstalling Veeam Agent for Linux .....	148
<b>GETTING STARTED .....</b>	<b>149</b>
<b>GETTING TO KNOW USER INTERFACE .....</b>	<b>150</b>
Veeam Agent for Linux Control Panel .....	151
Command Line Interface .....	153
Viewing Help .....	155
<b>LICENSING .....</b>	<b>156</b>
Product Editions .....	157
License Agreement .....	158
Installing License .....	159
Viewing License Information .....	161
Removing License .....	162
License Expiration.....	163
Managing License with Command Line Interface .....	164
Accepting License Agreements .....	165
Installing License .....	166
Viewing License Information.....	167

Removing License .....	168
<b>PERFORMING BACKUP .....</b>	<b>169</b>
Creating Custom Veeam Recovery Media .....	170
Creating Custom Veeam Recovery Media with Control Panel .....	172
Creating Custom Veeam Recovery Media with Command Line Interface .....	175
Creating Backup Jobs .....	177
Before You Begin .....	178
Creating Backup Job with Backup Job Wizard .....	180
Creating Backup Job with Command Line Interface .....	234
Starting and Stopping Backup Jobs .....	279
Starting Backup Job from Control Panel .....	280
Starting Backup Job from Command Line Interface .....	283
Creating Active Full Backups .....	285
Stopping Backup Job .....	286
Managing Backup Jobs .....	288
Viewing List of Backup Jobs .....	289
Viewing Backup Job Settings .....	290
Editing Backup Job Settings .....	292
Deleting Backup Job .....	298
Managing Backup Repositories .....	300
Creating Backup Repository .....	301
Viewing List of Backup Repositories .....	309
Editing Backup Repository Settings .....	310
Rescanning Veeam Backup Repository .....	312
Deleting Backup Repository .....	313
Managing Veeam Backup & Replication Servers .....	314
Connecting to Veeam Backup Server .....	315
Viewing List of Veeam Backup Servers .....	317
Viewing Backup Server Details .....	318
Editing Connection to Veeam Backup Server .....	319
Updating List of Veeam Backup Repositories .....	321
Deleting Connection to Veeam Backup Server .....	322
Managing Service Providers .....	323
Connecting to Service Provider .....	324
Viewing List of Service Providers .....	325
Editing Connection to Service Provider .....	326
Updating List of Cloud Repositories .....	329
Deleting Connection to Service Provider .....	330
Managing Backups .....	331
Viewing Backups .....	332



Viewing Backup Details .....	334
Viewing Restore Points in Backup .....	336
Importing Backups .....	337
Deleting Backups .....	340
<b>PERFORMING RESTORE .....</b>	<b>341</b>
Restoring from Veeam Recovery Media .....	342
Restoring Volumes .....	343
Restoring Files and Folders .....	392
Restoring Volumes with Command Line Interface .....	432
Before You Begin .....	433
Restoring from Backup .....	434
Restoring from Restore Point .....	440
Restoring Files and Folders with Recovery Wizard .....	447
Before You Begin .....	448
Step 1. Launch File Level Restore Wizard .....	449
Step 2. Select Backup and Restore Point .....	450
Step 3. Save Restored Files .....	454
Step 4. Stop Backup Mount Session .....	455
Restoring Files and Folders with Command Line Interface .....	456
Before You Begin .....	457
Restoring from Backup .....	458
Restoring from Restore Point .....	467
Exporting Backup to Virtual Disk .....	476
Exporting Backups .....	477
Exporting Restore Points .....	479
Restoring Data from Encrypted Backups .....	481
<b>REPORTING .....</b>	<b>484</b>
Viewing Job Session Progress .....	485
Viewing Real-Time Job Session Statistics .....	486
Viewing Job Session Result .....	488
Viewing Session Status .....	490
Viewing Session Logs .....	492
<b>EXPORTING PRODUCT LOGS .....</b>	<b>493</b>
Exporting Logs with Control Panel .....	494
Exporting Logs with Command Line Interface .....	496
<b>GETTING SUPPORT .....</b>	<b>497</b>
<b>USING WITH VEEAM BACKUP &amp; REPLICATION .....</b>	<b>498</b>
Setting Up User Permissions on Backup Repositories .....	500
Managing License .....	503
Managing Instance Consumption by Veeam Agents .....	504

Assigning License to Veeam Agent .....	505
Viewing Licensed Veeam Agents and Revoking License .....	506
Performing Data Protection Tasks.....	508
Backing Up to Backup Repositories .....	509
Backing Up to Cloud Repositories .....	510
Performing Backup Copy for Veeam Agent Backups .....	512
Using SureBackup .....	513
Archiving Veeam Agent Backups to Tape .....	515
Restoring Data from Veeam Agent Backups.....	516
Restoring Veeam Agent Backup to vSphere VM .....	517
Restoring Veeam Agent Backup to Hyper-V VM .....	519
Restoring Veeam Agent Backup to Nutanix VM .....	521
Restoring Veeam Agent Backup to Proxmox VM .....	522
Restoring to Microsoft Azure .....	523
Restoring to Amazon EC2 .....	524
Restoring to Google Compute Engine .....	525
Restoring Files and Folders.....	527
Restoring Application Items .....	528
Exporting Disks.....	529
Publishing Disks.....	538
Exporting Restore Point to Full Backup File .....	549
Performing Administration Tasks .....	550
Importing Veeam Agent Backups .....	551
Enabling and Disabling Veeam Agent Backup Jobs .....	553
Viewing Veeam Agent Backup Job Statistics .....	554
Deleting Veeam Agent Backup Jobs .....	555
Viewing Veeam Agent Backup Properties .....	556
Creating Recovery Token .....	557
Removing Veeam Agent Backups .....	559
Deleting Veeam Agent Backups from Disk .....	560
Configuring Global Settings.....	561
Assigning Roles to Users .....	562

# 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](http://veeam.com/documentation-guides-datasheets.html)
- Veeam R&D Forums: [forums.veeam.com](http://forums.veeam.com)

# About This Document

This user guide provides information about Veeam Agent for Linux version 6.2.

## Intended Audience

The user guide is intended for anyone who wants to use Veeam Agent for Linux to protect their computer.

# Overview

Veeam Agent for Linux is a data protection and disaster recovery solution for physical endpoints and virtual machines running Linux-based operating systems.

Veeam Agent can be used by IT administrators who run Linux infrastructure to protect different types of computers and devices: servers, desktops and laptops. The solution runs inside the guest OS and does not need access to virtualization infrastructure components. Thus, Veeam Agent can be used to protect Linux server instances deployed in the public cloud, for example, in Microsoft Azure environment.

## NOTE

Veeam Agent can operate in either standalone or managed mode. Depending on the mode, Veeam Agent provides different features and limitations. To learn more, see [Standalone and Managed Operation Modes](#).

Veeam Agent offers a variety of features to protect your data. You can create an entire system image backup, back up specific machine volumes or individual directories and files. Backups can be stored on a local hard drive, on an external hard drive, in a network shared folder, object storage repository or Veeam backup repository.

In case of a disaster, you can perform the following restore operations:

- Start the OS from the Veeam Recovery Media and use standard Linux command line tools to diagnose and fix problems.
- Perform bare metal restore.
- Restore necessary data from backups to its original location or a new location.

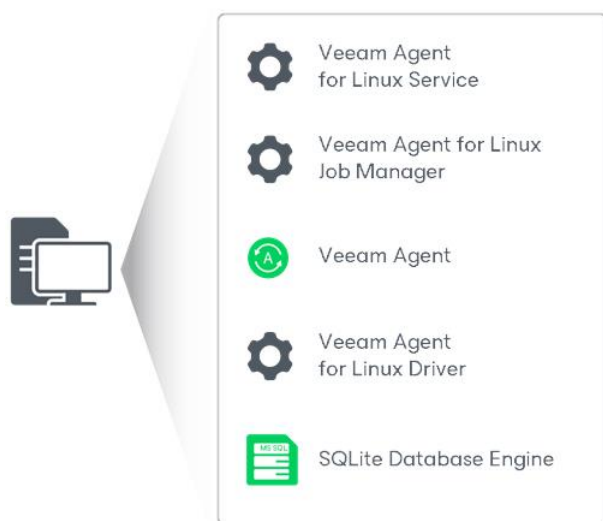
Veeam Agent integrates with Veeam Backup & Replication. Backup administrators who work with Veeam Backup & Replication can perform advanced tasks with Veeam Agent backups: restore files and disks from backups, manage backup jobs configured in Veeam Agent or backups created with these jobs.

# Solution Architecture

Veeam Agent for Linux is set up on a Linux-based physical endpoint or virtual machine whose data you want to protect.

When you install the product, Veeam Agent deploys the following components:

- *Veeam Agent for Linux Service* is a service responsible for managing all tasks and resources in Veeam Agent. The *veeamservice* component is registered as a daemon in the Linux OS upon the product installation. The service is started automatically when you start the OS and runs in the background.
- *Veeam Agent for Linux Job Manager* is a process started by *Veeam Agent for Linux Service* for every backup job session.
- *Veeam Agent* that communicates with the *Veeam Agent for Linux Service* and *Veeam Agent for Linux Job Manager*. *Veeam Agent* is started by *Veeam Agent for Linux Manager* to perform data transfer operations of any kind: copy data from the backed-up volume to the backup location during backup, from the backup location to the target volume during restore, perform data compression, and so on.
- *Veeam Agent for Linux Driver* is a Veeam driver (Linux kernel module) used to create volume snapshots in the Linux OS and keep track of changed data blocks.
- To store its configuration data, Veeam Agent uses the SQLite database engine. SQLite requires only few files to install and takes little resources to run on a Linux OS.



# Standalone and Managed Operation Modes

Veeam Agent can operate in two modes: *standalone mode* and *managed mode*. The current User Guide covers subjects related to Veeam Agent operating in the standalone mode only. Depending on the operation mode, Veeam Agent has different functionality and limitations.

## Standalone Mode

In this mode, Veeam Agent operates as a standalone product. To use Veeam Agent operating in the standalone mode, you must manually install the product directly on the computer whose data you want to protect.

For Veeam Agent operating in the standalone mode, data protection, disaster recovery and administration tasks are performed by the user. You can also use Veeam Agent operating in the standalone mode with Veeam Backup & Replication. In this scenario, you can use backup repositories managed by Veeam Backup & Replication as a target location for Veeam Agent backups and use the Veeam Backup & Replication console to perform a number of tasks with Veeam Agent backup jobs and backups. To learn more, see [Integration with Veeam Backup & Replication](#).

You can also use Veeam Backup & Replication as a gateway for creating backups targeted at the following types of repositories:

- Veeam Cloud Connect repository. To learn more, see [Backup to Veeam Cloud Connect](#).
- Object storage repository.

With Veeam Agent operating in the standalone mode, you can also back up data directly to an object storage repository. To learn more about both options, see [Backup to Object Storage](#).

## Managed Mode

In this mode, Veeam Agent operates under control from one of the following Veeam products:

- **Veeam Backup & Replication**

You can automate management of Veeam Agents on multiple computers in your infrastructure in the Veeam Backup & Replication console. You can configure Veeam Agent backup policies and perform other data protection and administration tasks on remote computers.

To use Veeam Agent operating in the managed mode, you must deploy the product in one of the following ways:

- From Veeam Backup & Replication
- Manually using external tools

To learn more about managed Veeam Agent deployment, see the [Protected Computers Discovery and Veeam Agent Deployment](#) section in the Veeam Agent Management User Guide.

For Veeam Agent managed by Veeam Backup & Replication, data protection, data restore and administration tasks are performed by a backup administrator in the Veeam Backup & Replication console. To learn about managing Veeam Agent in Veeam Backup & Replication, see the [Veeam Agent Management Guide](#).

- **Veeam Service Provider Console**

You can use Veeam Service Provider Console to manage Veeam Agents on multiple computers in your infrastructure. When Veeam Agent is managed by Veeam Service Provider Console, you can configure backup job settings, start and stop backup, change global settings, update and uninstall Veeam Agent and collect Veeam Agent data for monitoring and billing.

To manage Veeam Agent from Veeam Service Provider Console, you must install Veeam Service Provider Console management agent and Veeam Agent on the computer whose data you want to protect. After that, in Veeam Service Provider Console, you must activate Veeam Agent on the protected computer to set it into the managed operation mode.

For Veeam Agent managed by Veeam Service Provider Console, data protection, data restore and administration tasks are performed by a backup administrator in Veeam Service Provider Console.

Backup administrator can enable a read-only access mode for Veeam Agent installed on the protected computer. When you work directly with Veeam Agent operating in the read-only access mode, you can perform a limited set of operations, including:

- Running the backup job manually.
- Viewing backup session statistics.
- Restoring individual files.

To learn about deploying and managing Veeam Agent with Veeam Service Provider Console, see [Veeam Service Provider Console User Guides](#). Select the guide that suits your user role.



# Data Backup

It is recommended that you regularly back up data stored on your machine. Backup creates a safety copy of your data. If any kind of disaster strikes, you can restore your data from the backup and be sure that you will not lose the necessary information.

You can set up Veeam Agent to perform automatic scheduled backups (triggered at specific time of the day), or you can choose to back up data manually when needed. You can back up the entire computer image, specific computer volumes or individual directories and files.

You can set up Veeam Agent to create multiple backups — with individual backup scope, upon individual schedule or in different locations. This functionality is available if Veeam Agent operates in the Server edition. To learn more about editions, see [Product Editions](#).

Backups created with Veeam Agent can be saved to the following locations:

- Removable storage device
- Local computer drive
- NFS or SMB (CIFS) network shared folder
- Veeam backup repository managed by a Veeam backup server
- Veeam Cloud Connect repository
- Object storage repository

# Backup Types

Veeam Agent for Linux lets you create the following backup types:

- [Volume-level backup](#)
- [File-level backup](#)

## Volume-Level Backup

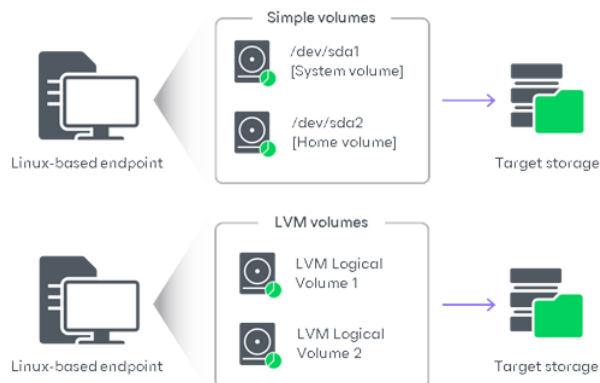
You can set up Veeam Agent for Linux to create volume-level backup. The volume-level backup captures the whole image of a data volume on your computer. You can use the volume-level backup to restore a computer volume, specific files and folders on the volume or perform bare metal recovery.

Veeam Agent for Linux supports backup of the following types of computer volumes:

- Simple volumes
- LVM logical volumes
- BTRFS subvolumes

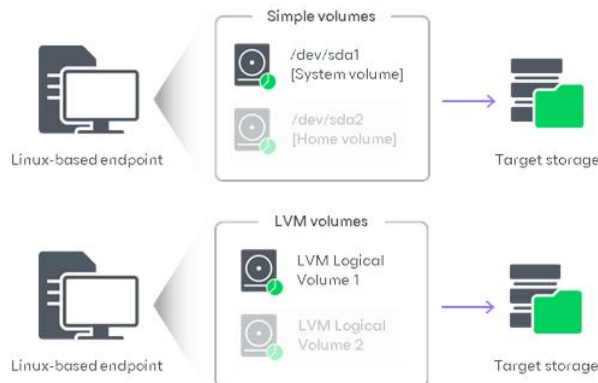
You can back up all computer volumes or specific computer volumes.

- When you back up the entire computer image, Veeam Agent captures the content of all volumes on your computer. The resulting backup file contains all volume data and Linux OS system data: system partition, partition table and bootloader.



- When you back up a specific computer volume, Veeam Agent captures only the data that resides on this specific volume: files, folder, application data and so on.

If you choose to back up the system volume (volume to which the root file system is mounted), Veeam Agent automatically includes the bootloader into the backup scope.

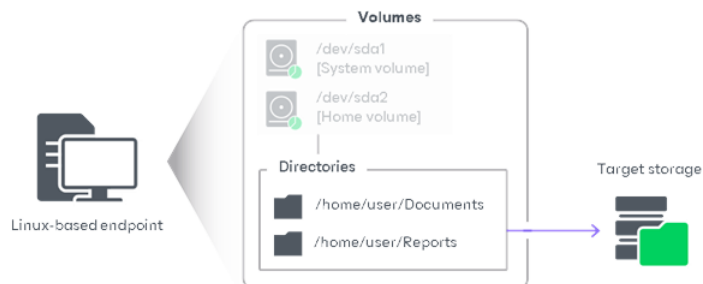


## File-Level Backup

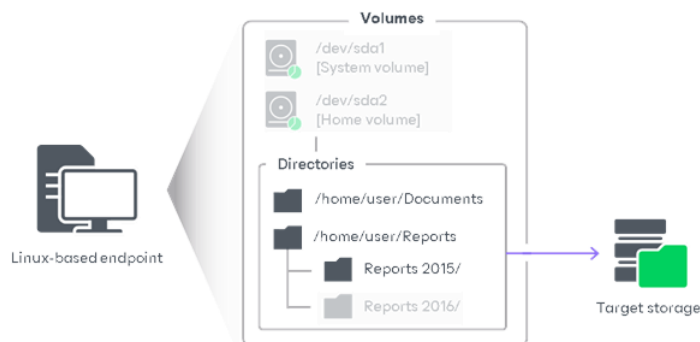
You can set up Veeam Agent for Linux to create file-level backup. The file-level backup captures only data of individual directories and files on the computer. You can use the file-level backup to restore files and directories that you have added to the backup scope.

With Veeam Agent for Linux, you can specify which files and directories to back up:

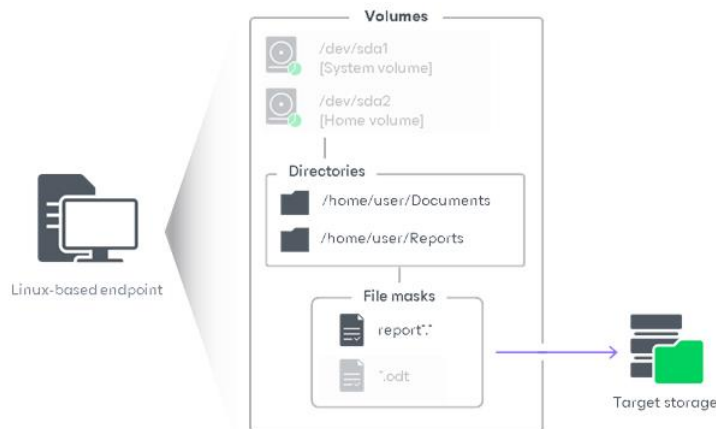
- You can include individual directories in the backup. When you include a directory in the backup, its subdirectories are automatically included in the backup too. When you recover from such backup, you will be able to restore directories that you have selected to back up, all subdirectories of these directories and files in these directories.



- You can exclude from the backup some subdirectories of the directories that are included in the backup. When you recover from such backup, you will be able to restore directories that you have selected to back up, specific subdirectories of these directories and files in these directories.



- You can include or exclude files of a specific type in/from the backup. You can specify file names explicitly or use UNIX wildcard characters to define include and exclude file name masks. When you recover from such backup, you will be able to restore directories that you have selected to back up with files whose names match the specified include masks.



## Snapshot-Less File-Level Backup

You can set up Veeam Agent for Linux to create file-level backup in the snapshot-less mode. This allows you to back up data that resides in any file system mounted to the root file system of the Veeam Agent computer. For example, you can use the snapshot-less mode to back up data that resides in a file system that is not supported for snapshot-based backup with Veeam Agent, such as UFS, ZFS, GFS, GFS2, OCFS2 or bcache fs. You can also use it to back up data that resides in an NFS or CIFS network shared folder.

To create backups in the snapshot-less mode, you must enable this mode in the properties of the file-level backup job. To learn more, see [Creating Backup Jobs](#).

In the snapshot-less mode, Veeam Agent does not create a snapshot of the backed-up volume. Instead, when the backup process starts, Veeam Agent reads files and directories that you selected to back up, and copies backed-up data to the target location.

### IMPORTANT

During backup in the snapshot-less mode, Veeam Agent does not track whether files and directories have changed in their original location since the time when the backup process started. To make sure that data in the backup is in the consistent state, you must not perform write operations in the file system that contains the backed-up data until the backup process completes.

# How Backup Works

Veeam Agent for Linux performs backup differently depending on the backup type:

- [Volume-level backup](#)
- [File-level backup](#)

## How Volume-Level Backup Works

During volume-level backup, Veeam Agent performs the following operations for every [backup job session](#):

1. When a new job session starts, Veeam Agent creates a backup file in the target location.
2. In the backup file, Veeam Agent creates a disk for each backed-up disk. In disks, Veeam Agent creates blank partitions that have the same size and location as partitions in backed-up disks.
3. Veeam Agent creates a snapshot of the volume whose data you want to back up. The snapshot is created on the volume that has enough free disk space to contain the snapshot data. To create a snapshot, Veeam Agent uses the *Veeam Agent for Linux Driver*.

The snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup. If a data block is about to change on disk during backup, Veeam Agent will copy this block to the snapshot. After the data block is overwritten on the source location, its original copy will remain intact in the snapshot.

### NOTE

Consider the following:

- If you instruct Veeam Agent to back up a database system, Veeam Agent prepares databases for backup before creating a snapshot of the volume. To learn more, see [Backup of Database Systems](#).
- During backup of data that resides in the BTRFS file system, Veeam Agent does not use its driver to create a snapshot. Instead, Veeam Agent leverages BTRFS capabilities to create a BTRFS snapshot.

4. [For incremental backup] Veeam Agent uses the *Veeam Agent for Linux Driver* to detect what blocks have changed on the volume since the previous job session. The driver keeps this information as a changed block tracking map in the RAM of your computer.

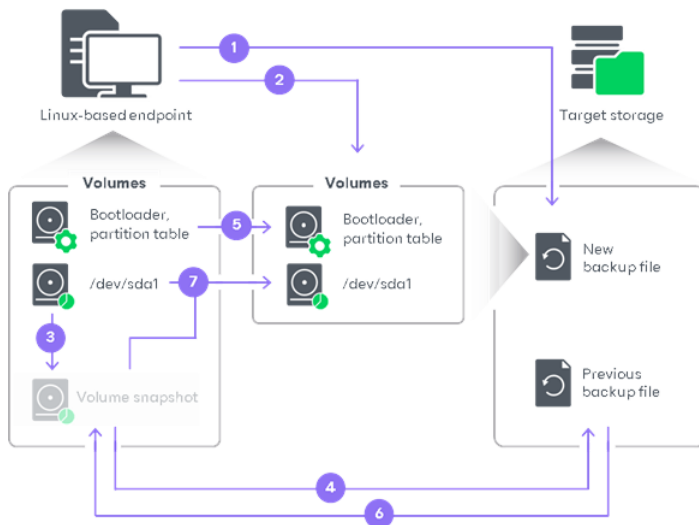
Mind that every time the driver is unloaded or the Veeam Agent computer is rebooted, the changed block tracking map is reset as well. In such case, to detect what data blocks have changed since the previous job session, Veeam Agent rescans the entire data added to the backup scope and creates a new changed block tracking map. In this case, backup requires greater time.

To learn about full and incremental backup, see [Backup Chain](#).

5. Veeam Agent copies the partition table and bootloader located on the hard disk to the backup file in the target location.
6. [For incremental backup] Veeam Agent calculates checksums for each data block and compares them with checksums from the backup file created during the previous job session. If checksums do not match, Veeam Agent will copy the data block to the target location during the next backup process step.
7. Veeam Agent copies data from the following sources:
  - Data that did not change on disk during backup is transferred from the source volume.

- Data that changed on disk during backup is transferred from the snapshot.

After all the data is transferred, Veeam Agent removes the snapshot.



## How File-Level Backup Works

During file-level backup, Veeam Agent performs the following operations for every [backup job session](#):

1. When a new job session starts, Veeam Agent creates a backup file in the target location.
2. In the backup file, Veeam Agent creates a disk. The disk contains a volume with the ext4 file system.
3. Veeam Agent creates a snapshot of the volume which data you want to back up. The snapshot is created on the volume that has enough free disk space to contain the snapshot data. To create a snapshot, Veeam Agent uses the *Veeam Agent for Linux Driver*.

The snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup. If a data block is about to change on disk during backup, Veeam Agent will copy this block to the snapshot. After the data block is overwritten on the source location, its original copy will remain intact in the snapshot.

### TIP

Consider the following:

- You can also set up Veeam Agent to create a file-level backup in the snapshot-less mode. This mode allows you to back up data that resides in any file system mounted to the root file system of the Veeam Agent computer. However, Veeam Agent does not track whether source files have changed since the backup process start. To learn more, see [Snapshot-Less File-Level Backup](#).
- Compared to the volume-level backup, the file-level backup, Veeam Agent does not provide changed block tracking mechanism and does not split source files into data blocks. As a result, if you plan to back up a significant amount of data, the file-level backup will require greater time, and created backup files will have greater size.

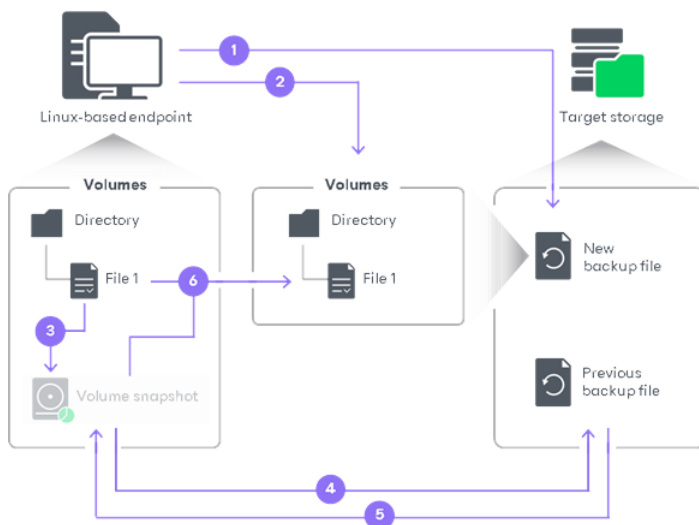
For example, you have a 1 GB file, and since the previous backup session only one data block of this file has changed. In case of the file-level backup, Veeam Agent will send the whole 1 GB file to the target again.

4. [For incremental backup] To detect files that changed on the Veeam Agent computer since the previous backup session, Veeam Agent reads file metadata and compares last modification time of files in the original location and files in the backup created during the previous job session. If the file has modification time later than the previous job session start time, Veeam Agent considers the file as changed.

To learn about full and incremental backup, see [Backup Chain](#).

5. [For incremental backup] Veeam Agent calculates checksums for each data block and compares them with checksums from the backup file created during the previous job session. If checksums do not match, Veeam Agent will copy the data block to the target location during the next backup process step.
6. Veeam Agent copies data that you selected for backup to the target location. As part of this process, Veeam Agent performs the following operations:
  - a. Enumerates all files in the source location.
  - b. For each enumerated file, creates a target file in the volume inside the backup file.
  - c. Opens the source and the target files.
  - d. Copies file data to the target location from the following sources:
    - Data blocks that did not change on disk during backup are transferred from the source volume.
    - Data blocks that changed on disk during backup are transferred from the snapshot.
  - e. Closes the source and target files.

After all backed-up files and directories are transferred, Veeam Agent removes the snapshot.



# Backup Job

To back up your data, you must configure a backup job. The backup job settings define what data you want to back up, what the target location and retention policy for created backups are and how to back up your data. If necessary, you can re-configure the backup job and change its settings at any time.

## NOTE

You cannot change the backup job type from volume-level to file-level, and vice versa.

In Veeam Agent for Linux, you can configure several backup jobs with different settings. For example, you can configure one backup job to create volume-level backup and another backup job to create file-level backup. You can configure backup jobs targeted at different backup locations to keep several copies of your backed-up data. You can also configure several backup jobs with individual schedule to fine-tune automatic backup creation process.

## NOTE

You can create more than one backup job only if Veeam Agent operates in the Workstation or Server edition. To learn more, see [Product Editions](#).

Veeam Agent launches the backup job according to the schedule you define. You can schedule the job to start at specific time daily or on specific week days. You can also start a backup job manually to perform backup on demand when needed.

Backup job scheduling settings are configured globally for all accounts of the Linux OS. As a result, Veeam Agent can start a backup job automatically regardless of the currently running user session.

## Backup Job Scripts

You can instruct Veeam Agent for Linux to run custom scripts within the backup job session:

- [Pre-job and post-job scripts](#) – Veeam Agent runs these scripts before the backup job starts and after the backup job completes. You can use pre-job and post-job scripts, for example, to configure email notifications about jobs performed by Veeam Agent.
- [Pre-freeze and post-thaw scripts](#) (in the Server edition only) – Veeam Agent runs these scripts before and after creating a snapshot. For example, the pre-freeze script may quiesce the file system and application data to bring the Linux OS to a consistent state before Veeam Agent creates a snapshot. After the snapshot is created, the post-thaw script may bring the file system and applications to their initial state.

Consider the following about using backup job scripts:

- Scripts must be created beforehand. You must specify paths to them in the job settings. Veeam Agent supports scripts in the SH file format.
- Scripts must have UNIX line endings (LF).
- Script settings are enabled at the job level. If Veeam Agent operates in the Server edition and you want to configure multiple backup jobs, you can specify individual scripts for each job.
- If you use relative paths in your scripts, during script execution such paths will refer to the root directory. For example, the script may have an output that must be saved to a new file. If you specify a relative path to that file or only a file name, the file will be created in the root directory. To specify a different location for a file, use a full absolute path.



# Pre-Job and Post-Job Scripts

You can instruct Veeam Agent for Linux to run custom pre-job and post-job scripts. Veeam Agent executes the pre-job script directly before the backup job starts. After the backup job completes, Veeam Agent executes the post-job script.

Veeam Agent starts the backup job regardless of the pre-job script result. If the pre-job script fails to execute, Veeam Agent will always start the backup job. Then, after the backup job completes, Veeam Agent will execute the post-job script.

The script is considered to be executed successfully if a "0" is returned.

The default time period for script execution is 10 minutes. After this period expires, Veeam Agent stops executing the script and displays a warning message in the job session. You can change the script timeouts in the `/etc/veeam/veeam.ini` configuration file with the `timeoutPrePost` parameter where the time period is set in seconds.

## Pre-Freeze and Post-Thaw Scripts

You can instruct Veeam Agent for Linux to run custom pre-freeze and post-thaw scripts. Veeam Agent executes the pre-freeze script before creating a snapshot. After the snapshot is created, Veeam Agent executes the post-thaw script.

### NOTE

Veeam Agent does not execute pre-freeze and post-freeze scripts if during the backup job a snapshot is not created.

The script is considered to be executed successfully if a "0" is returned.

By default, if the pre-freeze or post-thaw script fails to execute, Veeam Agent does not start the backup job. However, you can instruct Veeam Agent to ignore errors that occur during the script execution process. To allow Veeam Agent to start backup jobs regardless of the script execution result, in the `/etc/veeam/veeam.ini` configuration file, uncomment the `ignoreFreezeThawFailures` parameter and set its value to `true`.

If Veeam Agent is set up to ignore script errors, and the pre-freeze or post-thaw script fails to execute, Veeam Agent will start the backup job. After the job successfully completes, Veeam Agent will display the *Warning* status for the job session.

The default time period for script execution is 10 minutes. After this period expires, Veeam Agent stops executing the script. You can change the script timeouts in the `/etc/veeam/veeam.ini` configuration file with the `timeoutFreezeThaw` parameter where the time period is set in seconds.

### NOTE

You can specify pre-freeze and post-thaw scripts only if Veeam Agent for Linux operates in the Server edition. If these scripts were enabled for the job while Veeam Agent operated in the Server edition, and then Veeam Agent has switched to another edition (for example, to the Free edition after the license has expired), the backup job will fail. You will need to delete the existing job and create a new backup job without pre-freeze and post-thaw scripts enabled.

# File System Indexing

You can instruct Veeam Agent for Linux to create an index of files and directories located on the Veeam Agent computer during backup. File indexing allows you to search for specific files inside Veeam Agent backups and perform 1-click restore in Veeam Backup Enterprise Manager.

File indexing is enabled at the job level. You can specify granular indexing settings for each job.

## IMPORTANT

Indexing mechanism does not recognize file exclusion masks. If you specify masks to exclude certain files in a file-level backup job, Veeam Agent for Linux will nevertheless index all files located in the directories that have been selected for backup.

For example, you have included the `/home` directory into the backup and specified the `*.pdf` exclusion mask. The *Index everything* option is enabled for the backup job. In this case, when you browse the resulting backup in Veeam Backup Enterprise Manager, PDF files will be displayed in the `/home` directory as if they were backed up.

## Requirements for File System Indexing

- Veeam Agent for Linux must have either Workstation or Server license installed.
- The following utilities must be installed on the computer: `gzip` and `tar` (standard utilities for majority of Linux distributions). These utilities are provided along with the product in the product installation media.

## NOTE

Consider the following:

- File system indexing is optional. If you do not enable this option in the backup job settings, you will still be able to perform 1-click restore from the backup created with such backup job. For more information, see the [Preparing for File Browsing and Restore](#) section in the Veeam Backup Enterprise Manager User Guide.
- If SELinux is enabled in the Linux OS, file system indexing may fail.

## Automatic Job Retries

Veeam Agent supports automatic job retries if a scheduled backup job fails for any reason – for example, if the backup repository is not available or connection to it is interrupted during the backup job execution.

If the backup job fails, Veeam Agent will automatically create a new session for this backup job with the *Pending* status. By default, Veeam Agent for Linux retries a failed job 3 times with an interval of 10 minutes.

Veeam Agent automatically restarts the backup job under the following conditions:

- If the backup job was launched automatically according to a schedule and failed for any reason. Veeam Agent will not perform a backup job retry if the backup job ended with the *Success* or *Warning* status.
- [For object storage targets] If during the job session, a backup health check detects corrupted data in the backup that resides in an object storage repository. By default, Veeam Agent will launch a job retry. Veeam Agent will retry the backup job up to 3 times if the backup job was run on a schedule. Veeam Agent will retry the backup job only once if the backup job was launched manually. For more information, see [Health Check for Object Storage](#).

# Backup Repository

A backup job configured in Veeam Agent for Linux creates backup files in a backup repository. A backup repository is a directory on the storage where you want to keep backup files. You can use the following types of disk-based storage to create a backup repository:

- Local (internal) storage of the protected machine (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share.
- Veeam Backup & Replication 12.1 or later backup repository (including deduplication appliances).
- Veeam Cloud Connect 12.0 or later backup repository.
- Object storage repository, such as S3 Compatible storage, Amazon S3, Google Cloud or Microsoft Azure Blob.

## IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

Veeam Agent for Linux works with backup storage differently depending on the way you configure and start backup jobs – with the Veeam Agent control panel or command line interface.

## Backup Location and Control Panel

If you use the Veeam Agent control panel to perform backup tasks, you do not have to deal with backup repositories. When you specify a target location for backup in the Backup Job wizard, Veeam Agent configures the backup repository automatically. Veeam Agent saves path to the specified backup location, assigns to this location a unique name and ID and saves this information in the database. The information is used by Veeam Agent and is not displayed in the control panel.

If you target a backup job at the network shared folder, every time the backup job starts, Veeam Agent will automatically mount the shared folder to the `/tmp/veeam` directory in the computer file system and create a backup file in this directory. After the backup job completes, Veeam Agent will automatically unmount the network shared folder.

You can target several backup jobs to individual backup locations or use the same target location for several backup jobs. This may be useful if you want to back up different types of data to separate locations or to keep all backed-up data at one place.

## Backup Repository and Command Line Interface

If you work with Veeam Agent for Linux using the command line interface, you must deal with backup repositories depending on the target location selected for the backup job.

If you target a backup job at a local directory or network shared folder, you must create a repository before you configure a backup job:

- In case of a local directory, you specify a name for the repository and a local directory in which Veeam Agent will create backup files. To learn more, see [Creating Repository in Local Directory](#).

- In case of a network shared folder, you specify a name for the repository, a path to the network shared folder in which Veeam Agent will create backup files, a type of the network shared folder and additional mounting options.

Every time the backup job starts, Veeam Agent will automatically mount the shared folder to the `/tmp/veeam` directory in the computer file system and create a backup file in this directory. After the backup job completes, Veeam Agent will automatically unmount the network shared folder. To learn more, see [Creating Repository in NFS Share](#) and [Creating Repository in SMB Share](#).

If the directory to which the shared folder should be mounted resides on the backed-up volume, the backup job may fail.

- In case of object storage, you specify a name for the storage provider, a name for the repository and settings to access the storage account and bucket/container. To learn more, see [Creating Repository in Object Storage](#).


If you target a backup job at a Veeam backup repository or cloud repository, you do not need to create repositories. Before configuring the backup job, you must connect to the Veeam backup server or Veeam Cloud Connect service provider. To learn more, see [Connecting to Veeam Backup Server](#) and [Connecting to Service Provider](#).

You can configure several backup repositories and target different backup jobs at these repositories. This may be useful if you want to back up different types of data to separate locations or to keep several copies of your backed-up data.

# Backup Chain

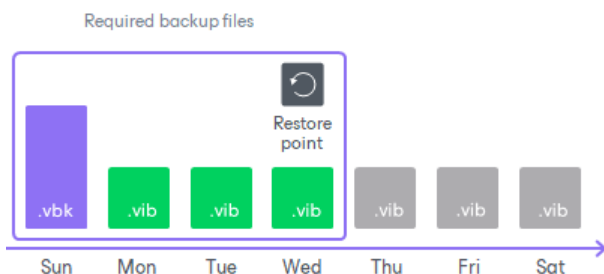
Every backup job session produces a new backup file in the target location. Backup files make up a backup chain. The backup chain can contain files of two types: full backups and incremental backups.

- During the first backup job session, Veeam Agent performs full backup. It copies all data that you have chosen to back up (entire volumes and folders) and stores the resulting full backup file (VBK) in the target location. The full backup takes significant time to complete and produces a large backup file: you have to copy the whole amount of data.
- During subsequent backup job sessions, Veeam Agent performs incremental backups. It copies only new or changed data relatively to the last backup job session and saves this data as an incremental backup file (VIB) in the target location. Incremental backups typically take less time than full backup: you have to copy only changes, not the whole amount of data.



After several backup cycles, you have a chain of backup files in the target location: the first full backup file and subsequent incremental backup files. Every backup file contains a restore point for backed-up data. A restore point is a "snapshot" of your data at a specific point in time. You can use restore points to roll back your data to the necessary state.

To recover data to a specific restore point, you need a chain of backup files: a full backup file plus a set of incremental backup files following this full backup file. If some file from the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, we recommend that you do not delete separate backup files manually. To learn more, see [Deleting Backups](#).



## Types of Backup Files

Veeam Agent produces backup files of the following types:

- VBK — full backup file.
- VIB — incremental backup file.
- VBM — backup metadata file. The backup metadata file is updated with every backup job session. It contains information about the computer on which the backup was created, every restore point in the backup chain, how restore points are linked to each other and so on. The backup metadata file is required for performing file-level and volume-level restore operations.

# Short-Term Retention Policy


Restore points in the backup chain are not kept forever. They are removed according to the retention policy. The retention policy helps maintain the life cycle of restore points and make sure that backup files do not consume the whole disk space.

Veeam Agent for Linux retains the number of latest restore points defined by the user. During every backup job session, Veeam Agent for Linux checks if there is any obsolete restore point in the backup chain. If some restore point is obsolete, it is removed from the chain.


## Removing Backups by Retention

When the obsolete restore points are removed by retention, Veeam Agent transforms the backup chain so it always contains a full backup file on which subsequent incremental backup files are dependent. To do so, Veeam Agent uses the following rotation scheme:


1. During every backup job session Veeam Agent adds a backup file to the backup chain and checks if there is an obsolete restore point.




2. If an obsolete restore point exists, Veeam Agent transforms the backup chain. As part of this process, it performs the following operations:
  - a. Veeam Agent rebuilds the full backup file to include in it data of the incremental backup file that follows the full backup file. To do this, Veeam Agent injects into the full backup file data blocks from the earliest incremental backup file in the chain. This way, a full backup 'moves' forward in the backup chain.



- b. Veeam Agent removes the earliest incremental backup file from the chain as redundant: its data has already been injected into the full backup file, and the full backup file includes data of this incremental backup file.



If the backup chain contains several obsolete restore points, the rebuild procedure is similar. Data from several restore points is injected to the rebuilt full backup file. This way, Veeam Agent makes sure that the backup chain is not broken, and you will be able to recover your data to any restore point.



## Long-Term Retention Policy

The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store backup files for long periods of time – for weeks, months and even years. For this purpose, Veeam Agent does not create any special new backup files – it uses backup files created while backup job runs and marks these backups with specific GFS flags.

To mark a backup file for long-term retention, Veeam Agent can assign to the file the following types of GFS flags: weekly (W), monthly (M) and yearly (Y). The types of GFS flags that Veeam Agent assigns depend on the configured [GFS retention policy settings](#).

### NOTE

Consider the following:

- GFS flags can be assigned only to full backup files created during the time period specified in GFS policy settings.
- If you store your backups in an object storage repository managed by Veeam Backup & Replication and connection to this repository is set up through a gateway server, configuring active full backups is not required, Veeam Agent will create a full backup based on the last incremental backup and will assign a GFS flag to this full backup. If some data blocks required to create the full backup already reside in the object storage repository, the full backup will contain links to such data blocks. To avoid extra costs, Veeam Agent does not retrieve actual data blocks from the object storage repository.

If Veeam Agent assigns a GFS flag to a full backup file, this backup file can no longer be deleted or modified. Veeam Agent does not apply short-term retention policy settings to the full backup file. For example, Veeam Agent ignores the backup file when determining whether the number of allowed backup files is exceeded.

When the specified retention period ends, Veeam Agent unassigns the GFS flag from the full backup file. If the backup file does not have any other GFS flags assigned, it can be modified and deleted according to the short-term retention policy.

Veeam Agent assigns GFS flags in the similar way as Veeam Backup & Replication does for VM backup files. To learn about logic behind GFS flags, see the [Assignment of GFS Flags](#) and [Removal of GFS Flags](#) sections in the Veeam Backup & Replication User Guide.

# Limitations


When planning to use GFS retention policy, consider the following limitations:

- [Applicable to all backup targets except object storage] While applying the GFS retention policy, Veeam Agent does not create new full backup files. You must configure your backup jobs in a way you do not lose any essential data due to an insufficient number of full backup files. For example, if you configure monthly GFS retention, you need at least one full backup file per month.
- If a GFS flag is assigned to a full backup file in an active backup chain, the following applies:
  - Veeam Agent cannot transform the backup chain according to the short-term retention policy.
  - Veeam Agent is not able to merge data from incremental backup files into the full backup file.
- Veeam Agent assigns GFS flags only after you save GFS retention policy settings. This means that GFS flags are assigned only to those backup files created after the configuration, while backup files created earlier are not affected and previously assigned flags are not modified.
- You cannot store full backups to which GFS flags are assigned in backup repositories with rotated drives.
- Retention policy for deleted items does not apply to full backup files to which GFS flags are assigned.

## Active Full Backup

When Veeam Agent performs active full backup, it produces a full backup file and adds this file to the backup chain.

The active full backup resets the backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file and its subsequent incremental backup files remain on the disk. After the last incremental backup file created prior to the active full backup becomes outdated, Veeam Agent automatically deletes the previous backup chain. To learn more, see [Retention Job for Active Full Backups](#).



You can create active full backups manually or schedule a backup job to create active full backups periodically. To do this, you can use the Veeam Agent for Linux control panel or command line interface.

- To learn how to configure active full backup schedule and create active full backups with the Veeam Agent for Linux control panel, see [Active Full Backup Settings](#) and [Starting Backup Job from Control Panel](#).
- To learn how to configure active full backup schedule and create active full backups with the Veeam Agent for Linux command line interface, see [Configuring Active Full Backup Schedule](#) and [Creating Active Full Backups](#).



# Active Full Backup Schedule

You can schedule a backup job to create active full backups periodically. Active full backup schedule depends on the regular backup schedule.

- In case active full backup is scheduled on a week day, Veeam Agent modifies the regular schedule of the backup job.

For example, the regular backup schedule is set to Monday and Tuesday at 15:00. Active full backup schedule is set to Friday. In this case, the backup job schedule will contain information that the job must start on Monday, Tuesday and Friday at 15:00.


- In case active full backup is scheduled on a day of the month, Veeam Agent runs the backup job on this day at the same time as it must run upon the regular schedule.

Keep in mind that if the job is not scheduled to run automatically, Veeam Agent will not run active full backup. For more information on how to configure backup job schedule, see [Configuring Backup Schedule](#) and [Configuring Active Full Backup Schedule](#).


To be able to restore data from a Veeam Agent backup, you need to have a full backup file and a chain of subsequent incremental backup files on the disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you create an active full backup, in some days there will be more restore points on the disk than specified by retention job settings. Veeam Agent will remove the full backup chain only after the last incremental backup file in the chain becomes outdated.

For example, the retention job is set to 3 restore points. A full backup file is created on Sunday, incremental backup files are created on Monday and Tuesday, and an active full backup is created on Wednesday. Although the backup chain now contains 4 restore points, Veeam Agent will not delete the previous backup chain. Veeam Agent will wait for the next 2 incremental backup files to be created, and only then will delete the whole previous chain, which will happen on Friday. As a result, although the retention job is set to 3 restore points, the actual number of backup files on the disk will be greater for some time.



Veeam Agent treats the active full backup in the same way as a regular full backup. If some restore point becomes obsolete, Veeam Agent will re-build the full backup file to include in it data of the incremental backup file that follows the full backup file. After that, Veeam Agent will remove the earliest incremental backup file from the chain as redundant.



# Data Compression

Veeam Agent provides mechanisms of data compression. Data compression lets you decrease traffic going over the network and disk space required for storing backup files.

## Data Compression

Data compression decreases the size of created backups but affects duration of the backup procedure. When you create a backup job in command line interface, Veeam Agent allows you to specify one of the following compression levels:

Compression Level	CLI Option	Compression Algorithm	Description
None	0	No compression	This compression level is recommended if you plan to store backup files on storage devices that support hardware compression and deduplication.
Dedupe-friendly	1	Rle	Optimized compression level for very low CPU usage. You can select this compression level if you want to decrease the load on the CPU of the Veeam Agent computer.
Optimal	2	Lz4	The default recommended compression level. It provides the best ratio between size of the backup file and time of the backup procedure.
High	3	Zstd 3	Provides up to 60% additional compression ratio over the Optimal level at the cost of 2x higher CPU usage and 2x slower restore.
Extreme	4	Zstd 9	Provides the smallest size of the backup file but reduces the backup performance. We recommend that you use the extreme compression level only on Veeam Agent computers with modern multi-core CPUs (6 cores recommended).

You can change data compression settings for existing backup jobs. New settings will not have any effect on previously created backup files in the backup chain. They will be applied to new backup files created after the settings were changed.

Compression settings are changed on the fly. You do not need to create a new full backup to use new settings — Veeam Agent will automatically apply the new compression level to newly created backup files.

## Storage Optimization

Depending on the type of storage you select as a backup target, Veeam Agent uses data blocks of different size, which optimizes the size of a backup file and job performance. You can choose one of the following storage optimization options:

- **4 MB** – select this option for backup jobs that can produce very large full backup files (larger than 16 TB). With this option selected, Veeam Agent will use data block size of 4096 KB.
- **1 MB (default)** – select this option for backup to SAN, DAS or local storage. With this option selected, Veeam Agent will use data block size of 1024 KB.

The SAN identifies larger blocks of data and therefore can process large amounts of data at a time. This option provides the fastest backup job performance.

- **512 KB** – select this option for backup to NAS and onsite backup. With this option selected, Veeam Agent will use data block size of 512 KB. This option reduces the size of an incremental backup file because of reduced data block sizes.
- **256 KB** – select this option if you plan to use WAN for offsite backup. With this option selected, Veeam Agent will use data block size of 256 KB. This results in the smallest size of backup files, allowing you to reduce the amount of traffic over WAN.

#### NOTE

If you change storage optimization settings, the new settings will be applied only after an active full backup is created. Veeam Agent will use the new block size for the active full backup and subsequent backup files in the backup chain. For more information on scheduling active full backups, see [Active Full Backup Settings](#).

# Data Encryption

Data security is an important part of the backup strategy. You must protect your information from unauthorized access, especially if you back up sensitive data to remote locations. To keep your data safe, you can use data encryption.

Data encryption transforms data to an unreadable, scrambled format with the help of a cryptographic algorithm and a secret key. If encrypted data is intercepted, it cannot be unlocked and read by the eavesdropper. Only intended recipients who know the secret key can reverse encrypted information back to a readable format.

In Veeam Agent, encryption works at the backup job level. Veeam Agent uses the block cipher encryption algorithm and stores data in the encrypted format to a backup file.

Encryption is performed on the trusted side depending on the backup target:

- Encryption is performed on the source side for all backup targets except the Veeam backup repository.
- Encryption is performed on the target side if you store backups in the Veeam backup repository.

Decryption is performed on the same side as encryption.

To create encrypted backups, you must enable the encryption option and specify a password that will be used for data encryption. To learn more, see [Data Encryption Settings](#).

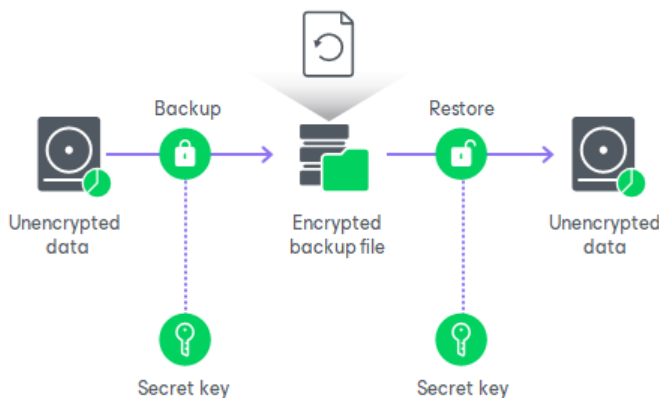
## NOTE

You cannot enable encryption options in the properties of the Veeam Agent backup job if you have chosen to create Veeam Agent backups in a Veeam backup repository. For such jobs, encryption options are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the [Data Encryption](#) section in the Veeam Backup & Replication User Guide.

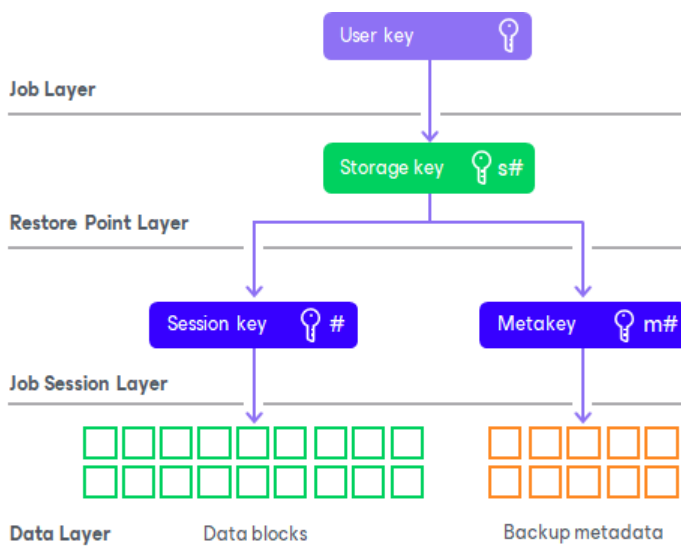
## Encryption Algorithms

To encrypt data in backups and files, Veeam Agent employs a symmetric key encryption algorithm.

The symmetric, or single-key encryption algorithm, uses a single, common secret key to encrypt and decrypt data on the trusted side. To restore encrypted data, you must have the same secret key. Users who do not have the secret key cannot decrypt data and get access to it.



Veeam Agent relies on a hierarchical encryption scheme. Each layer in the hierarchy encrypts the layer below with a key of specific type.

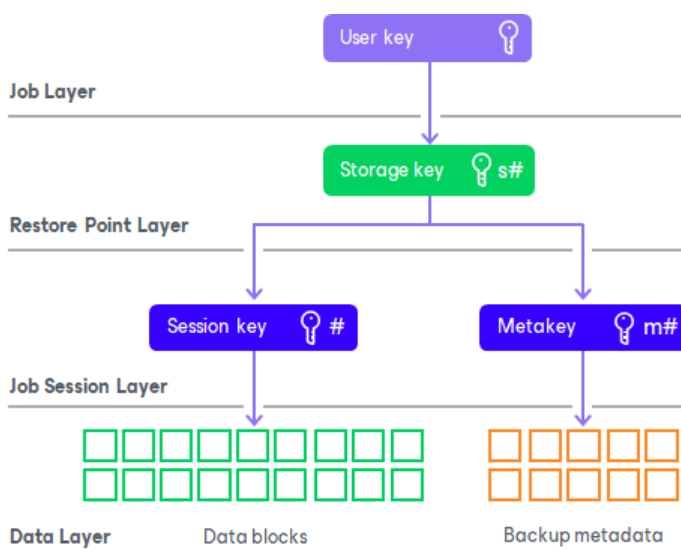


## Encryption Keys

An encryption key is a string of random characters that is used to bring data to a scrambled format and back to unscrambled. Encryption keys encode and decode initial data blocks or underlying keys in the key hierarchy.

Veeam Agent uses 4 types of keys:

- 3 service keys generated by Veeam Agent:
  - [Session Key](#)
  - [Metakey](#)
  - [Storage key](#)
- 1 key generated based on a user password: a [user key](#).




## Session Keys and Metakeys

The session key is the lowest layer in the encryption key hierarchy. When Veeam Agent encrypts data, it first encodes every data block in a file with a session key. For session keys, Veeam Agent uses the AES algorithm with a 256-bit key length in the CBC-mode.


Veeam Agent generates a new session key for every backup job session. For example, if you have created an encrypted backup job and run 3 job sessions, Veeam Agent will produce 3 backup files that will be encrypted with 3 different session keys:

- Full backup file encrypted with session key 1
- Incremental backup file encrypted with session key 2
- Incremental backup file encrypted with session key 3



The session key is used to encrypt only data blocks in backup files. To encrypt backup metadata, Veeam Agent applies a separate key – metakey. Use of a metakey for metadata raises the security level of encrypted backups.

For every job session, Veeam Agent generates a new metakey. For example, if you have run 3 job sessions, Veeam Agent will encrypt metadata with 3 metakeys.




In the encryption process, session keys and metakeys are encrypted with keys of a higher layer – storage keys. Cryptograms of session keys and metakeys are stored in the resulting file next to encrypted data blocks. Metakeys are additionally kept in the Veeam Agent database.

## Storage Keys

Backup files in the backup chain often need to be transformed, for example, when the earliest incremental backup file in the chain becomes obsolete and its data should be included into the full backup file. When Veeam Agent transforms a full backup file, it writes data blocks from several restore points to the full backup file. As a result, the full backup file contains data blocks that are encrypted in different job sessions with different session keys.

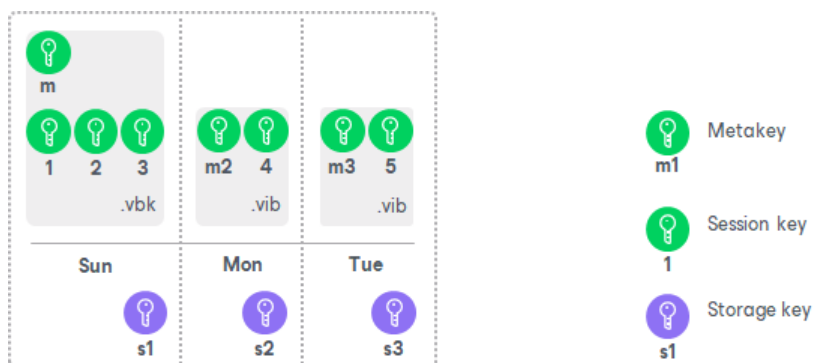
To restore data from such “composed” backup file, Veeam Agent would require a bunch of session keys. For example, if the backup chain contains restore points for 2 months, Veeam Agent would have to keep session keys for a 2-month period.



In such situation, storing and handling session keys would be resource consuming and complicated. To facilitate the encryption process, Veeam Agent uses another type of service key – a storage key.

For storage keys, Veeam Agent uses the AES algorithm with a 256-bit key length in the CBC-mode. A storage key is directly associated with one restore point in the backup chain. The storage key is used to encrypt the following keys in the encryption hierarchy:

- All session keys for all data blocks in one restore point
- Metafile encrypting backup metadata



During the restore process, Veeam Agent uses one storage key to decrypt all session keys for one restore point, no matter how many session keys were used to encrypt data blocks in this restore point. As a result, Veeam Agent does not need to keep the session keys history in the Veeam Agent database. Instead, it requires only one storage key to restore data from one file.

In the encryption process, storage keys are encrypted with a key of a higher layer – a user key. Cryptograms of storage keys are stored in the resulting file next to encrypted data blocks, and cryptograms of session keys and metafiles.

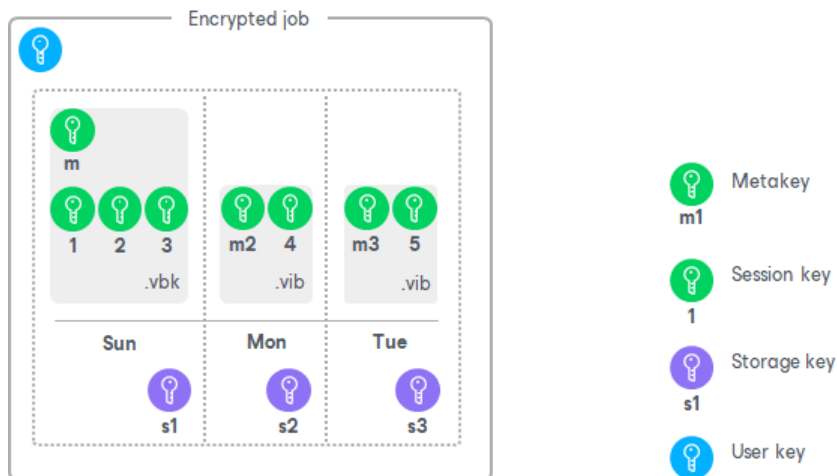
Storage keys are also kept in the Veeam Agent database. To maintain a set of valid storage keys in the database, Veeam Agent uses retention policy settings specified for the job. When some restore point is removed from the backup chain by retention, the storage key corresponding to this restore point is also removed from the Veeam Agent database.

## User Keys

When you enable encryption for a job, you must define a password to protect data processed by this job, and define a hint for the password. The password and the hint are saved in the job settings. Based on this password, Veeam Agent generates a user key.



The user key protects data at the job level. In the encryption hierarchy, the user key encrypts storage keys for all restore points in the backup chain.



Veeam Agent saves a hint for the password to its database and to the backup metadata file (VBM). When you decrypt a file, Veeam Agent displays a hint for the password that you must provide. After you enter a password, Veeam Agent derives a user key from the password and uses it to unlock the storage key for the encrypted file.

According to the security best practices, you should change passwords for encrypted jobs regularly. When you change a password for the job, Veeam Agent creates a new user key and uses it to encrypt new restore points in the backup chain. If you lose a password that was specified for encryption, you can change the password in the encryption settings. You can use the new password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

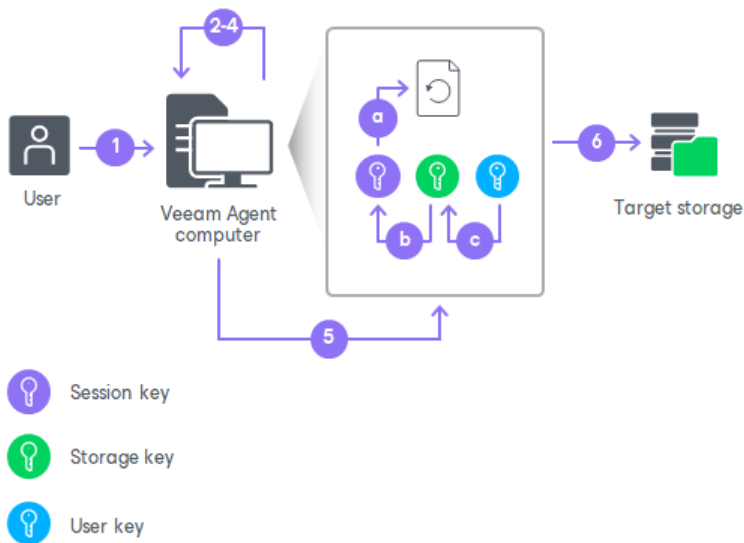
## How Data Encryption Works

Data encryption is performed as part of the backup process. Encryption works at the source or target side depending on the backup target. As a result, encryption keys are not passed to the untrusted side, which helps to avoid data interception.

In Veeam Agent, the encryption process includes the following steps:

1. When you create a backup job, you enable the encryption option for the job and enter a password to protect data at the job level.
2. Veeam Agent generates a user key based on the entered password.
3. When you start an encrypted job, Veeam Agent creates a storage key and stores this key in its database.
4. Veeam Agent creates a session key and a metakey. The metakey is stored in the Veeam Agent database.
5. Veeam Agent processes job data in the following way:
  - a. The session key encrypts data blocks in the backup file. The metakey encrypts backup metadata.
  - b. The storage key encrypts the session key and the metakey.
  - c. The user key encrypts the storage key.

6. Encrypted data blocks are stored to the target location. The cryptograms of the user key, storage key, session key and metakey are stored in the resulting file next to encrypted data blocks.



## How Data Decryption Works

When you restore data from an encrypted backup file, Veeam Agent performs data decryption automatically in the background or requires you to provide a password.

- If encryption keys required to unlock the backup file are available in the Veeam Agent database, you do not need to enter the password. Veeam Agent uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore does not differ from that from an unencrypted one.

Automatic data decryption can be performed when you encrypt and decrypt the backup file on the same Veeam Agent computer using the same Veeam Agent database.

- If encryption keys are not available in the Veeam Agent database, you need to provide a password to unlock the encrypted file.

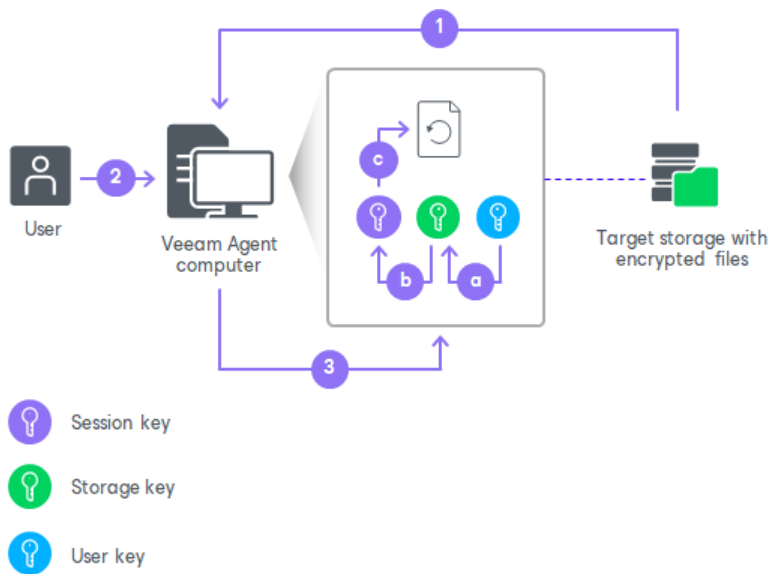
Data decryption is performed on the source or target side depending on the backup target. As a result, encryption keys are not passed to the untrusted side, which helps avoid data interception.

In Veeam Agent, the decryption process includes the following steps. Keep in mind that steps 1 and 2 are required only if you decrypt the file on the Veeam Agent computer other than the computer where the file was encrypted.

1. You select the backup from which you want to restore data. Veeam Agent notifies you that one or more files in the backup chain are encrypted and requires a password.
2. You specify a password for the imported file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.
3. Veeam Agent reads the entered password and generates the user key based on this password. With the user key available, Veeam Agent performs decryption in the following way:
  - a. Veeam Agent applies the user key to decrypt the storage key.
  - b. The storage key, in its turn, unlocks underlying session keys and a metakey.

c. Session keys decrypt data blocks in the encrypted file.

After the encrypted file is unlocked, you can work with it as usual.

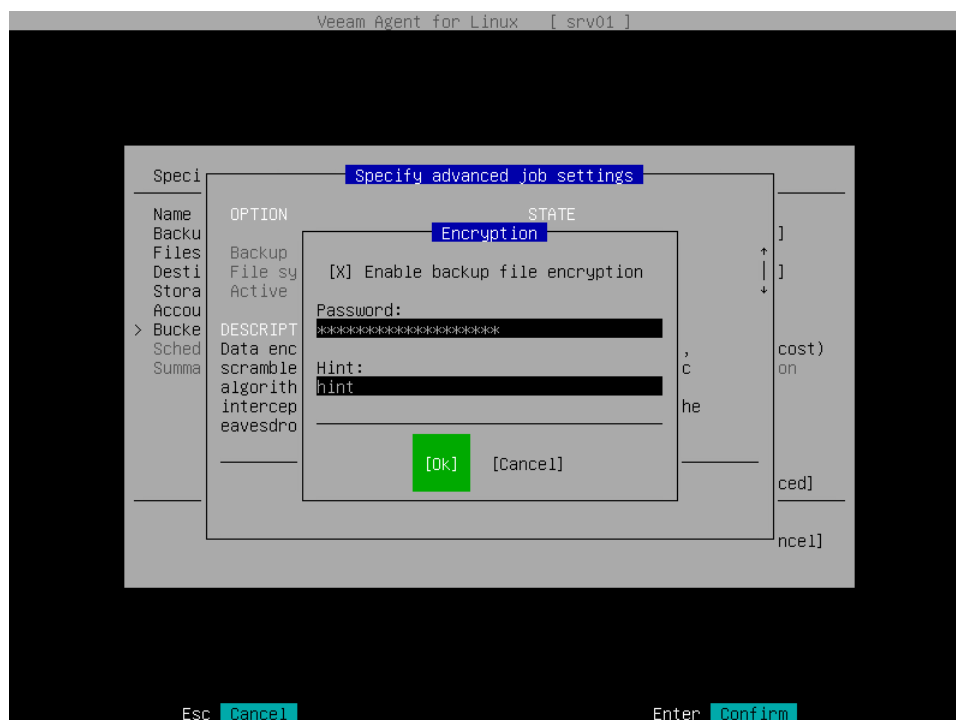


## Backup Job Encryption

Encryption for the backup job is configured in the advanced job settings. You should enable the encryption option and specify a password to protect data in backup files produced by the backup job.

### NOTE

You cannot specify encryption options for the backup job if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more, see the [Data Encryption](#) section of the Veeam Backup & Replication User Guide.



The backup job processing with encryption enabled includes the following steps:

1. You enable encryption for a backup job and specify a password.
2. Veeam Agent generates the necessary keys to protect backup data.
3. Veeam Agent encrypts data blocks and transfers them to the target location already encrypted.
4. On the target storage, encrypted data blocks are stored in a resulting backup file.



Restore of an encrypted backup file includes the following steps:

1. You select an encrypted backup file and define a password to decrypt the backup file. If the password has changed once or several times, you need to specify the latest password that was used to encrypt files in the backup chain.
2. Veeam Agent uses the provided password to generate user key and unlock the subsequent keys for backup file decryption.
3. Veeam Agent retrieves data blocks from the backup file, sends them to the target volume and decrypts them on the target volume.



## Encryption Best Practices

To guarantee the flawless process of data encryption and decryption, consider the following advice.

### Password

1. Use strong passwords that are hard to crack or guess. Consider the following recommendations:
  - a. The password must be at least 8 characters long.
  - b. The password must contain uppercase and lowercase characters.
  - c. The password must be a mixture of alphabetic, numeric and punctuation characters.
  - d. The password must significantly differ from the password you used previously.
  - e. The password must not contain any real information related to you, for example, date of birth, your pet's name, your logon name and so on.
2. Provide a meaningful hint for the password that will help you recall the password. The hint for the password must significantly differ from the password itself. The hint for the password is displayed when you select an encrypted backup server and attempt to unlock it.

3. Change passwords for encrypted jobs regularly. Use of different passwords helps increase the encryption security level.

## Encryption for Existing Job

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

Encryption is not retroactive. If you enable encryption for an existing backup job, Veeam Agent does not encrypt the previous backup chain created with this job. However, Veeam Agent encrypts backup metadata. As a result, you need to enter the password to restore data from unencrypted backup files in the backup chain as well as from encrypted backup files in this chain.

# Backup of Database Systems

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run one of the following database systems:

- [Oracle database system](#)
- [MySQL database system](#)
- [PostgreSQL database system](#)

To process database systems with Veeam Agent for Linux, enable application-aware processing for the backup job:

- If you work with Veeam Agent using the Veeam Agent control panel, configure application-specific settings at the **Advanced** step of the Backup Job wizard. To learn more, see [Specify Advanced Backup Settings](#).
- If you work with Veeam Agent using the command line interface, create the backup job, then specify application-specific settings for this job. To learn more, see [Creating Volume-Level Backup Job](#) and [Configuring Database Processing Settings](#).

## Considerations and Limitations

When you back up database systems, consider the following:

- Nosnap Veeam Agent for Linux and nosnap Veeam Agent for Linux on Power do not support application-aware processing and cannot be used to back up database systems.
- You can specify settings for database system processing only if Veeam Agent for Linux operates in the Server edition.
- Veeam Agent supports processing of database systems for the volume-level backup only.
- If there are multiple database systems on the Veeam Agent computer, consider the following:
  - Veeam Agent supports processing of multiple PostgreSQL database systems on one Veeam Agent computer.
  - Veeam Agent supports processing of multiple Oracle Database systems on one Veeam Agent computer only if such systems are of the same major version.
  - Veeam Agent does not support processing of multiple MySQL database systems on one Veeam Agent computer.
  - Veeam Agent does not support processing of multiple database systems of different types on one Veeam Agent computer.
- Veeam Agent does not support 32-bit database systems installed on a 64-bit Linux OS.

## Oracle Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the Oracle database system.

## NOTE

You can use Veeam Explorer for Oracle to restore Oracle databases from a Veeam Agent for Linux backup. For information about item-level recovery of Oracle systems, see the [Restoring Oracle Items](#) section of the Veeam Backup & Replication User Guide.

## Requirements and Limitations for Oracle Processing

- Oracle Database versions 11g – 21c are supported for all operating systems supported by Veeam Agent for Linux. To learn more, see [System Requirements](#).
- Automatic Storage Management (ASM) is not supported.
- Oracle Real Application Clusters (RAC) are not supported.
- Oracle Grid Infrastructure is not supported.
- Oracle Database Express Edition (XE) is not supported.
- SAP on Oracle is not supported.
- Oracle Database architectures with Data Guard are not supported.

## Authentication Methods

Veeam Agent for Linux can connect to the Oracle database system and perform Oracle archived logs backup and delete operations using one of the following account types:

- *System account* – Veeam Agent uses the account of the machine OS. To connect to the Oracle database system, the account must be a member of the group that owns configuration files for the Oracle database (for example, the oinstall group).
- *Oracle account* – Veeam Agent uses the Oracle account. To connect to the Oracle database system, the account must have SYSDBA rights.

## How Oracle Processing Works

To ensure that the backed-up data is in the consistent state, Veeam Agent for Linux performs the Oracle database system processing using an internal component: *oralib*. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent obtains information about Oracle databases that run on the Veeam Agent machine.
2. Veeam Agent connects to the Oracle database and operates depending on the database state and mode:
  - [Shutdown state](#)
  - [Backup state](#)
  - [Running database in ARCHIVELOG mode](#)
  - [Running database in NOARCHIVELOG mode](#)

After Veeam Agent for Linux finishes database system processing, Veeam Agent proceeds to the next step of the backup process. To learn more, see [How Backup Works](#).

## Processing of Database in Shutdown State

If the Oracle database is shut down, Veeam Agent skips it and tries to connect to the next Oracle database if there are multiple database instances on the machine. The skipped Oracle database will be included in the backup. You cannot restore such Oracle database as an independent item using Veeam Explorer for Oracle. To restore such database, you must restore the entire volume that contains the database. To learn more about restoring volumes, see [Volume-Level Restore](#).

Veeam Agent displays a warning message about the database that is shut down in the job session logs. The backup job does not fail.

## Processing of Database in Backup State

If the database is in the backup state, depending on the selected Oracle Processing option, Veeam Agent performs application-aware processing differently:

- If the Oracle processing is set to **Require successful processing**, the backup job will fail.
- If the Oracle processing is set to **Try application processing, ignore failures**, Veeam Agent will skip the database that is in the backup state and if there are multiple databases in the system, will try to connect to the next database. The skipped database will not be included in the backup.

## Processing of Running Database in ARCHIVELOG Mode

If the Oracle database is running in the ARCHIVELOG mode, the Oracle database system keeps archived logs that allow to recover all committed transactions of the database. To learn more, see [Oracle documentation](#).

If the database operates in the ARCHIVELOG mode, Veeam Agent performs the following operations:

1. Veeam Agent switches the database to the backup mode. Veeam Agent changes the database state using the Oracle functionality.
2. Veeam Agent creates a snapshot of the volume.
3. Veeam Agent returns the database to the initial state.

## Processing of Running Database in NOARCHIVELOG Mode

If the Oracle database is running in the NOARCHIVELOG mode, the Oracle database does not create archived logs. Logs that are created before the database is switched to NOARCHIVELOG remain untouched. In this mode, you can restore the database only to the state in which the database is contained in the restore point. You cannot recover transactions subsequent to that full database backup.

If the database operates in the NOARCHIVELOG mode, Veeam Agent performs the following operations:

1. Shuts down the database using the Oracle functionality.
2. Creates a snapshot of the volume.
3. Returns the database to the initial state.



# Archived Log Processing

In the ARCHIVELOG mode, the Oracle database system stores database archived logs to a certain location on the machine that runs the database system, as specified by the database administrator. Veeam Agent allows you to set up the following ways of archived logs processing:

- *Delete logs older than the specified time (in hours).* After the backup job completes, Veeam Agent deletes archived logs that are older than the specified time from the Veeam Agent machine. This helps make sure that logs do not overflow the storage space on the processed machine.
- *Delete oldest logs larger than the specified size (in GB).* After the backup job completes, Veeam Agent checks whether the total size of archived logs exceeds the specified size. After that, Veeam Agent deletes oldest archived logs that exceed the specified size from the processed machine. This helps make sure that logs do not overflow the storage space on the Veeam Agent machine.
- *Do not delete archived logs.* Log files remain untouched on the Veeam Agent machine.

Veeam Agent processes archive logs via Oracle Call Interface (OCI).

## MySQL Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the MySQL database system.

## Requirements and Limitations of MySQL Processing

- Veeam Agent for Linux supports processing of MySQL database systems version 5.7 – 9.0.
- Configurations with multiple MySQL installations or instances on the same machine are not supported.
- MySQL Cluster versions are not supported.
- MySQL tables that use the MyISAM storage engine must be locked to keep them in consistent state while Veeam Agent is creating the system snapshot. To correctly process such tables, MySQL account must have the following instance-wide privileges:
  - `SELECT`. This privilege enables Veeam Agent to access tables' metadata and select for a lock the tables that use the MyISAM storage engine. Without this privilege, the processing of the MySQL database system will run successfully but MyISAM tables will not be locked, which may result in an inconsistent state of the backed up data.
  - `LOCK TABLES`. This privilege is required for locking the selected MyISAM tables. If some MyISAM tables are selected but the MySQL account does not have the `LOCK TABLES` privilege, the processing of the MySQL database system will fail.
  - `RELOAD` or `FLUSH TABLES`. If some MyISAM tables are selected but the MySQL account does not have either `RELOAD` or `FLUSH TABLES` privilege, the processing of the MySQL database system will fail.

To obtain information about the privileges that are assigned to an account, use MySQL functionality, for example, the `SHOW GRANTS` statement. To learn more, see [MySQL documentation](#).

## Authentication Methods

Veeam Agent for Linux can connect to the MySQL database system using one of the following methods:

- *Password* – Veeam Agent uses the MySQL account credentials that you specify in the backup job settings.

- *Password file* – Veeam Agent uses the MySQL account credentials that are stored in the `.my.cnf` password file. To learn more about password file configuration, see [Preparing Password File for MySQL Processing](#).

## How MySQL Processing Works

To ensure that the backed-up data is in the consistent state, Veeam Agent for Linux performs the MySQL database system processing. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent connects to the MySQL database system and obtains the list of tables.
2. Veeam Agent locks the base tables that use the MyISAM storage engine. Veeam Agent changes the table state using the MySQL functionality. Tables that use the InnoDB storage engine do not require locking.  
  
Keep in mind that Veeam Agent supports processing of tables based on the MyISAM and InnoDB storage engines only. Veeam Agent does not support tables that use other storage engines.
3. Veeam Agent creates a snapshot of the volume.
4. Veeam Agent unlocks tables locked at Step 2.

After Veeam Agent unlocks tables, Veeam Agent proceeds to the next step of the backup process. To learn more, see [How Backup Works](#).

## PostgreSQL Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the PostgreSQL database system.

### NOTE

You can use Veeam Explorer for PostgreSQL to restore PostgreSQL instances from a Veeam Agent for Linux backup. For information about item-level recovery of PostgreSQL systems, see the [Restoring PostgreSQL Items](#) section of the Veeam Backup & Replication User Guide.

## Requirements and Limitations for PostgreSQL Processing

- Veeam Agent supports processing of PostgreSQL database systems version 12, 13, 14, 15 and 16.
- Veeam Agent does not support backup of PostgreSQL clusters.

## Authentication Methods

Veeam Agent for Linux can connect to the PostgreSQL database system using one of the following methods:

- *Database user with password* – Veeam Agent uses the PostgreSQL account credentials that you specify in the backup job settings.
- *Database user with password file* – Veeam Agent the PostgreSQL database system to use account credentials that are stored in the `.pgpass` password file. To learn more about password file configuration, see [Preparing Password File for PostgreSQL Processing](#).
- *System user without password* – Veeam Agent uses the peer authentication. In the peer authentication method, Veeam Agent for Linux uses the OS account as the PostgreSQL database user name.

# How PostgreSQL Processing Works

After Veeam Agent for Linux finishes database system processing, Veeam Agent proceeds to the next step of the backup process. To learn more, see [How Backup Works](#).

To ensure that the backed-up data is in the consistent state, Veeam Agent performs the PostgreSQL database processing using an internal component: *pgsqlagent*. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent scans the Veeam Agent computer for PostgreSQL instances.

By default, Veeam Agent recursively scans the `/etc/postgresql`, `/var/lib/postgresql` and `/var/lib/pgsql` directories for the configuration files of PostgreSQL instances. If your instance is stored in a custom location, you must specify its location in the PostgreSQL configuration file — `VeeamPostgreSQLAgent.xml`. You must create this file in the `/etc/veeam/` directory. To explicitly include or exclude specific directories in/from processing, you can use the following elements in the configuration file:

- `AddConfigDirs` — use this element to specify paths to directories that you want Veeam Agent to scan.
- `ExcludeConfigDirs` — use this element to specify paths to directories that you do not want Veeam Agent to scan.

## TIP

You can specify directories that you want to include and directories that you want to exclude in the same configuration file.

An example of the `VeeamPostgreSQLAgent.xml` file:

```
<config AddConfigDirs="/opt/psql/" ExcludeConfigDirs="/var/lib/postgresql/13/main45/,/var/lib/postgresql/13/maindd/" />
```

## IMPORTANT

The configuration file must be formatted as a single line XML.

2. If a PostgreSQL instance is detected, Veeam Agent collects information about its state and configuration settings. The following Veeam Agent behavior depends on the collected information:
  - [Shutdown state](#)
  - [Backup state](#)
  - [Running instance with WAL level set as minimal](#)
  - [Running instance with WAL level set as archival, replica or logical](#)

To learn more about the WAL level setting, see [PostgreSQL documentation](#).

## TIP

Veeam Agent stores all collected data about PostgreSQL instances in the `.VBM` file, which allows Veeam Agent to restore PostgreSQL instance as an application item. To learn more, see [Restoring PostgreSQL Items](#) in the Veeam Backup & Replication User Guide.

3. Veeam Agent creates a snapshot of the volume and proceeds to the next step of the backup process.

To learn more about backup process, see [How Backup Works](#).

## Processing of Instance in Shutdown State

If the database instance is shut down, Veeam Agent skips it and tries to connect to the next instance. The skipped database instance will be included in the backup. You cannot restore such PostgreSQL instance as an independent item using Veeam Explorer for PostgreSQL. You can restore such database instance only using either volume-level or file-level restore. To learn more about restoring volumes and files, see [Data Restore](#).

## Processing of Instance in Backup State

If the database instance is in the backup state, depending on the selected PostgreSQL Processing option, Veeam Agent performs the backup job differently:

- If the PostgreSQL processing is set to **Require successful processing**, the backup job will fail.
- If the PostgreSQL processing is set to **Try application processing, ignore failures**, Veeam Agent will skip the instance that is in the backup state and will try to connect to the next instance. The skipped database instance will not be included in the backup.

## Processing of Instance with WAL Level Set as Minimal

If the database is running and the WAL level is set as minimal, Veeam Agent forces a WAL checkpoint. This command fastens the database system restore. To learn more, see [PostgreSQL documentation](#).

Keep in mind that the backup of a PostgreSQL instance with the minimal WAL level does not contain logs. As a result, you can restore your instance only to image-level backup state.

## Processing of Instance with WAL Level Set as Archival, Replica or Logical

If the database instance is running and the WAL level is set as archival, replica or logical, Veeam Agent performs the following operations:

- Prepares the PostgreSQL instance and starts the on-line backup.
- Creates a snapshot of the instance.
- Stops the on-line backup.

# Backup to Object Storage

If you want to store your data in a cloud-based or on-premises storage, you can connect to the cloud storage service and create Veeam Agent backups in the object storage repositories provided by this service.

You can store Veeam Agent backups in the following types of object storage:

- Amazon S3
- Google Cloud Storage
- Microsoft Azure Blob Storage
- S3 compatible (including WasabiCloud and IBM Cloud)
- Veeam Data Cloud Vault added as a Veeam backup repository or Veeam Cloud Connect repository.

Depending on your backup infrastructure, object storage can be available in different configurations. To learn more, see the following subsections:

- [Backup destinations](#)
- [Types of Connection to Object Storage in Veeam Backup & Replication](#)
- [Considerations and Limitations](#)

## Backup Destinations

You can back up Veeam Agent computer data to an external cloud storage in the following ways:

- Directly to object storage. In this case, Veeam Agent connects to an object storage account and creates a backup repository in this storage.

Keep in mind that to connect to object storage, you need to have an account with access permissions to read and write data.

To learn more, see [Object Storage Settings](#).

- To object storage added as a Veeam backup repository. In this case, Veeam Agent connects to the Veeam backup repository and Veeam Backup & Replication connects to object storage and creates a backup repository in this storage.

To learn more, see [Veeam Backup Repository Settings](#).

- To object storage added as a Veeam Cloud Connect repository. In this case, Veeam Agent connects to the cloud backup repository and Veeam Backup & Replication connects to the object storage and creates a backup repository in this storage.

To learn more, see [Veeam Cloud Connect Repository Settings](#).

# Types of Connection to Object Storage in Veeam Backup & Replication

If you back up data to object storage added as a Veeam backup repository or Veeam Cloud Connect repository, you must configure a repository beforehand on the Veeam Backup & Replication side. Depending on the repository configuration, Veeam Backup & Replication provides one of the following connection types to the repository in the object storage:

- Connection through a gateway server. With this connection type, Veeam Agent connects to the repository using a proxy component — a gateway server that is assigned in the Veeam Backup & Replication console. The backup data is transferred from the Veeam Agent computer to the gateway server, then it is transferred from the gateway server to the repository.
- Direct connection. With this connection type, Veeam Agent connects directly to the repository. The backup data is transferred from the Veeam Agent computer to the repository without proxy components. The access to this repository is managed by Application Programming Interface (API) that is provided by the cloud service provider.

## Considerations and Limitations

Before you configure a backup job to store backups in an object storage repository, consider the following:

- Veeam Agent does not support direct backup to the Microsoft Azure Blob Storage under the general-purpose V1 storage account type.
- You can store backups only in those S3 compatible storage repositories that are accessible over the HTTPs protocol.
- [For object storage added as a Veeam backup repository or Veeam Cloud Connect repository] If you want to back up your data directly to the S3 compatible storage, you must additionally specify access permissions settings for the storage. For direct access, enable the **Agents share credentials to object storage repository** or the **Provided by IAM/STS object storage capabilities** access control option. For more information, see the [Managing Permissions for S3 Compatible Object Storage](#) section in the Veeam Backup & Replication User Guide.
- [For object storage added as a Veeam backup repository or Veeam Cloud Connect repository] Data recovery options are not available if you access the object storage repository using credentials with the read-only access permissions.
- Veeam Agent does not support backup to object storage for which lifecycle rules are enabled. Enabling lifecycle rules may result in backup and restore failures.

## Health Check for Object Storage

If you keep the backups of your Linux computer in an object storage repository, you can schedule regular health checks to validate integrity of the backups in the repository.

Consider the following about health check for object storage:

- Veeam Agent verifies metadata of the whole backup, not just the latest restore point.
- Veeam Agent does not read data from data blocks in the storage; Veeam Agent only lists data blocks to make sure all blocks in the storage are available for rebuilding every restore point in the active backup chain. This mechanism reduces the number of requests to the storage, which makes health check for object storage cost-efficient.

# Configuring Health Check Schedule

If you want to run health checks for a backup that resides in an object storage repository, you must set a schedule according to which Veeam Agent will perform health checks. You can set the schedule [in the Backup Job wizard](#) or [in command line interface](#) to run health checks weekly or monthly on specific days.

When you configure a health check schedule, consider the following:

- Health check is run automatically during incremental backup job session on the days specified in the health check schedule. If the backup job runs several times on a specified day, health check is performed only with the first run of the backup job on that day.

Health check is not performed during the first full backup or subsequent active full backup job sessions.

- If Veeam Agent does not run any backup jobs on the day specified in the health check schedule, health check will be performed during the first backup job session following that day.

For example, you may have scheduled to run a health check every last day of a month, while the backup job is scheduled to run every day and to create an active full backup on Sundays. If the last day of a month falls on a Sunday, the health check will be performed on the following Monday with the first incremental backup job session on that day.

## How Health Check Works

Veeam Agent performs a health check of a backup in the following way:

1. During the backup job session after a new incremental backup file is created, Veeam Agent starts the health check of the whole backup. Veeam Agent checks if the metadata of the backup is consistent, and no metadata is missing. Veeam Agent also checks if all data blocks for every restore point are available on the storage. Veeam Agent does not read data from data blocks.
2. If Veeam Agent does not find any corrupted data, the health check completes successfully. Otherwise, the health check completes with an error.

You can view the health check result in the session log. If during the health check, Veeam Agent finds corrupted data, it will also display information on where corrupt data has been detected – in metadata or blockstore, as well as list all restore points that share the corrupted data blocks.

Depending on the detected data inconsistency, Veeam Agent behaves in one of the following ways:

- If the health check detects corrupted metadata, Veeam Agent will mark the backup chain as corrupted in the Veeam Agent configuration database; the backup job session will fail. During the next scheduled or manual backup job session, Veeam Agent will create a full backup and will start a new backup chain. The corrupted backup chain will become orphaned and will remain in the repository – you can keep or delete it.
- If the health check detects corrupted data blocks in the latest restore point of the active backup chain, Veeam Agent launches a health check retry.

During the health check retry, Veeam Agent restarts the backup job to create a new restore point and transports data blocks from the Veeam Agent computer including the blocks that were corrupted in the object storage repository and the blocks that changed since the start of the backup job session that triggered the health check. Veeam Agent will not perform another health check after the job retry is finished successfully. The next health check will be run according to the defined schedule.

- If the health check detects corrupted data blocks in an inactive backup chain, Veeam Agent will not launch a health check retry. Veeam Agent will mark the backup and all related restore points as corrupted; the backup job session will end with a warning message.

#### NOTE

If you try to restore data from a corrupted backup, Veeam Agent will display a warning message informing you that the restore operation may fail or the restored data may be corrupted.

## Backup Immutability

If you store your backup files in an object storage repository, Veeam Agent allows you to protect backup data from deletion or modification by making that data temporarily immutable. It is done for increased security: immutability protects data in your recent backups from loss as a result of attacks, malware activity or any other injurious actions.

#### IMPORTANT

Backup immutability uses native object storage capabilities. You may incur additional API and storage charges from the storage provider.

## Supported Object Storage Types

Veeam Agent supports backup immutability for the following object storage types:

- Amazon S3
- S3 compatible storage that supports S3 Object Lock (including Wasabi)
- Microsoft Azure Blob Storage
- Veeam Data Cloud Vault

#### NOTE

Veeam Agent does not support backup immutability for the Google Cloud storage.

## Before You Begin

Before you configure immutability for Veeam Agent backups, you must prepare the target storage account. Depending on the selected object storage type, perform the following actions:

- [S3 Compatible and Amazon S3 storage] When you create the S3 bucket, you must enable versioning and the S3 Object Lock feature for the bucket. For more information, see [AWS documentation](#).
- [S3 Compatible and Amazon S3 storage] After you create the S3 bucket with Object Lock enabled, make sure that the default retention is disabled to avoid unpredictable system behavior and data loss. To disable the default retention, edit the Object Lock retention settings as described in [AWS documentation](#).
- [Microsoft Azure Blob storage] You must enable blob versioning and version-level immutability support in the storage account. For more information, see [Microsoft documentation](#).

Consider the following about backup immutability:

- The effective immutability period consists of the user-defined immutability period and the block generation period automatically appended by Veeam Agent. For more information, see [How Backup Immutability Works](#) and [Block Generation](#).



- [S3 Compatible and Amazon S3 storage] Veeam Agent will use the *compliance* retention mode for each uploaded object. For more information on retention modes of S3 Object Lock, see [AWS documentation](#).
- [Microsoft Azure Blob storage] Do not enable immutability for already existing containers in the Microsoft Azure Portal. Otherwise, Veeam Agent will not be able to process these containers properly and it may result in data loss.

## Configuring Backup Immutability

Depending on how you create the backup job and configure [connection to an object storage repository](#), you can define backup immutability settings in one of the following ways:

- [Backup Job wizard] You must specify the immutability period at the Bucket step of the wizard. For more information, see [Object Storage Settings](#).
- [Command line interface] You must specify the immutability period in the advanced options of the command for creating the backup job. For more information, see [Creating Backup Job with Command Line Interface](#).

### NOTE

If you want to create the backup job in command line interface, you must create the object storage repository first. For details, see [Creating Repository in Object Storage](#).

- If you create the backup job that is targeted at an object storage repository configured as a Veeam backup repository or Veeam Cloud Connect repository, the immutability period in the settings of the repository must be specified in Veeam Backup & Replication. For details, see the [Adding Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

## Backup Immutability and Retention Policy

Backup immutability operates with backup data and related metadata (checkpoints) on the object storage side. Retention policy operates with logical representation of the stored data, or restore points, on the Veeam Agent side. These two mechanisms act independently from each other.

Veeam Agent will remove the irrelevant restore points per the defined backup retention policy. If the data associated with the removed restore point is still immutable, such data will remain in the repository until expiration of the immutability period. After that it will be automatically removed from the storage.

## Limitation of Backup Immutability

If you use Veeam Agent in the standalone mode, you can restore the immutable data that is associated with a restore point removed by retention policy only in Veeam Backup & Replication console. In Veeam Backup & Replication, you must perform the following actions:

1. Add the object storage repository that contains the necessary data to Veeam Backup & Replication. For more information, see the [Adding Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.
2. Roll back to the necessary checkpoint. For more information, see the [Immutability](#) section in the Veeam PowerShell Reference.

3. Remove the repository from the Veeam Backup & Replication infrastructure. For more information, see the [Removing Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

After that, you will be able to use Veeam Agent to restore data from the object repository in a regular manner.

## How Backup Immutability Works

After you specify the immutability period for a backup and run the backup job for the first time, Veeam Agent will append an additional period of 10 days to the specified immutability period. This additional period is called *block generation*. The resulting effective immutability period is the sum of the user-defined immutability period and the block generation period. All data blocks transferred to the target repository within the block generation period will have the same immutability expiration date. For example, data block *a* added on day 1 of the block generation period will have the same immutability expiration date as block *b* added on day 9. For more information, see [Block Generation](#).

During the effective immutability period, the following operations with backup data in the object storage repository will be prohibited:

- Manual removal of data from the backup repository.
- Removal of data by backup retention policy.
- Removal of data using any object storage provider tools.
- Removal of data by the technical support department of the object storage provider.

## Extension of Effective Immutability Period

During each transfer of data to the object storage repository, Veeam Agent creates a new checkpoint file with metadata that describes the latest state of the backup in the storage. The immutable blocks of data from a previous checkpoint may be reused in the newly created checkpoint. Veeam Agent keeps reused, or dependent, blocks of data locked by continuously assigning them to new generations and extending their effective immutability period. This guarantees that the effective immutability period is no less than the immutability period defined by user.

During data transfer, the effective immutability period for the backup is set as follows:

- [For new data blocks in the checkpoint] Immutability is set anew. The user-defined immutability period is appended with a 10-day block generation period.
- [For data blocks reused from the previous checkpoint] Immutability is extended to the immutability expiration date set for the new blocks.
- [For data blocks that are not reused in the checkpoint] Immutability is not extended. Such data blocks will remain in the repository until their immutability period is over. After that Veeam Agent will automatically remove them from the repository.

## Block Generation

When you specify an immutability period for the recent backups, Veeam Agent will automatically add 10 days to the immutability expiration date. This period is called *block generation*. The block generation period serves to reduce the number of requests to the object storage repository, which results in lower traffic and reduced storage costs. You do not have to configure it, the block generation period is applied automatically.

When the block generation period is appended to the user-defined immutability period, it means there is no need to extend the immutability period for old data blocks when adding new data blocks to the backup during that block generation period.

Consider this example. When you create a full backup to start a backup chain, all data blocks transferred to the object storage repository are new. For these new blocks of data, Veeam Agent will add the block generation period of 10 days to the specified immutability period. If the immutability period is set by user to the default period of 30 days, the effective immutability period with the added block generation period will become 40 days. The first full backup starts its generation that will last for 10 days. All new and reused data blocks within this block generation period will have the same immutability expiration date. For instance, a data block that was transferred to the target repository on day 9 will have the same immutability expiration date as a data block transferred on day 1. This mechanism guarantees that the effective immutability period for all the data blocks within a generation is no less than 30 days.

If a block generation period is over but data blocks from that generation are reused in the newly created checkpoint, their effective immutability period is automatically extended to ensure that the effective immutability period for all the data blocks in the new checkpoint is no less than the user-defined immutability period. For more information, see [How Backup Immutability Works](#).

# Data Restore

Veeam Agent for Linux offers two data restore scenarios:

- You can perform volume-level restore to recover the entire system image of your computer or specific computer volumes. To learn more, see [Volume-Level Restore](#).
- You can perform file-level restore to recover individual files and directories. To learn more, see [File-Level Restore](#).

# Volume-Level Restore

If data on a computer volume gets corrupted, you can restore this volume from the backup. For volume-level restore, you can use backups that were created at the volume level. File-level backups cannot be used for volume restore.

When you perform volume-level restore, Veeam Agent for Linux restores the entire content of the volume. It retrieves from the backup data blocks pertaining to a specific volume and copies them to the necessary location.

Keep in mind that you cannot browse the volume in the backup and select individual files and directories for restore. For granular file-level restore, you can use the [File-Level Restore](#) option.

A volume can be restored to its original location or new location. If you restore the volume to its original location, Veeam Agent for Linux overwrites data on the original volume. If you restore the volume to a new location, and the target disk contains any data, Veeam Agent for Linux overwrites data in the target location with data retrieved from the backup.

## Limitations for Volume-Level Restore

Volume restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Linux swap space is hosted.
- You cannot restore a volume to the volume where the backup file used for restore is located.

To overcome the first two limitations, you can create a Veeam Recovery Media and use the **Volume Restore** wizard for volume-level restore. To learn more, see [Veeam Recovery Media](#).

# File-Level Restore

If you have lost or modified files and directories on your computer by mistake, you can restore a copy of the necessary objects from the backup. For file-level restore, you can use a backup of any type:

- Volume-level backup
- File-level backup

Veeam Agent for Linux does not simply extract files and folders from the backup file. During file-level restore, Veeam Agent for Linux performs the following operations:

1. Veeam Agent for Linux associates the backup file with a loop device, for example, `/dev/loop0`, to make the backup file accessible as a block device.
2. Veeam Agent for Linux mounts the loop device to the mount point directory in the computer's file system.
  - For file-level restore with the Veeam Agent for Linux control panel or Veeam Recovery Media, Veeam Agent for Linux mounts the backup content to the `/mnt/backup` directory.
  - For file-level restore with the command line interface, you can specify a directory in which Veeam Agent for Linux should mount the backup content.

After the backup content is mounted, you can use Linux command line utilities or preferred file browser to work with restored files and directories. You can browse for files and directories in the mounted backup and copy them to their initial location or to a new location.

# Veeam Recovery Media

Veeam Agent for Linux lets you use the Veeam Recovery Media – a recovery image of the Linux OS that provides an alternative way to boot your computer.

The recovery image includes a custom Linux OS with the limited functionality. It comprises Linux kernel and a set of GNU/Linux utilities necessary to boot the computer and perform basic administration tasks. If the OS installed on the computer fails to start for some reason, you can boot the recovery image OS. After booting, you can do the following:

- You can restore data from a backup to your computer. For this scenario, you must have a backup created with Veeam Agent for Linux.
- You can use Linux OS tools to diagnose problems and fix errors on your computer.

The recovery image can be helpful if one of the following errors occur:

- The OS on the computer fails to start.
- You want to perform bare metal restore from the backup on the computer without the OS and other software installed.
- You want to restore the system volume of the computer and so on.

Veeam Recovery Media is distributed as an ISO image. You can download the ISO image file from [this Veeam webpage](#): select an operating system to display the download links for the product and recovery ISO. You can burn the ISO image file to the following types of media:

- Removable storage devices such as USB drives or SD cards
- CD/DVD/BD

## NOTE

Consider the following:

- You can also download the Veeam Recovery Media ISO image from the [Veeam software repository](#).
- For information about how to burn the ISO image to a removable storage device, as well as workaround for accessing the Veeam recovery UI, see [this Veeam KB article](#).

When you boot from the Veeam Recovery Media, you can use the recovery environment to fix the OS system errors on your computer or restore data from the backup. Veeam Agent for Linux offers a set of tools for the computer system image and data recovery:

- Restore volumes – the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
- Restore files – the File Level Restore wizard to restore files and folders to the original location or to a new location.
- Exit to shell – Linux shell prompt with standard utilities to diagnose problems and fix errors.

# Veeam Recovery Media Versions

Veeam Agent for Linux offers Veeam Recovery Media for computers based on the x86 and x64 architecture with Linux kernel version 3.10 and later.

You cannot create custom Veeam Recovery Media for Veeam Agent computers that run Linux kernel version earlier than 3.10. For a workaround, see [this Veeam KB article](#).

You can download the recovery image from the following sources:

- [Veeam website](#)

Recovery image ISO files downloaded from the Veeam website have the following names:

- `veeam-recovery-media-6.2.0.101_i386.iso` – for Veeam Agent computers based on the x86 architecture.
- `veeam-recovery-media-6.2.0.101_x86_64.iso` – for Veeam Agent computers based on the x64 architecture.
- `veeam-recovery-media-6.2.0.101_ppc64le.iso` – for Veeam Agent computers based on the IBM Power architecture.

- [Veeam software repository](#)

Recovery image ISO files downloaded from the Veeam software repository have the following names:

- `veeam-recovery-i386-6.0.0.iso` – for Veeam Agent computers based on the x86 architecture.
- `veeam-recovery-amd64-6.0.0.iso` – for Veeam Agent computers based on the x64 architecture.
- `veeam-recovery-ppc64le-6.0.0.iso` – for Veeam Agent computers based on the IBM Power architecture.

The size of the regular recovery image file depends on the Veeam Agent computer architecture: 561 MB for x86 computers and 649 MB for x64 computers and 639 MB for IBM Power machines.



# Drivers in Veeam Recovery Media

The generic Veeam Recovery Media available for download from [the Veeam website](#) or [Veeam software repository](#) contains the following data:

1. Set of files required to start the recovery image OS from the recovery media.
2. Set of Veeam tools for the computer system image and data recovery.
3. Set of Linux command line tools to diagnose problems and fix errors on your computer. For the regular recovery image, in addition to the standard set of tools, you can install custom software from a software repository.
4. Drivers required to run hardware and devices on your computer in a regular way. The regular recovery image contains drivers included in the Linux kernel versions 6.1.0 and 6.5.0.

When you boot your computer from the Veeam Recovery Media, drivers from the Veeam Recovery Media are automatically loaded on the recovery image OS.

If your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media, you can create a custom recovery image. Veeam Agent will copy the Linux kernel running on your computer with its currently loaded modules and include them into the custom Veeam Recovery Media. To learn more, see [Creating Custom Veeam Recovery Media](#).

# Integration with Veeam Backup & Replication

## IMPORTANT

To use Veeam Agent for Linux 6.2 with Veeam Backup & Replication, you must install Veeam Backup & Replication 12.2 on the Veeam backup server.

You can store backup files created with Veeam Agent for Linux on backup repositories managed by Veeam Backup & Replication. To do this, you must select a Veeam Backup & Replication backup repository as a target location in the properties of the backup job. To learn more about supported backup repositories, see the [Backup Repositories](#) and [Scale-Out Backup Repositories](#) sections in the Veeam Backup & Replication User Guide.

## NOTE

Consider the following:

- The current guide covers subjects related to Veeam Agent for Linux operating in the standalone mode.
- You can also use Veeam Backup & Replication to manage Veeam Agent for Linux on computers in your infrastructure. As part of the Veeam Agent management scenario, you can remotely deploy Veeam Agent to your computers, as well as configure and manage Veeam Agent backup jobs in Veeam Backup & Replication. To learn more, see [Veeam Agent Management Guide](#).
- If you create a backup job with the Veeam Agent command line interface, you need to specify a Veeam backup repository in the backup job settings. Veeam backup repository appears in the list of backup repositories after you connect to a Veeam backup server. To learn more, see [Managing Veeam Backup & Replication Servers](#).

Veeam Agent for Linux works with the Veeam Backup & Replication backup repository as with any other backup repository. Backup files are stored to a separate folder; you can perform standard restore operations using these files.

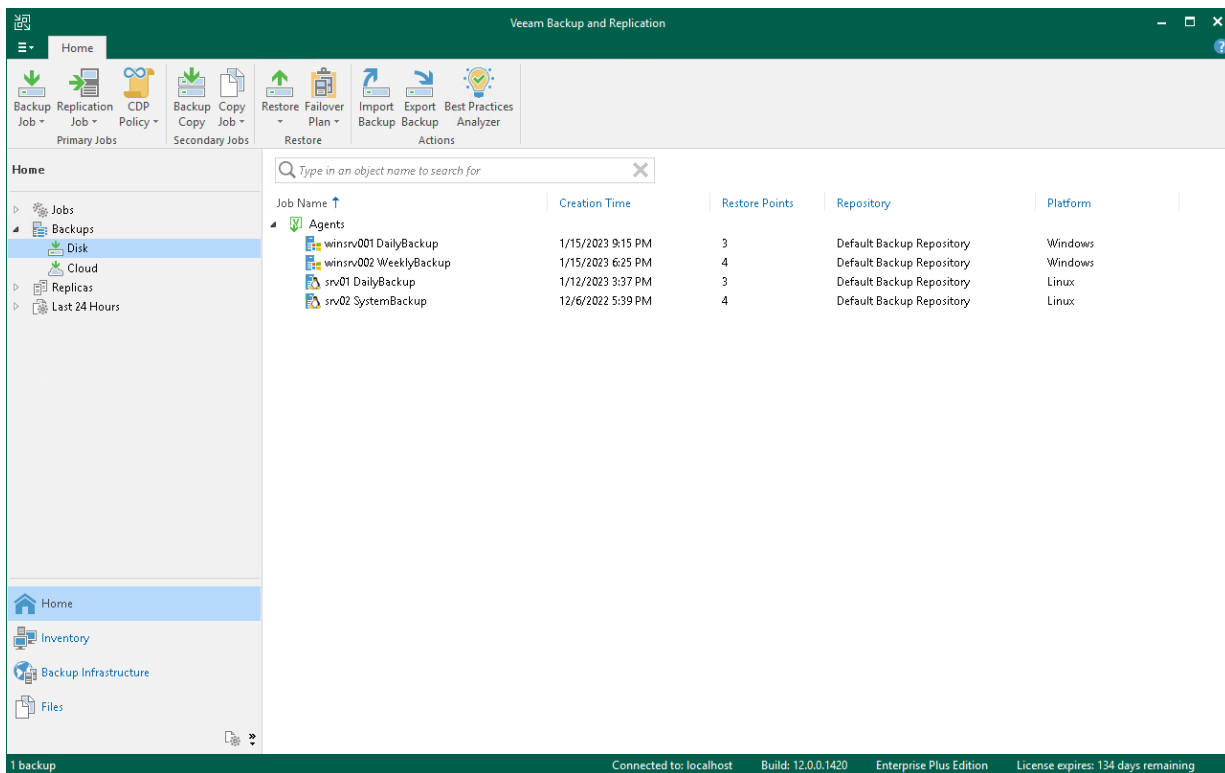
Information about Veeam Agent backups stored on the Veeam Backup & Replication backup repositories, backup jobs and sessions becomes available in the Veeam Backup & Replication console:

- The Veeam Agent for Linux backup job is displayed in the list of jobs in Veeam Backup & Replication.
- Backup files created with Veeam Agent for Linux are displayed in the list of backups, under the **Agents** node.
- Performed job sessions are available in the **History** view of Veeam Backup & Replication.

Backup administrators working with Veeam Backup & Replication can perform a set of operations with Veeam Agent backups:

- Perform data protection operations: copy Veeam Agent backups to secondary backup repositories and archive these backups to tape.
- Perform restore operations: restore individual files and directories, application items from Veeam Agent backups; restore computer disks and convert them to the VMDK, VHD or VHDX format; restore to Microsoft Azure and Amazon EC2.

- Perform administrative tasks: disable and delete Veeam Agent backup jobs, remove Veeam Agent backups and so on.



# Backup to Veeam Cloud Connect Repository

If you want to store your data in the cloud, you can connect to a Veeam Cloud Connect service provider (SP) and create Veeam Agent backups in a cloud repository. To do this, you must provide credentials of the tenant (or subtenant) account that you obtained from the SP and select a cloud repository as a target for backup files in the properties of the backup job. To learn more, see [Veeam Cloud Connect Repository Settings](#).

## NOTE

Consider the following:

- You can create Veeam Agent backups in a cloud repository if the SP backup server runs Veeam Backup & Replication 12.0 or later.
- Backup to a cloud repository is available if Veeam Agent for Linux operates in the Workstation or Server edition.

# Managing Veeam Agent in Veeam Backup & Replication

Veeam Backup & Replication lets you automate management of Veeam Agent on multiple computers in your infrastructure. You can deploy Veeam Agent for Linux, configure Veeam Agent backup jobs and perform other data protection and administration tasks on remote computers. To use the Veeam Agent management functionality in Veeam Backup & Replication, you must install Veeam Backup & Replication on the Veeam backup server.

To learn more, see [Veeam Agent Management Guide](#).

# Planning and Preparation

Before you install Veeam Agent for Linux, make sure that the target computer meets the system requirements, and all required ports are open.

# System Requirements

The protected Linux computer must meet requirements listed in the table below.

## NOTE

The following system requirements apply to the following Veeam Agent for Linux configuration:

- Veeam Agent for Linux version is 6.2.
- Veeam Agent for Linux is operating in the standalone mode.  
To learn about system requirements for Veeam Agent managed by Veeam Backup & Replication, see the [System Requirements](#) section in the Veeam Agent Management Guide.
- Veeam Agent for Linux is installed with the `veeam-libs`, `veeam`, and Veeam kernel module packages.  
To learn about system requirements for nosnap Veeam Agent for Linux, see [System Requirements for Nosnap Veeam Agent for Linux](#).  
To learn about system requirements for nosnap Veeam Agent for Linux on Power, see [System Requirements for Nosnap Veeam Agent for Linux on Power](#).

Specification	Requirement
Hardware	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the supported hardware.</p> <p>CPU: x86 or x64.</p> <p>Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. For more information, see <a href="#">RAM Requirements for Backup of Large Number of Files</a>.</p> <p>Disk Space: 100 MB free disk space for product installation.</p> <p>Network: 10 Mbps or faster network connection to a backup target.</p> <p>System firmware: BIOS or UEFI.</p> <p>Disk layout: MBR or GPT.</p>

Specification	Requirement
OS	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of supported OSes.</p> <p>Linux kernel version 2.6.32 to version 6.10 is supported.</p> <p>Veeam Agent supports the 64-bit versions of the following Linux distributions:</p> <ul style="list-style-type: none"> <li>• Debian 10.13 – 12.6</li> <li>• Ubuntu 16.04, 18.04, 20.04, 22.04, 22.10, 23.04, 23.10 and 24.04</li> <li>• RHEL 6.4 – 9.4</li> <li>• CentOS 7</li> <li>• Oracle Linux 6 – 9.4 (RHCK)</li> <li>• Oracle Linux 6 (starting from UEK R2) – Oracle Linux 8 (up to UEK R6)</li> <li>• Oracle Linux 8 (UEK R7) – for information on installation, see <a href="#">this Veeam KB article</a>.</li> <li>• Oracle Linux 9 (up to 5.15.0-209.161.7.2.el9uek)</li> <li>• SLES 12 SP4, 12 SP5, 15 SP1 – 15 SP6</li> <li>• SLES for SAP 12 SP4, 12 SP5, 15 SP1 – 15 SP6</li> <li>• Fedora 36, 37, 38 and 39</li> <li>• openSUSE Leap 15.3 – 15.6</li> <li>• openSUSE Tumbleweed has an <a href="#">experimental support</a> status.</li> <li>• Rocky Linux 9.3 and 9.4</li> <li>• AlmaLinux 9.3 and 9.4</li> </ul> <p>Veeam Agent supports 32-bit versions for RHEL 6 and Oracle Linux 6 distributions only.</p>



Specification	Requirement
File System	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of supported file systems.</p> <p>Veeam Agent for Linux supports consistent snapshot-based data backup for the following file systems:</p> <ul style="list-style-type: none"> <li>• BTRFS (for OSES that run Linux kernel 3.16 or later)</li> <li>• Ext 2/3/4</li> <li>• F2FS</li> <li>• FAT16</li> <li>• FAT32</li> <li>• HFS</li> <li>• HFS+</li> <li>• JFS</li> <li>• NILFS2</li> <li>• NTFS</li> <li>• ReiserFS</li> <li>• XFS</li> </ul> <p>The supported file system (except for BTRFS) can reside on a simple volume or LVM2 volume; volumes protected with encryption software such as dm-crypt are supported. BTRFS is supported only if it resides directly on a physical device with no additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) below or above it.</p> <p>Other file systems, file systems that are not located on logical volumes, as well as network file systems like NFS or SMB shares can be backed up using the snapshot-less mode only. For details, see <a href="#">Snapshot-Less File-Level Backup</a>.</p>

Specification	Requirement
Software	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of supported components.</p> <p>Protected computer must have the following components installed:</p> <ul style="list-style-type: none"> <li>• dkms</li> <li>• gcc</li> <li>• make</li> <li>• perl</li> <li>• linux-headers (for Debian-based systems)</li> <li>• kernel-headers (for RedHat-based systems)</li> <li>• kernel-devel (for RedHat-based systems)</li> <li>• kernel-uek-devel (for Oracle Linux with UEK)</li> <li>• libudev</li> <li>• libacl</li> <li>• libattr</li> <li>• lvm2</li> <li>• libfuse2 (FUSE libraries for Debian-based and SLES-based systems)</li> <li>• fuse-libs (FUSE libraries for RedHat-based and Fedora systems)</li> <li>• libncurses5</li> <li>• dmidecode</li> <li>• libmysqlclient</li> <li>• libpq5</li> <li>• python3</li> <li>• efibootmgr (for UEFI-based systems)</li> <li>• isolinux (for Debian-based systems)</li> <li>• syslinux (for RedHat-based systems)</li> <li>• btrfs-progs (for backup of BTRFS file system)</li> <li>• mksquashfs (for custom Veeam Recovery Media)</li> <li>• unsquashfs (for custom Veeam Recovery Media)</li> <li>• wget (for custom Veeam Recovery Media)</li> <li>• xorriso (for custom Veeam Recovery Media with EFI support)</li> <li>• tar (for file system indexing, log export and rotation)</li> <li>• gzip (for file system indexing, log export and rotation)</li> </ul>

## Considerations and Limitations

### Hardware

- For virtual machines, only full virtualization type is supported. Oracle VM virtual machines are supported with [limitations](#). Virtual I/O (VirtIO) devices have [experimental support](#) status. Other containers and paravirtualized instances are not supported; backup of such devices may result in corruption of the source file system – for more information, see [this Veeam KB article](#).
- Devices managed by Veritas Volume Manager are not supported.

## OS

- Linux kernel version 2.6.32 to version 6.10 is supported as long as you use kernels supplied by your distribution.

Consider the following limitations:

- Fedora 38, 39 and openSUSE Tumbleweed are supported up to kernel 6.10.
- Linux kernel 2.6.32-754.6.3 in CentOS / RHEL and Oracle Linux (RHCK) is not supported.
- Only GA versions of the [supported distributions](#) that have been released before the current version of Veeam Agent for Linux are supported.

If a new version of a supported Linux distribution is released after the release of the current version of Veeam Agent, Veeam Agent may require a patch to support this new OS version. For details on Veeam Agent compatibility with Linux OS versions, see [this Veeam KB article](#). Customers with a valid contract can request a patch from Veeam Support; for other customers, the support of the new Linux distribution will be provided with the next release of Veeam Agent.

- For the following distributions, we recommend installing Veeam kernel modules from pre-built binary packages provided by Veeam:
  - CentOS 7
  - RHEL 6.4 - 9.4
  - Rocky Linux 9.3 and 9.4
  - AlmaLinux 9.3 and 9.4
  - SLES 12 SP4, 12 SP5 15 SP1 - 15 SP6
  - SLES for SAP 12 SP4, 12 SP5 15 SP1 - 15 SP6
  - openSUSE Leap 15.3 - 15.6

For other supported distributions, use the `dkms` packages instead of the pre-built binary packages with Veeam kernel modules.

Consider the following about Veeam kernel modules from pre-built binary packages:

- Pre-built binary `veeamsnap` kernel module packages require kernel 2.6.32-131.0.15 or later for RHEL 6 (excluding 2.6.32-279.el6.i686) and 3.10.0-123 or later for CentOS / RHEL 7.0 - 7.9.
- Pre-built binary `blksnap` kernel module packages require kernel 5.3.18 or later.

For details on installing Veeam Agent on every supported distribution, see [Installing Veeam Agent for Linux](#).

- To ensure proper functioning of the Veeam kernel module, verify that your system does not have any of the following modules installed: `hcdpdriver`, `snapi26`, `snapi`, `snapper`, `dattobd`, `dattobd-dkms`, `dkms-dattobd`, `cdr` or `cxbf`.
- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.
- For cloud-based installations that use customized kernels (such as Linux distributions deployed from AWS Marketplace or Azure Marketplace), the Veeam kernel module has experimental support status. For details about experimental support, see [this Veeam KB article](#).
- RHEL, CentOS and Oracle Linux (RHCK) are supported up to certain kernel versions. For details, see [this Veeam KB article](#).

- Ubuntu with Linux kernel for KVM (Kernel-based Virtual Machine) is not supported. For the list of linux-kvm kernels for Ubuntu, see [Ubuntu documentation](#).
- [Oracle Linux (UEK) 6.6 – 7.4] If the operating system has the FIPS mode enabled, you must sign the DKMS Veeam kernel module. For more information on automating the process of signing DKMS kernel modules, see [Linux documentation](#).

## File System

- File-level backup has the following limitations:
    - Total size of all file systems must not exceed 216 TiB. This limitation applies to all file systems where files you plan to back up are located.
    - Size of a file included in a file-level backup must not exceed 16 TB.
    - Name of a file must not be larger than 254 bytes.

Keep in mind that characters that you can use in the file name may be encoded in 2 bytes or more.
  - To store volume snapshots, the `blksnap` kernel module requires an Ext4, BTRFS or XFS file system.
  - Veeam Agent supports backup of extended attributes with the following limitations:
    - Veeam Agent backs up extended attributes only with the following public namespaces: `system`, `security`, `trusted`, and `user`.
    - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the `ea_inodes` feature. Backups created using the `ea_inodes` feature cannot be mounted on kernel versions up to 4.12.
  - Each volume included in a backup must have a unique UUID.
  - Veeam kernel module provide a RAM-based changed block tracking (CBT) mechanism. Every time the module is unloaded or Veeam Agent for Linux computer is rebooted, CBT data is reset. As a result, Veeam Agent reads the entire data added to the backup scope to detect what blocks have changed since the last job session, and incremental backup requires greater time.
  - Backup of computers used as cluster nodes can be performed by Veeam Agent for Linux in the [Snapshot-Less File-Level Backup](#) mode only.
- Backup of computers used as cluster nodes can be also performed by nosnap Veeam Agent for Linux. For details, see [System Requirements for Nosnap Veeam Agent for Linux](#).
- Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).
  - Backup of file and directory attributes (for example, `a` – append only, `c` – compressed, and so on) is not supported.
  - Veeam Agent for Linux does not back up volumes that reside on USB devices and SD cards.
  - Veeam Agent for Linux does not back up LVM snapshots.
  - BFQ I/O scheduler is not supported.
  - Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.
  - Backup of pseudo file systems, such as `/proc`, `/sys`, `tmpfs`, `devfs` and others, is not supported.
  - During backup, network file systems are skipped unless explicitly included into the backup scope.

- Backup of BTRFS volumes and subvolumes with enabled file-system compression is not supported.

## Software

### IMPORTANT

Linux user account used to work with Veeam Agent for Linux must have the `/bin/bash` shell set as the default shell.

- The following packages are not required for CentOS, RHEL and SLES distributions if a pre-built binary package with Veeam kernel module is to be installed.
  - dkms
  - gcc
  - make
  - perl
  - kernel-headers (for RedHat-based systems)
  - kernel-devel (for RedHat-based systems)

For details, see [Installing Veeam Agent for Linux](#).

- Version of the following packages varies according to the Linux kernel version that you use:
  - linux-headers (for Debian-based systems)
  - kernel-headers (for RedHat-based systems)
  - kernel-devel (for RedHat-based systems)
  - kernel-uek-devel (for Oracle Linux systems with UEK)
- For openSUSE and SLES distributions, either of the following packages is required: `libncurses5` or `libncurses6`.
- The `dmidecode` package is required for Veeam Agent management — a valid BIOS UUID must be obtainable either from `dmidecode | grep -i uuid` or from `/sys/class/dmi/id/product_uuid`. Each Veeam Agent that consumes a license installed in Veeam Backup & Replication must have a unique BIOS UUID. If a valid UUID cannot be obtained, Veeam Agent will generate it automatically.
- The `libmysqlclient` package is required to process MySQL database system located on the Veeam Agent server. For details, see [Backup of MySQL Database](#). Package version varies according to the MySQL database system version that you use.
- The `libpq5` package is required to process PostgreSQL database system located on the Veeam Agent server. For details, see [Backup of PostgreSQL Database](#).
- The `python3` package or another RPM package providing a `/usr/bin/python3` binary is required for CentOS, RHEL 7.0 and later distributions if a pre-built binary Veeam kernel module package is to be installed.
- The `btrfs-progs` package version 3.16 or later is required.

## Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see [File System](#).

# Backup Target

Backup can be performed to the following types of storage:

- On-premises or cloud-based object storage.
- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.
- Veeam Backup & Replication 12.1 or later backup repository (including deduplication appliances).
- Veeam Cloud Connect 12.0 or later backup repository.

## IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

# Network

Consider the following:

- If you back up to a repository managed by a Veeam backup server, Veeam Agent for Linux must be able to establish a direct IP connection to the Veeam Backup & Replication server. Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 or IPv6 addresses.

# System Requirements for Nosnap Veeam Agent for Linux

You can install Veeam Agent for Linux using a `veeam-nosnap` package. This package allows Veeam Agent to operate without Veeam kernel module.

The `veeam-nosnap` package can be useful in the following cases:

- You do not want to install kernel sources and compilers on your computer.
- You want to use third-party tools to create data snapshots.
- You want to perform bare metal restore, but Veeam Recovery Media does not work with your computer. In this case you can install the `veeam-nosnap` package on LiveCD of your choice and access the Veeam recovery UI.
- You want to back up machines that are used as cluster nodes.

Before you install Veeam Agent using the `veeam-nosnap` package, consider the following limitations:

- The RAM-based changed block tracking (CBT) mechanism is not supported. As a result, if you plan to back up a significant amount of data, the backup will require greater time.
- Veeam Agent can create a snapshot of LVM logical volumes and BTRFS subvolumes. To back up data that resides on other file systems and volumes, you can use only file-level backup in the snapshot-less mode. For details, see [Snapshot-Less File-Level Backup](#).
- For a successful backup, Veeam Agent requires unallocated extents on volume groups.
- For a successful bare metal restore, all disks of the Veeam Agent computer you want to restore must be available in the backup.

# System Requirements for nosnap Veeam Agent for Linux

If you plan to use the `veeam-nosnap` package to install Veeam Agent, the protected Linux computer meet the requirements listed in the table below. To learn about system requirements for Veeam Agent installed using a Veeam kernel module package, see [System Requirements](#).

Specification	Requirement
Hardware	<p>CPU: x86 or x64.</p> <p>Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. For more information, see <a href="#">RAM Requirements for Backup of Large Number of Files</a>.</p> <p>Disk Space: 100 MB free disk space for product installation.</p> <p>Network: 10 Mbps or faster network connection to a backup target.</p> <p>System firmware: BIOS or UEFI.</p> <p>Disk layout: MBR or GPT.</p> <p>For virtual machines: Only full virtualization type is supported. Containers and paravirtualized instances are not supported. Oracle VM virtual machines are supported with <a href="#">limitations</a>.</p>
OS	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of supported OSes.</p> <p>Veeam Agent supports the 64-bit versions of the following distributions:</p> <ul style="list-style-type: none"><li>• Debian 10.13 – 12.6</li><li>• Ubuntu 16.04, 18.04, 20.04, 22.04, 22.10, 23.04, 23.10 and 24.04</li><li>• RHEL 6.4 – 9.4</li><li>• CentOS 7</li><li>• Oracle Linux 6 – 9.4 (RHCK)</li><li>• Oracle Linux 6 (starting from UEK R2) – Oracle Linux 9 (up to 5.15.0-209.161.7.2.el9uek)</li><li>• SLES 12 SP4, 12 SP5, 15 SP1 – 15 SP6</li><li>• SLES for SAP 12 SP4, 12 SP5, 15 SP1 – 15 SP6</li><li>• openSUSE Leap 15.3 – 15.6</li><li>• openSUSE Tumbleweed has an experimental support status. For details about experimental support, see <a href="#">this Veeam KB article</a>.</li><li>• Rocky Linux 9.3 and 9.4</li><li>• AlmaLinux 9.3 and 9.4</li></ul> <p>Veeam Agent supports 32-bit versions for RHEL 6 and Oracle Linux 6 distributions only.</p>



Specification	Requirement
File System	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of the supported file systems.</p> <p>Veeam Agent for Linux supports consistent snapshot-based data backup for the following file systems:</p> <ul style="list-style-type: none"> <li>• All <a href="#">supported file systems</a> that are built on top of LVM logical volumes.</li> <li>• BTRFS (for OSes that run Linux kernel 3.16 or later).</li> </ul> <p>BTRFS is supported only if it resides directly on a physical device with no additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) below or above it.</p> <p>Supported file systems that are not located on logical volumes, other file systems and network file systems like NFS or SMB shares can be backed up using the snapshot-less mode only. For details, see <a href="#">Snapshot-Less File-Level Backup</a>.</p>
Software	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of supported components.</p> <p>Protected computer must have the following components installed:</p> <ul style="list-style-type: none"> <li>• libacl</li> <li>• libattr</li> <li>• lvm2</li> <li>• libfuse</li> <li>• dmidecode</li> <li>• efibootmgr (for UEFI-based systems)</li> <li>• isolinux (for Debian-based systems)</li> <li>• syslinux (for RedHat-based systems)</li> <li>• btrfs-progs (for backup of BTRFS file system)</li> <li>• mksquashfs (for custom Veeam Recovery Media)</li> <li>• unsquashfs (for custom Veeam Recovery Media)</li> <li>• wget (for custom Veeam Recovery Media)</li> <li>• xorriso (for custom Veeam Recovery Media with EFI support)</li> <li>• tar (for file system indexing, log export and rotation)</li> <li>• gzip (for file system indexing, log export and rotation)</li> </ul>

# Considerations and Limitations

## OS

- Only GA versions of the [supported distributions](#) that have been released before the current version of Veeam Agent for Linux are supported.

If a new version of a supported Linux distribution is released after the release of the current version of Veeam Agent, Veeam Agent may require a patch to support this new OS version. Customers with a valid contract can request a patch from Veeam Support; for other customers, the support of the new Linux distribution will be provided with the next release of Veeam Agent.

- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.

## File System

- Veeam Agent for Linux does not back up volumes that reside on USB devices and SD cards.
- LVM volumes encrypted with dm-crypt software are not supported.
- Total size of all file systems must not exceed 216 TiB. This limitation applies to all file systems where files you plan to back up are located.
- Size of a file included in a file-level backup must not exceed 16 TB.
- Name of a file must not be larger than 254 bytes.

Keep in mind that characters that you can use in the file name may be encoded in 2 bytes or more.

- The amount of space required for LVM snapshots largely depends on the IO intensity. Generally, from 10% to 20% of the system's occupied space should be enough for storing an LVM snapshot.
- Veeam Agent supports backup of extended attributes with the following limitations:
  - Veeam Agent backs up extended attributes only with the following public namespaces: `system`, `security`, `trusted`, and `user`.
  - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the `ea_inodes` feature. Backups created using the `ea_inodes` feature cannot be mounted on kernel versions up to 4.12.

- Backup of file and directory attributes (for example, `a` – append only, `c` – compressed, and so on) is not supported.
- Each volume included in a backup must have a unique UUID.
- Consider the following about the backup of computers used as cluster nodes:
  - To back up data on local LVM volumes, you can use [file-level backup](#) or [volume-level backup](#).

## NOTE

Consider the following:

- During volume-level backup, data from shared disks, clustered file systems or clustered LVM will not be backed up.
  - To perform volume-level backup, Veeam Agent for Linux will create an LVM snapshot, which may cause instability of the cluster or cluster software. This can happen due to the failover conditions configured for the cluster. However, if the cluster instability is caused by creation of an LVM snapshot only during backup, please contact Veeam support for assistance.
- Backup of clustered file systems using a native file system snapshot is not supported. This includes snapshots created with the help of custom pre-job or post-job scripts.
  - The following objects can be backed up only by [snapshot-less file-level backup](#):
    - Files on shared disks, clustered file systems or clustered LVM.
    - Files on local file systems that are not located on LVM logical volumes.
  - Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).
  - Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.
  - Backup of pseudo file systems, such as `/proc`, `/sys`, `tmpfs`, `devfs` and others, is not supported.

## Software

### IMPORTANT

Linux user account used to work with Veeam Agent for Linux must have the `/bin/bash` shell set as the default shell.

- The `dmidecode` package is required for Veeam Agent management — a valid BIOS UUID must be obtainable either from `dmidecode | grep -i uuid` or from `/sys/class/dmi/id/product_uuid`. Each Veeam Agent that consumes a license installed in Veeam Backup & Replication must have a unique BIOS UUID. If a valid UUID cannot be obtained, Veeam Agent will generate it automatically.
- The `btrfs-progs` package version 3.16 or later is required.

## Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see [File System](#).

## Backup Target

Backup can be performed to the following types of storage:

- On-premises or cloud-based object storage.
- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.

- Veeam Backup & Replication 12.1 or later backup repository (including deduplication appliances).
- Veeam Cloud Connect 12.0 or later cloud repository.

### **IMPORTANT**

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

## **Network**

Consider the following:

- Veeam Agent for Linux should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 or IPv6 addresses.

# System Requirements for Nosnap Veeam Agent for Linux on Power

If you want to back up Linux computers running on IBM Power Systems, you can install Veeam Agent for Linux using the `veeam-nosnap` package for Linux on Power. This package allows Veeam Agent to operate without Veeam kernel module using the native file system snapshots instead.

Before you install Veeam Agent using the `veeam-nosnap` package for Linux on Power, consider the following limitations:

- The RAM-based changed block tracking (CBT) mechanism is not supported. As a result, if you plan to back up a significant amount of data, the backup will require greater time.
- Veeam Agent can create a snapshot of LVM logical volumes and BTRFS subvolumes. To back up data that resides on other file systems and volumes, you can use only file-level backup in the snapshot-less mode. For details, see [Snapshot-Less File-Level Backup](#).
- For a successful backup, Veeam Agent requires unallocated extents on volume groups.
- For a successful bare metal restore, all disks of the Veeam Agent computer you want to restore must be available in the backup.

## System Requirements for Veeam Agent for Linux on Power

If you plan to install Veeam Agent for Linux on Power, make sure the protected Linux computer meets the requirements listed in the table below. To learn about system requirements for Veeam Agent for Linux installed using Veeam kernel module package, see [System Requirements](#).

Specification	Requirement
Hardware	<p>System: IBM Power System</p> <p>CPU: IBM POWER9 or POWER10</p> <p>Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. For more information, see <a href="#">RAM Requirements for Backup of Large Number of Files</a>.</p> <p>Disk Space: 100 MB free disk space for product installation</p> <p>Network: 10 Mbps or faster network connection to a backup target</p> <p>Disk layout: MBR or GPT</p>
OS	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of supported OSes.</p> <p>Veeam Agent supports little endian versions of the following Linux distributions for IBM Power:</p> <ul style="list-style-type: none"><li>• SLES 15 SP3 and 15 SP4</li><li>• SLES for SAP 12 SP5, 15 SP3 and 15 SP4</li><li>• RHEL 8.4 and 8.6</li><li>• RHEL for SAP 8.4</li></ul>

Specification	Requirement
File System	<p><b>Important!</b> Check <a href="#">considerations and limitations</a> that apply to the list of supported file systems.</p> <p>Veeam Agent for Linux on Power supports consistent snapshot-based data backup for the following file systems:</p> <ul style="list-style-type: none"> <li>• All <a href="#">supported file systems</a> that are built on top of LVM logical volumes.</li> <li>• BTRFS (for OSes that run Linux kernel 3.16 or later)</li> </ul> <p>If BTRFS has additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) above it, only file-level restore operations are supported. Instant Recovery, restore verification (SureBackup), bare metal recovery and volume-level restore are not supported.</p> <p>Supported file systems that are not located on logical volumes, other file systems and network file systems like NFS or SMB shares can be backed up using the snapshot-less mode only. For details, see <a href="#">Snapshot-Less File-Level Backup</a>.</p>
Software	<p><b>Important!</b> Linux user account used to work with Veeam Agent for Linux on Power must have the <code>/bin/bash</code> shell set as the default shell.</p> <p>Protected computer must have the following components installed:</p> <ul style="list-style-type: none"> <li>• libacl</li> <li>• libattr</li> <li>• lvm2 <ul style="list-style-type: none"> <li>• libfuse2 (FUSE libraries for SLES-based systems)</li> <li>• fuse-libs (FUSE libraries for RedHat-based systems)</li> <li>• syslinux (for RedHat-based systems)</li> <li>• btrfs-progs (version 3.16 or later, for backup of BTRFS file system)</li> <li>• tar (for file system indexing, log export and rotation)</li> <li>• gzip (for file system indexing, log export and rotation)</li> </ul> </li> </ul>

## Considerations and Limitations

### OS

- Only GA versions of the [supported distributions](#) that have been released before the current version of Veeam Agent for Linux for Power are supported.
- If a new version of a supported Linux distribution is released after the release of the current version of Veeam Agent, Veeam Agent may require a patch to support this new OS version. Customers with a valid contract can request a patch from Veeam Support; for other customers, the support of the new Linux distribution will be provided with the next release of Veeam Agent.
- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.

## File System

- Veeam Agent does not back up volumes that reside on USB devices and SD cards.
- LVM volumes encrypted with dm-crypt software are not supported.
- Total size of all file systems must not exceed 216 TiB. This limitation applies to all file systems where files you plan to back up are located.
- Size of a file included in a file-level backup must not exceed 16 TB.
- Name of a file must not be larger than 254 bytes.

Keep in mind that characters that you can use in the file name may be encoded in 2 bytes or more.

- The amount of space required for LVM snapshots largely depends on the IO intensity. Generally, from 10% to 20% of the system's occupied space should be enough for storing an LVM snapshot.
- Veeam Agent supports backup of extended attributes with the following limitations:
  - Veeam Agent backs up extended attributes only with the following public namespaces: `system`, `security`, `trusted`, and `user`.
  - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the `ea_inodes` feature. Backups created using the `ea_inodes` feature cannot be mounted on kernel versions up to 4.12.
- Backup of file and directory attributes (for example, `a` – append only, `c` – compressed, and so on) is not supported.
- Each volume included in a backup must have a unique UUID.
- Consider the following about the backup of machines used as cluster nodes:
  - To back up data on local LVM volumes, you can use [file-level backup](#) or [volume-level backup](#).

### NOTE

Consider the following:

- During volume-level backup, data from shared disks, clustered file systems or clustered LVM will not be backed up.
- To perform volume-level backup, Veeam Agent for Linux will create an LVM snapshot, which can cause instability of the cluster or cluster software. This can happen due to the failover conditions configured for the cluster. However, if the cluster instability is caused by creation of an LVM snapshot only during backup, please contact Veeam support for assistance.
- Backup of clustered file systems using a native file system snapshot is not supported. This includes snapshots created with the help of custom pre-job or post-job scripts.
- The following objects can be backed up only by [snapshot-less file-level backup](#):
  - Files on shared disks, clustered file systems or clustered LVM
  - Files on local file systems that are not hosted by LVM
- Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).

- Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.
- Backup of pseudo file systems, such as `/proc`, `/sys`, `tmpfs`, `devfs` and others, is not supported.

## Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see [File System](#).

## Backup Target

Backup can be performed to the following types of storage:

- On-premises or cloud-based object storage.
- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.
- Veeam Backup & Replication 12.1 or later backup repository (including deduplication appliances).

### IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

## Network

Consider the following:

- Veeam Agent for Linux should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 or IPv6 addresses.



# RAM Requirements for Backup of Large Number of Files

Amount of RAM used by Veeam Agent to process backed-up files depends on the number of files included in the backup, the length of file names and depth of the directory structure. For large environments with great number of backed-up files, consider the following RAM sizing recommendations:

*For large number of backed-up files with short names (7 characters) under a single-level root directory*

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	1.6
1,000,000 Incremental + 100,000	2.1
2,000,000 Full	2.0
2,000,000 Incremental + 100,000	3.5
5,000,000 Full	4.4
5,000,000 Incremental + 100,000	7.2
10,000,000 Full	7.3
10,000,000 Incremental + 100,000	13.6

*For large number of backed-up files with long names (254 characters) under a single-level root directory*

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	2.8
1,000,000 Incremental + 100,000	4.6
2,000,000 Full	4.6
2,000,000 Incremental + 100,000	8.2
5,000,000 Full	10.0
5,000,000 Incremental + 100,000	19.2

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
10,000,000 Full	19.0
10,000,000 Incremental + 100,000	36.7

*For large number of backed-up files with short names (7 characters) in a 7-level directory structure*

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	2.3
1,000,000 Incremental + 100,000	2.7
2,000,000 Full	2.7
2,000,000 Incremental + 100,000	3.4
5,000,000 Full	3.9
5,000,000 Incremental + 100,000	5.3
10,000,000 Full	5.8
10,000,000 Incremental + 100,000	8.5

*For large number of backed-up files with long names (254 characters) in a 7-level directory structure*

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	2.8
1,000,000 Incremental + 100,000	3.5
2,000,000 Full	3.7
2,000,000 Incremental + 100,000	4.9
5,000,000 Full	6.4
5,000,000 Incremental + 100,000	9.3

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
10,000,000 Full	11.1
10,000,000 Incremental + 100,000	16.2

# Permissions

Depending on the scenario, the user accounts must have the permissions listed in the following subsections:

- [Permissions for Backup to Object Storage](#)
- [Permissions for Guest Processing](#)

## Permissions for Backup to Object Storage

If you plan to back up data to object storage, make sure that the user account that you use to connect to the object storage has the required permissions. The list of required permissions differs depending on the selected object storage:

- [Amazon S3 or S3 compatible](#)
- [Google Cloud Storage](#)

### Amazon S3 or S3 compatible

If you plan to back up data to the Amazon S3 or S3 compatible storage, make sure the user account that you plan to use has the following permissions:

Identity-based permission:

```
{
  "s3:ListAllMyBuckets"
}
```

Resource-based permissions:

```
{
  "s3:DeleteObject",
  "s3:GetBucketLocation",
  "s3:GetBucketObjectLockConfiguration",
  "s3:GetBucketVersioning",
  "s3:GetObject",
  "s3:ListBucket",
  "s3:PutObject"
}
```

#### TIP

For information about required permissions for Amazon S3 storage with immutability enabled, see the [Using Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

## Google Cloud Storage

If you plan to back up data to the Google Cloud storage, make sure the user account that you plan to use has the following permissions:

```
{
  "storage.buckets.get",
  "storage.buckets.list",
  "storage.objects.create",
  "storage.objects.delete",
  "storage.objects.get",
  "storage.objects.list"
}
```

## Permissions for Guest Processing

To use guest processing, make sure to configure user accounts according to the requirements listed in this section.

Consider the following general requirements when choosing a user account:

- The user account must have root privileges.
- The user account must have the home directory created.

Depending on the application you need to back up, the user account must have the permissions listed in the table below:

Application	Required Permission
MySQL	<p>To process the MySQL database system, the MySQL user account must have the following privileges:</p> <ul style="list-style-type: none"><li>○ SELECT for all tables. This privilege is required to allow Veeam Agent to access table metadata. To learn more, see <a href="#">MySQL documentation</a>.</li><li>○ LOCK TABLES. This privilege is required to allow Veeam Agent to process tables based on the MyISAM storage engine.</li><li>○ RELOAD. This privilege is required to allow the MySQL user account to perform FLUSH operations.</li></ul>
Oracle	<p>To back up Oracle data, the user account must be granted <i>SYSDBA</i> privileges. You can use either the same account that was specified at the <b>Guest Processing</b> step if such an account is a member of the <i>OSDBA</i> and <i>OINSTALL</i> groups, or you can use any other account that has <i>SYSDBA</i> privileges.</p> <p>To perform guest processing for Oracle databases on Linux servers, make sure that the <code>/tmp</code> directory is mounted with the <code>exec</code> option. Otherwise, you will get a "<i>Permission denied</i>" error.</p>
PostgreSQL	<p>To back up PostgreSQL instances, the user account must have the superuser privileges for the PostgreSQL instance. For more information, see <a href="#">PostgreSQL documentation</a>.</p>

# Ports

The following tables describe network ports that must be opened to ensure proper communication of Veeam Agent operating in the standalone mode with other infrastructure components.

To learn about ports required to enable proper work of Veeam Agent for Linux managed by Veeam Backup & Replication, see the [Ports](#) section in the Veeam Agent Management Guide.

## IMPORTANT

The list of ports required for computers booted from the Veeam Recovery Media is the same as the list of ports required for Veeam Agent computers.

## Communication Between Veeam Agent Components

The following table describes network ports that must be opened to enable proper communication between Veeam Agent for Linux components.

From	To	Protocol	Port	Notes
Veeam Agent computer	Veeam backup server	TCP	10002, 10006	Default ports used for communication with the Veeam backup server.  Data between the Veeam Agent for Linux computer and backup repositories is transferred directly, bypassing Veeam backup servers.
	Shared folder SMB (CIFS) share	TCP UDP	137 to 139, 445	Ports used as a data transmission channel from the Veeam Agent for Linux computer to the target SMB (CIFS) share.  Ports 137 to 139 are used by backup infrastructure components to communicate using NetBIOS.
	Shared folder NFS share	TCP UDP	111, 2049	Standard NFS ports used as a data transmission channel from the Veeam Agent for Linux computer to the target NFS share.
	Veeam Agent computer	TCP	2500 to 3300	Default range of ports used for communication between Veeam Agent for Linux components during data transmission. For every TCP connection that a backup job uses, one port from this range is assigned.  Ports must be open for incoming and outgoing traffic. Established connections must be allowed.

From	To	Protocol	Port	Notes
		TCP	10808	Port used locally on the Veeam Agent computer for communication via REST API between Veeam Agent components (such as control panel and command line interface) and Veeam Agent for Linux Service.

## Communication with Veeam Backup & Replication Repositories

The following table describes network ports that must be opened to ensure proper communication with Veeam backup repositories.

From	To	Protocol	Port	Notes
Veeam Agent computer	Linux server performing the role of a backup repository	TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a backup job uses, one port from this range is assigned.
	Microsoft Windows server performing the role of a backup repository	TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic RPC port range. For more information, see <a href="#">this Microsoft article</a> .
		TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a backup job uses, one port from this range is assigned.

## Communication with Veeam Cloud Connect Repositories

The following table describes network ports that must be opened to ensure proper communication with Veeam Cloud Connect repositories.

From	To	Protocol	Port	Notes
Veeam Agent computer	Cloud gateway	TCP	6180	Port on the cloud gateway used to transport Veeam Agent data to the Veeam Cloud Connect repository.

From	To	Protocol	Port	Notes
	Certificate Revocation Lists	TCP	80 or 443 (most popular)	<p>Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider.</p> <p>Generally, information about CRL locations can be found on the CA website.</p>

## Communication with Object Storage

The following table describes network ports that must be opened to ensure proper communication with object storage if you back up data to object storage directly or to object storage added as a Veeam backup repository with the direct connection mode. For more information about object storage connection modes, see [Types of Connection to Object Storage in Veeam Backup & Replication](#).

From	To	Protocol	Port	Notes
Veeam Agent Computer	Amazon S3 object storage	TCP	443	<p>Used to communicate with the Amazon S3 object storage through the following endpoints:</p> <ul style="list-style-type: none"> <li>• *.amazonaws.com (for both <i>Global</i> and <i>Government</i> regions)</li> <li>• *.amazonaws.com.cn (for <i>China</i> region)</li> </ul> <p>All AWS service endpoints are specified in the <a href="#">AWS documentation</a>.</p>
			80	<p>Used to verify the certificate status through the following endpoints:</p> <ul style="list-style-type: none"> <li>• *.amazontrust.com</li> <li>• *.cloudfront.net</li> </ul> <p>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.</p>



From	To	Protocol	Port	Notes
	Microsoft Azure object storage	TCP	443	<p>Used to communicate with the Microsoft Azure object storage through the following endpoints:</p> <ul style="list-style-type: none"> <li>• <code>xxx.blob.core.windows.net</code> (for <i>Global</i> region)</li> <li>• <code>xxx.blob.core.chinacloudapi.cn</code> (for <i>China</i> region)</li> <li>• <code>xxx.blob.core.usgovcloudapi.net</code> (for <i>Government</i> region)</li> </ul> <p>Consider that the &lt;xxx&gt; part of the address must be replaced with your actual storage account URL that can be found in the Azure management portal.</p>
			80	<p>Used to verify the certificate status through the following endpoints:</p> <ul style="list-style-type: none"> <li>• <code>ocsp.digicert.com</code></li> <li>• <code>ocsp.msocsp.com</code></li> </ul> <p>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself. For more details, see also <a href="#">Microsoft documentation</a>.</p>
	Google Cloud storage	TCP	443	<p>Used to communicate with Google Cloud storage through the following endpoints:</p> <ul style="list-style-type: none"> <li>• <code>storage.googleapis.com</code></li> </ul> <p>All cloud endpoints are specified in <a href="#">this Google article</a>.</p>
			80	<p>Used to verify the certificate status through the following endpoints:</p> <ul style="list-style-type: none"> <li>• <code>ocsp.pki.goog</code></li> <li>• <code>pki.goog</code></li> <li>• <code>crl.pki.goog</code></li> </ul> <p>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.</p>

From	To	Protocol	Port	Notes
	IBM Cloud object storage	TCP	Depends on device configuration	Used to communicate with IBM Cloud object storage.
	S3 compatible object storage	TCP	Depends on device configuration	Used to communicate with S3 compatible object storage.
	Veeam Data Cloud Vault storage	TCP	443	Used to communicate with the Veeam Data Cloud Vault storage through the <code>xxx.blob.core.windows.net</code> endpoint.

# Live Patching Support

The live patching technology allows you to patch a running Linux kernel without the need to reboot or change kernel-related files on the disk.

If you plan to use live patching, consider the following limitations:

- Only kGraft, kpatch and Ksplice kernel extensions are supported.
- Live patching is supported only for the following Linux distributions:
  - SLES 12 SP4, 12 SP5, 15 SP1 – 15 SP5
  - SLES for SAP 12 SP4, 12 SP5, 15 SP1 – 15 SP5
  - RHEL 6.4 – 9.4
  - Oracle Linux 6.4 – 9.4 (RHCK)
  - Oracle Linux 6 (starting from UEK R2) – Oracle Linux 9 (up to 5.15.0-207.156.6.el9uek)

Before live patching, check the following prerequisites:

- Before you start live patching on a production environment, make sure that a kernel patch does not harm your system in any way using a spare Veeam Agent computer.
- Back up an entire Veeam Agent computer before live patching.
- Make sure that there are no backup jobs running on the Veeam Agent computer during live patching.

# Installation and Configuration

You can install Veeam Agent for Linux on any Linux-based endpoint whose data you plan to protect – virtual machine or physical device (server, desktop or laptop).

# Before You Begin

Before you start the installation process, review the following information and prerequisites.

## Types of Veeam Agent for Linux Installation Packages

You can install Veeam Agent using one of the available installation packages:

- Veeam Agent for Linux — this set of packages depends on the Veeam kernel module for creating system snapshots. It works with the widest range of Linux distributions and file systems. For more information on system requirements and limitations, see [System Requirements](#).
- Nosnap Veeam Agent for Linux — this set of packages does not depend on the Veeam kernel module for creating system snapshots. Nosnap Veeam Agent for Linux leverages native file system snapshot capabilities on select Linux distributions. For more information on system requirements and limitations, see [System Requirements for Nosnap Veeam Agent for Linux](#).
- Nosnap Veeam Agent for Linux on Power — this set of nosnap packages is specifically designed for IBM Power Systems. For more information on system requirements and limitations, see [System Requirements for Nosnap Veeam Agent for Linux on Power](#).

For information on installation, see [Installing Veeam Agent for Linux](#) and [Installing Veeam Agent for Linux in Offline Mode](#).

## General Prerequisites

Before you start the installation process, consider the following:

- The computer on which you plan to install Veeam Agent must satisfy system requirements specified in this document. To learn more, see [System Requirements](#).
- To install Veeam Agent software packages, you must use the `root` account or any user account that has super user (root) privileges on the computer where you plan to install the product.
- If you install Veeam Agent in a UEFI system with Secure Boot, you must configure UEFI Secure Boot to enable your system to work with Veeam Agent. You can configure UEFI Secure Boot before or after the installation of Veeam Agent, but before you run a backup job. For more information, see [Configuring UEFI Secure Boot](#).
- You must not install Veeam Agent on the server that is used as a [hardened repository](#) in the Veeam Backup & Replication infrastructure.
- If you have used the Beta version of Veeam Agent, you must remove Veeam Agent software packages prior to installing the release version of the product. To learn more, see [Uninstalling Veeam Agent for Linux](#).

# Installing Veeam Agent for Linux

To install Veeam Agent for Linux on a computer with a connection to the internet, you must perform the following steps:

1. [Connect to the Veeam software repository.](#)
2. Install Veeam Agent for Linux packages from the Veeam software repository.

Installation instructions depend on the type of the [packages](#) you want to use for Veeam Agent installation:

- [Installing Veeam Agent for Linux \(with Kernel Module\)](#)
- [Installing Nosnap Veeam Agent for Linux](#)
- [Installing Nosnap Veeam Agent for Linux on Power.](#)

## TIP

If the computer where you want to install Veeam Agent for Linux is not connected to the internet, you can download and install Veeam Agent for Linux packages manually. To learn more, see [Installing Veeam Agent for Linux in Offline Mode](#).

# Connecting to Veeam Software Repository

To install Veeam Agent for Linux on a Linux computer, you must first connect the computer to the Veeam software repository. The Veeam software repository contains the Veeam Agent installation packages specific to the Linux distribution, version and architecture of the computer where you plan to install the product.

To connect to the Veeam software repository, do the following:

1. Download the Veeam software repository installation package (`veeam-release`) from the [this Veeam webpage](#), and save the downloaded package on the computer.
2. Navigate to the directory where you have saved the `veeam-release` package and install the package using the command for your Linux distribution.

## TIP

If the user account you use for Veeam Agent installation does not have root privileges, you can temporarily elevate this user account to root by using the `sudo` prefix in the install commands. If you run multiple commands in one command line, you must use the `sudo` prefix before each command – for example, `sudo rpm -ivh ./veeam-release* && sudo yum check-update`. Make sure the `sudo` user has sufficient privileges to run these commands.

To install the `veeam-release` package, use the following commands:

*For CentOS 7 / RHEL / Oracle Linux / Fedora / Rocky Linux / AlmaLinux*

```
rpm -ivh ./veeam-release* && yum check-update
```

*For openSUSE / SLES*

```
zypper in ./veeam-release* && zypper refresh
```

*For Debian / Ubuntu*

```
dpkg -i ./veeam-release* && apt-get update
```

# Installing Veeam Agent for Linux with Kernel Module

To install Veeam Agent, you can use a package manager of your choice that works with software packages in your Linux distribution.

## NOTE

Some dependency packages of the prerequisite software may require special handling. For details, see [Managing Package Dependencies](#).

To install Veeam Agent for Linux, use the following commands:

*For CentOS 7 / RHEL / Fedora / Rocky Linux / AlmaLinux*

```
yum install veeam
```

## NOTE

[For CentOS 7 / RHEL] If the `dkms` package is already installed in the OS, you can install Veeam Agent with one of the following commands:

- `yum install veeam`

With this command, the Veeam kernel module will be installed from the source RPM package using `dkms`.

- [For CentOS 7 / RHEL 6 - 8] `yum install kmod-veeamsnap veeam` / [For RHEL 9] `yum install kmod-blksnap veeam`

With this command, the non-DKMS version of the Veeam kernel module will be installed from the pre-built `kmod` binary package.

*For Oracle Linux 6 - 8*

```
yum install veeamsnap  
yum install veeam
```

## NOTE

If your system runs on Oracle Linux 8.x with UEK R7 kernel, you may need to rebuild the Veeam kernel module prior to its installation. For more information, see [this Veeam KB article](#).

*For Oracle Linux 9*

```
yum install blksnap  
yum install veeam
```

*For openSUSE Tumbleweed*

```
zypper in veeam
```



*For openSUSE Leap 15.3 with default kernel, Leap 15.4 and 15.5*

```
zypper in blksnap-kmp-default  
zypper in veeam
```

*For openSUSE Leap 15.3 with preemptive kernel*

```
zypper in blksnap-kmp-preempt  
zypper in veeam
```

*For SLES 12 SP4 - SP5, 15 SP1 - SP2 with default kernel*

```
zypper in veeamsnap-kmp-default  
zypper in veeam
```

*For SLES 12 SP4 - SP5, 15 SP1 - SP2 with preemptive kernel*

```
zypper in veeamsnap-kmp-preempt  
zypper in veeam
```

*For SLES 15 SP3 with default kernel, 15 SP4 and SP5*

```
zypper in blksnap-kmp-default  
zypper in veeam
```

*For SLES 15 SP3 with preemptive kernel*

```
zypper in blksnap-kmp-preempt  
zypper in veeam
```

*For Debian 10 / Ubuntu 16.04, 18.04 and 20.04 (kernel 5.4)*

```
apt-get install veeam
```

*For Debian 11 - 12.2 / Ubuntu 22.04, 22.10, 23.04 and 23.10*

```
apt-get install blksnap veeam
```

# Managing Package Dependencies

The following dependency packages may require special handling in case you see installation errors:

- The `dkms` package is not present in default repositories for some Linux distributions. You should obtain it from third-party repositories:
  - EPEL repository (for CentOS / RHEL / Oracle Linux / Fedora / Rocky Linux / AlmaLinux)
  - Packman repository (for openSUSE). To learn more, see [Installing dkms in openSUSE](#).

For SLES, the `dkms` package is not available in the Packman repository. You must use the package intended for openSUSE. To learn more, see [this Veeam KB article](#).

- Extended kernels, such as `kernel-pae`, `kernel-uek` and other, require appropriate `kernel-devel` packages to be installed, for example, `kernel-pae-devel`, `kernel-uek-devel`, and so on.

Version of the `kernel-devel` package must match your current kernel version. To check your current kernel version, run the `uname -r` command.

[For RHEL and derivatives] If the `yum` package manager installs packages that do not match your current kernel version, you should either update your system or fetch older versions of the required packages from the [CentOS Vault repository](#).

## Installing dkms in openSUSE

In openSUSE systems, while installing the `dkms` package, you may see an error similar to the following:

```
Problem: nothing provides kernel-devel needed by dkms-2.2.0.3-14.1.noarch
Solution 1: do not install dkms-2.2.0.3-14.1.noarch
Solution 2: break dkms-2.2.0.3-14.1.noarch by ignoring some of its dependencies
```

To install the `dkms` package, do the following:

1. Make sure that you have an appropriate `kernel-devel` package installed and its version matches your kernel version. For example:

```
root@localhost:~> rpm -qa | grep kernel-default
kernel-default-devel-3.0.101-91.1
kernel-default-3.0.101-91.1
```

2. Install the `dkms` package ignoring dependencies:

```
zypper -n install --force dkms
```

3. Make sure that you have allowed unsupported modules. To learn more, see [SUSE documentation](#).

# Installing Nosnap Veeam Agent for Linux

To install Veeam Agent, you can use a package manager of your choice that works with software packages in your Linux distribution. For example, use the following commands:

*For CentOS 7 / RHEL / Oracle Linux / Rocky Linux / AlmaLinux*

```
yum install veeam-nosnap
```

*For openSUSE Tumbleweed / OpenSUSE Leap / SLES*

```
zypper in veeam-nosnap
```

*For Debian / Ubuntu*

```
apt-get install veeam-nosnap
```

# Installing Nosnap Veeam Agent for Linux on Power

To install Veeam Agent, you can use a package manager of your choice that works with software packages in your Linux distribution. For example, use the following commands:

*For RHEL*

```
yum install veeam-nosnap
```

*For SLES*

```
zypper in veeam-nosnap
```

# Installing Veeam Agent for Linux in Offline Mode

If the computer where you want to install Veeam Agent for Linux has no connection to the internet, for example, for security reasons, you can install Veeam Agent in the offline mode. In this scenario, you do not need to download and install the Veeam software repository installation package (`veeam-release`). Instead, you need to download all Veeam Agent packages from the Veeam software repository and install them on the target computer.

Installation instructions depend on the type of the [packages](#) you want to use for Veeam Agent installation:

- [Install Veeam Agent for Linux \(with Kernel Module\) in Offline Mode](#)
- [Install Nosnap Veeam Agent for Linux in Offline Mode](#)
- [Install Nosnap Veeam Agent for Linux on Power in Offline Mode](#)

# Installing Veeam Agent for Linux with Kernel Module in Offline Mode

To install Veeam Agent for Linux, do the following:

1. On a computer that is connected to the internet, download Veeam Agent packages intended for your Linux distribution from the [Veeam software repository](#).
  - Veeam Agent for Linux packages in the Debian format reside in the following folders of the Veeam software repository:
    - [/backup/linux/agent/dpkg/debian/public/pool/veeam/b/blksnap-dkms/](#)
    - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeamsnap/](#)
    - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeam-libs/](#)
    - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeam/](#)
  - For RPM packages, the Veeam Agent directory has the following structure: *Package format > Distribution > Version > Architecture*.  
  
For example, Veeam Agent packages for 64-bit RHEL 9 reside in the [/rpm/el/9/x86\\_64/](#) folder of the Veeam software repository, and packages for 64-bit SLES 15 SP5 reside in the [/rpm/sles/SLE\\_15\\_SP5/x86\\_64/](#) folder.
2. Save Veeam Agent packages to a directory that can be accessed from the computer where you want to install the product, for example, a directory on a local drive or USB drive, or a network shared folder.
3. On the computer where you want to install Veeam Agent, navigate to the directory where you have saved the packages and install Veeam Agent:
  - [Installing Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux](#)
  - [Installing Veeam Agent for Linux in Oracle Linux](#)
  - [Installing Veeam Agent for Linux in Fedora](#)
  - [Installing Veeam Agent for Linux in SLES](#)
  - [Installing Veeam Agent for Linux in openSUSE](#)
  - [Installing Veeam Agent for Linux in Debian / Ubuntu](#)

## TIP

You can also set up a local mirror of the Veeam software repository in your internal network and add this repository to the list of software sources on a computer where you want to install the product. These operations may differ depending on the Linux distribution and package manager that you use. To learn more, refer to the documentation of your Linux distribution.

After you add a local repository to the list of software sources on a computer, you will be able to install and upgrade Veeam Agent in a regular way. To learn more, see [Installing Veeam Agent for Linux](#) and [Upgrading Veeam Agent for Linux](#).

# Installing Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux

To install Veeam Agent for Linux, use the following commands:

*For 32-bit RHEL 6*

```
rpm -i <...>/kmod-veeamsnap-6.2.0.101-2.6.32_131.0.15.el6.i386.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.i386.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el6.i386.rpm
```

*For 64-bit RHEL 6*

```
rpm -i <...>/kmod-veeamsnap-6.2.0.101-2.6.32_131.0.15.el6.x86_64.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el6.x86_64.rpm
```

*For CentOS 7 / RHEL 7*

```
rpm -i <...>/kmod-veeamsnap-6.2.0.101-1.el7.x86_64.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el7.x86_64.rpm
```

*RHEL 8*

```
rpm -i <...>/kmod-veeamsnap-6.2.0.101-1.el8.x86_64.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el8.x86_64.rpm
```

*RHEL 9 / Rocky Linux / AlmaLinux*

```
rpm -i <...>/kmod-blksnap-6.2.0.101-1.el9.x86_64.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el9.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## NOTE

The pre-built `veeamsnap` binaries require kernel 2.6.32-131.0.15 or later for RHEL 6 (excluding 2.6.32-279.el6.i686) and kernel 3.10.0-123 or later for CentOS / RHEL 7.0 - 7.7 to operate.

# Installing Veeam Agent for Linux in Oracle Linux

To install Veeam Agent for Linux, use the following commands:

### *For 32-bit Oracle Linux 6*

```
rpm -i <...>/veeamsnap-6.2.0.101-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.i386.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el6.i386.rpm
```

### *For 64-bit Oracle Linux 6*

```
rpm -i <...>/veeamsnap-6.2.0.101-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el6.x86_64.rpm
```

### *For Oracle Linux 7*

```
rpm -i <...>/veeamsnap-6.2.0.101-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el7.x86_64.rpm
```

### *For Oracle Linux 8*

```
rpm -i <...>/veeamsnap-6.2.0.101-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el8.x86_64.rpm
```

### *For Oracle Linux 9*

```
rpm -i <...>/blksnap-6.2.0.101-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.el9.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Veeam Agent for Linux in Fedora

To install Veeam Agent for Linux, use the following commands:

```
rpm -i <...>/blksnap-6.2.0.101-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-6.2.0.101-1.fc34.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.



# Installing Veeam Agent for Linux in SLES

To install Veeam Agent for Linux, use the following commands:

*For SLES 12 SP4*

```
zypper in <...>/veeamsnap-kmp-default-6.2.0.101_k4.12.14_94.41-sles12.4.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle12.x86_64.rpm
```

*For SLES 12 SP5*

```
zypper in <...>/veeamsnap-kmp-default-6.2.0.101_k4.12.14_120-sles12.5.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle12.x86_64.rpm
```

*For SLES 15 SP1*

```
zypper in <...>/veeamsnap-kmp-default-6.2.0.101_k4.12.14_195-sles15.1.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For SLES 15 SP2 with default kernel*

```
zypper in <...>/veeamsnap-kmp-default-6.2.0.101_k5.3.18_22-sles15.2.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For SLES 15 SP2 with preemptive kernel*

```
zypper in <...>/veeamsnap-kmp-preempt-6.2.0.101_k5.3.18_22-sles15.2.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For SLES 15 SP3 with default kernel*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k5.3.18_57-sles15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For SLES 15 SP3 with preemptive kernel*

```
zypper in <...>/blksnap-kmp-preempt-6.2.0.101_k5.3.18_57-sles15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For SLES 15 SP4*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k5.14.21_150400.22-sles15.4.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For SLES 15 SP5*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k5.14.21_150500.53-sles15.5.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For SLES 15 SP6*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k6.4.0_150600.21-sles15.6.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle16.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Veeam Agent for Linux in openSUSE

To install Veeam Agent for Linux, use the following commands:

*For openSUSE Tumbleweed*

```
zypper in <...>/blksnap-6.2.0.101-1.sle.noarch.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.suse.x86_64.rpm
```

*For openSUSE Leap 15.3 with default kernel*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k5.3.18_59.10-opensuse_leap15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For openSUSE Leap 15.3 with preemptive kernel*

```
zypper in <...>/blksnap-kmp-preempt-6.2.0.101_k5.3.18_59.10-opensuse_leap15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For openSUSE Leap 15.4*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k5.14.21_150400.22-opensuse_leap15.4.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For openSUSE Leap 15.5*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k5.14.21_150500.53-opensuse_leap15.5.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

*For openSUSE Leap 15.6*

```
zypper in <...>/blksnap-kmp-default-6.2.0.101_k6.4.0_150600.21-opensuse_leap15.6.x86_64.rpm
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm
zypper in <...>/veeam-6.2.0.101-1.sle15.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Veeam Agent for Linux in Debian / Ubuntu

To install Veeam Agent for Linux, use the following commands:

*For Debian 10 / Ubuntu 16.04, 18.04, 20.04 (kernel 5.4)*

```
apt-get install <...>/veeamsnap_6.2.0.101_all.deb
apt-get install <...>/veeam-libs_6.2.0.101_amd64.deb
apt-get install <...>/veeam_6.2.0.101_amd64.deb
```

*For Debian 11 – 12.6 / Ubuntu 22.04, 22.10, 23.04, 23.10 and 24.04*

```
apt-get install <...>/blksnap_6.2.0.101_all.deb
apt-get install <...>/veeam-libs_6.2.0.101_amd64.deb
apt-get install <...>/veeam_6.2.0.101_amd64.deb
```

where:

< . . . > – path to a directory where you have saved Veeam Agent packages.

# Installing Nosnap Veeam Agent for Linux in Offline Mode

To install nosnap Veeam Agent for Linux, do the following:

1. On a computer that is connected to the internet, download Veeam Agent packages intended for your Linux distribution from the [Veeam software repository](#).
  - For RPM packages of nosnap Veeam Agent for Linux, the Veeam Agent directory has the following structure: *Package format > Distribution > Version > Architecture*.  
  
For example, Veeam Agent packages for 64-bit RHEL 9 reside in the [/rpm/el/9/x86\\_64/](#) folder of the Veeam software repository, and packages for 64-bit SLES 15 SP5 reside in the [/rpm/sles/SLE\\_15\\_SP5/x86\\_64/](#) folder.
  - Nosnap Veeam Agent for Linux packages in the Debian format reside in the following folders of the Veeam software repository:
    - [/backup/linux/agent/dpkg/debian/public//pool/veeam/v/veeam-nosnap/](#)
    - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeam-lib/](#)
2. Save the `veeam-nosnap` and `veeam-lib` packages to a directory that can be accessed from the computer where you want to install the product, for example, a directory on a local drive or USB drive, or a network shared folder.
3. On the computer where you want to install Veeam Agent, navigate to the directory where you have saved the packages and install Veeam Agent:
  - [Installing nosnap Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux](#)
  - [Installing nosnap Veeam Agent for Linux in Oracle Linux](#)
  - [Installing nosnap Veeam Agent for Linux in SLES](#)
  - [Installing nosnap Veeam Agent for Linux in openSUSE](#)
  - [Installing nosnap Veeam Agent for Linux in Debian / Ubuntu](#)

## TIP

You can also set up a local mirror of the Veeam software repository in your internal network and add this repository to the list of software sources on a computer where you want to install the product. These operations may differ depending on the Linux distribution and package manager that you use. To learn more, refer to the documentation of your Linux distribution.

After you add a local repository to the list of software sources on a computer, you will be able to install and upgrade Veeam Agent in a regular way. To learn more, see [Installing Veeam Agent for Linux](#) and [Upgrading Veeam Agent for Linux](#).

## Installing Nosnap Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux

To install nosnap Veeam Agent for Linux, use the following commands:

#### *For 32-bit RHEL 6*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.i386.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el6.i386.rpm
```

#### *For 64-bit RHEL 6*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el6.x86_64.rpm
```

#### *For CentOS 7 / RHEL 7*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el7.x86_64.rpm
```

#### *RHEL 8*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el8.x86_64.rpm
```

#### *RHEL 9 / Rocky Linux / AlmaLinux*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el9.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Nosnap Veeam Agent for Linux in Oracle Linux

To install nosnap Veeam Agent for Linux, use the following commands:

#### *For 32-bit Oracle Linux 6*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.i386.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el6.i386.rpm
```

#### *For 64-bit Oracle Linux 6*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el6.x86_64.rpm
```

#### *For Oracle Linux 7*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el7.x86_64.rpm
```

#### *For Oracle Linux 8*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el8.x86_64.rpm
```

#### *For Oracle Linux 9*

```
rpm -i <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el9.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Nosnap Veeam Agent for Linux in SLES

To install nosnap Veeam Agent for Linux, use the following commands:

#### *For SLES 12 SP4 – SP5*

```
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.2.0.101-1.sle12.x86_64.rpm
```

#### *For SLES 15 SP1 – SP6*

```
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.2.0.101-1.sle15.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Nosnap Veeam Agent for Linux in openSUSE

To install nosnap Veeam Agent for Linux, use the following commands:

#### *For openSUSE Tumbleweed*

```
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.2.0.101-1.suse.x86_64.rpm
```

*For openSUSE Leap 15.3 - 15.6*

```
zypper in <...>/veeam-libs-6.2.0.101-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.2.0.101-1.sle15.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Nosnap Veeam Agent for Linux in Debian / Ubuntu

To install nosnap Veeam Agent for Linux, use the following commands:

```
apt-get install <...>/veeam-libs_6.2.0.101_amd64.deb  
apt-get install <...>/veeam-nosnap_6.2.0.101_amd64.deb
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.



# Installing Nosnap Veeam Agent for Linux on Power in Offline Mode

To install nosnap Veeam Agent for Linux on Power, do the following:

1. On a computer that is connected to the internet, download Veeam Agent packages intended for your Linux distribution from the [Veeam software repository](#).

For RPM packages of nosnap Veeam Agent for Linux on Power, the Veeam Agent directory has the following structure: *Package format* > *Distribution* > *Version* > *Architecture*.

For example, Veeam Agent packages for RHEL 8 reside in the [/rpm/el/8/ppc64le/](#) folder of the Veeam software repository, and packages for SLES 15 SP4 reside in the [/rpm/sles/SLE\\_15\\_SP4/ppc64le/](#) folder.

2. Save the `veeam-nosnap` and `veeam-libs` packages to a directory that can be accessed from the computer where you want to install the product, for example, a directory on a local drive or USB drive, or a network shared folder.
3. On the computer where you want to install Veeam Agent, navigate to the directory where you have saved the packages and install Veeam Agent:
  - [Installing nosnap Veeam Agent for Linux on Power in RHEL](#)
  - [Installing nosnap Veeam Agent for Linux on Power in SLES](#)

## TIP

You can also set up a local mirror of the Veeam software repository in your internal network and add this repository to the list of software sources on a computer where you want to install the product. These operations may differ depending on the Linux distribution and package manager that you use. To learn more, refer to the documentation of your Linux distribution.

After you add a local repository to the list of software sources on a computer, you will be able to install and upgrade Veeam Agent in a regular way. To learn more, see [Installing Veeam Agent for Linux](#) and [Upgrading Veeam Agent for Linux](#).

## Installing Nosnap Veeam Agent for Linux on Power in RHEL

To install nosnap Veeam Agent for Linux on Power, use the following commands:

```
rpm -i <...>/veeam-libs-6.2.0.101-1.ppc64le.rpm
rpm -i <...>/veeam-nosnap-6.2.0.101-1.el8.ppc64le.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

## Installing Nosnap Veeam Agent for Linux on Power in SLES

To install nosnap Veeam Agent for Linux on Power, use the following commands:

*For SLES for SAP 12 SP5*

```
zypper in <...>/veeam-libs-6.2.0.101-1.ppc64le.rpm  
zypper in <...>/veeam-nosnap-6.2.0.101-1.sle12.ppc64le.rpm
```

*For SLES / SLES for SAP 15 SP3 - SP4*

```
zypper in <...>/veeam-libs-6.2.0.101-1.ppc64le.rpm  
zypper in <...>/veeam-nosnap-6.2.0.101-1.sle15.ppc64le.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

# Configuring UEFI Secure Boot

When you install Veeam Agent on a UEFI system with Secure Boot enabled, you must configure the UEFI Secure Boot to allow your system to run Veeam Agent and perform backups. You do this by enrolling a Machine Owner Key (MOK) for the Veeam kernel module in your system's firmware. To enroll MOK, perform the following steps:

1. Request enrollment of the key. Depending on the kernel module type – [pre-built](#) or [DKMS](#), the key is either provided by Veeam or generated by DKMS:
  - [Pre-built kernel module] To make UEFI system with Secure Boot work with pre-built Veeam kernel module, Veeam Agent requires Veeam public key to be enrolled to the system's MOK list. For more information on requesting enrollment of the Veeam kernel module key to your system, see [Importing MOK for Pre-Built Kernel Module](#).
  - [DKMS kernel module] If you install Veeam Agent in Ubuntu 22.04 and later or Debian 12.0 and later, DKMS generates a Machine Owner Key that allows third-party modules to be run on the system's firmware. Such key must also be enrolled to the system's MOK list. For more information on requesting enrollment of the key for the Veeam DKMS module, see [Importing MOK for Veeam DKMS Module](#).

## NOTE

If UEFI system with Secure Boot enabled does not support automatic generation of the key for DKMS modules, you must either sign the Veeam kernel module yourself and enroll the Machine Owner Key to your system or disable Secure Boot.

2. Enroll the key using MOK management. For more information, see [Enrolling MOK](#).

## Importing MOK for Pre-Built Kernel Module

The Veeam kernel module key is provided within the `ueficert` package that resides in the [Veeam software repository](#). Depending on the Linux distribution version, the full name of the package can be `veeamsnap-ueficert-6.2.0.101-1.noarch` or `blksnap-ueficert-6.2.0.101-1.noarch`.

Install the package that contains the public key for pre-built Veeam kernel module by using the following command:

```
rpm -i <...>/veeamsnap-ueficert-6.2.0.101-1.noarch.rpm
```

or

```
rpm -i <...>/blksnap-ueficert-6.2.0.101-1.noarch.rpm
```

After you install the `ueficert` package, the key is automatically imported into the enrollment request. You can now [confirm the key enrollment](#).

### TIP

After the package is installed, you can verify that the key enrollment is planned for the next reboot using the following command: `mokutil -N`. If the command output shows that the key enrollment is not planned, request the enrollment of the public key manually with the following command: `mokutil --import veeamsnap-ueficert.crt` or `mokutil --import blksnap-ueficert.crt`.

By default, the key is stored in the `/etc/uefi/certs` directory.

## Importing MOK for DKMS Kernel Module

Veeam does not provide a `ueficert` package for the DKMS module because it is not possible to sign such module automatically. Depending on the Linux distribution and version, you may have several options to make your system load the Veeam DKMS module properly – for more information, see [Linux documentation](#).

If your system runs on Ubuntu 22.04 and later or Debian 12.0 and later, after you install Veeam kernel module using DKMS, a new Machine Owner Key is generated. Depending on the Linux distribution, perform the following steps to request enrollment of the key to your system's firmware:

- [Debian 12.0 and later] By default, the key is stored in the `/var/lib/dkms/` directory. To import the key, run the following command:

```
mokutil --import /var/lib/dkms/mok.pub
```

- [Ubuntu 22.04 and later] After you install the Veeam kernel module, the key is generated and imported into your system automatically. By default, the key is stored in the `/var/lib/shim-signed/mok` directory.

When the key is imported into the enrollment request, you will be prompted to enter a password that you will use to confirm the enrollment of the key during MOK management. After you set the password, you can [confirm the key enrollment](#).

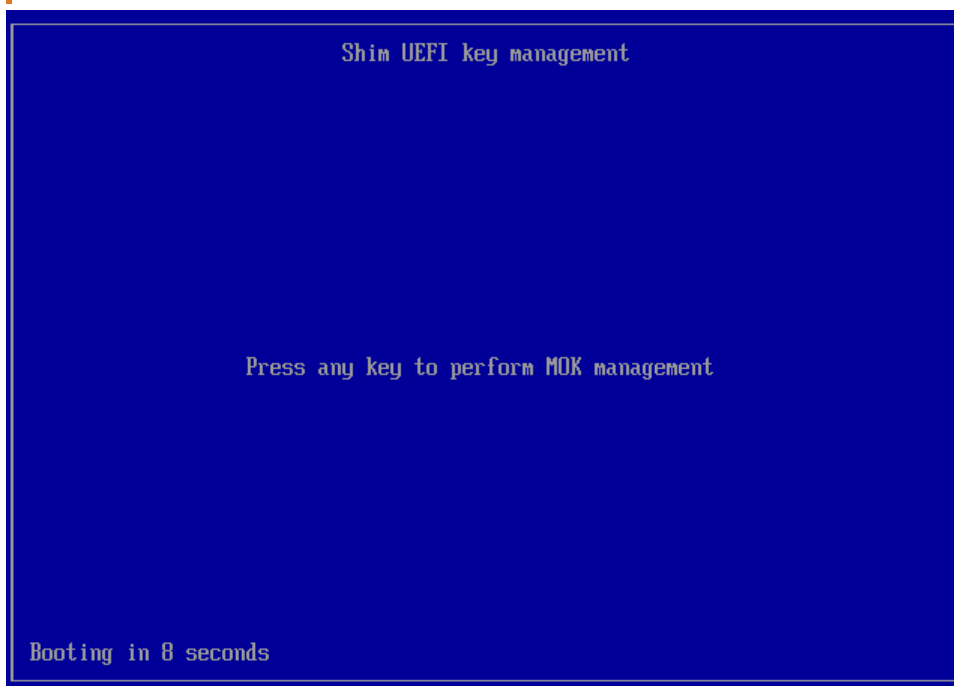
## Enrolling MOK

To enroll the Veeam or DKMS-generated key to the MOK list, do the following:

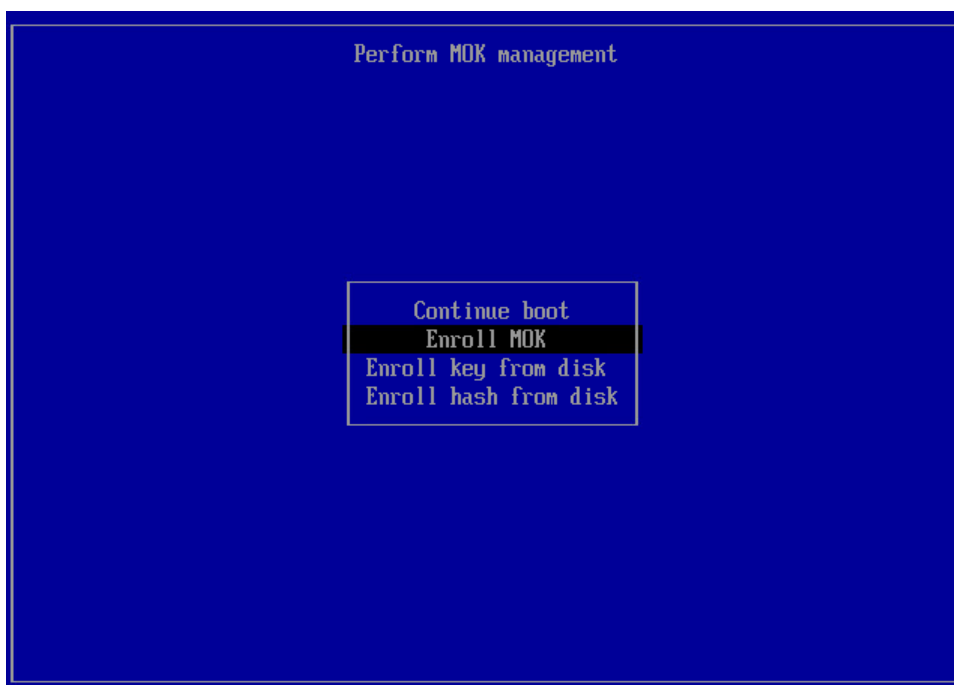
1. Reboot the computer.
2. During reboot, when prompted, press any key to perform MOK management.

## IMPORTANT


The prompt will time out in 10 seconds. If you don't press any key, the system will continue booting without enrolling the key. If you don't enroll the key at reboot, you will have to reconfigure the key by reinstalling the `ueficer` package and reboot again.



- At the first step of the wizard, select **Enroll MOK** and press [Enter].



4. At the **Enroll MOK** step, select **Continue** and press [Enter].

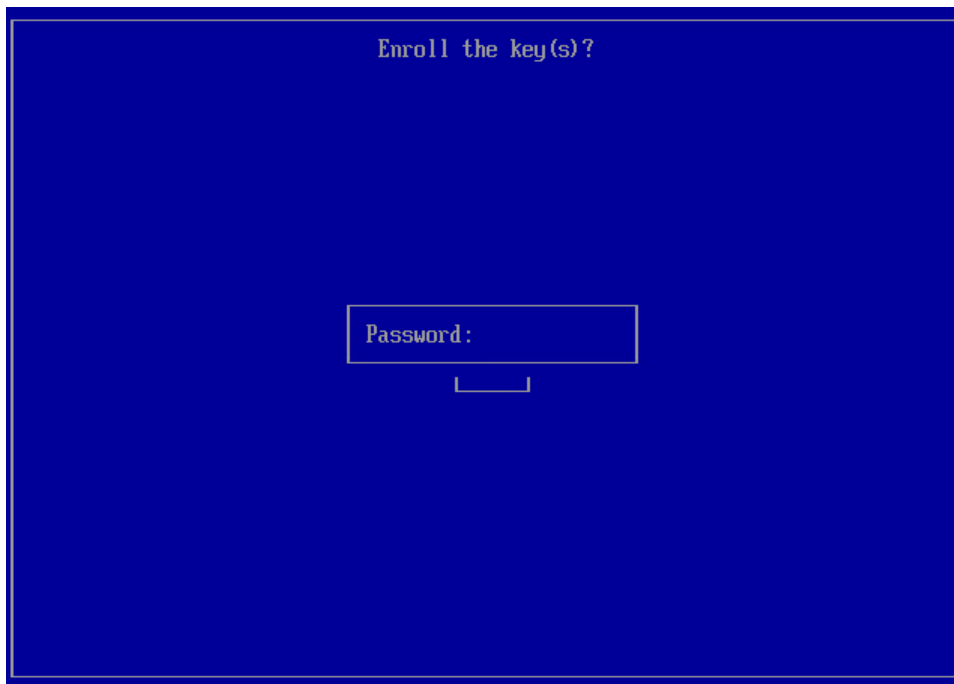


5. At the **Enroll the key(s)** step, select **Yes** and press [Enter].

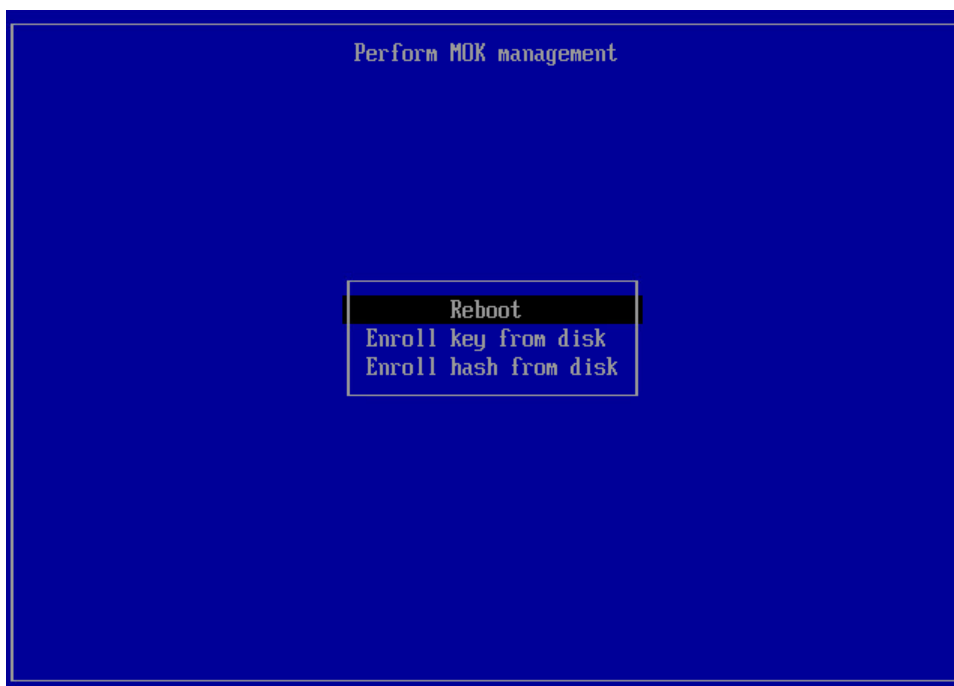


6. Depending on the type of key you enroll – Veeam or DKMS-generated, do the following:
  - [For Veeam public key for pre-built kernel module] Provide the password for the root account and press [Enter].

- [For DKMS-generated key for Veeam kernel module] Provide the password you set when you imported the key and press [Enter].



7. At the final step, select **Reboot** and press [Enter].



8. After the system reboots, verify that the key is successfully enrolled with the following command:  
`mokutil -l`. The system will list the enrolled keys.

# Upgrading Veeam Agent for Linux

For Veeam Agent for Linux, upgrade to newer versions is supported. You can start the upgrade process when the new version becomes available.

During the upgrade process, configuration and backup files that were created with the previous version of Veeam Agent are not impacted in any way.

## IMPORTANT

Before starting the upgrade process, make sure that there are no jobs running on the Veeam Agent computer.

Depending on the type of the [packages](#) you used for Veeam Agent installation, you can use the following upgrade procedures:

- [Upgrading Veeam Agent for Linux with kernel module](#)
- [Upgrading nosnap Veeam Agent for Linux](#)
- [Upgrading nosnap Veeam Agent for Linux on Power.](#)

## TIP

If the computer where you want to upgrade Veeam Agent for Linux is not connected to the internet and does not have access to a local mirror of the [Veeam software repository](#), you can download and re-install Veeam Agent for Linux packages manually. To learn more, see [Installing Veeam Agent for Linux in Offline Mode](#).



# Upgrading Veeam Agent for Linux with Kernel Module

The commands for the upgrade of Veeam Agent for Linux differ depending on the Linux distribution:

- [Upgrading Veeam Agent for Linux in CentOS 7 / RHEL 6 – 8](#)
- [Upgrading Veeam Agent for Linux in RHEL 9 / Rocky Linux / AlmaLinux](#)
- [Upgrading Veeam Agent for Linux in Oracle Linux 6 – 8](#)
- [Upgrading Veeam Agent for Linux in Fedora / Oracle Linux 9](#)
- [Upgrading Veeam Agent for Linux in openSUSE](#)
- [Upgrading Veeam Agent for Linux in SLES 12 SP4 – SP5, 15 SP1 – SP2](#)
- [Upgrading Veeam Agent for Linux in SLES 15 SP3 – SP5](#)
- [Upgrading Veeam Agent for Linux in Debian 10 / Ubuntu 16.04, 18.04, 20.04 \(kernel 5.4\)](#)
- [Upgrading Veeam Agent for Linux in Debian 11 – 12.0 / Ubuntu 22.04, 22.10 and 23.04](#)

## Upgrading Veeam Agent for Linux in CentOS 7 / RHEL 6 – 8

To upgrade Veeam Agent for Linux, use the following command:

```
yum update veeam
```

With these commands, a pre-built binary package with Veeam kernel module will be installed in your system. To stay on the DKMS version of the Veeam kernel module, use the following command for upgrade:

```
yum update veeamsnap && yum update veeam
```

## Upgrading Veeam Agent for Linux in RHEL 9 / Rocky Linux / AlmaLinux

To upgrade Veeam Agent for Linux, use the following command:

```
yum install kmod-blksnap veeam --allowerase
```

With this command, a pre-built binary package with the Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
yum install blksnap veeam --allowerase
```

## Upgrading Veeam Agent for Linux in Oracle Linux 6 – 8

To upgrade Veeam Agent for Linux, use the following command:

```
yum update veeam
```

## Upgrading Veeam Agent for Linux in Fedora / Oracle Linux 9

To upgrade Veeam Agent for Linux, use the following command:

```
yum update veeam --allowerase
```

## Upgrading Veeam Agent for Linux in openSUSE

To upgrade Veeam Agent for Linux, use the following commands:

*For openSUSE Tumbleweed*

```
zypper update veeam
```

*For openSUSE Leap 15.3 with default kernel*

```
zypper in -- replacefiles blksnap-kmp-default veeam
```

*For openSUSE Leap 15.3 with preemptive kernel*

```
zypper in -- replacefiles blksnap-kmp-preempt veeam
```

*For openSUSE Leap 15.4 and 15.5*

```
zypper in -- replacefiles blksnap-kmp-default veeam
```

With these commands, a pre-built binary package with the Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper update veeam
```

# Upgrading Veeam Agent for Linux in SLES 12 SP4 – SP5, 15 SP1 – SP2

To upgrade Veeam Agent for Linux, use the following commands:

*For default kernel*

```
zypper in veeamsnap-kmp-default veeam
```

*For preemptive kernel*

```
zypper in veeamsnap-kmp-preempt veeam
```

With these commands, a pre-built binary package with Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper update veeam
```

## Upgrading Veeam Agent for Linux in SLES 15 SP3 – SP5

To upgrade Veeam Agent for Linux, use the following commands:

*For SLES 15 SP3 with default kernel*

```
zypper in --replacefiles blksnap-kmp-default veeam
```

*For SLES 15 SP3 with preemptive kernel*

```
zypper in --replacefiles blksnap-kmp-preempt veeam
```

*For SLES 15 SP4 and SP5*

```
zypper in --replacefiles blksnap-kmp-default veeam
```

With these commands, a pre-built binary package with the Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper in --replacefiles blksnap veeam
```

## Upgrading Veeam Agent for Linux in Debian 10 / Ubuntu 16.04, 18.04, 20.04 (kernel 5.4)

To upgrade Veeam Agent for Linux, use the following commands:

```
apt-get update  
apt-get install veeam
```

## Upgrading Veeam Agent for Linux in Debian 11 – 12.0 / Ubuntu 22.04, 22.10 and 23.04

To upgrade Veeam Agent for Linux, use the following commands:

```
apt-get update  
apt-get install blksnap veeam
```

# Upgrading Nosnap Veeam Agent for Linux

The commands for the upgrade of nosnap Veeam Agent for Linux differ depending on the Linux distribution:

*For CentOS 7 / RHEL / Oracle Linux / Rocky Linux / AlmaLinux*

```
yum update veeam-nosnap
```

*For openSUSE Tumbleweed / OpenSUSE Leap / SLES*

```
zypper in veeam-nosnap
```

*For Debian / Ubuntu*

```
apt-get update  
apt-get install veeam-nosnap
```

# Upgrading Nosnap Veeam Agent for Linux on Power

The commands for the upgrade of nosnap Veeam Agent for Linux on Power differ depending on the Linux distribution:

*For RHEL*

```
yum update veeam-nosnap
```

*For SLES*

```
zypper in veeam-nosnap
```

# Granting Permissions to Users

When you install Veeam Agent for Linux, the product program files are placed to the folders on the system volume. For full access to Veeam Agent files, super user (root) privileges are required. Rights to execute product files and run commands are also granted to users that belong to the `veeam` group.

The `veeam` group is automatically created by Veeam Agent at the process of the product installation. To let regular users work with Veeam Agent without the need to gain root privileges, you can add the necessary users to this group. Users in the `veeam` group will be able to execute Veeam Agent commands and perform backup and restore tasks under regular user account.

To add a user to the `veeam` group, in most of Linux distributions you can use the following command:

```
usermod -a -G veeam <username>
```

where:

`<username>` — name of the account to which you want to grant access to Veeam Agent.

For example:

```
root@srv01:~# usermod -a -G veeam user
```

## IMPORTANT

Consider the following:

- To add a user to the `veeam` group, you must have super user (root) privileges in the Linux OS.
- After the user is added to the `veeam` group, the user must re-login to the Linux OS.
- Add only trusted users to the `veeam` group. Veeam Agent for Linux daemon runs and executes commands and scripts with the super user privileges. Thus, users who belong to this group can potentially escalate their privileges through the creative use of pre-freeze/post-thaw or pre-job/post-job scripts.

To check whether the user who is currently logged in to the Linux OS is added to the `veeam` group, you can use the following command:

```
groups
```

For example:

```
user@srv01:~$ groups
user adm cdrom sudo dip plugdev lpadmin sambashare veeam
```

# Performing Initial Setup

After you install Veeam Agent for Linux, you can use the [Veeam Agent for Linux control panel](#) to perform initial product setup. When you launch the control panel for the first time, Veeam Agent displays the initial setup wizard. The wizard offers you to accept license agreements, install a license and create a custom Veeam Recovery Media that will include drivers of your Veeam Agent computer.

To perform initial setup, launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command. Then use the initial setup wizard to complete the following steps:

1. [Accept Veeam and third-party license agreements.](#)
2. [Create a custom Veeam Recovery Media.](#)
3. [Install a license.](#)

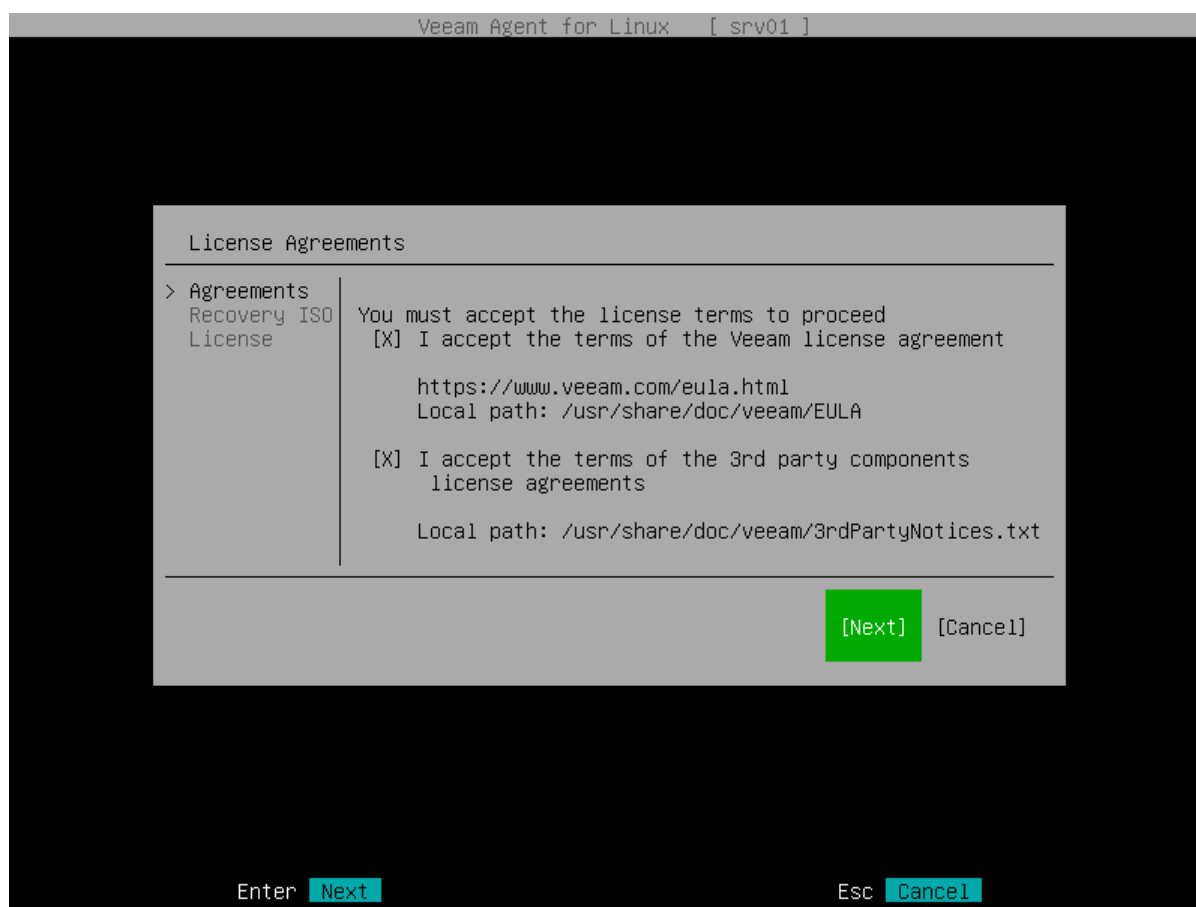


# Step 1. Accept License Agreements

At the **Agreements** step of the initial setup wizard, accept the terms of the product license agreement and license agreements for third-party components of the product. You must accept the license agreements to start using the product. Until you accept the license agreements, you will not be able to perform backup and data recovery tasks with the Veeam Agent for Linux control panel and command line interface.

To accept the license agreements:

1. Make sure that the **I accept the terms of the Veeam license agreement** option is selected and press [Space].
2. Select the **I accept the terms of the 3rd party components license agreements** option with the [Down] and [Up] key and press [Space].
3. Press [Enter].



## Step 2. Create Custom Veeam Recovery Media

At the **Recovery ISO** step of the initial setup wizard, specify settings for the custom Veeam Recovery Media.

In addition to the generic Veeam Recovery Media that is available for download at the Veeam website, you can create a custom Veeam Recovery Media. This option may be helpful if your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media. When you create the custom Veeam Recovery Media, Veeam Agent for Linux copies the Linux kernel running on your computer with its currently loaded modules and includes them into the custom recovery media.

### IMPORTANT

Recovery Media patching is not supported by Veeam Agent for Linux on Power. Switch to the **Next** button with the [Tab] key and press [Enter]. You will proceed immediately to the [License](#) step of the initial setup wizard.

Before you create a custom Veeam Recovery Media, check the following prerequisites:

- The Linux system must have the `genisoimage` package installed. For openSUSE and SLES 15 SP1 – 15 SP5 distributions, the Linux system must have the `mkisofs` package installed instead.
- The Linux system must have the `mksquashfs` and `unsquashfs` utilities installed.
- For custom Veeam Recovery Media with EFI support, the Linux system must have the following packages installed:
  - `xorriso`
  - `isolinux` (or `syslinux`, if the software package repository of your Linux distribution lacks the `isolinux` package).
- For the scenario where you create a custom Veeam Recovery Media using the **Download and patch ISO** option, the Linux system must have the `wget` utility installed.

### TIP

If you do not want to create a custom Veeam Recovery Media at the process of initial product setup, switch to the **Next** button with the [Tab] key and press [Enter]. You will proceed immediately to the [License](#) step of the initial setup wizard.

You can create the custom Veeam Recovery Media later, at any time you need, using the Veeam Agent for Linux command line interface. To learn more, see [Creating Custom Veeam Recovery Media](#).

To specify settings for the custom Veeam Recovery Media:

1. Make sure that the **Patch Veeam Recovery Media ISO** option is selected and press [Space].
2. If you want the Veeam Recovery Media to be able to boot on EFI-based systems, select the **EFI system** option with the [Tab] key and press [Space].

If you do not enable this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

4. Press [Tab] and select how you want to create a custom Veeam Recovery Media depending on the location of the generic recovery media ISO file:

- If you have not downloaded the generic Veeam Recovery Media, make sure that the **Download and patch ISO** option is selected and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer and use this image to create the custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you want only to download the generic Veeam Recovery Media, select the **Only download ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer. You can use the downloaded ISO file later to boot your Veeam Agent computer or to create a custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you have already downloaded the generic Veeam Recovery Media to a local directory on the Veeam Agent computer or to a network shared folder, select the **Patch local ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will use the generic Veeam Recovery Media ISO file to create the custom Veeam Recovery Media.

The name of the generic Veeam Recovery Media ISO file depends on the recovery image version, Veeam Agent computer architecture and the source from which you downloaded the ISO file: from the product download page or Veeam software repository. To learn more, see [Veeam Recovery Media Versions](#).

5. If you selected the **Download and patch ISO** or **Patch local ISO** option, the **EFI system** option is available. If you want to boot the Veeam Recovery Media on EFI-based systems, select the **EFI system** option with the [Tab] key and press [Space].

If you do not enable this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

6. If you selected the **Patch local ISO** option, in the **Path to local ISO** field, specify a path to the ISO file of the generic Veeam Recovery Media:

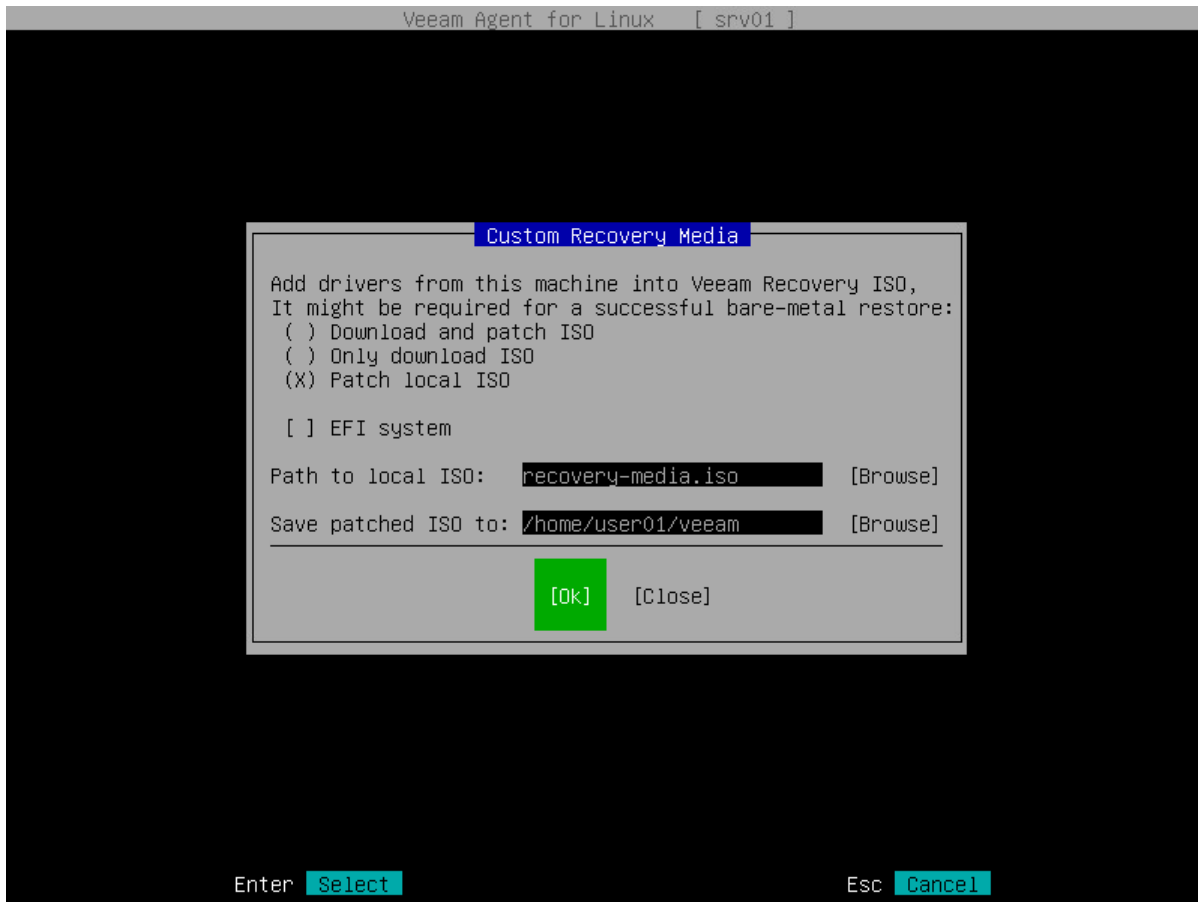
- a. Select the **Browse** option with the [Tab] key and press [Enter].
- b. In the **Path to ISO** window, select the necessary directory and press [Enter].
- c. Repeat the step 'b' until a path to the directory in which the recovery media ISO file resides appears in the **Current directory** field.
- d. In the directory where the recovery media ISO file resides, select the ISO file and press [Enter].

7. Specify a path to the resulting ISO file of the Veeam Recovery Media.

If you selected the **Download and patch ISO** or **Patch local ISO** option, in the **Save patched ISO to** field, you can specify a path to the resulting ISO file of the custom Veeam Recovery Media; if you selected the **Only download ISO** option, in the **Save ISO to** field, specify a path to the resulting ISO file of the generic Veeam Recovery Media:

- a. Select the **Browse** option with the [Tab] key and press [Enter].
- b. In the **Save patched ISO to** window, select the necessary directory and press [Enter].
- c. Repeat the step 'b' until a path to the directory where you want to save the resulting custom recovery media ISO file appears in the **Current directory** field.

- d. Select the **OK** button with the [Tab] key and press [Enter].
8. To start the custom recovery media creation process, select the **Next** button with the [Tab] key and press [Enter].



## Step 3. Install Product License

At the **License** step of the initial setup wizard, install the license. You can choose to install the license immediately or postpone this operation.

- If you choose to install the license, you can immediately browse for the license key on your computer and complete the license installation process.
- If you choose to postpone the license installation process, you will be able to install a license later at any time you need.

Until you install a license, Veeam Agent for Linux will operate in the Free edition. To learn more, see [Product Editions](#).

### NOTE

If you choose not to install a license and use Veeam Agent in the Free edition, Veeam Agent will display a notification offering to install a license every time you open the control panel. The notification will appear in the control panel until Veeam Agent completes the first backup job session.

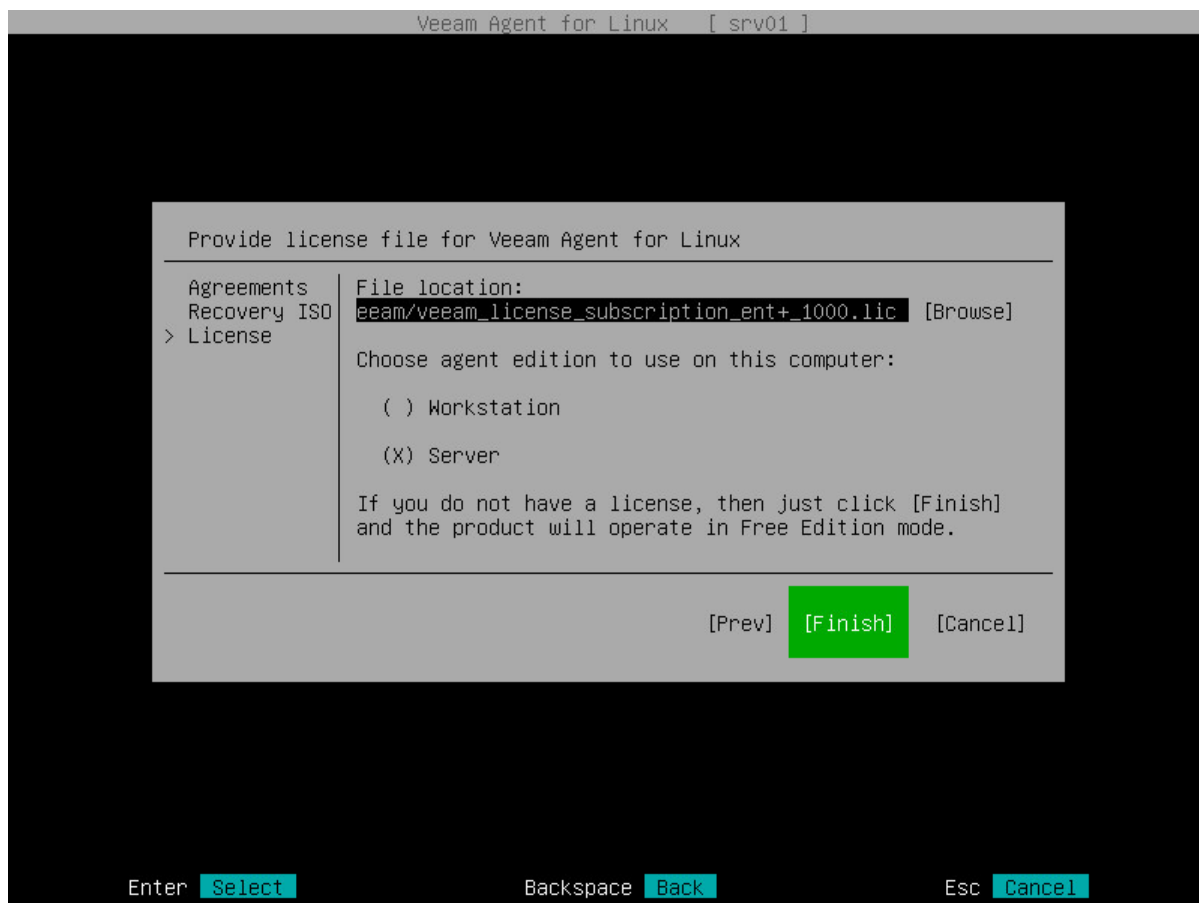
To install a license:

1. In the **File location** field, specify a path to the license key:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Choose license file location** window, select the necessary directory and press [Enter].
  - c. Repeat the step 'a' until a path to the directory in which the license key resides appears in the **Current directory** field.
  - d. In the directory where the license key resides, select the license key and press [Enter].
2. In the **Choose agent edition to use on this computer** section, select the product edition in which Veeam Agent will operate and press [Enter] to install the license and finish working with the initial setup wizard.

### TIP

Consider the following:

- If you do not want to install a license, to finish working with the initial setup wizard, switch to the **Finish** button with the [Tab] key and press [Enter].
- You can view information about the installed license (expiration date, status of the license, current edition of the product and so on) in the Veeam Agent control panel or using the Veeam Agent command line interface. To learn more, see [Viewing License](#).



# Configuring Advanced Settings

Veeam Agent for Linux allows you to configure the following settings:

- [HTTP proxy settings for Veeam Cloud Connect repository](#)
- [Connection settings for Veeam backup server](#)

## HTTP Proxy Settings for Veeam Cloud Connect Repository

If you want to use Veeam Agent for Linux to back up your data to a Veeam Cloud Connect repository, it might be required that you specify HTTP proxy settings for Veeam Agent.

Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider. In case it is not possible to establish a direct connection to CRLs, you must configure an HTTP proxy and specify settings to connect to the proxy in Veeam Agent.

To specify settings for an HTTP proxy, uncomment and edit the following lines in the `[cloudconnect]` section of the `/etc/veeam/veeam.ini` configuration file:

```
[cloudconnect]
...
# httpproxylogin= <username>
...
# httpproxypasswd= <password>
...
# httpproxyurl= <URL>
```

where:

- `<username>` – name of the account used to connect to the HTTP proxy.
- `<password>` – password of the account used to connect to the HTTP proxy.
- `<URL>` – URL of a proxy used for CRL checks.

### NOTE

If the proxy does not require authentication, you do not need to specify the account name and password. Keep in mind that only the basic authentication method is supported for connection to a proxy.

For example:

```
[cloudconnect]
...
# HTTP proxy login
httpproxylogin= user01
# HTTP proxy password
httpproxypasswd= P@ssw0rd
# HTTP proxy URL for CRL checks
httpproxyurl= http://proxy.company.lan:3128
```

# Connection Settings for Veeam Backup Server

If you want to connect Veeam Agent computer to Veeam backup server as a member of the protection group for pre-installed Veeam Agents, you must apply connection settings from the configuration file. The configuration file is one of the Veeam Agent for Linux setup files that you must obtain from your System Administrator. To learn more about protection group for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

To connect Veeam Agent for Linux to Veeam backup server:

1. Get the configuration file from your System Administrator and upload this file on the Veeam Agent computer.
2. Navigate to the directory where you have saved the configuration file and run the following command:

```
veeamconfig mode setvbrsettings --cfg <file_name>.xml
```

where <file\_name> is a configuration file name.

Alternatively, you can specify the full path to the configuration file with the `--cfg` option.

For example:

```
user@srv01:~# veeamconfig mode setvbrsettings --cfg /home/Linux\ Servers\ Distrib  
ibs/Linux/LinuxServers.xml
```

Mind that the connection between Veeam backup server and Veeam Agent computer added as a member of the protection group for pre-installed Veeam Agents is not persistent. Veeam Agent synchronizes with Veeam Backup & Replication every 6 hours. To synchronize Veeam Agent immediately, run the following command from the Veeam Agent computer:

```
veeamconfig mode syncnow
```



# Managing Veeam Agent Operation Mode

Veeam Agent for Linux can operate in different modes. Depending on the selected mode, Veeam Agent offers different features and limitations. To learn more, see [Standalone and Managed Operation Modes](#).

Veeam Agent allows you to perform the following actions to manage the operation mode:

- [View operation mode details](#)
- [Reset to the standalone operation mode](#)
- [Connect to Veeam backup server](#)
- [Synchronize with Veeam backup server](#)
- [Export logs to Veeam backup server](#)

# Viewing Operation Mode

To view the current Veeam Agent operation mode, use the following command:

```
veeamconfig mode info
```

Veeam Agent displays the operation mode details:

Parameter	Description
<b>Owner</b>	<p>Name of the backup repository that manages Veeam Agent.</p> <p>If Veeam Agent operates in the standalone mode, Veeam Agent will display the <i>Not Set</i> value.</p>
<b>Mode</b>	<p>Current Veeam Agent operating mode. Possible values:</p> <ul style="list-style-type: none"><li>• <i>Not Set</i> – Veeam Agent operates in the standalone mode.</li><li>• <i>Job</i> – Veeam Agent operates in the managed mode. Veeam Agent computer is protected by a backup job managed by backup server.</li><li>• <i>Policy</i> – Veeam Agent operates in the managed mode. Veeam Agent computer is protected by a backup job managed by Veeam Agent for Linux. Veeam Agent computer is connected to the Veeam backup server as a member of any protection group excluding protection group for pre-installed Veeam Agents.</li><li>• <i>Pre-installed</i> – Veeam Agent operates in the managed mode. Veeam Agent computer is protected by a backup job managed by Veeam Agent for Linux. Veeam Agent computer is connected to the Veeam backup server as a member of a protection group for pre-installed Veeam Agents.</li></ul> <p>Keep in mind that features and limitations of Veeam Agent operating in the managed mode are different from those in the standalone mode. To learn more about managed mode, see the <a href="#">Veeam Agent Management Guide</a>.</p>

For example:

```
user@srv01:~$ veeamconfig mode info
Owner: Backup server (backupserver001.tech.local)
Mode: Pre-installed
```

If Veeam Agent operates in the managed mode, you can reset it to the standalone mode at any time. To learn more, see [Resetting to Standalone Operation Mode](#).

# Resetting to Standalone Operation Mode

If Veeam Agent operates in the managed mode, you can manually reset it to the standalone mode from the Veeam Agent side. To learn more about operation modes, see the [Standalone and Managed Operation Modes](#).

Before you reset Veeam Agent to the standalone mode, consider the following:

- All backup jobs configured on the Veeam Agent computer will be deleted. If you plan to protect this computer with a standalone Veeam Agent, you will need to create new backup jobs.
- Veeam backup server settings including protection group configuration settings will be deleted.
- Previously created backup files will remain in the target backup repository. If the target repository is managed by the Veeam backup server, in the Veeam Backup & Replication console, they will be marked as *Orphaned*.
- If you want to reset Veeam Agent that operates in the *Job* or *Policy* mode, we recommend that you do the following:
  - Remove Veeam Agent computer from the protection group using the Veeam Backup & Replication console. To learn more about removing computers from a protection group in the Veeam Backup & Replication console, see the [Removing Computer from Protection Group](#) section in the Veeam Agent Management Guide. Veeam Agent on the protected computer will automatically switch to the standalone operation mode.
  - If Veeam Agent does not automatically switch to the standalone mode after you remove the Veeam Agent computer from the Veeam Backup & Replication configuration database, reset the operating mode on the Veeam Agent computer.
  - If Veeam Agent operates in the *Pre-installed* mode, Veeam Agent computer will be automatically removed from the protection group for pre-installed Veeam Agents in Veeam Backup & Replication.

To reset Veeam Agent to the standalone operating mode, run the following command:

```
veeamconfig mode reset
```

You can use the `--force` option to override additional input prompts and error messages:

```
veeamconfig mode reset --force
```

# Connecting to Veeam Backup & Replication

If you want to connect a Veeam Agent computer as a member of the protection group for pre-installed Veeam Agents to a Veeam backup server, you must apply connection settings from the protection group configuration file to Veeam Agent. The configuration file is one of the Veeam Agent setup files that you must obtain from your System Administrator. To learn more about deployment using external tools, see the [Deploying Veeam Agent for Linux](#) section in the Veeam Agent Management Guide.

To connect Veeam Agent to Veeam backup server:

1. Get the configuration file from your System Administrator and upload this file to the Veeam Agent computer.
2. Navigate to the directory where you have saved the configuration file and run the following command:

```
veeamconfig mode setvbrsettings --cfg <file_name>.xml --force
```

where:

- o `<file_name>` – configuration file name. Alternatively, you can specify the full path to the configuration file with the `--cfg` option.
- o `--force` – with this option enabled, Veeam Agent will override additional input prompts and error messages. This parameter is optional.

For example:

```
user@srv01:~$ veeamconfig mode setvbrsettings --cfg /home/Linux\ Servers\ Distribs/Linux/LinuxServers.xml
```

# Synchronizing with Veeam Backup Server

When Veeam Agent is managed by Veeam backup server, the connection between Veeam backup server and Veeam Agent computer added to a protection group is not persistent. Veeam Agent synchronizes with Veeam Backup & Replication every 6 hours. During the synchronization, Veeam Agent gets updated backup policies and configuration settings from the Veeam backup server, the Veeam backup server gets certificate details and session logs from Veeam Agent. To synchronize Veeam Agent immediately, run the following command:

```
veeamconfig mode syncnow
```

# Exporting Logs to Veeam Backup Server

If Veeam Agent is connected to the Veeam backup server as a member of the protection group for pre-installed Veeam Agents, Veeam Agent can collect the required logs, export them to an archive file and send to the Veeam backup server. This operation may be required if you want to report an issue and need to attach log files to the support case.

To export logs, use the following command:

```
veeamconfig mode exportdebuglogs
```

Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<agent>_<date>_<time>.tar.gz` and save the archive to the following folder on the Veeam backup server:

```
C:\ProgramData\Veeam\Backup\Endpoint\Other\AgentLogs\<computer_name>
```

where `<computer_name>` – name of the computer with Veeam Agent installed.

## TIP

If Veeam Agent operates in the standalone mode, you can export product logs only to a local directory on the Veeam Agent computer. To learn more, see [Exporting Product Logs](#).

# Uninstalling Veeam Agent for Linux

To uninstall Veeam Agent for Linux, you need to remove the `veeam-libs`, `veeam` and Veeam kernel module packages. To do this, run the following command with the name of the Veeam kernel module you used during installation – `veeamsnap` or `blksnap`:

*For CentOS 7 / RHEL / Oracle Linux / Fedora*

```
yum remove veeam veeam-libs veeamsnap
```

or

```
yum remove veeam veeam-libs blksnap
```

*For Rocky Linux / AlmaLinux*

```
yum remove veeam veeam-libs blksnap
```

*For openSUSE / SLES*

```
zypper rm veeam veeam-libs veeamsnap
```

or

```
zypper rm veeam veeam-libs blksnap
```

*For Debian / Ubuntu*

```
apt-get remove veeam veeam-libs veeamsnap
```

or

```
apt-get remove veeam veeam-libs blksnap
```

# Getting Started

To protect your computer from a disaster of any kind, you must perform the following operations in Veeam Agent for Linux:

1. Define what data you want to back up and configure the backup job.

Before you configure the backup job, you should decide on the following backup details:

- Backup destination: where you want to store your backed-up data.
- Backup scope: entire computer image, individual computer volumes or specific computer folders and files.
- Backup schedule: how often you want to back up your data.

After that, you can configure one or several backup jobs. The backup job captures the data that you have added to the backup scope and creates a chain of restore points in the target location. If your data gets lost or corrupted, you can restore it from the necessary restore point.

In Veeam Agent, you can configure the backup job in one of the following ways:

- [With the Backup Job wizard](#)
- [With the command line interface](#)

2. Monitor backup task performance.

You can use the Veeam Agent Control Panel to check how backup tasks are being performed, what errors have occurred during backup job sessions and so on. You can also use Veeam Agent command line interface to get information on backup and restore sessions status and view session logs. To learn more, see [Reporting](#).

3. In case of a disaster, you can restore the entire computer image or specific data on the computer. With Veeam Agent, you can perform data recovery operations in several ways:
  - You can boot from the Veeam Recovery Media and perform volume-level restore or file-level restore.
  - You can perform volume-level restore with Veeam Agent command line interface.
  - You can perform file-level restore with the Veeam Agent File Level Restore wizard.
  - You can export backup to a VHD virtual disk and attach this disk to a virtual machine to recover your computer in virtual environment.

To learn more, see [Performing Restore](#).



# Getting to Know User Interface

With Veeam Agent for Linux, you can perform backup, restore and configuration tasks in the following ways:

- [Using Veeam Agent control panel](#)

Veeam Agent control panel is a GUI-like user interface based on the `ncurses` programming library. With Veeam Agent control panel, you can perform all basic data protection tasks. You can configure a backup job, start and stop backup jobs, monitor backup job session performance and recover files and folders. When you perform restore tasks after booting from the Veeam Recovery Media, you can also perform volume-level restore with the Veeam Recovery Media wizard.

- [Using command line interface](#)

With Veeam Agent command line interface, in addition to operations that can be performed with the Veeam Agent control panel, you can perform a set of advanced tasks. For example, you can:

- Configure advanced settings for backup jobs: specify compression level and data block size.
- Perform operations with backup repositories.
- Perform volume-level restore without the need to boot from the Veeam Recovery Media.
- Export backups to VHD virtual disks.
- Monitor performance and status of any backup, restore and other data transfer session that was started in Veeam Agent.
- View detailed information on every backup that was created with Veeam Agent.
- Export/import Veeam Agent configuration database to/from a configuration file.

# Veeam Agent for Linux Control Panel

Veeam Agent for Linux control panel is a GUI-like user interface that lets users perform main backup and restore tasks in an easy way. With Veeam Agent for Linux control panel, you do not need to work with Linux shell and remember numerous commands. However, some advanced Veeam Agent for Linux operations are not supported by the control panel and can be performed with the command line interface only.

## IMPORTANT

You cannot use Veeam Agent for Linux control panel on terminals that do not support colors (for example, VT100).

To launch the Veeam Agent for Linux control panel, you can use the following commands:

```
veeamconfig ui
```

or

```
veeam
```

## NOTE

Veeam Agent for Linux control panel is based on the `ncurses` programming library. To use the Veeam Agent for Linux control panel, you must have the `ncurses` library installed in your Linux OS. To learn more, see [System Requirements](#).

When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent for Linux offers you to perform initial product setup. To learn more, see [Performing Initial Setup](#).

After you perform initial product setup, before you configure the first backup job, you can use the Veeam Agent for Linux control panel to perform the following operations:

- [Configure a new backup job](#).
- [Restore files and folders from existing backup](#).
- Manage [license](#) and [product logs](#).
- [Create a custom Veeam Recovery Media](#)

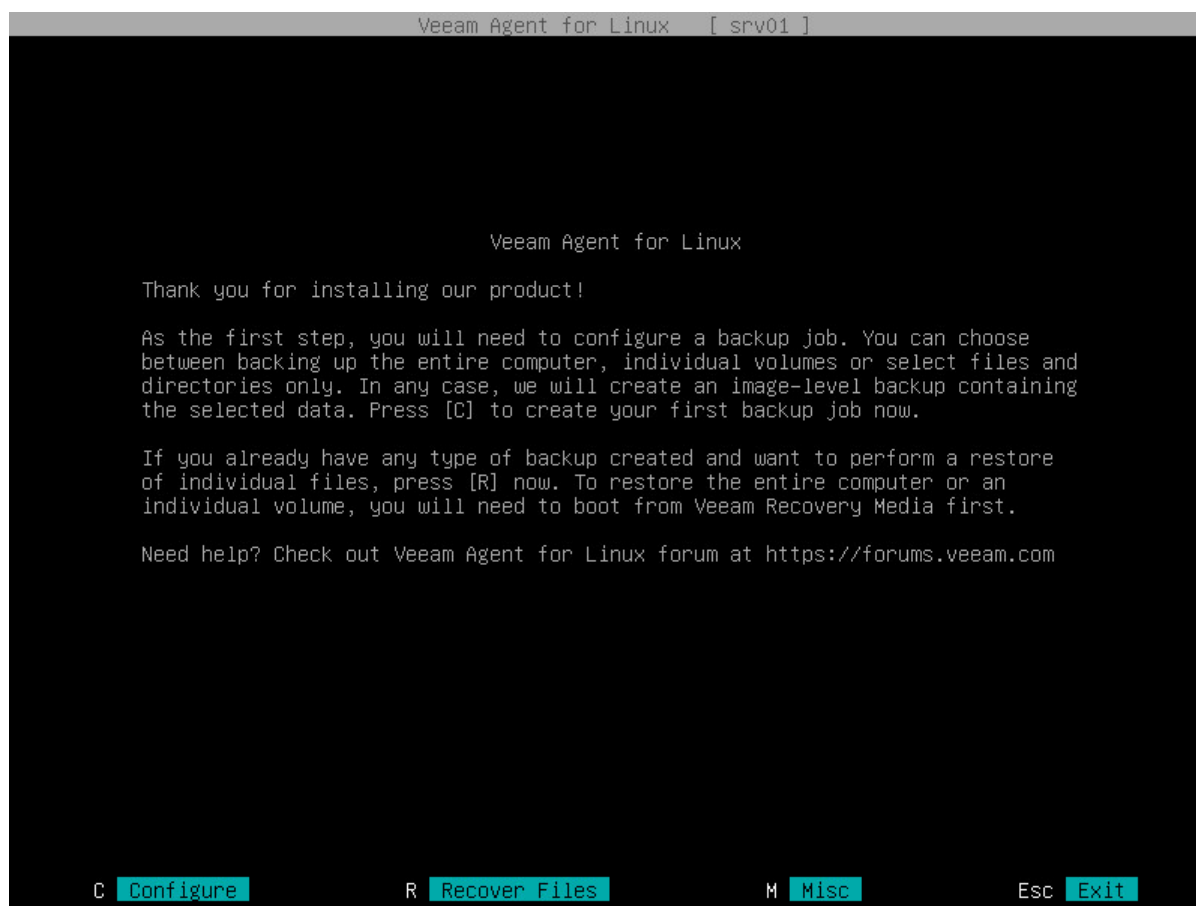
After you configure one or several backup jobs, you can also use the control panel to [start a backup job](#) and [work with backup job sessions](#).

## Navigating Veeam Agent for Linux Control Panel

In the Veeam Agent for Linux control panel, the use of a mouse is not supported. To start an operation, you need to use a specific key on your keyboard. For example, you can press the [C] key to start the backup job configuration, press the [S] key to start a backup job or press the [R] key to start the file-level restore process. Short help information on the currently available operations and keys is displayed at the bottom of the control panel.

To navigate the control panel, backup job configuration and file-level restore wizards, you can use the following keys:

- [Tab] – to switch between controls and buttons in the Backup Job wizard.
- [Up] and [Down] – to switch between items in a scrollable list.
- [Space] – to select the necessary item in a list. The selected item's mark may vary in different steps of the wizard.
- [Enter] – to proceed to the next step of a wizard or to view details of the backup job session selected in the list of sessions.
- [Backspace] – to return to the previous step of a wizard (you cannot use this button to change wizard steps when a text field is selected).
- [Esc] – to exit the currently used wizard or close the Veeam Agent for Linux control panel.



# Command Line Interface

Veeam Agent command line interface is a powerful tool that lets users perform advanced operations that are not supported by the Veeam Agent control panel.

To work with Veeam Agent using command line interface, you can use a terminal console (TTY) or a terminal emulator of your choice. All tasks in Veeam Agent are performed with the `veeamconfig` command-line utility. To perform tasks with Veeam Agent, you should construct the necessary command and type it in the Linux shell prompt.

You can view short help information on every Veeam Agent command at any time you need. To learn more, see [Viewing Help](#).

You should construct a command in the following format:

```
veeamconfig <command_1> <command_2> --<parameter_1> --<parameter_2> --<parameter_n>
```

where:

- `<command_1>` – command that defines a type of an object with which you want to perform a task. Currently, the following commands are available in Veeam Agent:
  - `aap`
  - `agreement`
  - `backup`
  - `cloud`
  - `config`
  - `downloadiso`
  - `gfs`
  - `grablogs`
  - `help`
  - `job`
  - `license`
  - `mode`
  - `objectstorage`
  - `patchiso`
  - `point`
  - `repository`
  - `schedule`
  - `session`
  - `ui`

- version
- vbrserver
- `<command_2>` – command that defines a task that you want to perform with an object of the specified type. For example, you can perform the following commands with backup repositories:
  - create
  - delete
  - edit
  - help
  - list
  - rescan
- `<parameter_1>`, `<parameter_2>`, `<parameter_n>` – parameters for the command that you want to execute. Commands may require one or several mandatory or optional parameters. Some commands, for example, `veeamconfig ui` and `veeamconfig [<command>] help` do not require parameters.

The following example shows the command that displays a list of backup repositories configured in Veeam Agent and the output of this command:

```
user@srv01:~$ veeamconfig repository list
```

Name	ID	Location	Typ
e Backup server			
Repository_1	{818e3a0f-8155-4a51-9430-248a203a43d1}	/home/backups	loca
1			
Repository_2	{2155a2e7-a1e9-4347-9d8b-cf8f3a6f3fcb}	172.17.53.47/veeam	cif
s			

# Viewing Help

You can view short help information on the specific Veeam Agent command. To view help, use the following command:

```
veeamconfig <command> help
```

where:

**<command>** – name of the command for which you want to view help information.

For example:

```
user@srv01:~$ veeamconfig help
```

or

```
user@srv01:~$ veeamconfig job help
```

or

```
user@srv01:~$ veeamconfig job create help
```

You can also view the manual page for the `veeamconfig` utility. Use the following command:

```
man veeamconfig
```

# Licensing

To work with Veeam Agent, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product.

If you want to use a commercial version of Veeam Agent, you must obtain a license and install it on the protected computer. If you do not install a license, the product will operate in the Free edition.

You can use the Veeam Agent control panel or Veeam Agent command line interface to install a license, monitor status of the installed license or remove the license if necessary.

# Product Editions

Veeam Agent for Linux offers three product editions that define product functionality and operation modes:

- *Server* – a commercial edition that provides access to all product functions and is intended for performing data protection tasks on servers that run Linux OS. Veeam Agent for Linux can operate in the server edition if a commercial license that supports this edition is installed on the protected computer.
- *Workstation* – a commercial edition that offers limited capabilities that are sufficient for performing data protection tasks on desktop computers and laptops that run Linux OS. Veeam Agent for Linux can operate in the workstation edition if a commercial license that supports this edition is installed on the protected computer.
- *Free* – a free edition that offers the same capabilities as the Workstation edition but does not come with a commercial support program. In contrast to the workstation and server editions, the Free edition does not require a license.

For more information about product editions, pricing and features available for them, see [this Veeam webpage](#).

## TIP

To check in which edition Veeam Agent for Linux currently operates, you can use the Veeam Agent for Linux control panel or command line interface. To learn more, see [Viewing License Information](#).

When you install a license on the protected computer, you can select in which edition Veeam Agent for Linux will operate: server edition or workstation edition (if both editions are supported by the license). If you use Veeam Agent for Linux with Veeam Backup & Replication, you must manage product licenses and editions from the Veeam Backup & Replication console. To learn more, see [Managing License with Veeam Backup & Replication](#).

After the license expires, Veeam Agent for Linux automatically switches to the Free edition. To learn more, see [License Expiration](#).

## Limitations for Free and Workstation Editions

Compared to the Server edition of Veeam Agent for Linux, Free and Workstation editions have the following limitations:

1. [Free edition] The number of backup jobs that you can configure in Veeam Agent for Linux is limited to one.
2. [Free edition] You cannot use a Veeam Cloud Connect repository as a target location for backup files.
3. [Free edition] You cannot perform direct backup to an object storage repository.
4. [Workstation edition] The number of backup jobs that you can configure in Veeam Agent for Linux is limited to one backup job targeted at a local drive, network shared folder, object storage repository or Veeam backup repository plus unlimited number of backup jobs targeted at a Veeam Cloud Connect repository.
5. [Free and Workstation editions] You cannot specify pre-freeze and post-thaw scripts in the backup job settings.
6. [Free and Workstation editions] You cannot specify database system processing settings.



# License Agreement

After you install Veeam Agent for Linux, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product. Until you accept the license agreements, you will not be able to perform backup and data recovery tasks with the Veeam Agent control panel and command line interface.

License agreements are located in the `/usr/share/doc/veeam` directory of the machine where you installed the product.

The process of accepting license agreements differs depending on the way you work with Veeam Agent – using the control panel or command line interface.

- When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent prompts you to accept the license agreements at the **Agreements** step of the initial setup wizard. To learn more, see [Accept License Agreements](#).
- When you run a Veeam Agent for Linux command, for example, `veeamconfig repository create`, Veeam Agent prompts you to accept license agreements. To accept the license agreement, type `y` or `yes` in the command prompt and press [Enter].

Alternatively, you can accept license agreements using the dedicated commands. To learn more, see [Accepting License Agreements](#).

# Installing License

When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent for Linux offers you to install a license at the [License](#) step of the initial setup wizard. You can choose to install the license immediately or postpone this operation.


If you choose to postpone the license installation process, you can install a license later at any time you need. Until you install a license, Veeam Agent for Linux will operate in the Free edition. To learn more, see [Product Editions](#).

## NOTE

If you choose not to install a license and use Veeam Agent for Linux in the Free edition, Veeam Agent for Linux will display a notification offering to install a license every time you open the control panel. The notification will appear in the control panel until Veeam Agent for Linux completes the first backup job session.

To install a license:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press [Enter].



4. In the **Manage license** window, make sure that the **Install** button is selected and press [Enter].

5. In the **Choose license** window, in the **File location** field, specify a path to the license key:
  - a. Select the **Browse** option with the [Tab] key and press [Space] or [Enter].
  - b. In the **Choose license file location** window, select the necessary directory and press [Enter].
  - c. Repeat the step 'b' until a path to the directory in which the license key resides appears in the **Current directory** field.

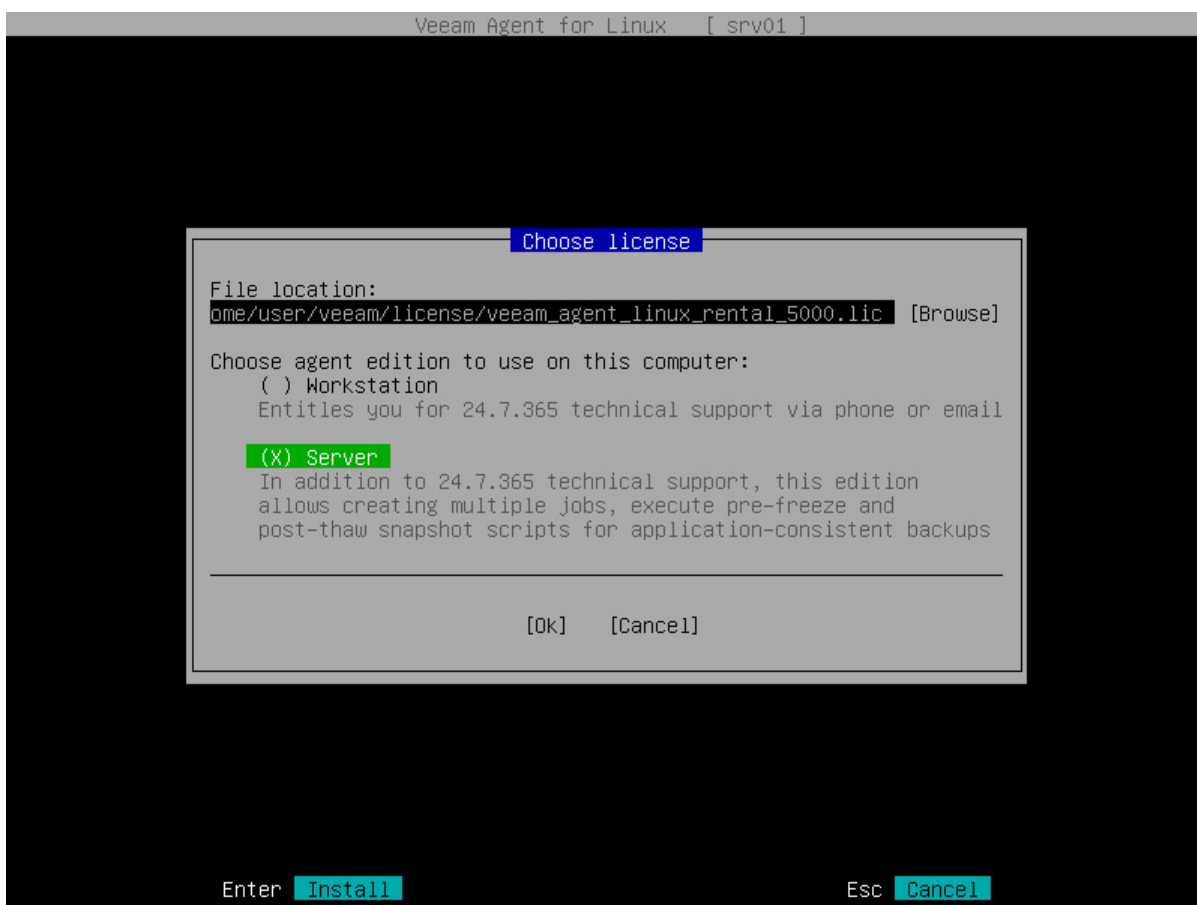
#### TIP

If you chose to install the license immediately from the Veeam Agent for Linux welcome screen notification, you will pass to the **Choose license** step right from the notification window.

6. In the **Choose agent edition to use on this computer** section, select the product edition in which Veeam Agent for Linux will operate and press [Enter]. To learn more about editions, see [Product Editions](#).
7. Veeam Agent for Linux will install the license and display a window notifying that the license is successfully installed. Press [Enter] to finish the license installation process.

#### TIP

After you install a license, you can view information about the license (expiration date, status of the license, current edition of the product and so on) in the **Manage license** window. You can also check information about the license using the Veeam Agent for Linux command line interface. To learn more, see [Viewing License](#).

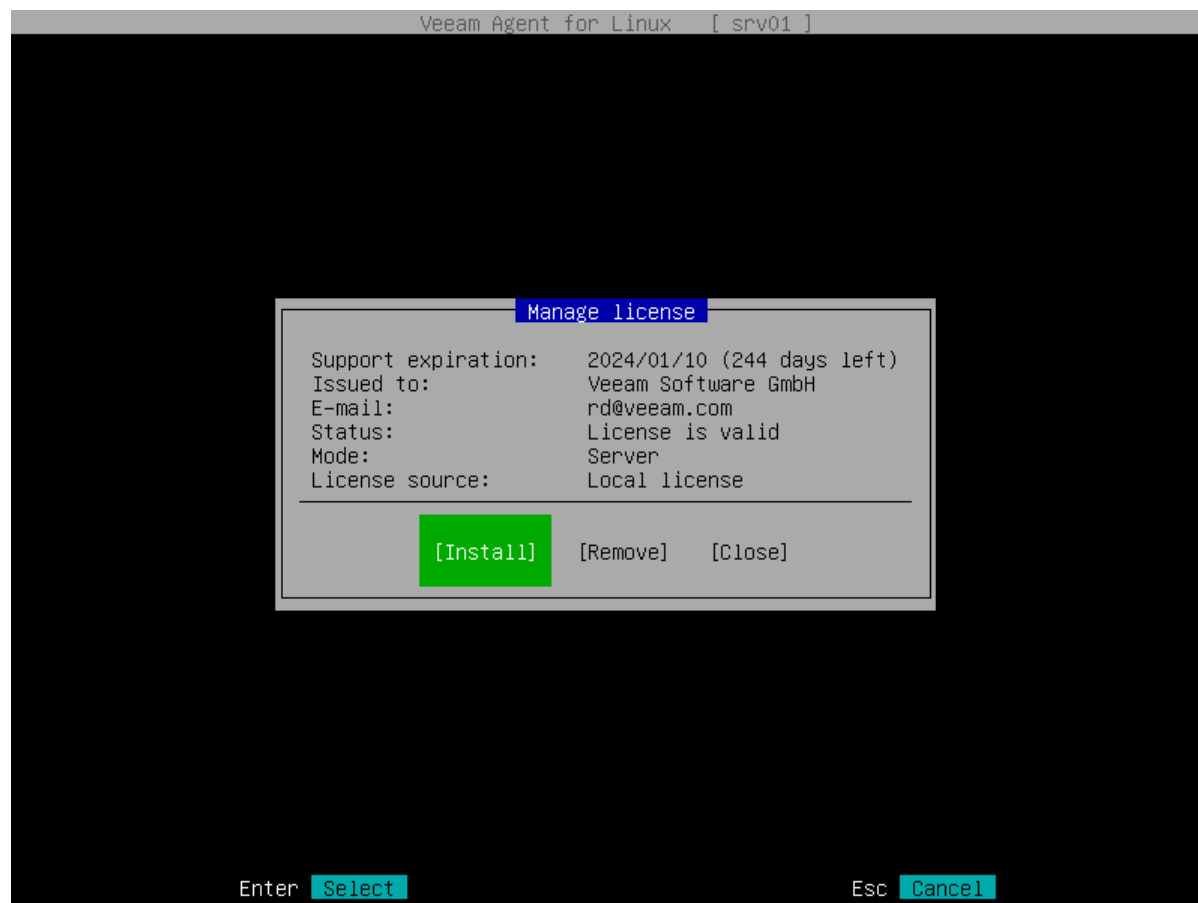


# Viewing License Information

To view information about the installed license, do the following:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press [Enter].

Veeam Agent for Linux will display information about the license.



# Removing License

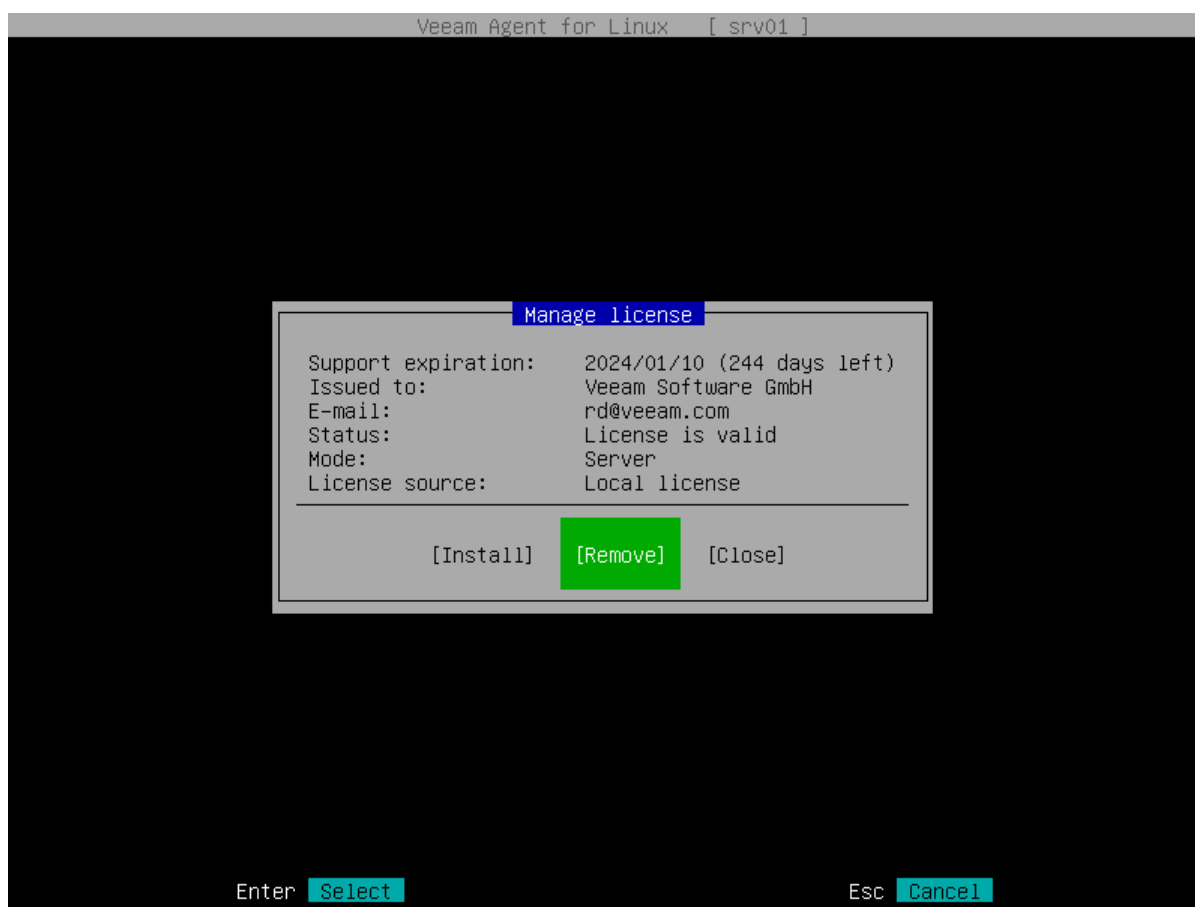
You can remove the license if necessary. To remove a license:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press [Enter].
4. In the **Manage license** window, press [Tab] to select the **Remove** button, then press [Enter].
5. Veeam Agent for Linux will remove the license and display a window notifying that the license is successfully removed. Press [Enter] to finish the license removal process.

## NOTE

After you remove the license, Veeam Agent for Linux will continue to operate in the Free edition. Consider the following:

- If Veeam Agent for Linux operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will fail.
- If pre-freeze or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will fail.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will fail.




# License Expiration

30 days before the license expiration date, Veeam Agent for Linux will display a warning at the top of the control panel. After the license expires, Veeam Agent for Linux will switch to the Free edition.

Consider the following:

- If Veeam Agent for Linux operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will be failing.
- If pre-freeze or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will be failing.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will be failing.

You can switch to the Free edition manually at any time if necessary. To learn more, see [Removing License](#).



# Managing License with Command Line Interface

You can use the Veeam Agent for Linux command line interface to perform the following operations with the license:

- [Accept license agreements for the product itself and its third-party components.](#)
- [Install a license on the protected computer.](#)
- [View information about the license.](#)
- [Remove the license.](#)

# Accepting License Agreements

To work with Veeam Agent for Linux, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product. Until you accept license agreements, you can use the `veeamconfig` utility to run the following commands only:

- `veeamconfig agreement show`
- `veeamconfig help` (or `veeamconfig -h` or `veeamconfig --help`)
- `veeamconfig mode info`
- `veeamconfig mode reset`
- `veeamconfig version` (or `veeamconfig -v` or `veeamconfig --version`)
- `veeamconfig ui`

To accept license agreements, use the following command:

```
veeamconfig agreement accepteula && veeamconfig agreement acceptthirdpartylicenses
```

## TIP

To check whether license agreements are accepted, use the following command: `veeamconfig agreement show`.



# Installing License

To install a license, use the following command:

```
veeamconfig license install --path <path> --workstation
```

or

```
veeamconfig license install --path <path> --server
```

where:

- `<path>` – path to the license key file in the local file system of your computer.
- `workstation` or `server` – edition in which Veeam Agent will operate. To learn more about editions, see [Product Editions](#).

Veeam Agent for Linux will install the license and display information about the license. You can also view this information later at any time. To learn more, see [Viewing License Information](#).

For example:

```
user@srv01:~$ veeamconfig license install --path /home/user/veeam/license/veeam
.lic --server
License was installed successfully.
License information:
  License source: Local license
  Mode: Server
  Support expiration: 2019/09/20 (649 days left)
  Status: License is valid.
  Issued to: TechCompany
  E-mail: administrators@tech.com
```

## TIP

You can also install a license using the Veeam Agent control panel. To learn more, see [Installing License](#).

# Viewing License Information

You can view information about the installed license. Use the following command:

```
veeamconfig license show
```

Veeam Agent for Linux will display information about the license. For example:

```
user@srv01:~$ veeamconfig license show
License information:
  License source: Local license
  Mode: Server
  Support expiration: 2019/09/20 (649 days left)
  Status: License is valid.
  Issued to: TechCompany
  E-mail: administrators@tech.com
```

# Removing License

You can remove a license with the following command:

```
veeamconfig license remove
```

After you remove the license, Veeam Agent for Linux will continue to operate in the Free edition. Consider the following:

- If Veeam Agent operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will fail.
- If pre-freeze or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will fail.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will fail.

# Performing Backup

You can back up your data to protect the entire computer image, individual volumes or folders and files on your computer. To back up your data, you must configure a backup job. Depending on the product edition, Veeam Agent lets you configure one or several backup jobs targeted at the same or different backup repositories.

You can configure a backup job that will automatically back up your data by the defined schedule. You can also start a backup job manually at any time.

# Creating Custom Veeam Recovery Media

In addition to the generic Veeam Recovery Media that is available for download at the Veeam website, you can create a custom Veeam Recovery Media. This option may be helpful if your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media. When you create a custom Veeam Recovery Media, Veeam Agent updates the generic Veeam Recovery Media: copies the Linux kernel running on your computer with its currently loaded modules and includes them into the custom recovery image.

## Prerequisite Software Requirements

Before you create custom Veeam Recovery Media, check the following prerequisites:

- The Linux system must have the `genisoimage` package installed. For openSUSE and SLES 15 SP1 – 15 SP5 distributions, the Linux system must have the `mkisofs` package installed.
- The Linux system must have the `mksquashfs` and `unsquashfs` utilities installed.
- For custom Veeam Recovery Media with EFI support, the Linux system must have the following packages installed:
  - `xorriso`
  - `isolinux` (or `syslinux`, if the software package repository of your Linux distribution lacks the `isolinux` package)
- For the scenario where you create a custom Veeam Recovery Media using the **Download and patch ISO** option, the Linux system must have the `wget` utility installed.

## Considerations and Limitations

Before you create custom Veeam Recovery Media, consider the following:

- The custom recovery image contains an unsigned Linux kernel. As a result, you cannot use it for UEFI systems with Secure Boot enabled.
- If you plan to use live patching to create a custom Veeam Recovery Media, consider the following [limitations](#).
- You cannot create a custom Veeam Recovery Media for Veeam Agent computers that run Linux kernel version earlier than 3.10. For a workaround, see [this Veeam KB article](#).
- You cannot create a custom Veeam Recovery Media for Veeam Agent computers running on IBM Power Systems.

## Creating Custom Veeam Recovery Media

You can create a custom Veeam Recovery Media in one of the following ways:

- With the Veeam Agent control panel. You can perform this operation in the following conditions:
  - During the process of initial product setup, at the [Recovery ISO](#) step of the initial setup wizard.
  - Any time you need, in the **Miscellaneous** menu. For details, see [Creating Custom Veeam Recovery Media with Control Panel](#).

- With the Veeam Agent command line interface. For details, see [Creating Custom Veeam Recovery Media with Command Line Interface](#).

If you create a custom Veeam Recovery Media using the command line interface, you can also specify a directory that contains additional drivers that you want to include in the recovery media.

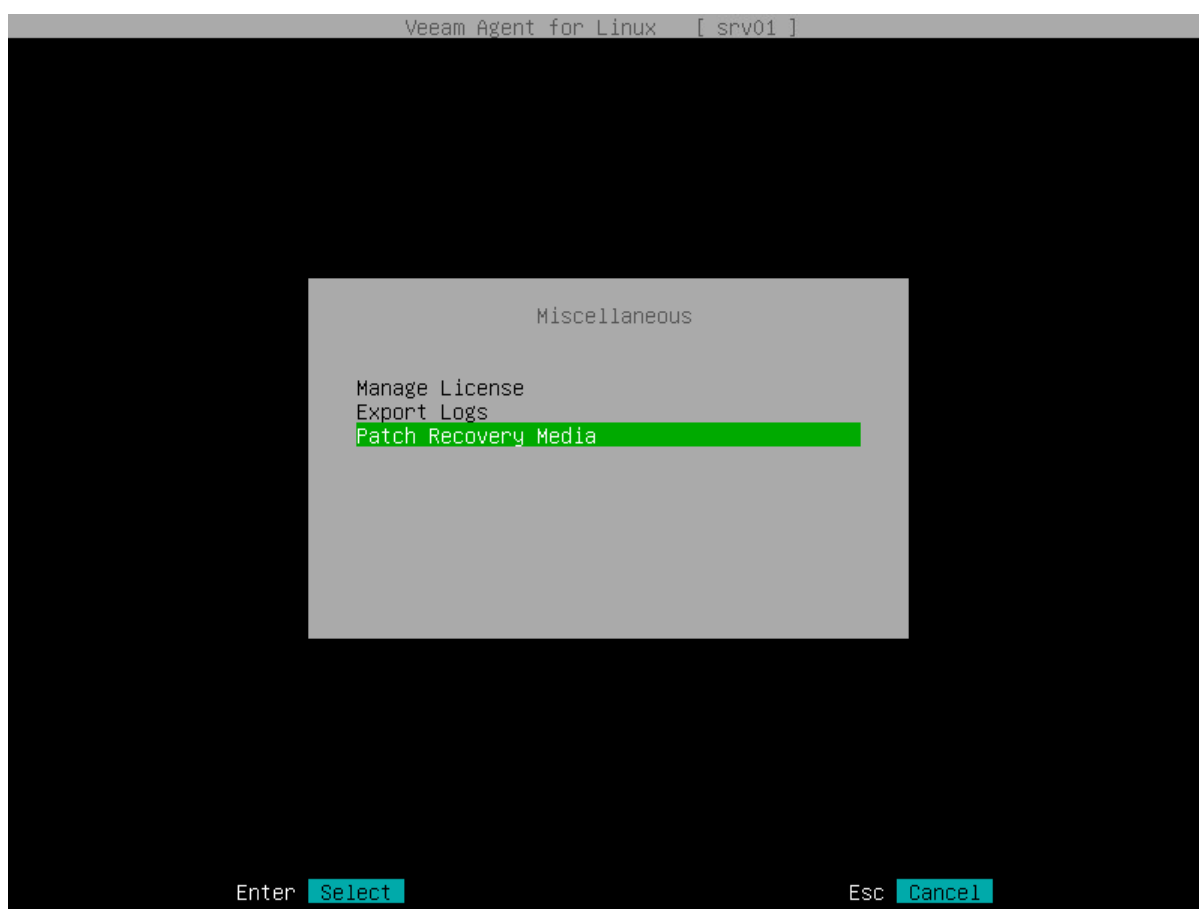
# Creating Custom Veeam Recovery Media with Control Panel

To create custom Veeam Recovery Media with the Veeam control panel, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, select the **Patch Recovery Media** option and press [Enter].

## IMPORTANT

Recovery Media patching is not supported by Veeam Agent for Linux on Power.



4. Press [Tab] and select how you want to create a custom Veeam Recovery Media depending on the location of the generic recovery media ISO file:
  - If you have not downloaded the generic Veeam Recovery Media, make sure that the **Download and patch ISO** option is selected and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer and use this image to create the custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you want only to download the generic Veeam Recovery Media, select the **Only download ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer. You can use the downloaded ISO file later to boot your Veeam Agent computer or to create a custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you have already downloaded the generic Veeam Recovery Media to a local directory on the Veeam Agent computer or to a network shared folder, select the **Patch local ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will use the generic Veeam Recovery Media ISO file to create the custom Veeam Recovery Media.

The name of the generic Veeam Recovery Media ISO file depends on the recovery image version, Veeam Agent computer architecture and the source from which you downloaded the ISO file: from the product download page or Veeam software repository. To learn more, see [Veeam Recovery Media Versions](#).

5. If you selected the **Download and patch ISO** or **Patch local ISO** option, the **EFI system** option is available. If you want to boot the Veeam Recovery Media on EFI-based systems, select the **EFI system** option with the [Tab] key and press [Space].

If you do not enable this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

6. If you selected the **Patch local ISO** option, in the **Path to local ISO** field, specify a path to the ISO file of the generic Veeam Recovery Media:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Path to ISO** window, select the necessary directory and press [Enter].
  - c. Repeat the step 'b' until a path to the directory in which the recovery media ISO file resides appears in the **Current directory** field.
  - d. In the directory where the recovery media ISO file resides, select the ISO file and press [Enter].

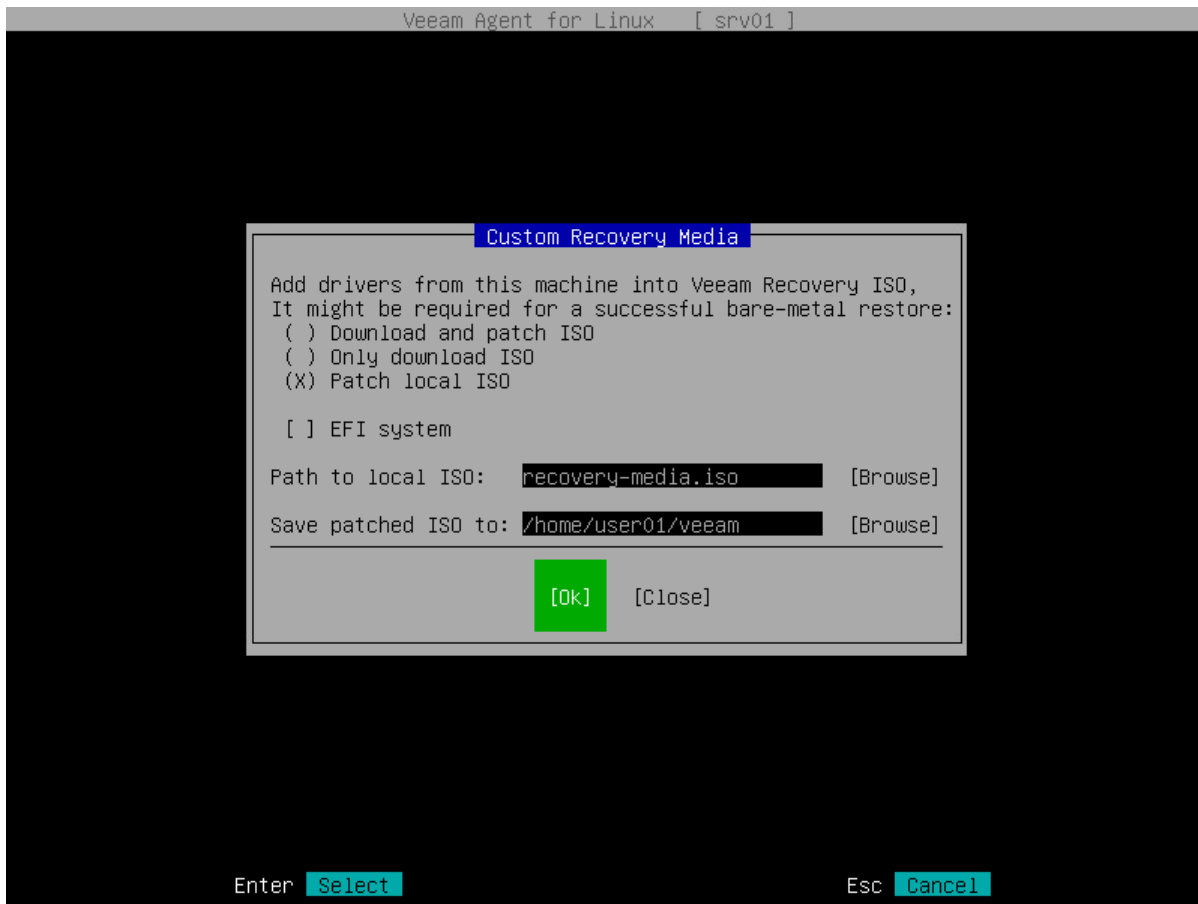
7. Specify a path to the resulting ISO file of the Veeam Recovery Media.

If you selected the **Download and patch ISO** or **Patch local ISO** option, in the **Save patched ISO to** field, you can specify a path to the resulting ISO file of the custom Veeam Recovery Media; if you selected the **Only download ISO** option, in the **Save ISO to** field, specify a path to the resulting ISO file of the generic Veeam Recovery Media:

- a. Select the **Browse** option with the [Tab] key and press [Enter].
- b. In the **Save patched ISO to** window, select the necessary directory and press [Enter].
- c. Repeat step 'b' until a path to the directory where you want to save the resulting custom recovery media ISO file appears in the **Current directory** field.
- d. Select the **OK** button with the [Tab] key and press [Enter].



8. To start the custom recovery media creation process, select the **Next** button with the [Tab] key and press [Enter].



# Creating Custom Veeam Recovery Media with Command Line Interface

To create a custom Veeam Recovery Media, you need to perform the following operations:

- [Download the ISO file of the generic Veeam Recovery Media](#). You can download this image from the [Veeam software repository](#) or from [this Veeam webpage](#)..
- [Using the downloaded ISO file, create the Custom Veeam Recovery Media](#).

## IMPORTANT

Recovery Media patching is not supported by Veeam Agent for Linux on Power.

## Downloading Generic Recovery Media

To download the generic Veeam Recovery Media with the command line interface, use the following command:

```
veeamconfig downloadiso --output <output_path>
```

where:

<output\_path> – path to the downloaded ISO file of the generic Veeam Recovery Media.

Veeam Agent downloads the ISO file of the generic Veeam Recovery Media depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

For example:

```
$ veeamconfig downloadiso --output /mnt/veeam/iso
```

## Creating Custom Recovery Media

To create the custom Veeam Recovery Media with the command line interface, use the following command:

```
veeamconfig patchiso --input <input_path> --output <output_path> --copy <additional_path>
```

or

```
veeamconfig patchiso --efi --input <input_path> --output <output_path> --copy <additional_path>
```

where:

- <input\_path> – path to the ISO file of the generic Veeam Recovery Media.
- <output\_path> – path to the resulting ISO file of the custom Veeam Recovery Media.

- `<additional_path>` – path to a directory with additional drivers that you want to include in the Veeam Recovery Media.

When you boot from the custom Veeam Recovery Media, the content of the directory specified with the `<additional_path>` parameter will be available in the root folder of the recovery environment.

- `--efi` – option that defines whether custom Veeam Recovery Media should be able to boot on EFI-based systems. Without this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

For example:

```
$ veeamconfig patchiso --input /mnt/veeam/iso/veeam-recovery-amd64-6.0.0.iso --  
output /mnt/veeam/iso/veeam-recovery-media-srv01.iso --copy /tmp/template --efi
```

# Creating Backup Jobs

You can choose one of the following backup modes:

- Backup of an entire computer image
- Backup of specific computer volumes, for example, a system volume or secondary volume
- Backup of individual files and folders

[For Server Edition] You can configure one or several backup jobs to back up your data. Configuring several backup jobs may be useful in the following situations:

- You can configure separate backup jobs for volume-level backup and file-level backup.
- You can configure backup jobs targeted at different backup repositories to keep several copies of your backed-up data at different locations.
- You can configure several backup jobs and define individual schedule for every job to back up necessary data at the desired time.

With Veeam Agent, you can configure the backup job in one of the following ways:

- [With the Backup Job wizard](#)
- [In the command line interface](#)

# Before You Begin

## Prerequisites

Before you configure the backup job, check the following prerequisites:

- The target location where you plan to store backup files must have enough free space.
- [For Veeam Backup & Replication repository targets] You can store created backups in a backup repository only if the backup server runs Veeam Backup & Replication 12.0 or later.
- [For Veeam Backup & Replication repository targets] If you plan to use a Veeam Backup & Replication repository as a target for backups, you must pre-configure user access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

## Considerations and Limitations

Before you configure the backup job, consider the following:

- When you configure the backup job with the Backup Job wizard, Veeam Agent creates the job with default values for the following advanced settings: compression level and data block size. You can specify custom values for these settings in the command line interface when you create a new backup job or edit an existing backup job. For more information, see [Creating Backup Job in Command Line Interface](#) and [Editing Backup Job Settings](#).
- You cannot save the backup of entire computer on the local computer disk. Use an external hard drive or USB drive, network shared folder or backup repository as a target location.
- [For nosnap Veeam Agent for Linux and Veeam Agent for Linux on Power] When you back up the entire computer or create a volume-level backup, Veeam Agent backs up the boot partition in the snapshot-less mode. To ensure consistency of the boot partition data during backup, Veeam Agent uses the `fsfreeze` or `umount` commands.
- Veeam Agent stops running the backup job after 21 days (504 hours).
- Veeam Agent does not back up data to which symbolic links are targeted. It only backs up the path information that the symbolic links contain. After restore, identical symbolic links are created in the restore destination.
- Veeam Agent does not support backup of bind mount points. In the scope of the backup job, you must specify the path to the original mount point instead.

## Navigating Backup Job Wizard

The Backup Job wizard window comprises the following areas:

- The navigation pane, located on the left of the window, displays the list of wizard steps and currently selected step of the wizard.
- The working area displays controls relating to a specific step of the wizard.
- The buttons area, located at the bottom of the window, displays buttons that you can use to switch between steps of the wizard (**Previous** and **Next**) and close the wizard (**Cancel** and **Finish**).

In the Backup Job wizard, the use of a mouse is not supported. To navigate the Backup Job wizard and associated dialog windows, you can use the following keys:

- [Tab] – to switch between displayed controls in the working area and buttons in the buttons area. The currently selected control or button is highlighted with a green color.
- [Up] and [Down] – to switch between items in a scrollable list.
- [Space] – to select the necessary item in a list. The selected item's mark may vary in different steps of the wizard.
- [Enter] – to proceed to the next step of the wizard or to open a directory.
- [Backspace] – to return to the previous step of a wizard.
- [Esc] – to cancel the backup job configuration and exit the wizard.

#### TIP

You can switch between steps of the Backup Job wizard in two ways. The easier and more comfortable way is to use the [Enter] key to proceed to the next step and [Backspace] key to return to the previous step of the wizard. You can also use the [Tab] key to select the **Next** or **Previous** button in the buttons area and then press [Enter] to switch to the next or previous step of the wizard respectively.

# Creating Backup Job with Backup Job Wizard

You can configure volume-level and file-level backup jobs with the Backup Job wizard.

# Step 1. Launch Backup Job Wizard

To launch the **Backup Job** wizard, do the following:

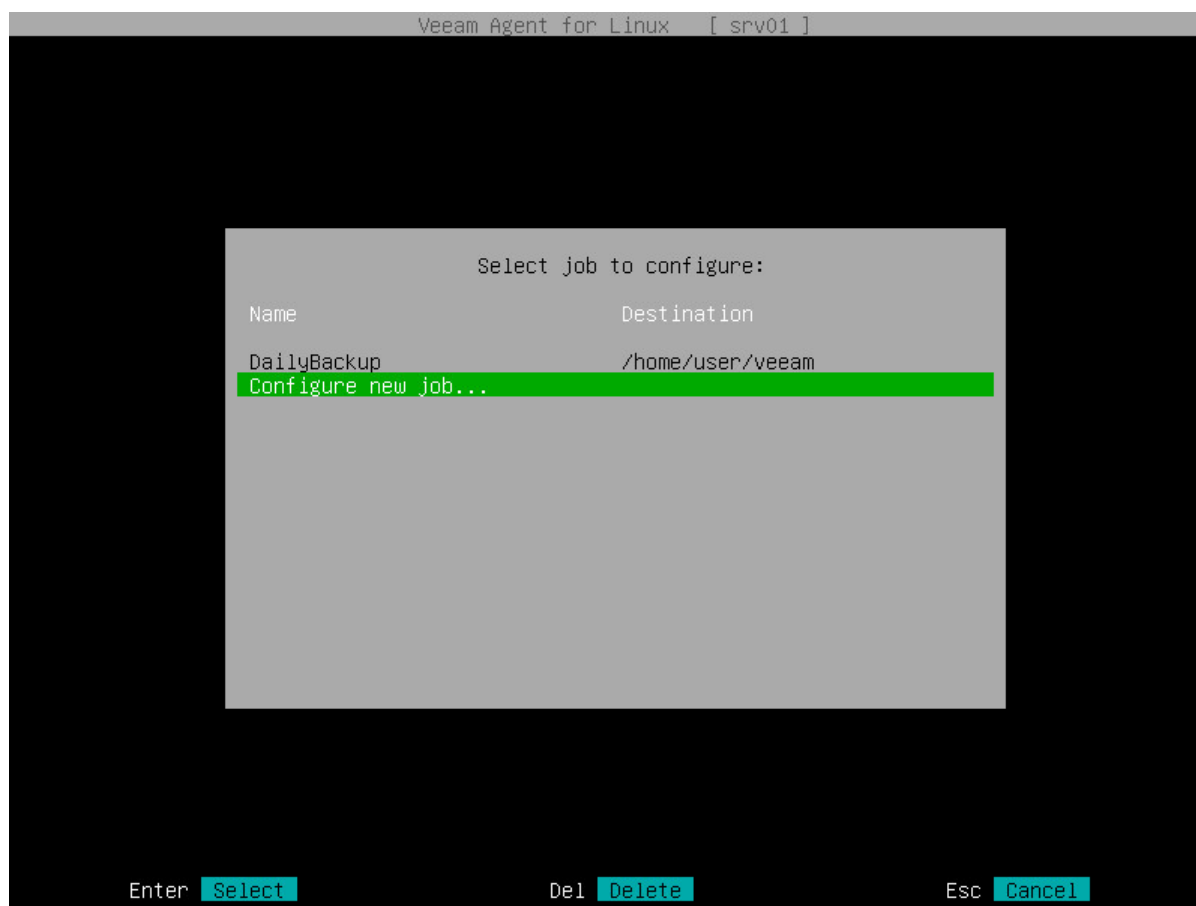
1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. If you have not configured any backup jobs yet, Veeam Agent will display a welcome screen. Press the [C] key to proceed to the Backup Job wizard and configure the backup job.
3. If you have already configured and performed a backup job, Veeam Agent will display the list of backup job sessions. When you press the [C] key to launch the Backup Job wizard, Veeam Agent will display a list of configured backup jobs. To configure a new backup job, select the **Configure new job** option and press [Enter].

## NOTE

The **Configure new job** option is not available if Veeam Agent for Linux operates in the Free edition and you have already configured one backup job.

To edit settings of a backup job that you have already configured, select the job in the list and press [Enter]. To learn more, see [Editing Backup Job Settings](#).

If you have decided not to create a backup job, press [Esc] to close the list of backup jobs and return to the welcome screen. After that, you can press [Esc] once again to return to the command line interface.



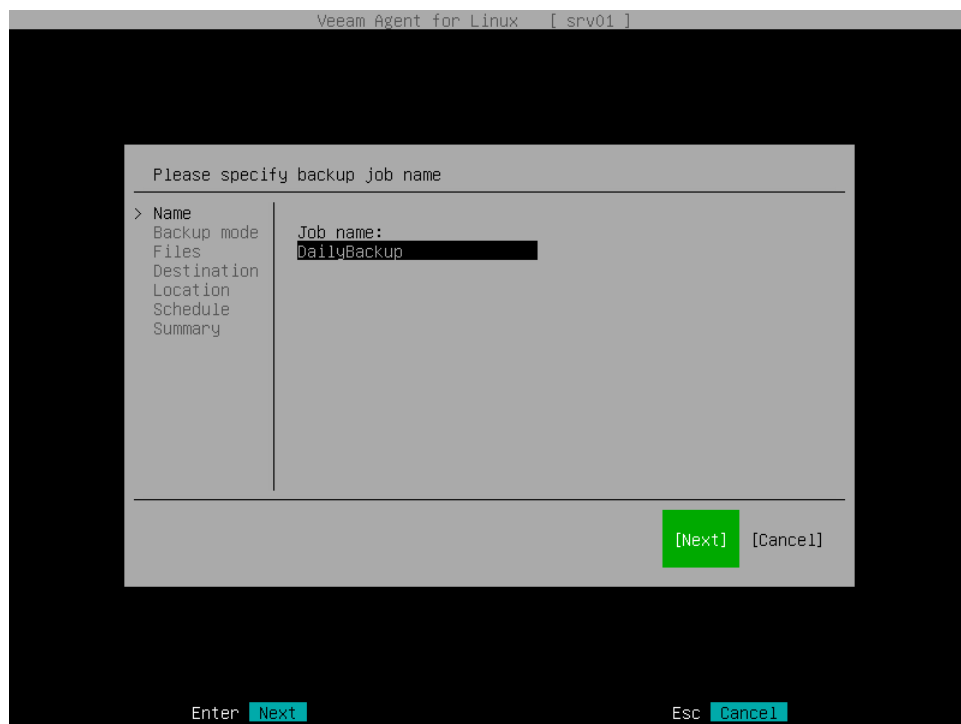


## Step 2. Specify Backup Job Name

At the **Name** step of the wizard, in the **Job name** field, type the name for the backup job and press [Enter].

### TIP

To proceed to the next step of the wizard, you can also select the **Next** button with the [Tab] key and then press [Enter].



## Step 3. Select Backup Mode

At the **Backup mode** step of the wizard, select the mode in which you want to create a backup:

1. Select the necessary backup mode. You can select one of the following options:
  - **Entire machine** – select this option if you want to create a backup of the entire computer image. When you restore data from such backup, you will be able to recover the entire computer image as well as data on specific computer volumes: files, folders, application data and so on. With this option selected, you will pass to the [Destination](#) step of the wizard.
  - **Volume level backup** – select this option if you want to create a backup of specific computer volumes, for example, the system volume. When you restore data from such backup, you will be able to recover data on these volumes only: files, folders, application data and so on. With this option selected, you will pass to the [Volumes](#) step of the wizard.
  - **File level backup** – select this option if you want to create a backup of individual directories on your computer. With this option selected, you will pass to the [Files](#) step of the wizard.
2. [For file-level backup] If you want to perform backup in the snapshot-less mode, select **Disable snapshot**. With this option selected, Veeam Agent will not create a snapshot of the backed-up volumes during backup. This allows Veeam Agent to back up data residing in file systems that are not supported for snapshot-based backup with Veeam Agent. To learn more, see [Snapshot-Less File-Level Backup](#).

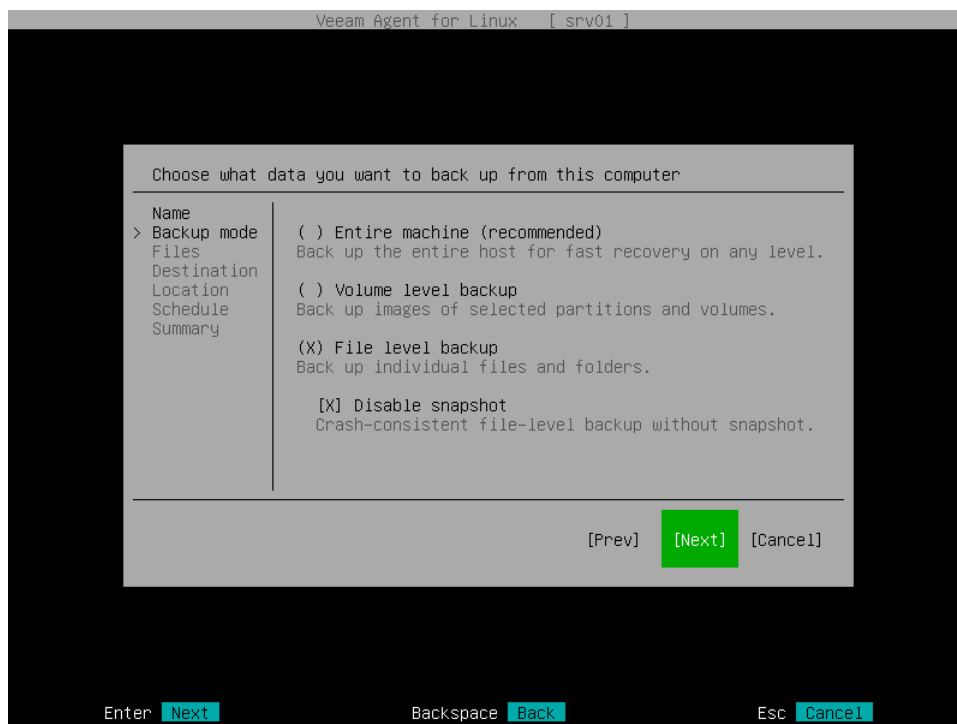
### IMPORTANT

Consider the following:

- [For entire machine backup] Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).
- [For volume-level backup] Volume-level backup job relies on a device name in the `/dev` directory. Device names in the `/dev` directory (for example, `/dev/md-127`, `/dev/dm-1`) must stay persistent for backed-up volumes. Otherwise, the job will back up the wrong volume.
- [For file-level backup] If the backed-up file system has a complex folder structure with many hierarchy levels, during incremental backup, the inbound network traffic on the Veeam Agent computer may exceed by far the outbound traffic. Significant amount of data can be transferred to the Veeam Agent computer from the target backup location even if few files are changed since the previous job session.

### TIP

File-level backup is typically slower than volume-level backup. If you plan to back up all folders with files on a specific volume, it is recommended that you configure volume-level backup instead of file-level backup.



## Step 4. Specify Backup Scope Settings

Specify backup scope for the backup job:

- [Select volumes to back up](#) — if you have selected the **Volume level backup** option at the [Backup Mode](#) step of the wizard.
- [Select folders to back up](#) — if you have selected the **File level backup** option at the [Backup Mode](#) step of the wizard.

### Selecting Volumes to Back Up

The **Volumes** step of the wizard is available if you have chosen to create a volume-level backup.

At this step of the wizard, you must specify the backup scope — define what volumes you want to include in the backup. Veeam Agent lets you include the following types of objects in the volume-level backup:

- [Block devices \(entire disks and individual volumes\)](#)
- [Mount points](#)
- [LVM logical volumes and volume groups](#)
- [BTRFS storage pools and subvolumes](#)

### Selecting Devices

To add a block device to the backup scope, do the following:

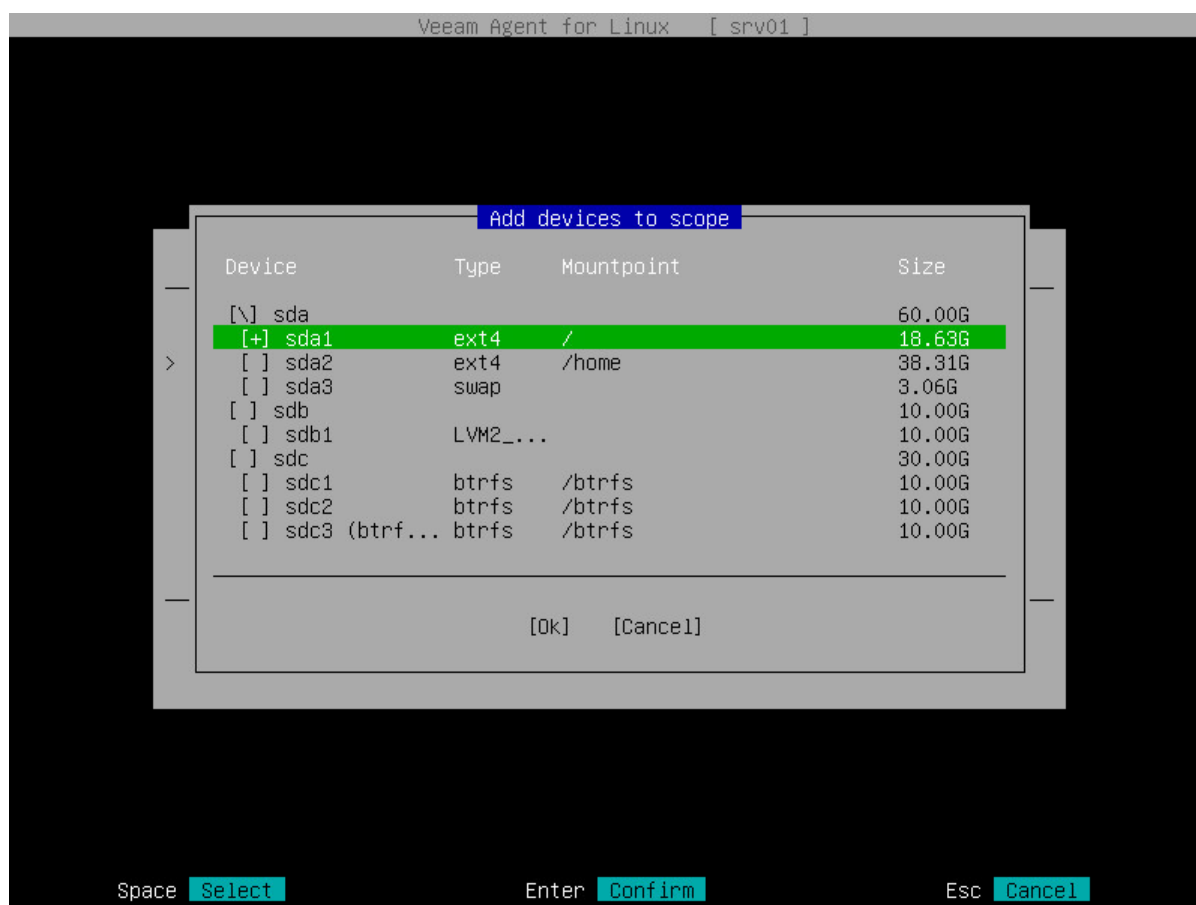
1. At the **Volumes** step of the wizard, make sure that the **Device** option is selected and press [Enter].
2. In the **Add devices to scope** window, select individual volumes or entire computer disks that you want to include in the backup and press [Enter].
  - To include individual volumes of your computer in the backup, select block devices that represent volumes that you want to back up, for example: *sda1* or *sda6*.
  - To include all volumes on a computer disk in the backup, select block devices that represent disks whose volumes you want to back up, for example: *sda* or *sdb*. All volumes on the selected disk will be automatically selected, too.

To navigate the list of volumes and select the necessary items, use the [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).

If you have created several system partitions, for example, a separate partition for the `/boot` directory, you should remember to include all of these partitions in the backup. Otherwise, Veeam Agent does not guarantee that the OS will boot properly when you attempt to recover from such backup.

## NOTE

If you include a block device in the backup, and this block device is a physical volume assigned to an LVM volume group, Veeam Agent will include the whole LVM volume group in the backup.



## Selecting Mount Points

### IMPORTANT

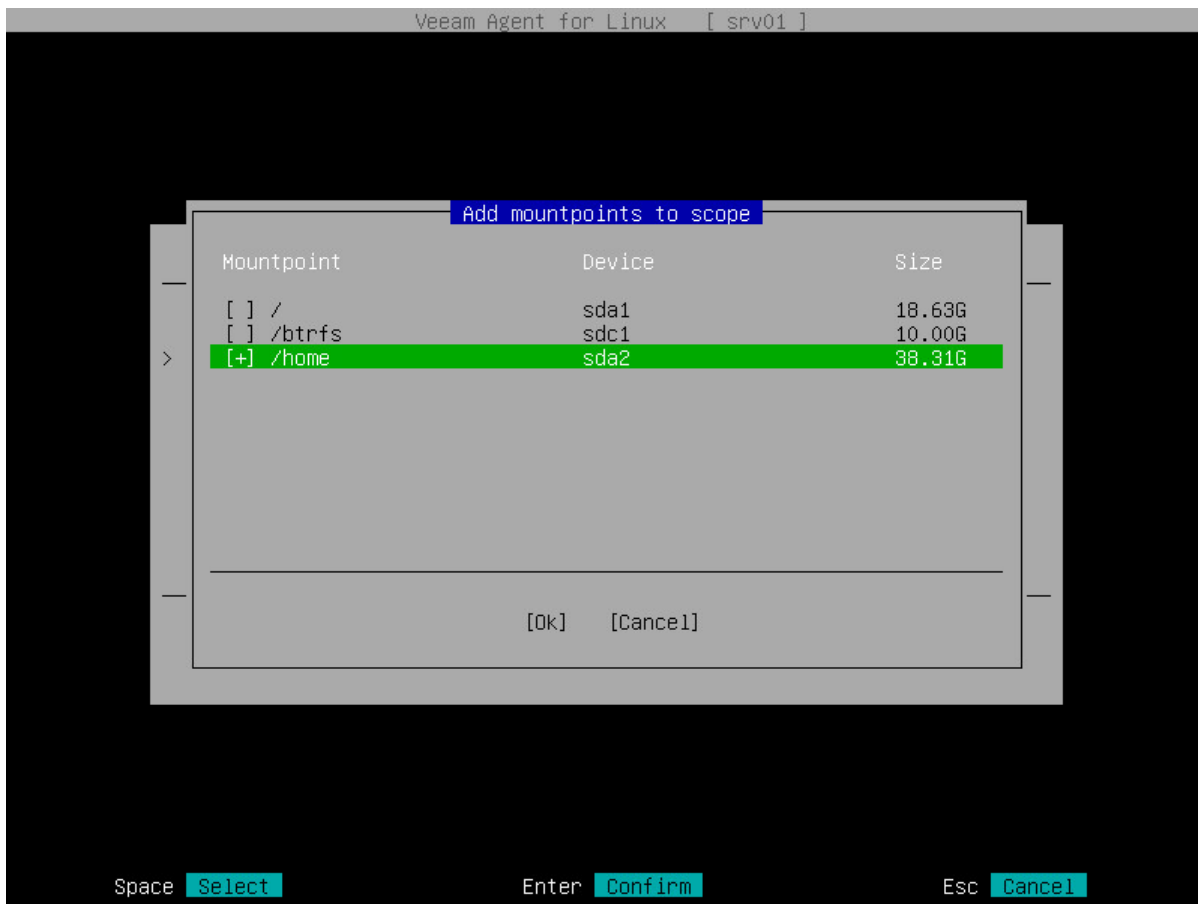
Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.

To add a mount point to the backup scope, do the following:

1. At the **Volume** step of the wizard, select the **Mountpoint** option and press [Enter].

2. In the **Add mountpoints to scope** window, select mount points that you want to include in the backup and press [Enter].

To navigate the list of mount points and select the necessary mount points, use [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).



## Selecting LVM Volumes

To add an LVM logical volume or volume group to the backup scope, do the following:


1. At the **Volume** step of the wizard, select the **LVM** option and press [Enter].
2. In the **Add LVM to scope** window, select LVM logical volumes or volume groups that you want to include in the backup and press [Enter].

To navigate the list of LVM volumes and select the necessary items, use [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).

If you include an LVM volume group in the backup, all LVM logical volumes in the selected volume group will be automatically selected, too.

### NOTE

Veeam Agent does not back up LVM snapshots.



## Selecting BTRFS Volumes

To add a BTRFS storage pool or subvolume to the backup scope, do the following:


1. At the **Volume** step of the wizard, select the **BTRFS** option and press [Enter].
2. In the **Add BTRFS to scope** window, select BTRFS storage pools or subvolumes that you want to include in the backup and press [Enter].

To navigate the list of BTRFS pools and subvolumes and select the necessary items, use [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).

Veeam Agent identifies BTRFS storage pools by UUIDs. If you include a BTRFS pool in the backup, all BTRFS subvolumes in the selected pool will be automatically selected, too.

### NOTE

You cannot add read-only BTRFS snapshots to the backup scope.



## Selecting Files and Directories to Back Up

The **Files** step of the wizard is available if you have chosen to create a file-level backup.

At this step of the wizard, you must specify the backup scope – define what directories with files you want to include in the backup.

In the file-level backup mode, you must include in the backup at least one directory. If you do not want to back up some subdirectories of the specified directory, you can exclude these directories from the backup.

You can also include or exclude files of a specific type in/from the backup. You can specify file names explicitly or use UNIX wildcard characters to define file name masks. Veeam Agent will apply the specified file name masks to files in directories that are included in the backup.

To specify the backup scope:

1. At the **Files** step of the wizard, make sure that the **Add directories** option is selected and press [Enter].
2. In the **Choose directories** window, select one or several directories that you want to include in the file-level backup.

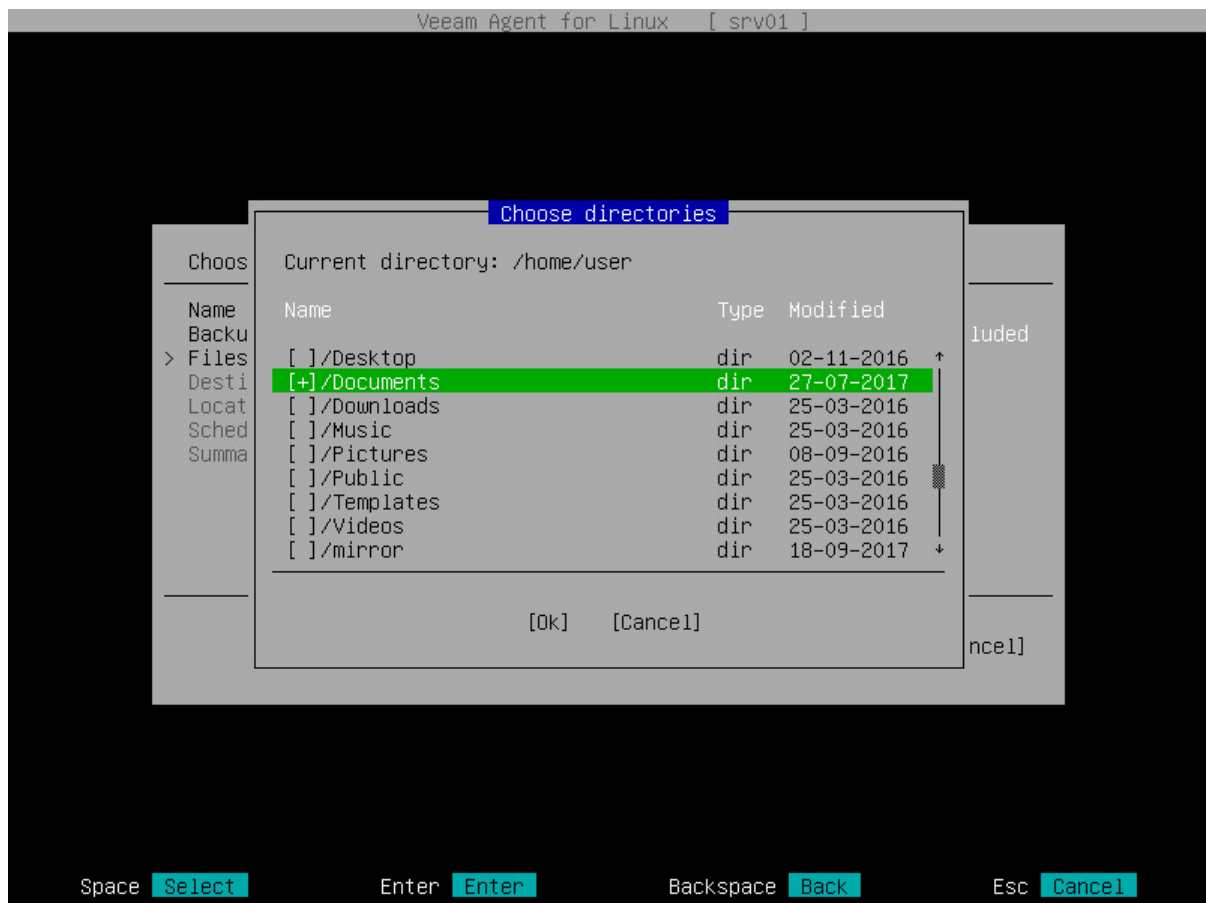
### IMPORTANT

Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.

- To navigate the list of directories, use the [Up] and [Down] keys.
- To browse for subdirectories, navigate to the necessary directory and press [Enter].

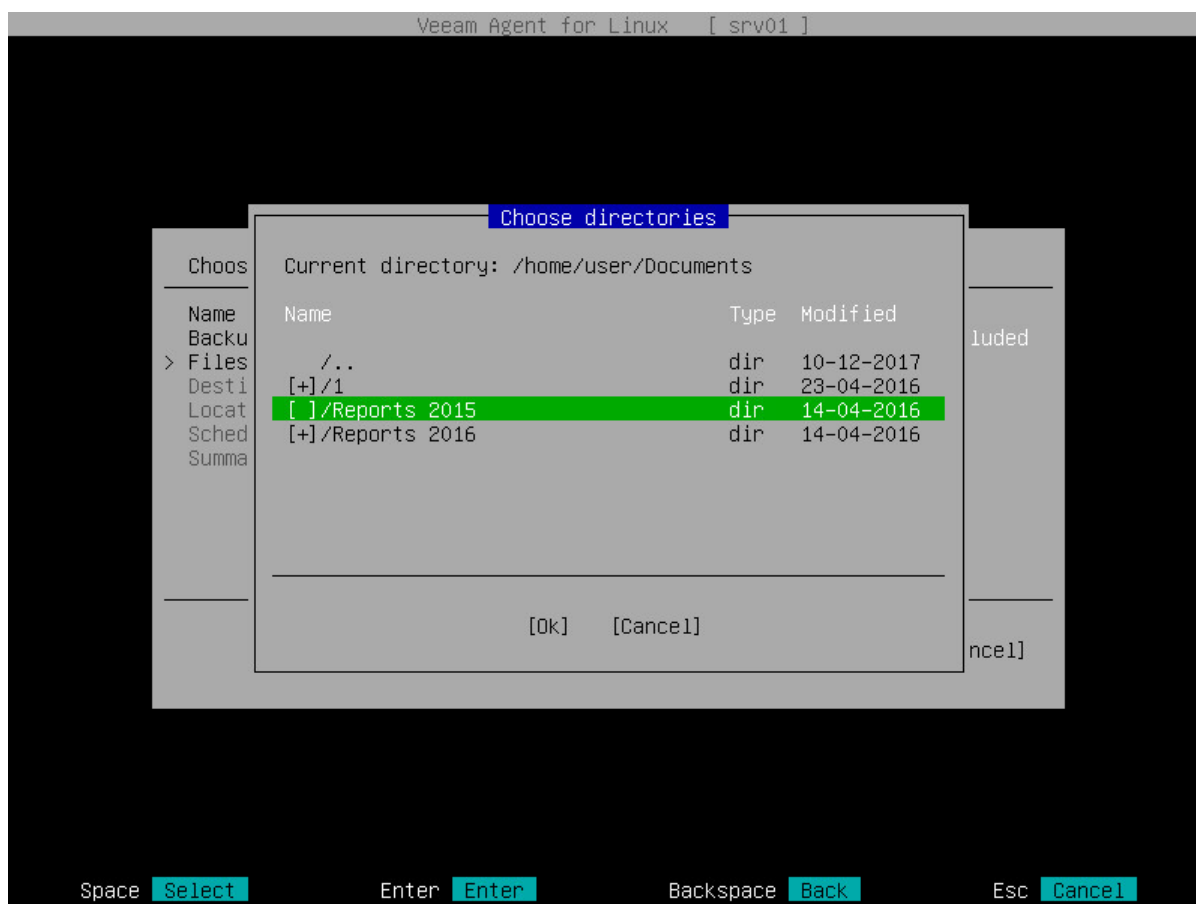


- To include a directory in the backup, navigate to the necessary directory and press [Space]. The included directory will be marked with the '+' character. All subdirectories of the selected directory will be included in the backup too.



3. Specify directories that you want to exclude from the file-level backup. To exclude a directory:
  - a. Browse for subdirectories of a directory that you have included in the backup.

- b. Navigate to the directory that you want to exclude from the backup and press [Space]. The excluded directory will not be marked with the '+' character.




4. Switch to the **OK** button and press [Enter]. Veeam Agent will display a list of paths to the selected directories and the number of excluded subdirectories for each directory in the list.
5. Specify file name masks for files that you want to include or exclude in/from the backup:
  - a. Select the **File Masks** option with the [Tab] key and press [Enter].
  - b. In the **File masks** window, make sure that the **Create Mask** button is selected and press [Enter].
  - c. In the **Mask** field, enter the file name mask, for example, `report.pdf`, `*filename*` or `*.odt`.  
Keep in mind that you must specify all names with masks in double quotation marks ("").
  - d. In the **Type** field, select one of the following options:
    - **Exclude** – if you do not want to back up files whose names match the specified mask. Veeam Agent will back up all files in the directories selected for backup except for such files.
    - **Include** – if you want to back up files whose names match the specified mask. Veeam Agent will create a backup only for such files in the directories selected for backup.

You can use a combination of include and exclude masks. Keep in mind that exclude masks have a higher priority than include masks. For example, you can specify masks in the following way:

- Include mask: `report*.*`

- Exclude mask: \*.odt


Veeam Agent will include in the backup all files whose name begins with `report` except for the files of the ODT format.



- Press [Enter]. Veeam Agent will display in the **File masks** window the specified file mask and its type: *Include* or *Exclude*.
- Repeat Steps **b** to **e** for each mask that you want to specify.
- After you specify all file masks, switch to the **OK** button and press [Enter].

## TIP

To remove a file name mask, in the **File masks** window, select the necessary mask and press [Delete].



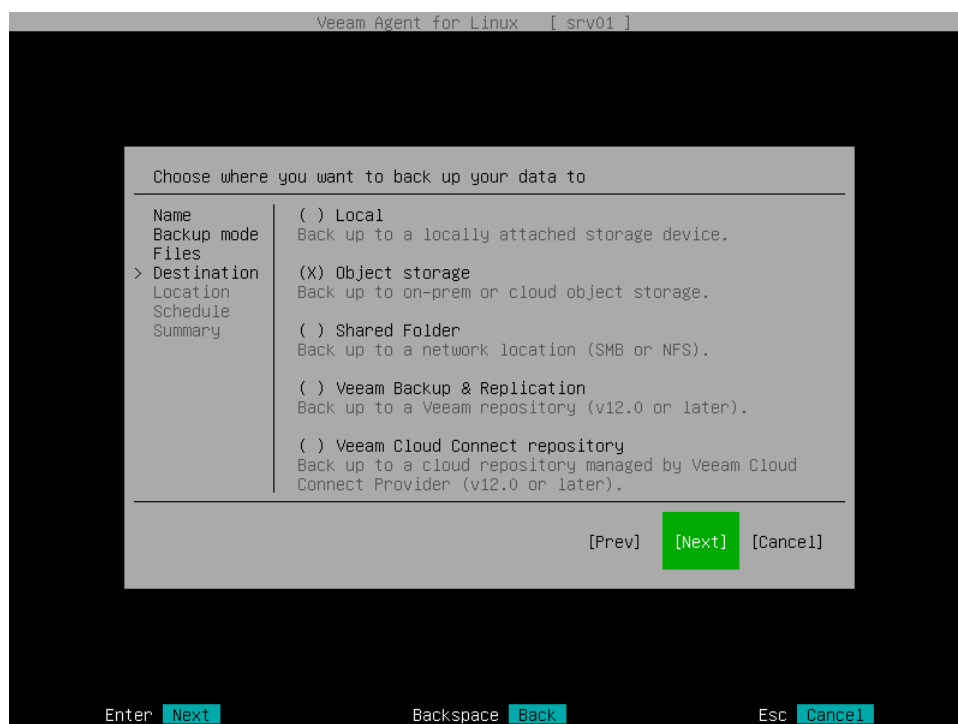
## Step 5. Select Backup Destination

At the **Destination** step of the wizard, select a target location for the created backup.

You can select one of the following options:

- **Local** – select this option if you want to save the backup in a removable storage device attached to the computer or on a local computer drive. With this option selected, you will pass to the [Location](#) step of the wizard.
- **Object storage** – select this option if you want to create the backup in an object storage exposed to you by cloud service provider. With this option selected, you will pass to the [Storage](#) step of the wizard.
- **Shared Folder** – select this option if you want to save the backup in a network shared folder. With this option selected, you will pass to the [Network](#) step of the wizard.
- **Veeam Backup & Replication** – select this option if you want to save the backup in a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Veeam](#) step of the wizard.
- **Veeam Cloud Connect repository** – select this option if you want to create the backup in a cloud repository exposed to you by the Veeam Cloud Connect service provider. With this option selected, you will pass to the [Service Provider](#) step of the wizard.

It is recommended that you store backups in the external location like USB storage device or network shared folder. You can also keep your backup files on the separate non-system local drive.



## Step 6. Specify Backup Storage Settings

Specify backup storage settings for the backup job:

- [Local storage settings](#) – if you have selected the **Local storage** option at the [Destination](#) step of the wizard.
- [Object storage settings](#) – if you have selected the **Object storage** option at the [Destination](#) step of the wizard.
- [Shared folder settings](#) – if you have selected the **Shared folder** option at the [Destination](#) step of the wizard.
- [Veeam backup repository settings](#) – if you have selected the **Veeam backup repository** option at the [Destination](#) step of the wizard.
- [Veeam Cloud Connect repository settings](#) – if you have selected the **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.

### NOTE

The **Veeam Cloud Connect repository** option is available if Veeam Agent operates in the Workstation or Server edition.

## Local Storage Settings

The **Location** step of the wizard is available if you have selected the **Local** option at the [Destination](#) step of the wizard. Specify location for the backup file and retention policy for the backup job:

1. To specify location for the backup file, browse to the directory where backup files must be saved:
  - a. Select the **Browse** option with the [Tab] key and press [Space] or [Enter].
  - b. In the **Choose backup location** window, select the necessary directory and press [Enter].
  - c. Repeat the step 'b' until a path to the directory in which you want to save backup files appears in the **Current directory** field.
  - d. To create a new directory, switch to the **Create Dir** button, press [Enter], then type a name for the new directory and press [Enter].
  - e. Switch to the **OK** button and press [Enter]. Veeam Agent will display the path to the specified directory in the **Location** field.

Alternatively, you can type a path to the directory in which you want to save backup files in the **Location** field.

After you specify location for the backup, Veeam Agent will display the following information on the volume where the directory selected for backup storage resides:

- **Space** – total size of the volume on which the selected directory resides.
  - **Free** – free space on the volume where the selected directory resides.
  - **Type** – filesystem type of the volume on which the selected directory resides.
2. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).

3. In the **Restore points** field, specify the number of backup files that you want to keep in the target location. By default, Veeam Agent keeps 7 latest backup files. When the number of restore points is exceeded, Veeam Agent for Linux will remove the earliest restore point from the backup chain.

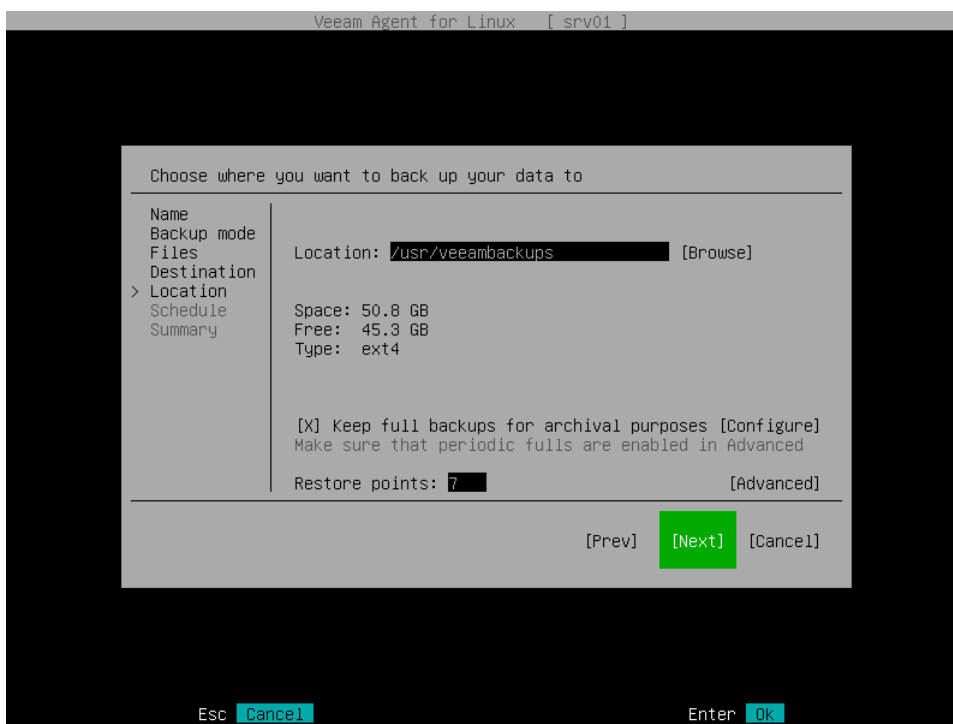
To learn more, see the [Short-Term Retention Policy](#).

4. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

## IMPORTANT

Consider the following:

- The backup location must reside on a separate volume from a volume whose data you plan to back up.
- USB storage devices formatted as FAT32 do not allow storing files larger than 4 GB in size. For this reason, it is recommended that you do not use such USB storage devices as a backup target.



## Object Storage Settings

The **Cloud Type** step of the wizard is available if you have selected the **Object storage** option at the [Destination](#) step of the wizard.

At the **Storage** step of the wizard, select the object storage. You can select one of the following options:

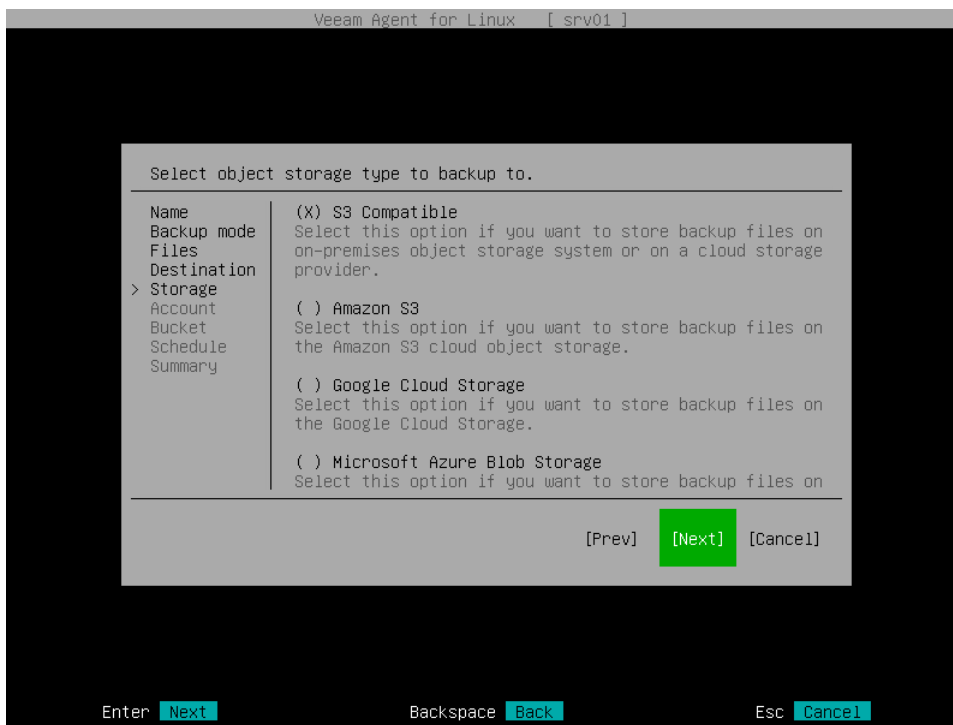
- **S3 compatible** – select this option if you want to create a backup in the S3 compatible storage. With this option selected, you will pass to the [Account](#) step of the wizard.

### TIP

If you plan to store backups in an IBM or Wasabcloud storage, use the **S3 compatible** option.

- **Amazon S3** – select this option if you want to create a backup in the Amazon S3 storage. With this option selected, you will pass to the [Account](#) step of the wizard.

- **Google Cloud storage** – select this option if you want to create a backup in the Google Cloud storage. With this option selected, you will pass to the [Account](#) step of the wizard.
- **Microsoft Azure Blob storage** – select this option if you want to create a backup in the Microsoft Azure storage. With this option selected, you will pass to the [Account](#) step of the wizard.



## S3 Compatible Settings

If you have selected to store backup files in the S3 compatible storage, specify the following settings:

1. [Account settings](#).
2. [Bucket settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the S3 compatible storage.

### NOTE

You can store backups only in the S3 compatible storage repositories that are accessible over the HTTPs protocol.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.


### NOTE

If you want to connect to the repository using an IPv6 address and port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is the IPv6 address of the object storage.
- `port` is the number of the port that Veeam Agent will use to connect to the object storage.



2. In the **Region** field, specify the storage region based on your regulatory and compliance requirements.
3. In the **Access key** field, enter the access key ID.
4. In the **Secret key** field, enter the secret access key.



## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in the S3 compatible storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

5. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Bucket** window, select the necessary bucket and press [Enter].
6. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Folder** window, select the necessary folder and press [Enter].

### TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

7. To prohibit modification and deletion of blocks of data in the object storage repository, select the **Make recent backups immutable for** check box and specify the immutability period in days. For more information, see [Backup Immutability](#).
8. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).

## NOTE

If you use the GFS retention scheme and enable immutability for the backup, the restore points with GFS flags will become immutable for the whole GFS retention period. You will not be able to delete such restore points until the GFS retention period is over.

5. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

6. Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).

The screenshot shows the 'Specify S3 compatible storage bucket to use' dialog box within the Veeam Agent for Linux configuration interface. The dialog has a sidebar on the left with a tree view containing: Name, Backup mode, Files, Destination, Storage, Account, > Bucket (selected), Schedule, and Summary. The main area contains the following fields and options:

- Bucket:** veeam-backups [Browse]
- Folder:** folder01 [Browse]
- ☐ Make backups immutable for 30 days (increases cost)  
Protects recent backups from modification or deletion by ransomware, malicious insiders, or hackers.
- ☒ Keep full backups for archival purposes [Configure]  
Make sure that periodic fulls are enabled in Advanced
- Restore points:** 7 [Advanced]

At the bottom of the dialog are three buttons: [Prev], [Next] (highlighted in green), and [Cancel]. Below the dialog, at the bottom of the terminal window, are three keyboard shortcuts: Enter Next, Backspace Back, and Esc Cancel.

After that, Veeam Agent will create a new repository in the object storage where you can store backups.

## Amazon S3 Settings

If you have selected to store backup files in the Amazon S3 storage, specify the following settings:

1. [Account settings](#).
2. [Bucket settings](#).


## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the Amazon S3 storage.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter the access key ID.
2. In the **Secret key** field, enter the secret access key.

3. In the **AWS region** window, select the AWS region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Global** region. Switch to the **Ok** button and press [Enter].



## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in the Amazon S3 storage and specified account settings to connect to the storage.

### IMPORTANT

You must create the bucket where you want to store your backup data beforehand. When you create a bucket, consider Amazon bucket naming rules. For example, it is not recommended that you use dots (.) in the bucket name. For more information on bucket naming rules, see [this AWS documentation article](#).

Specify settings for the bucket in the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups. Switch to the **Ok** button and press [Enter].
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Bucket** window, select the necessary bucket and press [Enter].
3. In the **Folder** field, specify the folder in the bucket:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Folder** window, select the necessary folder and press [Enter].

### TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

- To prohibit modification and deletion of blocks of data in the object storage repository, select the **Make recent backups immutable for** check box and specify the immutability period in days. For more information, see [Backup Immutability](#).
- To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).


#### NOTE

If you use the GFS retention scheme and enable immutability for the backup, the restore points with GFS flags will become immutable for the whole GFS retention period. You will not be able to delete such restore points until the GFS retention period is over.

- In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

- Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).



After that, Veeam Agent will create a new repository in the object storage where you can store backups.

## Google Cloud Storage Settings

If you have selected to store backup files in a Google Cloud storage repository, specify the following settings:

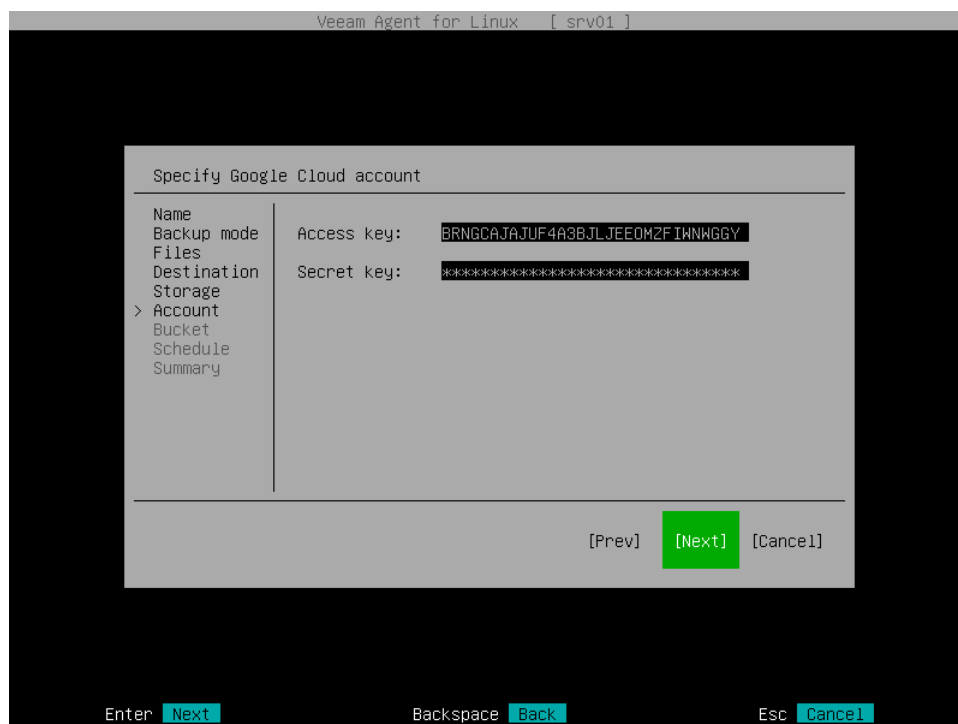
- [Account settings](#).
- [Bucket settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the Google Cloud storage.

To connect to the Google Cloud storage, in the **Access Key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Agent will use the HMAC key to authenticate requests to the Google Cloud storage. For more information about the Google Cloud accounts, see the [Google Cloud documentation](#).

If you have not created the HMAC key beforehand, you can create the key in the Google Cloud console, as described in [this Google Cloud documentation article](#).



## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in the Google Cloud storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups. Switch to the **Ok** button and press [Enter].
2. In the **Bucket** field, specify the bucket in the storage:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Bucket** window, select the necessary bucket and press [Enter].
3. In the **Folder** field, specify the folder in the bucket:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Folder** window, select the necessary folder and press [Enter].

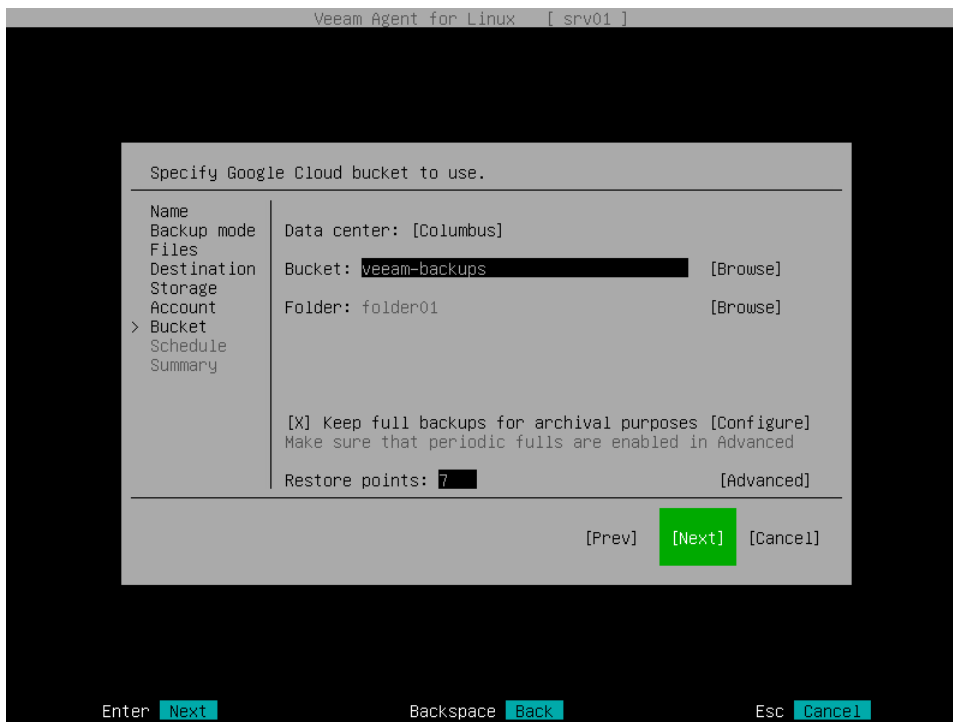
## TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

4. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).
5. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

6. Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).



After that, Veeam Agent will create a new repository in the object storage where you can store backups.

## Microsoft Azure Storage Settings

If you have selected to store backup files in the Microsoft Azure storage, specify settings to connect to the storage and container in this storage:

1. [Account settings](#).
2. [Container settings](#).


## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the Microsoft Azure storage.

## NOTE

The **Allow storage account key access** option for Shared Key authorization must be enabled in the storage account. To learn how to find this option, see [this Microsoft Docs article](#).

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.
3. In the **Azure Region** window, select the Microsoft Azure region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Azure Global (Standard)** region. Switch to the **Ok** button and press [Enter].



## Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to save backup files in the Microsoft Azure storage and specified account settings to connect to the storage.

Specify settings for the container in the storage:

1. In the **Container** field, specify the container in the storage:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Azure Container** window, select the necessary container and press [Enter].
2. In the **Folder** field, specify the folder in the container:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Specify Folder** window, select the necessary folder and press [Enter].

## TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

3. To prohibit modification and deletion of blocks of data in the object storage repository, select the **Make recent backups immutable for** check box and specify the immutability period in days. For more information, see [Backup Immutability](#).
4. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).

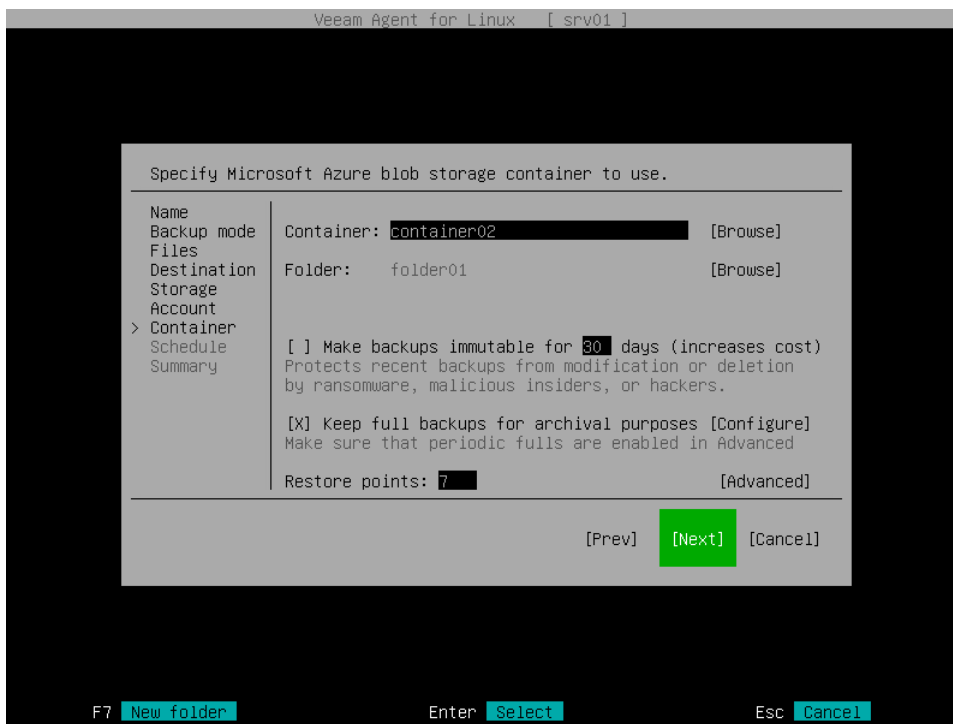
#### NOTE

If you use the GFS retention scheme and enable immutability for the backup, the restore points with GFS flags will become immutable for the whole GFS retention period. You will not be able to delete such restore points until the GFS retention period is over.

5. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

6. Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).



After that, Veeam Agent will create a new repository in the object storage where you can store backups.

## Shared Folder Settings

The **Network** step of the wizard is available if you have selected the **Shared Folder** option at the [Destination](#) step of the wizard.

To save backup files in a remote network location, Veeam Agent mounts to the local file system of your computer the network shared folder that you specify as a location for the backup. When you specify the network shared folder settings, Veeam Agent saves information about the network shared folder and its mount point in the database.



You do not need to mount the network shared folder in advance before every backup job run. Veeam Agent will do it automatically when the backup job is started manually or upon schedule.

After the backup job completes, Veeam Agent will automatically unmount the network shared folder.

Specify shared folder settings:

1. Select the type of a network shared folder:
  - **NFS** – to connect to a network shared folder using the NFS protocol.
  - **SMB** – to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Server** field, type the IP address or domain name of the server.
3. In the **Folder** field, type the name of the network shared folder in which you want to store backup files.

Every time the backup job starts, Veeam Agent will automatically mount the specified network shared folder to the `/tmp/veeam` directory in the computer file system. After the backup job completes, Veeam Agent will unmount the network shared folder.
4. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
5. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
6. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.
7. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see [Specify GFS Retention Settings](#).
8. In the **Restore points** field, specify the number of backup files that you want to keep in the target location. By default, Veeam Agent keeps 7 latest backup files. When the number of restore points is exceeded, Veeam Agent will remove the earliest restore point from the backup chain.

To learn more, see the [Short-Term Retention Policy](#).

9. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

The screenshot shows the 'Specify a network location to backup to' dialog box in the Veeam Agent for Linux configuration utility. The window title is 'Veeam Agent for Linux [ srv01 ]'. On the left is a sidebar with a tree view containing 'Name', 'Backup mode', 'Files', 'Destination', 'Network' (selected), 'Schedule', and 'Summary'. The main area contains the following fields and options:

- Name:** ( ) NFS, (X) SMB
- Backup mode:** ( ) NFS, (X) SMB
- Files:** (empty text field)
- Destination:** Server: 172.25.164.25
- Network:** Folder: VeeamBackups
- Schedule:** Domain: tech, Username: richard.olson
- Summary:** Password: (masked with asterisks)
- Advanced options:** [X] Keep full backups for archival purposes [Configure], Make sure that periodic fulls are enabled in Advanced
- Restore points:** 7 [Advanced]

At the bottom right are buttons for [Prev], [Next] (highlighted in green), and [Cancel]. At the bottom left are 'Esc' and 'Cancel' buttons, and at the bottom right are 'Enter' and 'Ok' buttons.

## Veeam Backup Repository Settings

If you have selected to store backup files on a Veeam Backup & Replication repository, specify settings to connect to the backup repository:

1. [Specify backup server settings](#).
2. [Select the Veeam backup repository](#).

## Specifying Backup Server Settings

The **Veeam** step of the wizard is available if you have chosen to store backup files on a Veeam Backup & Replication repository.

Specify settings for the Veeam backup server that manages the target backup repository:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
4. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.

5. In the **Password** field, type a password of the account that has access to the Veeam backup repository.

Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

The screenshot shows a terminal window titled 'Veeam Agent for Linux [ srv01 ]'. Inside, a configuration wizard is running. The current step is 'Specify a Veeam Backup & Replication server to backup to'. On the left, a sidebar lists the wizard steps: Name, Backup mode, Destination, > Veeam, Repository, Schedule, and Summary. The main area contains fields for: Address (172.24.31.136), Port (10006), Login (Administrator), Domain (empty), and Password (masked with asterisks). At the bottom of the wizard are buttons for [Prev], [Next] (highlighted in green), and [Cancel]. At the very bottom of the terminal window, there are keyboard shortcuts: Enter Next, Backspace Back, and Esc Cancel.

## Selecting Backup Repository

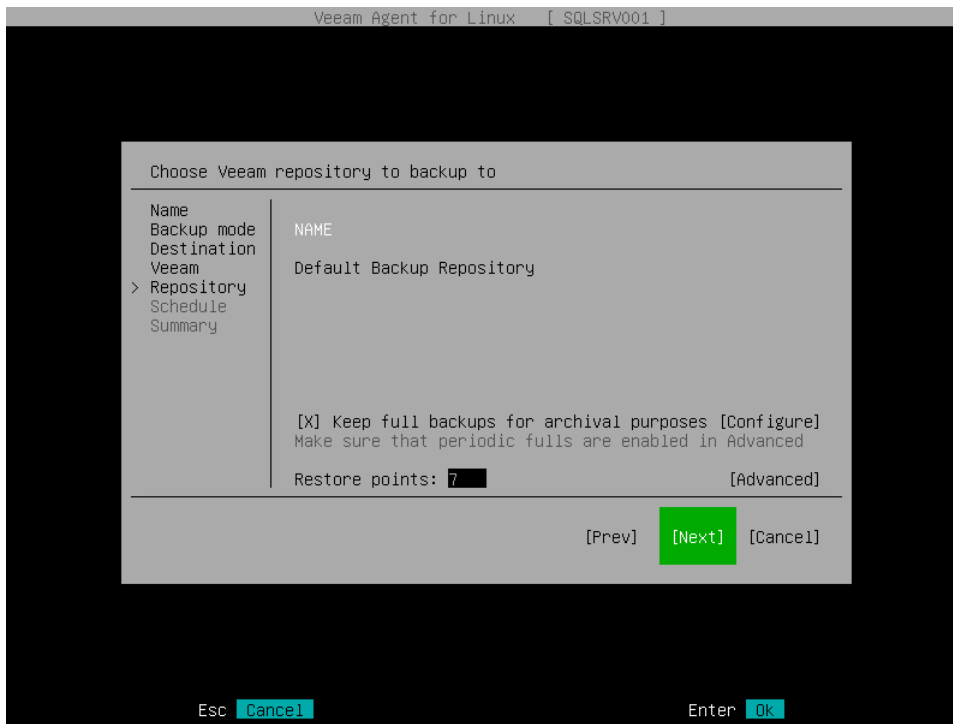
The **Repository** step of the wizard is available if you have chosen to save backup files on a Veeam Backup & Replication repository.

Specify settings for the target backup repository:

1. From the list of available backup repositories, select a backup repository where you want to store backups. The list of backup repositories displays only those repositories on which you have permissions to store data. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
2. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).
3. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

4. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).



## Veeam Cloud Connect Repository Settings

If you have selected to store backup files on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings](#).
2. [Verify the TLS certificate and specify user account settings](#).
3. [Select the cloud repository](#).

### NOTE

The **Veeam Cloud Connect repository** option is available if Veeam Agent operates in the Workstation or Server edition.

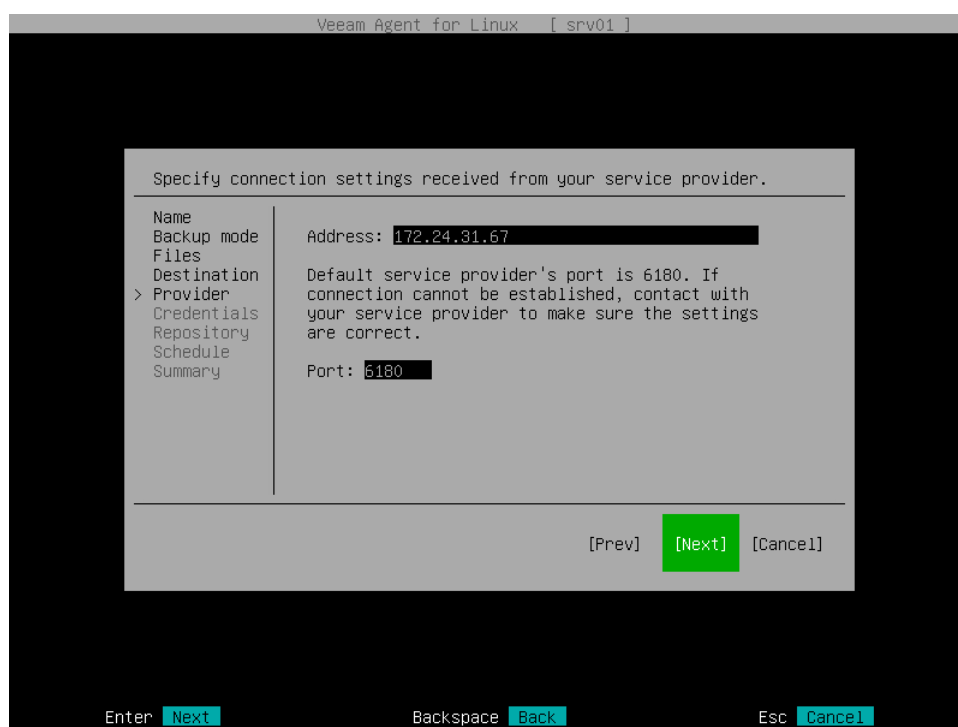
## Specifying Service Provider Settings

The **Provider** step of the wizard is available if you have chosen to save backup files on a Veeam Cloud Connect repository.

Specify settings for the cloud gateway that the Veeam Cloud Connect service provider (SP) or your backup administrator has provided to you:

1. In the **Address** field, enter a full DNS name or IP address of the cloud gateway.

2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.



## Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to save backup files in a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the cloud repository.

1. In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate:
  - [Optional] To verify the TLS certificate with a thumbprint, do the following:
    - i. Select the **Verify thumbprint** button with the [Tab] key and press [Enter].
    - ii. Copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Thumbprint verification** field.
    - iii. Switch to the **Verify** button and press [Enter]. Veeam Agent will check if the thumbprint you entered matches the thumbprint of the obtained TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

  - To accept the TLS certificate, select the **Accept** button with the [Tab] key and press [Enter].
2. In the **Username** field, enter the name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.

The screenshot shows a terminal window titled "Veeam Agent for Linux [ srv01 ]". Inside, a configuration wizard is running. The main window has a title bar "Specify credentials, and validate the certificate." and a sidebar on the left with the following menu items: Name, Backup mode, Files, Destination, Provider, > Credentials (highlighted), Repository, Schedule, and Summary. The main area contains a "Username:" field with the value "ABC\_Company\User\_02" and a "Password:" field with masked characters "\*\*\*\*\*". Below these fields is a "[View certificate]" link. At the bottom of the main area are three buttons: "[Prev]", "[Next]" (highlighted in green), and "[Cancel]". At the very bottom of the terminal window, there are three keyboard shortcuts: "Enter [Next]", "Backspace [Back]", and "Esc [Cancel]".

## Selecting Cloud Repository


The **Repository** step of the wizard is available if you have chosen to save backup files on a cloud repository and specified settings to connect to the SP.

Specify settings for the cloud repository:

1. From the **Repository** list, select a cloud repository where you want to store created backups. The **Repository** list displays only those cloud repositories that can be accessed by the tenant or subtenant account that you use to connect to the service provider.
2. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).
3. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

4. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).



## Step 7. Specify GFS Retention Settings

This step of the wizard is available if you have chosen to use a long-term, or Grandfather-Father-Son (GFS), retention policy.

To configure GFS retention policy settings for the backup job:

1. Select the **Keep full backups for archival purposes** option and click **Configure** at one of the following steps of the wizard:
  - **Location** — if you have selected the **Local storage** option at the **Destination** step of the wizard.
  - **Network** — if you have selected the **Shared folder** option at the **Destination** step of the wizard.
  - **Repository** — if you have selected the **Veeam backup repository** option at the **Destination** step of the wizard.
  - **Repository** — if you have selected the **Veeam Cloud Connect repository** option at the **Destination** step of the wizard.
  - **Bucket** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **S3 compatible** option at the **Storage** step of the wizard.
  - **Bucket** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **Amazon S3** option at the **Storage** step of the wizard.
  - **Bucket** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **Google Cloud storage** option at the **Storage** step of the wizard.
  - **Container** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **Microsoft Azure Blob storage** option at the **Storage** step of the wizard.
2. In the **Configure GFS** window, do the following:
  - a. If you want to create weekly restore points for archival purposes, select the **Keep weekly full backups for** check box. Then specify the number of weeks during which you want to prevent restore points from being modified and deleted.

In the **If multiple full backups exist, use the one from** list, select a week day when Veeam Agent must assign the weekly GFS flag to a full restore point.
  - b. If you want to create monthly restore points for archival purposes, select the **Keep monthly full backups for** check box. Then specify the number of months during which you want to prevent restore points from being modified and deleted.

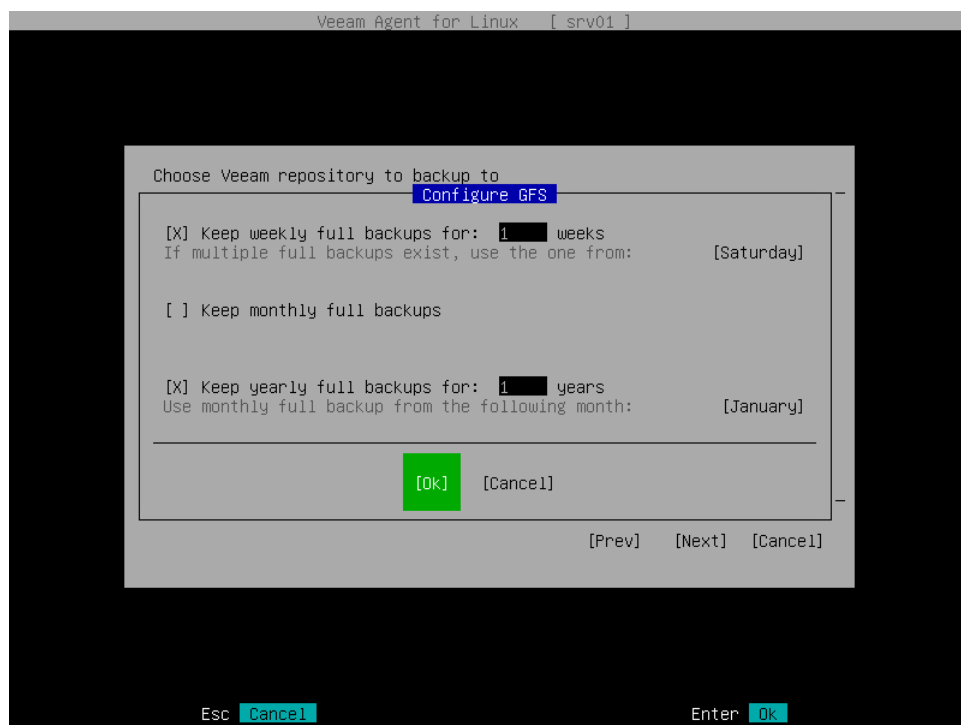
In the **Use weekly full backup for the following week of a month** list, select a week when Veeam Agent must assign the monthly GFS flag to a full restore point. A week equals 7 calendar days; for example, the first week of May is days 1–7, and the last week of May is days 25–31.
  - c. If you want to create yearly restore points for archival purposes, select the **Keep yearly full backups for** check box. Then specify the number of years during which you want to prevent restore points from being modified and deleted.

In the **Use monthly full backup for the following month** list, select a month when Veeam Agent must assign the yearly GFS flag to a full restore point.



## NOTE

- If you select to assign multiple types of GFS flags, the flags begin to depend on each other. For more information on this dependency, see [Assignment of GFS Flags](#) section in the Veeam Backup & Replication User Guide.
- To use a GFS retention policy, you must set Veeam Agent to create full backups. To learn more, see [Active Full Backup Settings](#).



## Step 8. Specify Advanced Backup Settings

To configure advanced settings for the backup job, select **Advanced** at one of the following steps of the wizard:

- [Location](#) — if you have selected the Local storage option at the [Destination](#) step of the wizard.
- [Network](#) — if you have selected the Shared folder option at the [Destination](#) step of the wizard.
- [Repository](#) — if you have selected the Veeam backup repository option at the [Destination](#) step of the wizard.
- [Repository](#) — if you have selected the Veeam Cloud Connect repository option at the [Destination](#) step of the wizard.
- [Bucket](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the S3 compatible option at the Storage step of the wizard.
- [Bucket](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the Amazon S3 option at the Storage step of the wizard.
- [Bucket](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the Google Cloud storage option at the Storage step of the wizard.
- [Container](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the Microsoft Azure Blob storage option at the Storage step of the wizard.

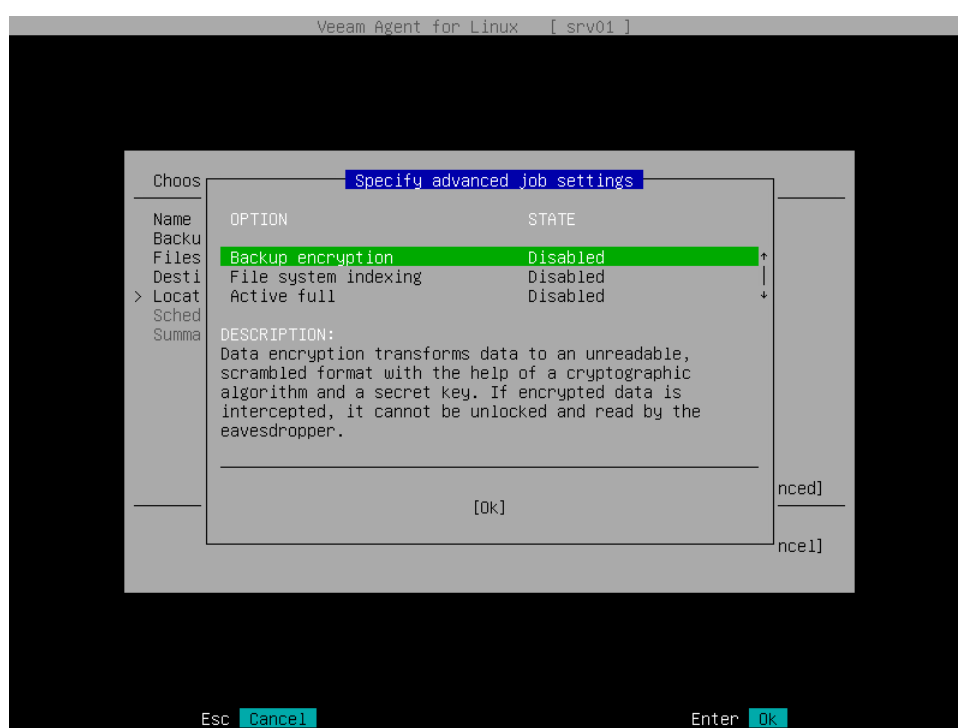
In the **Specify advanced job settings** window, specify advanced settings for the backup job:

- [Data encryption settings](#)
- [File indexing settings](#)
- [Oracle database system processing settings](#)
- [MySQL database system processing settings](#)
- [PostgreSQL database system processing settings](#)
- [Active full backup settings](#)
- [Backup maintenance settings](#)
- [Script settings](#)
- [Health check settings](#)

## NOTE

Consider the following:

- You cannot specify encryption settings for the backup job if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the [Data Encryption](#) section in the Veeam Backup & Replication User Guide.
- You can specify file indexing settings only if Veeam Agent operates in the Workstation or Server edition.
- Nosnap Veeam Agent for Linux and nosnap Veeam Agent for Linux on Power do not support application-aware processing and cannot be used to back up database systems.
- You can specify settings for Oracle, MySQL or PostgreSQL database system processing only if Veeam Agent operates in the Server edition. The settings are available for a volume-level backup job only.
- You can specify backup maintenance settings only if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.
- You can specify backup health check settings only if you have selected the **Object storage repository** option at the [Destination](#) step of the wizard.
- You cannot specify data compression settings when you configure a backup job with the Backup Job wizard. If you want to specify these settings, consider creating the backup job with the Veeam Agent command line interface. To learn more, see [Advanced Backup Job Settings](#).



## Data Encryption Settings

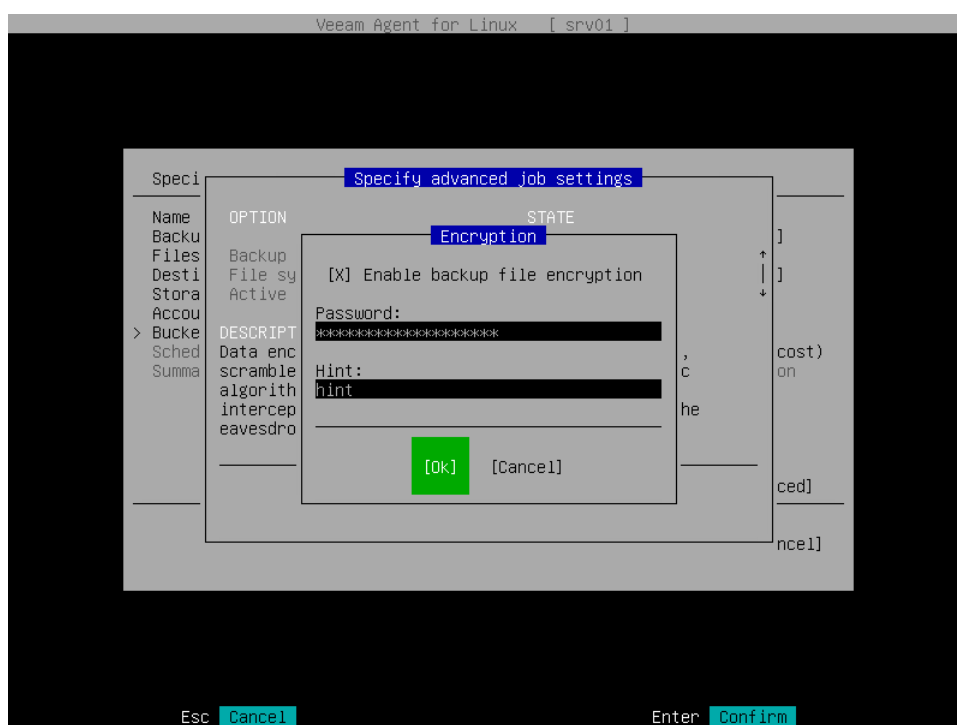
If you want to encrypt the content of backup files, specify data encryption settings for the backup job:

1. In the **Specify advanced job settings** window, select the **Backup encryption** option with the [Tab] key and press [Enter].

## NOTE

The **Backup encryption** option is unavailable if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the [Data Encryption](#) section in the Veeam Backup & Replication User Guide.

2. In the **Encryption** window, make sure that the **Enable backup file encryption** option is selected and press [Space].
3. In the **Password** field, type a password that you want to use for encryption.
4. In the **Hint** field, type a hint for the password. In case you lose the password, the specified hint will help you to remember the lost password.
5. Switch to the **Ok** button and press [Enter].



## File Indexing Settings

To specify file indexing settings for the backup job, do the following:

1. In the **Specify advanced job settings** window, select the **File system indexing** option with the [Tab] and [Down] keys and press [Enter].

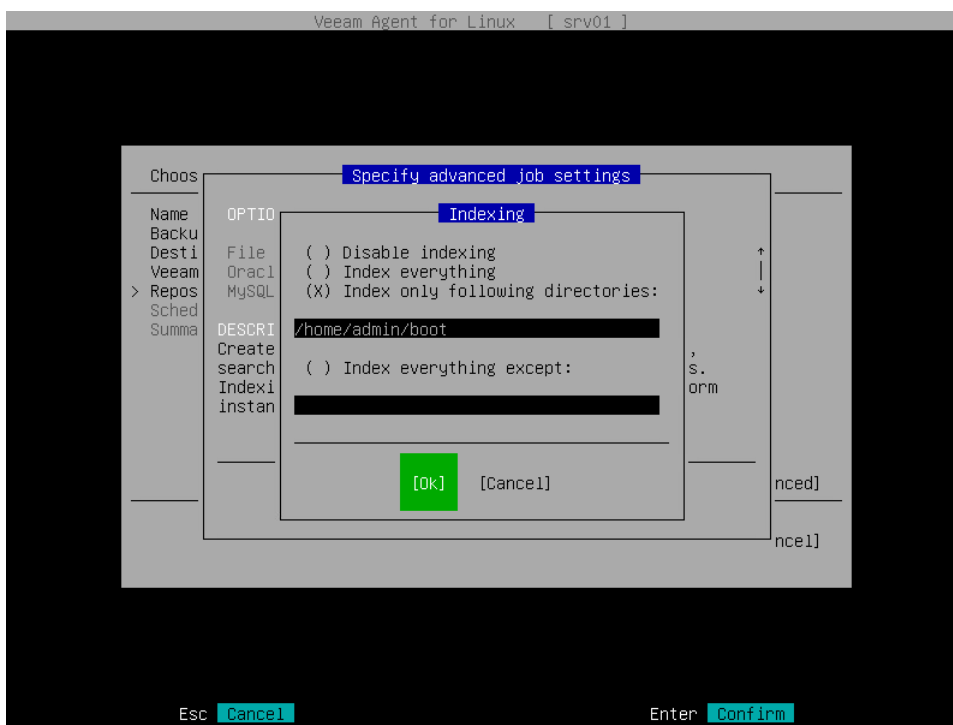
## NOTE

The **File system indexing** option is available if Veeam Agent for Linux operates in the Workstation or Server edition.

2. In the **Indexing** window, specify the indexing scope:
  - Select **Index everything** if you want to index all files within the backup scope that you have specified at the [Backup mode](#) step of the wizard. Veeam Agent for Linux will index all files that reside:

- On your computer OS (for entire machine backup)
  - On the volumes that you have selected for backup (for volume-level backup)
  - In the directories that you have selected for backup (for file-level backup)
- [For entire machine and volume-level backups] Select **Index only following directories** to define directories that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character.
  - [For entire machine and volume-level backups] Select **Index everything except** if you want to index all files within the specified backup scope except those files that reside in specific directories. Enter paths to directories whose files you do not want to index. To separate several paths, use the ',' (comma) character.

3. Switch to the **Ok** button and press [Enter].



## Oracle Database Processing Settings

To specify processing settings for the Oracle database system, do the following

1. In the **Specify advanced job settings** window, select the **Oracle processing** option with the [Tab] and [Down] keys and press [Enter].
2. In the **Oracle processing** section, select one of the following options:
  - **Require successful processing.** With this option selected, Veeam Agent will stop the backup process if an error occurs while processing the Oracle database system.
  - **Try application processing, ignore failures.** With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the Oracle database system.

3. In the **Archived logs processing** section, specify how Veeam Agent will process archived logs on the Oracle database:

- Select **Do not delete archived logs** if you want Veeam Agent to keep archived logs. When the backup job completes, Veeam Agent will not delete archived logs.


It is recommended that you select this option when you do not have databases running in the ARCHIVELOG mode. If the database is running in the ARCHIVELOG mode, archived logs may grow large and consume all disk space. In this case, the database administrator must take care of archived logs themselves.

- Select **Delete logs older than <N> hours** or **Delete oldest logs exceeding <N> GB** if you want Veeam Agent to delete archived logs that are older than <N> hours or larger than <N> GB. Veeam Agent will wait for the backup job to complete successfully and then trigger archived logs truncation via Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next successful backup job session.

#### TIP

If you configure backup job to back up archived logs, Veeam Agent for Linux will not trigger archived logs deletion after each log backup job session. To prevent Oracle database logs from overgrowing, run the backup job for the Veeam Agent computer more often.

4. In the **Specify Oracle account with SYSDBA privileges** section, specify which account type Veeam Agent will use to connect to the database system.
- Select **System account** if you want Veeam Agent to use an account of the Veeam Agent machine OS. The account must be a member of the group that owns configuration files for the Oracle database (for example, the install group).
  - Select **Oracle account** if you want Veeam Agent to use an Oracle account. The account must have SYSDBA rights.



# MySQL Database Processing Settings

## IMPORTANT

MySQL tables that use the MyISAM storage engine must be locked to keep them in consistent state while Veeam Agent is creating the system snapshot. To correctly process such tables, MySQL account must have the following instance-wide privileges:

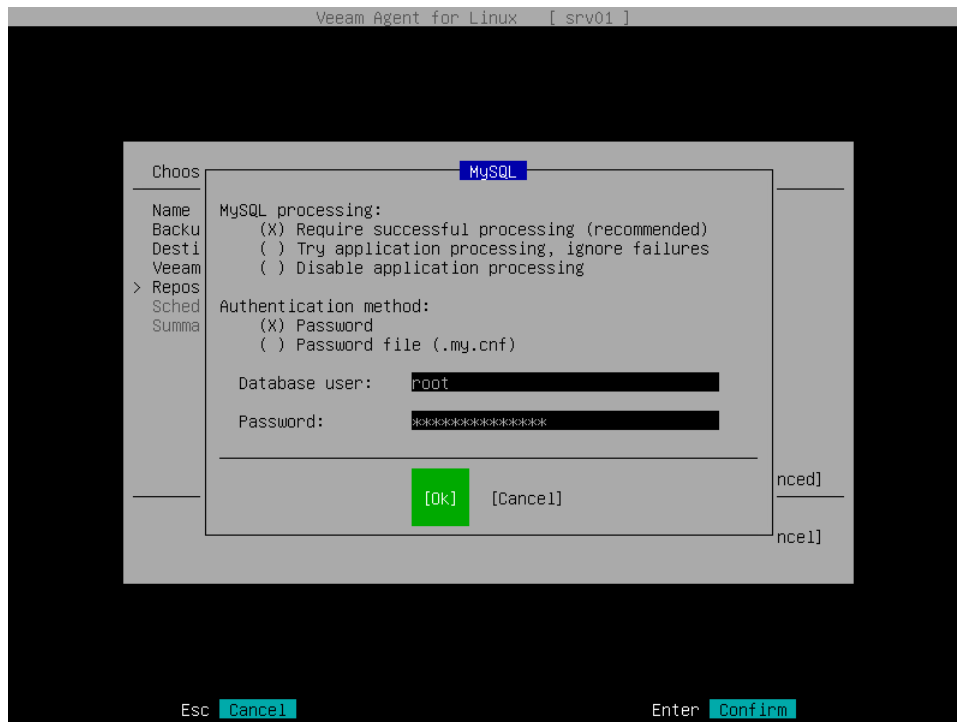
- **SELECT.** This privilege enables Veeam Agent to access tables' metadata and select for a lock the tables that use the MyISAM storage engine. Without this privilege, the processing of the MySQL database system will run successfully but MyISAM tables will not be locked, which may result in an inconsistent state of the backed up data.
- **LOCK TABLES.** This privilege is required for locking the selected MyISAM tables. If some MyISAM tables are selected but the MySQL account does not have the **LOCK TABLES** privilege, the processing of the MySQL database system will fail.
- **RELOAD or FLUSH\_TABLES.** If some MyISAM tables are selected but the MySQL account does not have either **RELOAD** or **FLUSH\_TABLES** privilege, the processing of the MySQL database system will fail.

To obtain information about the privileges that are assigned to an account, use MySQL functionality, for example, the **SHOW GRANTS** statement. To learn more, see [MySQL documentation](#).

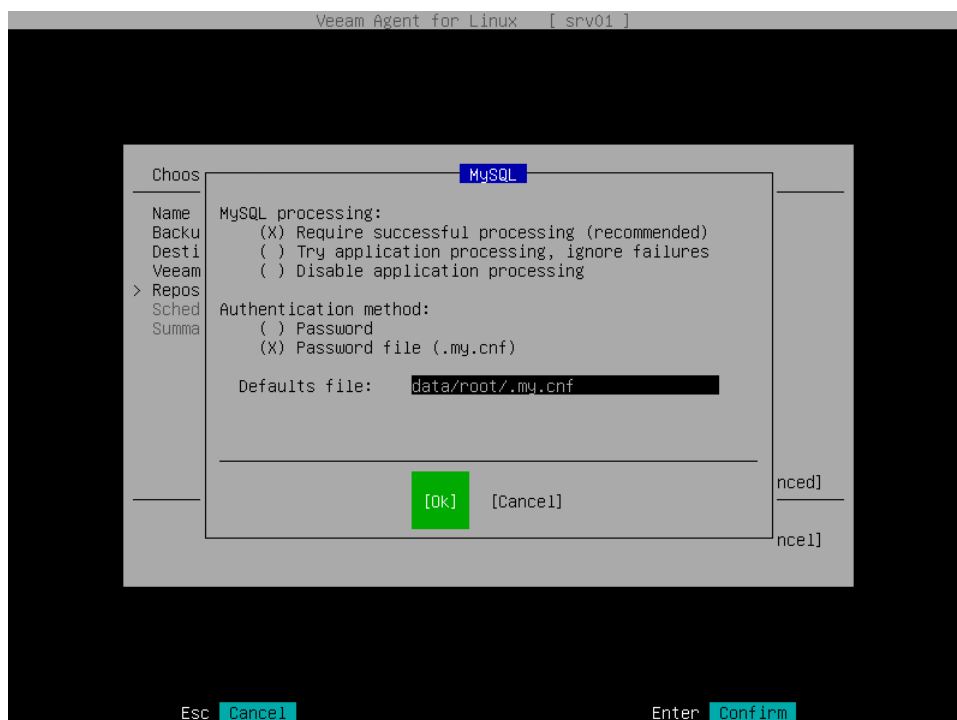
To specify processing settings for the MySQL database system, do the following:

1. In the **Specify advanced job settings** window, select the **MySQL processing** option with the [Tab] and [Down] keys and press [Enter].
2. In the **MySQL processing** section, select one of the following options:
  - **Require successful processing.** With this option selected, Veeam Agent will stop the backup process if an error occurs when processing the MySQL database system.
  - **Try application processing, ignore failures.** With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the MySQL database system.

3. In the **Authentication method** section, specify how Veeam Agent will connect to the MySQL database:
- Select **Password** if you want Veeam Agent to connect with the MySQL account name and password. With this option selected, you must specify account name and password in the backup job settings.



- Select **Password file** if you want Veeam Agent to connect with the MySQL account name and password that are stored in the `.my.cnf` password file. With this option selected, you must specify a path to the password file, but do not need to specify account credentials in the backup job settings. To learn more about password file configuration, see [Preparing Password File for MySQL Processing](#).





## Preparing Password File for MySQL Processing

You can use MySQL account credentials that are stored in the password file to connect Veeam Agent for Linux to the MySQL database system.

### NOTE

Consider the following:

- If you specify a custom path to the password file, specify a full path. Specifying relative paths is not supported.
- The password file can also contain user-specific connection settings that Veeam Agent will apply to connect to the MySQL database system. For example, if you want to connect to the MySQL database system using the custom socket, specify the socket path in the password file. To learn more, see [MySQL documentation](#).

If you want to use a password file for authentication, create a file. By default, Veeam Agent expects the password file to have the `.my.cnf` name and to be in the home directory of the `root` user. If the password file has a custom name or is stored in another directory, you can specify a custom path.

The password file must have the following contents:

```
[client]
user=<username>
password=<password>
```

where:

- `<username>` – name of the account that Veeam Agent will use to connect to the MySQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the MySQL database system.

For example:

```
[client]
user=root
password=P@ssw0rd
```

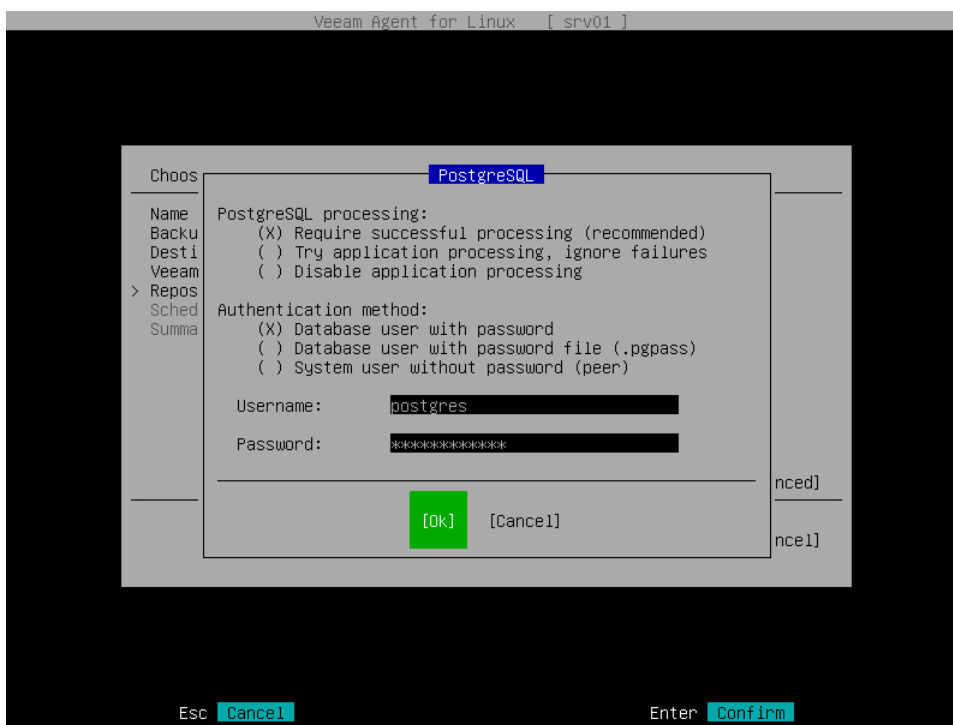
## PostgreSQL Database Processing Settings

To specify processing settings for the PostgreSQL database system, do the following:

1. In the **Specify advanced job settings** window, select the **PostgreSQL processing** option with the [Tab] and [Down] keys and press [Enter].
2. In the **PostgreSQL processing** section, select one of the following options:
  - **Require successful processing.** With this option selected, Veeam Agent will stop the backup process if an error occurs when processing the PostgreSQL database system.
  - **Try application processing, ignore failures.** With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the PostgreSQL database system.

3. In the **Authentication method** section, specify how Veeam Agent will connect to the PostgreSQL database:

- Select **Database user with password** if you want Veeam Agent to connect with the PostgreSQL account name and password. With this method selected, you must specify account name and password in the backup job settings.
- Select **Database user with password file** if you want Veeam Agent to connect with the PostgreSQL account password that is stored in the `.pgpass` password file. With this method selected, you must specify account name only in the backup job settings. To learn more about password file configuration, see [Password File for PostgreSQL](#).
- Select **System user without password** if you want Veeam Agent to connect using a peer authentication method. In the peer authentication method, Veeam Agent uses the OS account as the PostgreSQL database user name. With this option selected, you must specify OS account in the backup job settings. To learn more about peer authentication, see [PostgreSQL documentation](#).



## Preparing Password File for PostgreSQL Processing

You can use PostgreSQL account credentials that are stored in the password file to connect Veeam Agent to the PostgreSQL database system.

If you want to use a password file for authentication, create the `.pgpass` file in the home directory of the `root` user.

The password file must have the following contents:

```
<hostname>:<port>:<database>:<username>:<password>
```

where:

- `<hostname>` – name of the host where the PostgreSQL database system is located.

- `<port>` – number of the free port that Veeam Agent will use to connect to the PostgreSQL database system.
- `<database>` – name of the PostgreSQL database.
- `<username>` – name of the account that Veeam Agent will use to connect to the PostgreSQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the PostgreSQL database system.

For example:

```
srv01:5432:mydb:postgres:P@ssw0rd
```

For more information about the password file, see [PostgreSQL documentation](#).

## Active Full Backup Settings

To specify active full backup settings for the backup job, do the following:

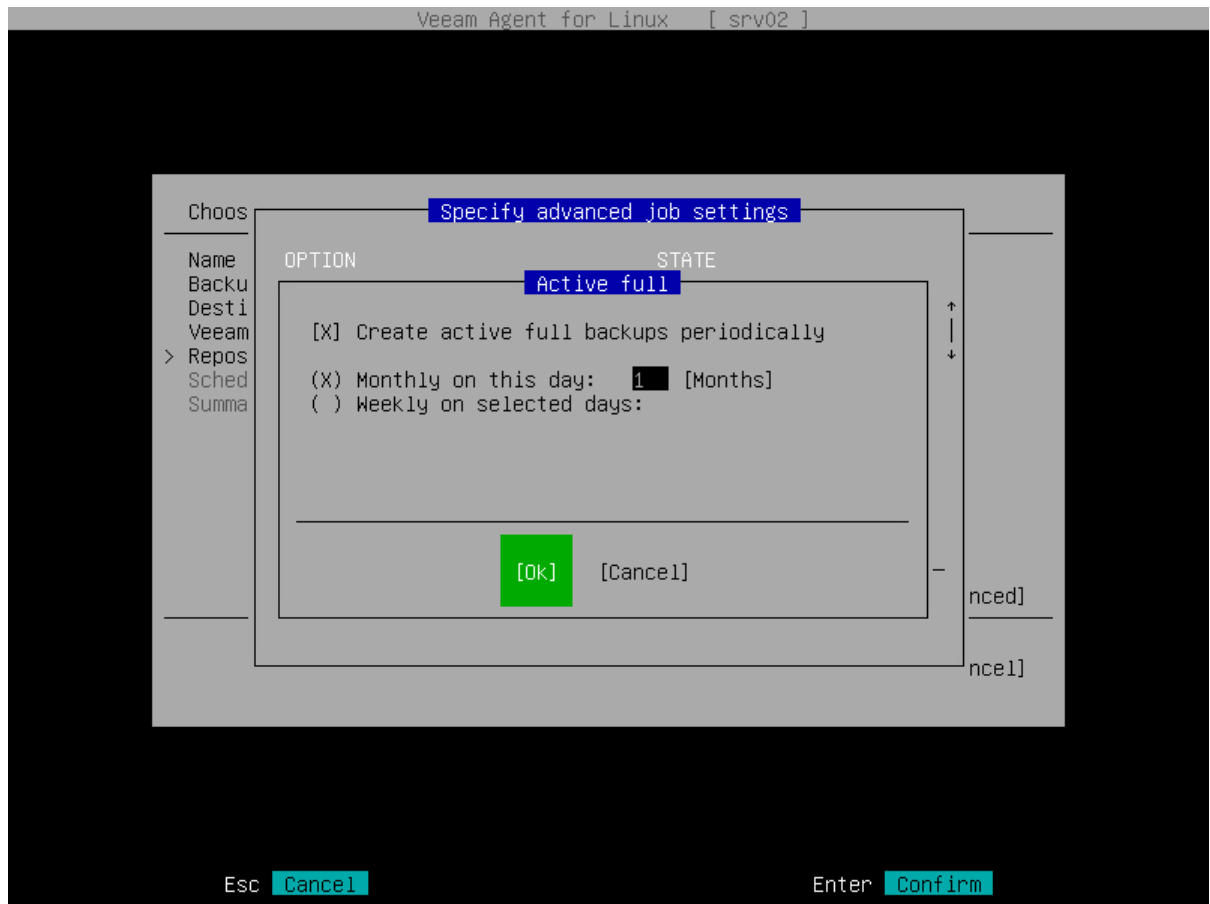
1. In the **Specify advanced job settings** window, select the **Active full** option with the [Tab] and [Down] keys and press [Enter].
2. In the **Active full** window, make sure that the **Create active full backups periodically** option is selected and press [Space].

### NOTE

If you plan to use a GFS retention policy, you must select the **Create active full backups periodically** option. Otherwise, Veeam Agent will not have full backups to mark with GFS flags. To learn more, see [Long-Term Retention Policy](#).

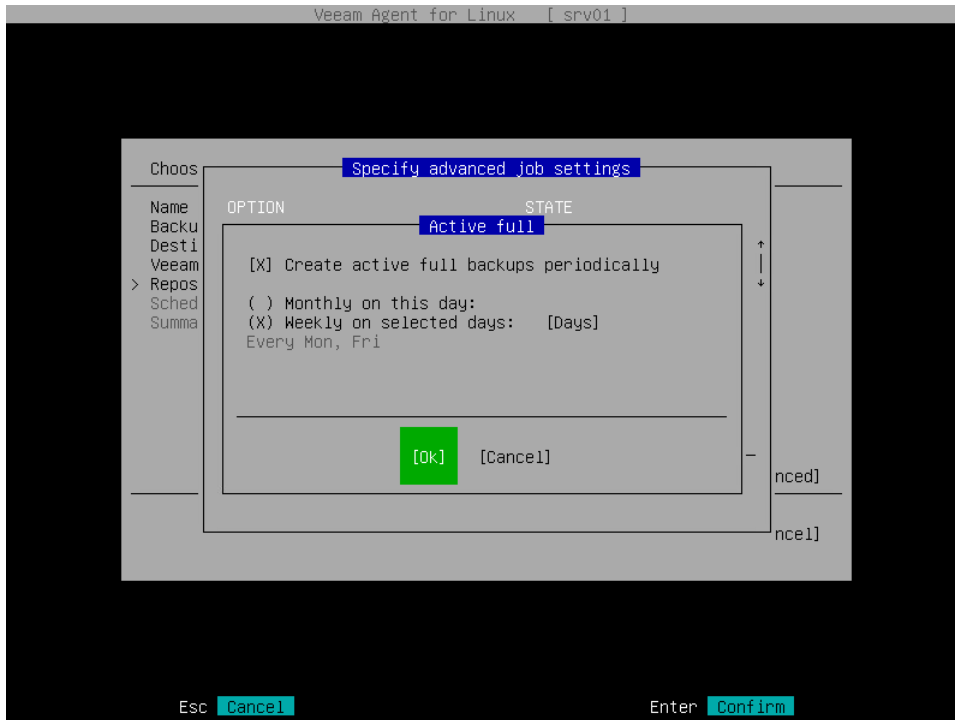
3. Specify schedule for periodic active full backups:
  - If you want active full backups to run monthly, do the following:
    - i. Select the **Monthly on this day** option and specify a day of a month when Veeam Agent will perform active full backup.
    - ii. To specify the months on which Veeam Agent will perform active full backups, select **Months** with the [Tab] key and press [Enter].
    - iii. In the **Months** window, specify the months on which Veeam Agent will perform active full backup. By default, Veeam Agent performs active full backup every month. To select months, use the [Up], [Down], [Right], [Left] and [Space] keys.

iv. Switch to the **Ok** button with the [Tab] key and press [Enter].



- If you want active full backups to run weekly, do the following:
  - i. Select the **Weekly on selected days** option, then select **Days** with the [Tab] key and press [Enter].
  - ii. In the **Days** window, specify the days on which Veeam Agent will perform active full backup. By default, Veeam Agent performs active full backup every Saturday. To select days, use the [Up], [Down], [Right], [Left] and [Space] keys.

- iii. Switch to the **Ok** button with the [Tab] key and press [Enter].



## Maintenance Settings

You can specify the number of days for which you want to keep the backup created with the backup job in the target location. To do this:

1. In the **Specify advanced job settings** window, select the **Maintenance** option with the [Tab] and [Down] keys and press [Enter].

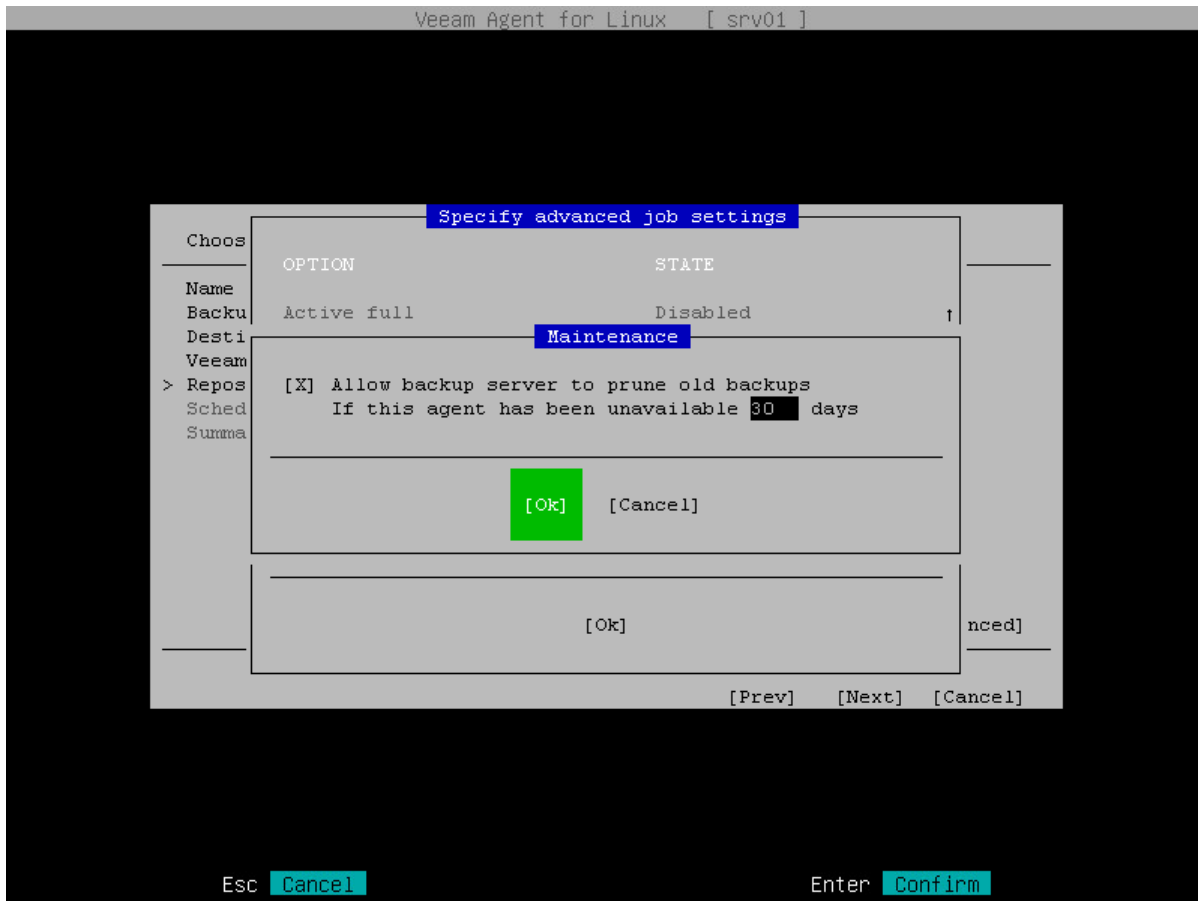
### NOTE

The **Maintenance** option is available if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.

2. In the **Maintenance** window, make sure that the **Allow backup server to prune old backups** option is selected and press [Space].

3. In the **If this agent has been unavailable <N> days** field, specify the number of days for which you want to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the period that you have specified. When this period is over, the backup will be removed from the target location.

By default, the retention period for old backups is 30 days. Do not set this retention period to 1 day or a similar short interval. In the opposite case, the backup job may work not as expected and remove data that you still require.



## Script Settings

To specify script settings for the backup job, do the following:

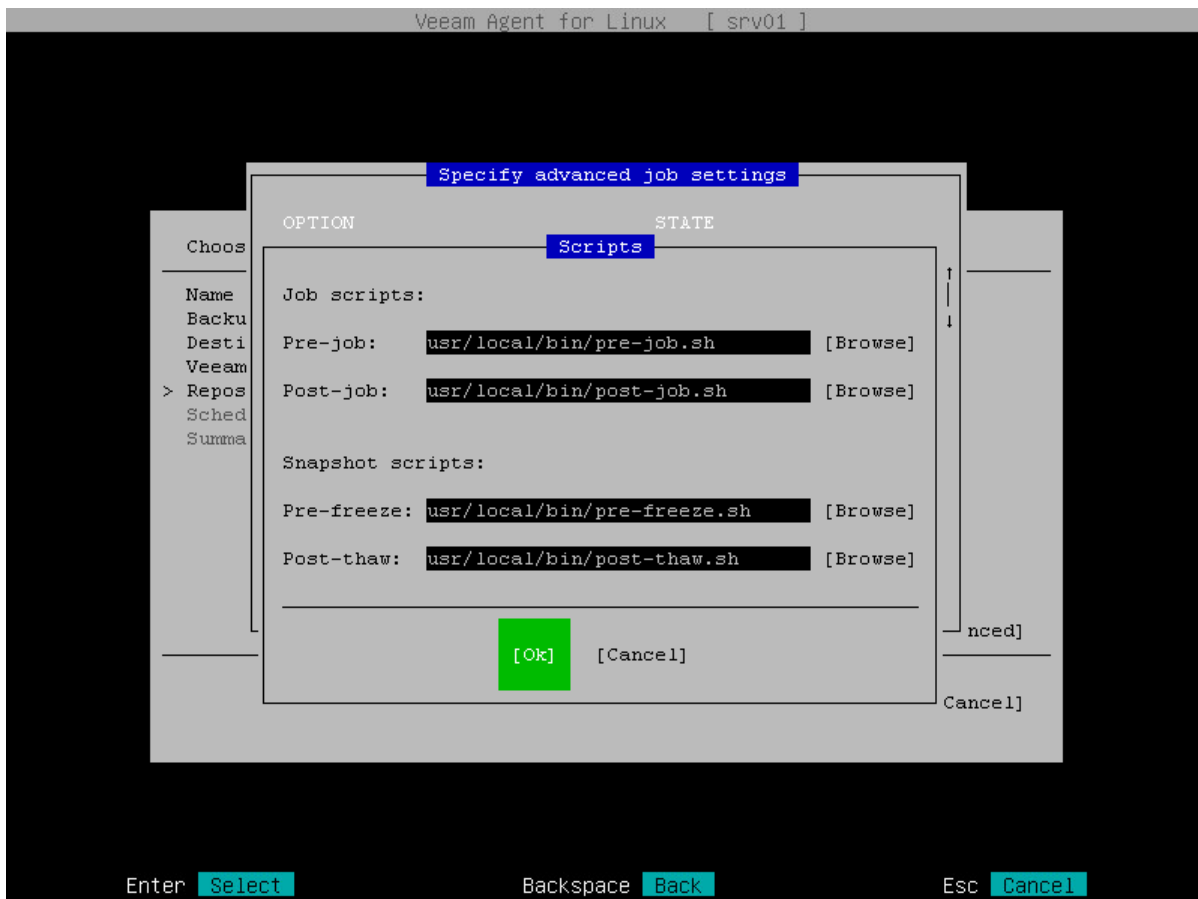
1. In the **Specify advanced job settings** window, select the **Scripts** option with the [Tab] and [Down] keys and press [Enter].
2. In the **Scripts** window, in the **Job scripts** section, specify custom scripts that you want to execute before and after the backup job:
  - In the **Pre-job** field, specify a path to the script that should be executed before the backup job starts.
  - In the **Post-job** field, specify a path to the script that should be executed after the backup job completes.
3. In the **Scripts** window, in the **Snapshot scripts** section, specify custom scripts that you want to execute before Veeam Agent creates a snapshot of the backed-up volume and after the snapshot is created:
  - In the **Pre-freeze** field, specify a path to the script that should be executed before Veeam Agent creates a volume snapshot.

- In the **Post-thaw** field, specify a path to the script that should be executed after Veeam Agent creates a volume snapshot.

4. Switch to the **Ok** button and press [Enter].

## IMPORTANT

You can specify snapshot script settings only if Veeam Agent for Linux operates in the Server edition. To learn more about editions, see [Product Editions](#).



## Specifying Path to Script

You can specify a path to the executable file of the job or snapshot script in one of the following ways:

1. Type a path to the executable file.
2. Browse to the executable file:
  - a. Select the **Browse** option with the [Tab] key and press [Enter].
  - b. In the **Choose script location** window, select the directory being a part of the path to the script and press [Enter].
  - c. Repeat the step 'b' until a path to the directory in which the executable file resides appears in the **Current directory** field.
  - d. Select the necessary executable file and press [Enter].

Alternatively, you can switch to the **Ok** button and press [Enter].

## TIP

If you do not want to execute a script, you can leave the corresponding field blank and proceed to the next step of the wizard.

# Health Check Settings

When you store backup files in an object storage repository, an automatic health check can help you avoid a situation when a restore point gets corrupted, making all dependent restore points corrupted, too. For more information, see [Health Check for Object Storage](#).

## NOTE

When you schedule a health check, consider the following:

- Health check runs automatically during incremental backup job session on the days specified in the health check schedule. If the backup job runs several times on a specified day, health check is performed only with the first run of the backup job on that day.  
Health check is not performed during the first full backup or subsequent active full backup jobs.
- If Veeam Agent does not run any backup jobs on the day specified in the health check schedule, health check will be performed during the first backup job session following that day.  
For example, you may have scheduled to run health check every last day of a month, while the backup job is scheduled to run every day and create an active full backup on Sundays. If the last day of a month falls on a Sunday, health check will be performed on the following Monday with the first incremental backup job session on that day.

To specify backup health check settings, do the following:

1. In the **Specify advanced job settings** window, select the **Health check** option with the [Tab] and [Down] keys and press [Enter].

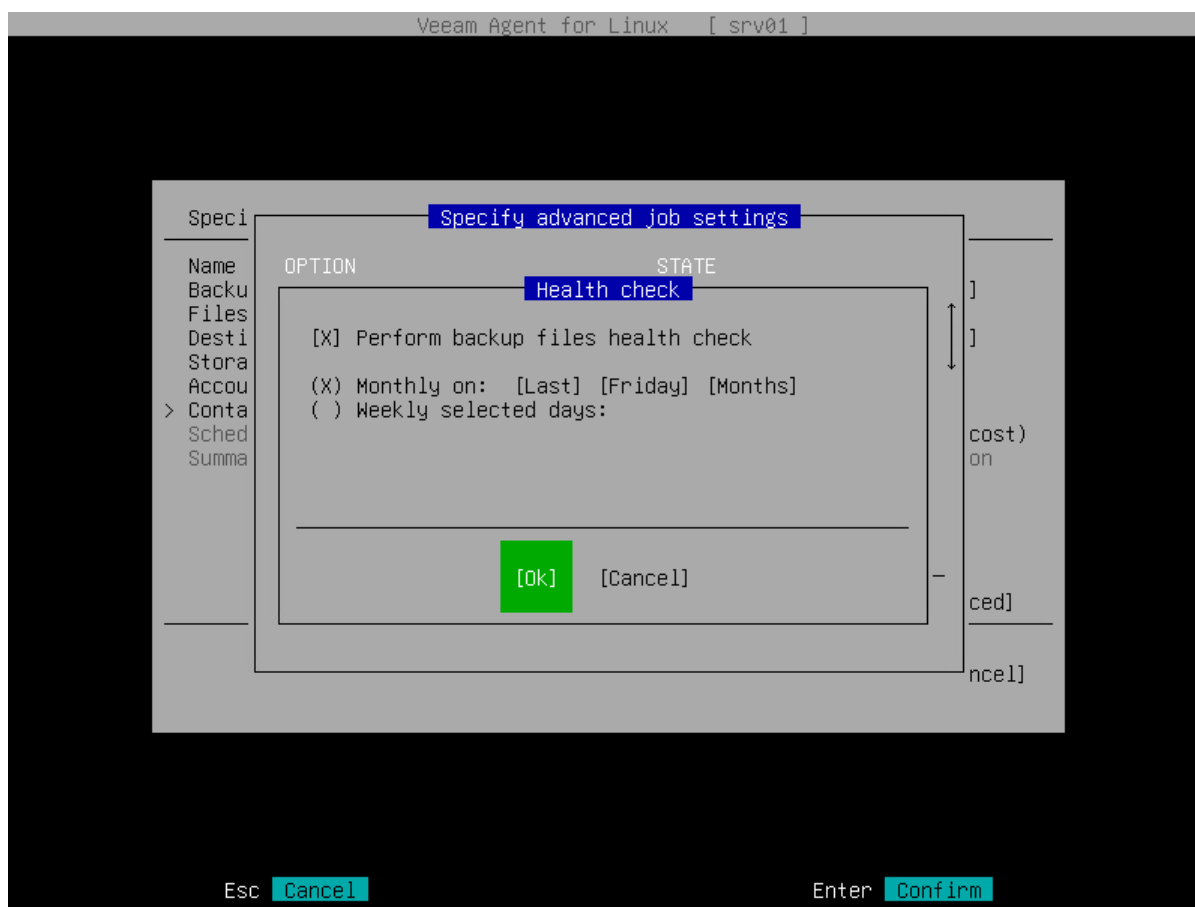
## NOTE

The **Health check** option is available if you have selected the **Object storage repository** option at the [Destination](#) step of the wizard.

2. In the **Health check** window, make sure that the **Perform backup files health check** option is highlighted and select it by pressing [Space].



3. Use the **Monthly on** or **Weekly selected days** settings to define the schedule for the health check of the backup in the repository.



## Step 9. Specify Backup Schedule

At the **Schedule** step of the wizard, specify the schedule according to which you want to perform backup.

Depending on the product edition, Veeam Agent provides the following scheduling options:

- [For Free and Workstation editions] You can set the backup job to run automatically on specific days of the week.
- [For Server edition] You can schedule the backup job to run on specific days of the week or month, as well as periodically.

To specify the schedule, do the following:

1. Make sure that the **Run the job automatically** check box is selected.

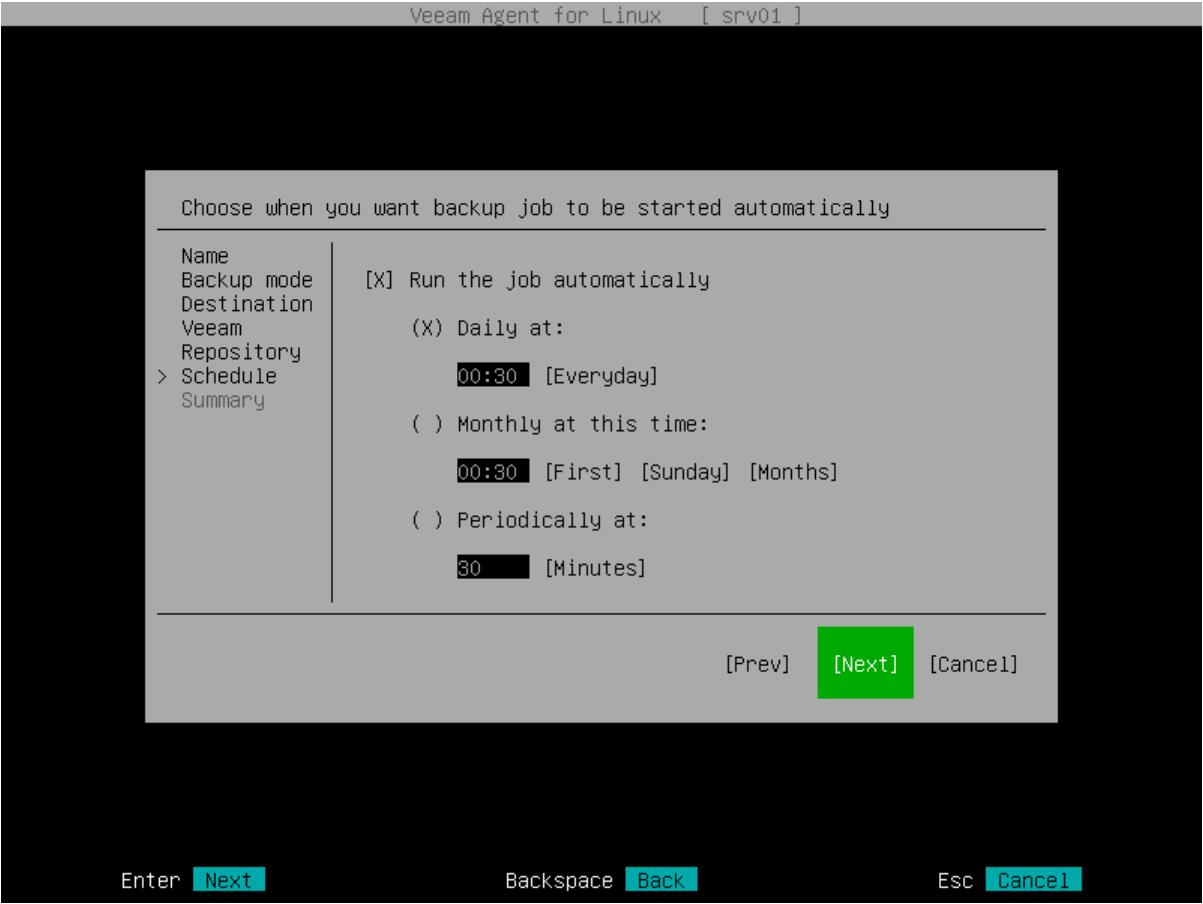
If you want to configure the backup job without schedule, you can clear the **Run the job automatically** check box. In this case you will be able to start the configured backup job manually at any time you need.

2. Define the schedule for the backup job:
  - To run the job at specific time daily or on specific weekdays, select the **Daily at** option. Use the fields of this option to configure the necessary schedule.
  - To run the job once a month on a specific day, select the **Monthly at this time** option. Use the fields of this option to configure the necessary schedule.
  - To run the job repeatedly throughout a day with a specific time interval, select the **Periodically at** option. Use the fields of this option to specify the time interval in hours or minutes.

### NOTE

Veeam Agent always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.

Veeam Agent for Linux will save the scheduling settings for the backup job in its database Veeam Agent can start a backup job automatically regardless of the currently running user session. You can change schedule settings at any time in Veeam Agent.



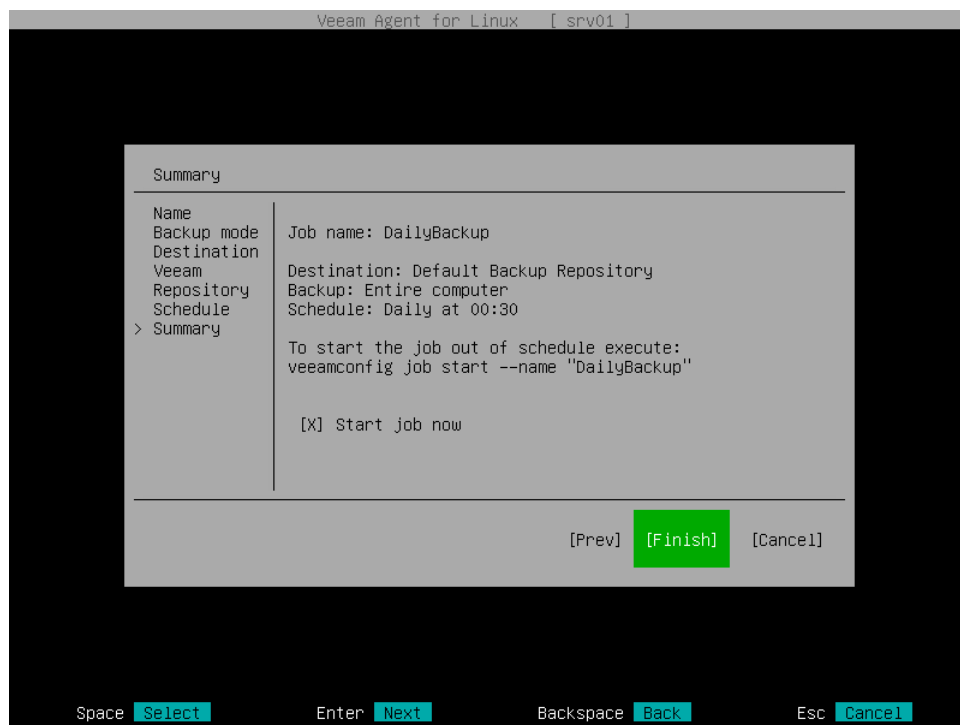
## Step 10. Review Backup Job Settings

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

1. Review settings of the configured backup job.
2. To start the job after you close the wizard, make sure that the **Start job now** check box is selected.

If you want to start the backup job later, you can clear the **Start job now** check box. You will be able to start the backup job manually at any time you need. To learn more, see [Starting Backup Job](#).

3. Press [Enter] to exit the wizard.



## What You Do Next

After you configure the backup job, you can start the backup job at any time you need. To learn more, see [Starting Backup Job](#).

If some of your data gets lost or corrupted, you can do the following:

- [Recover all computer volumes or specific volumes from the backup.](#)
- [Recover individual files and folders from the backup.](#)

# Creating Backup Job with Command Line Interface

You can configure the backup job with the command line interface. Using Veeam Agent for Linux commands, you can create volume-level and file-level backup jobs, specify advanced settings for the created backup job, define backup schedule and enable backup encryption.

## Creating Volume-Level Backup Job

### IMPORTANT

Volume-level backup job relies on a device name in the `/dev` directory. Device names in the `/dev` directory (for example, `/dev/md-127`, `/dev/dm-1`) must stay persistent for backed-up volumes. Otherwise, the job will back up the wrong volume.

You can create a volume-level backup of the entire computer image or specific volumes.

To back up the entire computer image, use the following command:

```
veeamconfig job create volumelevel --name <job_name> --reponame <repository_name> --backupallsystem  
<advanced_options> <schedule_options> <active_full_backup_options> <indexing_options>
```

To back up specific volumes, use the following command:

```
veeamconfig job create volumelevel --name <job_name> --reponame <repository_name> --objects <volume_to_backup> <advanced_options> <schedule_options> <active_full_backup_options> <indexing_options>
```

where:

- `<job_name>` – name for the created backup job.
- `<repository_name>` – name of the backup repository that should be used as a target location for the backup job. The backup repository must be created in advance.

If you want to create Veeam Agent backups in local directory or network shared folder, you need to create a repository. To learn more, see [Creating Backup Repository](#).

If you want to create Veeam Agent backups in a Veeam backup repository or cloud repository, you need to connect to the Veeam backup server or Veeam Cloud Connect service provider in advance, before configuring the backup job. To learn more, see [Connecting to Veeam Backup Server](#) and [Connecting to Service Provider](#).

If you want to create Veeam Agent backups in the object storage, you need to connect to an object storage and create a repository on this storage. To learn more, see [Creating Repository in Object Storage](#).

- `<volume_to_backup>` – object that should be included in backup:
  - For simple volumes – name of a block device that represents a volume or an entire disk that should be included in backup. You can specify entire disk to create backup of the entire computer image or individual computer volumes to create backup of specific volumes. If you want to back-up several disks or volumes, specify them one after another using the ',' (comma) character as a separator.

### IMPORTANT

Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.

### NOTE

If you include a block device in the backup, and this block device is a physical volume assigned to an LVM volume group, Veeam Agent will include the whole LVM volume group in the backup.

- For LVM volumes – name of an LVM logical volume that should be included in backup. If you want to back-up several LVM logical volumes, specify them one after another using the ',' (comma) character as a separator.
- `<advanced_options>` – advanced options for the backup job. To learn more, see [Advanced Backup Job Settings](#).
- `<schedule_options>` – schedule options for the backup job. To learn more, see [Schedule Settings](#).
- `<active_full_backup_options>` – active full backup schedule options for the backup job. To learn more, see [Active Full Backup Schedule Settings](#).
- `<indexing_options>` – file system indexing options for the backup job. To learn more, see [File System Indexing Settings](#).

For example:

```
$ veeamconfig job create --name SystemBackup --reponame Repository_01 --objects /dev/sda1 --weekdays Mon,Sun --weekdays-full Thu
```

### TIP

After you create the backup job, you can additionally configure the following backup job settings:

- Backup schedule. For details, see [Configuring Backup Schedule](#).
- Active full backup schedule. For details, see [Configuring Active Full Backup Schedule](#).
- Long-term retention policy. For details, see [Configuring Long-Term Retention Policy](#).
- Database processing settings for a volume-level backup job. For details, see [Configuring Database Processing Settings](#).
- [For job targeted at an object storage repository] Schedule for backup health check. For details, see [Configuring Health Check Schedule](#).

# Advanced Backup Job Settings

You can specify the following advanced options for the backup job:

Option	Description and values
<b>--compressionlevel</b>	<p>Data compression level. Possible values are:</p> <ul style="list-style-type: none"><li>• 0 – No compression</li><li>• 1 – Rle</li><li>• 2 – Lz4</li><li>• 3 – Zstd 3</li><li>• 4 – Zstd 9</li></ul> <p>The default value is 2.</p>
<b>--blocksize</b>	<p>Data block size in kilobytes. Possible values are 256, 512, 1024, 4096 or 8192.</p> <p>The default value is 1024.</p>
<b>--maxpoints</b>	<p>The number of restore points that you want to store in the backup location. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.</p>
<b>--immutabledays</b>	<p>The time period in days during which the backup stored in an object storage repository will be immutable to modification or deletion. For more information, see <a href="#">Backup Immutability</a>.</p>
<b>--prefreeze</b>	<p>Path to the script that should be executed before the snapshot creation.</p> <p>This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see <a href="#">Product Editions</a>.</p>
<b>--postthaw</b>	<p>Path to the script that should be executed after the snapshot creation.</p> <p>This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see <a href="#">Product Editions</a>.</p>
<b>--prejob</b>	<p>Path to the script that should be executed at the start of the backup job.</p>
<b>--postjob</b>	<p>Path to the script that should be executed after the backup job completes.</p>
<b>--setencryption</b>	<p>Defines that data encryption option is enabled for the job. When you use the <code>veeamconfig job create</code> command with the <code>--setencryption</code> option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.</p>

Option	Description and values
<b>--deleteold</b>	<p>The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999.</p> <p>If you do not specify the <code>--deleteold</code> option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually.</p>

## Schedule Settings

You can specify a daily, monthly or periodic schedule for the backup job.

Option	Description and values
<b>--weekdays</b>	<p>[For weekly schedules] Specifies the weekdays when the backup job must run. If you want to run the backup job more than once during the week, the list of weekdays must be separated by a comma (',' ). Possible values are:</p> <ul style="list-style-type: none"> <li><i>Mon</i> – Monday</li> <li><i>Tue</i> – Tuesday</li> <li><i>Wed</i> – Wednesday</li> <li><i>Thu</i> – Thursday</li> <li><i>Fri</i> – Friday</li> <li><i>Sat</i> – Saturday</li> <li><i>Sun</i> – Sunday</li> </ul>
<b>--daily</b>	[For weekly schedules] Defines that the backup job must start daily at specific time.
<b>--thisday</b>	[For monthly schedules] Specifies the day of the month when the backup job must run. Possible values: from 1 to 31 or <i>Last</i> .
<b>--weeknumber</b>	[For monthly schedules] Specifies the week of the month when the backup job must run. Possible values: <i>First</i> , <i>Second</i> , <i>Third</i> , <i>Fourth</i> or <i>Last</i> . This option must be used in combination with the <code>--monthlyweekday</code> option.



Option	Description and values
<b>--monthlyweekday</b>	<p>[For monthly schedules] Specifies the day of the week when the backup job must run. You can select only one weekday. Possible values are:</p> <ul style="list-style-type: none"> <li>• <i>Mon</i> – Monday</li> <li>• <i>Tue</i> – Tuesday</li> <li>• <i>Wed</i> – Wednesday</li> <li>• <i>Thu</i> – Thursday</li> <li>• <i>Fri</i> – Friday</li> <li>• <i>Sat</i> – Saturday</li> <li>• <i>Sun</i> – Sunday</li> </ul>
<b>--months</b>	<p>[For monthly schedules] Specifies the months when the backup job must run. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan, Apr, Jul, Oct</i>. If you do not specify this option, the backup job will run every month.</p>
<b>--every</b>	<p>[For periodic schedules] Specifies the period of time in minutes or hours between the runs of the backup job. The period must be specified in the <i>HH:MM</i> format – for example, <i>06:00</i>.</p>
<b>--at</b>	<p>[For weekly and monthly schedules] Specifies the time of day in the <i>HH:MM</i> format when the backup job must start – for example: <i>20:00</i>.</p>

After the backup job is created, Veeam Agent for Linux automatically enables backup schedule. To learn about how to configure backup schedule for an existing backup job, see [Configuring Backup Schedule](#).

## Active Full Backup Schedule Settings

You can specify schedule options for the backup job to create active full backups on specific weekdays or days of the month.

You can specify schedule options for the backup job to create active full backups on specific days of the week or month.

Option	Description and values
<b>--weekdays-full</b>	<p>[For weekly schedules] Specifies the weekdays when the backup job must create an active full backup. If you want to create an active full backup more than once during the week, the list of weekdays must be separated by a comma (',' ). Possible values are:</p> <ul style="list-style-type: none"> <li>• <i>Mon</i> – Monday</li> <li>• <i>Tue</i> – Tuesday</li> <li>• <i>Wed</i> – Wednesday</li> <li>• <i>Thu</i> – Thursday</li> <li>• <i>Fri</i> – Friday</li> <li>• <i>Sat</i> – Saturday</li> <li>• <i>Sun</i> – Sunday</li> </ul>
<b>--thisday-full</b>	<p>[For monthly schedules] Specifies the day of the month when the backup job must create an active full backup. Possible values: from 1 to 31 or <i>Last</i>.</p>
<b>--weeknumber-full</b>	<p>[For monthly schedules] Specifies the week of the month when the backup job must create an active full backup.. Possible values: <i>First</i>, <i>Second</i>, <i>Third</i>, <i>Fourth</i> or <i>Last</i>. This option must be used in combination with the <code>--monthlyweekday</code> option.</p>
<b>--monthlyweekday-full</b>	<p>[For monthly schedules] Specifies the day of the week when the backup job must create an active full backup. You can select only one weekday. Possible values are:</p> <ul style="list-style-type: none"> <li>• <i>Mon</i> – Monday</li> <li>• <i>Tue</i> – Tuesday</li> <li>• <i>Wed</i> – Wednesday</li> <li>• <i>Thu</i> – Thursday</li> <li>• <i>Fri</i> – Friday</li> <li>• <i>Sat</i> – Saturday</li> <li>• <i>Sun</i> – Sunday</li> </ul>
<b>--months-full</b>	<p>[For monthly schedules] Specifies the months when the backup job must create an active full backup. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan</i>, <i>Apr</i>, <i>Jul</i>, <i>Oct</i>. If you do not specify this option, the backup job will create an active full backup. every month.</p>

After the backup job is created, Veeam Agent automatically enables active full backup schedule. To learn about how to configure active full backup schedule for an existing backup job, see [Configuring Active Full Backup Schedule](#).

# File System Indexing Settings

You can specify one the following file system indexing options for the backup job:

Option	Description and values
<b>--indexall</b>	Defines that Veeam Agent for Linux must index all files on the volumes included in backup.
<b>--indexonly</b>	Path to a directory that contains files that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character.
<b>--indexexcept</b>	Path to a directory that contains files that you do not want to index. You can specify one or more paths. To separate several paths, use the ',' (comma) character.

To learn more about file indexing, see [File System Indexing](#).

## Creating File-Level Backup Job

To create a file-level backup job, use the following command:

```
veeamconfig job create filelevel --name <job_name> --reponame <repository_name>  
<objects> <advanced_options> <schedule_options> <active_full_backup_options> <i  
ndexing_options> --nosnap
```

where:

- **<job\_name>** – name for the created backup job.
- **<repository\_name>** – name of the backup repository that should be used as a target location for the backup job. The backup repository must be created in advance. To learn more, see [Creating Backup Repository](#).  
  
If you want to create Veeam Agent backups in the Veeam backup repository, you should connect to the Veeam backup server in advance, before configuring the backup job. To learn more, see [Connecting to Veeam Backup Server](#).
- **<objects>** – files and directories inclusion/exclusion options. To learn more, see [File Inclusion Options](#).
- **<advanced\_options>** – advanced options for the backup job. To learn more, see [Advanced Backup Job Settings](#).
- **<schedule\_options>** – schedule options for the backup job. To learn more, see [Schedule Settings](#).
- **<active\_full\_backup\_options>** – active full backup schedule options for the backup job. To learn more, see [Active Full Backup Schedule Settings](#).
- **<indexing\_options>** – file system indexing options for the backup job. To learn more, see [File System Indexing Settings](#).

- `--nosnap` — option that instructs Veeam Agent for Linux to perform backup in the snapshot-less mode. With this option enabled, Veeam Agent for Linux will not create a snapshot of the backed-up volumes during backup. This allows Veeam Agent to back up data residing in file systems that are not supported for snapshot-based backup with Veeam Agent for Linux. Keep in mind that the snapshot-less file-level backup does not guarantee that data in the backup is consistent. To learn more, see [Snapshot-Less File-Level Backup](#).

For example:

```
$ veeamconfig job create filelevel --name HomeFolderBackup --reponame NetworkRepository --includedirs /home/user --excludedirs /home/user/temp --excludemasks "*.pdf"
```

#### TIP

After you create the backup job, you can additionally configure the following backup job settings:

- Backup schedule. For details, see [Configuring Backup Schedule](#).
- Active full backup schedule. For details, see [Configuring Active Full Backup Schedule](#).
- Long-term retention policy. For details, see [Configuring Long-Term Retention Policy](#).
- [For job targeted at an object storage repository] Schedule for backup health check. For details, see [Configuring Health Check Schedule](#).

## File Inclusion Options

When you create a file-level backup job, you must specify at least one directory that should be included in backup. If you do not want to back up some files and directories in the specified directory, you can exclude specific files and directories from backup.

#### IMPORTANT

Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.

To define the backup scope for the file-level backup job, you can use the following command-line options:

Option	Description and values
<b>--includedirs</b>	<p>Full path to a directory that should be included in backup, for example: <code>/home/user</code>.</p> <p>You can specify one or several paths to directories in the computer file system. To separate several paths, use the <code>' , '</code> (comma) character, for example: <code>/home/user/Documents, /home/user/reports</code>.</p> <p><b>Tip:</b> If you want to backup the root directory and specify the <code>'/'</code> (slash) character, Veeam Agent will not automatically include the mount points in the backup scope. To include the mount points, you can do either of the following:</p> <ul style="list-style-type: none"><li>• Enable automatic inclusion of local mount points when the root directory is added into the backup scope. To do this, in the Veeam Agent configuration file, enable the <code>rootRecursion</code> option and set it to <i>true</i>: <code>rootRecursion = true</code>.</li></ul> <p>Note that even if you enable this configuration option, network file systems will not be included into backup automatically; you will need to specify paths to such mount points manually.</p> <ul style="list-style-type: none"><li>• Specify paths to the mount points manually.</li></ul> <p>For example, you have a network file system mounted to the <code>/home/media</code> directory. If you add <code>'/'</code> as an object to the backup scope, Veeam Agent will not back up the mounted network file system. To back up the root directory and the mounted network file system, add the following objects to the backup scope: <code>/, /home/media</code>.</p>
<b>--excludedirs</b>	<p>Full path to a directory that should be excluded from backup. The directory specified with this option must be a subdirectory of the directory specified with the <code>--includedirs</code> option. To separate several paths, use the <code>' , '</code> (comma) character, for example, <code>/home/user/Documents, /home/user/reports</code>.</p>

Option	Description and values
<b>--includemasks</b>	<p>A name mask for the files that should be included in the backup. You can use the following UNIX wildcard characters for file name masks:</p> <ul style="list-style-type: none"> <li>• '*' – a substitution for one or more characters in the file name. Can be used for any sequence of characters (including no characters). For example, *.pdf.</li> <li>• '?' – a substitution of one character in the file name. For example, repor?.pdf.</li> <li>• '[' – a substitution of one character in the file name with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: report_201[3456].pdf or report_201[3-6].pdf.</li> </ul> <p>Keep in mind that you must specify each name mask in double quotation marks (""). For example: --includemasks "*.bak".</p> <p>If you want to use several file name masks, you must specify them in double quotation marks ("") and separate them with a comma (,). For example: --includemasks "*.bak,*.pdf".</p> <p>File inclusion option is applied to all directories that are specified with the --includedirs option. For example, if you include in backup the /home/user/Documents directory and files that match the repor?.pdf file name mask, Veeam Agent for Linux will back up the /home/user/Documents/report.pdf file and will not back up the /home/user/reports/report.pdf file.</p>

Option	Description and values
<b>--excludemasks</b>	<p>A name mask for the files that should be excluded from the backup. You can use the following UNIX wildcard characters for file name masks:</p> <ul style="list-style-type: none"> <li>• '*' – a substitution for one or more characters in the file name. Can be used for any sequence of characters (including no characters). For example, *.pdf.</li> <li>• '?' – a substitution of one character in the file name. For example, repor?.pdf.</li> <li>• '[' – a substitution of one character in the file name with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: report_201[3456].pdf or report_201[3-6].pdf.</li> </ul> <p>Keep in mind that you must specify each name mask in double quotation marks (""). For example: --excludemasks "*.bak".</p> <p>If you want to use several file name masks, you must specify them in double quotation marks ("") and separate them with a comma (,). For example: --excludemasks *.bak, *.pdf".</p> <p>File exclusion option is applied to all directories that are specified with the --includedirs option and files that match file name masks specified with the --includemasks option. For example, you may want to specify the following backup scope for the backup job:</p> <ul style="list-style-type: none"> <li>• Include in backup the /home/user/Documents directory</li> <li>• Include files that match the report.* file name mask</li> <li>• Exclude files that match the *.odt file name mask.</li> </ul> <p>In this case, Veeam Agent for Linux will back up the /home/user/Documents/report.pdf file and will not back up /home/user/Documents/report.odt and /home/user/reports/report.pdf files.</p>

# Advanced Backup Job Settings

You can specify the following advanced options for the backup job:

Option	Description and values
<b>--compressionlevel</b>	<p>Data compression level. Possible values are:</p> <ul style="list-style-type: none"><li>• 0 – No compression</li><li>• 1 – Rle</li><li>• 2 – Lz4</li><li>• 3 – Zstd 3</li><li>• 4 – Zstd 9</li></ul> <p>To learn more about available data compression levels, see <a href="#">Data Compression</a>.</p>
<b>--blocksize</b>	<p>Data block size in kilobytes. The default value is 1024. Possible values are 256, 512, 1024 or 4096. To learn more about available sizes of data blocks, see <a href="#">Data Compression</a>.</p>
<b>--maxpoints</b>	<p>The number of restore points that you want to store in the backup location. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.</p>
<b>--immutabledays</b>	<p>The time period in days during which the backup stored in an object storage repository will be immutable to modification or deletion. For more information, see <a href="#">Backup Immutability</a>.</p>
<b>--prefreeze</b>	<p>Path to the pre-freeze script that should be executed before the snapshot creation.</p> <p>This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see <a href="#">Product Editions</a>.</p>
<b>--postthaw</b>	<p>Path to the post-thaw script that should be executed after the snapshot creation.</p> <p>This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see <a href="#">Product Editions</a>.</p>
<b>--prejob</b>	<p>Path to the script that should be executed at the start of the backup job.</p>
<b>--postjob</b>	<p>Path to the script that should be executed after the backup job completes.</p>
<b>--setencryption</b>	<p>Defines that data encryption option is enabled for the job. When you use the <code>veeamconfig job create</code> command with the <code>--setencryption</code> option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.</p>



Option	Description and values
<b>--deleteold</b>	<p>The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999.</p> <p>If you do not specify the <code>--deleteold</code> option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually.</p>

## Schedule Settings

You can specify a daily, monthly or periodic schedule for the backup job.

Option	Description and values
<b>--weekdays</b>	<p>[For weekly schedules] Specifies the weekdays when the backup job must run. If you want to run the backup job more than once during the week, the list of weekdays must be separated by a comma (',' ). Possible values are:</p> <ul style="list-style-type: none"> <li><i>Mon</i> – Monday</li> <li><i>Tue</i> – Tuesday</li> <li><i>Wed</i> – Wednesday</li> <li><i>Thu</i> – Thursday</li> <li><i>Fri</i> – Friday</li> <li><i>Sat</i> – Saturday</li> <li><i>Sun</i> – Sunday</li> </ul>
<b>--daily</b>	[For weekly schedules] Defines that the backup job must start daily at specific time.
<b>--thisday</b>	[For monthly schedules] Specifies the day of the month when the backup job must run. Possible values: from 1 to 31 or <i>Last</i> .
<b>--weeknumber</b>	[For monthly schedules] Specifies the week of the month when the backup job must run. Possible values: <i>First</i> , <i>Second</i> , <i>Third</i> , <i>Fourth</i> or <i>Last</i> . This option must be used in combination with the <code>--monthlyweekday</code> option.

Option	Description and values
<b>--monthlyweekday</b>	<p>[For monthly schedules] Specifies the day of the week when the backup job must run. You can select only one weekday. Possible values are:</p> <ul style="list-style-type: none"> <li>• <i>Mon</i> – Monday</li> <li>• <i>Tue</i> – Tuesday</li> <li>• <i>Wed</i> – Wednesday</li> <li>• <i>Thu</i> – Thursday</li> <li>• <i>Fri</i> – Friday</li> <li>• <i>Sat</i> – Saturday</li> <li>• <i>Sun</i> – Sunday</li> </ul>
<b>--months</b>	<p>[For monthly schedules] Specifies the months when the backup job must run. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan, Apr, Jul, Oct</i>. If you do not specify this option, the backup job will run every month.</p>
<b>--every</b>	<p>[For periodic schedules] Specifies the period of time in minutes or hours between the runs of the backup job. The period must be specified in the <i>HH:MM</i> format – for example, <i>06:00</i>.</p>
<b>--at</b>	<p>[For weekly and monthly schedules] Specifies the time of day in the <i>HH:MM</i> format when the backup job must start – for example: <i>20:00</i>.</p>

After the backup job is created, Veeam Agent for Linux automatically enables backup schedule. To learn about how to configure backup schedule for an existing backup job, see [Configuring Backup Schedule](#).

## Active Full Backup Schedule Settings

You can specify schedule options for the backup job to create active full backups on specific weekdays or days of the month.

You can specify schedule options for the backup job to create active full backups on specific days of the week or month.

Option	Description and values
<b>--weekdays-full</b>	<p>[For weekly schedules] Specifies the weekdays when the backup job must create an active full backup. If you want to create an active full backup more than once during the week, the list of weekdays must be separated by a comma (',' ). Possible values are:</p> <ul style="list-style-type: none"> <li>• <i>Mon</i> – Monday</li> <li>• <i>Tue</i> – Tuesday</li> <li>• <i>Wed</i> – Wednesday</li> <li>• <i>Thu</i> – Thursday</li> <li>• <i>Fri</i> – Friday</li> <li>• <i>Sat</i> – Saturday</li> <li>• <i>Sun</i> – Sunday</li> </ul>
<b>--thisday-full</b>	<p>[For monthly schedules] Specifies the day of the month when the backup job must create an active full backup. Possible values: from 1 to 31 or <i>Last</i>.</p>
<b>--weeknumber-full</b>	<p>[For monthly schedules] Specifies the week of the month when the backup job must create an active full backup.. Possible values: <i>First</i>, <i>Second</i>, <i>Third</i>, <i>Fourth</i> or <i>Last</i>. This option must be used in combination with the <code>--monthlyweekday</code> option.</p>
<b>--monthlyweekday-full</b>	<p>[For monthly schedules] Specifies the day of the week when the backup job must create an active full backup. You can select only one weekday. Possible values are:</p> <ul style="list-style-type: none"> <li>• <i>Mon</i> – Monday</li> <li>• <i>Tue</i> – Tuesday</li> <li>• <i>Wed</i> – Wednesday</li> <li>• <i>Thu</i> – Thursday</li> <li>• <i>Fri</i> – Friday</li> <li>• <i>Sat</i> – Saturday</li> <li>• <i>Sun</i> – Sunday</li> </ul>
<b>--months-full</b>	<p>[For monthly schedules] Specifies the months when the backup job must create an active full backup. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan</i>, <i>Apr</i>, <i>Jul</i>, <i>Oct</i>. If you do not specify this option, the backup job will create an active full backup. every month.</p>

After the backup job is created, Veeam Agent automatically enables active full backup schedule. To learn about how to configure active full backup schedule for an existing backup job, see [Configuring Active Full Backup Schedule](#).

# File System Indexing Settings

You can specify the following file system indexing option for the backup job:

Option	Description and values
<b>--indexall</b>	Defines that Veeam Agent for Linux must index all files in the directories included in backup.

To learn more about file indexing, see [File System Indexing](#).

## Configuring Backup Schedule

To run a backup job periodically without the user intervention, you can schedule it to start automatically. You can specify schedule settings individually for every job created in Veeam Agent. You can perform the following actions with the backup job schedule via command line interface:

- [Specify schedule settings for the job.](#)
- [Enable schedule for the job.](#)
- [View the schedule defined for the job.](#)
- [Disable schedule for the job.](#)

### TIP

You can also specify backup schedule for the backup job when you create the job. For details, see [Creating Volume-Level Backup Job](#) and [Creating File-Level Backup Job](#).

## Specifying Backup Schedule

Depending on the product edition, Veeam Agent allows you to set [daily](#), [monthly](#) or [periodic](#) schedule for a backup job. Daily schedules are available for the Free and Workstation editions of Veeam Agent. In the Server edition of Veeam Agent, you can additionally set monthly and periodic schedules for backup jobs. For details on Veeam Agent editions, see [Product Editions](#).

After you define the schedule, Veeam Agent automatically enables this schedule for the specified backup job.

## Specifying Daily Schedules

You can set the backup job to run automatically on specific weekdays or every day.

- To run the backup job on specific days of the week, use the following command:

```
veeamconfig schedule set --jobid <job_id> --weekdays <days> --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --weekdays <days> --at <time>
```

where:

- <job\_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job\_name> – name of the backup job for which you want to configure the schedule.
- <days> – days when the backup job must start separated by a comma (',') – for example: Monday, Tuesday, Wednesday, Thursday, Friday or Mon, Tue, Wed, Thu, Fri.
- <time> – time of day when the backup job must start specified in the HH:MM format – for example, 20:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --weekdays Monday, Tuesday, Wednesday, Thursday, Friday --at 20:00
```

- To run the backup job every day, use the following command:

```
veeamconfig schedule set --jobid <job_id> --daily --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> <daily options> --at <time>
```

where:

- <job\_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job\_name> – name of the backup job for which you want to configure the schedule.
- <time> – time of day when the backup job must start specified in the HH:MM format – for example, 20:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f --daily --at 20:00
```

# Specifying Monthly Schedules

You can set the backup job to run automatically on specific months or every month.

- To run the backup job monthly on a specific day of the specific week, use the following command:

```
veeamconfig schedule set --jobid <job_id> --monthlyweekday <day> --weeknumber <week> [--months <months>] --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --monthlyweekday <day> --weeknumber <week> [--months <months>] --at <time>
```

where:

- o <job\_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- o <job\_name> – name of the backup job for which you want to configure the schedule.
- o <day> – day of the week when the backup job must start – for example, `Tuesday` or `Tue`.
- o <week> – week of the month when the backup job must run. Possible values: `First`, `Second`, `Third`, `Fourth` or `Last`.
- o <months> – months when the backup job must run separated by a comma (',' ) – for example: `Jan, Apr, Jul, Oct`. If you do not specify this option, the backup job will run every month.
- o <time> – time of day when the backup job must start specified in the `HH:MM` format, – for example, `20:00`.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --monthlyweekday Mon --weeknumber Second --months Jan,Jul --at 20:00
```

- To run the backup job monthly on a specific day of the month, use the following command:

```
veeamconfig schedule set --jobid <job_id> --thisday <day> [--months <months>] --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --thisday <day> [--months <months>] --at <time>
```

where:

- o <job\_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).

- `<job_name>` – name of the backup job for which you want to configure the schedule.
- `<day>` – day of the month when the backup job must start. Possible values range from 1 to 31 or Last.
- `<months>` – months when the backup job must run separated by a comma (','), – for example: Jan, Apr, Jul, Oct. If you do not specify this option, the backup job will run every month.
- `<time>` – time of day when the backup job must start specified in the HH:MM format, – for example, 20:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --thisday 21
--months Jan,Jul --at 20:00
```

## Specifying Periodic Schedules

To run the job periodically, run the following command:

```
veeamconfig schedule set --jobid <job_id> --every <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --every <time>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the schedule.
- `<time>` – period of time when the backup job must start specified in the HH:MM format, – for example, 06:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --every 12:00
```

## Viewing Backup Schedule

To view the schedule defined for the backup job, use the following command:

```
veeamconfig schedule show --jobid <job_id>
```

or

```
veeamconfig schedule show --jobname <job_name>
```

where:

- <job\_id> – ID of the backup job for which you want to view the schedule.
- <job\_name> – name of the backup job for which you want to view the schedule.

Veeam Agent will display the details and the status (enabled or disabled) of the job schedule – for example:

```
user@srv01:~$ veeamconfig schedule show --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
Days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday
At: 20:00
Run automatically: enabled
```

or

```
user@srv01:~$ veeamconfig schedule show --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
Every 12 hours
Run automatically: enabled
```

## Disabling Backup Schedule

To disable the schedule for the backup job, use the following command:

```
veeamconfig schedule disable --jobid <job_id>
```

or

```
veeamconfig schedule disable --jobname <job_name>
```

where:

- <job\_id> – ID of the backup job for which you want to disable the schedule.
- <job\_name> – name of the backup job for which you want to disable the schedule.

For example:

```
user@srv01:~$ veeamconfig schedule disable --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
```



## Enabling Backup Schedule

After you define the schedule, Veeam Agent automatically enables this schedule for the specified backup job. If you disable the schedule for the backup job, you can enable it again by using the following command:

```
veeamconfig schedule enable --jobid <job_id>
```

or

```
veeamconfig schedule enable --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to enable the schedule. You should look up the job ID in advance, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to enable the schedule.

For example:

```
user@srv01:~$ veeamconfig schedule enable --jobid 4849a3ae-1935-4969-98a3-d8acd  
2f6c73f
```

You can disable the schedule for the job at any time. To learn more, see [Disabling Backup Schedule](#).

## Configuring Long-Term Retention Policy

You can configure the backup job to store backup files for long periods of time – for weeks, months and even years, you can set the long-term or Grandfather-Father-Son (GFS) retention policy. This policy uses backup files created while backup job is enabled and marks these backups with specific GFS flags. You can perform the following actions with the long-term retention policy in command line interface:

- [Specify long-term retention policy for the job.](#)
- [View the long-term retention policy defined for the job.](#)
- [Disable long-term retention policy for the job.](#)
- [Enable long-term retention policy for the job.](#)

## Specifying Long-Term Retention Policy

You can configure long-term retention policy to keep [weekly](#), [monthly](#) or [yearly](#) full backups.

# Configuring Long-Term Retention Policy to Keep Weekly Full Backups

To configure long-term retention policy to keep weekly full backups, use the following command:

```
veeamconfig gfs set --jobid <job_id> weekly --on <weekday> --keep <weeks>
```

or

```
veeamconfig gfs set --jobname <job_name> weekly --on <weekday> --keep <weeks>
```

where:

- **<job\_id>** – ID of the backup job for which you want to configure the long-term retention policy. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job\_name>** – name of the backup job for which you want to configure the long-term retention policy.
- **<weekday>** – week day when Veeam Agent must assign a weekly GFS flag to a full restore point – for example, `Tue` or `Tuesday`.
- **<weeks>** – number of weeks to keep the weekly GFS flag on the full restore point.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f  
weekly --on Saturday --keep 1
```

# Configuring Long-Term Retention Policy to Keep Monthly Full Backups

To configure long-term retention policy to keep monthly full backups, use the following command:

```
veeamconfig gfs set --jobid <job_id> monthly --on <week_number> --keep <months>
```

or

```
veeamconfig gfs set --jobname <job_name> monthly --on <week_number> --keep <months>
```

where:

- **<job\_id>** – ID of the backup job for which you want to configure the long-term retention policy. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).

- `<job_name>` – name of the backup job for which you want to configure the long-term retention policy.
- `<week_number>` – number of the week when Veeam Agent must assign a monthly GFS flag to a full restore point. Possible values: `First`, `Second`, `Third`, `Fourth` or `Last`.
- `<months>` – number of months to keep the monthly GFS flag on the full restore point.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
monthly --on Second --keep 6
```

## Configuring Long-Term Retention Policy to Keep Yearly Full Backups

To configure long-term retention policy to keep yearly full backups, use the following command:

```
veeamconfig gfs set --jobid <job_id> yearly --on <month> --keep <years>
```

or

```
veeamconfig gfs set --jobname <job_name> yearly --on <month> --keep <years>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure the long-term retention policy. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the long-term retention policy.
- `<month>` – month when Veeam Agent must assign a yearly GFS flag to a full restore point – for example, `Jan` or `January`.
- `<years>` – number of years to keep the yearly GFS flag on the full restore point.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
yearly --on January --keep 3
```

## Enabling Long-Term Retention Policy

To start marking backups with specific GFS flags, you must enable the long-term retention policy for the job. Use the following command:

```
veeamconfig gfs enable --jobid <job_id> [--type <period>]
```

or

```
veeamconfig gfs enable --jobname <job_name> [--type <period>]
```

where:

- **<job\_id>** – ID of the backup job for which you want to enable the long-term retention policy. You should look up the job ID in advance, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job\_name>** – name of the backup job for which you want to enable the long-term retention policy.
- **<period>** – type of the long-term retention policy. Possible values: `weekly`, `monthly` or `yearly`. This parameter is optional. You can use it to enable a specific type of long-term retention policy. To enable several types of retention at once, specify all necessary retention types separated by a comma (',') – for example: `weekly,monthly`.

For example:

```
user@srv01:~$ veeamconfig gfs enable --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f --type monthly
```

You can disable the long-term retention policy for the job at any time. To learn more, see [Disabling Long-Term Retention Policy](#).

## Viewing Long-Term Retention Policy

To view the long-term retention policy defined for the backup job, use the following command:

```
veeamconfig gfs show --jobid <job_id>
```

or

```
veeamconfig gfs show --jobname <job_name>
```

where:

- **<job\_id>** – ID of the backup job for which you want to view the long-term retention policy.
- **<job\_name>** – name of the backup job for which you want to view the long-term retention policy.

Veeam Agent for Linux displays the following information about the backup job long-term retention policy:

Parameter	Description
<b>GFS state</b>	State of long-term retention policy. Possible values: <ul style="list-style-type: none"><li>• GFS is enabled</li><li>• GFS is disabled</li><li>• GFS is not set</li></ul>
<b>Enabled</b>	Possible values: <code>true</code> or <code>false</code> .
<b>Desired time</b>	Weekday, week number or month when Veeam Agent will set the GFS flag on the full restore point.
<b>Keep for</b>	Period of time for retaining the GFS flag on the full restore point.

The information listed in the table above is displayed for weekly, monthly and yearly retention policies.

For example:

```
user@srv01:~$ veeamconfig gfs show --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
GFS is enabled
Weekly:
  Enabled: true
  Desired time: Friday
  Keep for: 1 weeks
Monthly:
  Enabled: false
  Desired time: First
  Keep for: 1 months
Yearly:
  Enabled: false
  Desired time: January
  Keep for: 1 years
```

## Disabling Long-Term Retention Policy

You can disable all or specific types of the long-term retention policy: weekly, monthly or yearly.

## Disabling All Types of Long-Term Retention

To disable the long-term retention policy for the backup job, use the following command:

```
veeamconfig gfs disable --jobid <job_id>
```

or

```
veeamconfig gfs disable --jobname <job_name>
```

where:

- <job\_id> – ID of the backup job for which you want to disable the long-term retention policy.
- <job\_name> – name of the backup job for which you want to disable the long-term retention policy.

For example:

```
user@srv01:~$ veeamconfig gfs disable --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
```

## Disabling Specific Types of Long-Term Retention

To disable a specific type of the long-term retention policy for the backup job, use the following command:

```
veeamconfig gfs set --jobid <job_id> <period> --disable
```

or

```
veeamconfig gfs set --jobname <job_name> <period> --disable
```

where:

- <job\_id> – ID of the backup job for which you want to disable the long-term retention policy.
- <job\_name> – name of the backup job for which you want to disable the long-term retention policy.
- <period> – single type of the long-term retention policy you want to disable. Possible values: `weekly`, `monthly` or `yearly`.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f weekly --disable
```

## Configuring Active Full Backup Schedule

You can schedule a backup job to create active full backups periodically. You can specify active full schedule settings individually for every job created in Veeam Agent. You can perform the following actions with the active full backup schedule via the command-line interface:

- [Specify active full backup schedule.](#)
- [Enable active full backup schedule.](#)
- [View active full backup schedule.](#)

- [Disable active full backup schedule.](#)

#### TIP

You can also specify active full backup schedule for the backup job when you create the job. For details, see [Creating Volume-Level Backup Job](#) and [Creating File-Level Backup Job](#).

## Specifying Active Full Backup Schedule

You can configure the backup job to create active full backups on a [weekly](#) or [monthly](#) schedule.

After you define the active full backup schedule, Veeam Agent automatically enables this schedule for the specified backup job.

## Specifying Weekly Schedules

To instruct Veeam Agent to create an active full backup on specific week days, use the following command:

```
veeamconfig schedule set --jobid <job_id> --weekdays <days>
```

or

```
veeamconfig schedule set --jobname <job_name> --weekdays <days>
```

where:

- **<job\_id>** – ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job\_name>** – name of the backup job for which you want to configure the active full backup schedule.
- **<days>** – days when the backup job must create an active full backup separated by a comma (',') – for example: `Monday, Friday` or `Mon, Fri`.

For example:

```
user@srv01:~$ veeamconfig schedule activefull set --jobname DailyBackup --weekdays Monday, Friday
```

## Specifying Monthly Schedules

You can configure the backup job to create active full backups on specific months or every month.

- To create an active full backup monthly on a specific day of a specific week, use the following command:

```
veeamconfig schedule activefull set --jobid <job_id> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

or

```
veeamconfig schedule set --jobname <job_name> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

where:

- <job\_id> – ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job\_name> – name of the backup job for which you want to configure the active full backup schedule.
- <day> – days when the backup job must create an active full backup separated by a comma (',' ). For example: `Monday, Friday`. The backup job will create an active full backup on the specified days at the time specified in the backup job schedule settings.
- <week> – week of the month when the backup job must create an active full backup. Possible values: `First, Second, Third, Fourth` or `Last`.
- <months> – months when the backup job must create an active full backup separated by a comma (',' ) – for example: `Jan, Apr, Jul, Oct`. If you do not specify this option, the backup job will create an active full backup every month.

For example:

```
user@srv01:~$ veeamconfig schedule activefull set --jobname DailyBackup --monthlyweekday Mon --weeknumber Second --months Jan,Jul
```

- To configure the backup job to create an active full backup monthly on a specific day of the month, use the following command:

```
veeamconfig schedule set --jobid <job_id> --thisday <day> [--months <months>]
```

or

```
veeamconfig schedule set --jobname <job_name> --thisday <day> [--months <months>]
```

where:

- <job\_id> – ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job\_name> – name of the backup job for which you want to configure the active full backup schedule.
- <day> – day of the month when Veeam Agent must create an active full backup. Possible values range from 1 to 31 or `Last`.



- **<months>** – months when the backup job must create an active full backup separated by a comma (',') – for example: `Jan, Apr, Jul, Oct`. If you do not specify this option, the backup job will create an active full backup every month.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --thisday 21
--months Jan,Jul
```

## Viewing Active Full Backup Schedule

To view the active full backup schedule defined for the backup job, use the following command:

```
veeamconfig schedule activefull show --jobid <job_id>
```

or

```
veeamconfig schedule activefull show --jobname <job_name>
```

where:

- **<job\_id>** – ID of the backup job for which you want to view the active full backup schedule.
- **<job\_name>** – name of the backup job for which you want to view the active full backup schedule.

Veeam Agent for Linux displays the following information about the active full backup schedule:

Parameter	Description
<b>Every &lt;value&gt;</b>	Days on which the backup job creates active full backups. For example: <i>Every Sat</i> or <i>Every 1 day of every month</i> .
<b>Run automatically</b>	State of the active full backup schedule. Possible values: <ul style="list-style-type: none"> <li>• Enabled</li> <li>• Disabled</li> </ul>

For example:

```
user@srv01:~$ veeamconfig schedule activefull show --jobname DailyBackup
Every second Monday of every month
Run automatically: enabled
```

## Disabling Active Full Backup Schedule

To disable the active full backup schedule for the backup job, use the following command:

```
veeamconfig schedule activefull disable --jobid <job_id>
```

or

```
veeamconfig schedule activefull disable --jobname <job_name>
```

where:

- <job\_id> – ID of the backup job for which you want to disable the active full backup schedule.
- <job\_name> – name of the backup job for which you want to disable the active full backup schedule.

For example:

```
user@srv01:~$ veeamconfig schedule activefull disable --jobname DailyBackup
```

## Enabling Active Full Backup Schedule

After you specify active full backup schedule settings for the backup job, Veeam Agent automatically enables active full backup schedule for the job. You can also enable active full backup schedule manually, for example, if you previously disabled it for the backup job. To enable active full backup schedule, use the following command:

```
veeamconfig schedule activefull enable --jobid <job_id>
```

or

```
veeamconfig schedule activefull enable --jobname <job_name>
```

where:

- <job\_id> – ID of the backup job for which you want to enable the active full backup schedule. You should look up the job ID in advance, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job\_name> – name of the backup job for which you want to enable the active full backup schedule.

For example:

```
user@srv01:~$ veeamconfig schedule activefull enable --jobname DailyBackup
```

You can disable the schedule for the job at any time. To learn more, see [Disabling Backup Schedule](#).

# Configuring Health Check Schedule

You can schedule a periodic [health check](#) of a backup that resides in an object storage repository. You can specify backup health check schedule settings individually for every backup job created in Veeam Agent for Linux or backup policy created in Veeam Backup & Replication. For more information on configuring backup health check in a backup policy, see the [Maintenance Settings](#) topic of the Veeam Agent Management Guide.

You can perform the following actions with backup health check schedule in command line interface:

- [Specify health check schedule.](#)
- [Enable health check schedule.](#)
- [View health check schedule.](#)
- [Disable health check schedule.](#)

## Specifying Health Check Schedule

You can schedule a backup health check to run on a specific week day of a specific month or on specific days of the week.

## Specifying Monthly Health Check Schedule

To instruct Veeam Agent to perform backup health check on a specific week day of a month, use the following command:

```
veeamconfig healthcheck set --light --jobid <job_id> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

or

```
veeamconfig healthcheck set --light --jobname <job_name> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

where:

- **<job\_id>** – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you configure the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job\_name>** – name of the backup job for which you want to configure the health check schedule.

### TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

- **<day>** – day of the week when the backup job must perform health check – for example, `Tuesday` or `Tue`.
- **<week>** – week of the month when the backup job must perform health check. Possible values: `First`, `Second`, `Third`, `Fourth` or `Last`.

- `<months>` – months when the backup job must perform health check, separated by a comma (',' ) – for example: Jan, Apr, Jul, Oct. If you do not specify this option, the health check will run every month.

For example:

```
user@srv01:~$ veeamconfig healthcheck set --light --jobname SystemBackup --monthlyweekday Fri --weeknumber Last --months Mar,Jun,Sep,Dec
```

## Specifying Weekly Health Check Schedule

To instruct Veeam Agent to perform backup health check on specific week days, use the following command:

```
veeamconfig healthcheck set --light --jobid <job_id> --weekdays <days>
```

or

```
veeamconfig healthcheck set --light --jobname <job_name> --weekdays <days>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you configure the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the health check schedule.

### TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

- `<days>` – comma-separated list of days when the backup job must run backup health check. For example: Mon, Fri. The backup job will run the health check on the specified days at the time specified in the backup job schedule settings.

For example:

```
user@srv01:~$ veeamconfig healthcheck set --light --jobname "System Backup" --weekdays mon,fri
```

## Enabling Health Check Schedule

After you set health check schedule for a backup job, Veeam Agent automatically enables this backup health check schedule for the job. You can also enable health check schedule manually – for example, if you previously disabled it. To enable health check schedule, use the following command:

```
veeamconfig healthcheck enable --light --jobid <job_id>
```

or

```
veeamconfig healthcheck enable --light --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you run the command to enable the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the health check schedule.

#### TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

For example:

```
user@srv01:~$ veeamconfig healthcheck enable --light --jobname SystemBackup
```

You can disable health check schedule for a job at any time. To learn more, see [Disabling Health Check Schedule](#).

## Viewing Health Check Schedule

To view the health check schedule defined for a backup job, use the following command:

```
veeamconfig healthcheck show --jobid <job_id>
```

or

```
veeamconfig healthcheck show --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you run the command to view the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the health check schedule.

#### TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

Veeam Agent for Linux displays the following information about the health check schedule:

Parameter	Description
<b>Every &lt;value&gt;</b>	Days on which the backup job runs the health check – for example, <i>Every last Fri of every month</i> .
<b>Run health-check automatically</b>	State of the backup health check schedule. Possible values: <ul style="list-style-type: none"><li>• Enabled</li><li>• Disabled</li></ul>

For example:

```
user@srv01:~$ veeamconfig healthcheck show --jobname SystemBackup
Every last Fri of Mar, Jun, Sep, Dec
Run health check automatically: enabled (light)
```

## Disabling Health Check Schedule

To disable the health check schedule for a backup job, use the following command:

```
veeamconfig healthcheck disable --jobid <job_id>
```

or

```
veeamconfig healthcheck disable --jobname <job_name>
```

where:

- **<job\_id>** – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you run the command to disable the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job\_name>** – name of the backup job for which you want to configure the health check schedule.

### TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

For example:

```
user@srv01:~$ veeamconfig healthcheck disable --jobname SystemBackup
```

# Configuring Database Processing Settings

You can enable database processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux. With database processing settings enabled, Veeam Agent will create transactionally consistent backups of Veeam Agent machines that run database systems. For information on the limitations of database processing, see [Backup of Database Systems](#).

You can perform the following actions with database processing settings via the command-line interface:

- [Specify Oracle database processing settings](#)
- [Specify MySQL database processing settings](#)
- [Specify PostgreSQL database processing settings](#)
- [View database processing settings](#)
- [Disable database processing settings](#)

## Specifying Oracle Processing Settings

You can enable Oracle processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

To enable Oracle processing settings for the backup job, use the following command:

```
veeamconfig aap set oracle --jobid <job_id> <oracle_options>
```

or

```
veeamconfig aap set oracle --jobname <job_name> <oracle_options>
```

where:

- `<job_id>` – ID of the backup job for which you want to enable Oracle processing settings. You should look up the job ID in advance, before configuring Oracle processing settings, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to enable Oracle processing settings.
- `<oracle_options>` – Oracle processing settings for the backup job. To learn more, see [Oracle Processing Settings](#).

### TIP

To view IDs or names of all existent backup jobs, you can press the [Tab] key right after you type the `--jobid` or `--jobname` option.

# Oracle Processing Settings

You can specify the following Oracle processing settings for the backup job:

Option	Description and values
<b>--tryprocess</b>	Defines that Veeam Agent must continue the backup process if errors occur when processing the Oracle database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the Oracle database system.
<b>--prunelogs</b>	<p>The number of hours to keep Oracle archived logs or the size of archived logs to keep.</p> <ul style="list-style-type: none"><li>• If you want Veeam Agent to delete archived logs that are older than &lt;N&gt; hours, specify the necessary value in the &lt;N&gt;H format. For example, 10H.</li><li>• If you want Veeam Agent to delete archived logs that are larger than &lt;N&gt; GB, specify the necessary value in the &lt;N&gt;G format. For example: 10G.</li></ul> <p>Veeam Agent will wait for the backup job to complete successfully and then trigger archived logs truncation via Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next successful backup job session.</p>
<b>--usroracleos</b>	<p>Name of the Veeam Agent machine OS account. To connect to the Oracle database system, the account must be a member of the group that owns configuration files for the Oracle database (for example, the oinstall group).</p> <p>You do not need this option if you want to use the <a href="#">Oracle account</a> to connect to the database. Instead, specify the necessary account with the <code>--usroracledb</code> option.</p>
<b>--usroracledb</b>	<p>Name of the Oracle account. To connect to the Oracle database system, the account must have SYSDBA rights on the databases to be processed.</p> <p>You do not need this option if you want to use the <a href="#">OS account</a> to connect to the database. Instead, specify the necessary account with the <code>--usroracleos</code> option.</p>

For example:

```
user@srv01:~$ veeamconfig aap set oracle --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75 --tryprocess --prunelogs 10G --usroracledb system
```

## Specifying MySQL Processing Settings

You can enable MySQL processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.



## IMPORTANT

MySQL tables that use the MyISAM storage engine must be locked to keep them in consistent state while Veeam Agent is creating the system snapshot. To correctly process such tables, MySQL account must have the following instance-wide privileges:

- `SELECT`. This privilege enables Veeam Agent to access tables' metadata and select for a lock the tables that use the MyISAM storage engine. Without this privilege, the processing of the MySQL database system will run successfully but MyISAM tables will not be locked, which may result in an inconsistent state of the backed up data.
- `LOCK TABLES`. This privilege is required for locking the selected MyISAM tables. If some MyISAM tables are selected but the MySQL account does not have the `LOCK TABLES` privilege, the processing of the MySQL database system will fail.
- `RELOAD` or `FLUSH TABLES`. If some MyISAM tables are selected but the MySQL account does not have either `RELOAD` or `FLUSH TABLES` privilege, the processing of the MySQL database system will fail.

To obtain information about the privileges that are assigned to an account, use MySQL functionality, for example, the `SHOW GRANTS` statement. To learn more, see [MySQL documentation](#).

To enable MySQL processing settings for the backup job, use the following command:

```
veeamconfig aap set mysql --jobid <job_id> <mysql_options>
```

or

```
veeamconfig aap set mysql --jobname <job_name> <mysql_options>
```

where:

- `<job_id>` – ID of the backup job for which you want to enable MySQL processing settings. You should look up the job ID in advance, before configuring MySQL processing settings, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to enable MySQL processing settings.
- `<mysql_options>` – MySQL processing settings for the backup job. To learn more, see [MySQL Processing Settings](#).

## TIP

To view IDs or names of all existent backup jobs, you can press the [Tab] key right after you type the `--jobid` or `--jobname` option.

# MySQL Processing Settings

You can specify the following MySQL processing settings for the backup job:

Option	Description and values
<b>--tryprocess</b>	Defines that Veeam Agent must continue the backup process if errors occur when processing the MySQL database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the MySQL database system.
<b>--usrmysqldb</b>	Name of the MySQL account. Veeam Agent can connect to the MySQL database system in one of the following ways: <ul style="list-style-type: none"><li>• If you specify account name (<b>--usrmysqldb</b> option) only, Veeam Agent will prompt you to specify a password to access the MySQL database system.</li><li>• If you specify account name and <b>password</b> (<b>--usrmysqldb</b> and <b>--password</b> options), Veeam Agent will access the MySQL database system.</li><li>• If you do not specify account credentials (<b>--usrmysqldb</b> and <b>--password</b> options), Veeam Agent will use a password file to connect to the MySQL database system. To learn more about password file configuration, see <a href="#">Preparing Password File for MySQL Processing</a>.</li></ul>
<b>--password</b>	Password of the MySQL account. If you do not specify the <b>--password</b> value, Veeam Agent will prompt you to specify a password to access the MySQL database. Keep in mind, if you specify the password using the <b>--password</b> option, password is stored in terminal in plain text.
<b>--defaults-file</b>	<p>Path to a password file. You must specify a full path to a password file if you want Veeam Agent to use a password file located in specific directory. Specifying relative paths is not supported.</p> <p>With this method selected, you do not need to specify account credentials in the backup job settings.</p> <p>You do not need this option in the following cases:</p> <ul style="list-style-type: none"><li>• Veeam Agent uses account name and password that are specified in the backup job settings to connect to the MySQL database.</li><li>• Veeam Agent uses account credentials that are stored in the password file in <code>/root/.my.cnf</code>.</li></ul>

## Examples

Authentication with password:

```
user@srv01:~$ veeamconfig aap set mysql --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75 --tryprocess --usrmysqldb root --password P@ssw0rd
```

Authentication with password file:

```
user@srv01:~$ veeamconfig aap set mysql --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8
ea75 --tryprocess
--defaults-file /data/root/.my.cnf --password P@ssw0rd
```

## Preparing Password File for MySQL Processing

You can use MySQL account credentials that are stored in the password file to connect Veeam Agent for Linux to the MySQL database system.

### NOTE

Consider the following:

- If you specify a custom path to the password file, specify a full path. Specifying relative paths is not supported.
- The password file can also contain user-specific connection settings that Veeam Agent will apply to connect to the MySQL database system. For example, if you want to connect to the MySQL database system using the custom socket, specify the socket path in the password file. To learn more, see [MySQL documentation](#).

If you want to use a password file for authentication, create a file. By default, Veeam Agent expects the password file to have the `.my.cnf` name and to be in the home directory of the `root` user. If the password file has a custom name or is stored in another directory, you can specify a custom path.

The password file must have the following contents:

```
[client]
user=<username>
password=<password>
```

where:

- `<username>` – name of the account that Veeam Agent will use to connect to the MySQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the MySQL database system.

For example:

```
[client]
user=root
password=P@ssw0rd
```

## Specifying PostgreSQL Processing Settings

You can enable PostgreSQL processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

To enable PostgreSQL processing settings for the backup job, use the following command:

```
veeamconfig aap set postgres --jobid <job_id> <postgres_options>
```

or

```
veeamconfig aap set postgres --jobname <job_name> <postgres_options>
```

where:

- **<job\_id>** – ID of the backup job for which you want to enable PostgreSQL processing settings. You should look up the job ID in advance, before configuring PostgreSQL processing settings, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job\_name>** – name of the backup job for which you want to enable PostgreSQL processing settings.
- **<postgres\_options>** – PostgreSQL processing settings for the backup job. To learn more, see [PostgreSQL Processing Settings](#).

#### TIP

To view IDs or names of all existent backup jobs, you can press the [Tab] key right after you type the `--jobid` or `--jobname` option.

## PostgreSQL Processing Settings

You can specify the following PostgreSQL processing settings for the backup job:

Option	Description and values
<b>--tryprocess</b>	Defines that Veeam Agent must continue the backup process if errors occur when processing the PostgreSQL database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the PostgreSQL database system.
<b>--usrpgdb</b>	<p>Name of the PostgreSQL account.</p> <p>If you use a password file to connect to the PostgreSQL database system, the <code>--usrpgdb</code> option allows to select the user from the password file.</p> <p>You do not need this option if you want to use a Veeam Agent machine OS account to connect to the PostgreSQL database system. Instead, specify the <a href="#">OS account</a> with the <code>--usrpgos</code> option.</p>

Option	Description and values
<b>--password</b>	<p>Password of the PostgreSQL account.</p> <p>If you do not specify this option, Veeam Agent will prompt to enter the password. If you do not specify the password in prompt, Veeam Agent uses a password file to connect to the PostgreSQL database system. To learn more about password file configuration, see <a href="#">Preparing Password File for PostgreSQL Processing</a>.</p> <p>Keep in mind, if you specify the password using the <code>--password</code> option, password is stored in terminal in plain text.</p>
<b>--usrpgos</b>	<p>Name of the OS account. Veeam Agent will use the name to connect to the PostgreSQL database system using the peer authentication method. In the peer authentication method, Veeam Agent uses the OS account as the PostgreSQL database user name. With this option selected, you must specify OS account only. To learn more about peer authentication, see <a href="#">PostgreSQL documentation</a>.</p> <p>You do not need this option if you want to use a PostgreSQL account to connect to the database system. Instead, specify the <a href="#">PostgreSQL account</a> with the <code>--usrpgdb</code> option.</p>

For example:

```
user@srv01:~$ veeamconfig aap set postgres --jobid 29bc2e1a-e35c-4efb-8d37-b717
7b8ea75 --tryprocess --usrpgdb postgres --password P@ssw0rd
```

## Preparing Password File for PostgreSQL Processing

You can use PostgreSQL account credentials that are stored in the password file to connect Veeam Agent to the PostgreSQL database system.

If you want to use a password file for authentication, create the `.pgpass` file in the home directory of the `root` user.

The password file must have the following contents:

```
<hostname>:<port>:<database>:<username>:<password>
```

where:

- `<hostname>` – name of the host where the PostgreSQL database system is located.
- `<port>` – number of the free port that Veeam Agent will use to connect to the PostgreSQL database system.
- `<database>` – name of the PostgreSQL database.
- `<username>` – name of the account that Veeam Agent will use to connect to the PostgreSQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the PostgreSQL database system.

For example:

```
srv01:5432:mydb:postgres:P@ssw0rd
```

For more information about the password file, see [PostgreSQL documentation](#).

## Viewing Database Processing Settings

To view database processing settings defined for the backup job, use the following command:

```
veeamconfig aap show --jobid <job_id>
```

or

```
veeamconfig aap show --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to view database processing settings.
- `<job_name>` – name of the backup job for which you want to view database processing settings.

Veeam Agent for Linux displays the following information about database processing settings:

- [Oracle processing settings](#)
- [MySQL processing settings](#)
- [PostgreSQL processing settings](#)

## Oracle Processing Settings

Parameter	Description
Oracle processing	<p>Oracle processing settings status. Possible values:</p> <ul style="list-style-type: none"><li>• <i>Required</i> – Oracle processing settings are enabled for the job. If an error occurs when processing the Oracle database system, Veeam Agent will stop the backup process.</li><li>• <i>Try</i> – Oracle processing settings are enabled for the job. If an error occurs when processing the Oracle database system, Veeam Agent will continue the backup process.</li><li>• <i>Disabled</i> – Oracle processing settings are disabled for the job using command line interface.</li></ul>

Parameter	Description
<b>Account used for processing</b>	Account used to connect to the Oracle database. Possible values: <ul style="list-style-type: none"> <li><i>System account (username: &lt;username&gt;)</i> – if Veeam Agent connects to the Oracle database system with the account of the Veeam Agent machine OS.</li> <li><i>Oracle account (username: &lt;username&gt;)</i> – if Veeam Agent connects to the Oracle database system with the Oracle account.</li> </ul> where <username> is a name of the user account that Veeam Agent will use to connect to the Oracle database.
<b>Delete logs over &lt;N&gt; Gb</b>	Veeam Agent displays this information if Veeam Agent is set to delete archived logs that are larger than <N> GB.
<b>Delete logs older &lt;N&gt; Hr</b>	Veeam Agent displays this information if Veeam Agent is set to delete archived logs that are older than <N> hours.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
Oracle processing: required
  Account used for processing: Oracle account (username: sys)
  Delete logs over 10 Gb
```

## MySQL Processing Settings

Parameter	Description
<b>MySQL processing</b>	MySQL processing settings status. Possible values: <ul style="list-style-type: none"> <li><i>Required</i> – MySQL processing settings are enabled for the job. If an error occurs when processing a MySQL database, Veeam Agent will stop the backup process.</li> <li><i>Try</i> – MySQL processing settings are enabled for the job. If an error occurs when processing a MySQL database, Veeam Agent will continue the backup process.</li> <li><i>Disabled</i> – MySQL processing settings are disabled for the job using command line interface.</li> </ul>
<b>Account used for processing</b>	Veeam Agent displays this information if Veeam Agent is set to connect to the MySQL database system with the account name and password.
<b>Path to a password file</b>	Veeam Agent displays this information if Veeam Agent is set to connect to the MySQL database system with the account credentials that are stored in the password file.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
MySQL processing: required
Account used for processing: username: root
```

## PostgreSQL Processing Settings

Parameter	Description
<b>PostgreSQL processing</b>	<p>PostgreSQL processing settings status. Possible values:</p> <ul style="list-style-type: none"><li>• <i>Required</i> – PostgreSQL processing settings are enabled for the job. If an error occurs when processing a PostgreSQL database, Veeam Agent will stop the backup process.</li><li>• <i>Try</i> – PostgreSQL processing settings are enabled for the job. If an error occurs when processing a PostgreSQL database, Veeam Agent will continue the backup process.</li><li>• <i>Disabled</i> – PostgreSQL processing settings are disabled for the job using command line interface.</li></ul>
<b>Account used for processing</b>	<p>Account used to connect to the PostgreSQL database. Possible values:</p> <ul style="list-style-type: none"><li>• <i>&lt;username&gt; (password)</i> – Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the PostgreSQL account.</li><li>• <i>&lt;username&gt; (file)</i> – Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the password file.</li><li>• <i>&lt;username&gt; (peer)</i> – Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the Veeam Agent machine OS account.</li></ul> <p>where <i>&lt;username&gt;</i> is a name of the account that Veeam Agent will use to connect to the PostgreSQL database.</p>

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
PostgreSQL processing: required
Account used for processing: postgres (password)
```



## Disabling Database Processing Settings

To disable database processing settings defined for the backup job, use the following command:

```
veeamconfig aap disable <db_sys> --jobid <job_id>
```

or

```
veeamconfig aap disable <db_sys> --jobname <job_name>
```

where:

- `<db_sys>` – name of the database system that you want to disable. Possible values:
  - *oracle* – Oracle database processing to be disabled.
  - *mysql* – MySQL database processing to be disabled.
  - *postgres* – PostgreSQL database processing to be disabled.
- `<job_id>` – ID of the backup job for which you want to disable database processing settings.
- `<job_name>` – name of the backup job for which you want to disable database processing settings.

For example:

```
user@srv01:~$ veeamconfig aap disable oracle --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea759
Oracle processing was disabled.
```

# Starting and Stopping Backup Jobs

You can start a backup job manually at any time you need, for example, if you want to create an additional restore point for Veeam Agent backup and do not want to change the job schedule. You can also stop the running backup job before the job session completes, if necessary.

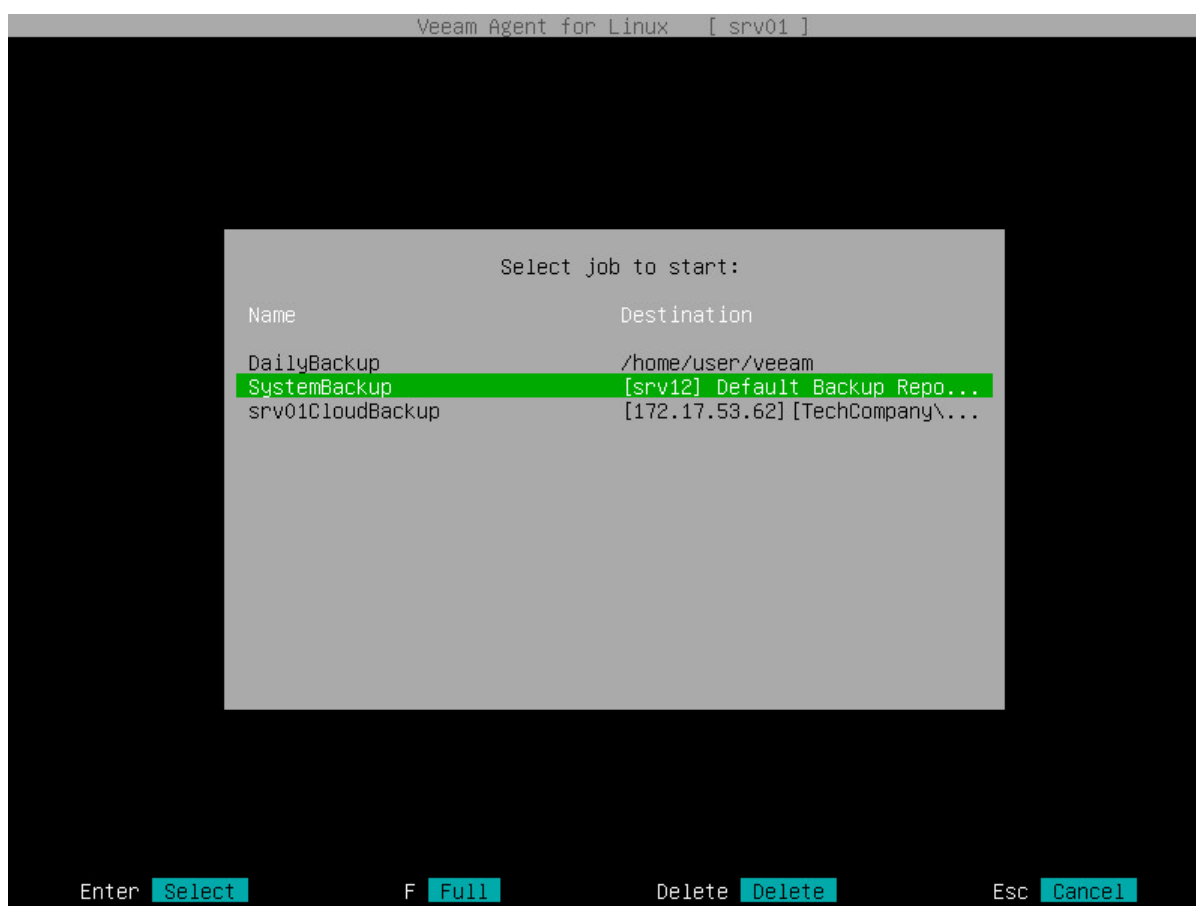
You can start and stop backup jobs in one of the following ways:

- [With the Veeam Agent control panel.](#)
- [With the Veeam Agent command line interface.](#)

# Starting Backup Job from Control Panel

To start a backup job with the Veeam Agent control panel, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. Press the [S] key to open the **Select job to start** dialog window.
3. Select the necessary backup job in the list and start the job in on of the following ways:
  - To start an incremental backup job session, press [Enter].
  - To create an active full backup, press [F].



4. Veeam Agent will immediately start the backup job and display a notification window informing that the job has been started. Press [Enter] to close the window and proceed to the list of backup job sessions.

You can monitor the backup job performance in the Veeam Agent control panel. To learn more, see [Viewing Real-Time Job Session Statistics](#).

If you start the backup job while another backup job is running, Veeam Agent will perform the backup job immediately after the current job is completed. For details, see [Job Queue](#).

# Job Queue

If another backup job is running when you start the backup job, Veeam Agent will submit this backup job to job queue. Veeam Agent will perform the job in the queue as soon as the previous job is completed.

Veeam Agent for Linux [ srv01 ]

Latest backup sessions:

Job name	State	Started at	Finished at
DailyBackup	Pending (0%)	---	---
SystemBackup	Running (16%)	2023-08-07 18:58:55	---

Info

The backup job has been added to the queue.

[Ok]

Enter Show

C Configure

S Start Job

R Recover Files

M Misc

Esc Quit

The queued backup job creates a new session with the *Pending* status. You can view all jobs in the queue in the **Latest backup sessions** list in the Veeam Agent control panel.

Veeam Agent for Linux [ srv01 ]

Latest backup sessions:

Job name	State	Started at	Finished at
DailyBackup	Pending (0%)	---	---
SystemBackup	Running (19%)	2023-08-07 18:58:55	---

Enter Show C Configure S Start Job R Recover Files M Misc Esc Quit

- NOTE**
- Consider the following about job queue:
- Job queue can contain up to 3 backup jobs besides the job that is already running.
  - You cannot submit the same backup job to the queue if it is already running.

# Starting Backup Job from Command Line Interface

You can start a backup job with the command line interface. When you start a backup job, Veeam Agent initiates a new backup job session and provides you with a Session ID. You can monitor the progress of the backup job session or view the session status.

To start a backup job, use the following command:

```
veeamconfig job start --name <job_name>
```

or

```
veeamconfig job start --id <job_id>
```

where:

- `<job_name>` – name of the backup job that you want to start.
- `<job_id>` – ID of the backup job that you want to start.

## TIP

Consider the following:

- You can use the `veeamconfig job start` command with the `--nosnap` option to start a file-level backup job. In this case, Veeam Agent will not create a snapshot of the backed-up volume during the backup job session. Keep in mind that the snapshot-less file-level backup does not guarantee that data in the backup is consistent. To learn more, see [Snapshot-Less File-Level Backup](#).
- You can use the `veeamconfig job start` command with the `--activefull` option to create active full backups. To learn more, see [Creating Active Full Backups](#).

For example:

```
$ veeamconfig job start --name SystemBackup
Backup job has been started.
Session ID: [{381532f7-426a-4e89-b9fc-43d98942c71a}].
Logs stored in: [/var/log/veeam/Backup/SystemBackup/Session_20161207_162608_{381532f7-426a-4e89-b9fc-43d98942c71a}].
```

You can [check the backup job session status](#) or [view the backup job session log](#) using the Veeam Agent command line interface.

You can also monitor the backup job performance in the Veeam Agent control panel. To learn more, see [Viewing Real-Time Job Session Statistics](#).

If you start the backup job while another backup job is running, Veeam Agent will perform the backup job immediately after the current job is completed. For details, see [Job Queue](#).

## Job Queue

If another backup job is running when you start the backup job, Veeam Agent will submit this backup job to job queue. Veeam Agent will perform the job in the queue as soon as the previous job is completed.

```
$ veeamconfig job start --name DailyBackup
The backup job has been added to the queue.
Session ID: [{10e8c599-b2aa-4008-89d9-af9b6e04aeba}].
Logs stored in: [/var/log/veeam/Backup/DailyBackup/Session_20230814_153342_{10e8c599-b2aa-4008-89d9-af9b6e04aeba}].
```

The queued backup job creates a new session with the *Pending* status. You can view all jobs in the queue by running the `veeamconfig session list` command.

```
$ veeamconfig session list
Job name      Type      ID                               State      Started
at           Finished at
SystemBackup  Backup    {37427202-b139-4b36-9982-e0c33894d0cc} Running    2023-08
-14 15:33
DailyBackup   Backup    {10e8c599-b2aa-4008-89d9-af9b6e04aeba} Pending
```

### NOTE

Consider the following about job queue:

- Job queue can contain up to 3 backup jobs besides the job that is already running.
- You cannot submit the same backup job to the queue if it is already running.

# Creating Active Full Backups

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

Before you create an active full backup, check the following prerequisites:

- The backup job must be configured.
- You cannot create an active full backup if a backup task of any type is currently running.

To perform active full backup, use the following command:

```
veeamconfig job start --name <job_name> --activefull
```

or

```
veeamconfig job start --id <job_id> --activefull
```

where:

- <job\_name> – name of the backup job that you want to start to create an active full backup.
- <job\_id> – ID of the backup job that you want to start to create an active full backup.

For example:

```
$ veeamconfig job start --name SystemBackup --activefull
Backup job has been started.
Session ID: [{ce864e24-8211-4df7-973a-741adce96fe7}].
Logs stored in: [/var/log/veeam/Backup/SystemBackup/Session_20180611_150046_{ce
864e24-8211-4df7-973a-741adce96fe7}].
```

You can view the progress for the active full backup session in the same way as for any other backup job session. In particular, you can [check the backup job session status](#) or [view the backup job session log](#) using the Veeam Agent command line interface.

You can also monitor the backup job performance in the Veeam Agent control panel. To learn more, see [Viewing Real-Time Job Session Statistics](#).



# Stopping Backup Job

You can stop the running backup job before the job session completes, for example, if the backup process is about to take long, and you do not want the job to produce workload on the production environment during business hours.

When you stop a backup job, the job session will finish immediately. Veeam Agent will not produce a new restore point during the session, and the session will finish with the *Failed* status.

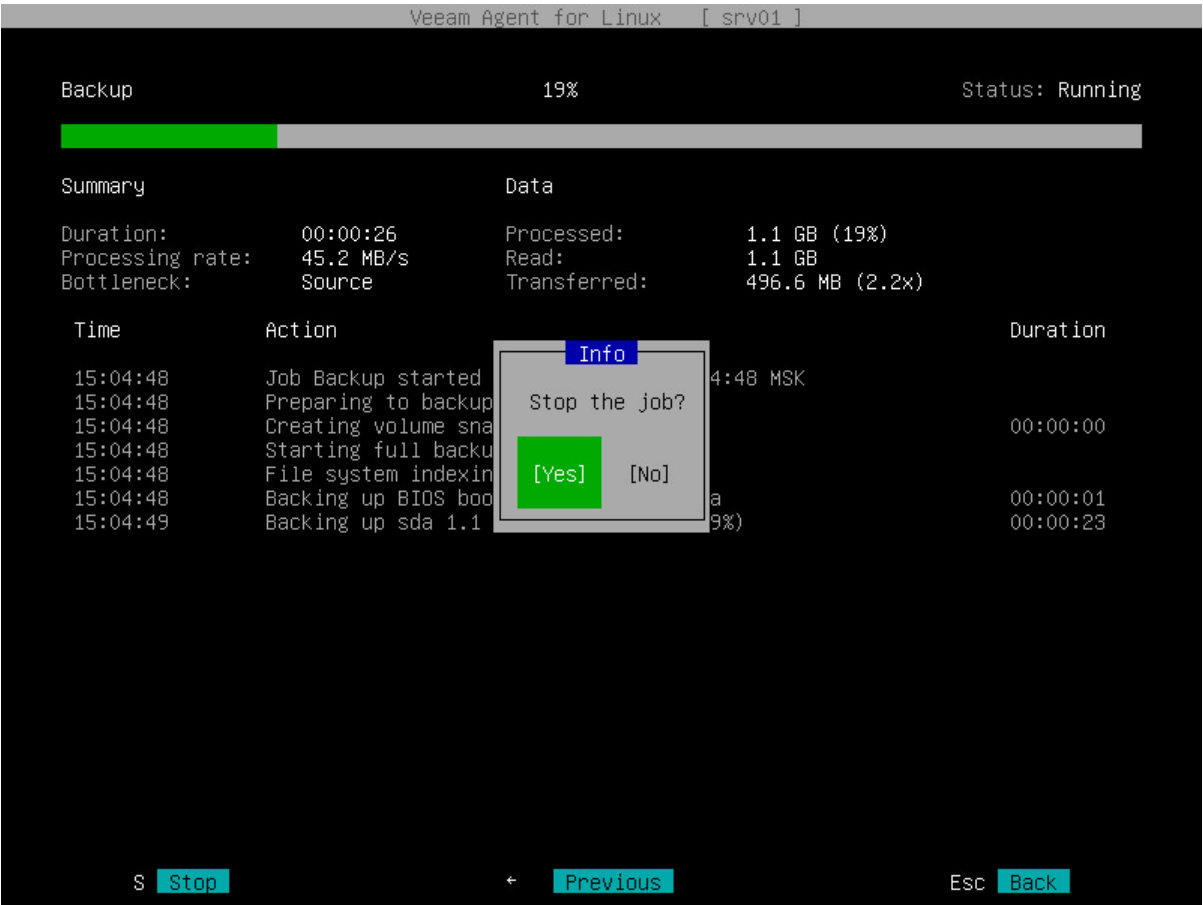
You can stop a job in one of the following ways:

- [With the control panel](#)
- [With the command line interface](#)

## Stopping Job from Control Panel

To stop a backup job:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, select the currently running session with the [Up] and [Down] keys and press [Enter].
3. In the session statistics window, press [S].
4. In the displayed window, make sure that the **Yes** button is selected and press [Enter].



# Stopping Job from Command Line Interface

To stop a backup job, use the following command:

```
veeamconfig session stop --id <session_id>
```

or

```
veeamconfig session stop --force --id <session_id>
```

where:

- `<session_id>` – ID of the currently running backup job session that you want to stop.
- `--force` – with this option enabled, Veeam Agent will immediately stop the backup session even if it is unable to stop the *veeamjobman* process for some reason.

For example:

```
$ veeamconfig session stop --id 381532f7-426a-4e89-b9fc-43d98942c71a  
Session has stopped.
```

# Managing Backup Jobs

You can perform the following actions with backup jobs configured in Veeam Agent for Linux:

- [View the list of configured backup jobs.](#)
- [View information about the backup job settings.](#)
- [Edit the backup job settings.](#)
- [Delete a backup job.](#)

# Viewing List of Backup Jobs

To view a list of backup jobs configured in Veeam Agent for Linux, use the following command:

```
veeamconfig job list
```

In the list of backup jobs, Veeam Agent for Linux displays the following information:

Parameter	Description
Name	Name of the backup job.
ID	ID of the backup job.
Repository	Name of the backup repository that is specified as a backup storage for the backup job.

For example:

```
user@srv01:~$ veeamconfig job list
Name                               ID                               Repository
SystemBackup                      {2495911e-58db-4452-b4d1-f53dcfbc600e} Repository_1
DocumentsBackup                   {bcf821e6-b35f-4d57-b1c3-d3a477605cb9} Repository_1
HomePartitionBackup               {2aaa8c71-2434-4f12-a168-3d8e225fa416} Repository_2
```

# Viewing Backup Job Settings

To view detailed information about the backup jobs settings, use the following command:

```
veeamconfig job info --name <job_name>
```

or

```
veeamconfig job info --id <job_id>
```

where:

- **<job\_name>** – name of the backup job for which you want to view settings.
- **<job\_id>** – ID of the backup job for which you want to view settings.

Veeam Agent for Linux displays the following information about the backup job:

Parameter	Description
<b>ID</b>	ID of the backup job.
<b>Name</b>	Name of the backup job.
<b>Repository ID</b>	ID of the backup repository that is specified as a backup storage for the backup job.
<b>Repository name</b>	Name of the backup repository that is specified as a backup storage for the backup job.
<b>Creation time</b>	Date and time of the backup job creation.
<b>Compression</b>	Data compression level. Possible values are: <ul style="list-style-type: none"><li>• <i>0</i> – No compression</li><li>• <i>1</i> – Rle</li><li>• <i>2</i> – Lz4</li><li>• <i>3</i> – Zstd 3</li><li>• <i>4</i> – Zstd 9</li></ul>
<b>Max Points</b>	Number of restore points to keep on disk. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.
<b>Index</b>	File system indexing options defined for the backup job.

Parameter	Description
<b>Objects for backup</b>	Backup scope specified for the backup job.

For example:

```

user@srv01:~$ veeamconfig job info --name SystemBackup
Backup job
  ID: {2495911e-58db-4452-b4d1-f53dcfbc600e}
  Name: SystemBackup
  Repository ID: {4557ef7a-9c44-4f28-b8d0-44d78e5ddd5d}
  Repository name: Repository_1
  Creation time: 2017-04-06 13:29:03
  Options:
    Compression: Lz4
    Max Points: 7
    Index all mounted filesystems on the volumes selected for backup
  Objects for backup:
    Include Disk: sda1

```

# Editing Backup Job Settings

If you want to change settings of the backup job, you can edit it at any time. For example, you may want to edit the backup job to add a new directory to the backup scope or change the target location.

To edit a backup job, use the following command:

*For volume-level backup jobs*

```
veeamconfig job edit volumelevel <option> for --name <job_name>
```

or

```
veeamconfig job edit volumelevel <option> for --id <job_id>
```

*For file-level backup jobs*

```
veeamconfig job edit filelevel <option> for --name <job_name>
```

or

```
veeamconfig job edit filelevel <option> for --id <job_id>
```

where:

- `<option>` – option that you want to edit for the job. You can specify one or several options at a time. To learn more about available options, see [Backup Job Options](#).
- `<job_name>` – name of the backup job that you want to edit.
- `<job_id>` – ID of the backup job that you want to edit.

For example:

```
user@srv01:~$ veeamconfig job edit volumelevel --name SystemVolumeBackup for --  
name SystemVolume
```

# Backup Job Options

You can use the following options to edit parameters for the backup job:

Option	Description and values
<b>--compressionlevel</b>	Data compression level. Possible values are: <ul style="list-style-type: none"><li>• <i>0</i> – No compression</li><li>• <i>1</i> – Rle</li><li>• <i>2</i> – Lz4</li><li>• <i>3</i> – ZlibLow</li><li>• <i>4</i> – ZlibHigh</li></ul>
<b>--blocksize</b>	Data block size in kilobytes. Possible values are 256, 512, 1024, 4096 or 8192.  The default value is <i>1024</i> .
<b>--maxpoints</b>	Number of restore points that you want to store in the backup location. By default, Veeam Agent keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent will remove the earliest restore point from the backup chain.
<b>--prefreeze</b>	Pre-freeze command that should be executed before the snapshot creation.
<b>--postthaw</b>	Post-thaw command that should be executed after the snapshot creation.
<b>--objects</b>	Object that should be included in backup: <ul style="list-style-type: none"><li>• For simple volumes – name of a block device that represents a volume or an entire disk that should be included in backup. You can specify entire disk to create backup of the entire computer image or individual computer volumes to create backup of specific volumes. If you want to back-up several disks or volumes, specify them one after another using a ',' (comma) character as a separator.</li><li>• For LVM volumes – name of an LVM logical volume that should be included in backup. If you want to back-up several LVM logical volumes, specify them one after another using a ',' (comma) character as a separator.</li></ul> This option is available for volume-level backup jobs only.
<b>--includedirs</b>	Full path to a directory that should be included in backup, for example: <i>/home/user</i> . The option is available for file-level backup jobs only.  You can specify one or several paths to directories in the computer file system. To separate several paths, use a ',' (comma) character, for example: <i>/home/user/Documents, /home/user/reports</i> .



Option	Description and values
<b>--excludedirs</b>	<p>Full path to a directory that should be excluded from backup. The option is available for file-level backup jobs only.</p> <p>The directory specified with this option must be a subdirectory of the directory specified with the <code>--includedirs</code> option. To separate several paths, use a ',' (comma) character, for example, <code>/home/user/Documents,/home/user/reports</code>.</p>
<b>--includemasks</b>	<p>Mask for file name or path that should be included in backup. The option is available for file-level backup jobs only.</p> <p>You can use the following UNIX wildcard characters for file name masks:</p> <ul style="list-style-type: none"> <li>'*' — a substitution for one or more characters in the file name or path. Can be used for any sequence of characters (including no characters). For example, <code>*.pdf</code>.</li> <li>'?' — a substitution of one character in the file name or path. For example, <code>repor?.pdf</code>.</li> <li>'[]' — a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: <code>report_201[3456].pdf</code> or <code>report_201[3-6].pdf</code>.</li> </ul> <p>To separate several masks, use a ',' (comma) character, for example, <code>report.*,reports.*</code>.</p> <p>File inclusion option is applied to all directories that are specified with the <code>--includedirs</code> option. For example, if you include in backup the <code>/home/user/Documents</code> directory and files that match the <code>repor?.pdf</code> file name mask, Veeam Agent will back up the <code>/home/user/Documents/report.pdf</code> file and will not back up the <code>/home/user/reports/report.pdf</code> file.</p>

Option	Description and values
<b>--excludemasks</b>	<p>Mask for file name or path that should be excluded from backup. The option is available for file-level backup jobs only.</p> <p>You can use the following UNIX wildcard characters for file name masks:</p> <ul style="list-style-type: none"> <li>'*' – a substitution for one or more characters in the file name or path. Can be used for any sequence of characters (including no characters). For example, *.pdf.</li> <li>'?' – a substitution of one character in the file name or path. For example, repor?.pdf.</li> <li>'[' – a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: report_201[3456].pdf or report_201[3-6].pdf.</li> </ul> <p>To separate several masks, use a ',' (comma) character, for example, report.*,reports.*.</p> <p>File exclusion option is applied to all directories that are specified with the --includedirs option and files that match file name masks specified with the --includemasks option. For example, you may want to specify the following backup scope for the backup job:</p> <ul style="list-style-type: none"> <li>Include in backup the /home/user/Documents directory</li> <li>Include files that match the report.* file name mask</li> <li>Exclude files that match the *.odt file name mask.</li> </ul> <p>In this case, Veeam Agent will backup the /home/user/Documents/report.pdf file and will not backup /home/user/Documents/report.odt and /home/user/reports/report.pdf files.</p> <p>If you want to use several name masks, you must specify them in double quotation marks, for example: veeamconfig job create filelevel --name BackupJob1 --reponame vault13 --includedirs /home --includemasks "*.bak,*.pdf".</p>
<b>--indexnothing</b>	Defines that file system indexing options are disabled for the backup job.
<b>--indexall</b>	Defines that Veeam Agent must index all files on the volumes included in backup.
<b>--indexonly</b>	Path to a directory that contains files that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character. The option is available for volume-level backup jobs only.
<b>--indexexcept</b>	Path to a directory that contains files that you do not want to index. You can specify one or more paths. To separate several paths, use the ',' (comma) character. The option is available for volume-level backup jobs only.

Option	Description and values
<b>--setencryption</b>	Defines that data encryption option is enabled for the job. You can use this option to enable encryption for the existing backup job or change a password used for encryption for the backup job. When you use the <code>veeamconfig job edit</code> command with the <code>--setencryption</code> option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.
<b>--resetencryption</b>	Defines that data encryption option is disabled for the job. You can use this option to disable encryption for the existing backup job.
<b>--deleteold</b>	<p>The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999.</p> <p>If you do not specify the <code>--deleteold</code> option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually.</p> <p>If you specified the value earlier and want to disable this setting, specify the <i>false</i> value for this option: <code>--deleteold false</code>. After the next successful backup session, this setting will be disabled for the backup in the target location.</p>
<b>--nosnap</b>	<p>Defines whether Veeam Agent must perform backup in the snapshot-less mode. Possible values:</p> <ul style="list-style-type: none"> <li><i>true</i> – if you use this option, Veeam Agent will create a snapshot of the backed-up volumes during file-level backup.</li> <li><i>false</i> – if you use this option, Veeam Agent will not create a snapshot of the backed-up volumes during file-level backup.</li> </ul> <p>Keep in mind that the snapshot-less file level backup does not guarantee that data in the backup is consistent. To learn more, see <a href="#">Snapshot-Less File-Level Backup</a>.</p>

## NOTE

Consider the following:

- If you change the target location for the backup job, during the next backup job session Veeam Agent for Linux will perform full data backup. All subsequent backup sessions will produce incremental backups – Veeam Agent for Linux will copy only changed data to the target location and add a new incremental backup file to the backup chain.
- If you change the backup scope for the backup job, during the next backup job session Veeam Agent for Linux will create a new incremental backup. The backup will contain all data blocks pertaining to new data added to the backup scope and changed data blocks pertaining to original data in the backup scope (data that was processed by the job at the time before you changed the backup scope).
- If you enable or disable encryption for the existing backup job that has already created one or more restore points, during the next job session, Veeam Agent for Linux will create active full backup.
- Full backup takes much more time than incremental backup. If you change the target location, you can copy an existing backup chain to the new location manually. In this case, the new backup job session will produce an incremental backup file and add it to the backup chain.

# Deleting Backup Job

You can delete a backup job configured in Veeam Agent for Linux. When you delete a backup job, backup files created by this job remain intact on the backup repository.

You can delete backup jobs in one of the following ways:


- With the Veeam Agent for Linux control panel
- With the Veeam Agent for Linux command line interface

## Deleting Backup Job with Control Panel

You can delete a backup job with the Veeam Agent control panel.

To delete a backup job:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. Press the [C] key to open the **Select job to configure** dialog window or the [S] key to open the **Select job to start** dialog window.
3. Select the necessary backup job in the list and press [Delete].
4. In the displayed notification window, make sure that the **Yes** button is selected and press [Enter].



# Deleting Backup Job with Command Line Interface

You can delete a backup job with the Veeam Agent command line interface. To delete a backup job, use the following command:

```
veeamconfig job delete --name <job_name>
```

or

```
veeamconfig job delete --id <job_id>
```

where:

- <job\_name> – name of the backup job that you want to delete.
- <job\_id> – ID of the backup job that you want to delete.

For example:

```
$ veeamconfig job delete --name SystemBackup
```

# Managing Backup Repositories

A backup repository is a storage location where Veeam Agent for Linux keeps backup files. You can use the following types of storage as a target location for a backup job:

- Local (internal) storage of the protected machine (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share.
- Object storage repository, such as S3 Compatible storage, Amazon S3, Google Cloud or Microsoft Azure Blob.
- 12.1 or later backup repository (including deduplication appliances).
- Veeam Backup & Replication 12.0 or later cloud repository.

## IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

# Creating Backup Repository

Veeam Agent for Linux creates backup files in a backup repository. When you create a backup job with the Backup Job wizard, you must specify a target location for the backup. Veeam Agent will create a backup repository in the specified location and save information about this repository in the database.

## IMPORTANT

A backup repository must be created on a separate volume from a volume whose data you plan to back up.

If you want to create backups in local directory, network shared folder or object storage, you must create a repository. To learn more, see the following sections:

- [Creating a repository in a local directory.](#)
- [Creating a repository in an NFS network shared folder.](#)
- [Creating a repository in an SMB network shared folder.](#)
- [Creating a repository in object storage.](#)

If you want to create Veeam Agent backups in a Veeam backup repository or cloud repository, you do not need to create repositories. Before configuring the backup job, you need to connect to the Veeam backup server or Veeam Cloud Connect service provider. To learn more, see the following sections:

- [Connecting to Veeam Backup Server](#)
- [Connecting to Service Provider](#)

## Creating Repository in Local Directory

To create a repository in a local directory, use the following command:

```
veeamconfig repository create --name <repository_name> --location <path_to_repository>
```

where:

- `<repository_name>` – name of the repository.
- `<path_to_repository>` – path to the directory in which backup files will be stored.

For example:

```
$ veeamconfig repository create --name VeeamBackup --location /home/backups
```



## Creating Repository in NFS Share

To create a repository in an NFS share, use the following command:

```
veeamconfig repository create --name <repository_name> --type nfs --location <path_to_repository>
--options <mounting_options>
```

where:

- `<repository_name>` – name of the backup repository.
- `<path_to_repository>` – path to the network shared folder where backup files will be stored in the *SERVER:/DIRECTORY* format.
- `<mounting_options>` – additional options that Veeam Agent will use to mount the network shared folder to the Veeam Agent machine file system. You can use the standard Linux `mount` command content as mounting options. This parameter is optional.

For example:

```
$ veeamconfig repository create --name VeeamBackup --type nfs --location srv01:/VeeamRepository --options vers=3,hard,retry=1
```

### TIP

If you mount a network shared folder to a directory in the Veeam Agent machine file system in advance, you can create the backup repository in the same way as in a local directory. For details, see [Creating Repository in Local Directory](#).

## Creating Repository in SMB Share

To create a repository in an SMB share, use the following command:

```
veeamconfig repository create --name <repository_name> --type smb --location <path_to_repository>
--username <user_name> --password --domain <domain> --options <mounting_options>
```

where:

- `<repository_name>` – name for the backup repository.
- `<path_to_repository>` – path to the network shared folder where backup files will be stored in the *//SERVER/DIRECTORY* format.
- `<user_name>` – account name that Veeam Agent will use to access the SMB network shared folder.
- `<domain>` – domain in which the account that has access permissions on the shared folder is registered.

- `<mounting_options>` – options that Veeam Agent will use to mount the network shared folder to the Veeam Agent machine file system. You can use the standard Linux `mount` command content as mounting options. This parameter is optional.

You can specify account name and domain for the SMB network shared folder using the `--username` and `--domain` parameters. If a password is required to access the network shared folder, you must also specify the `--password` parameter. When you run the `veeamconfig repository create` command, Veeam Agent will prompt you to type a password of the specified account.

Alternatively, you can specify account name, password and domain for the network shared folder as values for the `--options` parameter. Mind that these values will override values of the `--username`, `--password` and `--domain` parameters.

## Examples

Command with `--username`, `--password` and `--domain` parameters:

```
$ veeamconfig repository create --name VeeamBackup --type smb --location //srv02/VeeamRepository --username Administrator --password --domain srv02
```

Command with `--options` parameter:

```
$ veeamconfig repository create --name VeeamBackup --type smb --location //srv02/VeeamRepository --options username=Administrator,password=P@ssw0rd,domain=srv02,port=666
```

### TIP

If you mount a network shared folder to a directory in the Veeam Agent machine file system in advance, you can create the backup repository in the same way as in a local directory. For details, see [Creating Repository in Local Directory](#).

## Creating Repository in Object Storage

To create a repository in an object storage location, you must specify a storage provider name, a name for the backup repository and settings for the object storage account and bucket or container.

## Before You Begin

Before you start creating an object storage repository, consider the following:

- [Microsoft Azure Blob storage] The soft delete feature for blobs and containers must be disabled in the storage account.
- [Microsoft Azure Blob storage] To use the [Veeam backup immutability feature](#), you must enable blob versioning and version-level immutability support in the storage account. For more information, see [this Microsoft Azure documentation](#).
- [S3 Compatible and Amazon S3 storage] To use the [Veeam backup immutability feature](#), you must enable versioning and the S3 Object Lock feature in the storage account. For more information, see [this Amazon S3 documentation](#).

- [Google Cloud storage] The Veeam backup immutability feature is not supported for repositories configured in Google Cloud storage.

## Creating Object Storage Repository

To create an object storage repository, use the following command:

```
veeamconfig objectstorage createrepository <provider_type> <options>
```

where:

- `<provider_type>` – name of the object storage provider. Veeam Agent supports the following options:
  - `azureblob` – for creating a Microsoft Azure Blob repository.
  - `google` – for creating a Google Cloud repository.
  - `amazons3` – for creating an Amazon S3 repository.
  - `s3compatible` – for creating an S3 Compatible repository (including WasabiCloud and IBM Cloud repositories).
- `<options>` – options necessary to connect to the target object storage. For more information, see the following subsections:
  - [Specifying options for S3 Compatible repository](#)
  - [Specifying options for Amazon S3 repository](#)
  - [Specifying Options for Google Cloud repository](#)
  - [Specifying options for Microsoft Azure Blob repository](#)

After Veeam Agent creates a new backup repository in the object storage location, you can specify object storage as a destination for the backup job.

## Specifying Options for S3 Compatible Repository

To create a backup repository in an S3 compatible storage bucket, use the following command:

```
veeamconfig objectstorage createrepository s3compatible --name <repository_name> --servicepoint <address> --region <storage_region> --accesskeyid <id> [--fingerprint <ssl_thumbprint>] --bucketname <bucket_name> --folder <folder_name>
```

where:

- `<repository_name>` – name for the backup repository.
- `<address>` – address of the service point for the object storage.

## NOTE

If you want to connect to the repository using the IPv6 address and the port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is the IPv6 address of the object storage.
  - `port` is the number of the port that Veeam Agent will use to connect to the object storage.
- `<storage_region>` – region associated with the bucket.

## NOTE

You can find the list of supported regions in the documentation of the selected storage provider.

- `<id>` – access key associated with the object storage account.
- `<ssl_thumbprint>` – fingerprint to verify the SSL certificate.
- `<bucket_name>` – name of the bucket.
- `<folder_name>` – name of the folder in the bucket.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent will create a new folder in the bucket under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository s3compatible --name s3comp --servicepoint fd00:ca19:0:18b0:0:ac8a:abca:c942:9000 --accesskeyid S3ertlD9EIO9DjnZjuD4 --region us-east-1 --fingerprint <value> --bucketname backup01 --folder folder01
```

After you use the command, Veeam Agent will prompt you to specify a secret key associated with the object storage account. Enter the secret key to complete the creation of the repository.

## Specifying Options for Amazon S3 Repository

To create a backup repository in an Amazon S3 bucket, use the following command:

```
veeamconfig objectstorage createrepository amazons3 --name <repository_name> --accesskeyid <id> --region <storage_region> --bucketname <bucket_name> --folder <folder_name>
```

where:

- `<repository_name>` – name for the backup repository.
- `<id>` – access key associated with the Amazon S3 storage account.
- `<storage_region>` – region associated with the bucket.

## NOTE

You can find the full list of supported regions by storage provider in the `PublicCloudRegions.xml` file located in the `/Library/Application Support/Veeam` folder on Veeam Agent computer.

- `<bucket_name>` – name of the bucket.

## IMPORTANT

You must create the bucket where you want to store your backup data beforehand. When you create a bucket, consider Amazon bucket naming rules. For example, it is not recommended that you use dots (.) in the bucket name. For more information on bucket naming rules, see [this AWS documentation article](#).

- `<folder_name>` – name of the folder in the bucket.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent will create a new folder in the bucket under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository amazons3 --name amazon --accesskey  
id AMAZONKIAWHDY4BDYCJC --region us-east-1 --bucketname bucket01 --folder folde  
r01
```

After you use the command, Veeam Agent will prompt you to specify a secret key associated with the Amazon S3 storage account. Enter the secret key to complete the creation of the repository.

## Specifying Options for Google Cloud Repository

To create a backup repository in a Google Cloud storage bucket, use the following command:

```
veeamconfig objectstorage createrepository google --name <repository_name> --ac  
cesskeyid <id> --region <storage_region> --bucketname <bucket_name> --folder <f  
older_name>
```

where:

- `<repository_name>` – name for the backup repository.
- `<id>` – access key associated with the Google Cloud storage account.
- `<storage_region>` – region associated with the bucket.

## NOTE

You can find the full list of supported regions by storage provider in the `PublicCloudRegions.xml` file located in the `/Library/Application Support/Veeam` folder on Veeam Agent computer.

- `<bucket_name>` – name of the bucket.

- `<folder_name>` – name of the folder in the bucket.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent will create a new folder in the bucket under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository google --name google --accesskeyid
GOOGLE56L5ATTDKJCLWUQG3E --region europe-west3 --bucketname backup01 --folder f
older01
```

After you use the command, Veeam Agent will prompt you to specify a secret key associated with the Google Cloud storage account. Enter the secret key to complete the creation of the repository.

## Specifying Options for Microsoft Azure Blob Repository

To create a backup repository in a Microsoft Azure Blob container, use the following command:

```
veeamconfig objectstorage createrepository azureblob --name <repository_name> -
-account <storage_account_name> --region <storage_region> --bucketname <bucket_
name> --folder <folder_name>
```

- `<repository_name>` – name of the backup repository for the Veeam Agent database.
- `<account>` – name of the Microsoft Azure Blob storage account.
- `<storage_region>` – region associated with the container.

### NOTE

Veeam Agent supports specification of 3 generic Microsoft Azure Blob storage locations:

- **Azure Global (Standard)** – can be used for any data center region, except the regions in China and the regions intended for US governments. To specify this region in the command to create the repository, use the following value: `AzureCloud`.
- **Asia China** – can be used for any region in China. To specify this region in the command to create the repository, use the following value: `AzureChinaCloud`.
- **Azure Government** – can be used for Azure Government regions only. To specify this region in the command to create the repository, use the following value: `AzureGovernmentCloud`.

You can find the full list of supported regions by storage provider in the `PublicCloudRegions.xml` file located in the `/Library/Application Support/Veeam` folder on Veeam Agent computer.

- `<bucket_name>` – name of the container.
- `<folder_name>` – name of the folder in the container.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent creates a new folder in the container under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository azureblob --name azure --account my-account --region azurecloud --bucketname backup01 --folder folder01
```

After you use the command, Veeam Agent will prompt you to specify the shared key associated with the object storage account. Enter the shared key to complete the creation of the repository.

# Viewing List of Backup Repositories

To view backup repositories configured in Veeam Agent for Linux, use the following command:

```
veeamconfig repository list
```

Veeam Agent will display a list of backup repositories.

You can view the following information about backup repositories:

Parameter	Description
<b>Name</b>	Name of the backup repository.
<b>ID</b>	ID of the backup repository.
<b>Location</b>	Directory in the local file system specified as a target location for backup files.
<b>Type</b>	Type of the backup repository. Possible values: <ul style="list-style-type: none"><li>• Local</li><li>• Backup server</li></ul>
<b>Backup server</b>	Backup server on which Veeam backup repository added to Veeam Agent is configured.

For example:

```
user@srv01:~$ veeamconfig repository list
Name      ID                               Location      Type      Backup server
BackupVol01 {818e3a0f-8155-4a51-9430-248a203a43d1} /home/backups local
BackupVol02 {2155a2e7-a1e9-4347-9d8b-cf8f3a6f3fcb} /home/backups2 local
```



# Editing Backup Repository Settings

In command line interface, you can edit settings for a backup repository created with Veeam Agent for Linux in a local or network shared folder.

You can edit properties of the following repository types only in the backup job settings in the Veeam Agent control panel:

- Veeam backup repository.
- Veeam Cloud Connect repository.
- Object storage repository.

You can edit the following parameters for the backup repository:

- [Name of the backup repository](#)
- [Location of the backup repository](#)

## NOTE

Consider the following:

- If you change location for the backup repository that is already used by a backup job and contains backup files, during the next backup job run, Veeam Agent will create a new backup chain in the new repository location.
- You can temporarily change backup repository location if you want to create an ad hoc full backup in addition to the backup chain created by the backup job in the original repository location.

## Changing Backup Repository Name

To change a name for the backup repository, use the following command:

```
veeamconfig repository edit --name <new_name> for --name <old_name>
```

or

```
veeamconfig repository edit --name <new_name> for --id <id>
```

where:

- <old\_name> – current name of the backup repository.
- <new\_name> – desired name for the backup repository.
- <id> – ID of the backup repository.

For example:

```
user@srv01:~$ veeamconfig repository edit --name LocalRepository for --name Repository_1
```

## Changing Backup Repository Location

To change location for the backup repository, use the following command:

```
veeamconfig repository edit --location <path> for --name <name>
```

or

```
veeamconfig repository edit --location <path> for --id <id>
```

where:

- <path> – desired path for the backup repository.
- <name> – current name of the backup repository.
- <id> – ID of the backup repository.

For example:

```
user@srv01:~$ veeamconfig repository edit --location /home/veeam for --id 3458797-3ffe-45bc-870e-c5628643bbb3
```

## Changing Backup Repository Name and Location

You can change a name and location for the backup repository at the same time, for example:

```
user@srv01:~$ veeamconfig repository edit --name LocalRepository --location /home/veeam for --name Repository_1
```

# Rescanning Veeam Backup Repository

If Veeam Agent for Linux fails to display backups stored in the Veeam Backup & Replication backup repository for some reason, you can rescan the Veeam backup repository. Veeam Agent will try to reconnect to the Veeam backup server and refresh the list of backups in the backup repository.

To rescan a Veeam backup repository, use the following command:

```
veeamconfig repository rescan --id <repository_id>
```

or

```
veeamconfig repository rescan --name <repository_name>
```

where:

- `<repository_id>` – ID of the backup repository that you want to rescan.
- `<repository_name>` – name of the backup repository that you want to rescan.

For example:

```
user@srv01:~$ veeamconfig repository rescan --name [vbr01]BackupVol01
```

You can also rescan all Veeam backup repositories managed by the backup server to which Veeam Agent is connected with the following command:

```
veeamconfig repository rescan --all
```

## NOTE

When you use the `veeamconfig repository rescan` command with the `--all` option, consider the following:

- Rescanning can take significant amount of time if there are multiple repositories configured in Veeam Agent.
- Rescanning multiple object storage repositories may result in greater storage costs due to additional volume of data transactions.

## TIP

You can also use the `veeamconfig repository rescan` command to rescan local backup repositories. This may be useful, for example, after information about a backup stored in the local repository is deleted from the Veeam Agent configuration database, or after you copy a backup to the local repository.

# Deleting Backup Repository

You can delete a backup repository configured with Veeam Agent for Linux. When you delete a backup repository, Veeam Agent removes record of the deleted repository from its database. Backup files created by a backup job targeted at the deleted backup repository remain intact on the backup storage.

To delete a backup repository, use the following command:

```
veeamconfig repository delete --id <repository_id>
```

or

```
veeamconfig repository delete --name <repository_name>
```

where:

- `<repository_id>` – ID of the backup repository that you want to delete.
- `<repository_name>` – name of the backup repository that you want to delete.

For example:

```
user@srv01:~$ veeamconfig repository delete --name Repository_1
```

## NOTE

You cannot delete a backup repository that is specified as a backup storage location in the backup job settings.

# Managing Veeam Backup & Replication Servers

You can store backup files created with Veeam Agent for Linux on backup repositories managed by Veeam Backup & Replication. To do this, you must [connect to a Veeam backup server](#). After that, you can specify a Veeam backup repository as a target location for backup files [in the properties of the backup job](#).

# Connecting to Veeam Backup Server

To create Veeam Agent backups on a backup repository managed by Veeam Backup & Replication, you must connect to a Veeam backup server.

## IMPORTANT

Currently, Veeam Agent for Linux can be connected to one Veeam Backup & Replication server only. If you want to create backups on the backup repository managed by another Veeam backup server, you need to delete currently used backup server and all jobs targeted at backup repositories managed by this backup server. To learn more, see [Deleting Connection to Veeam Backup Server](#).

If you add a connection to another backup server, backup jobs targeted at the original backup server will fail, and backups created on the Veeam backup repository will become unavailable in Veeam Agent. To continue using the original backup server, you need to delete the connection to the new backup server and re-create all backup jobs that use the original backup server.

If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain on the backup repository.

To connect Veeam Agent for Linux to a Veeam backup server, use the following command:

```
veeamconfig vbrserver add --name <vbr_name> --address <vbr_address> --port <vbr_port> --login <username> --domain <domain> --password <password>
```

where:

- <vbr\_name> – name of the Veeam backup server that manages the backup repository.
- <vbr\_address> – DNS name or IP address of the Veeam backup server.
- <vbr\_port> – port over which Veeam Agent must communicate with Veeam Backup & Replication. The default port used for communication with the Veeam backup server is 10006.
- <username> – a name of the account that has access to the Veeam backup repository.
- <domain> – a name of the domain in which the account that has access to the Veeam backup repository is registered.
- <password> – password of the account that has access to the Veeam backup repository.

Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

For example:

```
user@srv01:~$ veeamconfig vbrserver add --name vbr01 --address 172.17.53.1 --port 10006 --login veeam --domain tech --password P@ssw0rd
```

When Veeam Agent for Linux connects to a Veeam Backup & Replication server, Veeam Agent retrieves information about backup repositories managed by this Veeam backup server and displays them in the list of available backup repositories. You can then specify a Veeam backup repository as a target for a backup job.

## TIP

To view the list of backup repositories, use the `veeamconfig repository list` command. To learn more, see [Viewing List of Backup Repositories](#).

# Viewing List of Veeam Backup Servers

To view a list of Veeam backup servers to which Veeam Agent for Linux is connected, use the following command:

```
veeamconfig vbrserver list
```

Veeam Agent will display the list of Veeam backup servers.

For the Veeam backup server in the list, Veeam Agent for Linux displays the following information:

Parameter	Description
Name	Name of the Veeam backup server.
ID	ID of the Veeam backup server in the Veeam Agent database.
Endpoint	IP address of the Veeam backup server and port over which Veeam Agent for Linux communicates with Veeam Backup & Replication.

For example:

```
user@srv01:~$ veeamconfig vbrserver list
Name          ID                                     Endpoint
vbr01         {0fc87c11-6a8d-48c1-8aeb-7f7655738796} 172.17.53.1:10006
```



# Viewing Backup Server Details

You can view detailed information about the Veeam backup server to which Veeam Agent for Linux is connected. Use the following command:

```
veeamconfig vbrserver info --name <vbr_name>
```

or

```
veeamconfig vbrserver info --id <vbr_id>
```

where:

- <vbr\_name> – name of the Veeam backup server.
- <vbr\_id> – ID of the Veeam backup server in the Veeam Agent database.

Veeam Agent for Linux displays the following information about the Veeam backup server:

Parameter	Description
<b>ID</b>	ID of the Veeam backup server in the Veeam Agent database.
<b>Name</b>	Display name of the Veeam backup server.
<b>Endpoint</b>	IP address of the Veeam backup server and port over which Veeam Agent for Linux communicates with Veeam Backup & Replication.
<b>Login</b>	Name of the account that has access to the Veeam backup repository.
<b>Domain</b>	Name of the domain in which the account that has access to the Veeam backup repository is registered.

For example:

```
user@srv01:~$ veeamconfig vbrserver info --name vbr01
VBR server
  ID: {0fc87c11-6a8d-48c1-8aeb-7f7655738796}
  Name: vbr01
  Endpoint: 172.17.53.1:10006
  Login: veeam
  Domain: tech
```

# Editing Connection to Veeam Backup Server

You can edit the following parameters for a connection to a Veeam backup server:

- [Display name of the Veeam backup server](#)
- [IP address and port used to connect to the Veeam backup server](#)
- [Account to connect to the Veeam backup server](#)

## Changing Veeam Backup Server Name

To change a name for the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --name <new_vbr_name>
```

where:

<new\_vbr\_name> – desired name for the backup server.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --name vbr01
```

## Changing IP Address and Port for Veeam Backup Server

To change the IP address and port used to connect to the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --address <vbr_address> --port <vbr_port>
```

where:

- <vbr\_address> – DNS name or IP address of the Veeam backup server.
- <vbr\_port> – port over which Veeam Agent for Linux must communicate with Veeam Backup & Replication.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --address 172.17.53.1 --port 10006
```

# Changing Account to Connect to Veeam Backup Server

## NOTE

If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain on the backup repository.

To change an account whose credentials will be used to connect to the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --login <username> --domain <domain> --password
```

where:

- `<username>` – name of the account that has access to the Veeam backup repository.
- `<domain>` – name of the domain in which the account that has access to the Veeam backup repository is registered.

When you run the command, Veeam Agent will prompt you to enter the password of the specified account.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --login veeam --domain tech --password  
Enter password:
```

## Changing Several Backup Server Parameters

You can change several parameters for the connection to the Veeam backup server simultaneously. For example, the following command changes the name and connection settings for the Veeam backup server:

```
user@srv01:~$ veeamconfig vbrserver edit --name vbr02 --address 172.17.53.2 --p  
ort 10006
```

# Updating List of Veeam Backup Repositories

When you connect to a Veeam backup server, Veeam Agent for Linux retrieves information about backup repositories managed by this Veeam backup server and displays them in the list of available backup repositories. You can refresh information about available Veeam backup repositories manually at any time. This may be useful, for example, after a new backup repository was added on the Veeam backup server.

To update the list of backup repositories managed by the Veeam backup server, use the following command:

```
veeamconfig vbrserver resync
```

## TIP

To view updated list of available Veeam backup repositories after resync, use the `veeamconfig repository list` command. To learn more, see [Viewing List of Backup Repositories](#).

# Deleting Connection to Veeam Backup Server

You can delete a connection to the Veeam backup server to which Veeam Agent is currently connected. When you delete a connection to a Veeam backup server, Veeam Agent removes record on the deleted backup server from its database. Veeam backup repositories managed by the deleted backup server are removed from the list of available backup repositories. Backup files created by backup jobs targeted these repositories remain intact on the backup storage.

You cannot delete a connection to a Veeam backup server in the following situations:

- Veeam Agent operates in the managed mode. To delete connection to a Veeam backup server, reset Veeam Agent to the standalone mode. For details, see [Resetting to Standalone Operation Mode](#).
- Veeam Agent has a backup job that saves backup files to a repository managed by this backup server. To remove such connection to a Veeam backup server, you first need to delete reference to the Veeam backup repository in the job settings.

To delete a connection to the Veeam backup server, use the following command:

```
veeamconfig vbrserver delete --name <vbr_name>
```

or

```
veeamconfig vbrserver delete --id <vbr_id>
```

where:

- <vbr\_name> – name of the Veeam backup server.
- <vbr\_id> – ID of the Veeam backup server.

For example:

```
user@srv01:~$ veeamconfig vbrserver delete --name vbr01
```

# Managing Service Providers

You can store backup files created with Veeam Agent for Linux on a cloud repository exposed to you by a Veeam Cloud Connect service provider. To do this, you must [connect to a service provider](#). After that, you can specify a cloud repository as a target location for backup files [in the properties of the backup job](#).

# Connecting to Service Provider

To create Veeam Agent backups on a cloud repository, you must connect to a Veeam Cloud Connect service provider.

To connect Veeam Agent for Linux to a service provider, use the following command:

```
veeamconfig cloud add --name <sp_name> --address <sp_address> --port <sp_port>
--login <username> --password <password> --fingerprint <sp_thumbprint>
```

where:

- `<sp_name>` – name of the service provider to which you want to connect.
- `<sp_address>` – IP address or full DNS name of the cloud gateway that the SP or your backup administrator has provided to you.
- `<sp_port>` – port over which Veeam Agent must communicate with the cloud gateway. The default port used for communication with the cloud gateway is 6180.
- `<username>` – name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT/SUBTENANT* format.
- `<password>` – password of the tenant or subtenant account used to connect to the service provider.
- `<sp_thumbprint>` – thumbprint used to verify the TLS certificate that the SP has provided to you.

For example:

```
user@srv01:~$ veeamconfig cloud add --name SP --address 172.17.53.15 --port 6180
--login TechCompany/User01 --password P@ssw0rd --fingerprint 92FA988A3D9E80EE
095DDAB75BF06B05DF6F205B
```

## NOTE

When you enter the `veeamconfig cloud add` command, Veeam Agent will display information about the TLS certificate obtained from the SP. To accept the certificate, type `yes` in the command prompt and press [Enter].

When Veeam Agent connects to the service provider, Veeam Agent retrieves information about cloud repositories available to the tenant or subtenant and displays them in the list of available backup repositories. You can then specify a cloud repository as a target for a backup job.

## TIP

To view the list of available cloud repositories, use the `veeamconfig repository list` command. To learn more, see [Viewing List of Backup Repositories](#).

# Viewing List of Service Providers

To view a list of service providers to which Veeam Agent is connected, use the following command:

```
veeamconfig cloud list
```

Veeam Agent will display the list of service providers.

For the service provider in the list, Veeam Agent for Linux displays the following information:

Parameter	Description
<b>Name</b>	Name of the service provider.
<b>ID</b>	ID of the service provider in the Veeam Agent database.
<b>Address</b>	IP address of the cloud gateway and port over which Veeam Agent communicates with the cloud gateway.
<b>Gate servers</b>	IP address of the cloud gateway and port over which Veeam Agent communicates with the cloud gateway.
<b>Username</b>	Name of the tenant or subtenant account used for connection to the service provider.

For example:

```
user@srv01:~$ veeamconfig cloud list
Name      ID                               Address      Gate
servers   Username
SP        {0840f770-354d-426a-b5ce-1aa80f56cc08}  172.17.53.15:618
0          TechCompany
```



# Editing Connection to Service Provider

You can edit the following parameters for a connection to a Veeam Cloud Connect service provider:

- [Name of the Veeam Cloud Connect service provider](#)
- [IP address and port used to connect to the cloud gateway](#)
- [Account to connect to the service provider](#)
- [Thumbprint to connect to the service provider](#)

## Changing SP Name

To change a name for the SP, use the following command:

```
veeamconfig cloud edit --name <new_sp_name> for --name <old_sp_name>
```

or

```
veeamconfig cloud edit --name <new_sp_name> for --id <sp_id>
```

where:

- `<old_sp_name>` – current name of the SP.
- `<new_sp_name>` – desired name for the SP.
- `<sp_id>` – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --name SP for --id 7d3022de-4f4d-4c70-85eb-e8a946a555cd
```

## Changing IP Address and Port for Cloud Gateway

To change the IP address and port of the cloud gateway provided by the SP, use the following command:

```
veeamconfig cloud edit --address <sp_address> --port <sp_port> for --name <sp_name>
```

or

```
veeamconfig cloud edit --address <sp_address> --port <sp_port> for --id <sp_id>
```

where:

- `<sp_address>` – IP address or full DNS name of the cloud gateway that the SP or your backup administrator has provided to you.
- `<sp_port>` – port over which Veeam Agent must communicate with the cloud gateway. The default port used for communication with the cloud gateway is 6180.
- `<sp_name>` – name of the SP.
- `<sp_id>` – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --address 172.17.53.67 --port 6180 for --name SP
```

## Changing Account to Connect to SP

To change an account whose credentials will be used to connect to the SP, use the following command:

```
veeamconfig cloud edit --login <username> --password <password> for --name <sp_name>
```

or

```
veeamconfig cloud edit --login <username> --password <password> for --id <sp_id>
```

where:

- `<username>` – name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT/SUBTENANT* format.
- `<password>` – password of the tenant or subtenant account used to connect to the service provider.
- `<sp_name>` – name of the SP.
- `<sp_id>` – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --login ABC_Compan/User01 --password P@ssw0rd for --name SP
```

# Changing Thumbprint to Connect to SP

To change a thumbprint that will be used to connect to the SP, use the following command:

```
veeamconfig cloud edit --fingerprint <sp_thumbprint> for --name <sp_name>
```

or

```
veeamconfig cloud edit --fingerprint <sp_thumbprint> for --id <sp_id>
```

where:

- <sp\_thumbprint> – thumbprint used to verify the TLS certificate and connect to the service provider.
- <sp\_name> – name of the SP.
- <sp\_id> – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --fingerprint 92FA988A3D9E80EE095DDAB75BF0  
6B05DF6F205B for --name SP
```

# Updating List of Cloud Repositories

When you connect to the Veeam Cloud Connect service provider, Veeam Agent for Linux retrieves and saves to the database information about cloud repositories available to the tenant or subtenant whose account you use to connect to the SP. You can refresh information about available cloud repositories manually at any time. This may be useful, for example, after the SP changes backup resource settings for the tenant.

To update the list of cloud repositories, use the following command:

```
veeamconfig cloud resync
```

If the cloud repository currently used as a target location for Veeam Agent backups becomes unavailable, and Veeam Agent fails to reflect this change in its database for some reason, the `veeamconfig cloud resync` command may finish with errors. In this case, you can use the `--force` option to refresh information about available cloud repositories. For example:

```
veeamconfig cloud resync --force
```

With the `--force` option, Veeam Agent will retrieve the list of available cloud repositories from the service provider and save the new information about cloud repositories in the Veeam Agent database.

## TIP

To view updated list of available cloud repositories after resync, use the `veeamconfig cloud list` command. To learn more, see [Viewing List of Service Providers](#).

# Deleting Connection to Service Provider

You can delete a connection to the service provider to which Veeam Agent for Linux is currently connected. When you delete a connection to a service provider, Veeam Agent removes the record on the deleted service provider from the database. Cloud repositories managed by the deleted service provider are removed from the list of available backup repositories. Backup files created by backup jobs targeted at these repositories remain intact on the cloud repository.

You cannot delete a connection to the service provider if a cloud repository managed by this service provider is used by a backup job. To remove such connection to a service provider, you first need to delete a reference to the cloud repository in the job settings.

To delete a connection to the service provider, use the following command:

```
veeamconfig cloud delete --name <sp_name>
```

or

```
veeamconfig cloud delete --id <sp_id>
```

where:

- <sp\_name> – name of the service provider.
- <sp\_id> – ID of the service provider.

For example:

```
user@srv01:~$ veeamconfig cloud delete --name SP
```

# Managing Backups

You can perform the following operations with backups created by backup jobs configured in Veeam Agent for Linux:

- [View backups](#)
- [View backup details](#)
- [View restore points in backup](#)
- [Export backup to a virtual disk](#)
- [Import backup to the Veeam Agent database](#)
- [Delete backup](#)

# Viewing Backups

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

## NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the <b>Repository</b> column. For information about the import procedure, see <a href="#">Importing Backups</a> .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```



# Viewing Backup Details

You can view detailed information about specific backup. To view backup details, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup\_id> – ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
<b>Machine name</b>	Host name of the machine on which the backup job is configured and the name of the job.
<b>Name</b>	Name of the volume in the backup.
<b>Device</b>	Path to the block device file that represents the volume.
<b>FS UUID</b>	File system ID.
<b>Offset</b>	Position of the volume on the computer disk.
<b>Size</b>	Size of the volume in the backup.

For example:

```
user@srv01:~$ veeamconfig backup show --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Machine name: srv01 SystemBackup
  Name:      [sda1]
  Device:    [/dev/sda1]
  FS UUID:   [6945f2eb-e8bb-48fe-a276-5ba67b9030a5]
  Offset:    [1048576] bytes (2048 sectors)
  Size:      [9999220736] bytes (19529728 sectors)
```

For a file-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
<b>Machine name</b>	Host name of the machine on which the backup job is configured and the name of the job.
<b>Backed up</b>	Backup scope for the file-level backup job.

For example:

```
user@srv01:~$ veeamconfig backup show --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
Machine name: srv01 DocsBackup
File-level backup
Backed up:
/home/user/Documents
```

# Viewing Restore Points in Backup

To view information about restore points in the backup, you can use one of the following commands:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where:

<backup\_id> – ID of the backup for which you want to view information on restore points.

For example:

```
user@srv01:~$ veeamconfig backup info --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
```

or

```
user@srv01:~$ veeamconfig point list --backupid 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
```

Veeam Agent for Linux displays the following information about restore points in the backup:

Parameter	Description
<b>Job name</b>	Name of the backup job by which the backup was created.
<b>OIB ID</b>	ID of the restore point in the backup.
<b>Type</b>	Type of the restore point. Possible values: <ul style="list-style-type: none"><li>• Full</li><li>• Increment</li></ul>
<b>Created at</b>	Date and time of the restore point creation.
<b>Is corrupt</b>	Indicates whether restore point in the backup is corrupted. Possible values: <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul>
<b>Retention</b>	Displays information about enabled long-term retention per each type: weekly (W), monthly (M) and yearly (Y).

# Importing Backups

You can import a backup created by Veeam Agent into the Veeam Agent database. For example, you may want to import a previously deleted backup or backup that was created in a network shared folder by Veeam Agent installed on another computer.

To import a backup:

1. Start the import process with the following command:

```
veeamconfig backup import --path <path>
```

where:

<path> – path to the VBM file of the backup that you want to import.

For example:

```
user@srv01:~$ veeamconfig backup import --path /home/share/BackupJob/BackupJob.vbm
Backup has been imported successfully.
Session ID: [{4031f058-766c-4f2c-a7ae-7257adb2929f}].
Logs stored in: [/var/log/veeam/Import/Session_{4031f058-766c-4f2c-a7ae-7257adb2929f}].
```

2. You can monitor the import process and result by viewing the import session log with the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the import session.

For example:

```
user@srv01:~$ veeamconfig session log --id 4031f058-766c-4f2c-a7ae-7257adb2929f
2023-11-19 13:21:33 UTC {765af178-a9cc-4596-8bf2-03850c5dalac} [info] Job started at 2023-11-19 16:21:33
2023-11-19 13:21:33 UTC {6ae2922d-454b-4a8d-a11b-2b5c7a85029d} [info] Importing backup
2023-11-19 13:21:33 UTC {783f40a7-ead7-4555-9c35-545d875990ee} [info] Backup has been imported.
```

3. Imported backup will be displayed in the list of backups. To view the list of backups, use the following command:

```
veeamconfig backup list
```

For example:

```
user@srv01:~$ veeamconfig backup list
Job name          Backup ID          Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocsBackup   {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomeBackup   {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
BackupJob          {64957b1d-d219-456c-a9cd-9598292c10cd} Importe
d    2023-11-19 19:12
```

## Importing Encrypted Backups

You can import an encrypted backup created by Veeam Agent into the Veeam Agent database. This operation is required if you want to use the Veeam Agent command line interface to restore data from an encrypted backup created by Veeam Agent running on another computer.

To import an encrypted backup:

1. Start the import process with the following command:

```
veeamconfig backup import --path <path>
```

where:

<path> – path to the VBM file of the backup that you want to import.

For example:

```
user@srv01:~$ veeamconfig backup import --path /home/share/srv15\ Backup/B
ackup.vbm
```

2. Veeam Agent will prompt you to provide a password for the backup file. Type in the password and press [Enter] **key** to import the backup.

Veeam Agent displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.

If you enter the correct password, Veeam Agent will decrypt the backup file and import it into the database.

```
user@srv01:~$ veeamconfig backup import --path /home/share/srv15\ Backup/B
ackup.vbm
[Info] Backup srv15 Backup encrypted
[Info] Press "Enter" to skip. Enter password to decrypt the backup:
[Info] Hint: Standard password
Password:
Backup imported successfully
```

3. Imported backup will be displayed in the list of backups. To view the list of backups, use the following command:

```
veeamconfig backup list
```

For example:

```
user@srv01:~$ veeamconfig backup list
Job name          Backup ID          Repository
y    Created at
srv15 Backup      {4b1f873c-857d-b984-4f22-6ce66bf62570} Importe
d      2018-06-12 20:20
srv01 ServerBackup {f212f641-54aa-40de-a0eb-8727be56760b} Importe
d      2018-06-12 20:04
```

# Deleting Backups

Backup files created with Veeam Agent are removed automatically according to the retention policy settings. You can also remove backups from the target location and Veeam Agent configuration database manually if necessary.

## Removing Backup from Configuration

To remove a backup from the Veeam Agent configuration database, use the following command:

```
veeamconfig backup delete --id <backup_id>
```

where `<backup_id>` is an ID of the backup that you want to delete.

The way Veeam Agent removes a backup from configuration depends on the backup location:

- If the backup resides in a local directory or network shared folder, Veeam Agent removes records about the deleted backup from the Veeam Agent database. Backup files themselves (VBK, VIB, VBM) remain in the backup repository.

You can import information about the removed backup later to Veeam Agent and perform restore operations with the imported backup. To import information about the removed backup, use the `veeamconfig repository rescan --all` command.

- If the backup resides in a Veeam Backup & Replication repository, Veeam Agent removes records about the deleted backup from the Veeam Agent database and Veeam Backup & Replication database. Backup files themselves (VBK, VIB, VBM) remain in the backup repository.

If you want to import information about the removed backup later to Veeam Agent and perform restore operations with this backup, you must contact backup administrator working with Veeam Backup & Replication. The administrator must rescan the backup repository that contained the backup in the Veeam Backup & Replication console. For details, see the [Rescanning Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

After rescan, the backup will be displayed in the list of backups on the Veeam Agent machine connected to the Veeam backup server.

## Deleting Backup Files

To delete backup files from the target location and Veeam Agent database, use the following command:

```
veeamconfig backup delete --id <backup_id> --purge
```

where `<backup_id>` is an ID of the backup that you want to delete.

Veeam Agent for Linux will remove records about the deleted backup from the Veeam Agent database and, additionally, delete backup files themselves from the destination storage.

# Performing Restore

If you experience a problem with your computer, your data gets lost or corrupted, you can use one of the following options to recover your data or bring the computer back to work:

- [Restore from the Veeam Recovery Media](#)
  - [Restore volumes](#)
  - [Restore files and folders](#)
- [Restore volumes with the command line interface](#)
- Restore files and folders:
  - [Restore files and folders with the File Level Restore wizard](#)
  - [Restore files and folders with the command line interface](#)
- [Export data as VHD disks](#)
- [Restore data from encrypted backups](#)



# Restoring from Veeam Recovery Media

If the OS on your computer fails to start, you can use the Veeam Recovery Media to recover your computer. The Veeam Recovery Media will help you boot the computer in the limited mode. After booting, you can use a backup created with Veeam Agent for Linux to restore the whole system image of your computer, specific volumes on your computer or specific files and folders. You can also use standard Linux command line utilities to diagnose problems and fix errors.

## IMPORTANT

If you plan to use the custom Veeam Recovery Media, Veeam Agent requires 3 GB RAM or more installed on the target computer or virtual machine. Memory consumption varies depending on size and number of modules included into the recovery media. To learn more, see [Creating Custom Veeam Recovery Media](#).

# Restoring Volumes

You can restore a specific computer volume or all volumes from the volume-level backup.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

## Before You Begin

Before you boot from the recovery image and restore your data, check the following prerequisites and limitations:

- You must have a recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For volume-level restore, you can use a volume-level backup created with Veeam Agent for Linux. Make sure that the backup or system image is available on the computer drive (local or external), on a network shared folder or on the backup repository managed by a Veeam backup server.
- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a Veeam backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- You cannot restore a volume to the volume where the backup file that you use for restore is located.
- If you restore to a virtual environment, note that the current version of Veeam Recovery Media supports only the VMware and Hyper-V virtualization solutions. To resolve possible issues during bare metal recovery of Oracle VM virtual machines, use instructions in the second section of [this Veeam KB article](#).

# Step 1. Boot from Veeam Recovery Media

To boot from the Veeam Recovery Media:


1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.  
  
[For removable storage device] Attach the removable storage device with the recovery image to your computer.
2. Start your computer.
3. [For regular recovery image] In the boot menu, select what Linux kernel version to use to boot your computer and specify boot options if necessary.

You can select a Linux kernel version if you use generic Veeam Recovery Media downloaded from [the Veeam website](#) or [Veeam software repository](#). If you created a custom Veeam Recovery Media, you will be prompted to boot using the Linux kernel of your Veeam Agent computer included in the recovery image.

To specify boot options, press the [Tab] key and type the necessary options in the command prompt.

## NOTE

For the [legacy recovery image](#), the boot menu is unavailable. After you start your computer, Veeam Agent will immediately start loading files from the Veeam Recovery Media.



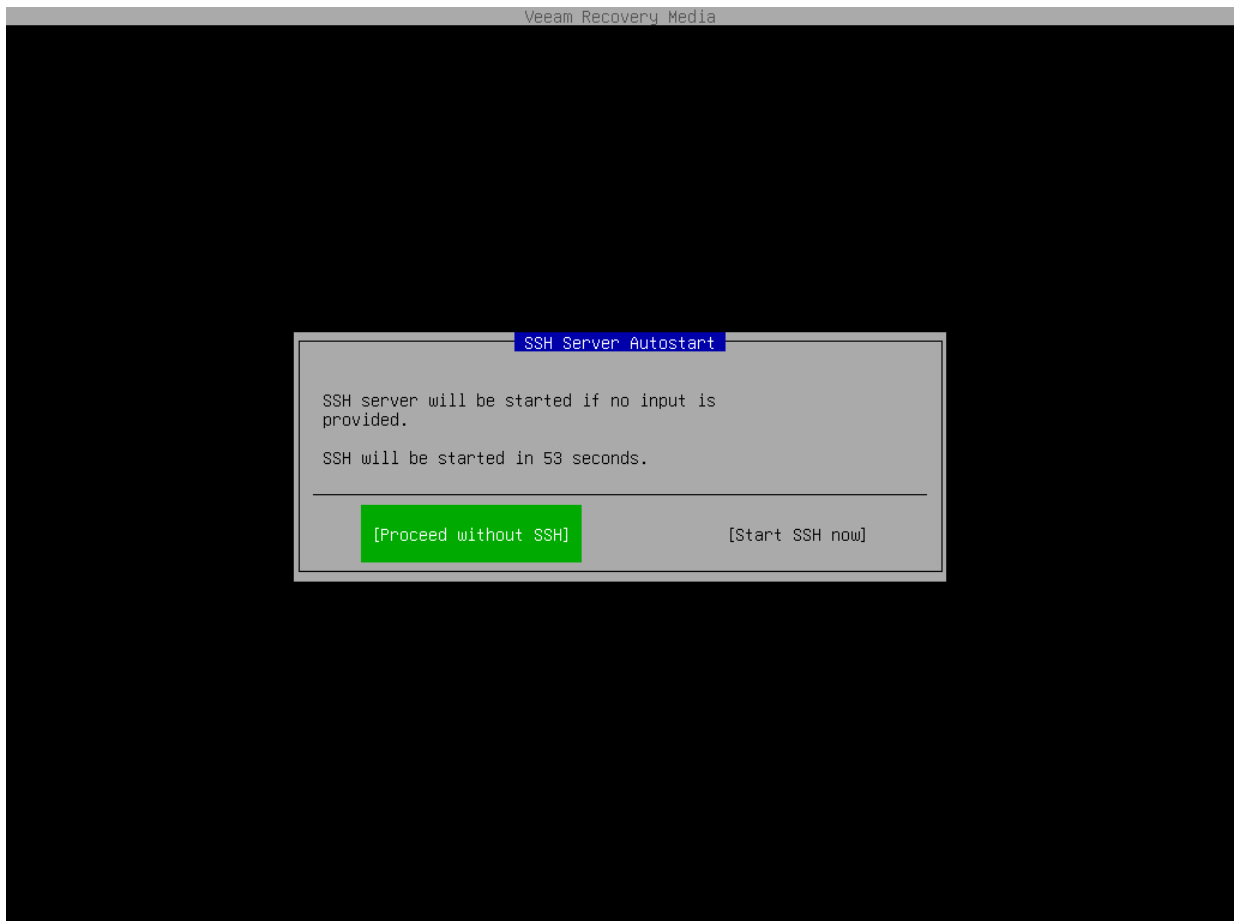
4. Wait for Veeam Agent to load files from the Veeam Recovery Media.

5. After the recovery image OS has loaded, choose whether you want to start the SSH server. The SSH server allows you to connect to the Veeam Recovery Media from a remote machine.

The Veeam Recovery Media starts the SSH server automatically after a time-out. The default value for the time-out is 60 seconds.

If you do not want to start the SSH server, make sure that the **Proceed without SSH** button is selected and press [Enter]. You will proceed immediately to the step 7.

To override the default time-out and start the SSH server immediately, select the **Start SSH now** button using the [Tab] key and press [Enter].



6. After the SSH server has started, review settings to connect to the Veeam Recovery Media and press [Enter].

The Veeam Recovery Media displays the following connection settings:

- IP address of the computer booted from the Veeam Recovery Media
- User name and password of the account used to connect to the Veeam Recovery Media
- Fingerprints of the computer booted from the Veeam Recovery Media

## NOTE

The user name of the account used to work with the Veeam Recovery Media is *veeamuser*.

If you want to use command-line utilities built in the regular recovery image, use the `sudo` command to provide the *veeamuser* account with privileges of the *root* account.

```
SSH Connection Info

Credentials
login: veeamuser
passwd: kaAnL


NetConfigs
ens160
IP: 172.24.28.72
IPv6: fd00:ac18:0:1810:0:b5f9:ab46:e4c5

Fingerprints
ecdsa-sha2-nistp256
SHA1:BaXFVwjaWKUf6Rvv2gAwR+g+knI
MD5:6d:9c:56:1d:62:d3:f6:56:f0:0e:62:25:31:da:3c:a2
ssh-ed25519
SHA1:618oSzFazLsSUaMDD/EQJCymqjc
MD5:2b:2b:5d:78:14:66:55:da:cc:7e:6a:bb:29:a3:01:da
ssh-rsa
SHA1:6PsfT1Vv+Gkn8dgdR7420HasGBQ
MD5:ff:96:15:0b:e4:30:86:67:08:8e:7b:21:47:0c:b4:0a
```

[Continue]

7. Accept the terms of the product license agreement and license agreements for third-party components of the product:
  - a. Make sure that the **I accept Veeam End User Software License Agreement** option is selected and press [Space].
  - b. Select the **I accept the terms of the following 3rd party software components license agreements** option with the [Tab] key and press [Space].


c. Switch to the **Continue** button with the [Tab] key and press [Enter].



8. Make sure that network settings are specified correctly and configure the network adapter if necessary. To learn more, see [Configure Network Settings](#).
9. Choose the necessary recovery option. Veeam Agent offers the following tools:
  - **Restore volumes** – the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
  - **Restore files** – the File Level Restore wizard to restore files and folders to the original location or to a new location.
  - **Exit to shell** – Linux shell prompt with standard utilities to diagnose problems and fix errors.

## TIP


To stop working with the Veeam Recovery Media and shut down or restart your computer, in the Veeam Recovery Media main menu, select the **Reboot** or **Shutdown** option and press [Enter].



## Step 2. Configure Network Settings


If there is a DHCP server in your network, Veeam Agent will configure the network settings automatically. To verify or configure network settings manually, use **nmtui**, a text-based user interface network manager tool provided with Veeam Recovery Media. To learn more about working with nmtui, see [Linux documentation](#).

1. In the Veeam Recovery Media main menu, select **Configure network** and press [Enter].





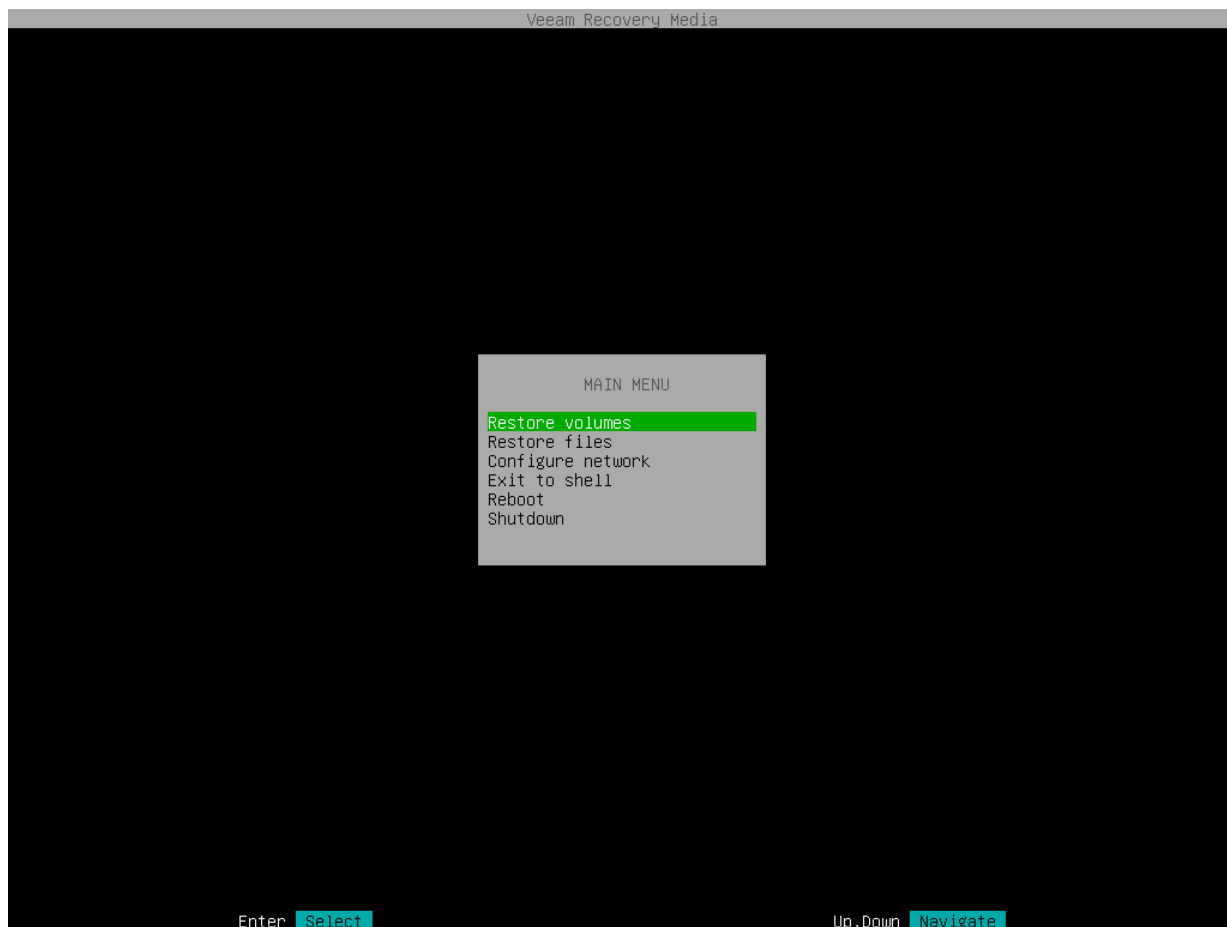
2. To add new or modify existing connection, in NetworkManager, select **Edit a connection**.



3. After you add or edit a connection, in the main menu of the NetworkManager, select **Activate a connection**.
  - a. If the connection is new, choose it in the list of connections; then select **Activate**.
  - b. If the connection was modified, you must reactivate it. To do this, choose it in the list of connections and select **Deactivate**; then choose the connection again and select **Activate**.
4. After you finish working with Network Manager, press [Esc] to return to the Veeam Recovery Media main menu and launch the Volume Restore wizard.

## Step 3. Launch Volume Restore Wizard

To launch the volume restore wizard, in the Veeam Recovery Media main menu, select **Restore volumes** and press [Enter].




## Step 4. Select Backup Location

At the **Select backup location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

To recover data from backup, you need to mount the backup storage on which the backup file resides to the recovery image OS file system. Veeam Agent for Linux automatically mounts external USB drives that are connected to the computer and displays them in the list of available backup locations. You can select the necessary device and press [Enter] to pass to the [Browse for backup files](#) step of the wizard.

If the backup file is located in a network shared folder, on a local drive or on a Veeam backup repository, select one of the following options:

- **Mount local disk** – select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. With this option selected, you will pass to the [Select local disk](#) step of the wizard.
- **Add object storage repository** – select this option if the backup file resides in an object storage repository. With this option selected, you will pass to the [Select cloud storage type](#) step of the wizard.
- **Add shared folder** – select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the [Mount shared folder](#) step of the wizard.
- **Add VBR server** – select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Specify backup server parameters](#) step of the wizard.
- **Add Cloud Connect provider** – select this option if the backup file resides in a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the [Specify Cloud provider parameters](#) step of the wizard.



## Step 5. Specify Backup Location Settings

Specify settings for the target storage that contains a backup file from which you plan to restore data:

- [Specify shared folder settings](#) — if you have selected the **Add shared folder** option at the [Select backup location](#) step of the wizard.
- [Select local drive](#) — if you have selected the **Mount local disk** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam backup repository settings](#) — if you have selected the **Add VBR server** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam Cloud Connect repository settings](#) — if you have selected the **Add Cloud provider** option at the [Select backup location](#) step of the wizard.
- [Specify object storage repository settings](#) - if you have selected the **Add object storage repository** option at the [Select backup location](#) step of the wizard.

### Shared Folder Settings

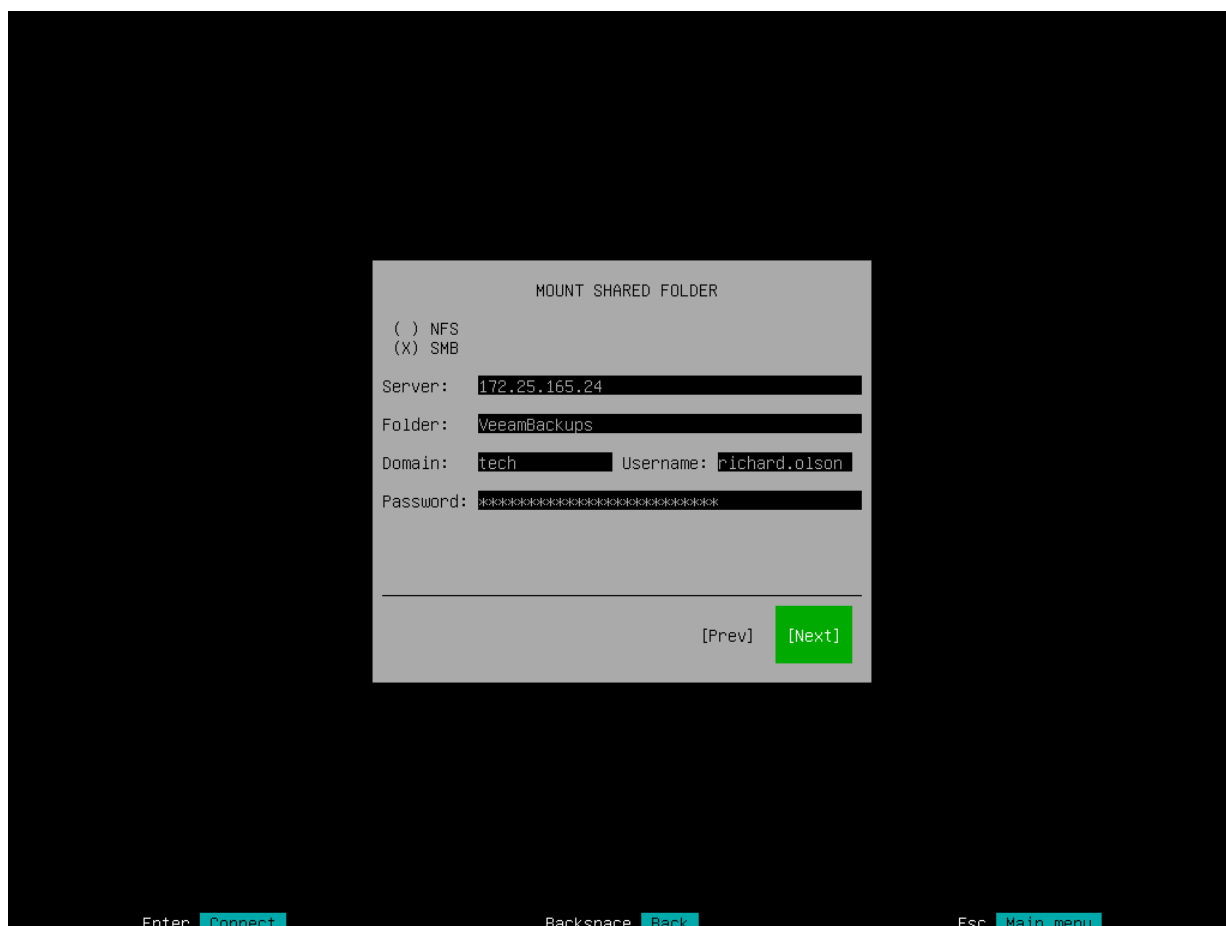
The **Mount shared folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. Select the type of a network shared folder:
  - **NFS** — to connect to a network shared folder using the NFS protocol.
  - **SMB** — to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Path** field, specify the network shared folder name in the *SERVER/DIRECTORY* format: type an IP address or domain name of the server and the name of the network shared folder in which the backup file resides.
3. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
4. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
5. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.
6. Press [Enter] to connect to the network shared folder. Veeam Agent will mount the specified network shared folder to the `/media` directory of the recovery image OS file system and display content of the network shared folder.

## TIP

You can mount several network shared folders to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Add shared folder** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted network shared folders and browse for a backup file located on the necessary storage.




## Local Backup Repository Settings

The **Select local disk** step of the wizard is available if you have selected to restore data from a backup file located on a computer drive.

In the list of devices, select the necessary disk or disk partition and press [Enter]. Veeam Agent will mount the selected device to the `/media` directory of the recovery image OS file system and display content of the directory.

## TIP

You can mount several devices to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Mount local disk** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted devices and browse for a backup file located on the necessary storage.



## Veeam Backup Repository Settings

The **Specify Backup Server parameters** step of the wizard is available if you have selected to restore data from a backup repository managed by the Veeam backup server.

Specify settings for the Veeam backup server that manages the backup repository where the backup file resides:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. Select the type of **Authentication** to access the Veeam backup server:
  - **Login and password.** With this option selected, specify the following settings:
    - i. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
    - ii. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.

- iii. In the **Password** field, type a password of the account that has access to the Veeam backup repository.

#### NOTE

If you want to perform restore from a backup created by Veeam Agent operating in the managed mode, you must use an account that has the Veeam Backup Administrator or Veeam Restore Operator role on the Veeam backup server. For more information about user roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

The screenshot shows the Veeam Recovery Media interface. At the top, it says 'Veeam Recovery Media'. The main area is a dark gray rectangle. In the center, there is a light gray dialog box titled 'Specify Backup Server parameters:'. Inside the dialog, the following fields are visible: 'Address: 172.24.31.136', 'Port: 10006', 'Authentication: (X) Login and password ( ) Recovery token', 'Login: Administrator', 'Domain: ', and 'Password: \*\*\*\*\*'. At the bottom right of the dialog are two buttons: '[Prev]' and '[Next]'. Below the dialog, at the bottom of the screen, there are three keyboard shortcuts: 'Enter Connect', 'Backspace Back', and 'Esc Main menu'.

- **Recovery token:** With this option selected, in the **Token** field, enter the value of the recovery token generated in the Veeam Backup & Replication console. For more information on generating recovery tokens, see [Creating Recovery Token](#) in the Veeam Agent Management guide.

Veeam Recovery Media

Specify Backup Server parameters:

Address: 172.24.31.136

Port: 10006

Authentication:

( ) Login and password

(X) Recovery token

Token: 5e6d-45aA-DBBd-d0ec

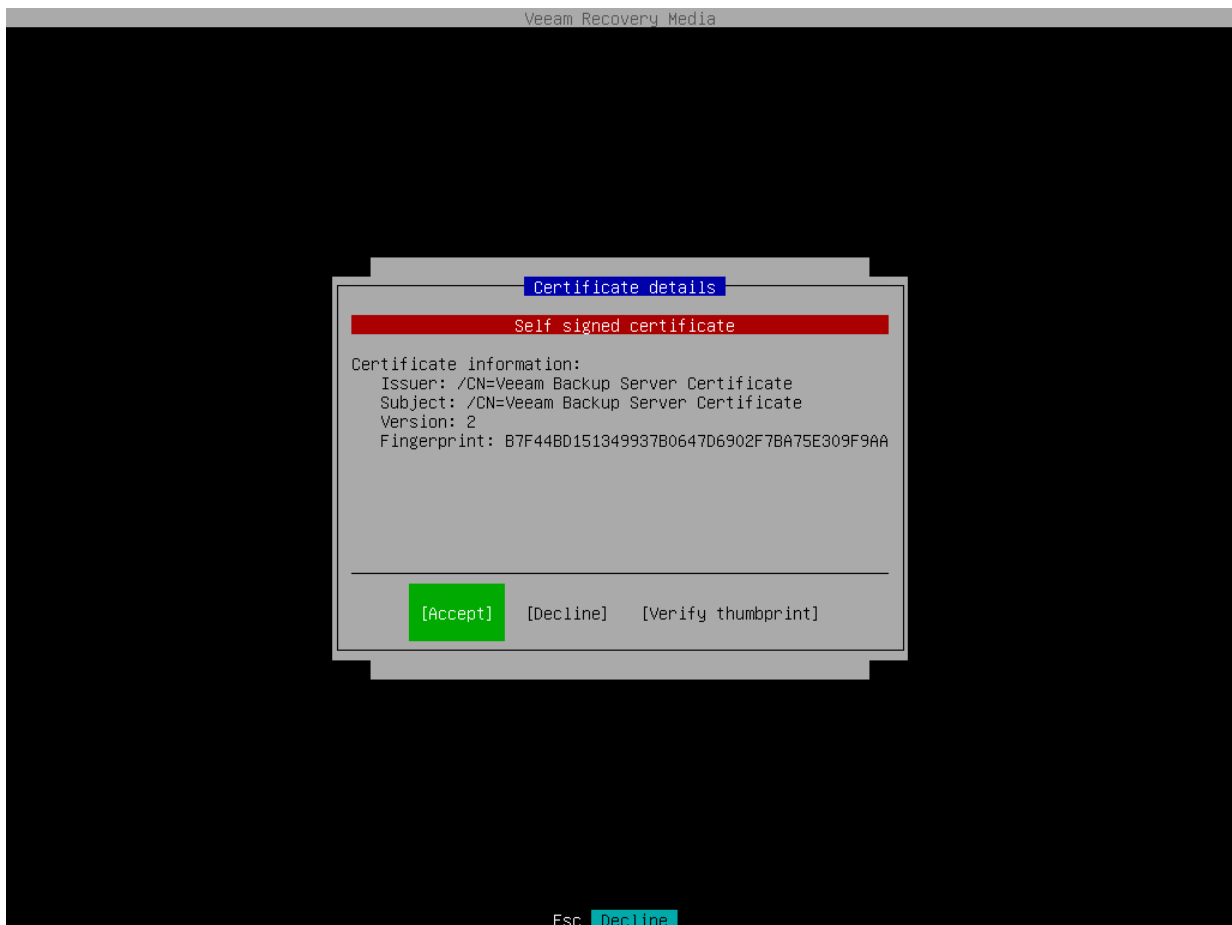
Recovery tokens can be created using the Veeam Backup & Replication console.

[Prev] [Next]

Enter .Connect Backspace .Back Esc .Main menu



4. Press [Enter]. Veeam Agent will connect to the Veeam backup server. If prompted, accept the self-signed certificate of the Veeam backup server to continue.



After successful connection to the Veeam backup server, you will pass immediately to the [Backup](#) step of the wizard.

## Veeam Cloud Connect Repository Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings.](#)
2. [Verify the TLS certificate.](#)
3. [Specify user account settings.](#)


## Specifying Service Provider Settings

The **Specify Cloud provider parameters** step of the wizard is available if you have selected to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.

3. Press [Enter]. Veeam Agent will connect to the service provider and display the [Certificate details](#) window.



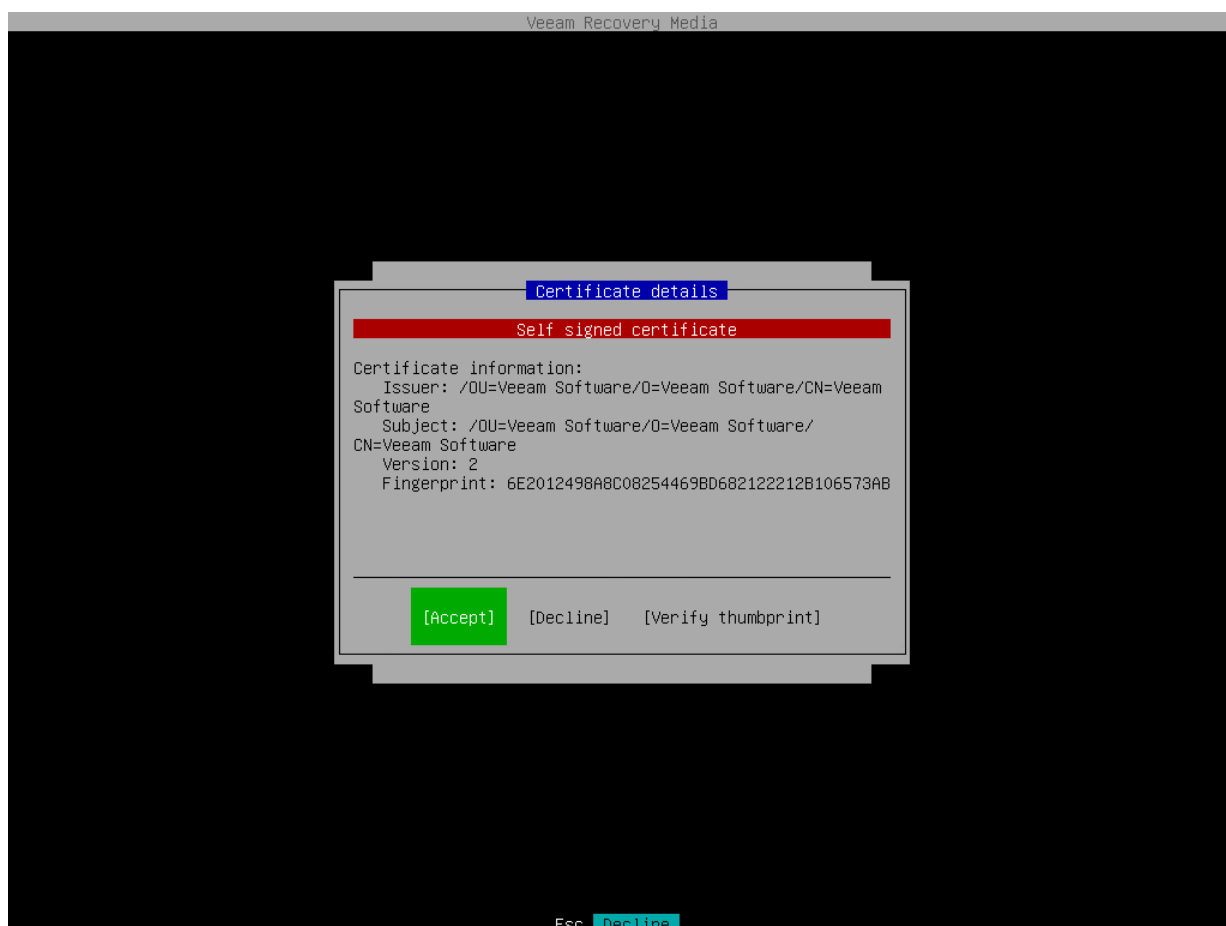
## Verifying TLS Certificate

In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate.

- To accept the TLS certificate, select the **Accept** button with the [Tab] key and press [Enter].
- [Optional] To verify the TLS certificate with a thumbprint, do the following:
  - a. Select the **Verify thumbprint** button with the [Tab] key and press [Enter].
  - b. In the **Thumbprint verification** field, enter the thumbprint that you obtained from the SP.

- c. Switch to the **Verify** button and press [Enter]. Veeam Agent will check if the thumbprint that you entered matches the thumbprint of the obtained TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.




## Specifying User Account Settings

The **Specify Cloud provider credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

1. In the **Username** field, enter the name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.
2. In the **Password** field, provide a password for the tenant or subtenant account.

3. Press [Enter]. Veeam Agent will connect to the cloud repository, and you will pass immediately to the [Backup](#) step of the wizard.



## Object Storage Repository Settings

If you have selected to restore data from a backup file located in a object storage repository, specify settings to connect to the repository:

At the **Select cloud storage type** step of the wizard, select one of the following options:

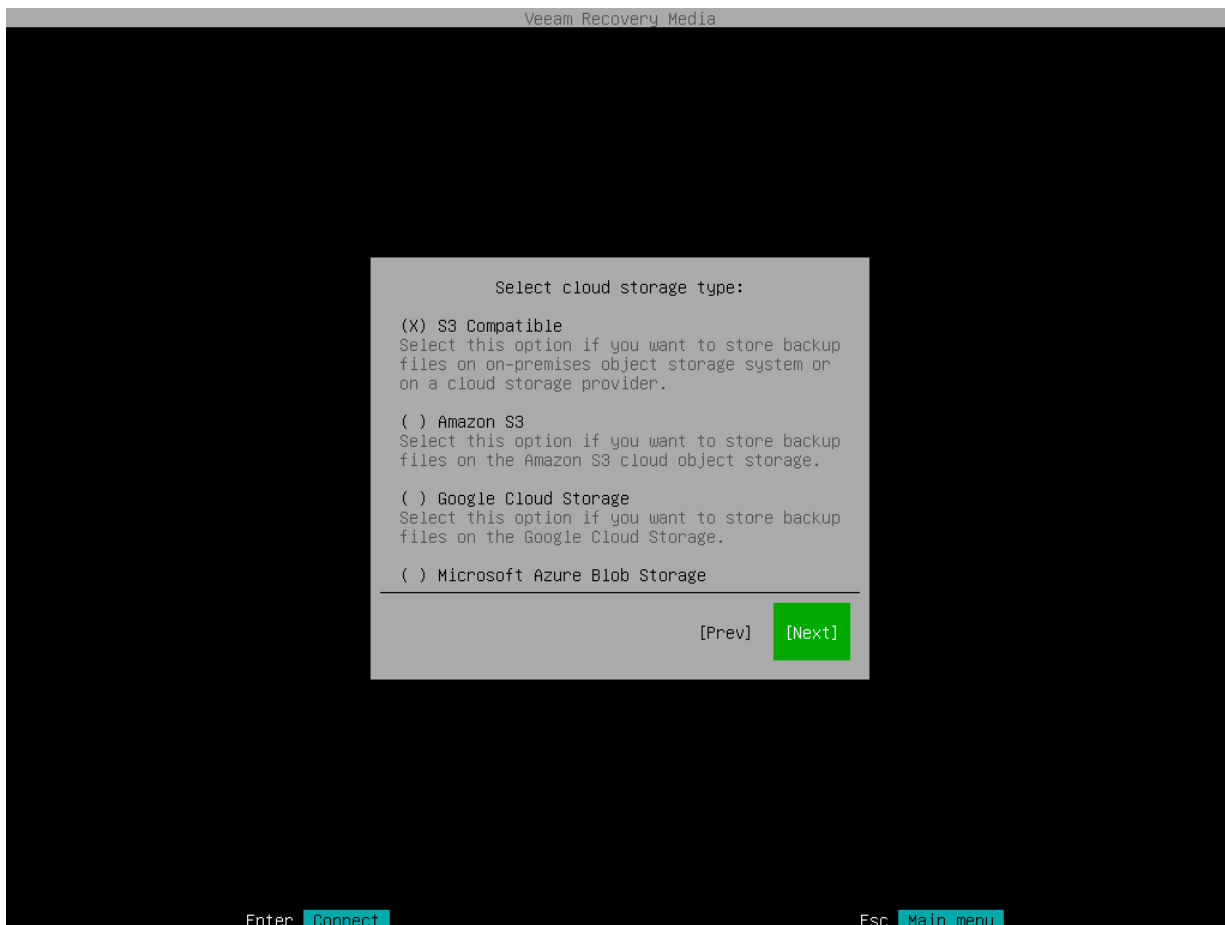
- **S3 Compatible** – select this option if you want to import a backup from an S3 compatible storage repository.

### TIP

If you plan to restore from backups in an IBM or Wasabi object storage, use the **S3 Compatible** storage option.

- **Amazon S3** – select this option if you want to import a backup from an Amazon S3 storage repository.
- **Google Cloud Storage** – select this option if you want to import a backup from a Google Cloud storage repository.

- **Microsoft Azure Blob Storage** – select this option if you want to import a backup from a Microsoft Azure storage repository.



## Specifying Settings for S3 Compatible Repository

If you have selected to import backup from an S3 Compatible storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.

### NOTE

If you want to connect to the repository using the IPv6 address and the port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is an IPv6 address of the cloud storage.
- `port` is a number of a port that Veeam Agent will use to connect to the cloud storage.

2. In the **Region** field, specify a storage region based on your regulatory and compliance requirements.
3. In the **Access key** field, enter an access key ID.
4. In the **Secret key** field, enter a secret access key.

The screenshot shows a terminal window titled 'Veeam Recovery Media'. In the center is a dialog box titled 'Specify S3 compatible storage:'. The dialog box contains the following fields and values:

- Service point: `https://myservicepoint.com:9000`
- Region: `reg-1`
- S3 compatible account:
- Access key: `Access_Key`
- Secret key: `*****`

At the bottom right of the dialog box are two buttons: '[Prev]' and '[Next]'. The '[Next]' button is highlighted in green. At the bottom of the terminal window, there are two keyboard shortcuts: 'Enter Connect' and 'Esc Main menu'.


## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Bucket** field, specify a bucket on the storage:
  - a. Click **Browse**.
  - b. In the **Buckets** window, select the necessary bucket and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



## Specifying Settings for Amazon S3 Repository

If you have selected to store backup files on an Amazon S3 storage, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)


## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter an access key ID.
2. In the **Secret key** field, enter a secret access key.

3. In the **AWS region** window, select an AWS region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Global** region.



## Specifying Bucket Settings


The **Bucket** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
  - a. Click **Browse**.
  - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.



b. In the **Folders** window, select the necessary folder and click **OK**.



## Specifying Settings for Google Cloud Repository


If you have selected to import backup from a Google Cloud storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository.

To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Agent will use the HMAC key to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see the [Google Cloud documentation](#).




## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
  - a. Click **Browse**.
  - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



## Specifying Settings for Microsoft Azure Repository

If you have selected to import backup from a Microsoft Azure storage repository, specify settings to connect to the storage:

1. [Specify account settings](#).
2. [Specify container settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository.


### NOTE

The **Allow storage account key access** option for Shared Key authorization must be enabled in the storage account. For more information on how to find this option, see [Microsoft Docs](#).

To connect to the Microsoft Azure storage, specify the following:

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.

3. In the **Region** window, select a Microsoft Azure region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Azure Global (Standard)** region.




## Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository and specified account settings to connect to the storage.

Specify settings for the container on the storage:

1. In the **Container** field, specify a container on the storage:
  - a. Click **Browse**.
  - b. In the **Containers** window, select the necessary container and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.


b. In the **Folders** window, select the necessary folder and click **OK**.



## Step 6. Browse for Backup File

At the **Browse for backup files** step of the wizard, select the backup file that you plan to use for volume-level restore:

1. In the file system tree, select a directory in which the backup file you plan to use for restore resides:
  - Use [Up] and [Down] arrow keys to select a directory.
  - Use the [Enter] key to open the necessary directory.
2. In the directory where the backup file resides, select the backup file and press [Enter].



## Step 7. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.


The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays information about backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported backups** pane, ensure that the backup from which you want to recover data is selected and press [Enter].

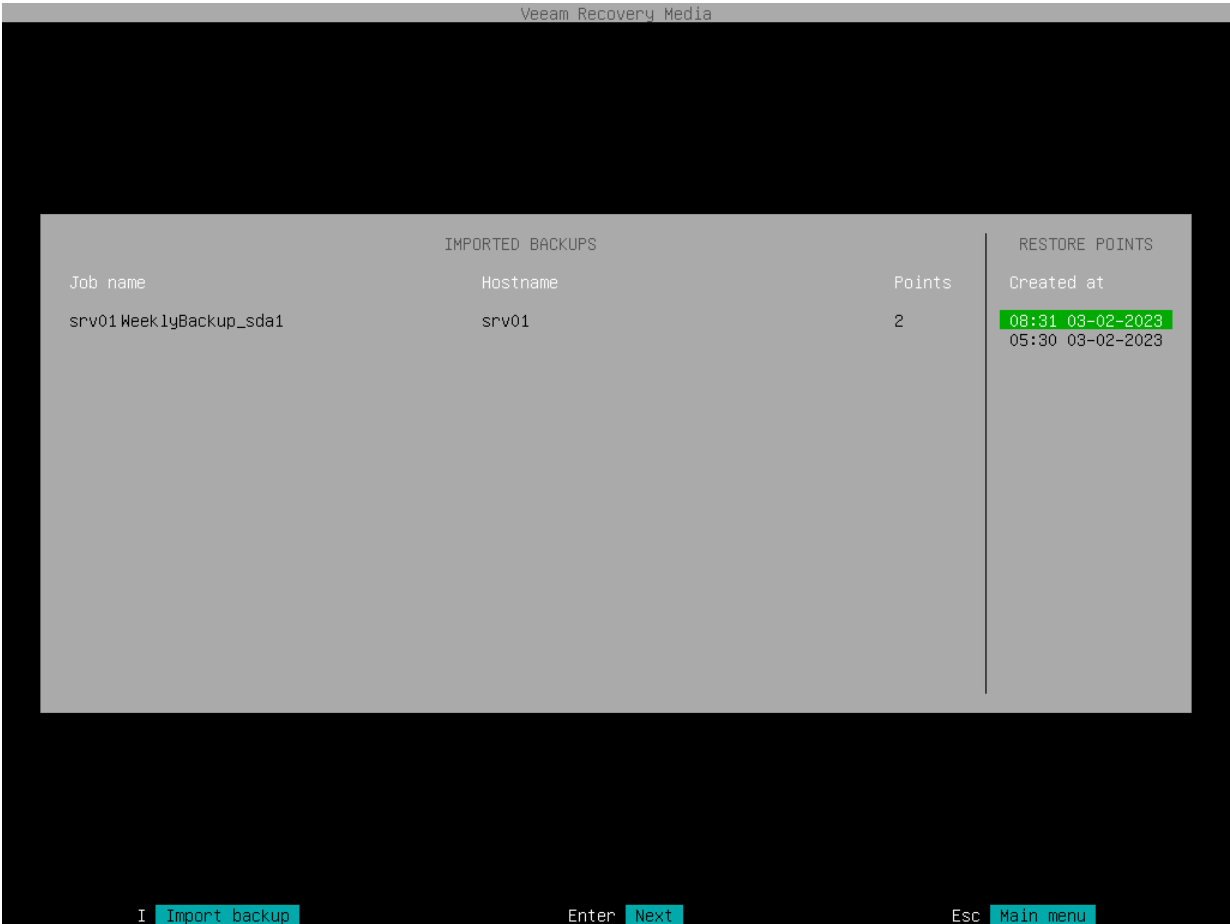
If you want to select another backup, press the [i] key and browse for the necessary backup file. To learn more, see [Locate Backup File](#).



2. In the **Restore points** pane, select with [Up] and [Down] keys the restore point from which you want to recover data and press [Enter].

NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).





## Step 8. Map Restored Disks

At the **Disk Mapping** step of the wizard, select what volumes you want to restore and map volumes from the backup to volumes on your computer.

### IMPORTANT

It is strongly recommended that you change disk mapping settings only if you have experience in working with Linux disks and partitions. If you make a mistake, your computer data may get corrupted.

You can map volumes in the backup (source volumes) and volumes on your computer (target volumes) in one of the following ways:

- [Map a source volume to a target volume](#)
- [Map a target volume to a source volume](#)

As well as individual volumes, you can also map entire disks:

- [Map a source disk to a target disk](#)
- [Map a target disk to a source disk](#)

If you choose to restore an entire disk, Veeam Agent will try to map all volumes that reside on this disk.

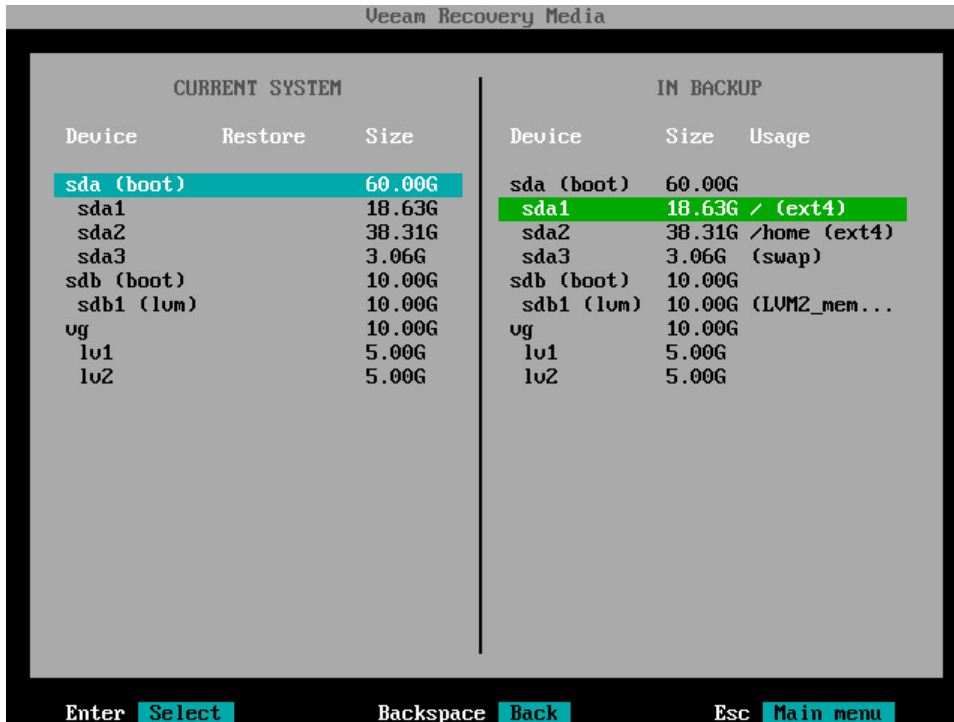
If you want to restore BTRFS subvolumes, you must map subvolumes in the backup to a BTRFS pool on the Veeam Agent computer. To learn more, see [Mapping BTRFS Subvolumes](#).

## Mapping Source Volume to Target Volume

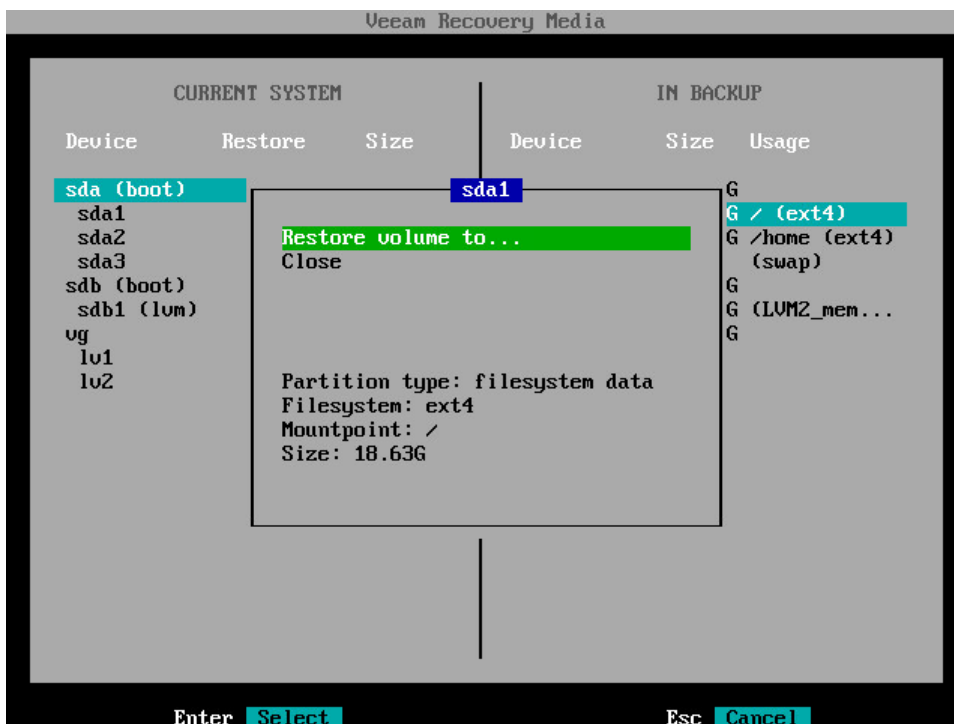
The **In backup** pane of the **Veeam Recovery Media** wizard contains a list of disks and volumes in the backup. You can select volumes in the backup that you want to restore to your computer and specify mapping rules for these volumes.

To map a source volume to a target volume:

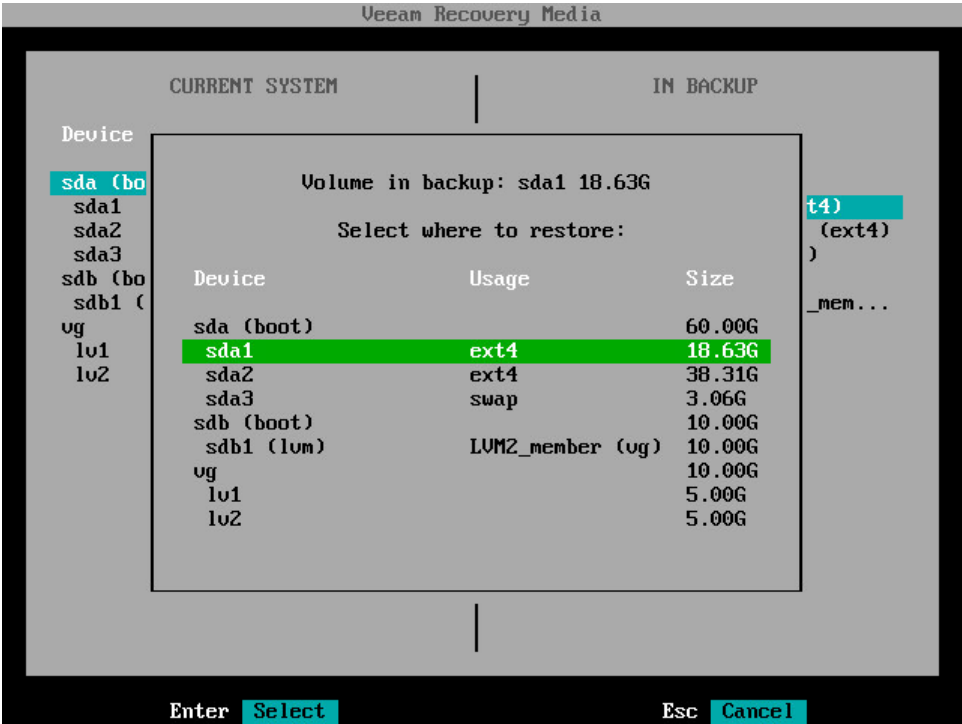
1. In the **In backup** pane, select a volume in the backup whose data you want to recover and press [Enter].



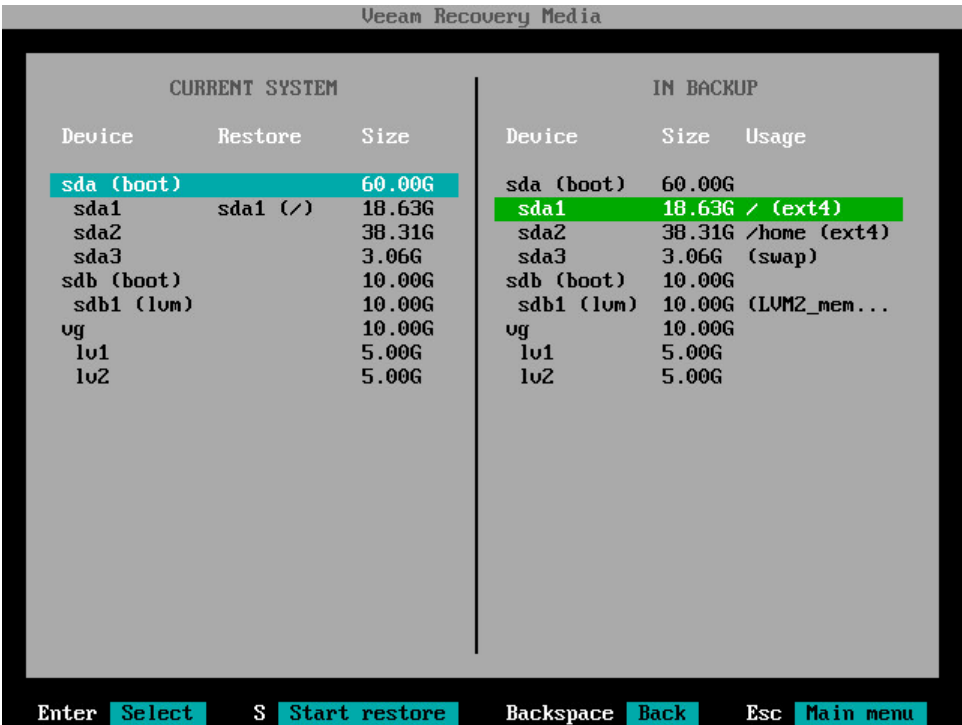
2. Veeam Agent for Linux will display a window with information on the selected volume (partition type, file system type, mount point and volume size) and a list of available operations:
  - **Restore volume to** – select this option if you want to restore the selected volume to your computer.
  - **Close** – select this option if you want to close the window and select another volume.
3. Select the **Restore volume to** option and press [Enter].



- Veeam Agent for Linux will display a list of volumes on your computer. Select the volume that you want to restore and press [Enter].



- In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volume from backup will be restored to the target volume.



- Repeat steps 1-5 for all volumes that you want to restore.
- Press [S] to start the restore process.

## Mapping Target Volume to Source Volume

The **Current system** pane of the **Veeam Recovery Media** wizard displays a partition table of your computer booted from the Veeam Recovery Media. In this pane, you can select volumes on your computer which you want to restore and specify mapping rules for these volumes. If necessary, you can edit the disk layout before restoring volumes.

To map a target volume to a source volume:

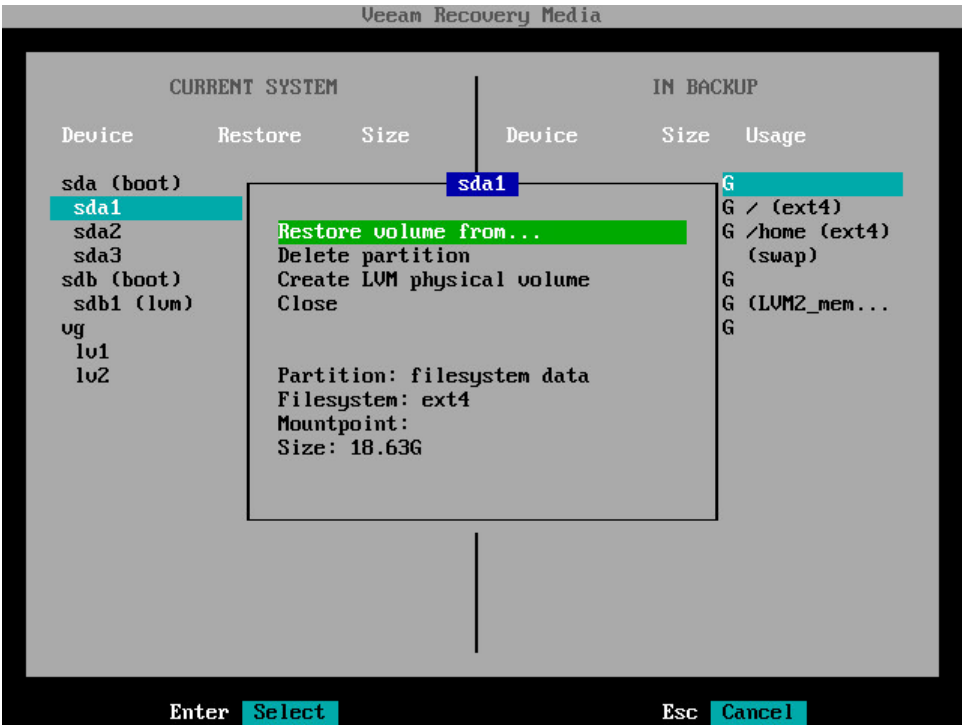
1. In the **Current system** pane, select a volume on your computer whose data you want to recover and press [Enter].

CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lvm)		10.00G	sdb1 (lvm)	10.00G	(LVM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	

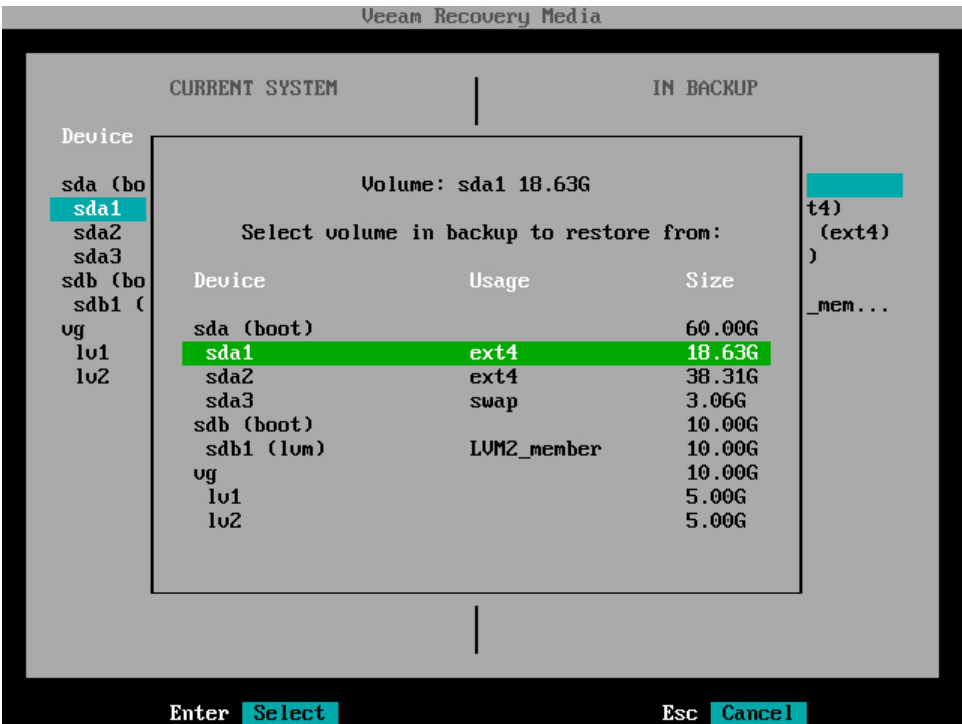
Enter **Select**      Backspace **Back**      Esc **Main menu**

2. Veeam Agent for Linux will display a window with information on the selected volume (partition type, file system type, mount point and volume size) and a list of available operations:
  - **Restore volume from** – select this option if you want to recover the selected volume from the backup.
  - **Delete partition** [for simple volumes] or **Delete volume** [for LVM volumes] – select this option if you want to change the disk layout before restoring a volume. After you delete a partition or volume, you will be able to create a new partition or volume of the desired size and map a volume in the backup to the volume on your computer.
  - [For simple volumes] **Create LVM physical volume** – select this option if you want to create an LVM physical volume on the selected disk partition. In the created physical volume, you will be able to create a volume group and restore to this volume group LVM logical volumes from the backup.
  - **Close** – select this option if you want to close the window and select another volume.

3. Select the **Restore volume from** option and press [Enter].



4. Veeam Agent for Linux will display a window with a list of volumes in the backup. Select the volume that you want to restore and press [Enter].



- In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volume from backup will be restored to the target volume.

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
<b>sda1</b>	<b>sda1 (/)</b>	<b>18.63G</b>	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lum)		10.00G	sdb1 (lum)	10.00G	(LVM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	

- Repeat steps 1-5 for all volumes that you want to restore.
- Press [S] to start the restore process.

## Mapping Source Disk to Target Disk

The **In backup** pane of the **Veeam Recovery Media** wizard contains a list of disks and volumes in the backup. As well as individual volumes, you can select for restore entire computer disks.

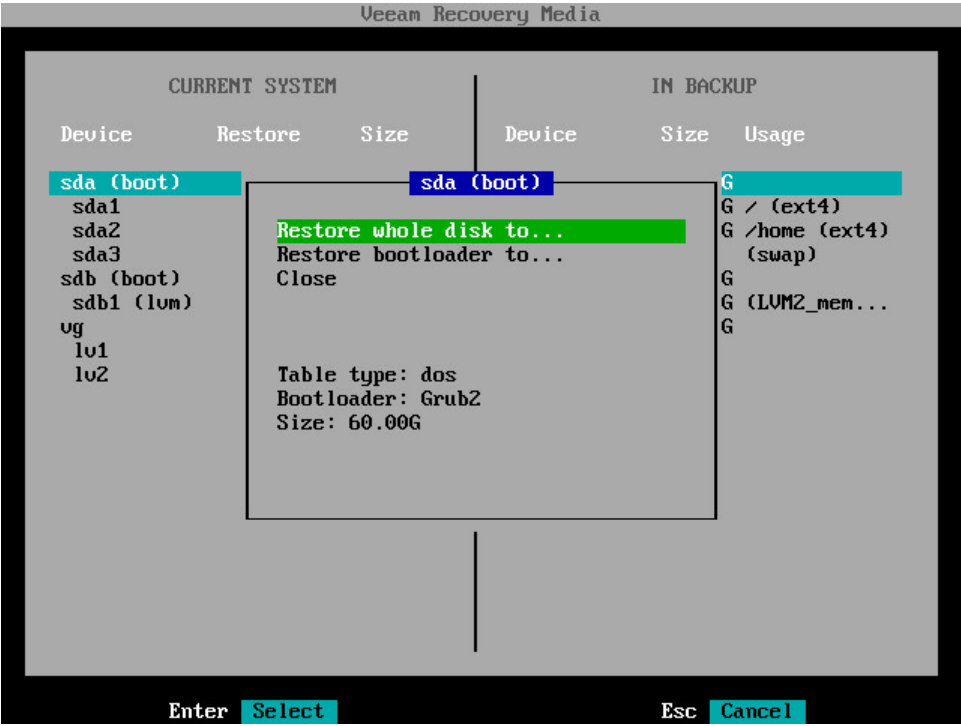
To map a source disk to a target disk:

1. In the **In backup** pane, select a disk in the backup volumes on which you want to recover and press [Enter].


Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lvm)		10.00G	sdb1 (lvm)	10.00G	(LVM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	

2. Veeam Agent for Linux will display a window with information on the selected disk (partition table type, bootloader type and disk size) and a list of available operations:
  - **Restore whole disk to** – select this option if you want to restore all volumes on the selected disk in the backup to your computer.
  - **Restore bootloader to** – select this option if you want to restore a bootloader from the disk in the backup to your computer.
  - **Close** – select this option if you want to close the window and select another disk or volume.

- To restore volumes that reside on the selected disk, select the **Restore whole disk to** option and press [Enter].



- Veeam Agent for Linux will display a list of disks and volumes on your computer. Select the disk whose volumes you want to restore and press [Enter].





- In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volumes from the disk in the backup will be restored to the target disk.

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)	loader (sda)	60.00G	sda (boot)	60.00G	
sda1	sda1 (/)	18.63G	sda1	18.63G	/ (ext4)
sda2	sda2 (/home)	38.31G	sda2	38.31G	/home (ext4)
sda3	sda3 (swap)	3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lum)		10.00G	sdb1 (lum)	10.00G	(LUM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	

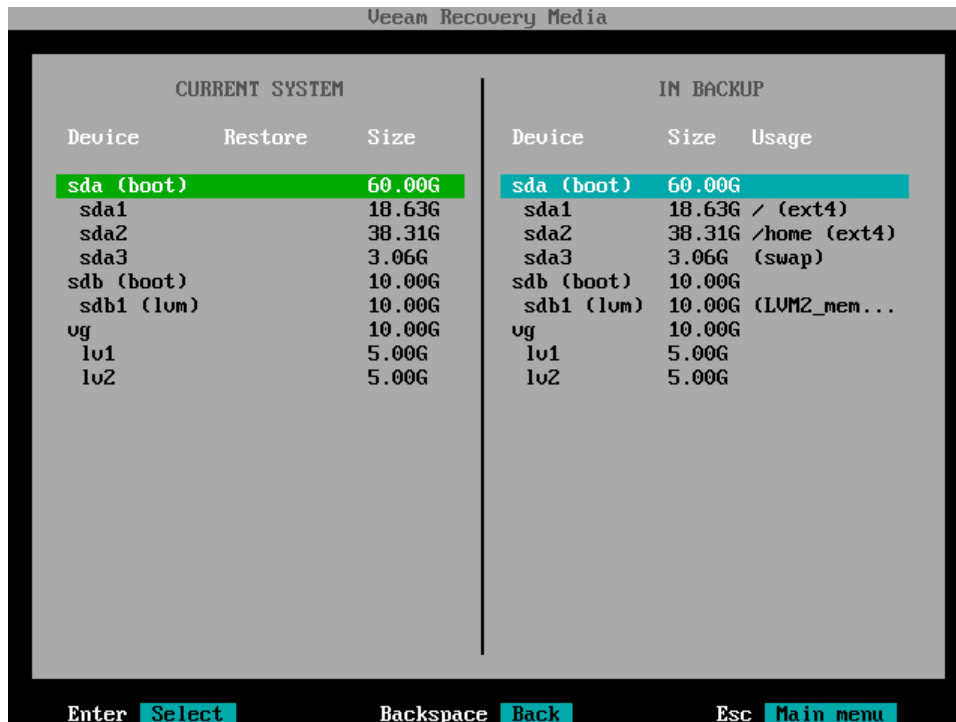
- Repeat steps 1-5 for all computer disks whose volumes you want to restore.
- Press [S] to start the restore process.

## Mapping Target Disk to Source Disk

The **Current system** pane of the **Veeam Recovery Media** wizard displays a partition table of your computer booted from the Veeam Recovery Media. As well as individual volumes, you can select for restore entire computer disks. If necessary, you can edit the disk layout before restoring volumes.

To map a target disk to a source disk:

1. In the **Current system** pane, select a disk on your computer to which you want to restore volumes and press [Enter].



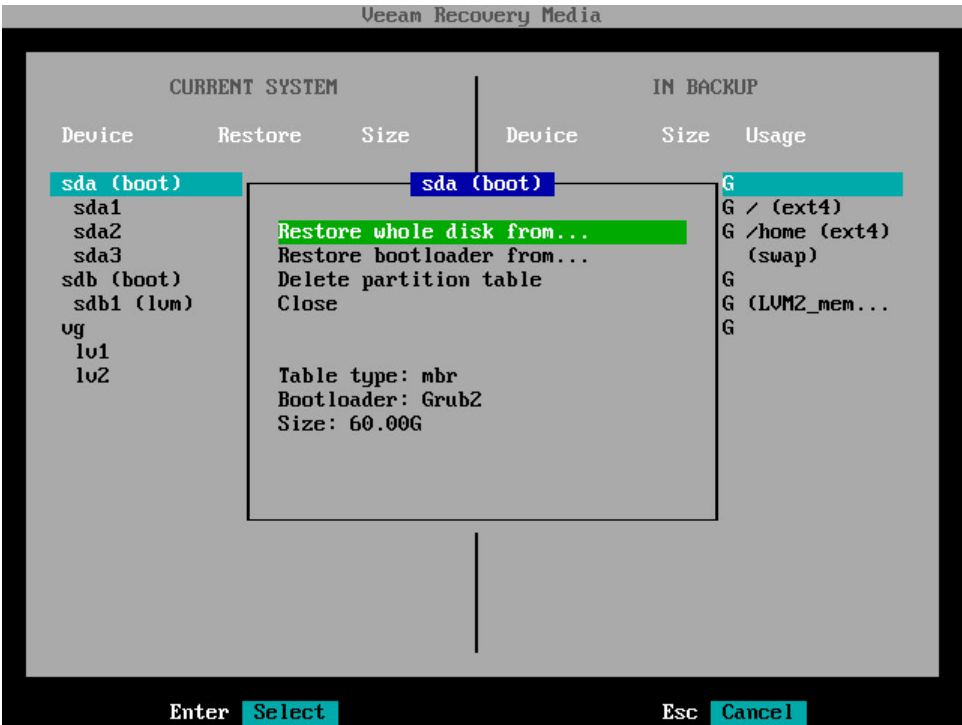
The screenshot shows the Veeam Recovery Media interface. It features two side-by-side tables. The left table, titled 'CURRENT SYSTEM', lists disks and their sizes. The right table, titled 'IN BACKUP', lists disks from a backup and their usage. At the bottom, there are navigation instructions: Enter Select, Backspace Back, and Esc Main menu.

CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lvm)		10.00G	sdb1 (lvm)	10.00G	(LVM2_mem...)
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	

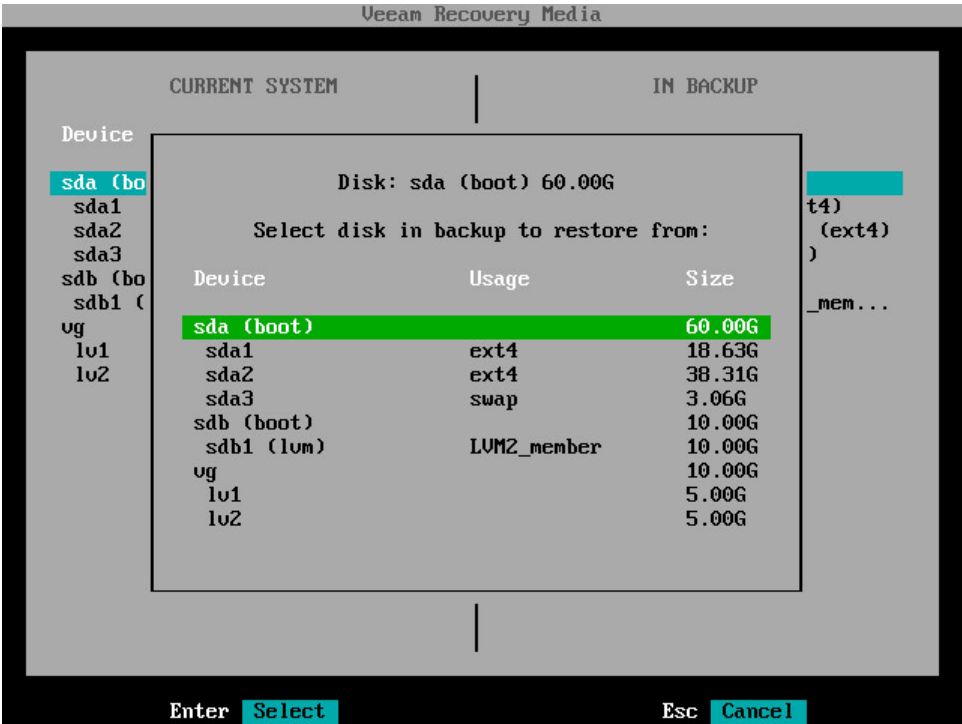
Enter Select      Backspace Back      Esc Main menu

2. Veeam Agent for Linux will display a window with information on the selected disk (partition table type, bootloader type and disk size) and a list of available operations:
  - **Restore whole disk from** – select this option if you want to restore to the selected disk all volumes from a disk in the backup.
  - **Restore bootloader from** – select this option if you want to restore to the selected disk a bootloader from a disk in the backup.
  - **Delete partition table** – select this option if you want to change the disk layout before restoring volumes. After you delete a partition table, you will be able to create a new partition table, create disk partitions and volumes of the desired size, and map volumes in the backup to volumes on your computer.
  - **Close** – select this option if you want to close the window and select another disk or volume.

3. To restore volumes to the selected disk, select the **Restore whole disk from** option and press [Enter].



4. Veeam Agent for Linux will display a list of disks and volumes in the backup. Select the disk whose volumes you want to restore and press [Enter].



5. In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volumes from the disk in the backup will be restored to the target disk.

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)	loader (sda)	60.00G	sda (boot)	60.00G	
sda1	sda1 (/)	18.63G	sda1	18.63G	/ (ext4)
sda2	sda2 (/home)	38.31G	sda2	38.31G	/home (ext4)
sda3	sda3 (swap)	3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lum)		10.00G	sdb1 (lum)	10.00G	(LUM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	

Enter Select S Start restore Backspace Back Esc Main menu

6. Repeat steps 1-5 for all disks whose volumes you want to restore.
7. Press [S] to start the restore process.

## Mapping Btrfs Subvolumes

If the backup contains BTRFS file system data, in the **In backup** pane of the **Veeam Recovery Media** wizard, Veeam Agent displays the list of backed-up BTRFS subvolumes. Information about the original BTRFS pool that contained these subvolumes is not included in the backup.

You can restore from the backup all BTRFS subvolumes or selected subvolumes. To restore a subvolume, you must specify a target BTRFS pool – a BTRFS pool on the computer where you perform restore using the Veeam Recovery Media.

You can restore BTRFS subvolumes to the original BTRFS pool or new BTRFS pool. If the target BTRFS pool contains a subvolume with the same name as the name of the subvolume that you selected for restore, Veeam Agent will automatically map these subvolumes. During the restore process, Veeam Agent will overwrite data on the target subvolume with the data retrieved from the backup.

### NOTE

Veeam Agent for Linux does not check whether the target BTRFS pool has enough disk space to restore the selected subvolumes. If the total size of the restored data is larger than the size of the target BTRFS pool, after the restore process completes, the restored data will be corrupted.

To map a source BTRFS subvolume to a target BTRFS pool:

1. In the **In backup** pane, select a subvolume in the backup whose data you want to restore and press [Enter].

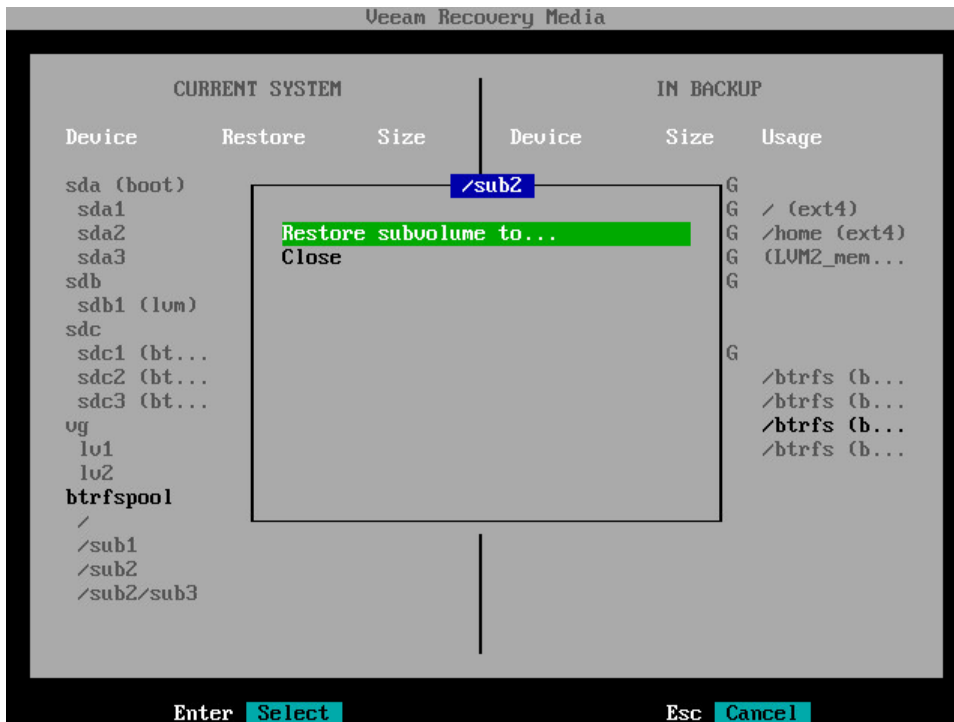
You can also choose to restore all subvolumes from the backup at once. To do this, in the **In backup** pane, select **btrfs** and press [Enter].

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sdb1 (lvm)	10.00G	(LVM2_men...
sdb		10.00G	vg	10.00G	
sdb1 (lvm)		10.00G	lv1	5.00G	
sdc		30.00G	lv2	5.00G	
sdc1 (bt...		10.00G	btrfs	30.00G	
sdc2 (bt...		10.00G	/		/btrfs (b...
sdc3 (bt...		10.00G	/sub1		/btrfs (b...
vg		10.00G	<b>/sub2</b>		<b>/btrfs (b...</b>
lv1		5.00G	/sub2/sub3		/btrfs (b...
lv2		5.00G			
<b>btrfspool</b>		<b>30.00G</b>			
/					
/sub1					
/sub2					
/sub2/sub3					

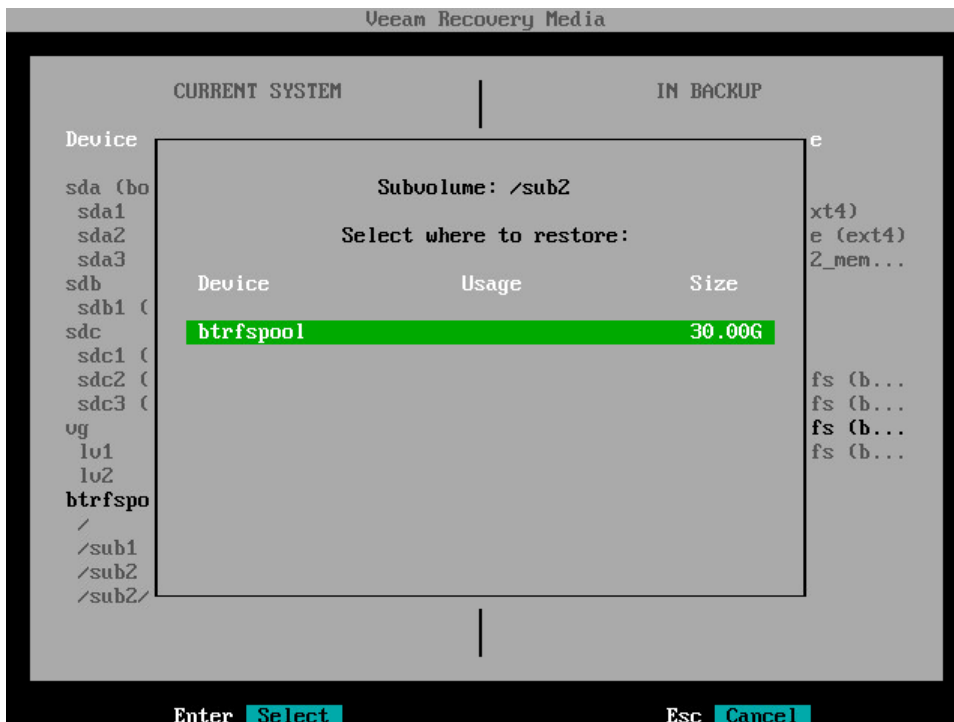
Enter **Select**      Backspace **Back**      Esc **Main menu**

2. In the displayed window, select the necessary option for BTRFS restore and press [Enter]. The available options depend on what BTRFS subvolumes you selected for restore: all subvolumes or specific subvolume.
  - **Restore subvolume to** – this option is available if you chose to restore a specific BTRFS subvolume from the backup. Select this option to restore the selected subvolume to your computer.
  - **Restore btrfs to** – this option is available if you chose to restore all BTRFS subvolumes from the backup. Select this option to restore subvolumes to your computer.

- **Close** – select this option if you want to close the window and select another subvolume.



3. Veeam Agent for Linux will display a list of BTRFS pools on your computer. Select the BTRFS pool where you want to restore data from the backup and press [Enter].



4. In the **Current system** pane, in the **Restore** column, Veeam Agent for Linux will display which subvolume from backup will be restored to the target BTRFS pool.

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sdb1 (lvm)	10.00G	(LVM2_men...
sdb		10.00G	vg	10.00G	
sdb1 (lvm)		10.00G	lv1	5.00G	
sdc		30.00G	lv2	5.00G	
sdc1 (bt...		10.00G	btrfs	30.00G	
sdc2 (bt...		10.00G	/		/btrfs (b...
sdc3 (bt...		10.00G	/sub1		/btrfs (b...
vg		10.00G	/sub2		/btrfs (b...
lv1		5.00G	/sub2/sub3		/btrfs (b...
lv2		5.00G			
btrfs pool		30.00G			
/					
/sub1					
/sub2	/sub2				

5. If you want to restore more than one subvolume, repeat steps 1-4 for all subvolumes that you want to restore.
6. Press [S] to start the restore process.

## Step 9. Complete Restore Process

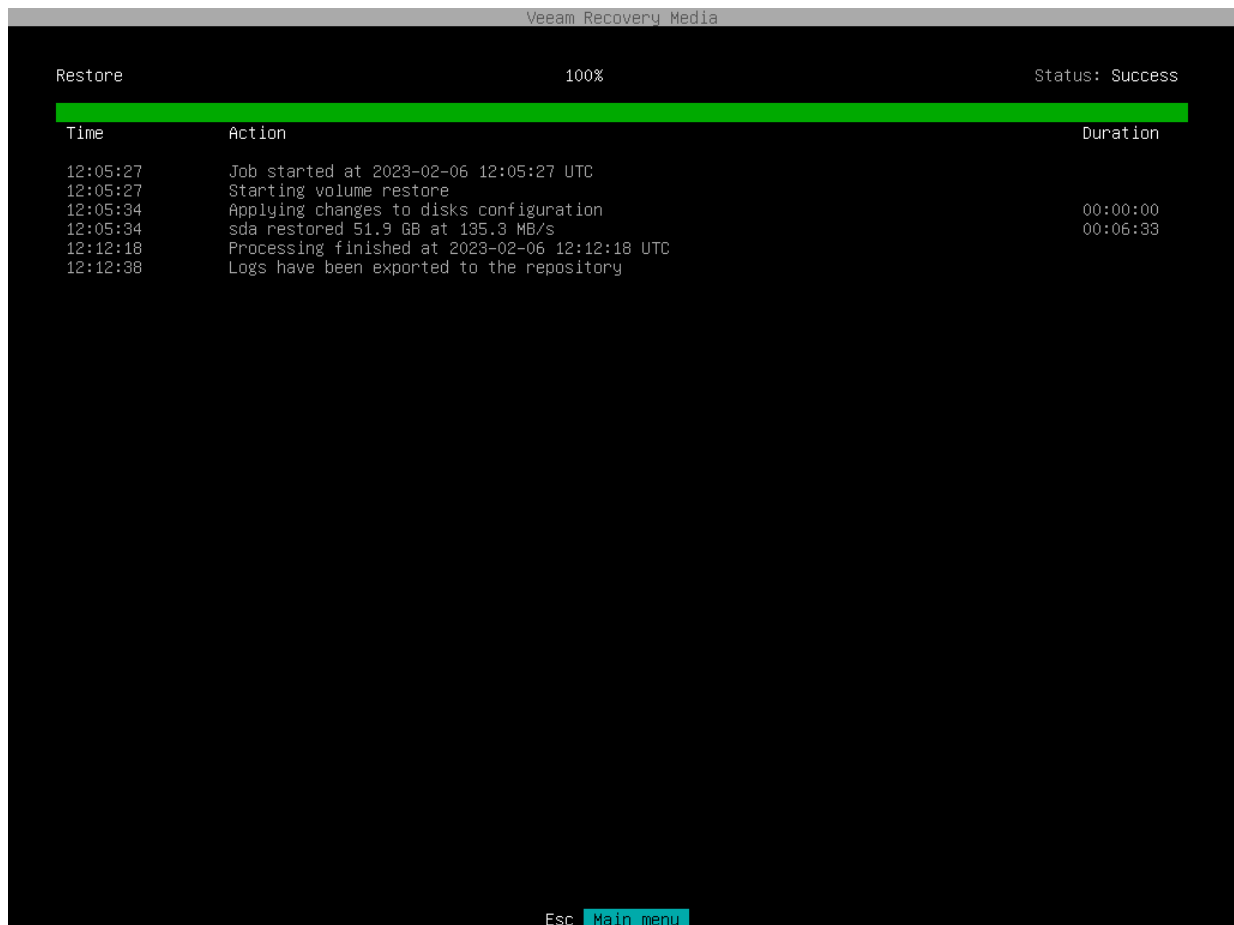
At the **Recovery summary** step of the wizard, complete the procedure of volume-level restore.

1. Review the specified recovery settings.






2. Press [Enter] to start the volume-level restore process. Veeam Agent for Linux will perform partition re-allocation operations if necessary, restore the necessary data from the backup and overwrite data on your computer with it.



## Step 10. Finish Working with Veeam Recovery Media

When the restore operation completes, finish working with the Veeam Recovery Media and start your operating system.

1. Press [Esc] to return to the Veeam Recovery Media main menu.
2. Eject the media or removable storage device with the recovery image.
3. In the Veeam Recovery Media main menu, select the **Reboot** option and press [Enter].



4. Wait for your Linux operating system to start.

# Restoring Files and Folders

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore with the Veeam Recovery Media, Veeam Agent publishes the backup content directly into the recovery image OS file system and displays it in the file browser. You can restore files and folders to their initial location or copy files and folders to a new location.

## Before You Begin

Before you boot from the recovery image and recover your data, check the following prerequisites:

- You must have a recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For data recovery, you can use a volume-level or file-level backup created with Veeam Agent for Linux. Make sure that the backup or system image is available on the computer drive (local or external), on a network shared folder or on the backup repository managed by a Veeam backup server.
- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a Veeam backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- If you restore to a virtual environment, note that the current version of Veeam Recovery Media supports only the VMware and Hyper-V virtualization solutions. To resolve possible issues during bare metal recovery of Oracle VM virtual machines, use instructions in the second section of [this Veeam KB article](#).

# Step 1. Boot from Veeam Recovery Media

To boot from the Veeam Recovery Media:


1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.  
  
[For removable storage device] Attach the removable storage device with the recovery image to your computer.
2. Start your computer.
3. [For regular recovery image] In the boot menu, select what Linux kernel version to use to boot your computer and specify boot options if necessary.

You can select a Linux kernel version if you use generic Veeam Recovery Media downloaded from [the Veeam website](#) or [Veeam software repository](#). If you created a custom Veeam Recovery Media, you will be prompted to boot using the Linux kernel of your Veeam Agent computer included in the recovery image.

To specify boot options, press the [Tab] key and type the necessary options in the command prompt.

## NOTE

For the [legacy recovery image](#), the boot menu is unavailable. After you start your computer, Veeam Agent will immediately start loading files from the Veeam Recovery Media.



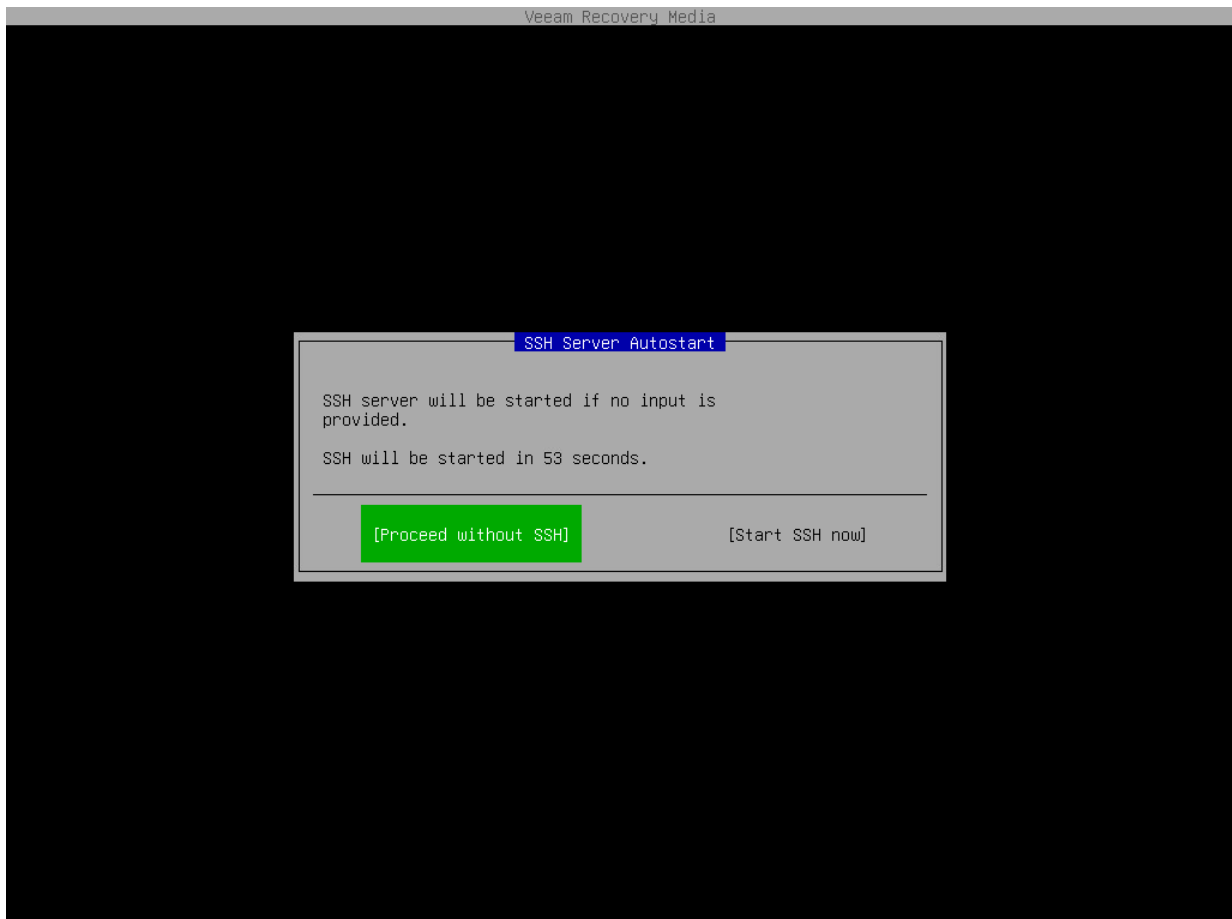
4. Wait for Veeam Agent to load files from the Veeam Recovery Media.

5. After the recovery image OS has loaded, choose whether you want to start the SSH server. The SSH server allows you to connect to the Veeam Recovery Media from a remote machine.

The Veeam Recovery Media starts the SSH server automatically after a time-out. The default value for the time-out is 60 seconds.

To override the default time-out and start the SSH server immediately, select the **Start SSH now** button and press [Enter].

If you do not want to start the SSH server, make sure that the **Proceed without SSH** button is selected and press [Enter]. You will proceed immediately to the step 7.



6. After the SSH server has started, review settings to connect to the Veeam Recovery Media and press [Enter].

The Veeam Recovery Media displays the following connection settings:

- IP address of the computer booted from the Veeam Recovery Media
- User name and password of the account used to connect to the Veeam Recovery Media
- Fingerprints of the computer booted from the Veeam Recovery Media

## NOTE

The user name of the account used to work with the Veeam Recovery Media is *veeamuser*.

If you want to use command-line utilities built in the regular recovery image, use the `sudo` command to provide the *veeamuser* account with privileges of the *root* account.

```
SSH Connection Info

Credentials
login: veeamuser
passwd: kaAnL

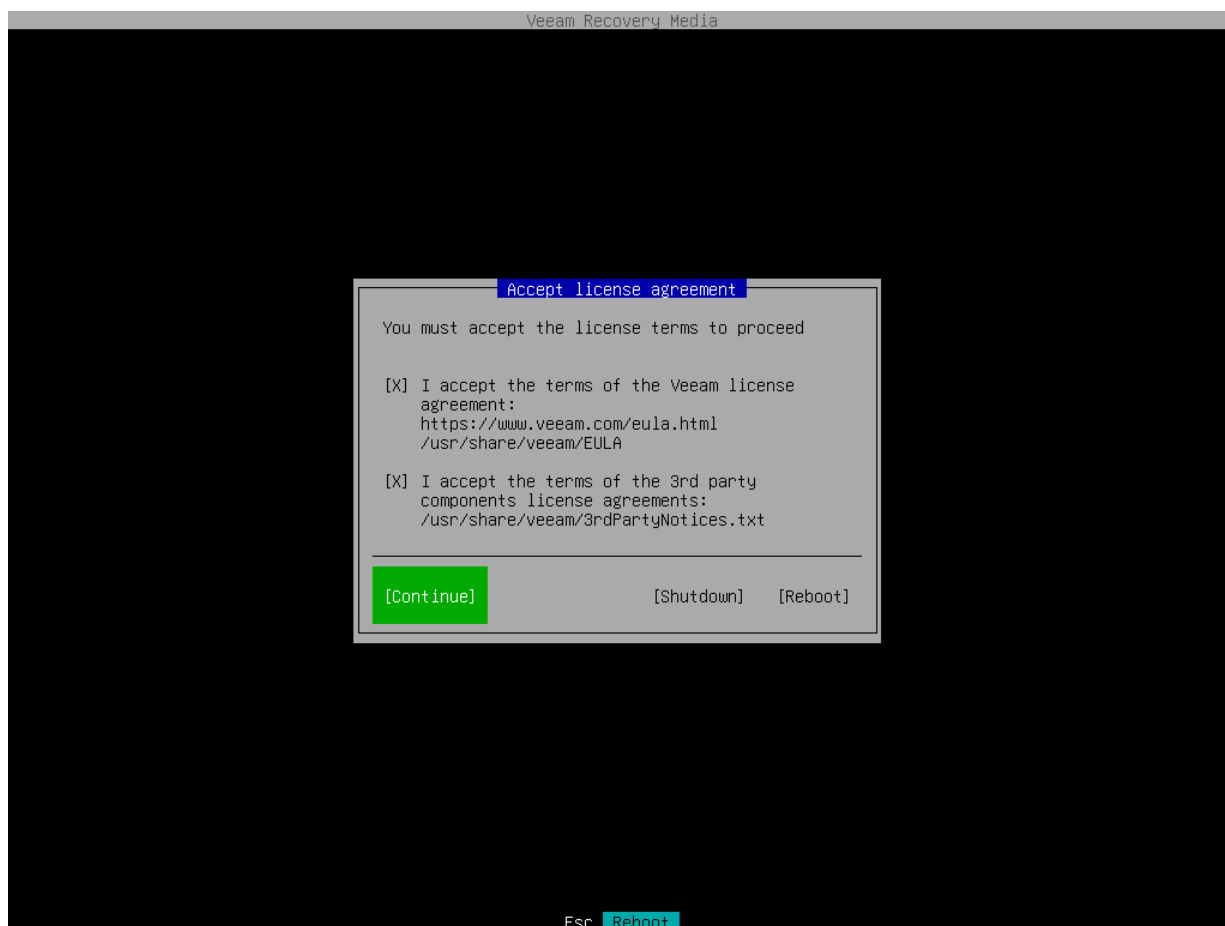
NetConfigs
ens160
IP: 172.24.28.72
IPv6: fd00:ac18:0:1810:0:b5f9:ab46:e4c5

Fingerprints
ecdsa-sha2-nistp256
SHA1:BaXFVwjaWKUf6Rvv2gAwR+g+knI
MD5:6d:9c:56:1d:62:d3:f6:56:f0:0e:62:25:31:da:3c:a2
ssh-ed25519
SHA1:618oSzFazLsSUaMDD/EQJCymqjc
MD5:2b:2b:5d:78:14:66:55:da:cc:7e:6a:bb:29:a3:01:da
ssh-rsa
SHA1:6PsfT1Vv+Gkn8dgdR7420HAsGBQ
MD5:ff:96:15:0b:e4:30:86:67:08:8e:7b:21:47:0c:b4:0a
```

[Continue]

7. Accept the terms of the product license agreement and license agreements for third-party components of the product:
  - a. Make sure that the **I accept Veeam End User Software License Agreement** option is selected and press [Space].
  - b. Select the **I accept the terms of the following 3rd party software components license agreements** option with the [Tab] key and press [Space].


c. Switch to the **Continue** button with the [Tab] key and press [Enter].



8. Make sure that network settings are specified correctly and configure the network adapter if necessary. To learn more, see [Configure Network Settings](#).
9. Choose the necessary recovery option. Veeam Agent offers the following tools:
  - **Restore volumes** – the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
  - **Restore files** – the File Level Restore wizard to restore files and folders to the original location or to a new location.
  - **Exit to shell** – Linux shell prompt with standard utilities to diagnose problems and fix errors.

## TIP

To stop working with the Veeam Recovery Media and shut down or restart your computer, in the Veeam Recovery Media main menu, select the **Reboot** or **Shutdown** option and press [Enter].







## Step 2. Configure Network Settings

If there is a DHCP server in your network, Veeam Agent will configure the network settings automatically. To verify or configure network settings manually, use **nmtui**, a text-based user interface network manager tool provided with Veeam. To learn more about working with nmtui, see [Linux documentation](#).

1. In the Veeam Recovery Media main menu, select **Configure network** and press [Enter].




2. To add new or modify existing connection, in NetworkManager, select **Edit a connection**.



3. After you add or edit a connection, in the main menu of the NetworkManager, select **Activate a connection**.
  - a. If the connection is new, choose it in the list of connections; then select **Activate**.
  - b. If the connection was modified, you must reactivate it. To do this, choose it in the list of connections and select **Deactivate**; then choose the connection again and select **Activate**.
4. After you finish working with Network Manager, press [Esc] to return to the Veeam Recovery Media main menu and launch the File Level Restore wizard.

## Step 3. Launch File Level Restore Wizard

To launch the file-level restore wizard, in the Veeam Recovery Media main menu, select **Restore files** and press [Enter].




## Step 4. Select Backup Location

At the **Select backup location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

To recover data from backup, you need to mount the backup storage on which the backup file resides to the recovery image OS file system. Veeam Agent automatically mounts external USB drives that are connected to the computer and displays them in the list of available backup locations. You can select the necessary device and press [Enter] to pass to the [Browse for backup files](#) step of the wizard.

If the backup file is located in a network shared folder or on a local drive, select one of the following options:

- **Mount local disk** – select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. With this option selected, you will pass to the [Select local disk](#) step of the wizard.
- **Add object storage repository** – select this option if the backup file resides in an object storage repository. With this option selected, you will pass to the [Select cloud storage type](#) step of the wizard.
- **Add shared folder** – select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the [Mount shared folder](#) step of the wizard.
- **Add VBR server** – select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Specify backup server parameters](#) step of the wizard.
- **Add Cloud provider** – select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the [Specify Cloud provider parameters](#) step of the wizard.



## Step 5. Specify Backup Location Settings

Specify settings for the target storage that contains a backup file from which you plan to restore data:

- [Specify shared folder settings](#) — if you have selected the **Add shared folder** option at the [Select backup location](#) step of the wizard.
- [Select local drive](#) — if you have selected the **Mount local disk** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam backup repository settings](#) — if you have selected the **Add VBR server** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam Cloud Connect repository settings](#) — if you have selected the **Add Cloud provider** option at the [Select backup location](#) step of the wizard.
- [Specify object storage repository settings](#) — if you have selected the **Add object storage repository** option at the [Select backup location](#) step of the wizard.

### Shared Folder Settings

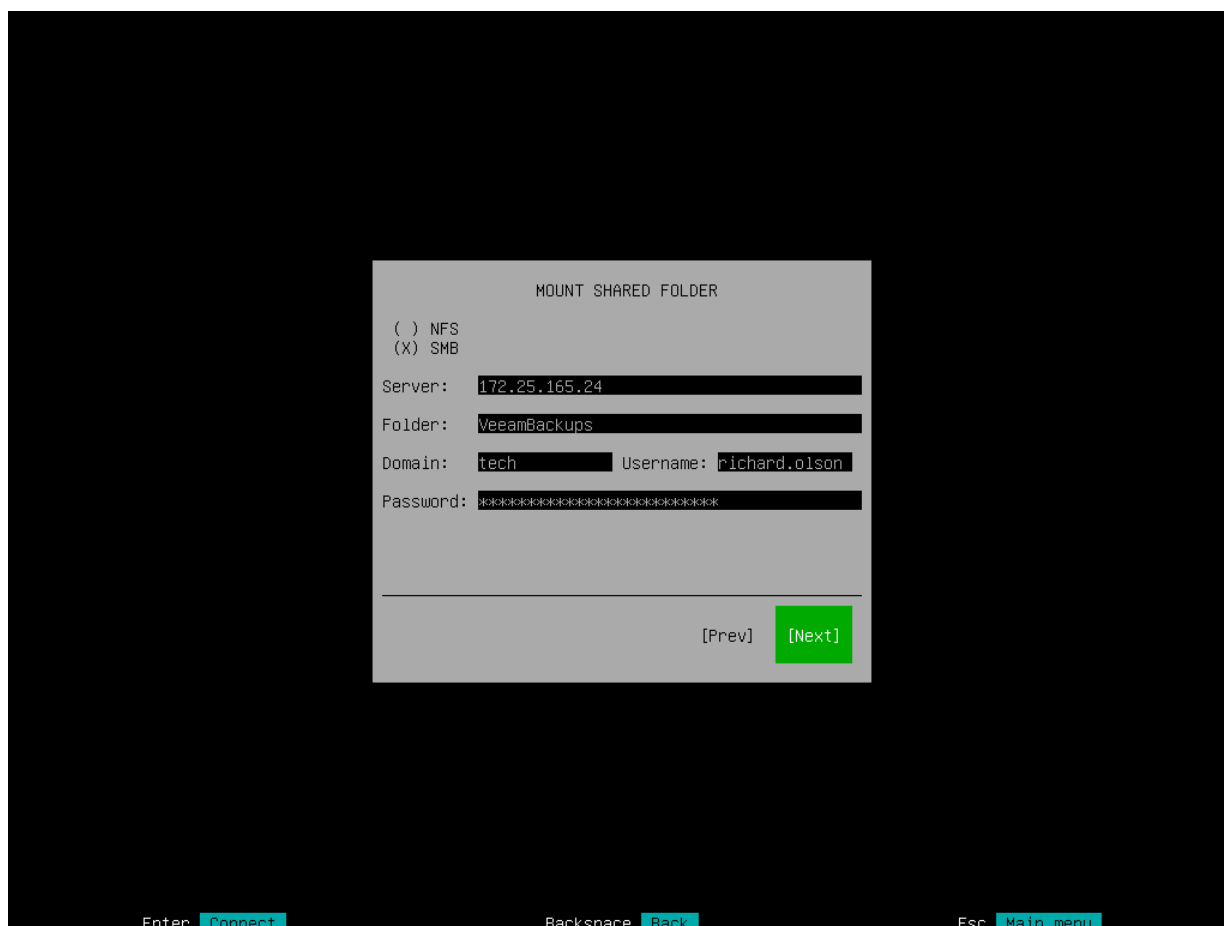
The **Mount shared folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. Select the type of a network shared folder:
  - **NFS** — to connect to a network shared folder using the NFS protocol.
  - **SMB** — to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Path** field, specify the network shared folder name in the *SERVER/DIRECTORY* format: type an IP address or domain name of the server and the name of the network shared folder in which the backup file resides.
3. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
4. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
5. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.

## TIP

You can mount several network shared folders to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Add shared folder** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted network shared folders and browse for a backup file located on the necessary storage.




## Local Backup Repository Settings

The **Select local disk** step of the wizard is available if you have selected to restore data from a backup file located on a computer drive.

In the list of devices, select the necessary disk or disk partition and press [Enter]. Veeam Agent will mount the selected device to the `/media` directory of the recovery image OS file system and display content of the directory.

## TIP

You can mount several devices to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Mount local disk** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted devices and browse for a backup file located on the necessary storage.



## Veeam Backup Repository Settings

The **Specify Backup Server parameters** step of the wizard is available if you have selected to restore data from a backup repository managed by the Veeam backup server.

Specify settings for the Veeam backup server that manages the backup repository where the backup file resides:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. Select the type of **Authentication** to access the Veeam backup server:
  - o **Login and password.** With this option selected, specify the following settings:
    - i. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
    - ii. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.

- iii. In the **Password** field, type a password of the account that has access to the Veeam backup repository.

#### NOTE

If you want to perform restore from a backup created by Veeam Agent operating in the managed mode, you must use an account that has the Veeam Backup Administrator or Veeam Restore Operator role on the Veeam backup server. For more information about user roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

The screenshot shows the Veeam Recovery Media interface. At the top, it says "Veeam Recovery Media". The main area is a dark gray rectangle. In the center, there is a light gray dialog box titled "Specify Backup Server parameters:". Inside the dialog box, the following fields are visible: "Address: 172.24.31.136", "Port: 10006", "Authentication:" with two options: "(X) Login and password" and "( ) Recovery token", "Login: Administrator", "Domain:", and "Password: \*\*\*\*\*". At the bottom right of the dialog box, there are two buttons: "[Prev]" and "[Next]". Below the dialog box, at the bottom of the screen, there are three keyboard shortcuts: "Enter Connect", "Backspace Back", and "Esc Main menu".



- **Recovery token:** With this option selected, in the **Token** field, enter the value of the recovery token generated in the Veeam Backup & Replication console. For more information on generating recovery tokens, see [Creating Recovery Token](#) in the Veeam Agent Management guide.

Veeam Recovery Media

Specify Backup Server parameters:

Address: 172.24.31.136

Port: 10006

Authentication:  
( ) Login and password  
(X) Recovery token


Token: 5e6d-45aA-DBBd-d0ec

Recovery tokens can be created using the Veeam Backup & Replication console.

[Prev] [Next]

Enter Connect Backspace Back Esc Main menu

4. Press [Enter]. Veeam Agent will connect to the Veeam backup server. If prompted, accept the self-signed certificate of the Veeam backup server to continue.



After successful connection to the Veeam backup server, you will pass immediately to the [Backup](#) step of the wizard.

## Veeam Cloud Connect Repository Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings.](#)
2. [Verify the TLS certificate.](#)
3. [Specify user account settings.](#)


## Specifying Service Provider Settings

The **Specify Cloud provider parameters** step of the wizard is available if you have selected to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.

3. Press [Enter]. Veeam Agent will connect to the service provider and display the [Certificate details](#) window.




## Verifying TLS Certificate

In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate.

- To accept the TLS certificate, select the **Accept** button with the [Tab] key and press [Enter].
- [Optional] To verify the TLS certificate with a thumbprint, do the following:
  - a. Select the **Verify thumbprint** button with the [Tab] key and press [Enter].
  - b. In the **Thumbprint verification** field, enter the thumbprint that you obtained from the SP.

- c. Switch to the **Verify** button and press [Enter]. Veeam Agent for Linux will check if the thumbprint that you entered matches the thumbprint of the obtained TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.




## Specifying User Account Settings

The **Specify Cloud provider credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

1. In the **Username** field, type a name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.
2. In the **Password** field, provide a password for the tenant or subtenant account.

3. Press [Enter]. Veeam Agent for Linux will connect to the cloud repository, and you will pass immediately to the [Backup](#) step of the wizard.



## Object Storage Repository Settings

If you have selected to restore data from a backup file located in a object storage repository, specify settings to connect to the repository:

At the **Select cloud storage type** step of the wizard, select one of the following options:


- **S3 Compatible** – select this option if you want to import a backup from an S3 compatible storage repository.

### TIP

If you plan to restore from backups in an IBM or Wasabi object storage, use the **S3 Compatible** storage option.

- **Amazon S3** – select this option if you want to import a backup from an Amazon S3 storage repository.
- **Google Cloud Storage** – select this option if you want to import a backup from a Google Cloud storage repository.

- **Microsoft Azure Blob Storage** – select this option if you want to import a backup from a Microsoft Azure storage repository.



## Specifying Settings for S3 Compatible Repository

If you have selected to import backup from an S3 Compatible storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.

### NOTE

If you want to connect to the repository using the IPv6 address and the port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is an IPv6 address of the cloud storage.
- `port` is a number of a port that Veeam Agent will use to connect to the cloud storage.

2. In the **Region** field, specify a storage region based on your regulatory and compliance requirements.
3. In the **Access key** field, enter an access key ID.
4. In the **Secret key** field, enter a secret access key.

The screenshot shows a terminal window titled "Veeam Recovery Media". In the center is a dialog box titled "Specify S3 compatible storage:". The dialog contains the following fields and values:

- Service point: `https://myservicepoint.com:9000`
- Region: `reg-1`
- S3 compatible account:
- Access key: `Access_Key`
- Secret key: `*****`

At the bottom right of the dialog are two buttons: "[Prev]" and "[Next]". The "[Next]" button is highlighted in green. At the bottom of the terminal window, there are two keyboard shortcuts: "Enter Connect" and "Esc Main menu".


## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Bucket** field, specify a bucket on the storage:
  - a. Click **Browse**.
  - b. In the **Buckets** window, select the necessary bucket and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



## Specifying Settings for Amazon S3 Repository

If you have selected to store backup files on an Amazon S3 storage, specify settings to connect to the storage:

1. [Specify account settings](#).
2. [Specify bucket settings](#).

## Specifying Account Settings


The **Account** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter an access key ID.
2. In the **Secret key** field, enter a secret access key.



3. In the **AWS region** window, select an AWS region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Global** region.




## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
  - a. Click **Browse**.
  - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



## Specifying Settings for Google Cloud Repository


If you have selected to import backup from a Google Cloud storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository.

To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Agent will use the HMAC key to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see the [Google Cloud documentation](#).




## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
  - a. Click **Browse**.
  - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



## Specifying Settings for Microsoft Azure Repository

If you have selected to import backup from a Microsoft Azure storage repository, specify settings to connect to the storage:

1. [Specify account settings](#).
2. [Specify container settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository.

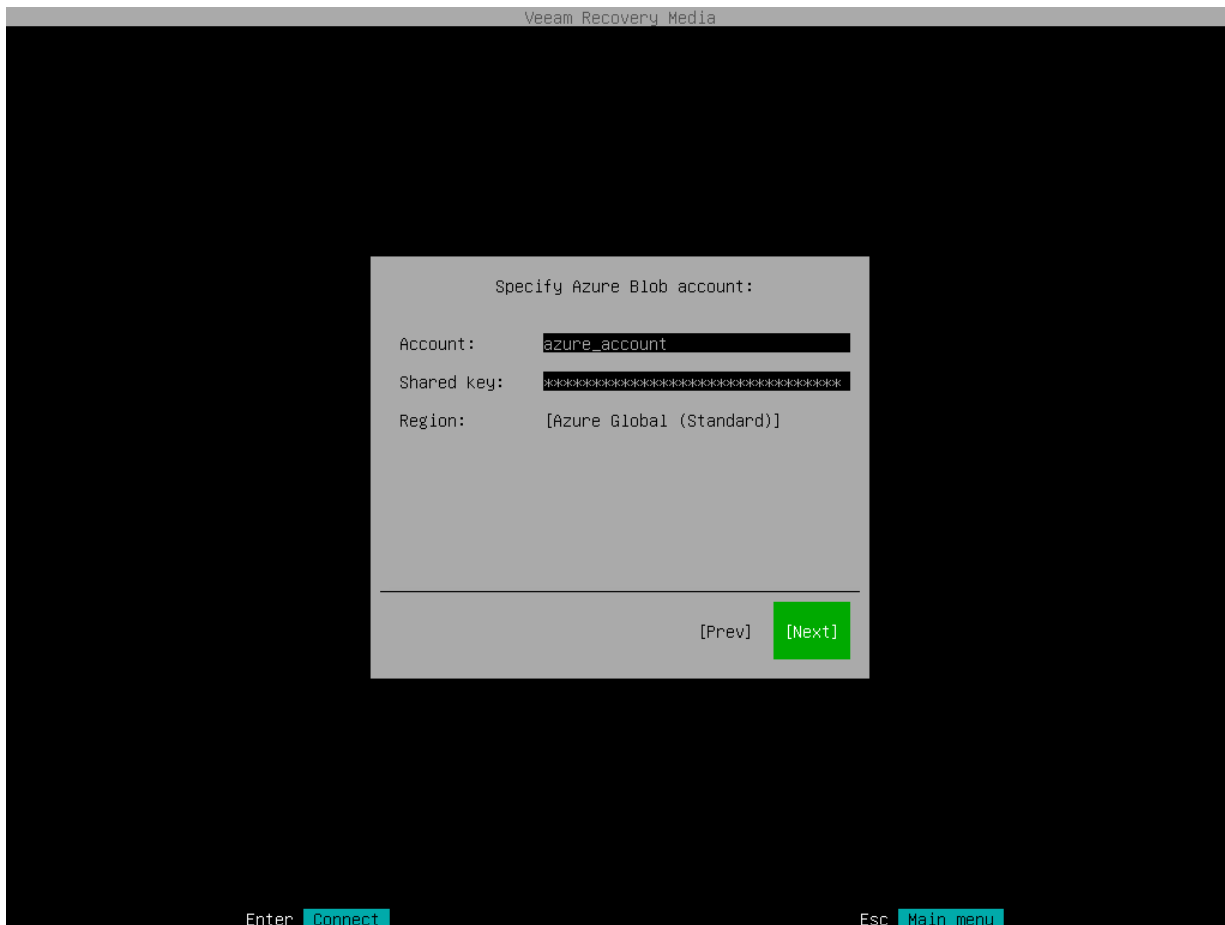
### NOTE

The **Allow storage account key access** option for Shared Key authorization must be enabled in the storage account. For more information on how to find this option, see [Microsoft Docs](#).

To connect to the Microsoft Azure storage, specify the following:

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.

3. In the **Region** window, select a Microsoft Azure region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Azure Global (Standard)** region.




## Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository and specified account settings to connect to the storage.

Specify settings for the container on the storage:

1. In the **Container** field, specify a container on the storage:
  - a. Click **Browse**.
  - b. In the **Containers** window, select the necessary container and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
  - a. Click **Browse**.


b. In the **Folders** window, select the necessary folder and click **OK**.



## Step 6. Browse for Backup File

At the **Browse for backup files** step of the wizard, select the backup file that you plan to use for volume-level restore:

1. In the file system tree, select a directory in which the backup file you plan to use for restore resides:
  - Use the [Up] and [Down] keys to select a directory.
  - Press [Enter] to open the necessary directory.
2. In the directory where the backup file resides, select the backup file and press [Enter].



## Step 7. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.


The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays information about backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported backups** pane, ensure that the backup from which you want to recover data is selected and press [Enter].

If you want to select another backup, press the [i] key and browse for the necessary backup file. To learn more, see [Locate Backup File](#).

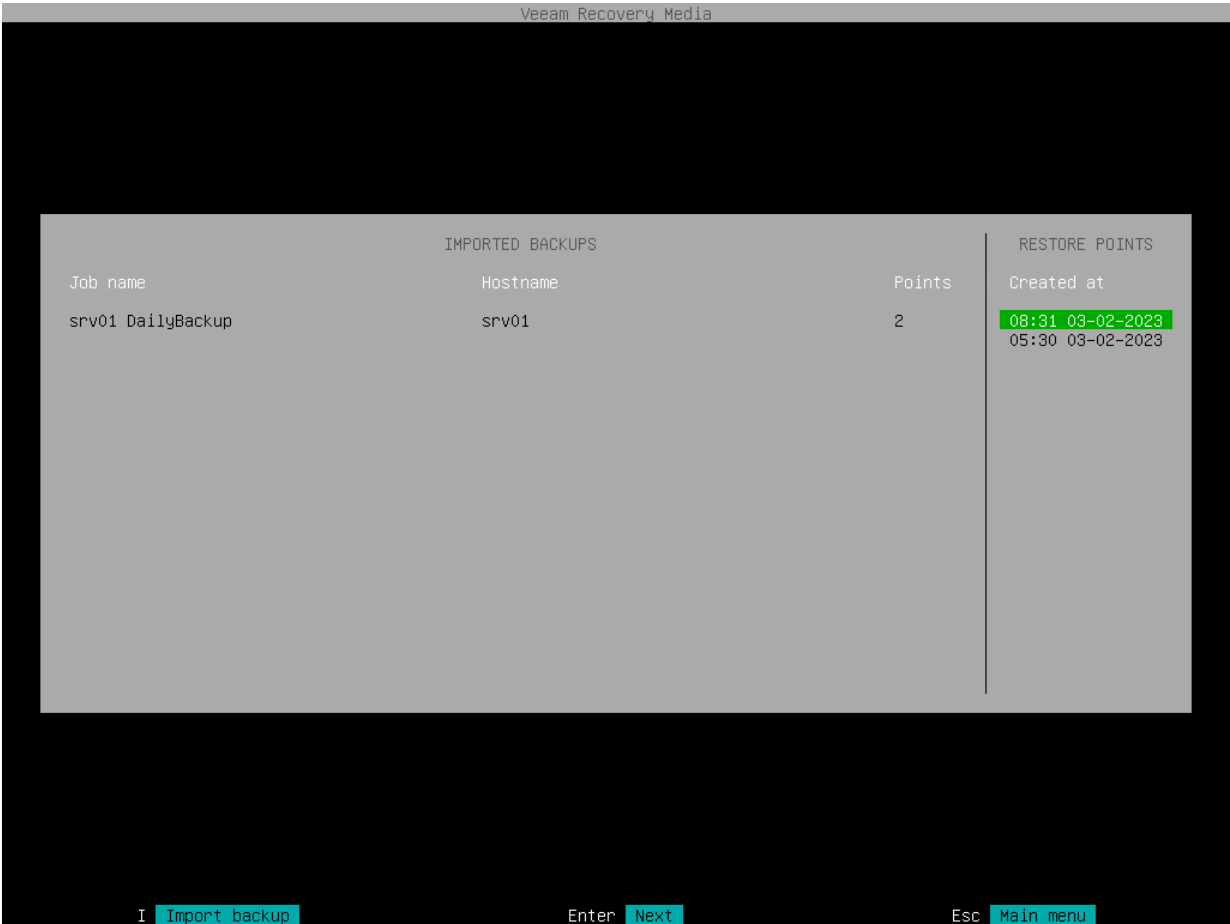


2. In the **Restore points** pane, select with the [Up] and [Down] keys the restore point from which you want to recover data and press [Enter].



NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).



3. Veeam Agent will mount the content of the backup file to the `/mnt/backup` directory in the recovery image OS file system and display a notification window with the corresponding message. Press [Enter] to proceed to the File Level Restore wizard menu, open the file manager and save restored files.

When you perform file-level restore with the File Level Restore wizard, Veeam Agent always mounts the backup to the `/mnt/backup` directory. If you want to specify another directory for backup mount, you can perform file-level restore with the Veeam Agent command line interface. To learn more, see [Restoring Files and Folders with Command Line Interface](#).



## Step 8. Save Restored Files

When the backup file content is mounted to the recovery image OS file system, Veeam Agent opens the File Level Restore wizard menu displaying a list of available operations.

### NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:


- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

You can perform the following operations with file-level backup:

- **Start file browser** – select this option if you want to start the file manager and work with restored files and folders. To learn more, see [Working with Midnight Commander](#).
- **Stop backup mount** – select this option if you want to stop the backup mount session and unmount the backup file content from the `/mnt/backup` directory of the recovery image OS file system. To learn more, see [Stopping Backup Mount Session](#).
- **Exit to shell** – select this option if you want to open the Linux shell prompt and use common Linux command-line tools.

## TIP

To stop working with the Veeam Recovery Media and shut down or restart your computer, in the File Level Restore wizard menu, select the **Reboot** or **Shutdown** option and press [Enter].



## Working with Midnight Commander

To work with restored files and folders, you can use Midnight Commander – a file manager that is included into the Veeam Recovery Media. With the Midnight Commander file manager, you can browse the mounted backup content and file system on your computer, and save restored files and folders to the original location or to a new location.

To launch the file manager, in the File Level Restore wizard menu, select **Start file browser** and press [Enter].

When you launch Midnight Commander, Veeam Agent displays in the file manager the directory with the backup content and your computer's file system:

- In the left pane, Veeam Agent displays a directory of your computer's file system mounted under the `/mnt/system` directory of the recovery image OS file system. By default, Veeam Agent mounts to the recovery image OS file system the following volumes of your computer:
  - If you use a volume-level backup for file-level restore, Veeam Agent detects the partition table in the backup, mounts to the `/mnt/system` directory block devices that represent volumes of your computer with the same names as volumes in the backup. For example, if your volume-level backup contains `/dev/sda1` and `/dev/sda6` volumes with `/` and `/home` mount points, Veeam Agent will mount to the `/mnt/system` directory both root (`/`) and `/home` partitions.

- If you use a file-level backup for file-level restore, Veeam Agent mounts to the `/mnt/system` directory only the system volume of your computer, for example, `/dev/sda1`. If you want to save restored files and folders to a directory on another computer volume or to a network shared folder, you need to mount this volume or folder manually. To mount a target storage for restored files:

- In Midnight Commander, press [F10] to close the file manager.
- In the **File Level Restore** wizard menu, select the **Exit to shell** option and press [Enter].
- Mount the target storage for the restored files and folders with the `mount` command.

- In the right pane, Veeam Agent displays a directory in which the backup content is mounted. Veeam Agent mounts the backup content under the `/mnt/backup` folder.

While the Midnight Commander file manager is open, you can perform the following operations with restored files and folders:

- [Save files to initial location](#)
- [Save files to a new location](#)

After you finish working with files and folders, [finish working with the Veeam Recovery Media](#).

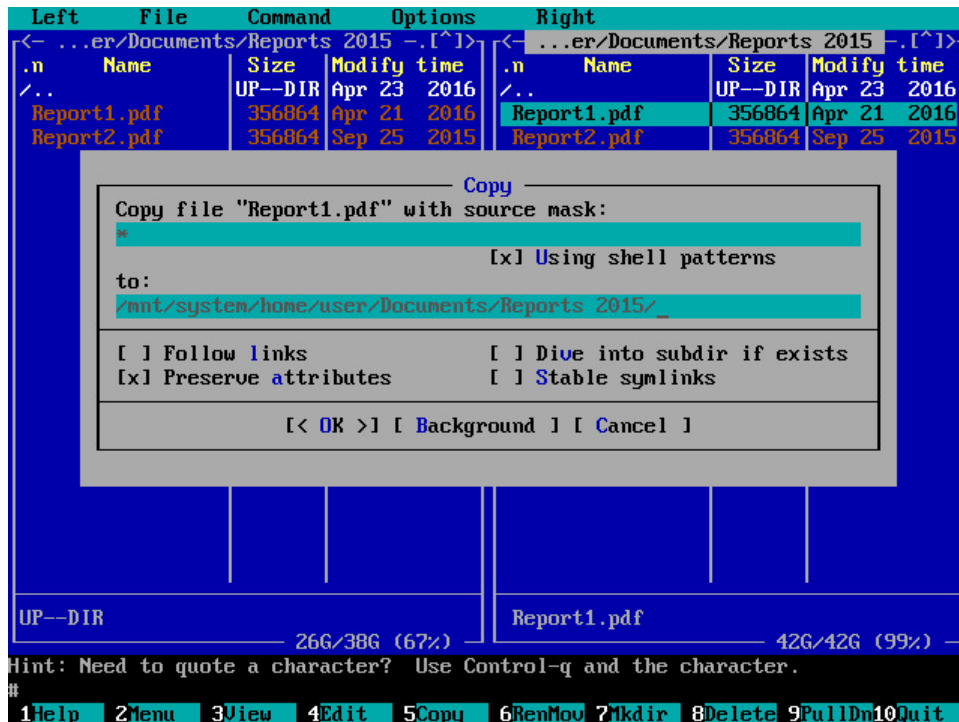
Left		File	Command	Options	Right														
<-		/mnt/system	.[^]>		<- /mnt/backup		.[^]>												
.n	Name	Size	Modify	time	.n	Name	Size	Modify	time										
/..		UP--DIR		Nov 23 20:09	/..		UP--DIR		Nov 23 20:09										
/bin		4096	Nov 22 14:40		/FileLevelBackup_0		4096	Nov 23 18:06											
/boot		4096	Nov 22 14:43																
/dev		4096	Mar 25 2016																
/etc		12288	Nov 23 17:48																
/home		4096	Nov 23 18:20																
/lib		4096	Mar 25 2016																
/lib64		4096	Nov 22 14:40																
/lost+found		16384	Mar 25 2016																
/media		4096	Nov 21 19:21																
/mnt		4096	Nov 2 20:23																
/opt		4096	Mar 25 2016																
/proc		4096	Aug 26 2015																
/root		4096	Nov 21 20:39																
/run		4096	Mar 25 2016																
/sbin		12288	Nov 22 14:41																
/sda5		4096	Apr 20 2016																
/srv		4096	Mar 25 2016																
/sys		4096	Apr 6 2015																
/tmp		4096	Nov 23 18:23																
/usr		4096	Mar 25 2016																
UP--DIR					UP--DIR														
		4867M/9258M (52%)					658M/955M (68%)												
Hint: Completion works on all input lines in all dialogs. Just press M-Tab.																			
#																			
1	Help	2	Menu	3	View	4	Edit	5	Copy	6	RenMov	7	Mkdir	8	Delete	9	PullDn	10	Quit

## Saving Files to Initial Location

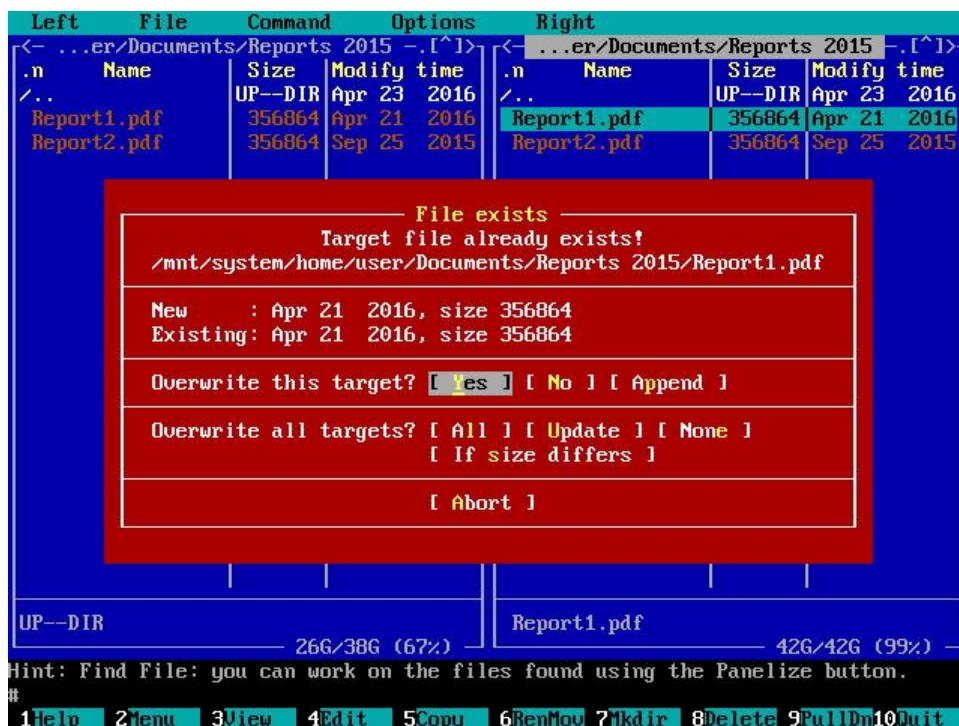
To save restored files or folders to their initial location on your computer, do the following:

1. In the left pane of the file manager window, open the directory in your computer's file system in which the backed-up file or folder that you want to restore originally resided.
2. In the right pane of the file manager window, open the directory that contains the file or folder in the backup that you want to restore to its original location.
3. Select the file or folder that you want to restore and press [F5].

4. In the **Copy** dialog window, review the file or folder copy settings, select **Ok** and press [Enter].



5. If the file or folder you want to restore exists in its original location, Midnight Commander will display a warning. In the warning window, select the necessary operation with the target file or folder and press [Enter]. Midnight Commander will save the file or folder in its original location.

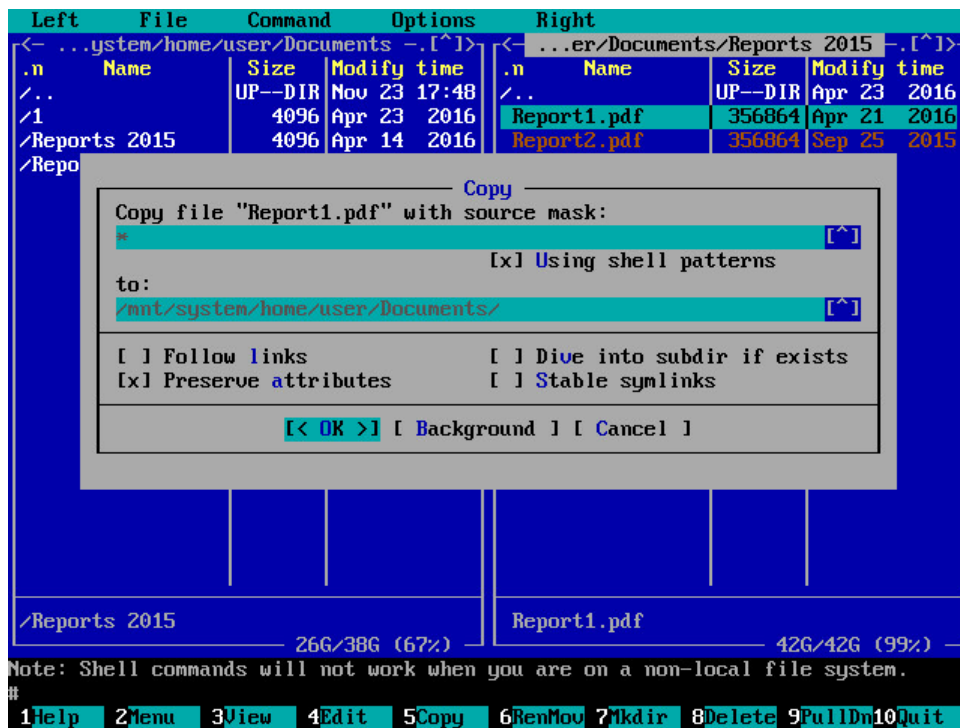


6. After you finish working with files and folders, press [F10] to close the file manager.

## Saving Files to New Location

To save restored files or folders to a new location on your computer or to a network shared folder, do the following:

1. In the left pane of the file manager window, open the directory in your computer's file system in which you want to restore a file or folder.
2. In the right pane of the file manager window, open the directory that contains the file or folder in the backup that you want to restore.
3. Select the file or folder that you want to restore and press [F5].
4. In the **Copy** dialog window, review the file or folder copy settings, select **Ok** and press [Enter].



- Midnight Commander will save the file or folder to the specified location.

Left	File	Command	Options	Right					
<- ...ystem/home/user/Documents -.[^]>				<- ...er/Documents/Reports 2015 -.[^]>					
.n	Name	Size	Modify time	.n	Name	Size	Modify time		
./..		UP--DIR	Nov 23 17:48	./..		UP--DIR	Apr 23 2016		
/1		4096	Apr 23 2016	Report1.pdf		356864	Apr 21 2016		
/Reports 2015		4096	Apr 14 2016	Report2.pdf		356864	Sep 25 2015		
/Reports 2016		4096	Apr 14 2016						
Report1.pdf		356864	Apr 21 2016						
/Reports 2015				Report1.pdf					
26G/38G (67%)				42G/42G (99%)					
Hint: M-t changes quickly the listing mode.									
#									
1Help	2Menu	3View	4Edit	5Copy	6RenMov	7Mkdir	8Delete	9PullDn	10Quit


- After you finish working with files and folders, press [F10] to close the file manager.

## Stopping Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. To unmount a backup, you need to stop the backup mount session. This may be required, for example, if you want to stop working with files and folders in one backup and mount another backup for file-level restore.




To stop the backup mount session with the Veeam Recovery Media, in the File Level Restore wizard menu, select the **Stop backup mount** option and press [Enter]. Veeam Agent will stop the backup mount session, unmount the backup from the `/mnt/backup` directory of the recovery image OS file system, exit the File Level Restore wizard and display the Veeam Recovery Media main menu.



## Step 9. Finish Working with Veeam Recovery Media

When the restore operation completes, finish working with the Veeam Recovery Media and start your operating system.

1. Eject the media or removable storage device with the recovery image.
2. In the File Level Recovery wizard menu or Veeam Recovery Media main menu, select the **Reboot** option and press [Enter].



3. Wait for your Linux operating system to start.

# Restoring Volumes with Command Line Interface

You can restore a specific computer volume or all volumes from the volume-level backup.

## NOTE

You cannot use the Veeam Agent for Linux command line interface to restore BTRFS subvolumes.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

You can use Veeam Agent commands to restore volumes from a backup or restore point:

- [Restore from backup](#)

When you restore a volume from the backup, Veeam Agent will automatically select the latest restore point in the backup. The volume will be restored to the state in which the volume was at the time when the latest restore point was created.

- [Restore from a restore point](#)

When you restore a volume from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

# Before You Begin

Before you begin the volume-level restore process, check the following prerequisites:

- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders and on Veeam backup repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

Volume-level restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Linux swap space is hosted.
- You cannot restore a volume to the volume where the backup file that you use for restore is located.

To overcome the first two limitations, you can boot from the recovery image and use the Veeam Recovery Media tools for volume-level restore. To learn more, see [Restoring from Veeam Recovery Media](#).

# Restoring from Backup

With Veeam Agent command line interface, you can restore volumes from the backup. When you restore a volume from the backup, Veeam Agent automatically selects the latest restore point in the backup and restores the volume to the state in which the volume was at the time when the latest restore point was created.

# Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

## NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the <b>Repository</b> column. For information about the import procedure, see <a href="#">Importing Backups</a> .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

## Step 2. Explore Backup Content

To view detailed information about specific backup, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup\_id> – ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Name	Name of the volume in the backup.
Device	Path to the block device that represents the volume.
FS UUID	File system ID.
Offset	Position of the volume on the computer disk.
Size	Size of the volume in the backup.

For example:

```
user@srv01:~$ veeamconfig backup show --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Machine name: srv01 HomePartitionBackup
  Name:      [sda6]
  Device:    [/dev/sda6]
  FS UUID:   [4967f2eb-e8bb-48fe-a694-5ba67b9030a5]
  Offset:    [11813257216] bytes (23072768 sectors)
  Size:      [41872785408] bytes (81782784 sectors)
```



## Step 3. Start Restore Process

To start the process of volume-level restore from the backup, use the following command:

```
veeamconfig backup restore --id <backup_id> --targetdev <target_volume> --backupdev <volume_in_backup>
```

where:

- `<backup_id>` – ID of the backup.
- `<target_volume>` – path to a block device that represents a volume on your computer that you want to recover.
- `<volume_in_backup>` – path to a block device that represents a volume in the backup.

This parameter is optional. If you do not specify this parameter, Veeam Agent will restore from the backup a volume that has the same name as a `<target_volume>`.

For example:

```
user@srv01:~$ veeamconfig backup restore --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b --targetdev /dev/sdb --backupdev /dev/sda6
Restoring backup.
Backup: 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Devices:
  Device in current system: [/dev/sdb]   In backup: [/dev/sda6];
You are sure? (y/n)
y
Volume restore from backup has been started.
Session ID: [{0b72ef45-4c88-4639-b940-ad3828b1cd4e}].
Logs stored in: [/var/log/veeam/Restore/Session_{0b72ef45-4c88-4639-b940-ad3828b1cd4e}].
```

### IMPORTANT

You can restore a backed-up volume only to a target volume that is not used by your Linux OS (that does not have file system mount points). For example, you can add a new disk to your computer and restore a volume in the backup to this disk.

If you want to restore a volume to the location that is crucial for the OS running, you should boot from the Veeam Recovery Media and perform volume-level restore with the Volume Restore wizard. For example, this approach is helpful when you restore the root (/) partition.

Alternatively, if the volume is backed-up in the unmounted state, it can be restored without booting from the Veeam Recovery Media.

## Step 4. Monitor Restore Process

You can monitor the restore process by viewing the restore session log in the command line interface.

To view Veeam Agent for Linux session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 0b72ef45-4c88-4639-b940-ad3828b1cd4e
2023-11-27 11:04:04 UTC {b141f32a-3e77-45a6-b55a-c100a1464d67} [info] Job started at 2023-11-27 14:04:04
2023-11-27 11:04:04 UTC {9b60ac03-2de0-4fe2-a00e-bec556d98ee8} [info] Starting volume restore
2023-11-27 11:04:07 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [processing] sdb
2023-11-27 11:04:15 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 512.0 kB at 58.6kB/s (0%)
...
2023-11-27 11:14:35 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB at 10.6MB/s (97%)
2023-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB at 10.6MB/s (100%)
2023-11-27 11:14:37 UTC {00add723-cbfa-4cc8-b299-d2349a051d6f} [warn] /dev/sdb has a duplicate filesystem UUID
2023-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb restored 6.5GB at 10.6MB/s
2023-11-27 11:14:37 UTC {8b8742a2-1c80-4e14-bbf1-45a3612bc3a7} [info] Volume restore completed
```

### TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

# Restoring from Restore Point

With Veeam Agent command line interface, you can restore volumes from the specific restore point. When you restore a volume from the restore point, you can select the necessary restore point in the backup to recover data to a desired point in time.

# Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

## NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the <b>Repository</b> column. For information about the import procedure, see <a href="#">Importing Backups</a> .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

## Step 2. Explore Restore Points

To view information about restore points in the backup, use the following command:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where

<backup\_id> – ID of the backup for which you want to view information on restore points.

You can view the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: <ul style="list-style-type: none"><li>• Full</li><li>• Increment.</li></ul>
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul>
Retention	Displays information about enabled long-term retention per each type: weekly (W), monthly (M) and yearly (Y).

For example:

```
user@srv01:~$ veeamconfig backup info --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Job name      OIB ID      Type      C
reated at    Is corrupt  Retention
srv01 HomePartitionBackup {23cb927d-5e2d-42fe-a4a4-e5f254a6413e} Full      2
023-11-15 11:28 false      WM
srv01 HomePartitionBackup {25e31075-4c30-4d67-86a6-293c0887f4eb} Increment 2
023-11-15 11:58 false
srv01 HomePartitionBackup {9375140d-720a-4d3e-a69b-ab9cf60d53fa} Increment 2
023-11-27 13:15 false
```

or

```
user@srv01:~$ veeamconfig point list --backupid 4f75bb20-a6b6-4323-9287-1c6c8ce
ccb6b
Job name          OIB ID          Type          C
reated at        Is corrupt      Retention
srv01 HomePartitionBackup {23cb927d-5e2d-42fe-a4a4-e5f254a6413e} Full          2
023-11-15 11:28 false      WM
srv01 HomePartitionBackup {25e31075-4c30-4d67-86a6-293c0887f4eb} Increment     2
023-11-15 11:58 false
srv01 HomePartitionBackup {9375140d-720a-4d3e-a69b-ab9cf60d53fa} Increment     2
023-11-27 13:15 false
```

## Step 3. Start Restore Process

To start the process of volume-level restore from the specific restore point, use the following command:

```
veeamconfig point restore --id <point_id> --targetdev <target_volume> --backupdev <volume_in_backup>
```

where:

- `<point_id>` – ID of the restore point.
- `<target_volume>` – path to a block device that represents a volume on your computer that you want to recover.
- `<volume_in_backup>` – path to a block device that represents a volume in the backup.

This parameter is optional. If you do not specify this parameter, Veeam Agent will restore from the backup a volume that has the same name as a `<target_volume>`.

For example:

```
user@srv01:~$ veeamconfig point restore --id 9375140d-720a-4d3e-a69b-ab9cf60d53fa --backupdev /dev/sda6 --targetdev /dev/sdb
Restoring point.
Restore point: 9375140d-720a-4d3e-a69b-ab9cf60d53fa
Devices:
  Device in current system: [/dev/sdb]   In backup: [/dev/sda6];
You are sure? (y/n)
y
Volume restore by point has been started.
Session ID: [{697d9348-9001-4845-8764-3cc4fb3f296b}].
Logs stored in: [/var/log/veeam/Restore/Session_{697d9348-9001-4845-8764-3cc4fb3f296b}].
```

### IMPORTANT

You can restore a backed-up volume only to a target volume that is not used by your Linux OS (that does not have file system mount points). For example, you can add a new disk to your computer and restore a volume in the backup to this disk.

If you want to restore a volume to the location that is crucial for the OS running, you should boot from the Veeam Recovery Media and perform volume-level restore with the Volume Restore wizard. For example, this approach is helpful when you restore the root (/) partition.

Alternatively, if the volume is backed-up in the unmounted state, it can be restored without booting from the Veeam Recovery Media.



## Step 4. Monitor Restore Process

You can monitor the restore process by viewing the restore session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 697d9348-9001-4845-8764-3cc4fb3f296b
2023-11-27 10:35:47 UTC {b9604775-d265-4537-b98e-848fd77c7375} [info] Job started at 2023-11-27 13:35:47
2023-11-27 10:35:47 UTC {ed66a1f6-5216-4596-a7b5-be10dd10c32f} [info] Starting volume restore
2023-11-27 10:35:50 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [processing] sdb
2023-11-27 10:35:59 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb 512.0 kB at 59.1kB/s (0%)
...
2023-11-27 10:46:27 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb 6.5GB at 10.5MB/s (100%)
2023-11-27 10:46:28 UTC {dae118c8-eb7c-4e14-9832-f0bfd089b329} [warn] /dev/sdb has a duplicate filesystem UUID
2023-11-27 10:46:28 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb restored 6.5GB at 10.5MB/s
2023-11-27 10:46:28 UTC {a21a89d9-d0ca-4f5c-8399-28ae599f2f1c} [info] Volume restore completed
```

### TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

# Restoring Files and Folders with Recovery Wizard

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent publishes the backup content directly into the computer file system. You can browse to files and folders in the backup, restore files and folders to their initial location, copy files and folders to a new location or simply target applications to restored files and work with them as usual.

# Before You Begin

Before you begin the file-level restore process, check the following prerequisites:


- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- [For backups of BTRFS file system] A machine on which you perform file-level restore must run the same or later Linux kernel version as the machine on which the backup was created.

For example, you created a backup of a machine that runs Linux kernel version 4.14. If you perform file-level restore from this backup on another machine that runs Linux kernel 2.6, the file-level restore process will fail.

# Step 1. Launch File Level Restore Wizard

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

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the [R] key to proceed to the File Level Restore wizard.



## Step 2. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.

The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays available backups and information about each backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:


1. In the **Imported backups** pane, select with [Up] and [Down] keys the backup from which you want to recover data and press [Enter].

In the list of backups, Veeam Agent displays backups that were created by backup jobs configured with Veeam Agent on your computer. If Veeam Agent for Linux is connected to a Veeam Backup & Replication server or a Veeam Cloud Connect service provider, backups created in the Veeam backup repository or cloud repository also appear in the list.

By default, Veeam Agent displays in the list only those backups in the Veeam backup repository that were created under your account. If you used an account to which the Veeam Backup Administrator role is assigned to connect to the Veeam backup server, you can also view all Veeam Agent backups that are stored in the Veeam backup repository to which Veeam Agent is connected. To view such backups, click the **Show all** link at the bottom of the list.

If Veeam Agent fails to display backups stored in the Veeam backup repository for some reason, you can press the [R] key to rescan the backup repository. Veeam Agent will try to reconnect to the Veeam backup server and refresh the list of backups.

If you want to recover data from a backup that is stored in another location, for example, a backup created with another instance of Veeam Agent in a network shared folder, you can import such backup. Press the [I] key, browse to the directory in which the backup file resides and select the necessary backup file. The selected backup file will be added to the list of backups.



2. In the **Restore points** pane, select with [Up] and [Down] keys the restore point from which you want to recover data and press [Enter].

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).

```

Veeam Agent for Linux [ srv02 ]

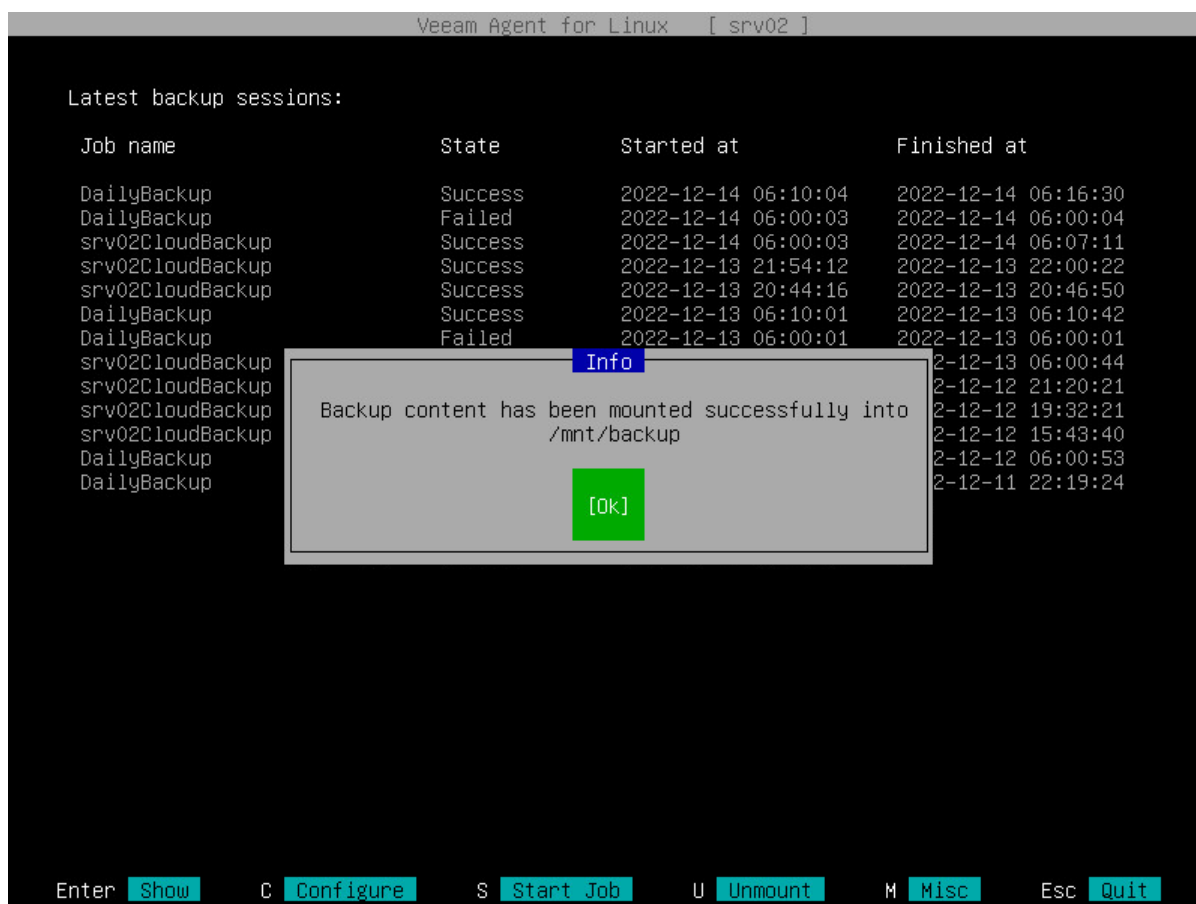
IMPORTED BACKUPS
Job name                               Hostname                               Points
srv02 DailyBackup                     srv02                                  4
srv02 srv02CloudBackup                 srv02                                  5
srv02 srv02CloudBackup(encrypted)srv02 2
Show all...

RESTORE POINTS
Created at
06:10 14-12-2022
06:10 13-12-2022
06:00 12-12-2022
22:14 11-12-2022

I Import backup Enter Next Esc Main menu

```

3. Veeam Agent will mount the content of the backup file to the `/mnt/backup` directory in the computer's file system and display a notification window with the corresponding message. Press [Enter] to close the window and return to the Veeam Agent control panel.



#### TIP

When you finish working with restored files and folders, you can unmount the backup from the `/mnt/backup` folder. To learn more, see [Stop Backup Mount Session](#).



## Step 3. Save Restored Files

When the backup file content is mounted to the `/mnt/backup` directory in the computer's file system, you can use Linux command line utilities or preferred file browser to work with restored files and directories. You can browse for files and directories in the mounted backup and copy files and directories that you want to restore to their initial location or to a new location.

### NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:

- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

In the following example, the restored file `Report1.pdf` is copied from the mounted backup to the new location with Linux command line utilities:


```
user@srv01:~$ ls Documents/
Reports
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/
Report1.pdf Report2.pdf
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Report1.pdf Documents/
user@srv01:~$ ls Documents/
Report1.pdf Reports
```

## Step 4. Stop Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. To unmount a backup, you need to stop the backup mount session. This may be required, for example, if you want to stop working with files and folders in one backup and mount another backup for file-level restore. You can also stop the backup mount session to unmount a backup after you have finished working with restored files and folders.

To stop the backup mount session, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command:
2. In the Veeam Agent control panel, press the [U] key to unmount a backup.
3. Veeam Agent will stop the backup mount session and display a notification window. Press [Enter] to close the window and return to the Veeam Agent control panel.



# Restoring Files and Folders with Command Line Interface

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent publishes the backup content directly into the computer file system. You can browse to files and folders in the backup, restore files and folders to their initial location, copy files and folders to a new location or simply target applications to restored files and work with them as usual.

With the Veeam Agent command line interface, you can restore files and folders in a more flexible way than with the use of the File Level Restore wizard. In particular, you can specify a directory in which Veeam Agent should mount the backup file content for file-level restore. You can also mount several backups to different directories to work with files and folders restored from different backups simultaneously.

You can use Veeam Agent commands to restore files and folders from backup or from specific restore point:

- [Restore from backup](#)

When you restore files and folders from the backup, Veeam Agent will automatically select the latest restore point in the backup. You can restore files and folders to the state in which they were at the time when the latest restore point was created.

- [Restore from a restore point](#)

When you restore files and folders from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

# Before You Begin

Before you begin the file-level restore process, check the following prerequisites:

- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- [For backups of BTRFS file system] A machine on which you perform file-level restore must run the same or later Linux kernel version as the machine on which the backup was created.

For example, you created a backup of a machine that runs Linux kernel version 4.14. If you perform file-level restore from this backup on another machine that runs Linux kernel 2.6, the file-level restore process will fail.

# Restoring from Backup

With Veeam Agent command line interface, you can restore files and folders from the backup. When you perform file-level restore from the backup, Veeam Agent for Linux automatically selects the latest restore point in the backup. You can restore files and folders to the state in which they were at the time when the latest restore point was created.

# Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

## NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the <b>Repository</b> column. For information about the import procedure, see <a href="#">Importing Backups</a> .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

## Step 2. Explore Backup Content

For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

To view detailed information about specific backup, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup\_id> – ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
<b>Machine name</b>	Host name of the machine on which the backup job is configured and the name of the job.
<b>Name</b>	Name of the volume in the backup.
<b>Device</b>	Path to the block device that represents the volume.
<b>FS UUID</b>	File system ID.
<b>Offset</b>	Position of the volume on the computer disk.
<b>Size</b>	Size of the volume in the backup.

For a file-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
<b>Machine name</b>	Host name of the machine on which the backup job is configured and the name of the job.
<b>Backed up</b>	Backup scope for the file-level backup job.



For example:

```
user@srv01:~$ veeamconfig backup show --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
Machine name: srv01 DocumentsBackup
File-level backup
Backed up:
/home/user/Documents
```

## Step 3. Mount Backup

To mount a backup for file-level restore, use the following command:

```
veeamconfig backup mount --id <backup_id> --mountdir <path>
```

where:

- `<backup_id>` – ID of the backup that you want to mount to the computer file system for file-level restore.
- `<path>` – path to the directory to which you want to mount the backup file content.

For example:

```
user@srv01:~$ veeamconfig backup mount --id ea64a7e5-038a-4c86-970a-6d59d4cf3968 --mountdir /mnt/backup
Backup is mounted.
Session ID: [{2a313184-32d0-4d3a-a1b0-2eebac986047}].
Logs stored in: [/var/log/veeam/Mount/Session_{2a313184-32d0-4d3a-a1b0-2eebac986047}].
```

## Step 4. Monitor Mount Process and Result

You can monitor the backup mount process by viewing the mount session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the backup mount session.

For example:

```
user@srv01:~$ veeamconfig session log --id 2a313184-32d0-4d3a-a1b0-2eebac986047
2023-11-22 17:30:34 UTC {30878c82-27d0-45dc-ab21-6f27d5082fd4} [info] Job started at 2023-11-22 20:30:34
2023-11-22 17:30:34 UTC {714b21d0-0d20-486e-b1e5-22d5fb5a8ee9} [info] Mounting restore point
2023-11-22 17:30:35 UTC {d331f038-5b7c-4549-85cf-5e1b54dbaf71} [info] Restore point has been mounted
```

To ensure that the backup is successfully mounted, you can browse to the directory that you specified in the `veeamconfig backup mount` command. For example:

```
user@srv01:~$ ls /mnt/backup/
FileLevelBackup_0
```

### TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

## Step 5. Save Restored Files

When the backup file content is mounted to the computer file system, you can use Linux command line utilities or preferred file browser to work with restored files and folders. You can browse for files and folders in the mounted backup and copy files and folders that you want to restore to their initial location or to a new location.

### NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:

- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

In the following example, the restored file `Report1.pdf` is copied from the mounted backup to a new location with the Linux command line utilities:

```
user@srv01:~$ ls Documents/  
Reports  
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/  
Report1.pdf  Report2.pdf  
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Report1.pdf Documents/  
user@srv01:~$ ls Documents/  
Report1.pdf  Reports
```

## Step 6. Stop Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. After you have finished working with restored files and folders, you can stop the backup mount session to unmount the backup.

To stop the backup mount session, use the following command:

```
veeamconfig session stop --id <session_id>
```

where:

<session\_id> – ID of the backup mount session that you want to stop.

Veeam Agent will stop the mount session and unmount the backup from the computer file system. For example:

```
user@srv01:~$ veeamconfig session stop --id 2a313184-32d0-4d3a-a1b0-2eebac98604
7
Session has stopped.
user@srv01:~$ ls /mnt
user@srv01:~$
```

# Restoring from Restore Point

With Veeam Agent command line interface, you can restore files and folders from the specific restore point. When you restore files and folders from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

## Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

### NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
<b>Job name</b>	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
<b>Backup ID</b>	ID of the backup.
<b>Repository</b>	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the <b>Repository</b> column. For information about the import procedure, see <a href="#">Importing Backups</a> .
<b>Created at</b>	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```



## Step 2. Explore Restore Points

To view information about restore points in the backup, use the following command:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where:

<backup\_id> – ID of the backup for which you want to view information on restore points.

You can view the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: <ul style="list-style-type: none"><li>• Full</li><li>• Increment.</li></ul>
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul>
Retention	Displays information about enabled long-term retention per each type: weekly (W), monthly (M) and yearly (Y).

For example:

```
user@srv01:~$ veeamconfig backup info --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
Job name      OIB ID      Type      Creat
ed at      Is corrupt  Retention
srv01 DocumentsBackup {0f3c9f3e-3985-4dc9-8cd6-979dba810c2f} Full      2023-
11-11 18:31 false      M
srv01 DocumentsBackup {ff0c6969-8b9b-4865-b4f9-d686faf41d50} Increment 2023-
11-14 13:35 false
srv01 DocumentsBackup {a9e420df-d749-4b9a-b675-19d8e94c3bf1} Increment 2023-
11-15 13:43 false
```

or

```
user@srv01:~$ veeamconfig point list --backupid ea64a7e5-038a-4c86-970a-6d59d4c
f3968
Job name          OIB ID          Type          Creat
ed at          Is corrupt Retention
srv01 DocumentsBackup {0f3c9f3e-3985-4dc9-8cd6-979dba810c2f} Full          2023-
11-11 18:31 false      M
srv01 DocumentsBackup {ff0c6969-8b9b-4865-b4f9-d686faf41d50} Increment      2023-
11-14 13:35 false
srv01 DocumentsBackup {a9e420df-d749-4b9a-b675-19d8e94c3bf1} Increment      2023-
11-15 13:43 false
```

## Step 3. Mount Restore Point

To mount a backup for file-level restore, use the following command:

```
veeamconfig point mount --id <point_id> --mountdir <path>
```

where:

- `<point_id>` – ID of the restore point that you want to mount to the computer file system for file-level restore.
- `<path>` – path to the directory to which you want to mount the backup file content.

For example:

```
user@srv01:~$ veeamconfig point mount --id b127e64e-1f1c-4e0b-bb36-b087761267b3
--mountdir /mnt/backup
Restore point is mounted.
Session ID: [{4d69dd85-ac60-4cff-883d-50f25f49a9c8}].
Logs stored in: [/var/log/veeam/Mount/Session_{4d69dd85-ac60-4cff-883d-50f25f49a9c8}].
```

## Step 4. Monitor Mount Process and Result

You can monitor the restore point mount process by viewing the mount session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the restore point mount session.

For example:

```
user@srv01:~$ veeamconfig session log --id 4d69dd85-ac60-4cff-883d-50f25f49a9c8
2023-11-23 12:44:55 UTC {9c5c8ece-cb88-4742-bb90-1f8ff79b4bdc} [info] Job started at 2023-11-23 15:44:55
2023-11-23 12:44:55 UTC {4ac10045-a74b-4a41-9c5e-53521cba1045} [info] Mounting restore point
2023-11-23 12:44:56 UTC {540a61f7-5d5c-47d5-a2b8-51daa694d5ec} [info] Restore point has been mounted
```

To ensure that the restore point is successfully mounted, you can browse to the directory that you specified in the `veeamconfig point mount` command. For example:

```
user@srv01:~$ ls /mnt/backup/
FileLevelBackup_0
```

### TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

## Step 5. Save Restored Files

When the restore point is mounted to the computer file system, you can use Linux command line utilities or preferred file browser to work with restored files and folders. You can browse for files and folders in the mounted backup and copy files and folders that you want to restore to their initial location or to a new location.

### NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:

- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

In the following example, the restored file `Report1.pdf` is copied from the mounted restore point to a new location with the Linux command line utilities:

```
user@srv01:~$ ls Documents/  
Reports  
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/  
Report1.pdf  Report2.pdf  
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Repor  
t1.pdf Documents/  
user@srv01:~$ ls Documents/  
Report1.pdf  Reports
```

## Step 6. Stop Backup Mount Session

When Veeam Agent mounts a restore point for file-level restore, Veeam Agent starts a new restore point mount session. After you have finished working with restored files and folders, you can stop the mount session to unmount the restore point.

To stop the restore point mount session, use the following command:

```
veeamconfig session stop --id <session_id>
```

where:

<session\_id> – ID of the restore point mount session that you want to stop.

Veeam Agent will stop the mount session and unmount the restore point from the computer file system. For example:

```
user@srv01:~$ veeamconfig session stop --id 4d69dd85-ac60-4cff-883d-50f25f49a9c
8
Session has stopped.
user@srv01:~$ ls /mnt
user@srv01:~$
```

# Exporting Backup to Virtual Disk

You can export a backup to a virtual disk in the VHD format. You can then attach the created VHD disk to a virtual machine to recover your computer in a virtual environment.

- [Exporting Backups](#)
- [Exporting Restore Points](#)

# Exporting Backups

You can export the backup file to a virtual disk in the VHD format. When you export a backup, you export to a virtual disk data pertaining to the latest restore point in the backup. The created VHD disk will reflect the state in which backed-up volumes were at the time when the latest restore point was created.

To export backup to a VHD disk:

1. Start the export process with the following command:

```
veeamconfig backup export --id <backup_id> --outdir <path>
```

where:

- o <backup\_id> – ID of the backup that you want to export to a virtual disk.
- o <path> – full path to a directory in which you want to save the created virtual disk. Specifying relative paths is not supported.

For example:

```
user@srv01:~$ veeamconfig backup export --id 45f074d2-d2d9-423d-84e9-8f179
8b08d4c --outdir /home/user/disk
Export has been started.
Session ID: [{5f001367-8937-46e0-a756-449bf9f1a182}].
Logs stored in: [/var/log/veeam/Export/Session_{5f001367-8937-46e0-a756-44
9bf9f1a182}].
```

2. You can monitor the export process and result by viewing the export session log with the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the export session.

For example:

```
user@srv01:~$ veeamconfig session log --id 5f001367-8937-46e0-a756-449bf9f
1a182
2023-11-27 11:20:56 UTC {b54af37c-35a6-4807-80d2-0f070f024e69} [info] Job
started at 2023-11-27 14:20:56
2023-11-27 11:20:56 UTC {48d699d2-86cf-4a32-b9c8-ab51b8325f3c} [info] Expo
rting virtual disks content
2023-11-27 11:20:57 UTC {0e2e7d97-f067-4823-8dde-084c401eb62b} [processing
] Restoring device: [30460cb5].
2023-11-27 11:22:59 UTC {0e2e7d97-f067-4823-8dde-084c401eb62b} [info] Devi
ce [30460cb5] has been exported
2023-11-27 11:23:00 UTC {36f0d0c5-2af7-48d8-abc2-c8ef9aaffc54} [info] Virt
ual disks content has been exported
```

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).



3. Exported backup will be saved as a virtual disk file in the specified directory. You can check this with a file browser or with the following command:

```
ls <path>
```

where:

<path> – path to the directory in which the virtual disk with the backup is saved.

For example:

```
user@srv01:~$ ls disk/  
dev_30460cb5.vhd
```

# Exporting Restore Points

You can export the specific restore point to a virtual disk in VHD format. When you export a restore point, you select the necessary restore point in the backup to recover data to a desired point in time. The created VHD disk will reflect the state in which backed-up volumes were at the time when the selected restore point was created.

To export restore point to a VHD disk:

1. Start the export process with the following command:

```
veeamconfig point export --id <point_id> --outdir <path>
```

where:

- o <point\_id> – ID of the restore point that you want to export to a virtual disk.
- o <path> – full path to a directory in which you want to save the created virtual disk. Specifying relative paths is not supported.

For example:

```
user@srv01:~$ veeamconfig point export --id b319ealf-59a2-41ea-9ca3-b668e86ac941 --outdir /home/user/veeam/  
Export has been started.  
Session ID: [{aeb9c549-a660-4a0e-b89c-cb076b8bfa85}].  
Logs stored in: [/var/log/veeam/Export/Session_{aeb9c549-a660-4a0e-b89c-cb076b8bfa85}].
```

2. You can monitor the export process and result by viewing the export session log with the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the export session.

For example:

```
user@srv01:~$ veeamconfig session log --id aeb9c549-a660-4a0e-b89c-cb076b8bfa85  
2023-05-05 11:15:21 UTC {b950503d-55c9-435f-946e-1078184f5a86} [info] Job started at 2023-05-05 14:15:21  
2023-05-05 11:15:21 UTC {32d56391-9002-431e-ae6b-2285537a67e5} [info] Exporting virtual disks content  
2023-05-05 11:15:22 UTC {ba3dabe0-0556-430c-9671-9448a6dc4bcb} [processing] Restoring device: [30460cb5].  
2023-05-05 11:17:26 UTC {ba3dabe0-0556-430c-9671-9448a6dc4bcb} [info] Device [30460cb5] has been exported  
2023-05-05 11:17:26 UTC {9e945c29-900e-4a07-9e3b-ccf7f156807d} [info] Virtual disks content has been exported
```

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

3. Exported backup will be saved as a virtual disk file in the specified directory. You can check this with a file browser or with the following command:

```
ls <path>
```

where

<path> – path to the directory in which the virtual disk with the backup is saved.

For example:

```
user@srv01:~$ ls /home/user/veeam/  
dev_30460cb5.vhd
```

# Restoring Data from Encrypted Backups

When you restore data from an encrypted backup, Veeam Agent performs data decryption automatically in the background or requires you to specify a password.

- If encryption keys required to unlock the backup file are available in the Veeam Agent database, that is, if you encrypt and decrypt the backup file on the same Veeam Agent computer, you do not need to specify the password. Veeam Agent uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore from the encrypted backup does not differ from that from an unencrypted one.
- If encryption keys are not available in the Veeam Agent database, you need to provide a password to unlock the encrypted file. The password must be the same as the password that was used to encrypt the backup file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including restore points that were encrypted with an old password and restore points that were created before you have enabled the encryption option for the job.

The process of unlocking an encrypted backup file differs depending on what Veeam Agent user interface you use for data restore.


- [Veeam Agent graphical user interface](#)
- [Veeam Agent command line interface](#)

## Restoring Data from Encrypted Backups Using GUI

To restore data from an encrypted backup using the Veeam Agent graphical user interface:

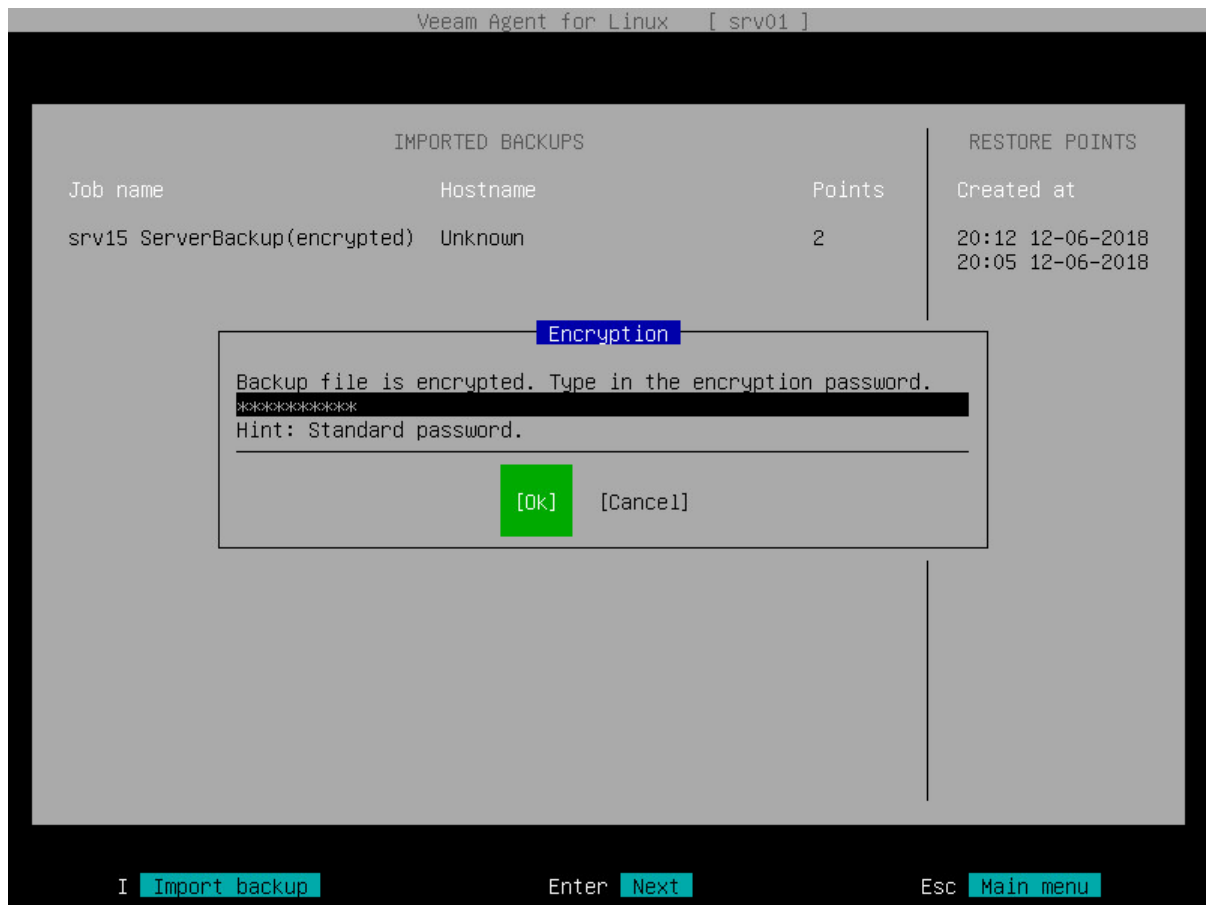
1. Launch the necessary data restore wizard:
  - If you want to perform file-level restore from an encrypted backup that was created on another Veeam Agent computer, launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command. To learn more, see [Restoring Files and Folders](#).
  - If you want to perform volume-level restore or file-level restore recovery from an encrypted backup, boot from the Veeam Recovery Media and launch the necessary data restore wizard. To learn more, see [Restoring from Veeam Recovery Media](#).
2. Follow the steps of the wizard to specify where the encrypted backup file that you plan to use for restore resides. If the backup file resides in a remote location, select the backup location type and specify settings to connect to the backup location.

3. Select the encrypted backup and restore point from which you want to restore data.



4. Veeam Agent will display the **Encryption** window. Enter the password for the backup file.  
In the **Hint** field of the **Encryption** window, Veeam Agent displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.  
If you changed the password one or several times while the backup chain was created, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

If you enter correct password, Veeam Agent will decrypt the backup metadata. You will be able to continue the restore operation in a regular manner.



## Restoring Data from Encrypted Backups Using Command Line Interface

To restore data from an encrypted backup using the Veeam Agent command line interface, complete the following steps:

1. Import the encrypted backup file to the Veeam Agent database. To learn more, see [Importing Encrypted Backups](#).
2. Perform the necessary restore operation in a regular manner. To learn more, see [Restoring Volumes with Command Line Interface](#) and [Restoring Files and Folders with Command Line Interface](#).

# Reporting

Veeam Agent for Linux provides several ways to get information about performed operations:

- With the Veeam Agent control panel
- With the Veeam Agent command line interface

For every data transfer operation, for example data backup and restore, backup import and export, Veeam Agent starts a new session. You can monitor performance of sessions started by Veeam Agent in the following ways:

- [Monitor backup job session progress with the control panel.](#)
- [View real-time backup job session statistics with the control panel.](#)
- [View backup job sessions results with the control panel.](#)
- [View the session status using the command line interface.](#)
- [View session logs.](#)

# Viewing Job Session Progress

You can monitor the backup job session progress in the list of sessions in the Veeam Agent control panel. For the currently running backup job session, Veeam Agent shows session status and percentage of session completion in the **State** column of the list of sessions.

To view backup job session progress, do the following:

1. If you have started the backup job from the command line, launch the Veeam Agent control panel with the `veeam` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, monitor progress of the currently running session.

If you have started the backup job from the Veeam Agent control panel, Veeam Agent will immediately display the list of backup job sessions with the currently running session.

## TIP

You can stop the backup job session at any time. To stop the backup job session, press the [S] key.

```
Veeam Agent for Linux [ srv01 ]

Latest backup sessions:

Job name          State          Started at      Finished at
-----
DailyBackup       Running (93%)  2017-12-10 20:23:25  ---

Enter Show  C Configure  S Start Job  R Recover Files  M Misc  Esc Quit
```



# Viewing Real-Time Job Session Statistics

You can view real-time statistics for a job session in the Veeam Agent control panel. Veeam Agent shows detailed data for every backup job session: job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data and details of the session performance, for example, warnings and errors that have occurred in the process of operation.

To view detailed information on the currently running backup job session, do the following:

1. If you have started the backup job from the command line, launch the Veeam Agent control panel with the `veeam` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, select the currently running session with the [Up] and [Down] keys and press [Enter].

If you have started the backup job from the Veeam Agent control panel, the current session will be already selected in the list of backup job sessions.

## TIP

You can stop the backup job session at any time. To stop the backup job session, press the [S] key.

## Statistics Counters

Veeam Agent for Linux displays jobs statistics for the following counters:

- The pane at the top of the control panel shows information on the job session type, percentage of the job completion and session status. If Veeam Agent operates in the Server edition and you have created more than one backup job, the job name also appears on the pane.
- The **Summary** box shows general information about the job:
  - **Duration** – time from the job start till the job end.
  - **Processing rate** – average speed of data processing. This counter is a ratio between the amount of processed data (**Processed** counter) and job duration (**Duration** counter).
  - **Bottleneck** – bottleneck in the data transmission process.
- The **Data** box shows information about processed data:
  - **Processed** – total size of all volumes processed by the job.
  - **Read** – amount of data read from the backed-up volume by Veeam Agent for Linux prior to applying compression. For incremental job runs, the value of this counter is typically lower than the value of the **Processed** counter. Veeam Agent reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target location.
  - **Transferred** – amount of data transferred from the backed-up volume to the backup location after applying compression. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Agent can perform additional activities with data, for example, decompress data prior to writing the file to disk. The activities can impact the size of the resulting file.
- The box in the center of the control panel shows a list of operations performed during the job session, their start time and duration time. To scroll the list of operations, use **Up** and **Down** arrow keys on the keyboard.

- The pane at the lower side of the control panel shows help information on how to navigate the control panel.

Veeam Agent for Linux [ srv01 ]

Backup [SystemBackup]
56%
Status: Running

**Summary**

Duration: 00:01:11

Processing rate: 42.1 MB/s

Bottleneck: Agent

**Data**

Processed: 2.8 GB (56%)

Read: 2.8 GB

Transferred: 1.3 GB (2.1x)

Time	Action	Duration
14:27:14	Job SystemBackup started at 2016-12-07 14:27:14 MSK	
14:27:14	Preparing to backup	
14:27:15	Creating volume snapshot	00:00:00
14:27:15	Starting full backup to Repository_1	
14:27:15	Backing up BIOS bootloader on /dev/sda	00:00:01
14:27:17	Backing up sda 2.8 GB at 43.3 MB/s (56%)	00:01:06

S Stop
Esc Back

# Viewing Job Session Result

You can view detailed statistics on every backup job session performed by Veeam Agent for Linux.

To view statistics for a specific job session:

1. Open the Veeam Agent control panel with one of the following commands:

```
veeam
```

or

```
veeamconfig ui
```

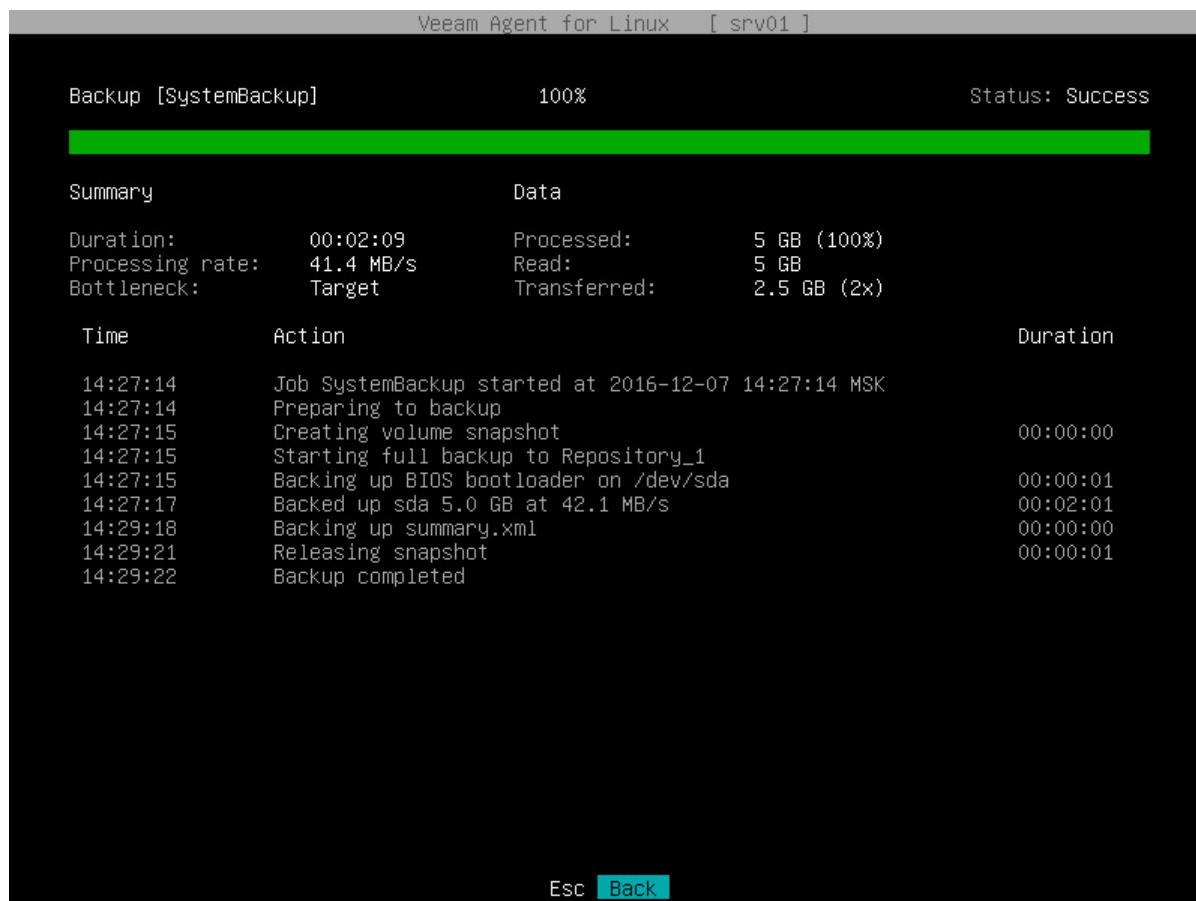
or

```
veeamconfig session ui
```

2. In the **Latest backup sessions** list, select the necessary backup job session with the [Up] and [Down] keys and press [Enter].

To return to the list of backup job sessions, press [Esc]. You can then select another backup job session or exit the Veeam Agent control panel in one of the following ways:

- with the [Esc] key – if you opened the control panel with the `veeam` or `veeamconfig ui` command.
- with the [Q] key – if you opened the control panel with the `veeamconfig session ui` command.



# Viewing Session Status

You can view status of every session that was started by Veeam Agent for Linux. To view the session status, use the following command:

```
veeamconfig session info --id <session_id>
```

where:

<session\_id> – ID of the session for which you want to check status.

Veeam Agent displays the following information about sessions:

Parameter	Description
ID	ID of the session.
Job name	Name of the backup job parent to the session. Veeam Agent displays value for this parameter only for backup job sessions.
Job ID	ID of the backup job parent to the session. Veeam Agent displays value for this parameter only for backup job sessions.
State	Current status of the session.
Start time	Date and time of the session start.
End time	Date and time of the session completion. Veeam Agent displays value for this parameter only for completed sessions.

The following example shows status information on the completed backup job session:

```
user@srv01:~$ veeamconfig session info --id 1592755d-3a2b-40a9-a036-5c81853b369e
Backup session
  ID: {1592755d-3a2b-40a9-a036-5c81853b369e}
  Job name: SystemBackup
  Job ID: {2495911e-58db-4452-b4d1-f53dcfbc600e}
  State: Success
  Start time: 2023-11-11 14:37:21 UTC
  End time: 2023-11-11 14:40:02 UTC
```

The following example shows status information on the running volume restore session:

```
user@srv01:~$ veeamconfig session info --id 697d9348-9001-4845-8764-3cc4fb3f296
b
Restore session
  ID: {697d9348-9001-4845-8764-3cc4fb3f296b}
  State: Running
  Start time: 2023-11-27 10:35:47 UTC
  End time:
```

# Viewing Session Logs

You can monitor the backup and restore process by viewing the backup job session and restore session logs in the Veeam Agent command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session\_id> – ID of the backup job or restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 0b72ef45-4c88-4639-b940-ad3828b1cd4e
2023-11-27 11:04:04 UTC [info] Job started at 2023-11-27 11:04:04
2023-11-27 11:04:04 UTC [info] Starting volume restore
2023-11-27 11:04:07 UTC [processing] sdb
2023-11-27 11:04:15 UTC [info] sdb 512.0kB at 58.6kB/s (0%)
2023-11-27 11:04:25 UTC [info] sdb 125.0MB at 6.7MB/s (0%)
2023-11-27 11:04:35 UTC [info] sdb 238.5MB at 8.3MB/s (1%)
...
2023-11-27 11:14:32 UTC [info] sdb 6.5GB at 10.7MB/s (92%)
2023-11-27 11:14:35 UTC [info] sdb 6.5GB at 10.6MB/s (97%)
2023-11-27 11:14:37 UTC [info] sdb 6.5GB at 10.6MB/s (100%)
2023-11-27 11:14:37 UTC [warn] /dev/sdb has a duplicate filesystem UUID
2023-11-27 11:14:37 UTC [info] sdb restored 6.5GB at 10.6MB/s
2023-11-27 11:14:37 UTC [info] Volume restore completed
```

# Exporting Product Logs

Veeam Agent offers a simple and convenient way to collect product logs and export them to an archive file. This operation may be required if you want to report an issue and need to attach log files to the support case.

When you export logs, Veeam Agent collects its log files and configuration files, exports them to an archive file in the `tar.gz` format and saves this archive file to a directory on the Veeam Agent computer.

You can perform the export logs operation in one of the following ways:

- [With the Veeam Agent control panel](#) — in this case, you can specify a directory to which Veeam Agent should save the log archive.
- [With the command line interface](#) — in this case, Veeam Agent will save the log archive to the current working directory.

## TIP

When you perform restore operations after booting from the Veeam Recovery Media, Veeam Agent also saves restore logs to the backup location. Restore logs are saved to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`. The archive is placed to the folder that contains the backup file from which you restored data.

If you encounter problems after restoring from the Veeam Recovery Media, it is recommended that you attach restore logs, as well as product logs collected by Veeam Agent, to the support case.




# Exporting Logs with Control Panel

You can use the Veeam Agent control panel to collect and export product logs. When you export logs with the control panel, you can choose where Veeam Agent should save the resulting log archive.

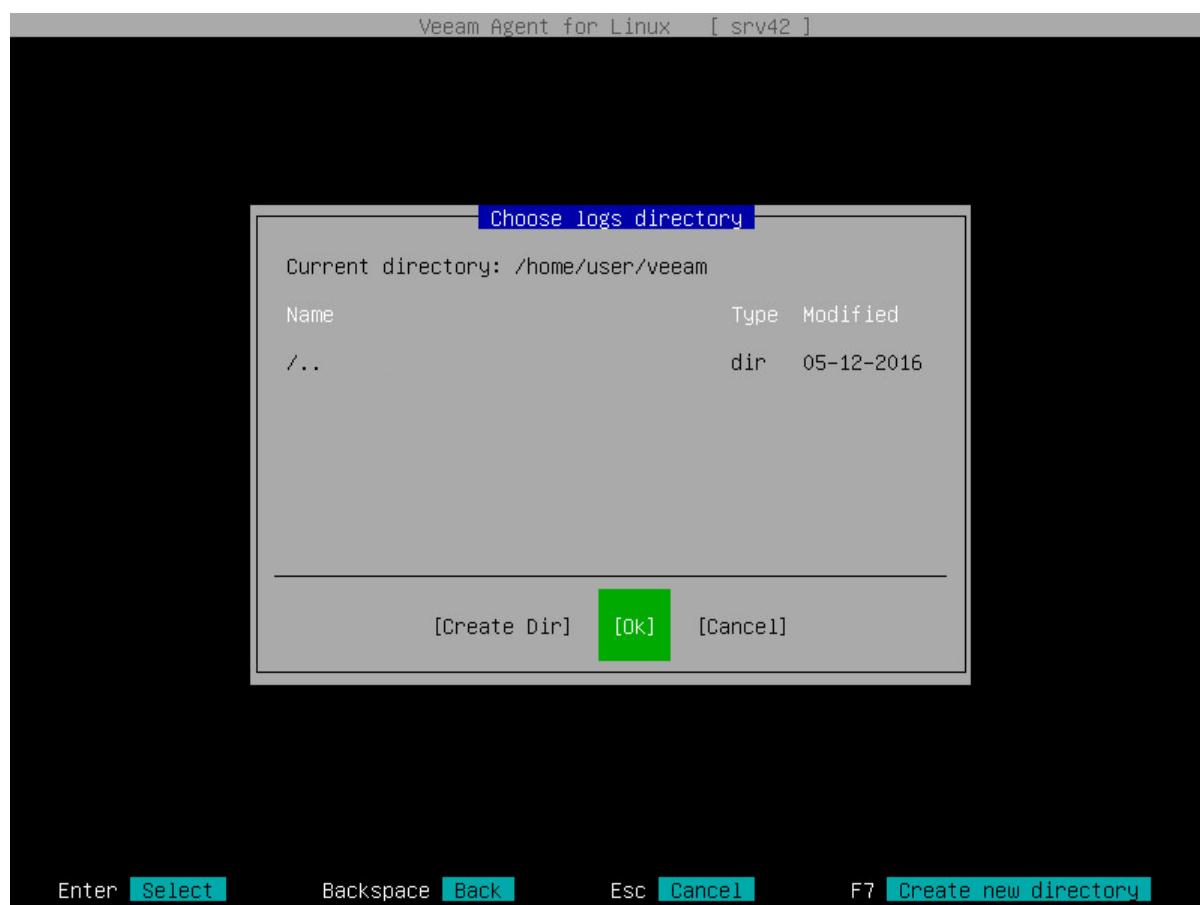
To export logs:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, select the **Export Logs** option and press [Enter].



4. In the **Choose logs directory** window, specify a directory to which you want to save the log archive:
  - a. In the **Choose logs directory** window, select the necessary directory and press [Enter].
  - b. Repeat the step 'a' until a path to the directory in which you want to save exported logs appears in the **Current directory** field.
  - c. To create a new directory, switch to the **Create Dir** button, press [Enter], then type a name for the new directory and press [Enter].

- d. Switch to the **Ok** button and press [Enter]. Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`, and save the archive to the specified directory.



# Exporting Logs with Command Line Interface

You can use the Veeam Agent command line interface to collect and export product logs. To export logs, use the following command:

```
veeamconfig grablogs
```

Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`, and save the archive to the current working directory.

For example:

```
user@srv01:~$ veeamconfig grablogs  
Logs have been exported successfully.
```

# Getting Support

If you have any questions or want to share your feedback about Veeam Agent, you can use one of the following options:

- You can search for the information on the necessary subject in the current Veeam Agent for Linux User Guide.
- You can visit [Veeam R&D Forums](#) and share your opinion or ask a question.
- If you use Veeam Agent with an active license installed, you can visit [Veeam Customer Support Portal](#) and submit a support case to the Veeam Customer Support Team.

# Using with Veeam Backup & Replication

If you have the Veeam backup infrastructure deployed in the production environment, you can use Veeam Agent together with Veeam Backup & Replication.

## IMPORTANT

If you plan to use Veeam Agent for Linux 6.2 with Veeam Backup & Replication, you must install Veeam Backup & Replication on the Veeam backup server.

For more information on managing connection to a Veeam backup server, see [Managing Veeam Backup & Replication Servers](#).

## NOTE

This and subsequent sections describe tasks with Veeam Backup & Replication available for Veeam Agent operating in the standalone mode. For information about tasks available in Veeam Backup & Replication within the Veeam Agent management scenario, see the [Veeam Agent Management Guide](#).

## Tasks with Veeam Backup & Replication

Veeam Backup & Replication lets you perform a number of additional data protection and disaster recovery tasks, as well as administrative actions with Veeam Agent backups. You can:

- [Grant access permissions on backup repositories](#).
- [Manage Veeam Agent licenses](#).

### *Data protection tasks*

- [Create Veeam Agent backups on backup repositories](#).
- [Create Veeam Agent backups on Veeam Cloud Connect repositories](#).
- [Copy Veeam Agent backups to secondary backup repositories](#).
- [Archive Veeam Agent backups to tape](#).

### *Restore tasks*

- [Restore Veeam Agent backups to Hyper-V VMs](#).
- [Restore Veeam Agent backups to VMware vSphere VMs](#).
- [Restore Veeam Agent backups to Nutanix VMs](#).
- [Restore files and folders from Veeam Agent backups](#).
- [Restore application items from Veeam Agent backups](#).
- [Restore disks from Veeam Agent backups](#).
- [Publish disks to analyze backup content](#).
- [Restore data from Veeam Agent backups to Amazon EC2](#).
- [Restore data from Veeam Agent backups to Microsoft Azure](#).

- [Restore data from Veeam Agent backups to Google Compute Engine.](#)
- [Export restore points of Veeam Agent backups to standalone full backup files.](#)

#### *Administrative tasks*

- [Import Veeam Agent backups.](#)
- [Enable and disable Veeam Agent backup jobs.](#)
- [Delete Veeam Agent backup jobs.](#)
- [View Veeam Agent backup properties.](#)
- [Create recovery tokens.](#)
- [Remove Veeam Agent backups.](#)
- [Delete Veeam Agent backups.](#)
- [Configure global settings.](#)
- [Assign roles to users.](#)

# Setting Up User Permissions on Backup Repositories

To be able to store backups in a backup repository managed by a Veeam backup server, the user must have access permissions on this backup repository.

## IMPORTANT

Veeam Agent for Linux does not support Veeam backup repositories with enabled KMS encryption. To learn more about KMS encryption for Veeam backup repositories, see the [Key Management System Keys](#) section in the Veeam Backup & Replication User Guide.

## NOTE

If you plan to create backups in a Veeam backup repository with Veeam Agent backup jobs configured in Veeam Backup & Replication, you do not need to grant access permissions on the backup repository to users. In the Veeam Agent management scenario, to establish a connection between the backup server and protected computers, Veeam Backup & Replication uses a TLS certificate. To learn more, see the [Configuring Security Settings](#) section in the Veeam Agent Management Guide.

Access permissions are granted to security principals such as users and AD groups by the backup administrator working with Veeam Backup & Replication. Users with granted access permissions can target Veeam Agent backup jobs at this backup repository and perform restore from backups located in this backup repository.

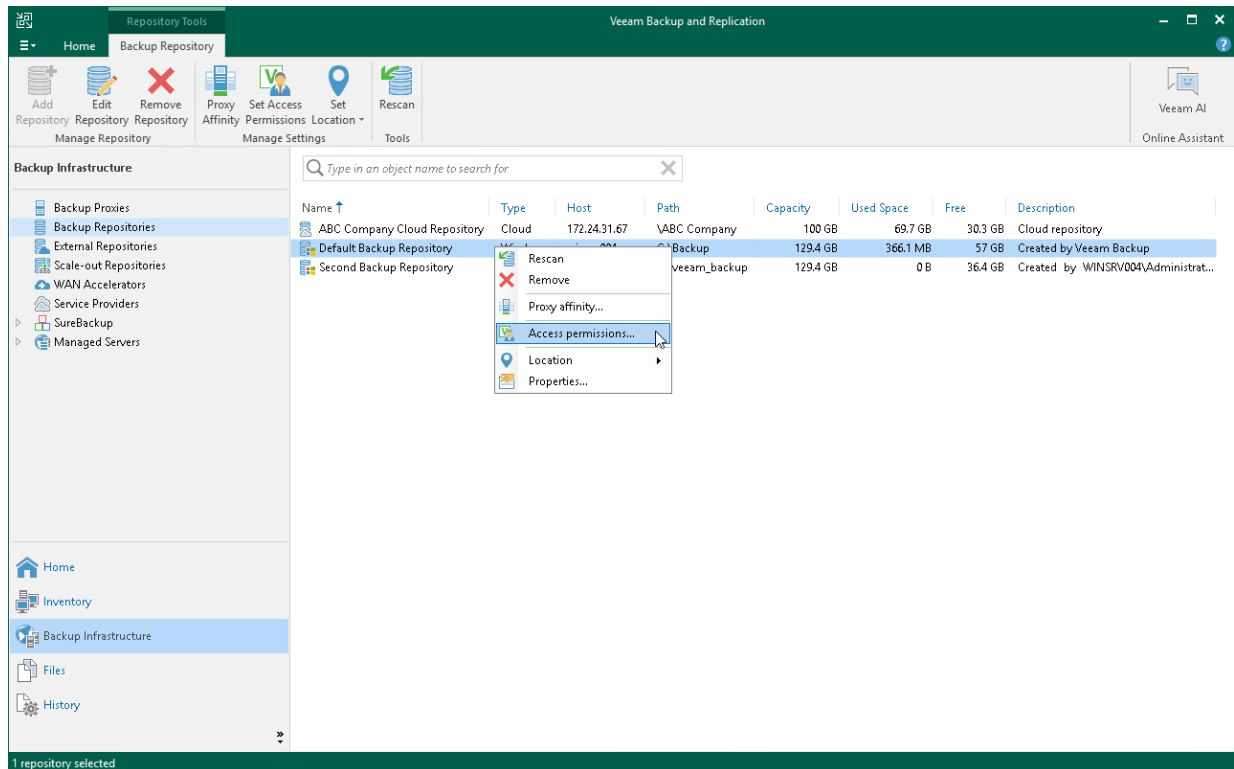
Right after installation, access permissions on the default backup repository are set to *Allow to everyone* for testing and evaluation purposes. If necessary, you can change these settings.

After you create a new backup repository, access permissions on this repository are set to *Deny to everyone*. To allow users to store backups in the backup repository, you must grant users with access permissions to this repository.

To grant access permissions to a security principal:

1. In Veeam Backup & Replication, open the **Backup Infrastructure** view.
2. In the inventory pane, click one of the following nodes:
  - The **Backup Repositories** node – if you want to grant access permissions on a regular backup repository to Veeam Agent users.
  - The **Scale-out Repositories** node – if you want to grant access permissions on a scale-out backup repository to Veeam Agent users.

3. In the working area, select the necessary backup repository and click **Set Access Permissions** on the ribbon, or right-click the backup repository and select **Access permissions**. If you do not see the **Set Access Permissions** button on the ribbon or the **Access permissions** command is not available in the shortcut menu, press and hold the [Ctrl] key, right-click the backup repository and select **Access permissions**.

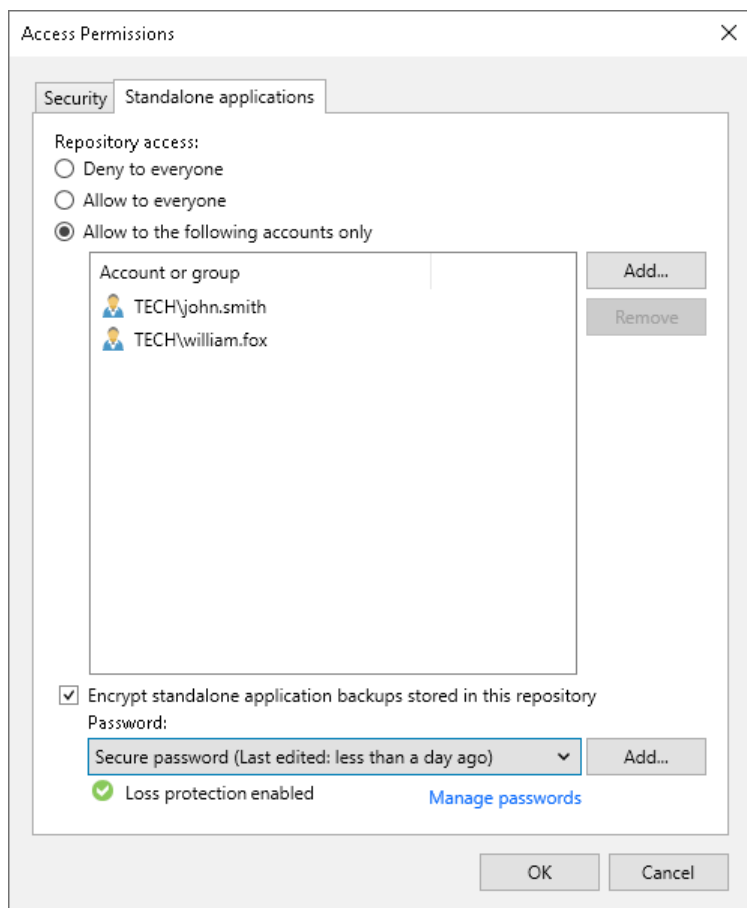


4. In the **Access Permissions** window, in the **Standalone applications** tab, specify to whom you want to grant access permissions on this backup repository:
  - **Allow to everyone** – select this option if you want all users to be able to store backups on this backup repository. Setting access permissions to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). However, we recommend this scenario for demo environments only.
  - **Allow to the following accounts or groups only** – select this option if you want only specific users to be able to store backups on this backup repository. Click **Add** to add the necessary users and groups to the list.
5. If you want to encrypt Veeam Agent backup files stored in the backup repository, select the **Encrypt backups stored in this repository** check box and choose the necessary password from the field below. If you have not specified a password beforehand, click **Add** on the right or the **Manage passwords** link to add a new password. Veeam Backup & Replication will encrypt files at the backup repository side using its built-in encryption mechanism. To learn more, see [Veeam Backup & Replication Documentation](#).



## IMPORTANT

If Veeam Agent is set up to use the backup cache, and the backup cache contains one or more restore points, Veeam Agent will automatically remove these restore points from the backup cache after you enable or disable the encryption option for the backup repository.



# Managing License

If you plan to use Veeam Agent with Veeam Backup & Replication, you must install a license in Veeam Backup & Replication or Veeam Backup Enterprise Manager. The license must have a total number of instances that is sufficient to protect machines (servers and workstations) on which you plan to install Veeam Agent. For more information, see [Veeam Licensing Policy](#).

After Veeam Agent connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming instances in the license. The product edition for Veeam Agent is selected depending on the type of the OS running on the Veeam Agent computer. You can switch to another commercial edition of Veeam Agent manually if needed. If you do not want Veeam Agents to consume instances, you can restrict instance consumption. For more information, see [Managing Instance Consumption by Veeam Agents](#).

The number of backup jobs configured in Veeam Agent does not impact instance consumption. For example, if 2 backup jobs are configured in Veeam Agent that operates in the Server edition, this Veeam Agent will consume instances required for 1 server.

Veeam Agent obtains information about the license from Veeam Backup & Replication and keeps it locally on the Veeam Agent computer. Information about the license is valid for 32 days. If Veeam Agent does not connect to Veeam Backup & Replication during this period, Veeam Backup & Replication will revoke its license.

## NOTE

In addition to managing Veeam Agent licenses, you can use the Veeam Backup & Replication console to manage Veeam Agent backup jobs and perform operations with backups created by these jobs.

If your backup server is connected to Veeam Backup Enterprise Manager, you can use Veeam Backup Enterprise Manager to manage licenses and perform restore tasks with Veeam Agent backups. You cannot manage Veeam Agent backup jobs with Veeam Backup Enterprise Manager.

For more information on Veeam Backup & Replication licensing, see the [Licensing](#) section in the Veeam Backup & Replication User Guide.

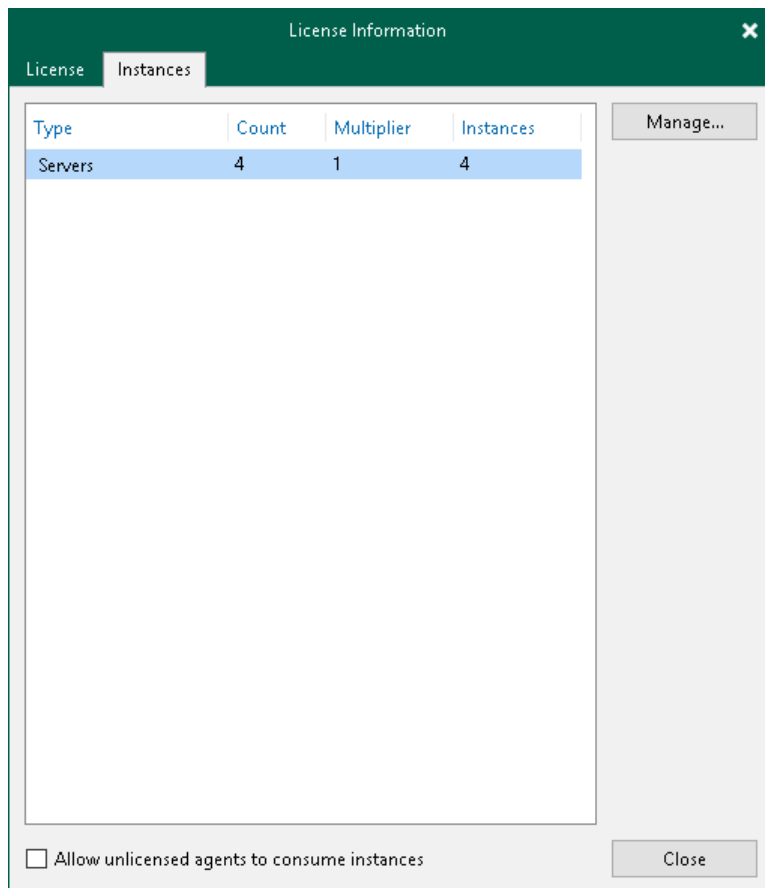
# Managing Instance Consumption by Veeam Agents

By default, Veeam Backup & Replication allows Veeam Agents to connect to the Veeam backup server and consume instances in the license. If you do not want Veeam Agents to consume instances, you can restrict instance consumption.

If you restrict instance consumption, Veeam Backup & Replication will switch all Veeam Agents connected to this Veeam backup server to the free edition that offers limited capabilities. For information about Veeam Agent editions, see [Product Editions](#).

To restrict instance consumption by Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, click the **Instances** tab.
3. On the **Instances** tab, clear the **Allow unlicensed agents to consume instances** check box.
4. Click **Close**.



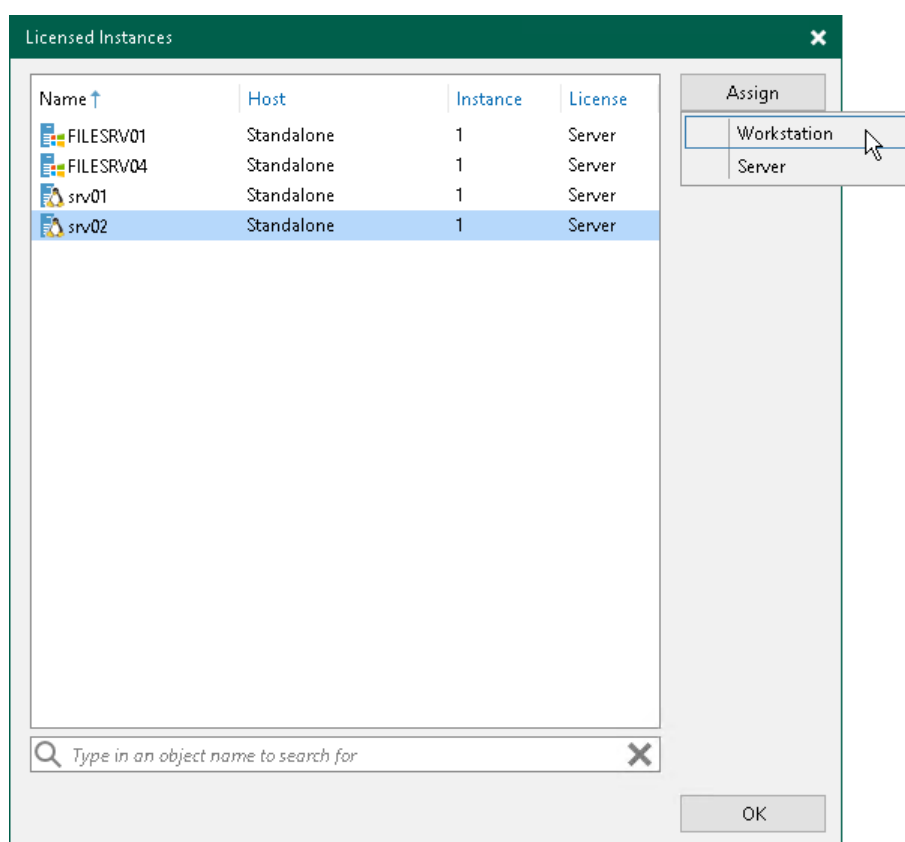
# Assigning License to Veeam Agent

After Veeam Agent connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming the license. The product edition for Veeam Agent is selected depending on the type of the OS running on the protected computer.

You can also assign a license to Veeam Agent manually if needed. When you assign a license, you can select the product edition, too.

To assign a license:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the **Licensed Instances** window, select the Veeam Agent to which you want to assign the license, click **Assign** and select the desired product edition: *Workstation* or *Server*.



# Viewing Licensed Veeam Agents and Revoking License

When Veeam Agent connects to the backup server, Veeam Backup & Replication applies a license to the Veeam Agent. You can view to which Veeam Agents the license is currently applied.

To view a list of licensed Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.

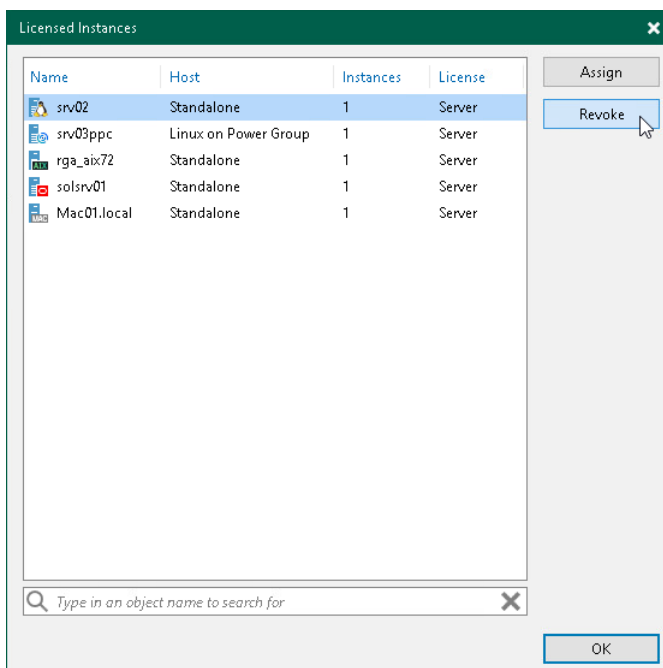
In the list of licensed instances, Veeam Backup & Replication displays Veeam Agents that have established a connection with the backup server when you created the backup job.

## Revoking License from Veeam Agents

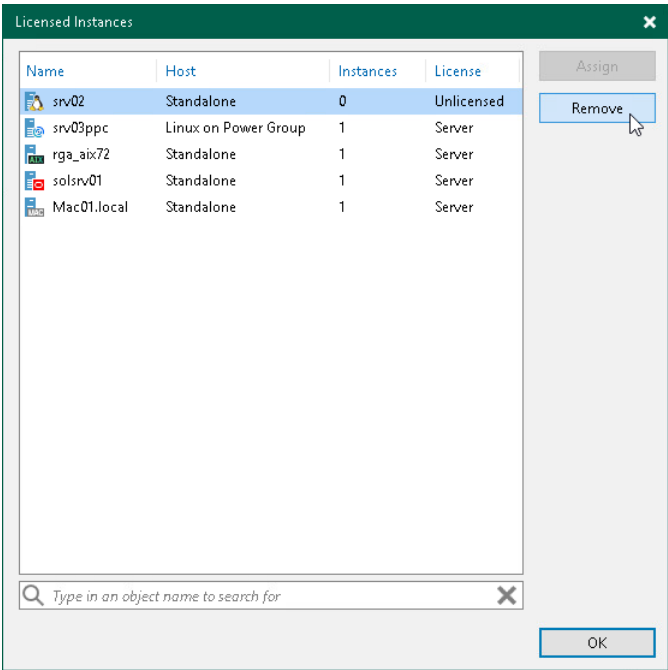
You can revoke the license from some Veeam Agents and re-apply it to other protected workloads. License revoking can be helpful, for example, if you do not want to use some Veeam Agents with Veeam Backup & Replication anymore.

To revoke a license from the Veeam Agent:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the Licensed Instances window, select a Veeam Agent and click **Revoke**. Veeam Backup & Replication will revoke the license from the Veeam Agent, and the license will be freed for other workloads that you want to protect with Veeam products.



The Veeam Agent from which you have revoked the license will become unable to connect to the Veeam backup server but will remain in the **Licensed Instances** list. To allow this Veeam Agent to create backups in the Veeam backup repository, select the Veeam Agent and click **Remove**. During the next backup job session, the Veeam Agent will connect to the Veeam backup server and start consuming the license.



# Performing Data Protection Tasks

You can perform the following data protection tasks:

- Back up your data and store the resulting backup files in one of the following types of Veeam backup repositories:
  - [In a backup repository managed by a Veeam backup server](#)
  - [In a Veeam Cloud Connect repository](#)
- [Copy Veeam Agent backups from the backup repository to a secondary backup repository with backup copy jobs.](#)
- [Use SureBackup.](#)
- [Archive Veeam Agent backups to tapes with backup to tape jobs.](#)

# Backing Up to Backup Repositories

You can store backups created with Veeam Agent in backup repositories connected to Veeam backup servers. To do this, you must perform the following actions:

1. [Set up user permissions at the backup repository side.](#)
2. [Point the Veeam Agent backup job to the backup repository.](#)

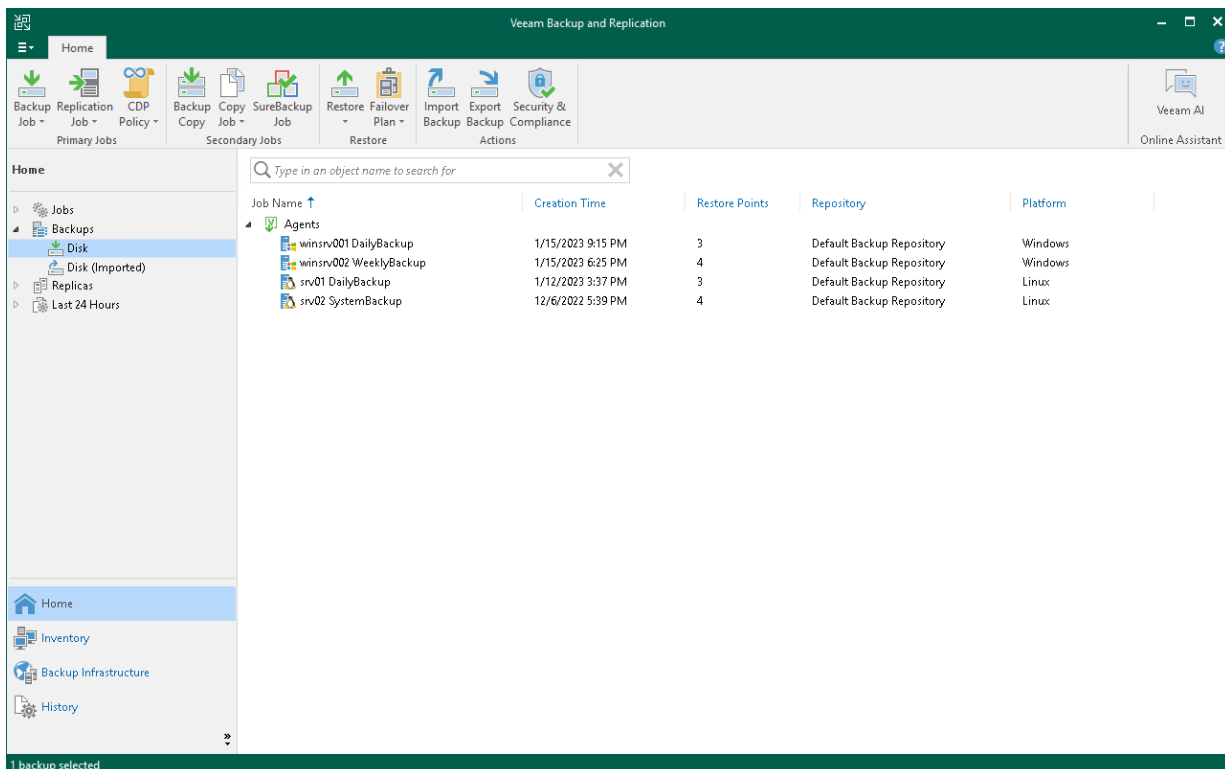
## NOTE

Consider the following:

- A Veeam Agent backup job can be started automatically upon the defined schedule or manually from the Veeam Agent computer. You cannot start, stop, retry or edit Veeam Agent backup jobs in the Veeam Backup & Replication console.
- If the user is granted restore permissions on the Veeam backup server, the user will be able to see all backups in the backup repository.
- The user who creates a Veeam Agent backup in the backup repository is set as the owner of the backup file. The backup file owner can access this file and restore data from it. If the user who is not the backup file owner needs to perform operations with the backup file, the user must have the Veeam Backup & Replication role that allows to perform these operations. To learn more about roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

Backup jobs targeted at the backup repository become visible in Veeam Backup & Replication under the **Jobs > Backup** node in the **Home** view. Backups created with Veeam Agent are available under the **Backups > Disk** node in the **Home** view.

The Veeam Backup Administrator working with Veeam Backup & Replication can manage Veeam Agent backup jobs and restore data from Veeam Agent backups. To learn more, see [Restoring Data from Veeam Agent Backups](#) and [Performing Administration Tasks](#).





# Backing Up to Cloud Repositories

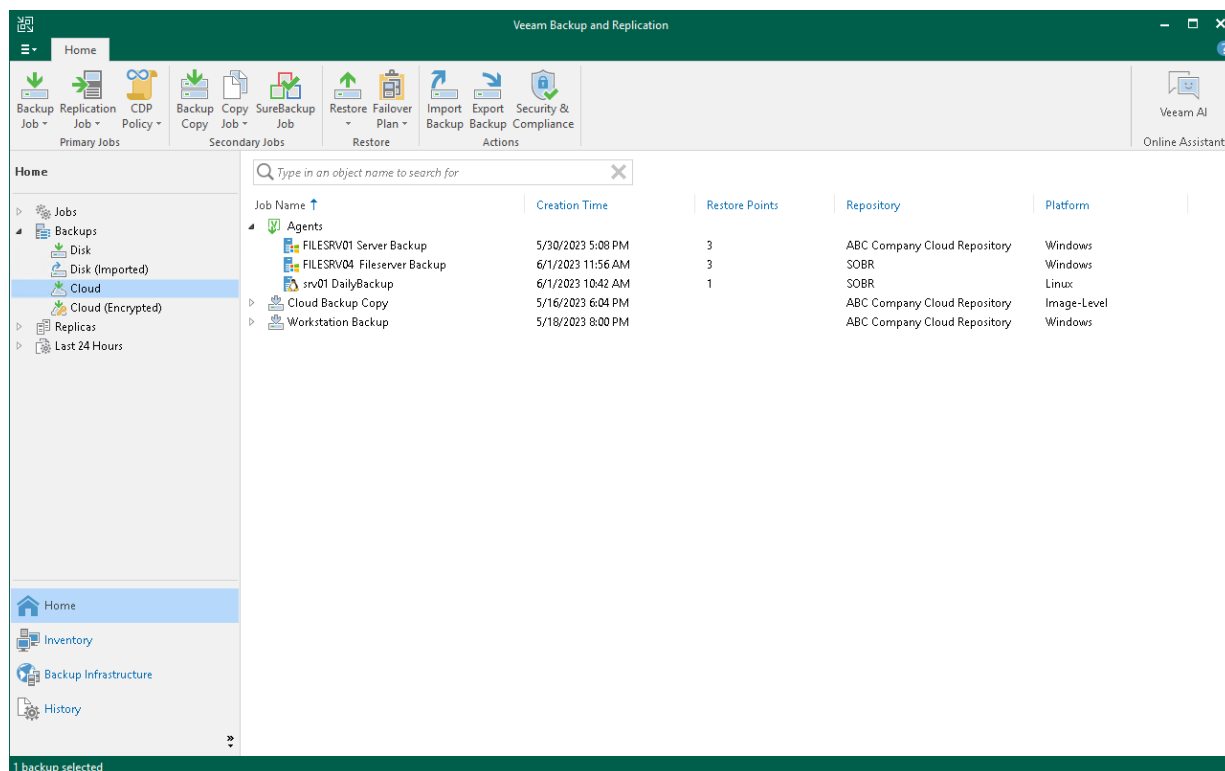
You can store backups created with Veeam Agent in cloud repositories provided to you by a Veeam Cloud Connect service provider. To do this, you must connect to the service provider and point the backup job to the cloud repository. To learn more, see [Specify Service Provider Settings](#).

## Veeam Agent Backups on Tenant Side

Backups created with Veeam Agent are available under the **Cloud** node in the **Home** view of the Veeam Backup & Replication console deployed on the tenant side.

The backup administrator working with Veeam Backup & Replication on the tenant side can manage Veeam Agent backups created in the cloud repository and restore data from such backups. To recover data from a Veeam Agent backup, you can perform the following operations:

- [Export computer disks as virtual disks](#).
- [Restore guest OS files](#).
- [Export restore points to standalone full backup files](#).



## Veeam Agent Backups on Service Provider Side

The service provider can view information about backup and restore sessions performed by Veeam Agent users. The full list of sessions is available in the **History** view of the Veeam backup console deployed on the service provider side. The list of sessions performed within the last 24 hours is available under the **Last 24 hours** node in the **Cloud Connect** view of the Veeam backup console on the service provider side. The service provider cannot view detailed statistics about individual sessions in the list.

The service provider cannot perform restore tasks with Veeam Agent backups that are stored in the cloud repository. The service provider can perform the following restore tasks with unencrypted Veeam Agent backups stored in the cloud repository:

- Instant recovery
- Disk restore
- Disk publish

To learn more, see the [Restoring Data from Tenant Backups](#) section in the Veeam Cloud Connect Guide.

The screenshot shows the Veeam Backup and Replication console interface. The top navigation bar includes 'Home' and 'View' tabs. Below the navigation bar is a toolbar with icons for Backup Job, Replication Job, CDP Policy, Backup SureBackup Copy, Restore, Import Backup, Export Backup, and Security & Compliance Actions. The main area is divided into a left sidebar and a central pane. The sidebar shows a tree view under 'Cloud Connect' with options like Cloud Gateways, Gateway Pools, Tenants, Backup Storage, Replica Resources, and 'Last 24 Hours'. The central pane displays a table of backup sessions.

Job Name	Session Type	Status	Start Time	End Time	Tenant	Data Sent	Data Received
srv001_daily_backup	Cloud Backup	Success	11/30/2023 11:24 AM	11/30/2023 11:25 AM	ABC Company	54.1 KB	655 MB
srv001_volume_backup	Cloud Backup	Success	11/30/2023 11:26 AM	11/30/2023 11:27 AM	ABC Company	15.6 KB	15.8 KB
srv002_system_backup	Cloud Backup	Success	11/30/2023 11:25 AM	11/30/2023 11:26 AM	ABC Company	97.8 KB	2.2 MB
srv003_system_backup	Cloud Backup	Success	11/30/2023 11:19 AM	11/30/2023 11:20 AM	ABC Company	40.6 KB	20.5 MB

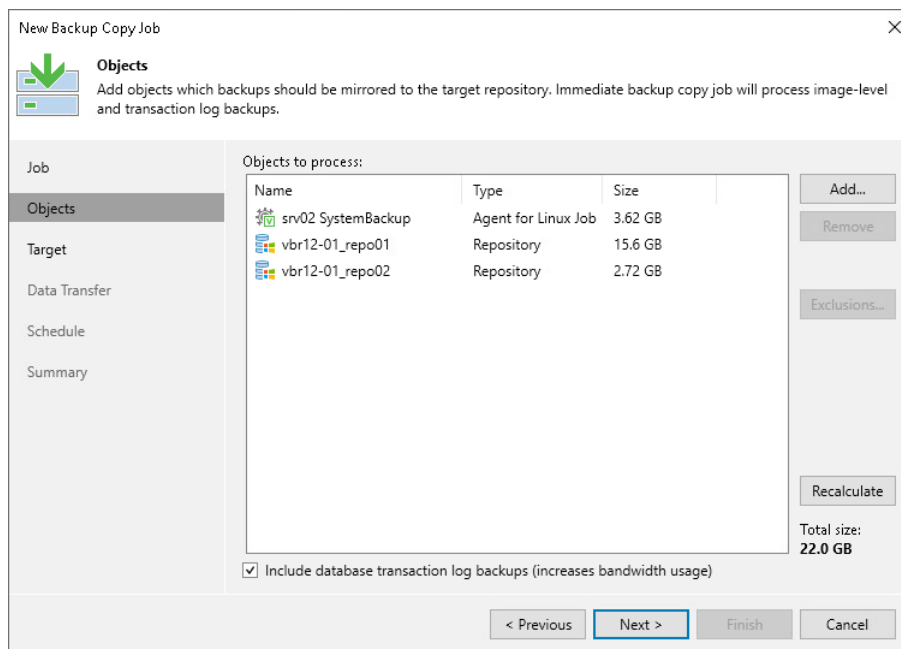
At the bottom of the console, there is a status bar showing '5 sessions'.

# Performing Backup Copy for Veeam Agent Backups

You can configure backup copy jobs that will copy backups created with Veeam Agent to a secondary backup repository.

Backup copy jobs treat Veeam Agent backups as usual backup files. The backup copy job setup and processing procedures practically do not differ from the same procedures for a backup copy job that processes VM backups. To learn more about backup copy jobs, see the [Backup Copy](#) section in the Veeam Backup & Replication User Guide.

When mapping a backup copy job to a Veeam Agent backup, consider the limitations listed in the [Map Backup File](#) section in the Veeam Backup & Replication User Guide.



## Restoring Data from Copies of Veeam Agent Backups

Backups copied to the secondary backup repository do not preserve user access permissions. At the same time, users who created backups do not have access permissions on these secondary repositories. For this reason, users cannot restore data from their backups residing in the secondary site.

To overcome this limitation, you can delegate the restore task to backup administrators who work with Veeam Backup & Replication. Backup administrators can use Veeam Backup & Replication options to recover data from such backups: for example, perform file-level restore or retrieve necessary application items with Veeam Explorers.

You can also restore data from the copied backup stored in the target repository using Veeam Agent.

# Using SureBackup

Veeam Backup & Replication offers the SureBackup technology to test backups and check if you can recover data from them. You can verify any restore point of a backed-up computer protected with Veeam Agent for Linux.

To learn more about the logic behind SureBackup, see the [How SureBackup Works](#) section in the Veeam Backup & Replication User Guide.

Before creating the SureBackup job, check limitations for Veeam Agent backups below. Then launch the **New SureBackup Job** wizard to create the SureBackup job. To learn more, see the [Creating SureBackup Job](#) section in the Veeam Backup & Replication User Guide.

## Limitations

For backups created with Veeam Agent for Linux, SureBackup has the following limitations:

- SureBackup job in the **Backup verification and content scan only** mode is not supported.
- SureBackup is not supported for backups stored in the Veeam Cloud Connect repository.
- SureBackup is not supported for backups stored in the archive tier of the the scale-out backup repository.
- SureBackup is not supported for backups containing drives greater than 64 TB.
- If you plan to verify computer recovery with VMware vSphere, consider the following:
  - SureBackup is not supported for backups of 4 KB sector drives.
  - SureBackup is not supported for backups of storage spaces.
  - SureBackup is not supported for backups containing more than 54 drives.
- When Veeam Backup & Replication publishes virtual machines based on backed-up Veeam Agent computers in the isolated virtual environment, all these virtual machines are included in the first isolated network added during the virtual lab configuration. To learn more, see the [Create Isolated Networks](#) section in the Veeam Backup & Replication User Guide.
- You cannot use SureBackup with backup files created with Veeam Agent for Linux on Power.
- The successful recovery verification is not guaranteed for the following Linux distributions:
  - Amazon Linux 2
  - Amazon Linux 2023
  - openSUSE Tumbleweed
- The successful recovery verification is not guaranteed for backups of Linux-based systems that contain encrypted devices.
- If you want Veeam Backup & Replication to connect the recovered VM to the virtual network, one of the following configuration utilities must be installed on the protected computer:
  - Netplan
  - NetworkManager

- `sysconfig`
- `systemd-networkd`
- `ifupdown/ifupdown2`
- SureBackup is not supported for file-level backups. You must use volume-level backup of the protected computer. The backup must include the `root` file system (`/`) and all partitions specified in the `/etc/fstab` file. To learn more about backup types, see [Backup Types](#).

# Archiving Veeam Agent Backups to Tape

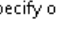
You can configure backup to tape jobs to archive Veeam Agent backups to tape.

Backup to tape jobs treat Veeam Agent backups as usual backup files. The archiving job setup and processing procedures practically do not differ from the regular ones. To learn more about backup to tape jobs, see the [Backup to Tape](#) section in the Veeam Backup & Replication User Guide.

## NOTE

For the **After this job** option in the backup to tape job schedule settings, you cannot select a backup job managed by Veeam Agent or a standalone Veeam Agent backup job as the preceding backup job.

New Backup to Tape Job

**Backups**  
Specify objects to be processed by this tape job.

Name

Backups

Media Pool

Incremental Backup

Options

Schedule

Summary

Backups:

Name	Type	Size	
Daily Backup	Linux Agent Ba...	3.17 GB	
System Backup	Linux Agent Ba...	333 GB	

Add...  
  
Remove  
  
  
Up  
  
Down  
  
  
Full:  
335 GB  
Incremental:  
1.58 GB

# Restoring Data from Veeam Agent Backups

You can perform the following restore operations:

- [Restore Veeam Agent backups to VMware vSphere VMs](#)
- [Restore Veeam Agent backups to Hyper-V VMs](#)
- [Restore Veeam Agent backups to Nutanix AHV VMs](#)
- [Restore Veeam Agent backups to Proxmox VE VMs](#)
- [Restore data from Veeam Agent backups to Microsoft Azure](#)
- [Restore data from Veeam Agent backups to Amazon EC2](#)
- [Restore data from Veeam Agent backups to Google Compute Engine](#)
- [Restore individual files and folders from Veeam Agent backups](#)
- [Restore application items from Veeam Agent backups with Veeam Explorers](#)
- [Export computer disks as VMDK, VHD or VHDX disks](#)
- [Publish disks to analyze backup content](#)
- [Export restore points of Veeam Agent backups to standalone full backup files](#)

# Restoring Veeam Agent Backup to vSphere VM

In the Veeam Backup & Replication console, you can use Instant Recovery to restore a Veeam Agent computer as a VMware vSphere VM in your virtualization environment.

A restored VMware vSphere VM will have the same settings as the backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves the settings of the Veeam Agent computer from the backup and applies them to the target VM. These settings include:

- Amount of RAM.
- Number of CPU cores.
- Number of network adapters.
- Network adapter settings.
- BIOS UUID.

If you do not want to preserve the backed-up machine UUID for a VMware vSphere VM, you can create a new UUID during the Instant Recovery configuration process.

- Number of disks and volumes.
- Size of volumes.

## Considerations and Limitations

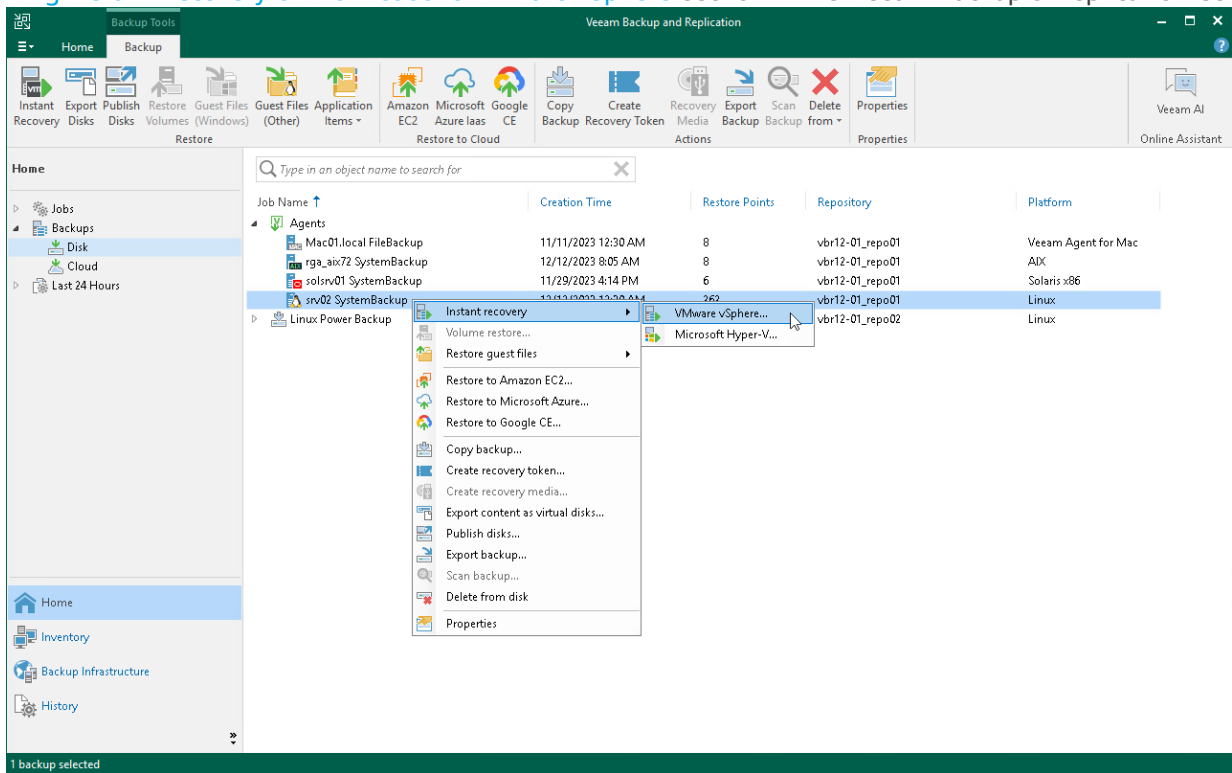
If you restore a Veeam Agent computer to a VMware vSphere VM, consider the following:

- You can use entire machine or volume-level backups of Linux computers. Volume-level backups must include the `root` file system (`/`) and all partitions specified in the `/etc/fstab` file.
- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.
- If you restore a workload to the production network, make sure that the original workload is powered off.
- If the disk you want to restore contains an LVM volume group, consider the following:
  - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Backup & Replication will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
  - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Backup & Replication will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.



# Restore to vSphere VM

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the [Performing Instant Recovery of Workloads to VMware vSphere](#) section in the Veeam Backup & Replication User



Guide. 1 backup selected

# Restoring Veeam Agent Backup to Hyper-V VM

In the Veeam Backup & Replication console, you can use Instant Recovery to restore a Veeam Agent computer as a Hyper-V VM in your virtualization environment.

A restored Hyper-V VM will have the same settings as the backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves settings of the Veeam Agent computer from the backup and applies them to the target VM.

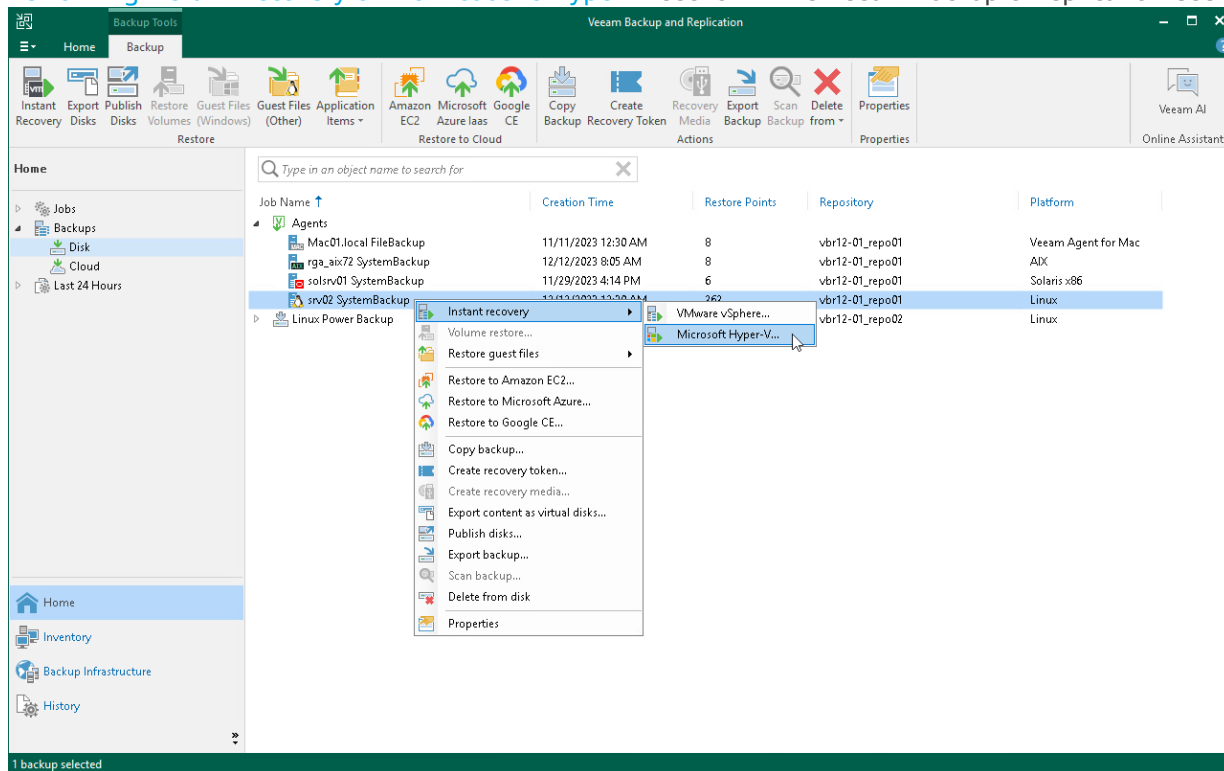
## Considerations and Limitations

If you restore a Veeam Agent computer to a Hyper-V VM, consider the following:

- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot use backups stored in a Veeam Cloud Connect repository for this operation.
- To restore to a Hyper-V VM from a backup of a Linux computer, you must consider the Hyper-V limitations. To learn more, see [this Microsoft article](#).
- 
- Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.
- Veeam Agent computer disks are recovered as dynamically expanding virtual disks.
- By default, Veeam Backup & Replication automatically powers on a VM after restore. If you do not want to power on a VM after restore, you can change this setting during the Instant Recovery configuration process.
- If the disk you want to restore contains an LVM volume group, consider the following:
  - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Backup & Replication will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
  - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Backup & Replication will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

# Restore to Hyper-V VM

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the [Performing Instant Recovery of Workloads to Hyper-V](#) section in the Veeam Backup & Replication User Guide.



# Restoring Veeam Agent Backup to Nutanix VM

You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a Nutanix AHV VM in your virtualization environment.

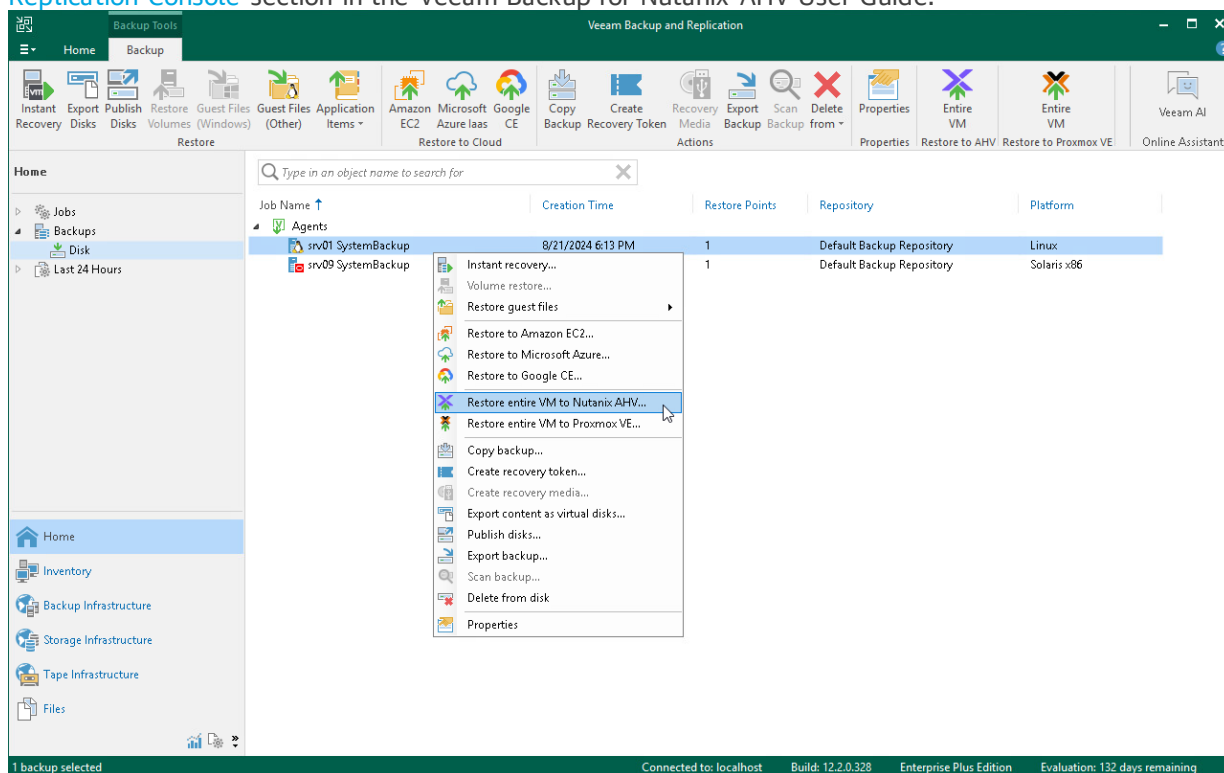
## Considerations and Limitations

If you restore a Veeam Agent computer to a Nutanix AHV VM, consider the following:

- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- If the disk you want to restore contains an LVM volume group, consider the following:
  - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
  - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

## Restore to Nutanix AHV

The procedure of restore to Nutanix AHV for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more about restore to Nutanix AHV, see the [Restoring VMs Using Veeam Backup & Replication Console](#) section in the Veeam Backup for Nutanix AHV User Guide.



# Restoring Veeam Agent Backup to Proxmox VM

You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a Proxmox VE VM in your virtualization environment.

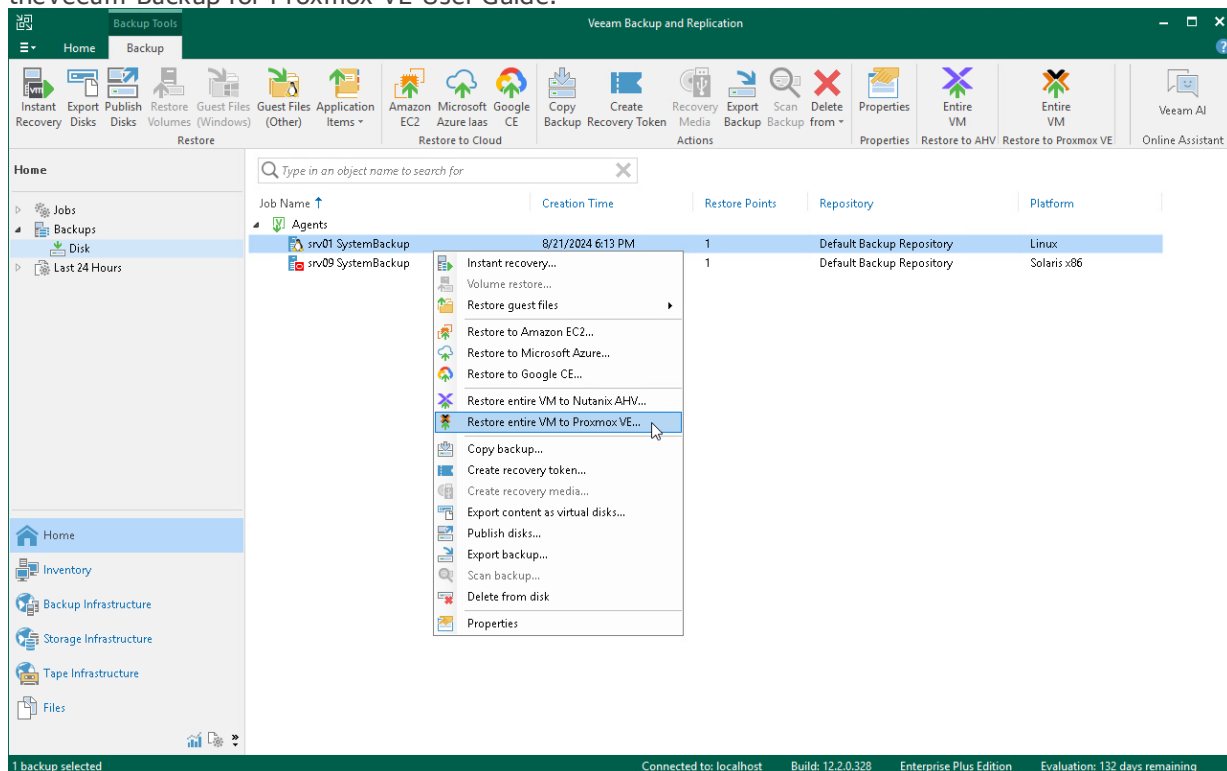
## Considerations and Limitations

If you restore a Veeam Agent computer to a Proxmox VE VM, consider the following:

- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- If the disk you want to restore contains an LVM volume group, consider the following:
  - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
  - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

## Restore to Proxmox VE

The procedure of restore to Proxmox VE for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more about restore to Proxmox VE, see the [Performing VM Restore](#) section in the Veeam Backup for Proxmox VE User Guide.



# Restoring to Microsoft Azure

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Microsoft Azure.

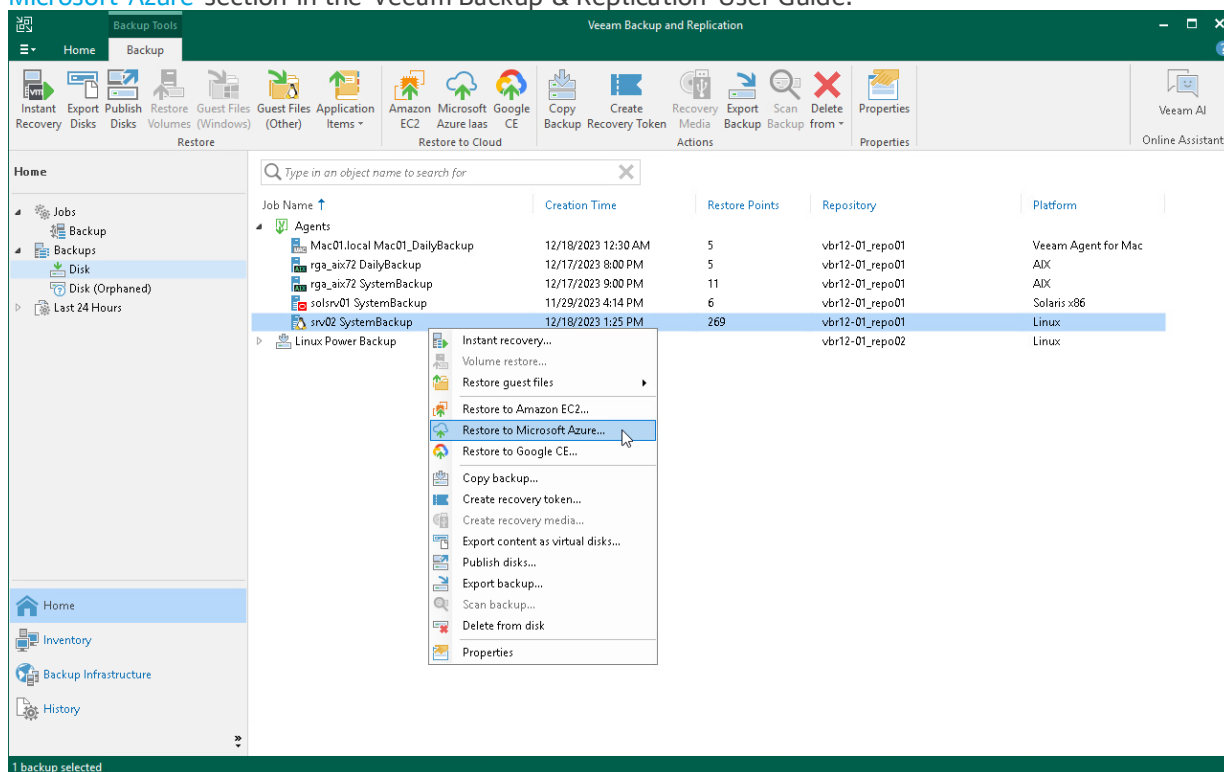
## Considerations and Limitations

If you restore a Veeam Agent computer to Microsoft Azure, consider the following:

- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level.
- If you recover an EFI-based system to Microsoft Azure, Veeam Agent will restore a BIOS-based Generation 1 VM.
- Veeam Backup & Replication offers experimental support for generation 2 VMs within restore to Microsoft Azure feature. To learn more, see the [Generation 2 VM Support](#) section in the Veeam Backup & Replication User Guide.

## Restore to Microsoft Azure

The procedure of restore to Microsoft Azure from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Microsoft Azure, see the [Restoring to Microsoft Azure](#) section in the Veeam Backup & Replication User Guide.



# Restoring to Amazon EC2

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Amazon EC2.

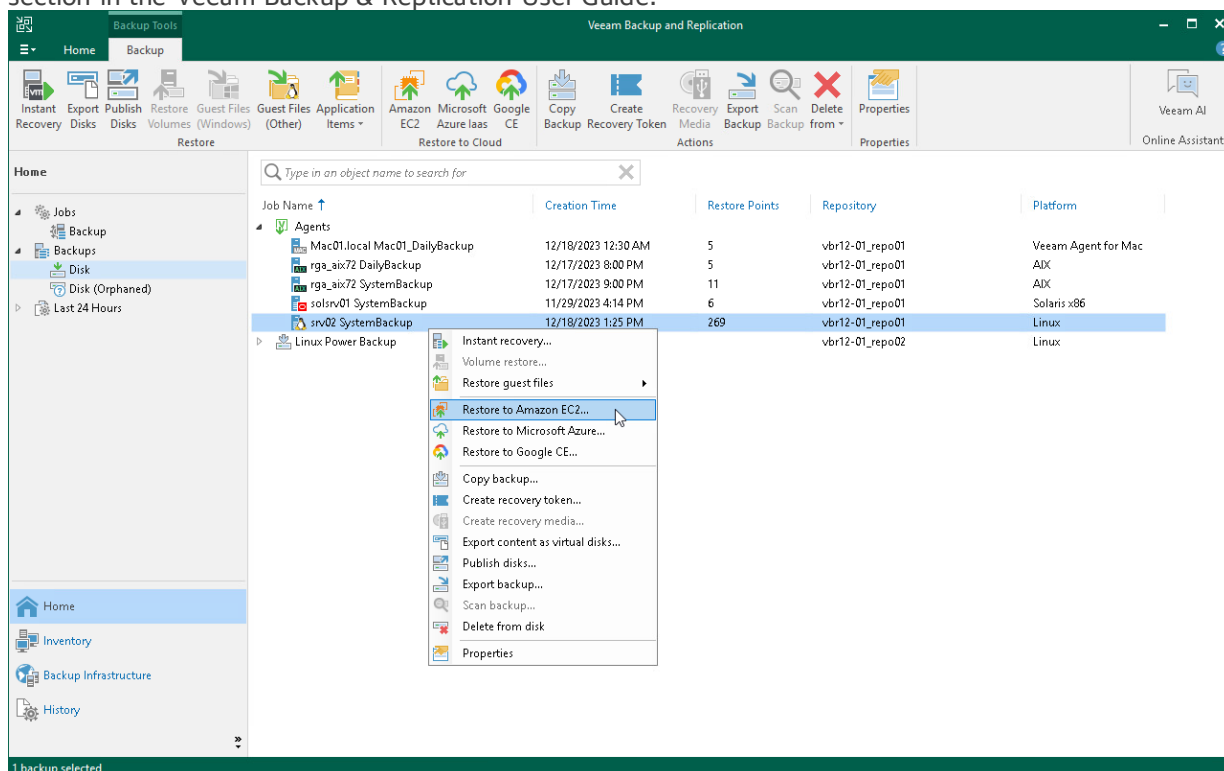
## Considerations and Limitations

If you restore a Veeam Agent computer to Amazon EC2, consider the following:

- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level. If the disk you want to restore contains an LVM volume group, consider the following:
  - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
  - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

## Restore to Amazon EC2

The procedure of restore to Amazon EC2 from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Amazon EC2, see the [Restoring to Amazon EC2](#) section in the Veeam Backup & Replication User Guide.



# Restoring to Google Compute Engine

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Google Compute Engine.

## Considerations and Limitations

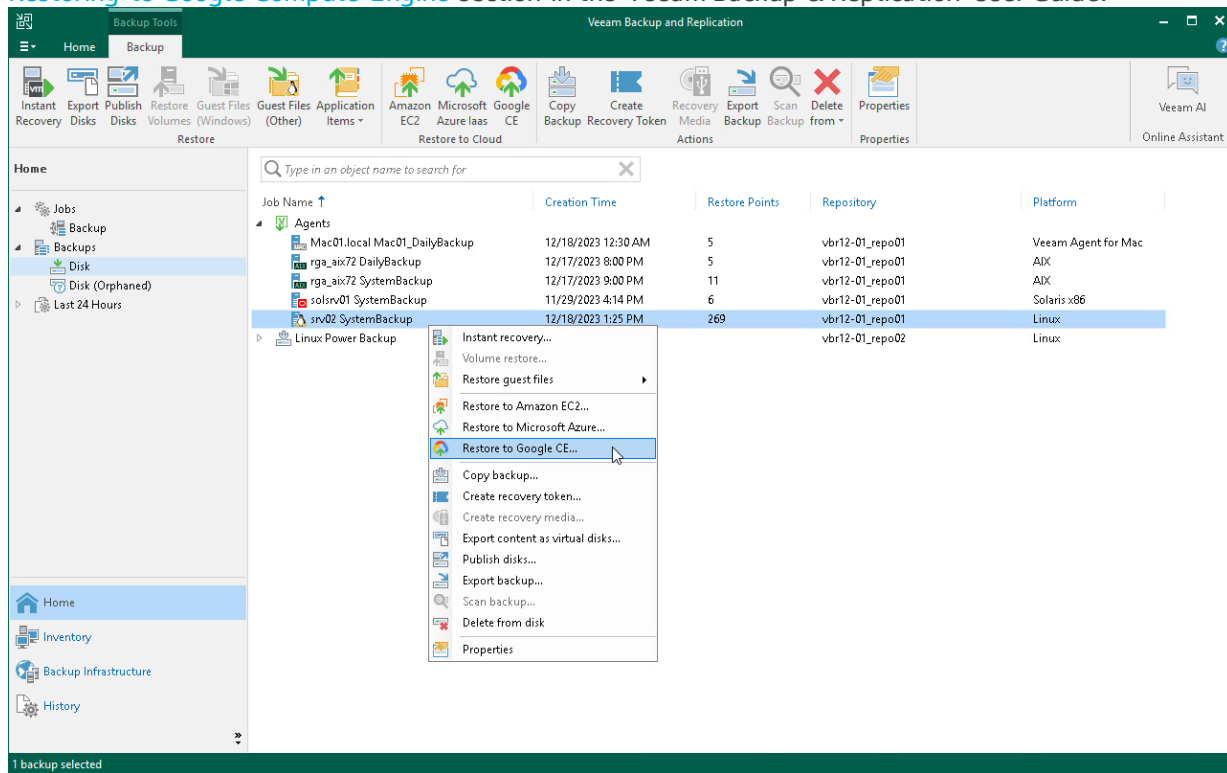
If you restore a Veeam Agent computer to Google Compute Engine, consider the following:

- You can use backups of Linux computers stored in a Veeam backup repository. You cannot perform this operation with Veeam Agent backups created on the Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level. If the disk you want to restore contains an LVM volume group, consider the following:
  - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
  - Root file system partition and boot partition must not be on LVM logical volumes. For more information on this limitation, see [Google documentation](#).
  - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.



# Restore to Google Compute Engine

The procedure of restore to Google Compute Engine from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Google Compute Engine, see the [Restoring to Google Compute Engine](#) section in the Veeam Backup & Replication User Guide.



# Restoring Files and Folders

You can use the Veeam Backup & Replication console to restore individual files and folders from Veeam Agent backups.

The procedure of file-level restore from a Veeam Agent backup is similar to the same procedure for a VM backup. To learn more about file-level restore, see the [Restore from Linux, Unix and Other File Systems](#) section in the Veeam Backup & Replication User Guide.

When you perform the file-level restore procedure, Veeam Backup & Replication provides the following options for mounting disks of the machine from the backup or replica:

- Mounting disks to a helper host — any Linux host in your infrastructure with a [supported operating system](#). You can also mount disks from the Veeam Agent for Linux backup to the original host.
- Mounting disks to a temporary helper appliance — a helper VM required to mount Linux computer disks from the backup.

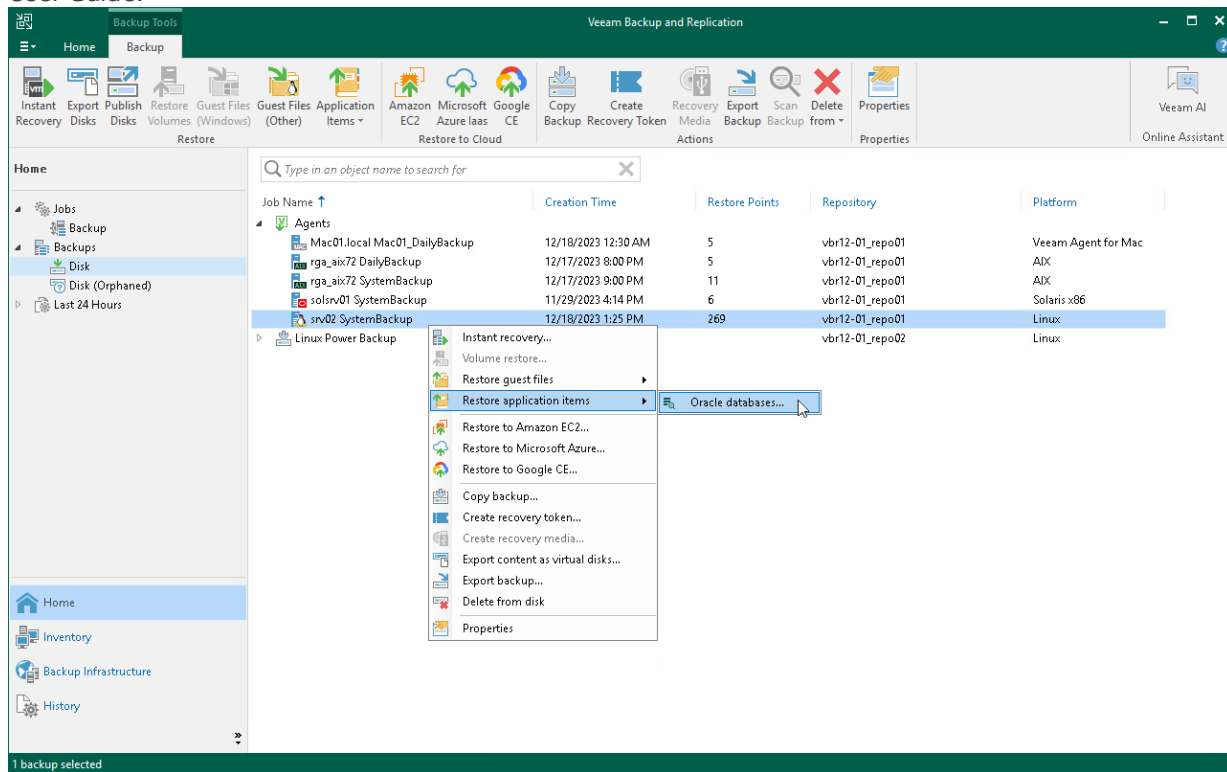
If you have selected to mount disks to a temporary helper appliance, it is recommended that you add a vCenter Server and not a standalone ESXi host in the Veeam backup console. If Veeam Backup & Replication is set up to deploy a helper appliance on a standalone ESXi host, after Veeam Backup & Replication removes the helper appliance, the helper VM will be displayed in vCenter as orphaned.

# Restoring Application Items

You can use Veeam Explorers to restore application items from backups created using Veeam Agent for Linux. Veeam Backup & Replication lets you restore items and objects from the following applications:

- Oracle
- PostgreSQL

The procedure of application item restore from a Veeam Agent backup does not differ from the same procedure for a VM backup. To learn more, see the [Application Item Restore](#) section in the Veeam Backup & Replication User Guide.



# Exporting Disks

You can restore computer disks from Veeam Agent backups created using Veeam Agent for Linux and convert them to disks of the VMDK, VHD or VHDX format.

During disks restore, Veeam Backup & Replication creates standard virtual disks that can be used by VMware vSphere and Microsoft Hyper-V VMs.

- When you restore a disk in the VMDK format, Veeam Backup & Replication creates a pair of files that make up the VM virtual disk: a descriptor file and file with the virtual disk content.
- When you restore a disk in the VHD/VHDX format, Veeam Backup & Replication creates a file of the VHD or VHDX format.

You can save converted disks locally on any server or SMB share added to the backup infrastructure or place disks on a datastore connected to an ESXi host (for VMDK disk format only). VMDK disks can be restored as thin provision and thick disks:

- Disks restored to a datastore are saved in the thin provisioned format.
- Disks restored to a server are saved in the thick provisioned format.

Veeam Backup & Replication supports batch disk restore. For example, if you choose to restore 2 computer disks, Veeam Backup & Replication will convert them to 2 virtual disks and store these disks in the specified location.

## IMPORTANT

Consider the following:

- If the backup from which you restore disks contains a Btrfs storage pool, during the disk restore process Veeam Backup & Replication will create a separate disk and restore the Btrfs pool to this disk.
- If the disk you want to restore contains an LVM volume group, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.

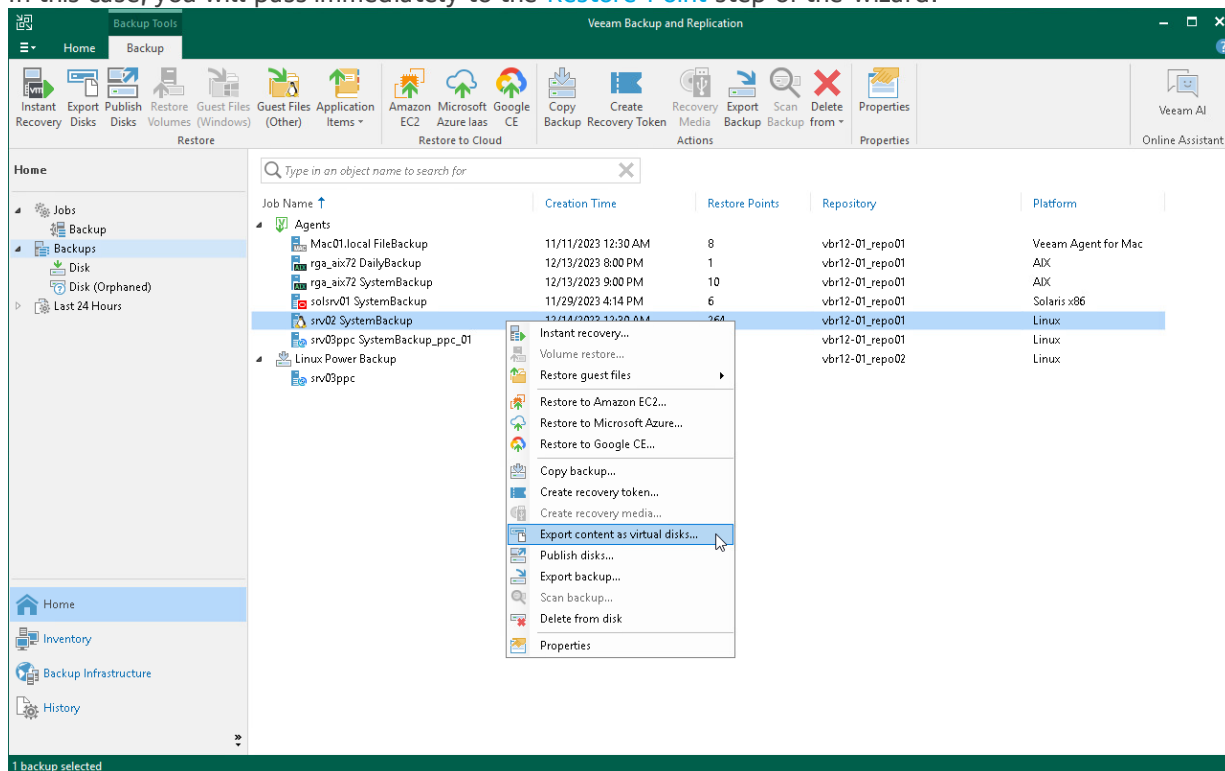
To restore disks and convert them to the VMDK, VHD or VHDX format, perform the following steps in the **Export Disk** wizard:

# Step 1. Launch Export Disk Wizard

To launch the **Export Disk** wizard, do either of the following:

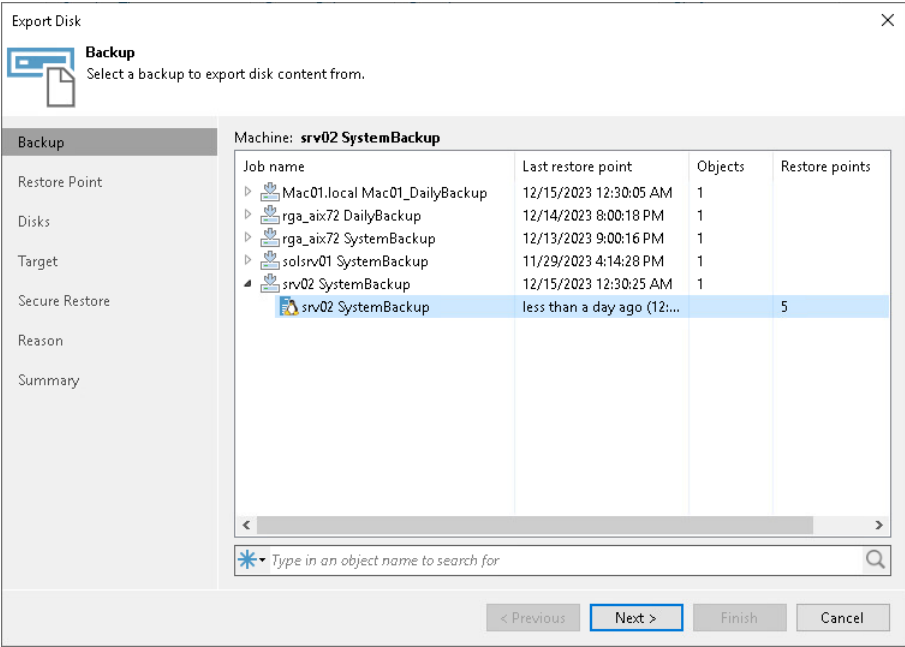
- Open the **Home** tab and click **Restore > Agent > Disk restore > Export disk**. In this case, you will be able to select a backup of the necessary Veeam Agent computer at the **Backup** step of the wizard.
- Open the **Home** view. In the inventory pane, click the **Backups** node. In the working area, expand the necessary Veeam Agent backup, select the necessary computer in the backup and click **Export Disks** on the ribbon or right-click a computer in the backup and select **Export content as virtual disks**.

In this case, you will pass immediately to the **Restore Point** step of the wizard.



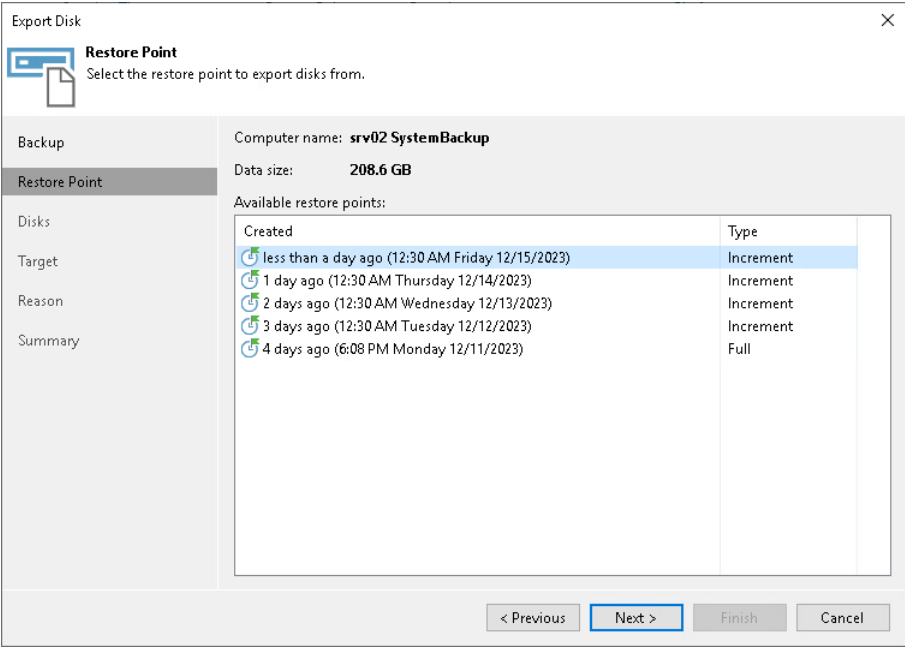
# Step 2. Select Backup

At the **Backup** step of the wizard, select a backup from which you want to restore disks. In the list of backups, Veeam Backup & Replication displays all backups that are currently hosted on the Veeam backup repository and Veeam Cloud Connect repository.



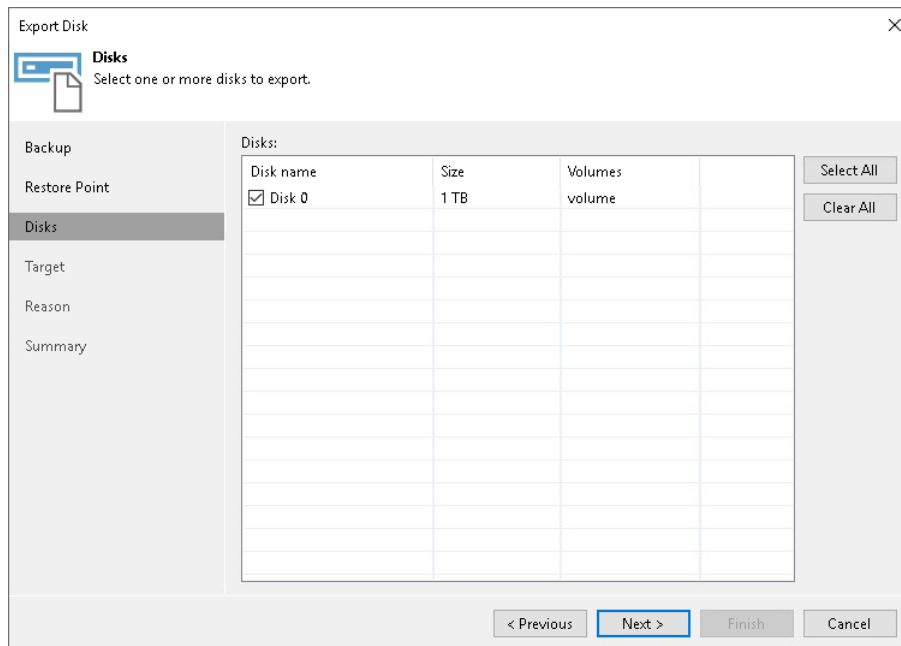
# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point from which you want to restore disks. In the list of points, Veeam Backup & Replication displays all restore points that have been created. Make sure that you select a restore point that relates to the selected backup.



## Step 4. Select Disks

At the **Disks** step of the wizard, select check boxes next to those disks that you want to export.





## Step 5. Select Destination and Disk Format

At the **Target** step of the wizard, select the destination for disk export and format in which you want to save the resulting virtual disk.

1. From the **Server** list, select a server on which the resulting virtual disks must be saved. If you plan to save the disks in the VMDK format on a datastore, select an ESXi host to which this datastore is connected.
2. In the **Path to folder** field, specify a folder on the server or datastore where the virtual disks must be placed.
3. Select the export format for disks:
  - **VMDK** – select this option if you want to save the resulting virtual disk in the VMware VMDK format.
  - **VHD** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHD format.
  - **VHDX** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHDX format (supported by Microsoft Windows Server 2012 and later).
4. Click **Disk type** to specify how the resulting disk must be saved:
  - [For VMDK disk format] in the thin provisioned, lazy zeroed thick provisioned, or eagerly zeroed thick provisioned format
  - [For VHD and VHDX disk formats] in the dynamic or fixed format
5. [For export of a VMDK disk to an ESXi host] Click the **Pick proxy to use** link to select backup proxies over which backup data must be transported to the target datastore.

## NOTE

Consider the following:

- If you have selected to store the resulting virtual disk in a datastore, you will be able to save the virtual disk in the VMDK format only. Other options will be disabled.
- If you have selected to store the resulting virtual disk on the server running Microsoft Windows Server OS and in the VMDK format, you will be able to save the virtual disk in the lazy zeroed thick provisioned format only.

The screenshot shows the 'Export Disk' dialog box with the 'Target' tab selected. The dialog has a sidebar on the left with options: Backup, Restore Point, Disks, Target (selected), Reason, and Summary. The main area is titled 'Target' and contains the following fields and options:

- Server:** A dropdown menu showing 'winsrv0042019.tech.local'.
- Path to folder:** A text box containing 'C:\WeeamBackup' and a 'Browse...' button.
- Export format:** Three radio button options:
  - VMDK** (selected): This virtual disk type is used by VMware products such as VMware Workstation, or VMware vSphere. Maximum VMDK disk size is 62TB. Pick proxy to use.
  - VHD**: This virtual disk type is used by Microsoft products such as Microsoft Hyper-V or Microsoft Azure. Maximum VHD disk size is 2TB.
  - VHDX**: This virtual disk type is used by more recent versions of Microsoft products such as Microsoft Hyper-V. Maximum VHDX disk size is 64TB.
- Disk type:** A dropdown menu showing 'Thick (lazy)'.

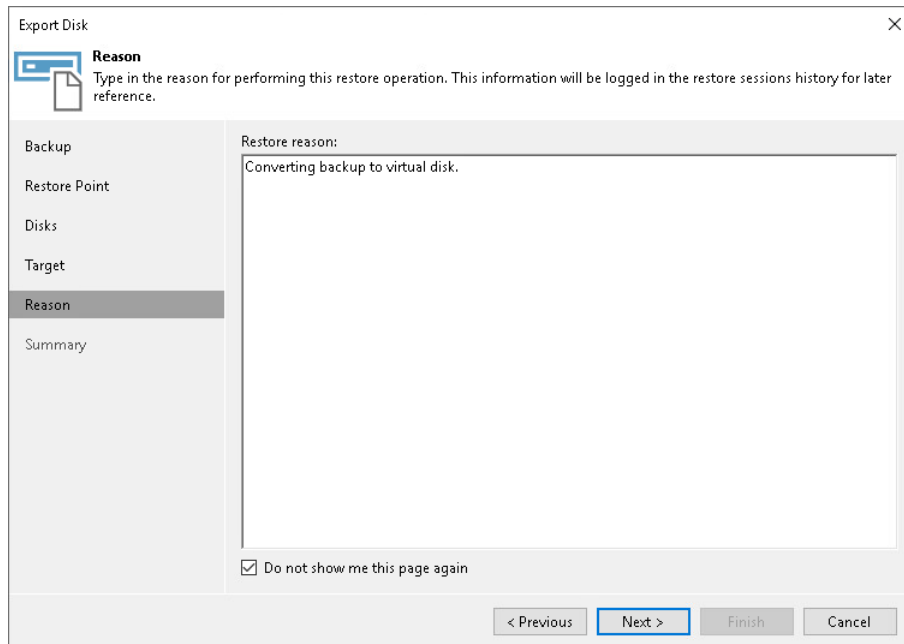
At the bottom of the dialog are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

## Step 6. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the computer volume.

### TIP

If you do not want to display the **Restore Reason** step of the wizard in future, select the **Do not show me this page again** check box.

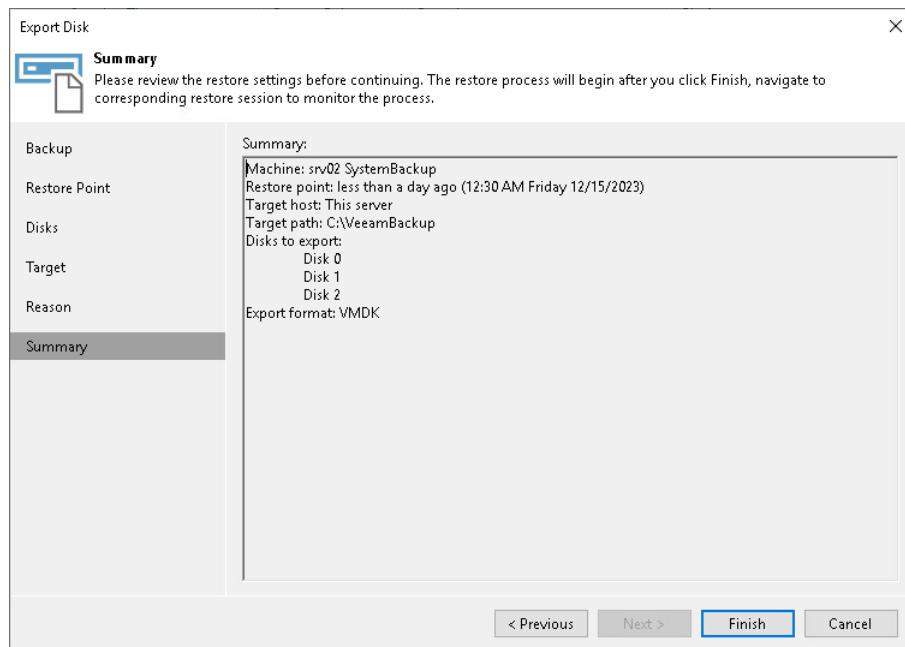


The screenshot shows the 'Export Disk' wizard window, specifically the 'Reason' step. The window has a title bar 'Export Disk' and a close button. On the left is a sidebar with steps: Backup, Restore Point, Disks, Target, Reason (selected), and Summary. The main area is titled 'Reason' and contains the instruction: 'Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.' Below this is a large text box labeled 'Restore reason:' containing the text 'Converting backup to virtual disk.' At the bottom left of the main area is a checkbox labeled 'Do not show me this page again' which is checked. At the bottom right are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 7. Complete Restore Process

At the **Summary** step of the wizard, complete the disk restore procedure.

1. Review details for the disk to be restored.
2. Click **Finish** to start the restore procedure and exit the wizard.



# Publishing Disks

You can use the Veeam backup console to publish disks from backups created by Veeam Agent backup jobs and backup copy jobs.

## TIP

You can publish disks using the PowerShell console. To learn more, see the [Disk Publishing \(Data Integration API\)](#) section in the Veeam PowerShell Reference.

Disk publishing allows you to save time by getting backup content of one or multiple disks instead of all disks from a backup. This technology gives read-only access to data and helps if you want to analyze data of your backup. For example, look for specific documents or usage patterns, or perform antivirus scan of backed-up data.

For Linux-based Veeam Agent computers, disk publishing uses the FUSE protocol. After the publishing, the target server can access the backup content using the FUSE protocol and read the necessary data from the disk.

To learn more, see the [Disk Publishing](#) section in the Veeam Backup & Replication User Guide.

## Performing Disk Publish

Before you publish disks, [check prerequisites](#). Then use the **Publish Disks** wizard.

1. [Launch the wizard](#).
2. [Select a Veeam Agent computer whose disks you want to publish](#).
3. [Select a restore point](#).
4. [Select disks](#).
5. [Specify the target server](#).
6. [Specify a reason for disk publishing](#).
7. [Finish working with the wizard](#).

## Before You Begin

Before you publish disks, check the following requirements and limitations:

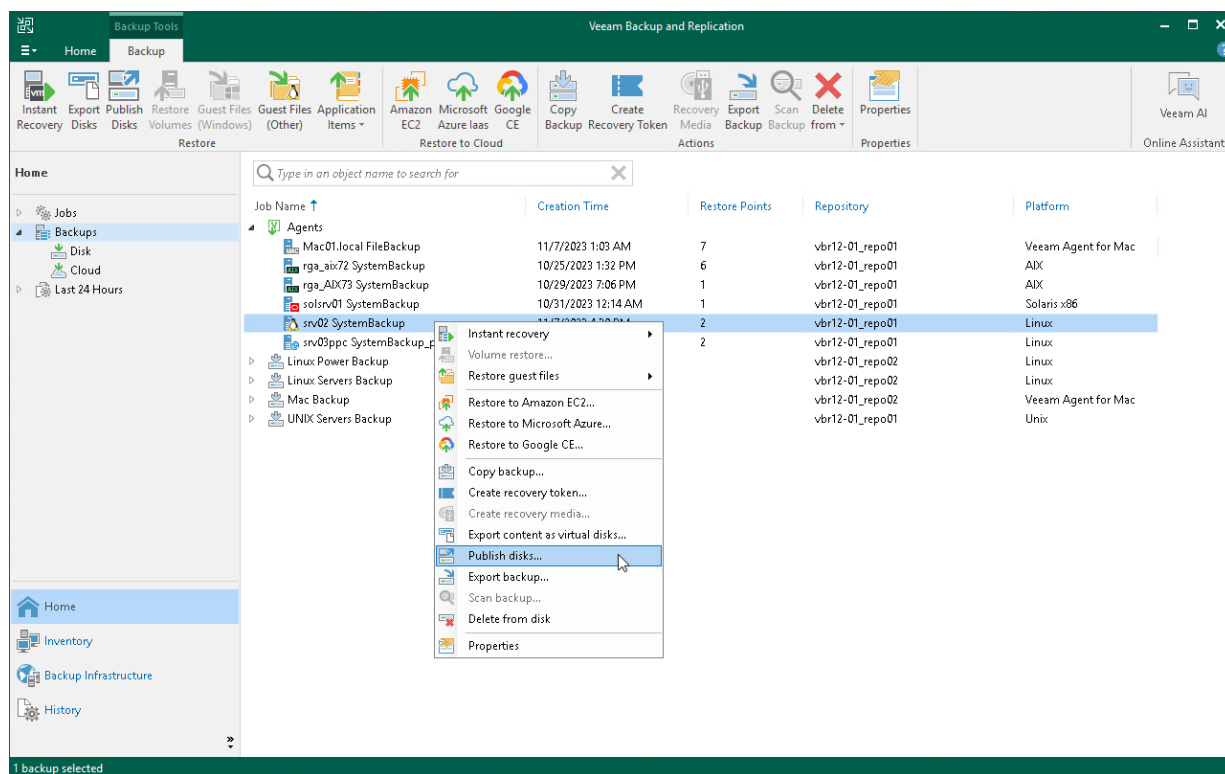
- The necessary ports must be opened on the target server. For more information, see [Ports](#).
- The target server must support the file system of the disk that you plan to publish.
- If data deduplication is enabled for some disks in a backup, data deduplication must be enabled on the target server.
- The 32-bit version of a Linux server is not supported as the target server.
- You cannot publish disks from backups stored in the Veeam Cloud Connect repository.

For the full list of limitations, see the [Considerations and Limitations](#) section in the Veeam Backup & Replication User Guide.

## Step 1. Launch Publish Disks Wizard

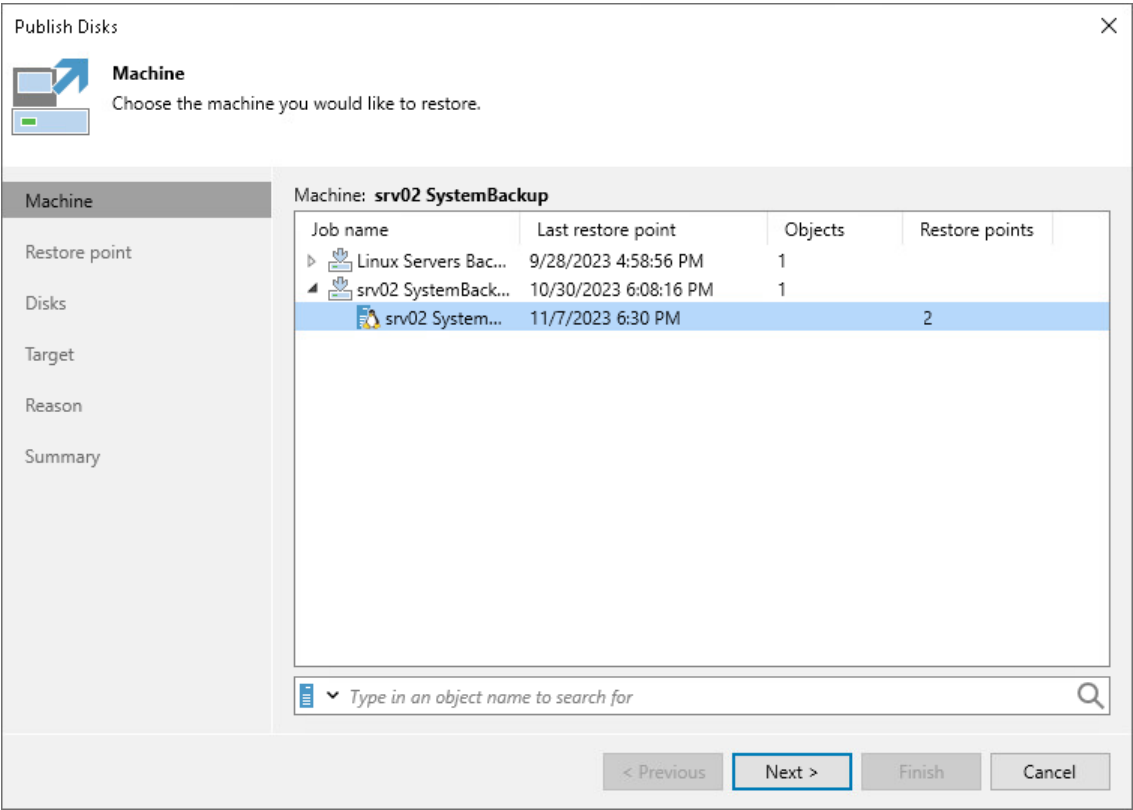
To launch the **Publish Disks** wizard, do either of the following:

- On the **Home** tab, click **Restore > Agent > Disk Restore > Publish disk**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary Veeam Agent backup, select a computer whose disks you want to publish and click **Publish Disks** on the ribbon. Alternatively, you can right-click the computer and select **Publish disks**. In this case, you will proceed to the [Restore point](#) step of the wizard.



## Step 2. Select Computer

At the **Machine** step of the wizard, expand a backup and select a Veeam Agent computer whose disks you want to publish.



### Step 3. Select Restore Point

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

Machine

Restore point

Disks

Target

Reason

Summary

Computer name: **srv02 SystemBackup**

Data size: **208 GB**

Available restore points:

Created	Type	Backup
5 days ago (5:30 PM Thursda...	Increment	srv02 SystemBackup
8 days ago (6:08 PM Monday...	Full	srv02 SystemBackup

< Previous

Next >

Finish

Cancel



## Step 4. Select Disks

At the **Disks** step of the wizard, select a check box next to the disks that you want to publish. Click **Select All** if you want to select all disks from the backup.

Machine

Restore point

**Disks**

Target

Reason

Summary

Disks:

Disk name	Size	Volumes
<input type="checkbox"/> Disk 0	160 GB	sda1
<input type="checkbox"/> Disk 1	158 GB	home; root; swap_1
<input checked="" type="checkbox"/> Disk 2	200 GB	sdb

Select All

Clear All

< Previous

Next >

Finish

Cancel

## Step 5. Select Target Server

At the **Target** step of the wizard, select a Linux server that will have access to disk content.

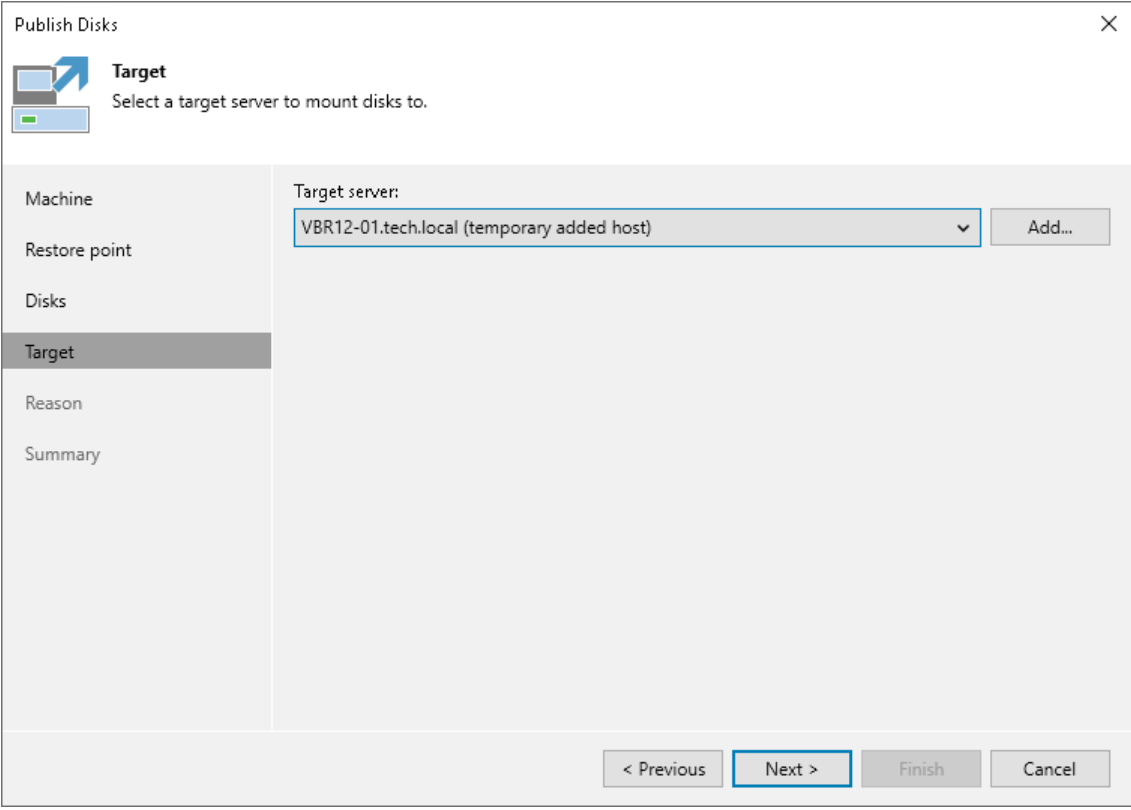
You can select one of the following types of servers:

- A server added to the backup infrastructure.

If you want to add a new backup server to the backup infrastructure at this step, click **Add**. In this case, you will be able to add a new Linux server. To learn more, see the [Adding Linux Servers](#) section in the Veeam Backup & Replication User Guide.

- A temporary server. In this case, select *Specify a different host* from the drop-down list. In the **Target Server** window, specify the following settings:
  - a. In the **Host name** field, specify a server name or IP address of the server.
  - b. Select the account from the **Credentials** list. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add a new account in the Credentials Manager. To learn more, see the [Credentials Manager](#) section in the Veeam Backup & Replication User Guide.
  - c. Click **Advanced** and customize connection settings in the **Network Settings window**. To learn more, see [Customizing Connection Settings](#).
- The original server. In this case, select *Original server* from the drop-down list.

If prompted, specify credentials for the target server.



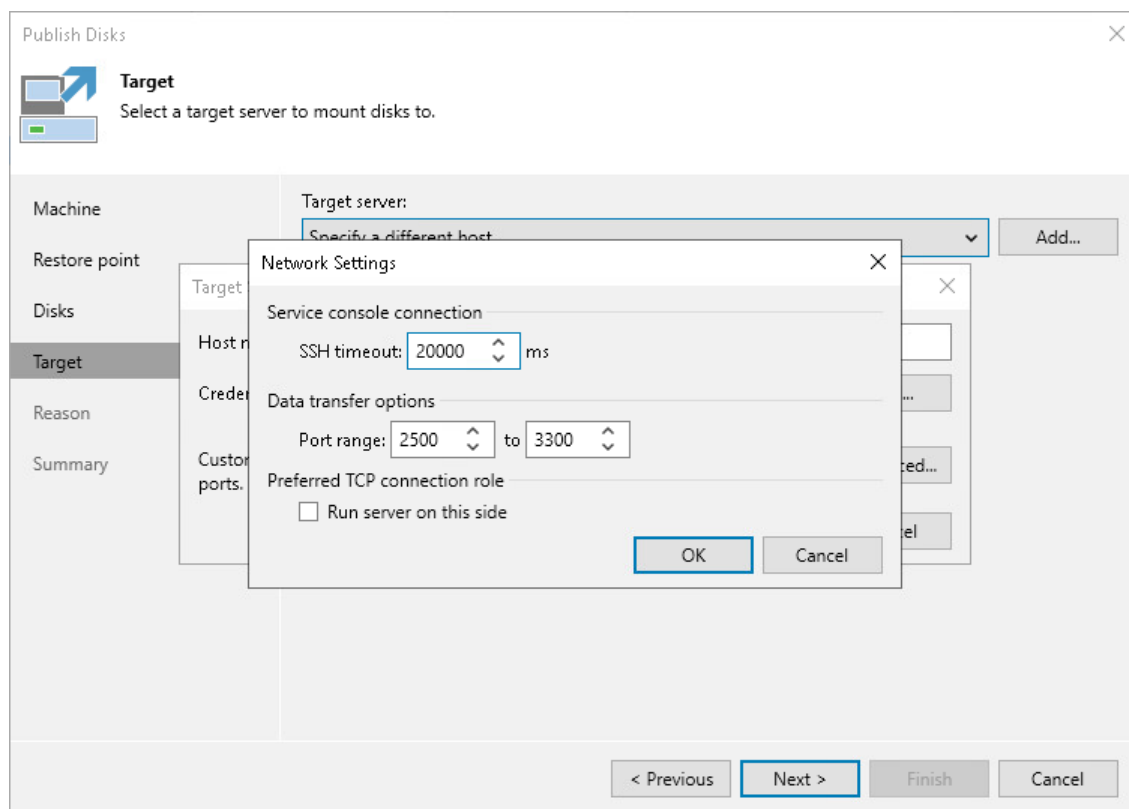
The screenshot shows the 'Publish Disks' wizard window with the 'Target' step selected. The window has a sidebar on the left with options: Machine, Restore point, Disks, Target (selected), Reason, and Summary. The main area is titled 'Target' and contains the instruction 'Select a target server to mount disks to.' Below this, there is a 'Target server:' label followed by a dropdown menu showing 'VBR12-01.tech.local (temporary added host)' and an 'Add...' button. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

## Customizing Connection Settings

If necessary, you can customize connection settings for a target Linux server at the **Target** step of the **Publish Disks** wizard. To do so, click **Advanced** in the **Target Server** window and specify settings in the **Network Settings** window:

1. In the **Service console connection** section, specify an SSH timeout.
2. In the **Data transfer options** section, specify connection settings for file copy operations.
3. [For Linux server deployed outside NAT] In the **Preferred TCP connection role** section, select the **Run server on this side** check box.

To learn more about these settings, see the [Specify Credentials and SSH Settings](#) section in the Veeam Backup & Replication User Guide.

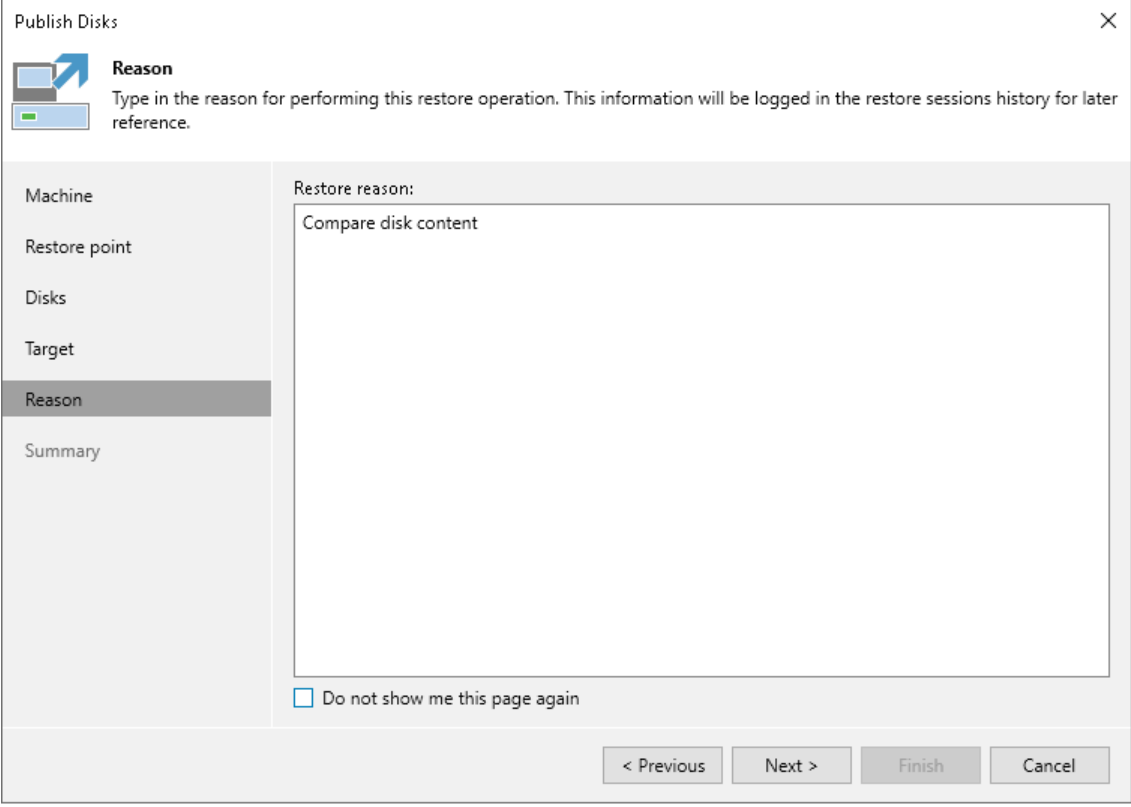


## Step 6. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for publishing disks.

### TIP

If you do not want to show this page, select the **Do not show me this page again** check box. If you further will want to return this page, follow the instructions described in [this Veeam KB article](#).



The screenshot shows the 'Publish Disks' wizard window. The title bar says 'Publish Disks' with a close button. On the left is a sidebar with a tree view containing: Machine, Restore point, Disks, Target, Reason (selected), and Summary. The main area has a header 'Reason' with a sub-header 'Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.' Below this is a large text box labeled 'Restore reason:' containing the text 'Compare disk content'. At the bottom left of the main area is a checkbox labeled 'Do not show me this page again'. At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Publish Disks

**Reason**  
Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

Machine  
Restore point  
Disks  
Target  
**Reason**  
Summary

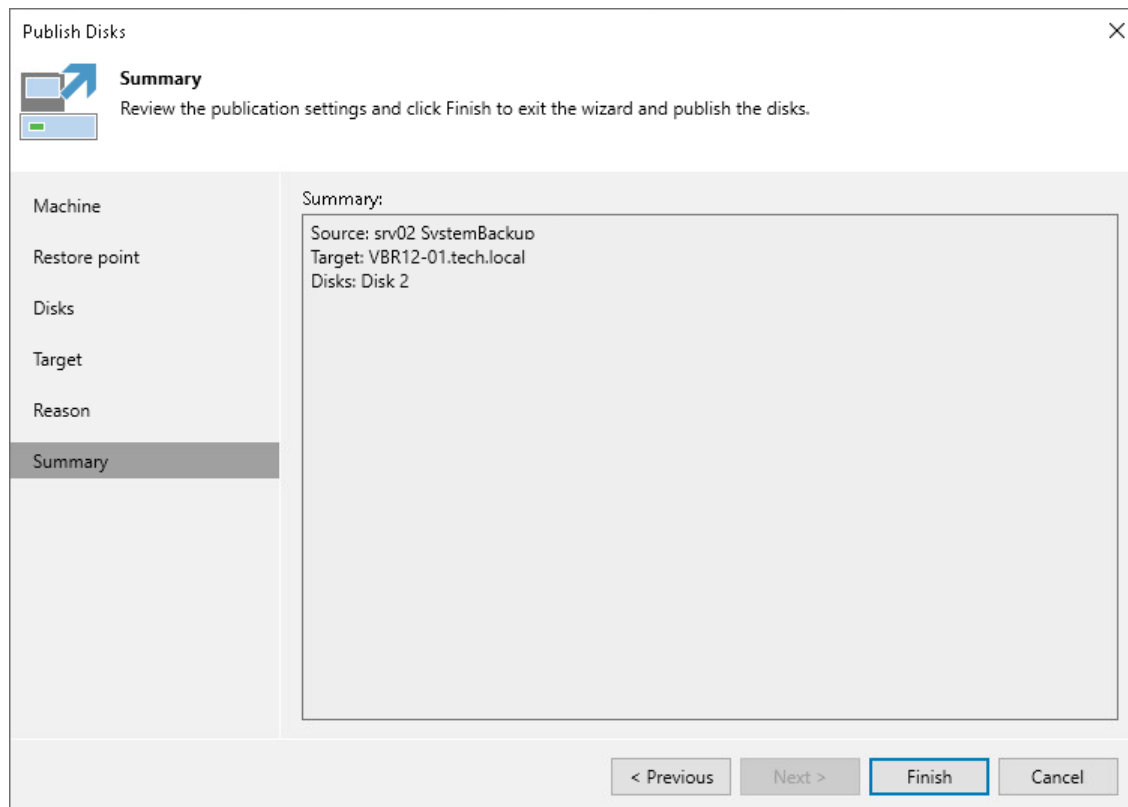
Restore reason:  
Compare disk content

☐ Do not show me this page again

< Previous   Next >   Finish   Cancel

## Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.



## What You Do Next

After the disks are published, go to the following locations on the target server to browse disks content:

- Go to the `/tmp/Veeam.Mount.Disks` location to browse disks images.
- Go to the `/tmp/Veeam.Mount.FS` location to browse disks content.

After you started a disks publishing session, you can view the session statistics or stop the session from the Veeam backup console. To learn more, see [Managing Publishing Disks Session](#).

## Managing Publishing Disks Session

After you started a publishing session, you can check details about the session or stop it.

### Viewing Statistics on Publishing Session

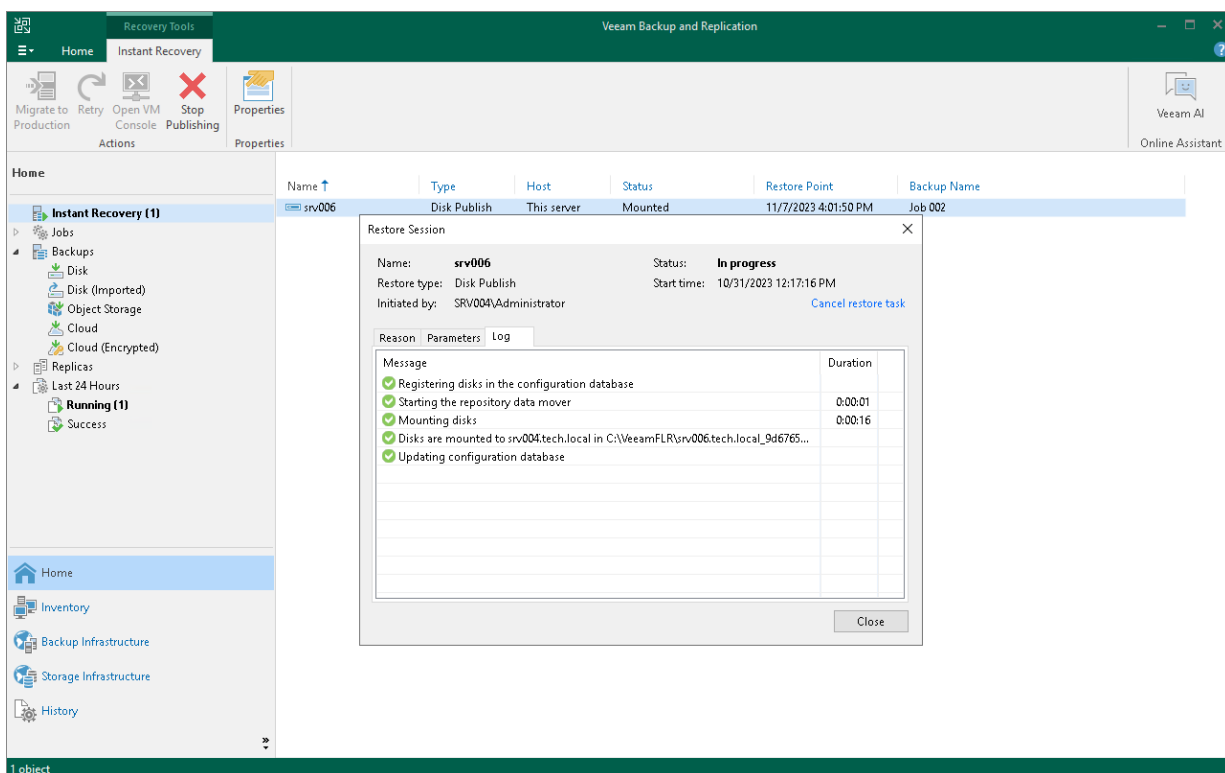
To view publishing session statistics, do one of the following:

- Open the **Home** view. In the inventory pane, select **Instant Recovery**. In the working area, select the necessary publishing session and click **Properties** on the ribbon. Alternatively, right-click the session and **Properties**.
- Open the **Home** view. In the inventory pane select **Last 24 hours**. In the working area, double-click the necessary publishing session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

- Open the **History** view. In the inventory pane select **Restore**. In the working area, double-click the necessary publishing session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

The publishing statistics provides the following data:

- At the top of the **Restore Session** window, Veeam Backup & Replication shows general session statistics. It includes a name of the Veeam Agent computer whose disk you want to publish, a name of the backup server which initiated the publishing session, a user name of the account under which the session was started, session status and duration details.
- The **Reason** tab shows the reason for the publishing session.
- The **Parameters** tab shows information about the target server, the Veeam Agent computer whose disks you publish and the restore point selected for publishing.
- The **Log** tab shows the list of operations performed during the session.




## Stopping Publishing Session

To stop a publishing session, do one of the following:

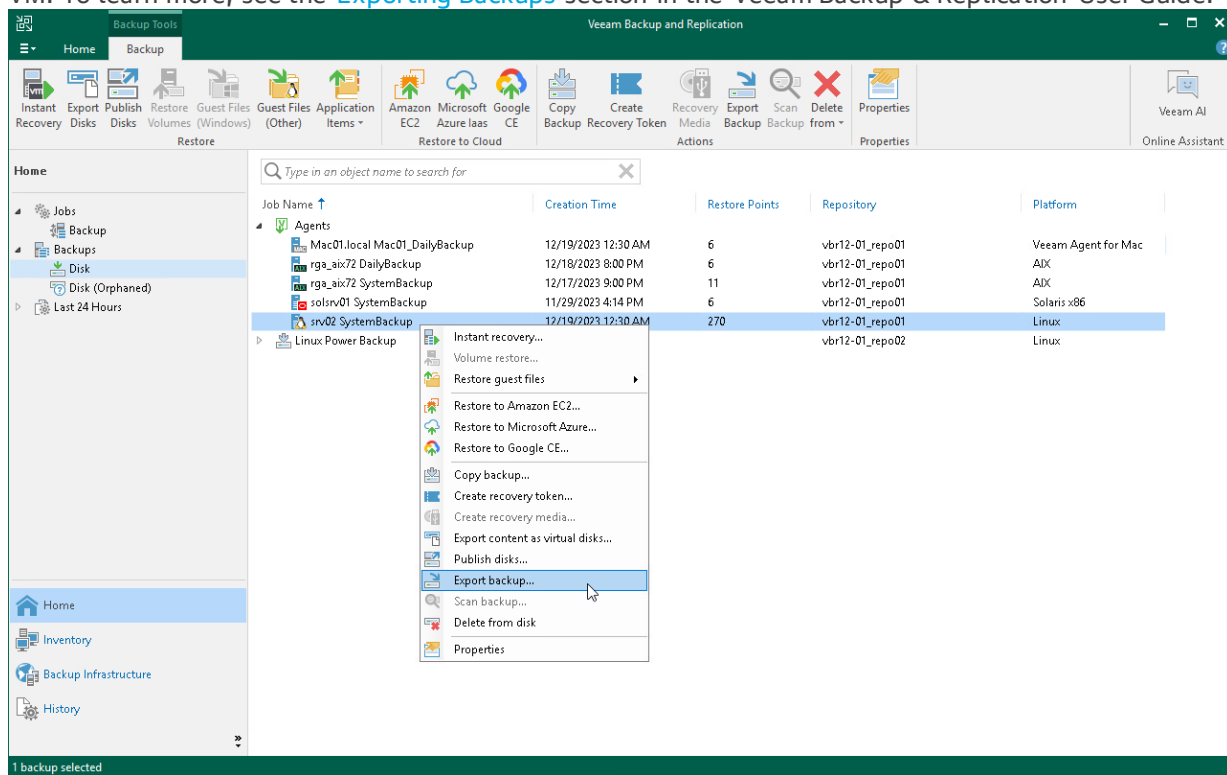
- Open the **Home** view. In the inventory pane select **Instant Recovery**. In the working area, double-click the necessary publishing session and click **Cancel restore task** in the **Restore Session** window. Alternatively, you can select the necessary publishing session and click **Stop Publishing** on the ribbon or right-click the session and click **Stop Publishing**.
- Open the **Home** view. In the inventory pane select **Last 24 hours**. In the working area, double-click the necessary publishing session and click **Cancel restore task** in the **Restore Session** window. Alternatively, you can select the necessary publishing session and click **Stop** on the ribbon or right-click the session and click **Stop session**.

- Open the **History** view. In the inventory pane select **Restore**. In the working area, select the necessary publishing session and double-click it. In the **Restore Session** window, click **Cancel restore task**. Alternatively, you can right-click the publishing session and click **Stop session**.



# Exporting Restore Point to Full Backup File

You can restore data from a specific restore point in a Veeam Agent backup and export this data to a standalone full backup file. The procedure of Veeam Agent backup export does not differ from the same procedure for a VM. To learn more, see the [Exporting Backups](#) section in the Veeam Backup & Replication User Guide.





# Performing Administration Tasks

You can manage Veeam Agent backup jobs and backups created with these jobs. Veeam Backup & Replication allows you to perform the following administration tasks:

- [Import Veeam Agent backups.](#)
- [Enable and disable Veeam Agent backup jobs.](#)
- [Delete Veeam Agent backup jobs.](#)
- [View Veeam Agent backup properties.](#)
- [Create recovery token.](#)
- [Remove Veeam Agent backups.](#)
- [Delete Veeam Agent backups.](#)
- [Configure global settings.](#)
- [Assign roles to users.](#)

# Importing Veeam Agent Backups

You may need to import a Veeam Agent backup in the Veeam Backup & Replication console in the following situations:

- The Veeam Agent backup is stored on a drive managed by another computer (not the Veeam backup server).
- The Veeam Agent backup is stored in a backup repository managed by another Veeam backup server.
- The Veeam Agent backup has been removed in the Veeam Backup & Replication console.

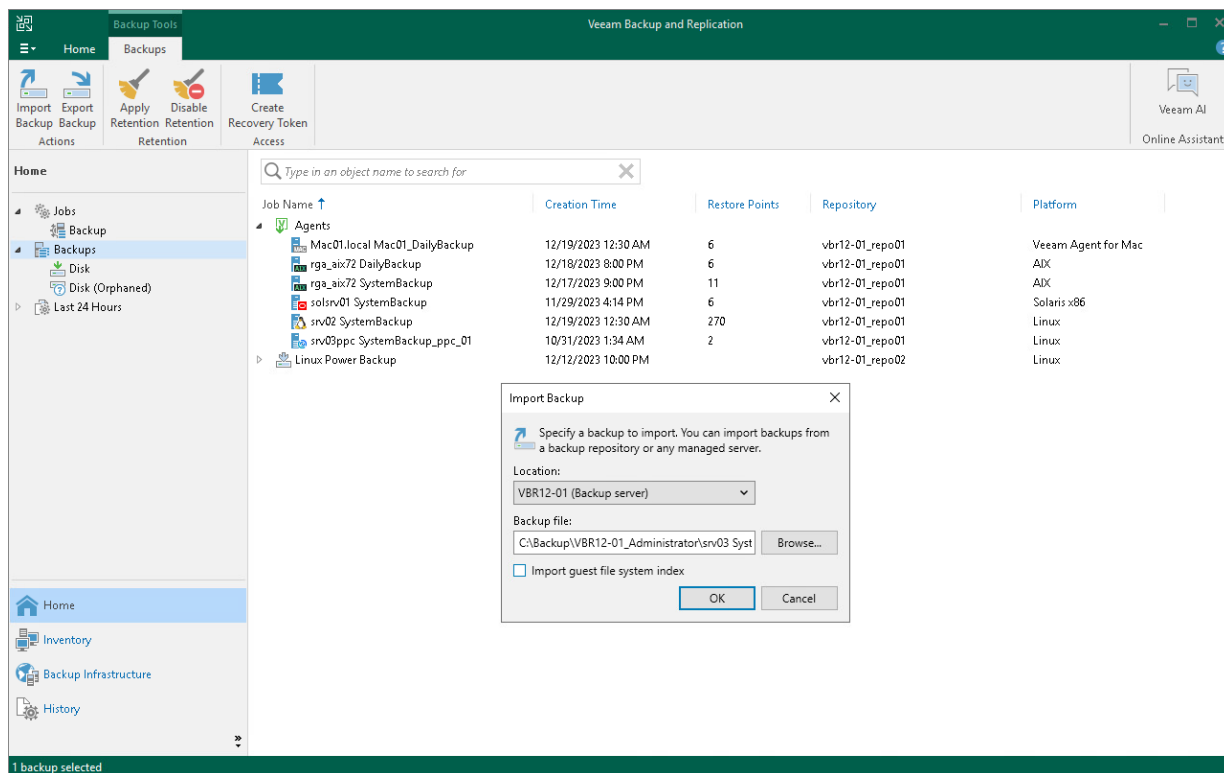
After importing, the Veeam Agent backup becomes available in the Veeam Backup & Replication console. You can restore data from such backup in a regular manner.

Before importing a backup, check the following prerequisites:

- The computer or server from which you plan to import the backup must be added to Veeam Backup & Replication. Otherwise you will not be able to access backup files.
- To be able to restore data from previous backup restore points, make sure that you have all incremental restore points in the same folder where the full backup file resides.

To import a Veeam Agent backup:

1. In Veeam Backup & Replication, click **Import Backup** on the **Home** tab.
2. From the **Computer** list, select the computer or server on which the backup you want to import is stored.
3. Click **Browse** and select the necessary VBM or VBK file. If you select the VBM file, the import process will be notably faster. We recommend that you use the VBK files for import only if a corresponding VBM file is not available.
4. Click **OK**. The imported backup will become available in the **Home** view, under the **Backups > Disk (imported)** node in the inventory pane.



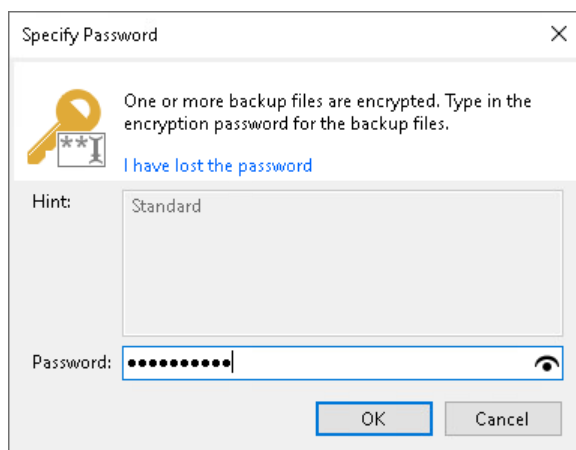
# Importing Encrypted Backups

You can import Veeam Agent backups that were encrypted by Veeam Backup & Replication or Veeam Agent for Microsoft Windows.

To import an encrypted backup file:

1. On the **Home** tab, click **Import Backup**.
2. From the **Computer** list, select the host on which the backup you want to import is stored.
3. Click **Browse** and select the VBM or VBK file.
4. Click **OK**. The encrypted backup will appear under the **Backups > Disk (encrypted)** node in the inventory pane.
5. In the working area, select the imported backup and click **Specify Password** on the ribbon, or right-click the backup and select **Specify password**.
6. In the **Password** field, enter the password for the backup file. If you changed the password one or several times while the backup chain was created, you need to specify the latest password. For Veeam Agent backups, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

If you enter correct password, Veeam Backup & Replication will decrypt the backup file. The backup will be moved under the **Backups > Disk (imported)** node in the inventory pane.



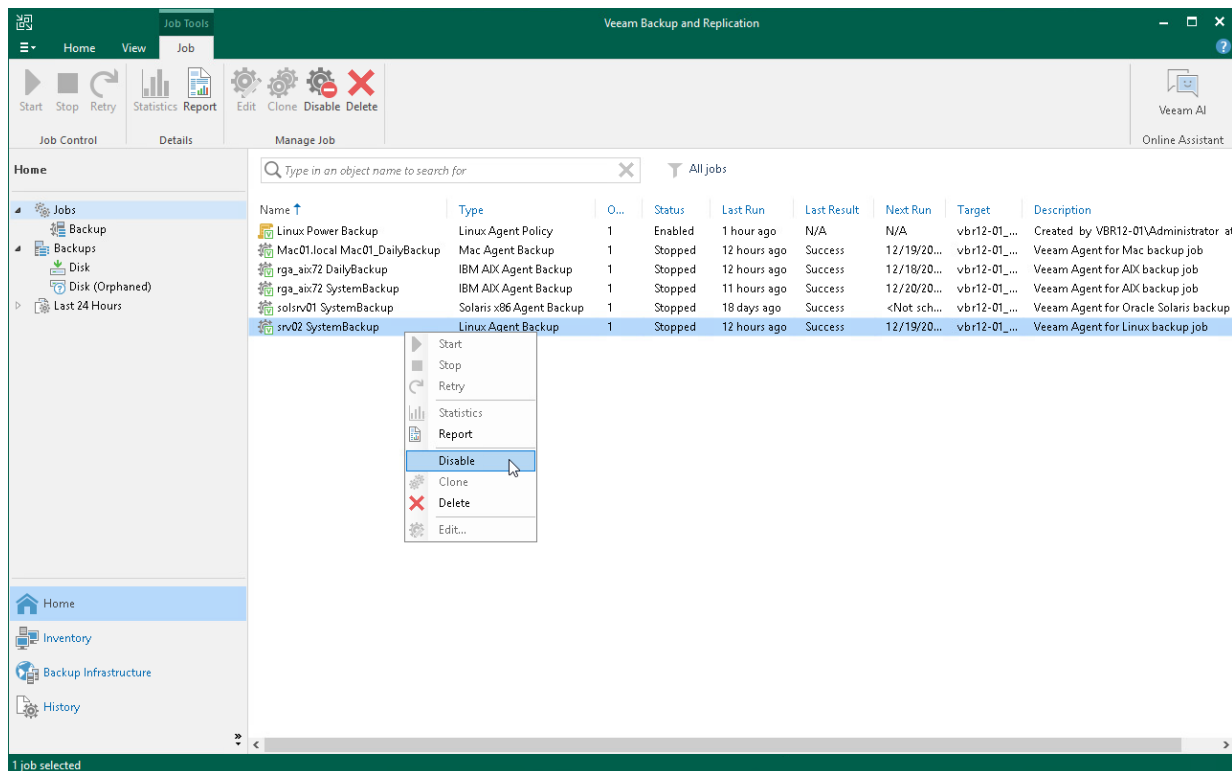
# Enabling and Disabling Veeam Agent Backup Jobs

You can disable and enable Veeam Agent jobs in Veeam Backup & Replication.

When you disable the job, you prohibit the user to store the resulting backup in the backup repository. If the user starts a disabled job manually or the job starts by schedule, the job session will fail and report the "*Job is disabled on backup server*" error. To let Veeam Agent store backups in the backup repository again, you must enable the disabled job.

To disable or enable the scheduled backup job in Veeam Backup & Replication:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Disable** on the ribbon, or right-click the necessary job in the working area and select **Disable**. To enable the disabled job, click **Disable** on the toolbar, or right-click the job and select **Disable** once again.

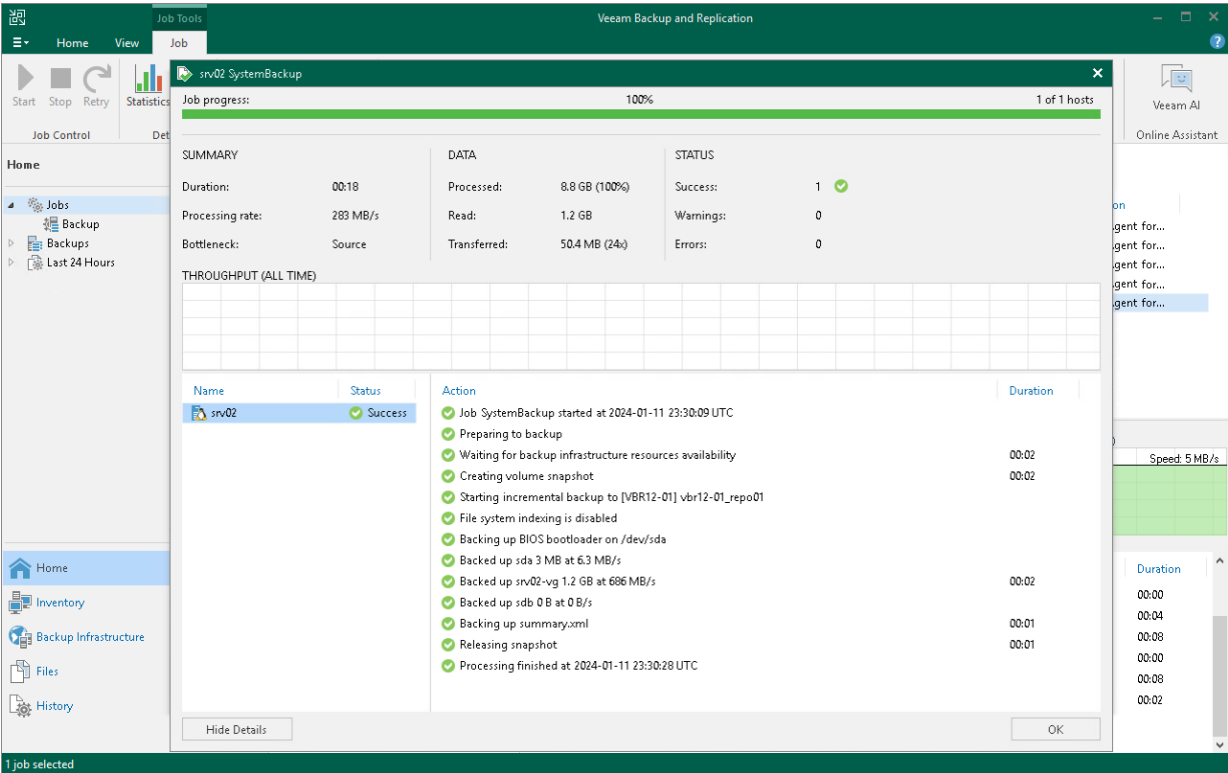


# Viewing Veeam Agent Backup Job Statistics

You can view statistics about Veeam Agent backup jobs in the Veeam Backup & Replication console. Veeam Backup & Replication displays statistics for Veeam Agent backup jobs in the similar way as for regular backup jobs. The difference is that the list of objects included in the job contains a Veeam Agent machine instead of one or several VMs.

To view Veeam Agent backup job statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. In the working area, select the necessary Veeam Agent backup job and click **Statistics** on the ribbon, or right-click the job and select **Statistics**.



# Deleting Veeam Agent Backup Jobs

You can delete Veeam Agent backup jobs.

When you delete a Veeam Agent backup job, Veeam Backup & Replication removes all records about the job from its database and console. When the user starts a new Veeam Agent backup job session manually or the job starts automatically by schedule, the job will appear in the Veeam Backup & Replication console again, and records about a new job session will be stored in the Veeam Backup & Replication database.

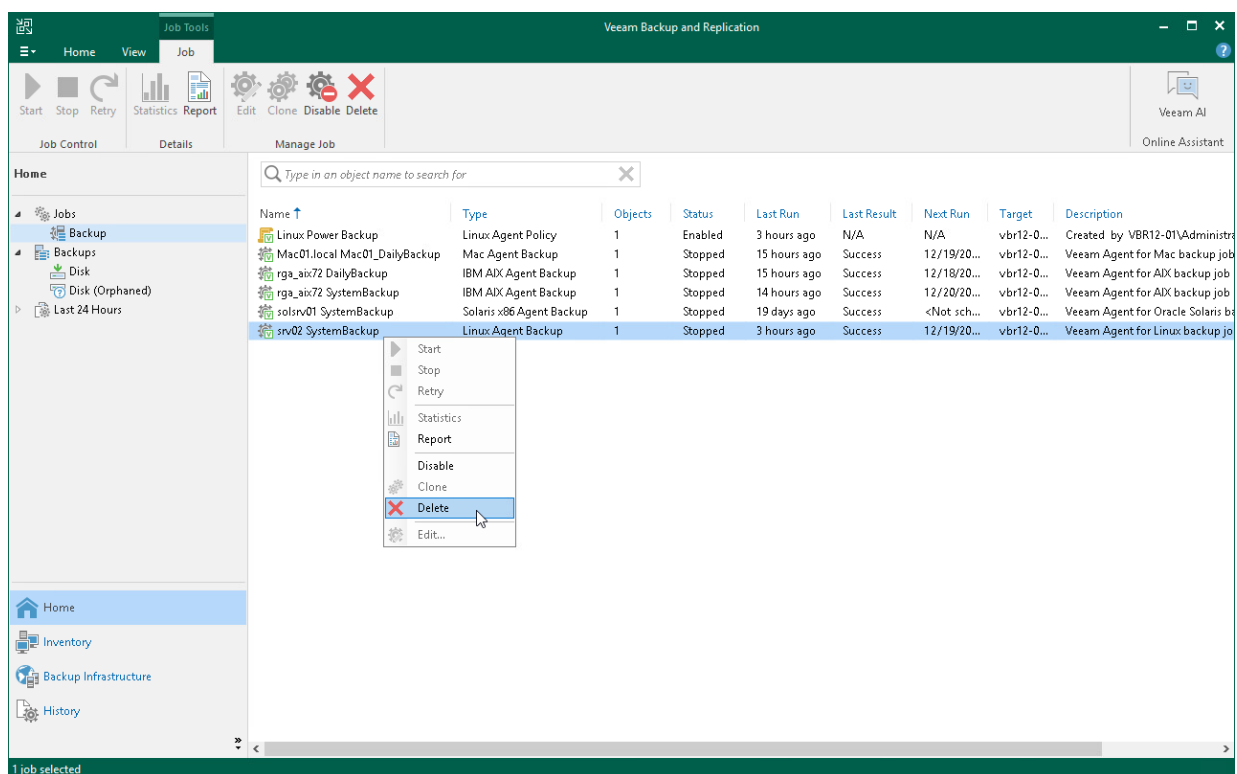
## NOTE

When you delete a Veeam Agent backup job, the backup files become orphaned and can be deleted by the background retention. For more information about the background retention, see the [Background Retention](#) section in the Veeam Backup & Replication User Guide.

To prevent the job from starting permanently, you must delete the job and unassign access rights permissions for this user from the backup repository. To completely delete the job, you must perform this operation in Veeam Agent on the Veeam Agent computer.

To remove a job:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Delete** on the ribbon, or right-click the necessary job in the working area and select **Delete**.



# Viewing Veeam Agent Backup Properties

You can view statistics about Veeam Agent backups.

To view Veeam Agent backup statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click **Disk** under the **Backups** node.
3. In the working area, expand the **Agents** node, select the necessary backup and click **Properties** on the ribbon, or right-click the backup and select **Properties**.

The screenshot shows the Veeam Backup & Replication interface. The 'Home' view is active, and the 'Agents' node is expanded in the left pane. The 'srv02 SystemBackup' is selected. The 'Agent Backup Properties' dialog is open, displaying the following information:

Object	Repository	Owner	Folder
srv02 SystemBackup	vbr12-01_repo01	VBR12-01\Administrator	C:\Backup\VBR12-01\Administrator\srv02 SystemBackup\

Name	Data Size	Backup Size	Date
SystemBackup_2023-12-18T132525.vib	263 MB	89.1 MB	12/18/2023 1:25:25 PM
SystemBackup_2023-12-18T003036.vib	149 MB	24.7 MB	12/18/2023 12:30:36 AM
SystemBackup_2023-12-17T003035.vib	161 MB	29.1 MB	12/17/2023 12:30:35 AM
SystemBackup_2023-12-16T003033.vib	227 MB	69.6 MB	12/16/2023 12:30:33 AM
SystemBackup_2023-12-15T003034.vib	251 MB	83.6 MB	12/15/2023 12:30:34 AM
SystemBackup_2023-12-14T003033.vib	355 MB	184 MB	12/14/2023 12:30:33 AM
SystemBackup_2023-12-13T003034.vib	233 MB	68.6 MB	12/13/2023 12:30:34 AM
SystemBackup_2023-12-12T003031.vib	219 MB	36.9 MB	12/12/2023 12:30:31 AM
SystemBackup_2023-12-11T003032.vib	223 MB	63.8 MB	12/11/2023 12:30:32 AM
SystemBackup_2023-12-10T003031.vib	187 MB	45.0 MB	12/10/2023 12:30:31 AM
SystemBackup_2023-12-09T003030.vib	127 MB	21.6 MB	12/9/2023 12:30:30 AM

Backup size: 8.52 GB

# Creating Recovery Token

If you want to recover volumes or an entire computer protected with Veeam Agent, you can use the **Create recovery token** operation.

You can generate the recovery token on the Veeam Backup & Replication side. Then, on the computer side, with this recovery token get access to the backup and recover data that are stored in the backup. To learn more, see [Veeam Backup Repository Settings](#).

## Considerations and Limitations

Before creating a recovery token, consider the following prerequisites and limitations:

- Recovery tokens stay valid for 24 hours.
- You can recover files and folders from the selected backups only.
- During recovery, Veeam Backup & Replication does not stop backup operations.
- You cannot create a recovery token for backups stored in Veeam Cloud Connect repository.

## Generating Recovery Token

To create a recovery token on the Veeam Backup & Replication side:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. In the working area, right-click the backup and select **Create recovery token**.

You can create a recovery token for several backups. To do this, press and hold [Ctrl], select multiple backups, right-click one of the selected backups and select **Create recovery token**.

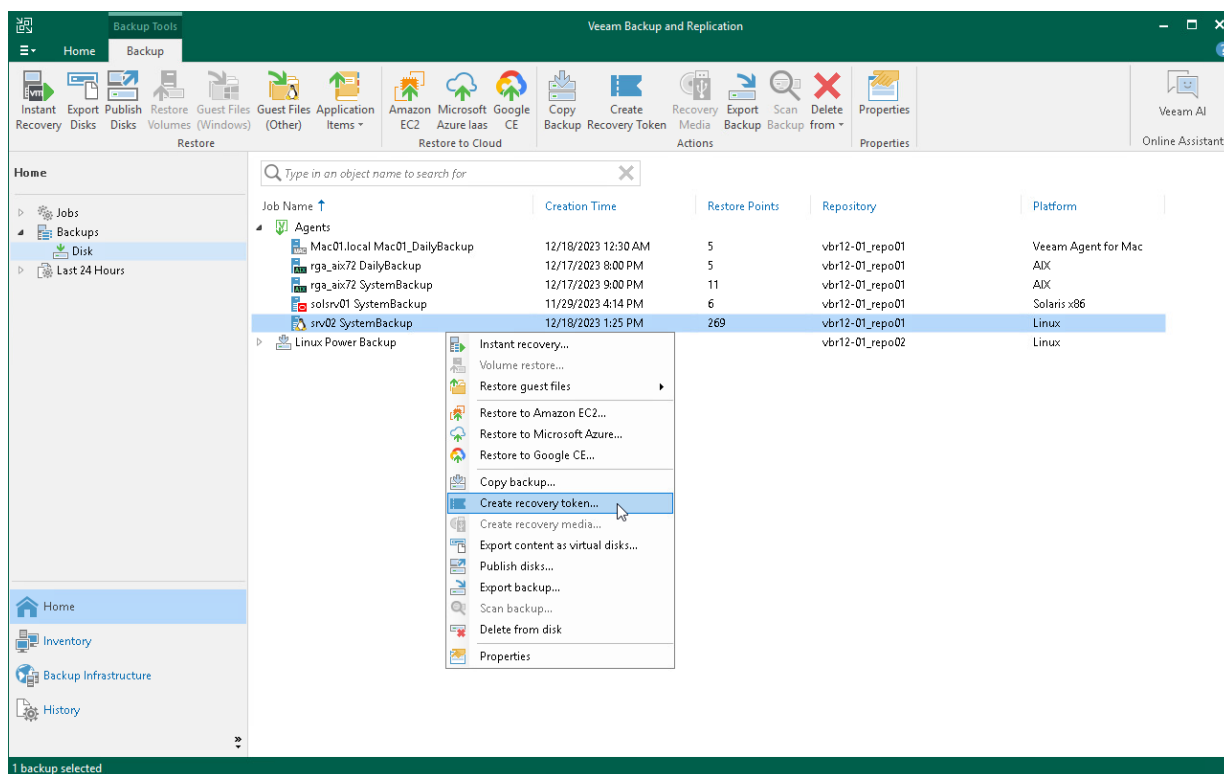
4. In the **Create Recovery Token** window, click **Create**.

You can modify the existing recovery token using the PowerShell console. To learn more, see the [Working with Tokens](#) section in the Veeam PowerShell Reference.



## TIP

Alternatively, you can get access to the backup using user credentials. To learn more, see [Veeam Backup Repository Settings](#).




# Removing Veeam Agent Backups

If you want to remove records about Veeam Agent backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation. When you remove a Veeam Agent backup from configuration, the actual backup files remain in the backup repository. You can import the backup to the Veeam Backup & Replication at any time later and restore data from it.

## IMPORTANT

Removing backups from configuration is designed for experienced users only. Consider using the [Delete from disk](#) operation instead.

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. Press and hold the [Ctrl] key, select the backup, right-click the backup and select **Remove from configuration**.

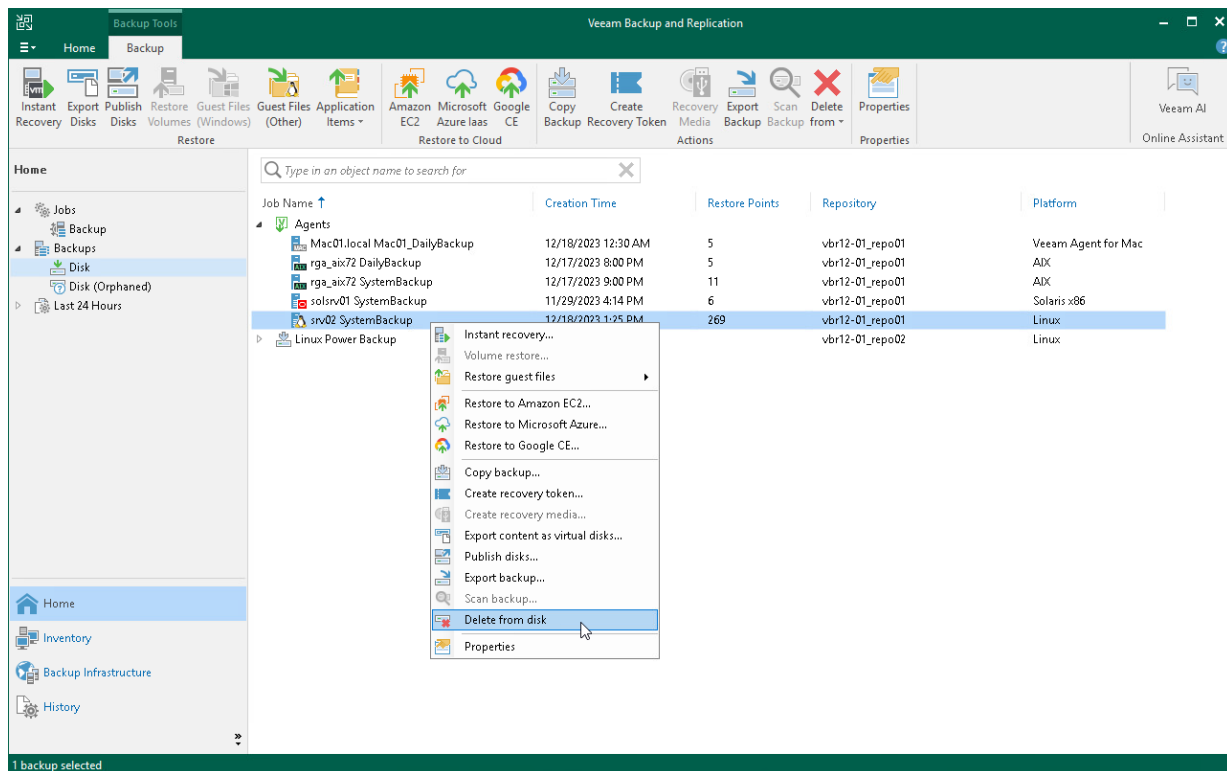


# Deleting Veeam Agent Backups from Disk

If you want to delete records about backups from the Veeam Backup & Replication console and configuration database and, additionally, delete backup files from the backup repository, you can use the **Delete from disk** operation.

To remove a Veeam Agent backup from the backup repository:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. Select the necessary computer backup and click **Delete from > Disk** on the ribbon or right-click the computer and select **Delete from disk**.



# Configuring Global Settings

Global settings configured on the Veeam backup server apply to Veeam Agent backup jobs as well. You can:

- Configure network throttling settings so that Veeam Agent backup job does not consume all network resources. To learn more, see the [Specifying I/O Settings](#) topic in the Veeam Backup & Replication User Guide.
- Configure the following global notification settings to get alerted about the Veeam Agent backup job results:
  - Email notifications. To learn more, see the [Specifying Email Notification Settings](#) section in the Veeam Backup & Replication User Guide.
  - SNMP notifications. To learn more, see the [Specifying SNMP Settings](#) section in the Veeam Backup & Replication User Guide.

# Assigning Roles to Users

User roles configured on the Veeam backup server apply to Veeam Agent backup jobs as well.

To learn more, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.



# Veeam Agent for Microsoft Windows

---

Version 6

User Guide

October, 2024

© 2024 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

<b>CONTACTING VEEAM SOFTWARE .....</b>	<b>9</b>
<b>ABOUT THIS DOCUMENT .....</b>	<b>10</b>
<b>OVERVIEW .....</b>	<b>11</b>
Solution Architecture .....	12
Standalone and Managed Operation Modes .....	13
Data Backup .....	15
Backup Types .....	16
How Backup Works .....	26
Backup Job .....	28
Backup Job Schedule .....	29
Ad-Hoc Backup .....	36
Parallel Disk Processing .....	39
Backup Chain .....	40
Backup to Object Storage .....	56
Backup to Deduplicating Storage Appliances .....	61
Health Check for Backup Files .....	63
Changed Block Tracking .....	70
Data Compression .....	75
Guest Processing .....	77
Data Encryption .....	91
Backup Cache .....	102
Backup to Rotated Drives .....	106
Backup of External Drives .....	109
Data Restore .....	110
Volume-Level Restore .....	111
File-Level Restore .....	115
Restore from Encrypted Backups .....	116
Veeam Recovery Media .....	117
Drivers in Veeam Recovery Media .....	118
BitLocker Encrypted Volumes Support .....	119
Integration with Veeam Backup & Replication .....	124
Backup to Veeam Cloud Connect Repository .....	126
<b>MANAGING VEEAM AGENT IN VEEAM BACKUP &amp; REPLICATION .....</b>	<b>127</b>
<b>PLANNING AND PREPARATION .....</b>	<b>128</b>
System Requirements .....	129
Permissions .....	133
Ports .....	136



<b>INSTALLATION AND CONFIGURATION .....</b>	<b>142</b>
Before You Begin .....	143
Installing Veeam Agent for Microsoft Windows .....	144
Installing Veeam Agent for Microsoft Windows in Unattended Mode .....	145
Using Sysprep and Veeam Agent for Microsoft Windows .....	147
Upgrading Veeam Agent for Microsoft Windows.....	148
Configuring Advanced Settings.....	150
Uninstalling Veeam Agent for Microsoft Windows .....	151
What You Do Next .....	152
<b>GETTING STARTED .....</b>	<b>153</b>
<b>LICENSING .....</b>	<b>154</b>
Product Editions .....	155
License Agreements.....	156
Installing License .....	157
Selecting Product Edition .....	159
Revoking License .....	160
<b>PERFORMING BACKUP .....</b>	<b>161</b>
Creating Veeam Recovery Media .....	162
Before You Begin .....	163
Step 1. Launch Create Recovery Media Wizard .....	165
Step 2. Specify Recovery Media Options .....	167
Step 3. Specify Path to ISO .....	169
Step 4. Review Recovery Image Settings .....	170
What You Do Next.....	172
Creating Veeam Recovery Media with Command Line Interface .....	173
Creating Backup Jobs .....	174
Before You Begin .....	175
Step 1. Launch New Backup Job Wizard .....	176
Step 2. Specify Job Name and Description .....	177
Step 3. Select Backup Mode .....	178
Step 4. Specify Backup Scope Settings .....	180
Step 5. Select Backup Destination .....	187
Step 6. Specify Backup Storage Settings .....	189
Step 7. Specify GFS Retention Policy .....	220
Step 8. Specify Advanced Backup Settings .....	222
Step 9. Specify Backup Cache Settings .....	228
Step 10. Specify Guest Processing Settings.....	229
Step 11. Specify Backup Schedule .....	238
Step 12. Review Backup Job Settings.....	243
What You Do Next.....	244

Managing Backup Jobs .....	245
Editing Backup Job Settings .....	246
Disabling and Enabling Scheduled Backups .....	247
Stopping Backup Job .....	249
Removing Backup Job .....	250
Controlling Backup Post-Job Action.....	251
Performing Ad-Hoc Backups.....	252
Creating Incremental Backups .....	253
Creating Active Full Backups .....	255
Creating Standalone Full Backups .....	257
Performing Backup to Another Location .....	259
Performing Backup with Command Line Interface .....	261
Deleting Backups.....	265
Managing Backup Cache.....	267
Monitoring Backup Cache Activity .....	268
Pausing Backup Cache Synchronization .....	271
Deleting Restore Points from Backup Cache .....	272
<b>PERFORMING RESTORE.....</b>	<b>273</b>
Restoring from Veeam Recovery Media .....	274
Before You Begin .....	275
Step 1. Boot from Veeam Recovery Media.....	276
Step 2. Select Network Adapter or Wireless Network .....	278
Step 3. Launch Veeam Recovery Media Wizard.....	280
Step 4. Specify Backup File Location .....	281
Step 5. Select Network Storage Type .....	283
Step 6. Specify Network Storage Settings .....	284
Step 7. Select Backup .....	298
Step 8. Select Restore Point .....	300
Step 9. Select Data Restore Mode .....	301
Step 10. Map Restored Disks.....	302
Step 11. Resize Restored Volumes .....	306
Step 12. Start Restore Process .....	307
Using Veeam Agent and Microsoft Windows Tools .....	308
Using Microsoft Windows Recovery Environment .....	310
Restoring Volumes.....	311
Before You Begin .....	312
Step 1. Launch Volume Level Restore Wizard .....	313
Step 2. Specify Backup File Location .....	314
Step 3. Select Network Storage Type .....	315
Step 4. Specify Network Storage Settings .....	316

Step 5. Select Backup .....	329
Step 6. Select Restore Point .....	331
Step 7. Map Restored Disks .....	332
Step 8. Resize Restored Volumes .....	336
Step 9. Complete Restore Process .....	337
Restoring Files and Folders .....	338
Before You Begin .....	339
Step 1. Launch File Level Restore Wizard .....	340
Step 2. Specify Backup File Location .....	341
Step 3. Select Remote Storage Type .....	342
Step 4. Specify Remote Storage Settings .....	343
Step 5. Select Backup .....	356
Step 6. Select Restore Point .....	358
Step 7. Complete Restore Process .....	359
Step 8. Save Restored Files .....	360
Restoring Data from Encrypted Backups .....	367
<b>REPORTING .....</b>	<b>369</b>
Viewing Statistics in Control Panel .....	370
Viewing Statistics for Separate Restore Points .....	373
Viewing Information About Job Retries .....	376
Viewing Status of Restore Points in Backup Cache .....	378
Monitoring Backup State with Tray Agent .....	379
Monitoring Backup Process in Taskbar Button .....	381
Viewing and Dismissing Veeam Agent Events .....	382
Viewing Events with Windows Notification Center .....	385
Viewing Job Session Results in Email Reports .....	387
<b>SPECIFYING SETTINGS .....</b>	<b>388</b>
Throttling Backup Activities .....	389
Restricting Network Connections Usage .....	390
Limiting Bandwidth Consumption .....	391
Disabling Backup over Metered Connections .....	392
Disabling Backup over VPN Connections .....	394
Selecting Wireless Networks for Backup .....	396
Managing Rotated Drives .....	398
Disabling Control Panel Notifications .....	399
Enabling Email Notifications .....	400
Custom SMTP Server Settings .....	401
Gmail Server Settings .....	403
Microsoft 365 Server Settings .....	406
Checking for New Product Versions and Updates .....	410

<b>MANAGING VEEAM CBT DRIVER.....</b>	<b>411</b>
Installing Veeam CBT Driver .....	412
Removing Veeam CBT Driver .....	414
Removing CBT Driver with Veeam Recovery Media .....	415
Resetting CBT.....	416
<b>GETTING SUPPORT .....</b>	<b>417</b>
Reporting Issues .....	418
<b>USING WITH VEEAM BACKUP &amp; REPLICATION.....</b>	<b>420</b>
Setting Up User Permissions on Backup Repositories .....	422
Managing License .....	425
Managing Instance Consumption by Veeam Agents.....	426
Assigning License to Veeam Agent.....	427
Viewing Licensed Veeam Agents and Revoking License .....	428
Performing Data Protection Tasks.....	430
Backing Up to Backup Repositories .....	431
Backing Up to Cloud Repositories .....	432
Performing Backup Copy for Veeam Agent Backups .....	434
Using SureBackup .....	436
Archiving Veeam Agent Backups to Tape.....	437
Scanning Backup.....	438
Restoring Data from Veeam Agent Backups.....	440
Restoring Veeam Agent Backup to vSphere VM .....	441
Restoring Veeam Agent Backup to Hyper-V VM .....	443
Restoring Veeam Agent Backup to Nutanix VM .....	445
Restoring Veeam Agent Backup to Proxmox VM .....	446
Restoring to Microsoft Azure .....	447
Restoring to Amazon EC2 .....	448
Restoring to Google Compute Engine .....	449
Restoring Volumes.....	450
Restoring Files and Folders.....	461
Restoring Application Items.....	480
Exporting Disks.....	481
Publishing Disks.....	491
Exporting Restore Point to Full Backup File .....	501
Performing Administration Tasks .....	502
Importing Veeam Agent Backups .....	503
Enabling and Disabling Veeam Agent Backup Jobs .....	505
Viewing Veeam Agent Backup Job Statistics .....	506
Deleting Veeam Agent Backup Jobs .....	507
Viewing Veeam Agent Backup Properties .....	508

Creating Recovery Token .....	509
Copying Veeam Agent Backups .....	511
Removing Veeam Agent Backups .....	512
Deleting Veeam Agent Backups from Disk .....	513
Configuring Global Settings .....	514
Assigning Roles to Users .....	515
<b>AUTOMATING VEEAM AGENT FOR WINDOWS OPERATIONS .....</b>	<b>516</b>
<b>APPENDIX A. VEEAM AGENT EVENTS.....</b>	<b>517</b>
<b>APPENDIX B. MOVING VEEAM AGENT BACKUPS.....</b>	<b>521</b>
Moving Veeam Agent Backups to Veeam Backup Repository .....	522
Moving Veeam Agent Backups to Veeam Cloud Connect Repository .....	525

# 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](http://veeam.com/documentation-guides-datasheets.html)
- Veeam R&D Forums: [forums.veeam.com](http://forums.veeam.com)

# About This Document

This user guide provides information about Veeam Agent for Microsoft Windows.

## Intended Audience

The user guide is intended for anyone who wants to use Veeam Agent for Microsoft Windows 6.2 to protect their computer.

# Overview

Veeam Agent for Microsoft Windows is a data protection and disaster recovery solution for physical and virtual machines. Veeam Agent for Microsoft Windows can be used to protect different types of computers and devices: desktops, laptops and tablets. The solution can be installed on Windows-based workstations, physical servers and virtual machines. To learn about supported OSes, see [System Requirements](#).

## NOTE

Veeam Agent for Microsoft Windows can operate in either standalone or managed mode. Depending on the mode, Veeam Agent provides different features and limitations. To learn more, see [Standalone and Managed Operation Modes](#).

Veeam Agent for Microsoft Windows offers a variety of features to protect your data. You can:

- Create a Veeam Recovery Media on an external hard drive, USB flash drive, CD/DVD/BD, or create an ISO file with the Veeam Recovery Media on disk.
- Create an entire system image backup, back up specific computer volumes or individual folders with files. Backups can be stored on an external hard drive, in a network shared folder, in object storage, in a Veeam backup repository or Veeam Cloud Connect repository.

In case of a disaster, you can perform the following restore operations:

- Start the OS from the Veeam Recovery Media and use Veeam Agent for Microsoft Windows and standard Microsoft Windows tools to diagnose and fix problems.
- Perform bare metal restore.
- Restore necessary data from backups to its original location or a new location.

Veeam Agent for Microsoft Windows integrates with Veeam Backup & Replication. Backup administrators who work with Veeam Backup & Replication can perform advanced tasks with Veeam Agent backups: perform data restore tasks with Veeam Agent backups, manage Veeam Agent backup jobs or backups created with these jobs.



# Solution Architecture

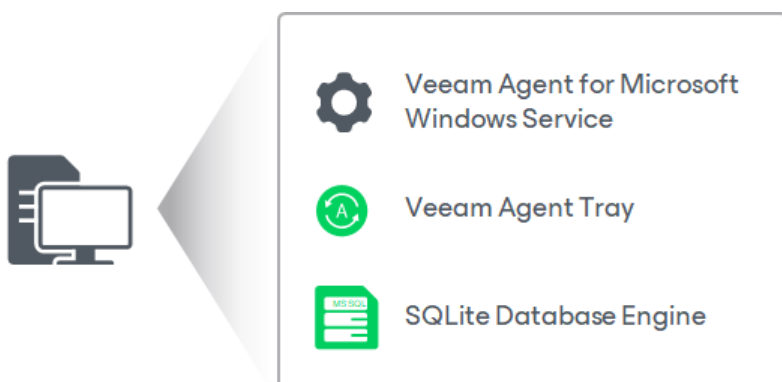
Veeam Agent for Microsoft Windows is set up on a computer whose data you want to protect.

Veeam Agent for Microsoft Windows has a one-service architecture. When you install the product, Veeam Agent deploys the following components on the computer:

- *Veeam Agent for Microsoft Windows Service* is a Microsoft Windows service responsible for performing all types of backup and restore tasks. The service is started automatically when you power on the computer, and runs in the background under the Local System account.
- *Veeam Agent Tray* is a tray agent that communicates with the Veeam Agent for Microsoft Windows Service to let you monitor the backup operation status and provide quick access to main functions of Veeam Agent for Microsoft Windows: starting backup and restore operations, viewing statistics for created backups and so on. The Veeam Agent Tray starts when you log on to the system and runs in the background.
- To store its configuration data, Veeam Agent uses the SQLite database engine. SQLite requires only few files to install and takes little resources to run.

## NOTE

The account under which Veeam Agent for Microsoft Windows Service runs should not be changed. Configurations with custom account are not supported.



# Standalone and Managed Operation Modes

Veeam Agent can operate in two modes: *standalone mode* and *managed mode*. The current User Guide covers subjects related to Veeam Agent operating in the standalone mode only. Depending on the operation mode, Veeam Agent has different functionality and limitations.

## Standalone Mode

In this mode, Veeam Agent operates as a standalone product. To use Veeam Agent operating in the standalone mode, you must manually install the product directly on the computer whose data you want to protect.

For Veeam Agent operating in the standalone mode, data protection, disaster recovery and administration tasks are performed by the user. You can also use Veeam Agent operating in the standalone mode with Veeam Backup & Replication. In this scenario, you can use backup repositories managed by Veeam Backup & Replication as a target location for Veeam Agent backups and use the Veeam Backup & Replication console to perform a number of tasks with Veeam Agent backup jobs and backups. To learn more, see [Integration with Veeam Backup & Replication](#).

You can also use Veeam Backup & Replication as a gateway for creating backups targeted at the following types of repositories:

- Veeam Cloud Connect repository. To learn more, see [Backup to Veeam Cloud Connect](#).
- Object storage repository.

With Veeam Agent operating in the standalone mode, you can also back up data directly to object storage. To learn more about both options, see [Backup to Object Storage](#).

## Managed Mode

In this mode, Veeam Agent operates under control from one of the following Veeam products:

- **Veeam Backup & Replication**

You can automate management of Veeam Agents on multiple computers in your infrastructure in the Veeam Backup & Replication console. You can configure Veeam Agent backup policies and perform other data protection and administration tasks on remote computers.

To use Veeam Agent operating in the managed mode, you must deploy the product in one of the following ways:

- From Veeam Backup & Replication
- Manually using external tools

To learn more about managed Veeam Agent deployment, see the [Protected Computers Discovery and Veeam Agent Deployment](#) section in the Veeam Agent Management User Guide.

For Veeam Agent managed by Veeam Backup & Replication, data protection, data restore and administration tasks are performed by a backup administrator in the Veeam Backup & Replication console. To learn about managing Veeam Agent in Veeam Backup & Replication, see the [Veeam Agent Management Guide](#).

- **Veeam Service Provider Console**

You can use Veeam Service Provider Console to manage Veeam Agents on multiple computers in your infrastructure. When Veeam Agent is managed by Veeam Service Provider Console, you can configure backup job settings, start and stop backup, change global settings, update and uninstall Veeam Agent and collect Veeam Agent data for monitoring and billing.

To manage Veeam Agent from Veeam Service Provider Console, you must install Veeam Service Provider Console management agent and Veeam Agent on the computer whose data you want to protect. After that, in Veeam Service Provider Console, you must activate Veeam Agent on the protected computer to set it into the managed operation mode.

For Veeam Agent managed by Veeam Service Provider Console, data protection, data restore and administration tasks are performed by a backup administrator in Veeam Service Provider Console.

Backup administrator can enable a read-only access mode for Veeam Agent installed on the protected computer. When you work directly with Veeam Agent operating in the read-only access mode, you can perform a limited set of operations, including:

- Running the backup job manually.
- Viewing backup session statistics.
- Restoring individual files.

To learn about deploying and managing Veeam Agent with Veeam Service Provider Console, see [Veeam Service Provider Console User Guides](#). Select the guide that suits your user role.

# Data Backup

We recommend that you regularly back up data stored on your computer. Backup creates a safety copy of your data. If any kind of disaster strikes, you can restore your data from the backup and be sure that you will not lose the necessary information.

You can set up Veeam Agent for Microsoft Windows to perform automatic scheduled backups (triggered at specific time of the day or on specific events), or you can choose to back up data manually when needed. You can back up the entire computer image, specific computer volumes or individual folders with files.

Backups created with Veeam Agent for Microsoft Windows can be saved to one of the following locations:

- Removable storage device
- Local computer drive
- Network shared folder
- On-premises or cloud-based object storage
- Backup repository managed by a Veeam backup server
- Cloud repository managed by a Veeam Cloud Connect service provider

# Backup Types

Veeam Agent for Microsoft Windows lets you create the following backup types:

- [Volume-level backup](#)
- [File-level backup](#)

In the file-level backup, you can select what personal data to back up. To learn more, see [Personal Data Backup](#).

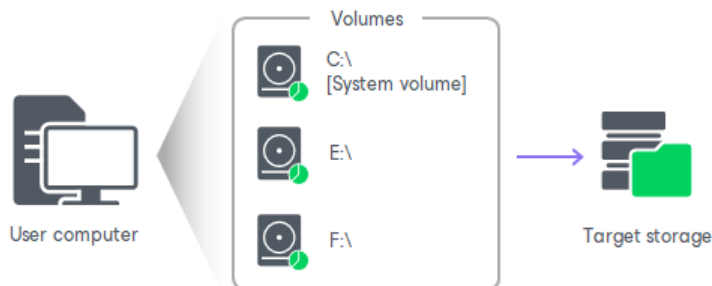
When backing up the operating system data, you can select the desired backup type. To learn more, see [System State Data Backup](#).

## Volume-Level Backup

You can set up Veeam Agent for Microsoft Windows to create a volume-level backup. The volume-level backup captures the whole image of a data volume (also called logical drive or partition) on your computer. You can use the volume-level backup to restore a computer volume, specific files and folders on the volume or perform bare metal recovery.

You can back up all computer volumes or specific computer volumes.

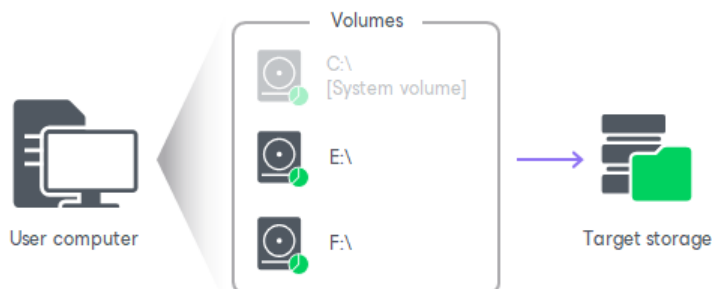
- When you back up the entire computer image, Veeam Agent captures the content of all volumes on your computer. The resulting backup file contains all volume data and Microsoft Windows OS system data: system partition and boot partition. For GPT disks on Microsoft Windows 8.1, 10, 11, 2012, 2012 R2, 2016, 2019 and 2022, Veeam Agent additionally backs up the recovery partition.



- When you back up a specific computer volume, Veeam Agent captures only that data that resides on this specific volume: files, folder, application data and so on.

If you choose to back up the system volume (volume on which Microsoft Windows is installed), Veeam Agent automatically includes the *System Reserved* partition into the backup scope. You can exclude the *System Reserved* partition from the backup if necessary. In this case, Veeam Agent will capture only data on the system volume.

To learn more, see [System State Data Backup](#).

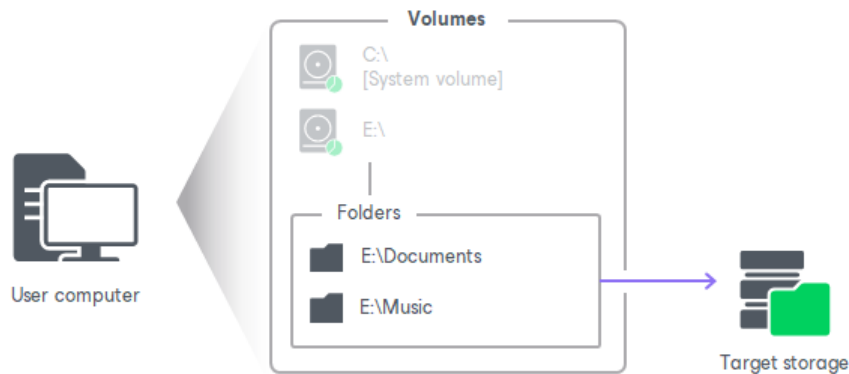


# File-Level Backup

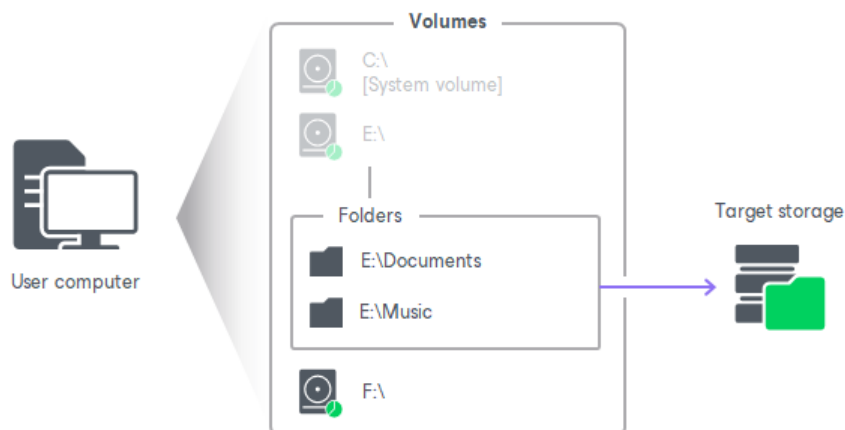
You can set up Veeam Agent for Microsoft Windows to create a file-level backup. The file-level backup captures only data of individual folders on the computer. You can use the file-level backup to restore files and folders that you have added to the backup scope.

Veeam Agent lets you create two types of file-level backups:

- You can include individual folders into the backup. When you recover from such backup, you will be able to restore folders that you have selected to back up, and files in these folders.



- You can create a hybrid backup that will include folders and specific computer volumes. When you recover from such backup, you will be able to restore the following components:
  - For backed-up volume: the entire volume and individual files and folders on this volume.
  - For backed-up folders: folders that you have selected to back up, and files in these folders.

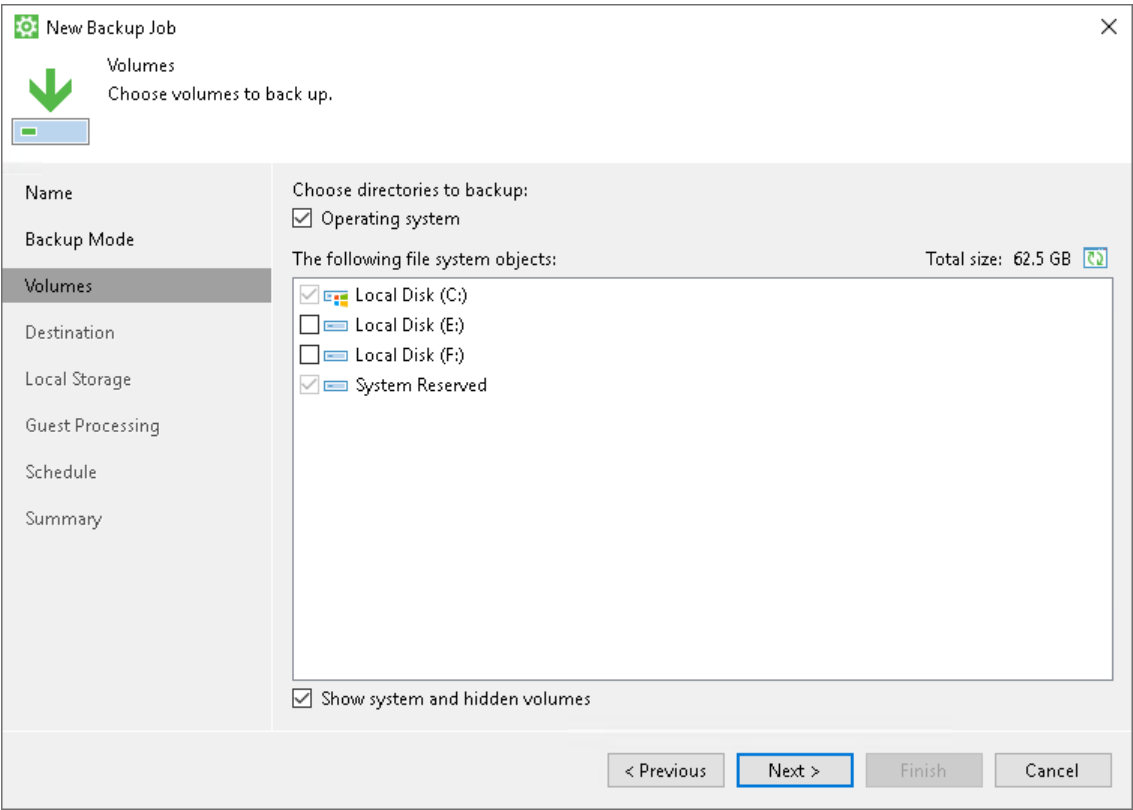


# System State Data Backup

To be able to restore critical components related to the OS and start the OS after recovery, you must include in the backup the system volume (volume on which the OS is installed) and the System Reserved/UEFI or other system partitions. To simplify this process, Veeam Agent for Microsoft Windows offers you to add the *Operating system* component to the backup scope. When you select to back up the operating system data, Veeam Agent automatically includes in the backup all data related to the OS. You can include the operating system data in the volume-level and file-level backup.

# System State Backup in Volume-Level Backup

To include the operating system data in the volume-level backup, select the **Operating system** check box. When you select to back up the operating system data, Veeam Agent automatically includes in the backup the system volume and existing system partitions. If some additional system partitions appear on the Veeam Agent computer in the future, for example, after the OS upgrade, Veeam Agent will add these partitions to the backup scope too.



Alternatively, you can explicitly select to back up the system volume. When you select to back up the system volume, Veeam Agent automatically includes existing system partitions in the backup. However, if additional system partitions appear on the Veeam Agent computer, Veeam Agent may be unable to back up such volumes. Thus, we recommend that you use the *Operating system* option to create system state data backup.

New Backup Job

Volumes

Choose volumes to back up.

Name

Backup Mode

Volumes

Destination

Local Storage

Guest Processing

Schedule

Summary

Choose directories to backup:

☐ Operating system

The following file system objects:

Total size: 62.2 GB

☒ Local Disk (C:)

☐ Local Disk (E:)

☐ Local Disk (F:)

☒ System Reserved

☒ Show system and hidden volumes

< Previous

Next >

Finish

Cancel



# System State Backup in File-Level Backup

To include the operating system data in the file-level backup, select the **Operating system** check box. When you select to back up the operating system data, Veeam Agent automatically includes in the backup all data related to the OS: the system volume and existing system partitions.

**New Backup Job**

**Files**  
Choose individual files and folders to back up. For best performance, when backing up thousands of files, select the entire volume, then uncheck all unnecessary items.

**Name**  
**Backup Mode**  
**Files**  
Destination  
Local Storage  
Guest Processing  
Schedule  
Summary

**Choose directories to backup:**  
☒ Operating system  
☐ Personal files

Include: Desktop, Documents, Pictures, Video, Music, Favorites, Downloads, Other **Choose...**

The following file system objects: **Total size: 62.5 GB**

- ☒ Local Disk (C:) > volume level backup
- ☐ Local Disk (E:)
- ☐ Local Disk (F:)
- ☒ System Reserved

To specify file inclusion and exclusion settings, click **Advanced**

**< Previous** **Next >** **Finish** **Cancel**

Alternatively, you can manually select to back up the system volume and system partitions.

In this case, you will be able to exclude specific folders related to the OS from the backup (for example, the *Users* folder and *Documents* and *Settings* folder). When you select to back up the *Operating system* data, you cannot choose which components related to the OS must be backed up and which must be excluded.

New Backup Job

Files

Choose individual files and folders to back up. For best performance, when backing up thousands of files, select the entire volume, then uncheck all unnecessary items.

Name

Backup Mode

Files

Destination

Local Storage

Guest Processing

Schedule

Summary

Choose directories to backup:

☐ Operating system

☐ Personal files

Include: Desktop, Documents, Pictures, Video, Music, Favorites, Downloads, Other

Choose...

The following file system objects:

Total size: 62.2 GB

>

☒

Local Disk (C:) > volume level backup

>

☐

Local Disk (E:)

>

☐

Local Disk (F:)

>

☒

System Reserved

To specify file inclusion and exclusion settings, click Advanced

Advanced

< Previous

Next >

Finish

Cancel

## Personal Data Backup

You can include personal data in the backup. In the file-level backup, to simplify this process, Veeam Agent for Microsoft Windows offers you to add the *Personal files* component to the backup scope. When you select to back up the personal data, Veeam Agent automatically includes in the backup data related to Veeam Agent computer user profiles.

# Personal Data Backup in File-Level Backup

To include personal data in the file-level backup, select the *Personal files* option. When you select to back up the personal data by default, Veeam Agent includes in the backup all data stored in the *Users* folder on the system volume excluding application data and data related to roaming user profiles. If you store personal folders of user profiles in custom locations, Veeam Agent backs up them too. If you do not want to back up default personal data, you can change the backup scope by selecting what folders Veeam Agent includes in the *Personal files* component.

**New Backup Job**

**Files**  
Choose individual files and folders to back up. For best performance, when backing up thousands of files, select the entire volume, then uncheck all unnecessary items.

**Name**  
**Backup Mode**  
**Files**  
**Destination**  
**Local Storage**  
**Guest Processing**  
**Schedule**  
**Summary**

Choose directories to backup:  
☐ Operating system  
☒ Personal files

Include: Desktop, Documents, Pictures, Video, Music, Favorites, Downloads, Other **Choose...**

The following file system objects: **Total size: 0.0 B**

- > ☐ Local Disk (C:)
- > ☐ Local Disk (E:)
- > ☐ Local Disk (F:)
- ☐ System Reserved

To specify file inclusion and exclusion settings, click **Advanced**

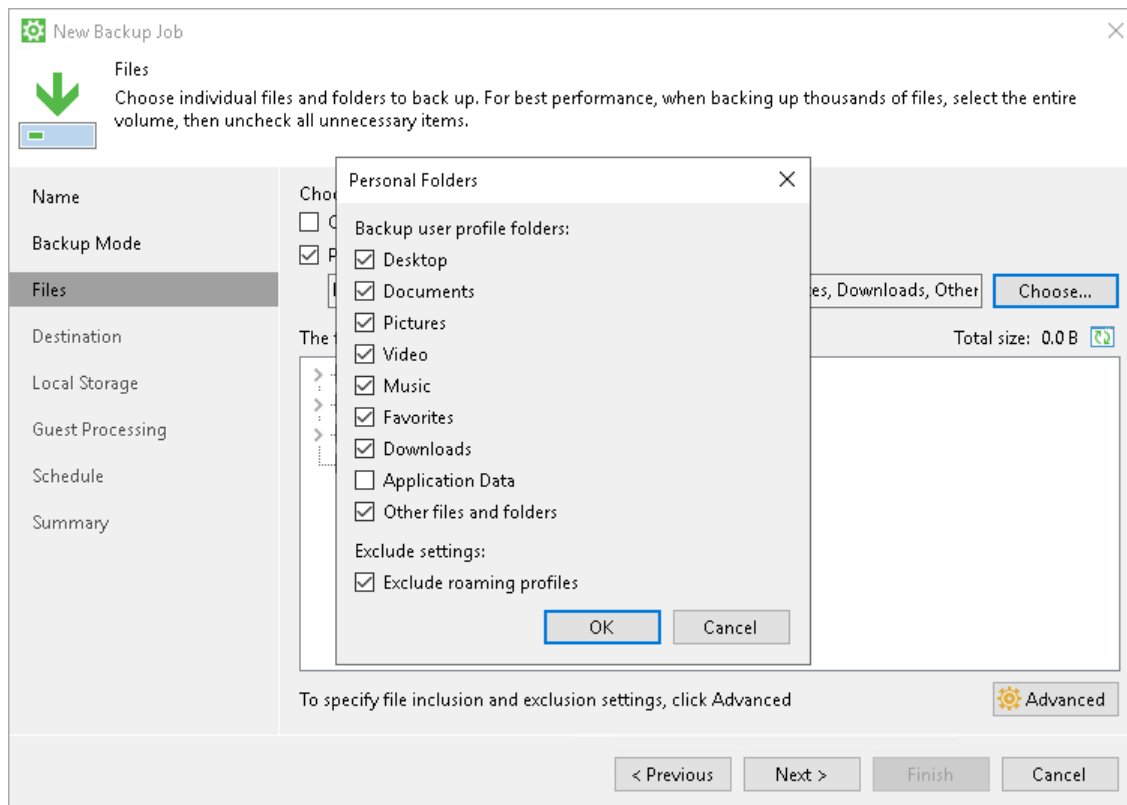
**< Previous** **Next >** **Finish** **Cancel**

## TIP

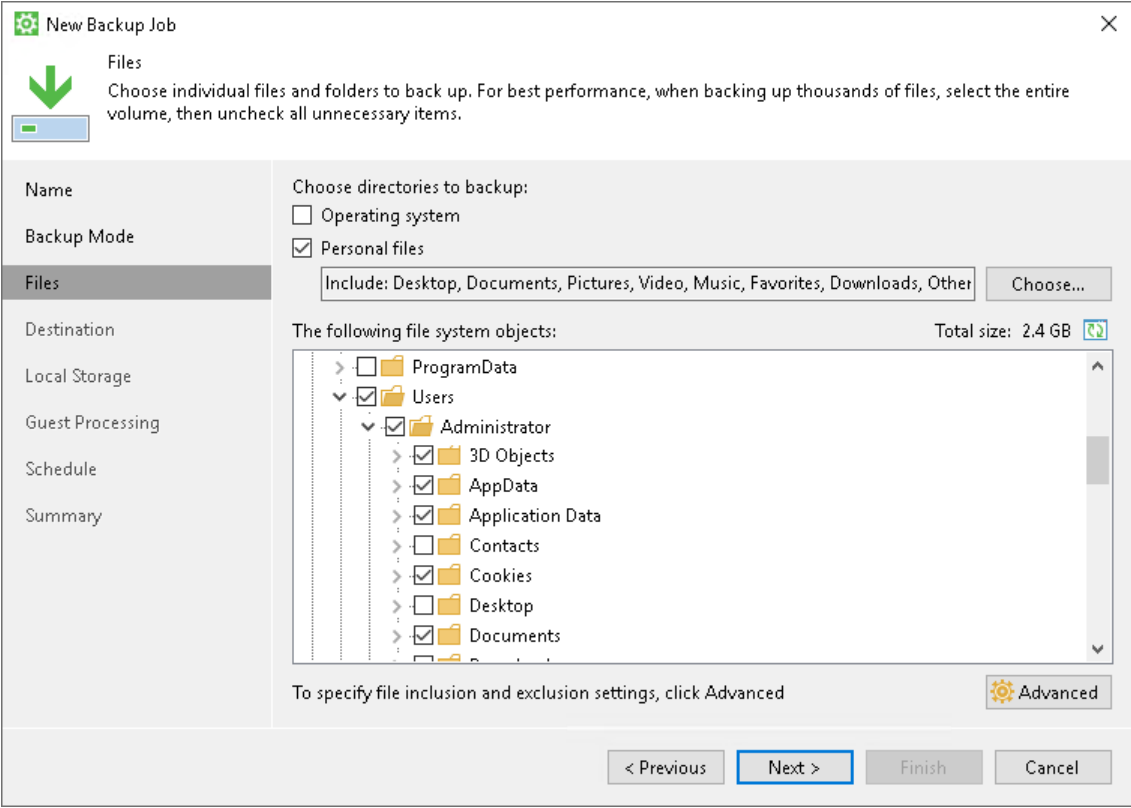
- By default, in the *Personal Folders* window, all options are selected in the **Backup user profile folders** list excluding the **Application Data** option and the **Exclude roaming profiles** option is selected.
- If you want Veeam Agent to back up personal data close to the way personal data was backed up in Veeam Agent 3.0 and 4.0, select all options in the **Backup user profile folders** list and clear the **Exclude roaming profiles** option. Note that in this case, the only difference is that Veeam Agent 6 also backs up selected folders of user profiles that are stored in custom locations.
- If Veeam Agent fails to back up personal data, Veeam Agent displays a warning message in the job session and a warning notification for each folder that Veeam Agent failed to back up. If you do not want to get these warnings, you can disable them with a registry value. To learn more, [contact Veeam Customer Support](#).

You can use the **Other files and folders** option to back up all folders and files that are located in the *Users* folder on the system volume, but are not available in the **Personal Folders** window. Keep in mind that depending on the **Other files and folders** option, Veeam Agent behaves in one of the following ways:

- If the option is cleared, Veeam Agent backs up only selected personal folders in the *Users* folder on the system volume and personal folders stored in custom locations. Keep in mind, that in such case Veeam Agent backs up only those user profiles that are related to authenticated users. As a result, for example, default and system user profiles will be excluded. Besides that, temporary and corrupted user profiles will be excluded too.
- If the option is selected, Veeam Agent backs up all files and folders in the *Users* folder on the system volume and personal folders stored in custom locations, but excludes personal data according to the options that you cleared.

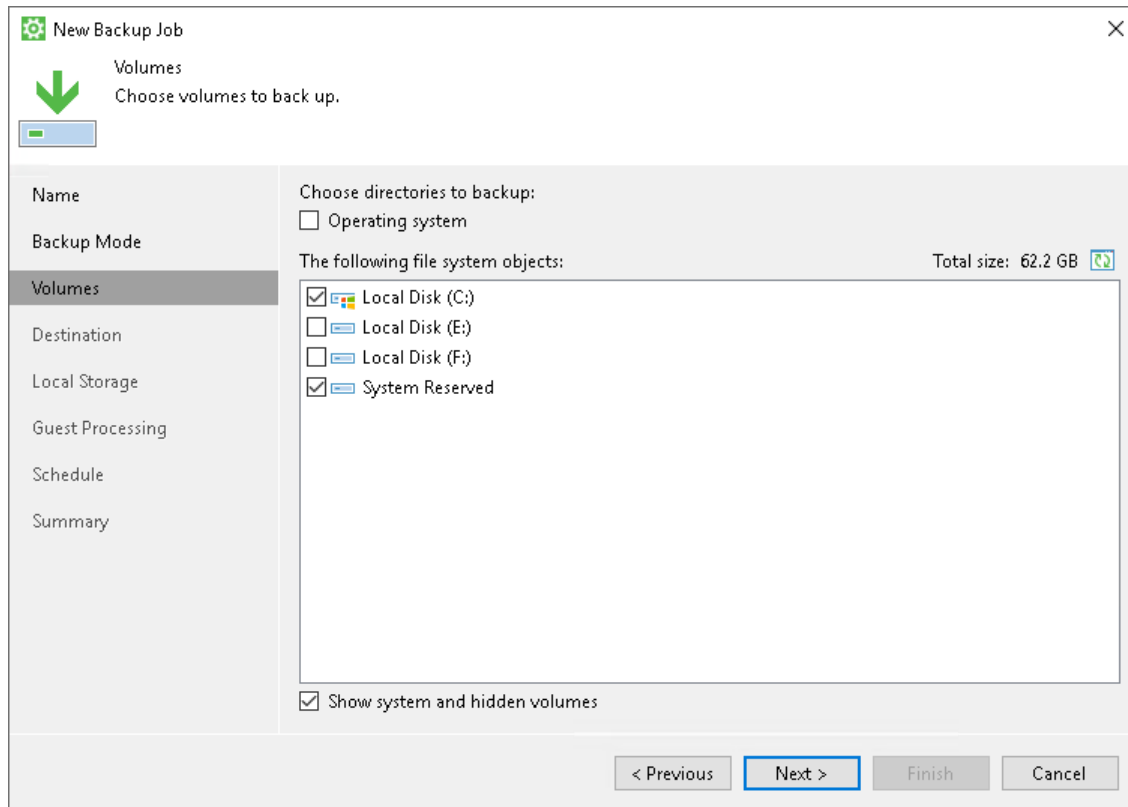


Alternatively, you can explicitly select to back up the *Users* folder. In this case, you still will be able to exclude specific subfolders of the *Users* folder from the backup. But you will need to exclude subfolders for each user manually. You will also need to find user folders stored in custom locations and include subfolders there if necessary. Thus, we recommend that you use the *Personal files* component to create personal data backup.



# Personal Data Backup in Volume-Level Backup

Typically, data and settings of the Veeam Agent computer users are located in the Users folder on the system volume, for example, `C:\Users`. Thus, to include the operating system data in the volume-level backup, you must explicitly select to back up the system volume.



# How Backup Works

During backup, Veeam Agent for Microsoft Windows performs the following operations:

1. Veeam Agent requests the creation of a Microsoft VSS snapshot of the volume whose data you want to back up. The VSS snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup.
  - In the Free and Workstation product editions, the VSS snapshot type depends on the OS of the Veeam Agent computer. On Microsoft Windows Client OSes, Veeam Agent requests a copy-only VSS snapshot. On Microsoft Windows Server OSes, Veeam Agent requests a full VSS snapshot.
  - In the Server product edition, the VSS snapshot type depends on the specified application-aware processing settings. To learn more, see [Application-Aware Processing](#).

Veeam Agent does not request a VSS snapshot for the EFI system partition on GPT disks, because its data does not change during backup. For the System Reserved and other system partitions, VSS snapshot can be created if there is enough free disk space on the partition.

## NOTE

Consider the following:

- By default, Microsoft Windows does not include offline Outlook Data Files (.ost) into a VSS snapshot. As a result, these files are not included into Veeam Agent backups, too.
- If the Microsoft VSS technology fails to create a VSS snapshot for some reason, Veeam Agent for Microsoft Windows resends the request up to 3 times.

2. Veeam Agent reads data from the created VSS snapshot, compresses it and copies it to the target location.
  - For volume-level backup, Veeam Agent copies data blocks of the whole volume.
  - For file-level backup, Veeam Agent creates a volume inside the backup file in the target location. The content of the volume in the backup file is synchronized with the volume on the source: Veeam Agent for Microsoft Windows copies only those data that you have selected to back up.

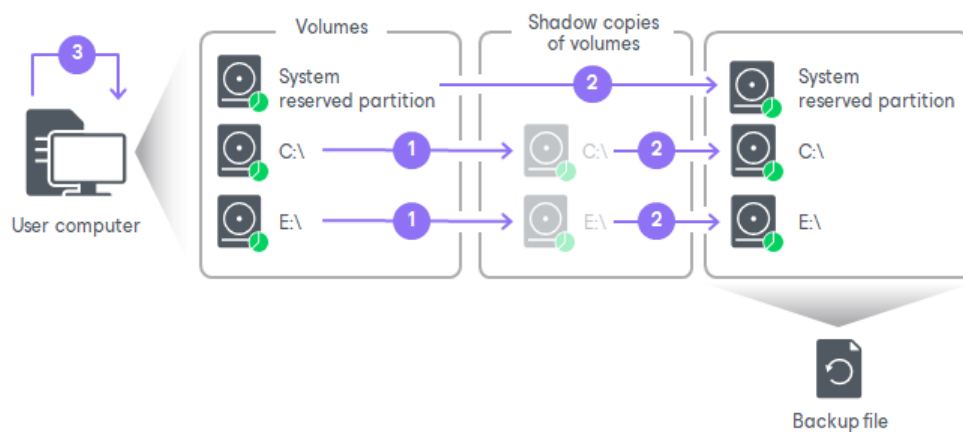
During incremental backup, Veeam Agent uses Changed Block Tracking (CBT) to retrieve only those data blocks that have changed since the previous backup session. To learn more, see [Change Block Tracking](#).

In the target location, Veeam Agent stores copied data to the backup file.

3. If an application on the computer uses transaction logs to maintain the database consistency, Veeam Agent can truncate transaction logs upon successful backup.
  - In the Free and Workstation product editions, Veeam Agent truncates transaction logs depending on the OS of the Veeam Agent computer. On Microsoft Windows Client OSes, Veeam Agent does not truncate logs. On Microsoft Windows Server OSes, Veeam Agent always truncates logs.
  - In the Server product edition, Veeam Agent truncates logs depending on the specified application-aware processing settings. To learn more, see [Application-Aware Processing](#).

## IMPORTANT

The Veeam Agent Service runs under the LocalSystem account. On Microsoft SQL Server 2012, this account does not have necessary permissions to truncate transaction logs. If you want Veeam Agent to automatically truncate transaction logs, you need to manually add the LocalSystem account to a group that has the SQL Server System Administrator rights.





# Backup Job

Veeam Agent for Microsoft Windows lets you configure a scheduled backup job that will perform backup automatically in a timely manner. You can set up the backup job once and forget about running the backup operation manually. Veeam Agent for Microsoft Windows will periodically launch the job to back up necessary data on your computer.

The backup job settings define what data you want to back up, what the target location and retention policy for created backups are and how often you want to back up your data. If necessary, you can re-configure the backup job at any time or remove the job.

Depending on the product edition, you can configure one or more backup jobs in Veeam Agent for Microsoft Windows:

- For the Free product edition, you can configure one backup job only.
- For the Workstation product edition, you can configure an unlimited number of backup jobs targeted at a Veeam Cloud Connect repository plus one backup job targeted at any other supported backup location (local or removable computer drive, network shared folder or Veeam backup repository).
- For the Server product edition, you can configure an unlimited number of backup jobs targeted at any supported backup location.

To learn more about product editions, see [Product Editions](#).

For example, you can configure one backup job to create volume-level backup and another backup job to create file-level backup. Or you can configure backup jobs targeted at different backup locations to keep several copies of your backed-up data. You can also configure multiple backup jobs with individual schedule to fine-tune automatic backup creation process.

Settings of the backup job apply to ad-hoc backups as well: standalone full backups, active full backups and incremental backups.

# Backup Job Schedule

Veeam Agent for Microsoft Windows launches the backup job according to the schedule you define. Scheduling options available for the backup job differ depending on the edition of Veeam Agent:

- For the Free and Workstation product editions, you can schedule the job to start at specific time daily or on specific week days. You can also instruct Veeam Agent to automatically perform backup on specific events. To learn more, see [Scheduling Options in Free and Workstation Editions](#).
- For the Server product edition, you can configure daily, monthly and periodic backup job schedule. You can also specify settings for automatic job retries and configure a backup window. To learn more, see [Scheduling Options in Server Edition](#).

If you configure multiple backup jobs in Veeam Agent, you can specify individual schedule for each job. Veeam Agent processes backup jobs sequentially, one after another. If the backup job must start upon schedule at the time when another backup job is already running, the job will wait for the running job to complete, and start afterwards.

For portable devices, Veeam Agent does not start a backup job on the defined schedule if a device is working on battery and the battery level is below 20%.

If a backup job fails, Veeam Agent automatically retries the job. In the Free and Workstation editions, Veeam Agent retries the job every 10 minutes within the next 23 hours. In the Server edition, you can specify retry settings along with other scheduling options. To learn more, see [Automatic Job Retries](#) and [Job Retry](#)

In any edition, Veeam Agent can wake your computer from sleep at the time when the backup job must start. To learn more, see [Computer Wake Up from Sleep](#).

## Scheduling Options in Free and Workstation Editions

You can schedule the backup job to start at specific time daily or on specific week days. You can also instruct Veeam Agent for Microsoft Windows to automatically perform backup on specific events.

### Missed Backup Schedule

Veeam Agent for Microsoft Windows does not perform scheduled backups if the computer is powered off. To handle situations of short power outage or computer restart, Veeam Agent provides a tolerance window of 15 minutes for scheduled backups.

For example, you have configured the backup job to run daily at 10:00 PM. At 9:55 PM, there is a power outage that lasts for 10 minutes. When the computer is on again at 10:05, Veeam Agent will automatically launch the scheduled job to back up your data.

Additionally, you can instruct Veeam Agent to resume missed daily backup. If the computer is powered off at the time when the scheduled backup job must start, and you power on the computer later, Veeam Agent will not wait for the next scheduled backup. Instead, Veeam Agent will start the backup job right after the computer is powered on to ensure no necessary data is lost because of the missed backup.

### Backup on Specific Events

In addition to the basic job schedule, you can instruct Veeam Agent to launch the backup job on specific events. Veeam Agent lets you trigger backup on the following events:

- Lock — the user locks the computer.
- Log off — the user performs a logout operation on the computer.

- When backup target is connected – the target backup location becomes available: the user attaches a known removable storage device to the computer or a network connection to the backup repository is established.

You can instruct Veeam Agent to eject the removable storage device after the backup job successfully completes. This helps to protect backup files in the target location from encrypting ransomware, such as CryptoLocker.

Backup on specific event helps you ensure that you capture all changes made within a specific time interval – for example, during a working day. When the necessary event occurs, Veeam Agent automatically launches the scheduled backup job. As a result, you can be sure that all changes made within some period of time are backed up, and you do not lose your data.

If you choose to perform backup on specific events, you can restrict the frequency of backup job sessions. You can instruct Veeam Agent not to start the backup job at specific events more often than once a specified time interval, for example, not more often than every 2 hours. This option does not affect daily schedule. Daily backups are performed according to the defined schedule regardless of the specified time interval.

Backup on specific events helps you fine-tune the backup job schedule. For example, you can specify the following scheduling settings for the backup job:

- The backup job must start automatically at 10:00 PM every day.
- The backup job must start at computer lock.
- The backup job must not run more often than every 2 hours.

Veeam Agent will launch the backup job at the end of the working day, when you lock your computer. In addition, Veeam Agent will perform backup at 10:00 PM regardless of the time interval between the computer lock and scheduled backup.

If you lock your computer later than at 10:00 PM, Veeam Agent will perform backup in the following order. At 10:00 PM, Veeam Agent will launch the backup job upon the daily schedule. If the time interval between the scheduled backup and computer lock is greater than 2 hours, Veeam Agent will additionally perform backup at computer lock. If the time interval between the scheduled backup and computer lock is not greater than 2 hours, Veeam Agent will not perform backup at computer lock.

## Automatic Job Retries

Veeam Agent for Microsoft Windows supports automatic retries for the scheduled backup job. If the backup job is started on the defined daily schedule and fails for some reason, Veeam Agent automatically retries the job every 10 minutes within the next 23 hours.

If the backup job fails and Veeam Agent retries the job session, Veeam Agent does not transfer all the data again. Instead, Veeam Agent continues data transfer that was started before the backup job fail. To do so, Veeam Agent compares hash values for data blocks on source and target. After the hash comparison, Veeam Agent transfers only those data blocks that were not transferred before the job fail. If data blocks on source were changed before the automatic retry, Veeam Agent transfers these data blocks as well.

If you edit certain job configuration options before the automatic retry, Veeam Agent may require to transfer all the data to target again. These options include backup scope, data compression, storage optimization, data encryption. Keep in mind that if you back up data to the Veeam backup repository and change data encryption options on the Veeam Backup & Replication side, Veeam Agent transfers all the data to the repository again.

Veeam Agent does not fail the backup job session in case Veeam Agent computer is put into sleep or hibernate mode during the job session, and the backup job is targeted at one of the following locations:

- Local computer drive
- Network shared folder

After the computer is turned on and the connection is restored, Veeam Agent operates in the same way as in case of automatic job retry. Veeam Agent does not transfer all the data again, but continues the job session from the point where it stopped.

## Limitations for Automatic Job Retry

- Veeam Agent continues data transfer after the backup job fail only if you configured a job that creates volume-level backups. To learn more about volume-level backups, see [Volume-Level Backup](#).
- Veeam Agent does not perform automatic retry for jobs started manually.
- Veeam Agent does not automatically retry the backup job if the job session is started when the computer is powered on after missed daily backup.
- Veeam Agent retries a job only if the previous job session has failed. Veeam Agent does not perform a retry if a job session has finished with the *Success* or *Warning* status.
- For portable devices, Veeam Agent does not automatically retry the backup job if a device is working on battery and the battery level is below 20%.

## Scheduling Options in Server Edition

You can schedule the backup job to start automatically at specific time. Veeam Agent for Microsoft Windows lets you configure the following settings for the job:

- [Scheduling settings](#)
- [Job retry settings](#)
- [Backup window settings](#)

## Automatic Startup Schedule


Veeam Agent for Microsoft Windows lets you configure the following scheduling settings for jobs:

- [You can schedule the backup job to run at specific time every day or on selected days](#)
- [You can schedule the backup job to run periodically at specific time intervals](#)
- [You can schedule the backup job to run continuously](#)

## Job Started at Specific Time

You can schedule the backup job to start at specific time daily, on specific week days or monthly on selected days.


This type of schedule requires that you define the exact time when the job must be started. For example, you can configure the job to start daily at 10:00 PM or every first Sunday of the month at 12:00 AM.




## Job Started at Specific Time Intervals

You can schedule the backup job to start periodically throughout a day at a specific time interval. The time interval between job sessions can be defined in minutes or hours. For example, you can configure a job to start every 30 minutes or every 2 hours.


For periodically run jobs, reference time is midnight (12:00 AM). Veeam Agent for Microsoft Windows always starts counting defined intervals from 12:00 AM, and the first job session will start at 12:00 AM. For example, if you configure a job to run with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.



If necessary, you can specify an offset for periodically run jobs. The offset is an exact time within an hour when the job must start. For example, you can configure the job to start with a 4-hour interval and specify offset equal to 15 minutes. In this case, the job will start at 12:15 AM, 4:15 AM, 8:15 AM, 12:15 PM, 4:15 PM and so on.




If a session of the periodically run job does not fit into the specified time interval and overlaps the next planned job session, Veeam Agent for Microsoft Windows starts the next backup job session at the nearest scheduled interval. For example, you set up the job to run with a 4-hour interval. The first job session starts at 12:00 AM, takes 5 hours and completes at 5:00 AM. In this case, Veeam Agent for Microsoft Windows will start a new job session at 8:00 AM.



## Job Run Continuously

You can schedule the job to run continuously — that is, in a non-stop manner. A new session of the continuously running job starts as soon as the previous job session completes. Continuously run job can help you implement near-continuous data protection (near-CDP) for the most critical applications.



## Job Retry

You can instruct Veeam Agent for Microsoft Windows to retry the backup job several times if the initial job session fails. By default, Veeam Agent for Microsoft Windows automatically retries a failed job for 3 times within one job session. If necessary, however, you can define a custom number of retries in the job settings.

If a backup job fails and Veeam Agent retries the job session, Veeam Agent does not transfer all the data again. Instead, Veeam Agent continues data transfer that was started before the backup job fail. To do so, Veeam Agent compares hash values for data blocks on source and target. After the hash comparison, Veeam Agent also transfers only those data blocks that were not transferred before the job fail. If data blocks on source were changed before the retry, Veeam Agent transfers these data blocks as well.

If you edit certain job configuration options before the retry, Veeam Agent may require to transfer all the data to target again. These options include backup scope, data compression, storage optimization, data encryption. Keep in mind that if you back up data to the Veeam backup repository and change data encryption options on the Veeam Backup & Replication side, Veeam Agent transfers all the data to the repository again.

Veeam Agent does not fail the backup job session in case Veeam Agent computer is put into sleep or hibernate mode during the job session, and the backup job is targeted at one of the following locations:

- Local computer drive
- Network shared folder

After the computer is turned on and the connection is restored, Veeam Agent operates in the same way as in case of job retry. Veeam Agent does not transfer all the data again, but continues the job session from the point where it stopped.

## Limitations for Job Retry

- Veeam Agent continues data transfer after the backup job fail only if you configured a job that creates volume-level backups. To learn more about volume-level backups, see [Volume-Level Backup](#).
- Veeam Agent does not perform retry for jobs started manually.
- Veeam Agent does not retry the backup job if the job session is started when the computer is powered on after missed daily backup.
- Veeam Agent retries a job only if the previous job session has failed. Veeam Agent does not perform a retry if a job session has finished with the *Success* or *Warning* status.
- For portable devices, Veeam Agent does not retry the backup job if a device is working on battery and the battery level is below 20%.

## Backup Window

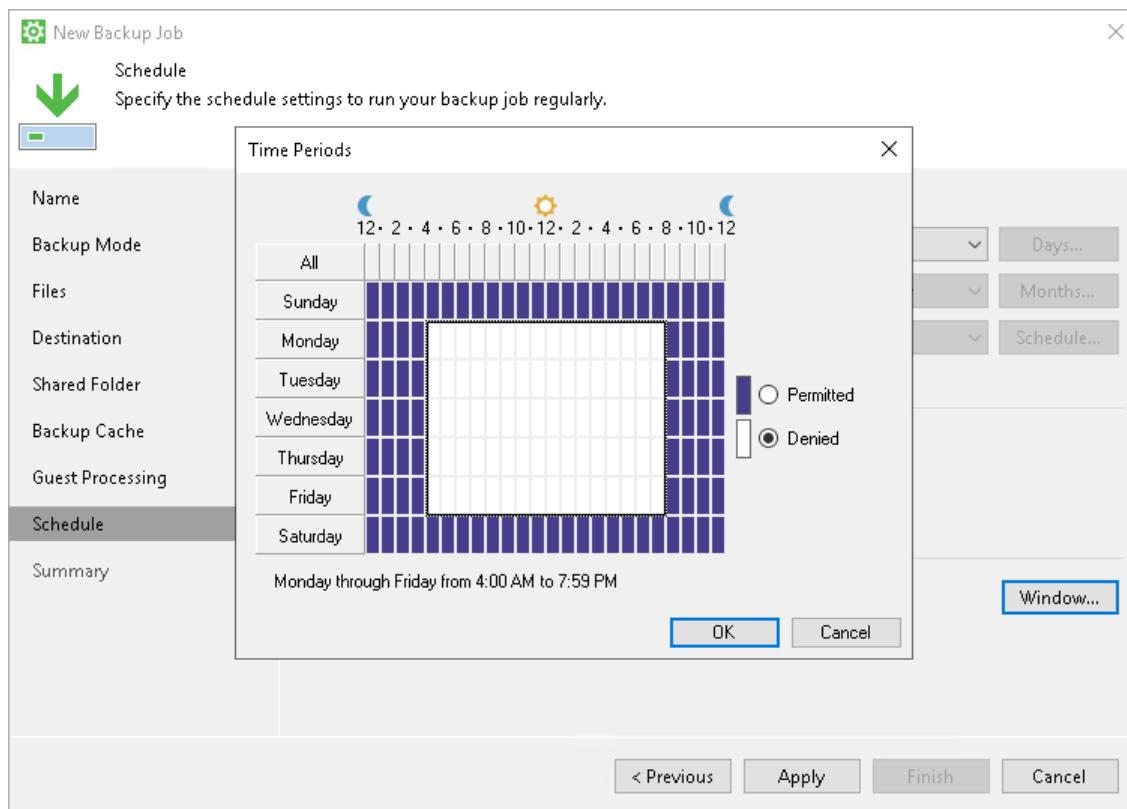
If necessary, you can specify a backup window for the backup job. The backup window is a period of time on week days when the job is permitted to run.

The backup window can be helpful if you do not want the data protection job to produce unwanted overhead for the production environment or do not want the job to overlap production hours. In this case, you can define the time interval during which the job must not run.

If the job exceeds the allowed window, it will be automatically terminated. In this case, data transport and backup chain transformation processes are stopped.

## IMPORTANT

The backup window does not affect the process of uploading backup files from the backup cache to the target storage. If Veeam Agent creates one or more backup files in the backup cache, and then the backup target becomes available, Veeam Agent uploads backup files to the target location immediately, regardless of the specified backup window.




## Backup Window for Periodically Run Job

If you define the backup window for the job that runs periodically at specific time intervals, Veeam Agent will immediately start the job after the denied window is over. All subsequent backup job sessions will be performed according to specified scheduling settings.

For example, you have configured the job to run with a 4-hour interval with an offset of 15 minutes. The allowed backup window for the job is 7:00 PM to 8:00 AM. Veeam Agent will run this job in the following way:

1. The first job session will start at 12:15 AM (since midnight is a reference time for periodically run jobs).
2. The next job session will start at 4:15 AM.
3. The job session at 8:15 AM will not be performed, because it falls into the denied period of the backup window.
4. The next job session will start immediately after the denied period is over: at 7:15 PM.

- After that, Veeam Agent for Microsoft Windows will run the job by the defined schedule: at 8:15 PM, 12:15 AM and so on.



## Computer Wake Up from Sleep

If your computer is in the standby mode at the time when the backup job must start, Veeam Agent for Microsoft Windows automatically wakes your computer from sleep. The wake-up feature lets you schedule your backup at night. At the defined time, Veeam Agent for Microsoft Windows will wake up the computer and perform a scheduled task. If necessary, you can additionally instruct Veeam Agent for Microsoft Windows to bring the computer back to the standby mode or power off the computer after the backup is finished.

Veeam Agent for Microsoft Windows wakes up the computer by default, unless the power saving settings on the computer prohibit this. If the wake up operation is not possible for some reason, the computer will remain in the standby mode, and the backup operation will not be performed. You can instruct Veeam Agent for Microsoft Windows to resume missed backup in such situations. To learn more, see [Missed Backup Schedule](#).

### IMPORTANT

[For tablets running Microsoft Windows 8.1 or later] If at the moment of backup a computer is in the Connected Standby power saving mode, Veeam Agent for Microsoft Windows will fail to wake it up due to limitations set by the OS itself.



# Ad-Hoc Backup

You can create ad-hoc backups of your data when you need.

Ad-hoc backups let you capture your data at a specific point in time. You can create ad-hoc backups before you perform some alterations on your computer: install new software or enable a new feature. Ad-hoc backups help you protect your computer from potential data corruption or data loss that can be caused by these operations. If an error occurs, you can always restore data from the ad-hoc backup and bring your computer system to a state before the alteration was made.


Veeam Agent for Microsoft Windows lets you create the following types of ad-hoc backups:

- [Incremental backup](#)
- [Standalone full backup](#)

## Ad-Hoc Incremental Backup

If you want to create a new backup of your data in addition to backups created with the scheduled backup job, you can perform ad-hoc incremental backup. Ad-hoc incremental backup adds a new restore point to the backup chain. For example, you may want to back up your data before you install new software on your computer or enable a new feature.

For ad-hoc incremental backup, Veeam Agent for Microsoft Windows uses settings specified for the backup job. For example, if you have configured the backup job to perform backup of a specific volume, the ad-hoc incremental backup operation will create an incremental backup of this volume and save it in the target location, next to existing backup files in the backup chain.



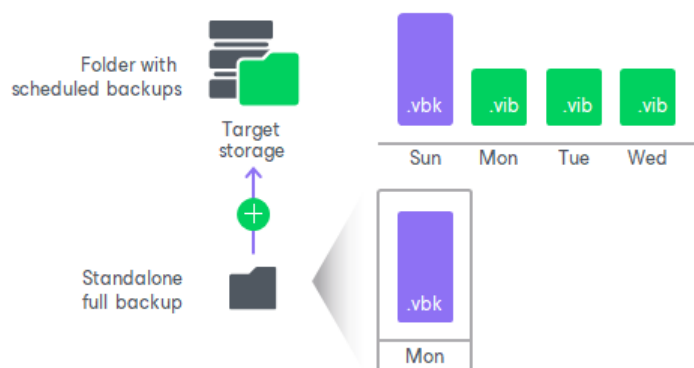
Unlike the backup job started upon schedule, the ad-hoc incremental backup task is not retried automatically. If the task fails for some reason, you will have to start it manually again.

Veeam Agent for Microsoft Windows treats restore points created by ad-hoc incremental backup as regular restore points, and applies to them retention policy settings specified for the backup job. To learn more, see [Backup Retention Policy](#).

## Standalone Full Backup

Sometimes you need to create a full backup of your data. For example, you may want to save a copy of your data on a CD or DVD or create a full backup of all data on your computer at some point in time. In these situations, you can perform standalone full backup.

When Veeam Agent for Microsoft Windows performs standalone full backup, it produces a full backup of your data in a separate folder in the target location. The standalone full backup is not associated with subsequent incremental backups. You can use it as an independent restore point for data recovery.



To create a standalone full backup, Veeam Agent for Microsoft Windows uses settings specified for the backup job. For example, if you have configured the backup job to perform backup of a specific volume, the standalone full backup will create a full backup of this volume in a separate folder in the target location.

Unlike the backup job started upon schedule, the standalone full backup task is not retried automatically. If standalone full backup fails for some reason, you will have to start the standalone full backup task manually again.

The standalone full backup is not removed by retention. To remove it, you must manually delete the full backup file from disk.

#### NOTE

You cannot perform standalone full backup if the backup job is targeted at one of the following locations:

- Veeam backup repository
- Veeam Cloud Connect repository
- Object storage

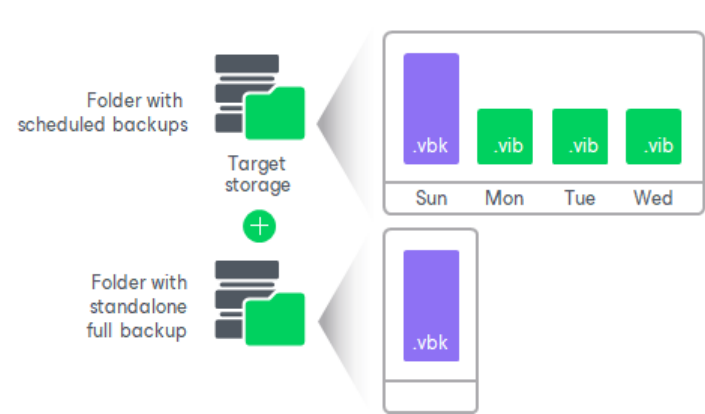
## Standalone Full Backup to Another Location

You can create a standalone full backup in a separate location that is not specified as a target location in the backup job settings. For example, you may want to save a copy of your data on a removable storage device while your scheduled backup job is targeted at the network shared folder.

Backup to another location practically does not differ from regular standalone full backup. The only difference is that you must manually select a target location in which Veeam Agent for Microsoft Windows will save the backup file. You can save backup files to the following locations:

- Removable storage device
- Local computer drive
- Network shared folder

You cannot use a Veeam backup repository or Veeam Cloud Connect repository as a target for backup to another location.



# Parallel Disk Processing

Veeam Agent for Microsoft Windows supports parallel disk processing. If you included several disks of the Veeam Agent computer in the backup scope, Veeam Agent will process disks simultaneously.

## How Parallel Disk Processing Works

During backup, Veeam Agent for Microsoft Windows requests the creation of a Microsoft VSS snapshot of the volume whose data you want to back up as described in section [How Backup Works](#). After the VSS snapshot is created, Veeam Agent transfers data blocks of disks to the target storage simultaneously.


The number of disks that are simultaneously processed by Veeam Agent is limited by repository settings on the Veeam Backup & Replication side. During the backup job session, Veeam Backup & Replication compares the number of currently processed disks to the limit specified in the properties of the backup repository. If the number of currently processed disks is greater than the limit, Veeam Agent will not transfer data of a new disk. Only after data transfer for at least one of the disks completes, Veeam Agent starts data transfer for the next disk. To learn more, see the [Limitation of Concurrent Tasks](#) section in the Veeam Backup & Replication User Guide.

You can also limit the number of disks that are processed simultaneously with a registry value. Keep in mind that if the number of simultaneously processed disks is limited by both repository settings and registry value, Veeam Agent uses the lower limit. To learn more about the registry value, see [this Veeam KB article](#).

If processing of at least one disk fails, the entire backup job session completes with an error.

## Limitations for Parallel Disk Processing


- Parallel disk processing is available only in the Server edition of Veeam Agent for Microsoft Windows.
- Veeam Agent offers parallel disk processing for backups that are created in the following types of target locations:
  - Veeam backup repository
  - Veeam Cloud Connect repository
- Parallel disk processing is not supported for dynamic disks.
- Parallel disk processing is not performed for file-level backups.



# Backup Chain

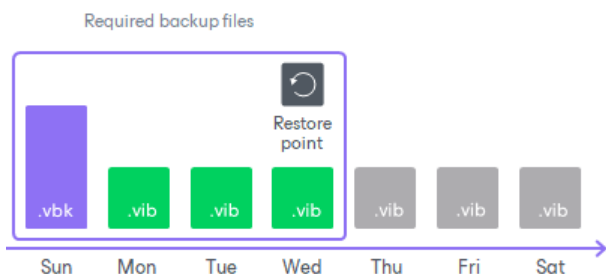
Every backup job session produces a new backup file in the target location. Backup files make up a backup chain. The backup chain can contain files of two types: full backups and incremental backups.

- During the first backup job session, Veeam Agent performs full backup. It copies all data that you have chosen to back up (entire volumes and folders) and stores the resulting full backup file (VBK) in the target location. The full backup takes significant time to complete and produces a large backup file: you have to copy the whole amount of data.
- During subsequent backup job sessions, Veeam Agent performs incremental backups. It copies only new or changed data relatively to the last backup job session and saves this data as an incremental backup file (VIB) in the target location. Incremental backups typically take less time than full backup: you have to copy only changes, not the whole amount of data.



After several backup cycles, you have a chain of backup files in the target location: the first full backup file and subsequent incremental backup files. Every backup file contains a restore point for backed-up data. A restore point is a "snapshot" of your data at a specific point in time. You can use restore points to roll back your data to the necessary state.

To recover data to a specific restore point, you need a chain of backup files: a full backup file plus a set of incremental backup files following this full backup file. If some file from the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, we recommend that you do not delete separate backup files manually. To learn more, see [Deleting Backups](#).



## Types of Backup Files

Veeam Agent produces backup files of the following types:

- VBK — full backup file.
- VIB — incremental backup file.
- VBM — backup metadata file. The backup metadata file is updated with every backup job session. It contains information about the computer on which the backup was created, every restore point in the backup chain, how restore points are linked to each other and so on. The backup metadata file is required for performing file-level and volume-level restore operations.

## NOTE

For backup jobs with database log backup options enabled, Veeam Agent for Microsoft Windows additionally produces backup files of the following types:

- VLB, VSM and VLM files — for Microsoft SQL Server transaction log backups
- VLB, VOM and VLM files — for Oracle archived log backups

Keep in mind that Veeam Agent creates VLM files only in backup chains started in Veeam Agent for Microsoft Windows 5.0 and later. If you upgrade Veeam Agent and continue the backup chain that was started in the earlier Veeam Agent version, Veeam Agent will not create VLM files in this backup chain.

## Short-Term Retention Policy

Restore points in the backup chain are not kept forever. They are removed according to the retention policy. The short-term retention policy helps maintain the life cycle of restore points and make sure that backup files do not consume the whole disk space.

Short-term retention policy depends on the edition of Veeam Agent for Microsoft Windows:

- In the Free and Workstation editions, Veeam Agent retains restore points for a certain number of days. To learn more, see [Retention Policy in Free and Workstation Editions](#).
- In the Server edition, you can select the logic behind the short-term retention policy. Veeam Agent can retain restore points for a certain number of days or retains a certain number of latest restore points that is defined by the user. To learn more, see [Retention Policy in Server Edition](#).

You can also specify the retention policy for outdated backups — backups for which Veeam Agent does not create new restore points within a specified time period. To learn more, see [Retention Policy for Outdated Backups](#).

## Retention Policy in Free and Workstation Editions

In the Free and Workstation editions, Veeam Agent for Microsoft Windows retains restore points for the last <N> days; the number of days is defined by the user. During every backup job session, Veeam Agent for Microsoft Windows checks if there is any obsolete restore point in the backup chain. If some restore point is obsolete, it is removed from the chain.

For retention policy settings, Veeam Agent for Microsoft Windows takes into account not calendar days but days on which backup files were successfully created. Veeam Agent ignores restore points created on the day when the retention policy is applied. In fact, Veeam Agent keeps restore points for the  $<N> + 1$  days, where <N> is the number of days that you specified in the backup job settings.

For example, you have configured the backup job in the following way:

- The backup job runs daily.
- The retention policy is set to 5 days.


The backup job has successfully run 3 times and created 3 restore points in the backup chain. After that, you have turned off your computer for 10 days. When you turn on your computer, Veeam Agent for Microsoft Windows runs a backup job by schedule and creates a new restore point. The earliest restore point, however, is not removed from the backup chain. At the end of a new backup job session, the backup chain will have only 4 restore points created during 4 days when the backup job was successfully run.




## Removing Backups by Retention

When the obsolete restore points are removed by retention, Veeam Agent transforms the backup chain so it always contains a full backup file on which subsequent incremental backup files are dependent. To do so, Veeam Agent uses the following rotation scheme:


1. During every backup job session Veeam Agent adds a backup file to the backup chain and checks if there is an obsolete restore point.




2. If an obsolete restore point exists, Veeam Agent transforms the backup chain. As part of this process, it performs the following operations:
  - a. Veeam Agent rebuilds the full backup file to include in it data of the incremental backup file that follows the full backup file. To do this, Veeam Agent injects into the full backup file data blocks from the earliest incremental backup file in the chain. This way, a full backup 'moves' forward in the backup chain.



- b. Veeam Agent removes the earliest incremental backup file from the chain as redundant: its data has already been injected into the full backup file, and the full backup file includes data of this incremental backup file.



If the backup chain contains several obsolete restore points, the rebuild procedure is similar. Data from several restore points is injected to the rebuilt full backup file. This way, Veeam Agent makes sure that the backup chain is not broken, and you will be able to recover your data to any restore point.



## Retention Policy in Server Edition

In the Server edition, you can select the logic behind the short-term retention policy. Veeam Agent offers the following options:


- Veeam Agent retains restore points for the last <N> days; the number of days is defined by the user. In this case, Veeam Agent works in a similar way as in the Free or Workstation editions. To learn more, see [Retention Policy in Free and Workstation Editions](#).
- Veeam Agent retains the number of latest restore points defined by the user.

During every backup job session, Veeam Agent checks if there is any obsolete restore point in the backup chain. If some restore point is obsolete, it is removed from the chain.

## Removing Backups by Retention


When the obsolete restore points are removed by retention, Veeam Agent transforms the backup chain so it always contains a full backup file on which subsequent incremental backup files are dependent. To do so, Veeam Agent uses the following rotation scheme:

1. During every backup job session Veeam Agent adds a backup file to the backup chain and checks if there is an obsolete restore point.







2. If an obsolete restore point exists, Veeam Agent transforms the backup chain. As part of this process, it performs the following operations:
  - a. Veeam Agent rebuilds the full backup file to include in it data of the incremental backup file that follows the full backup file. To do this, Veeam Agent injects into the full backup file data blocks from the earliest incremental backup file in the chain. This way, a full backup 'moves' forward in the backup chain.



- b. Veeam Agent removes the earliest incremental backup file from the chain as redundant: its data has already been injected into the full backup file, and the full backup file includes data of this incremental backup file.



If the backup chain contains several obsolete restore points, the rebuild procedure is similar. Data from several restore points is injected to the rebuilt full backup file. This way, Veeam Agent makes sure that the backup chain is not broken, and you will be able to recover your data to any restore point.



## Retention Policy for Outdated Backups

In addition to specifying general retention policy settings, you can define retention policy for outdated backups. An outdated backup is a backup for which no new restore points were created within the last <N> days. Retention policy for outdated backups helps to avoid keeping redundant data in the target location: once the outdated backup retention period is over, an outdated backup is automatically removed from the target location.

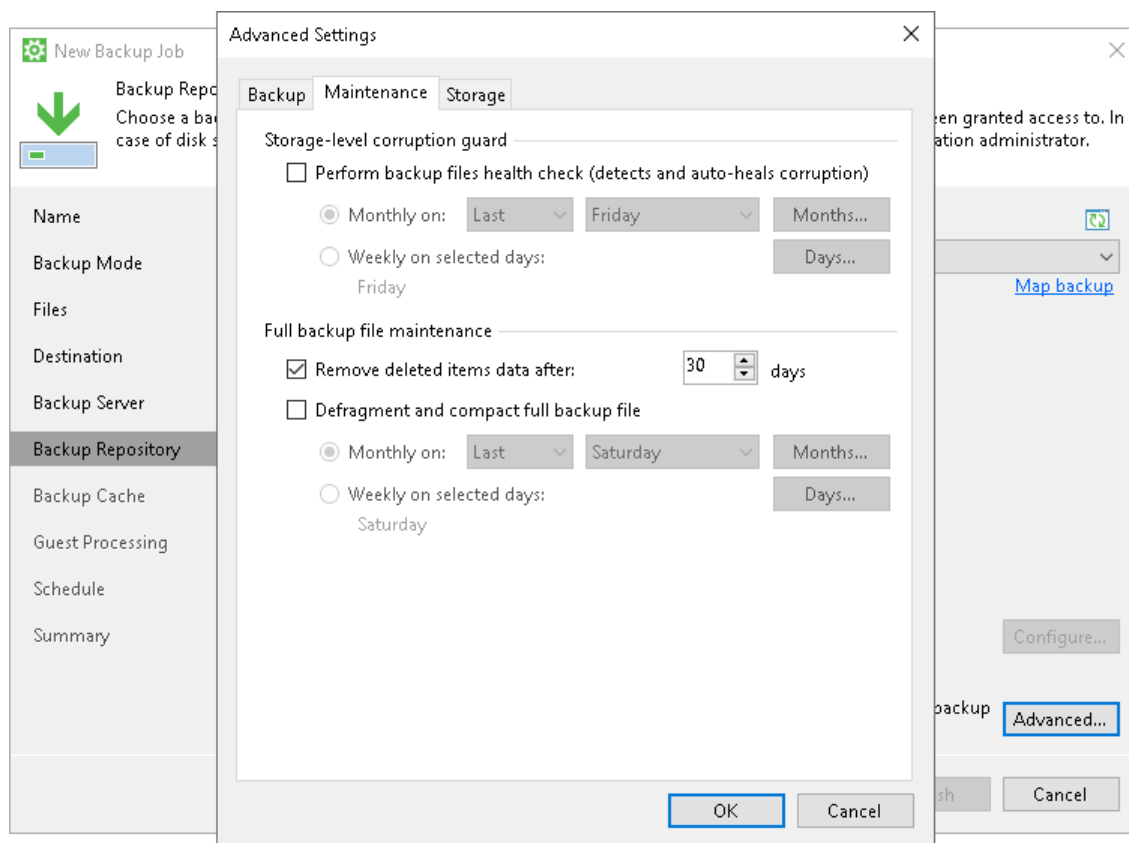
Veeam Agent for Microsoft Windows offers retention policy for outdated backups that are created in the following types of target locations:

- Veeam backup repository
- Veeam Cloud Connect repository

To specify retention policy for outdated backups, you must enable the **Remove deleted items data after <N> days** option in backup job settings and specify the number of days for which data of outdated backups must be retained in the backup repository. By default, Veeam Agent for Microsoft Windows retains outdated backups for 30 days. You can change the number of days to retain outdated backups if necessary.

Consider the following:

- You must use retention policy for outdated backups wisely. We strongly recommend that you set retention policy for outdated backups to 7 days or more to prevent unwanted data loss.
- The **Remove deleted items after <N> days** option lets you control data of backups for which Veeam Agent for Microsoft Windows does not produce new restore points for some time. In addition to it, Veeam Agent for Microsoft Windows applies general retention policy rules to maintain the necessary number of restore points in the backup chain. To learn more, see [Backup Retention Policy](#).



## How Retention Policy for Outdated Backups Works

Although you specify retention policy settings for outdated backups in Veeam Agent for Microsoft Windows, actions required to track and delete outdated Veeam Agent backups are taken on the Veeam Backup & Replication side.

To perform necessary actions, Veeam Backup & Replication performs background retention for backups on the backup server. The background retention starts automatically every 24 hours at 00:30.

### NOTE

You can change the background retention schedule with registry values. For more information, [contact Veeam Customer Support](#).

During the background retention session, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication checks the configuration database to detect Veeam Agent backups that reside in the backup repository.

2. For each Veeam Agent backup, Veeam Backup & Replication checks whether the following conditions are met:

- a. The configuration database contains information about the backup job that created the Veeam Agent backup.

The configuration database does not contain information about the backup job that created the backup in the following cases:

- The backup job that created the backup was removed in Veeam Agent for Microsoft Windows.
- The Veeam Agent backup was imported in the Veeam Backup & Replication console.
- The Veeam Agent backup is a standalone full backup.

If the configuration database does not contain information about the backup job that created the backup, Veeam Backup & Replication does not remove such a backup from the backup repository.

- b. The **Remove deleted items after <N> days** option is enabled in the backup job settings. If the option is enabled, Veeam Backup & Replication checks whether the following conditions are met:

- No new restore points were created by the backup job for the last <N> days.
- No new backup job sessions were detected by Veeam Backup & Replication for the last <N> days.
- [For backup jobs with database log backup enabled] No database log copies (Microsoft SQL Server transaction log or Oracle archive log copies) were created by the database log backup job for the last <N> days.

Where <N> is the number of days specified for the **Remove deleted items data after <N> days** option.

3. If all the conditions listed in the step 2 are met, Veeam Backup & Replication removes the Veeam Agent backup from the configuration database and deletes actual backup files from the backup repository.

## Long-Term Retention Policy (GFS)

The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store backup files for long periods of time — for weeks, months and even years. For this purpose, Veeam Agent does not create any special new backup files — it uses backup files created while backup job runs and marks these backups with specific GFS flags.

To mark a backup file for long-term retention, Veeam Agent can assign to the file the following types of GFS flags: weekly (W), monthly (M) and yearly (Y). The types of GFS flags that Veeam Agent assigns depend on the configured [GFS retention policy settings](#).

### NOTE

Consider the following:

- GFS flags can be assigned only to full backup files created during the time period specified in GFS policy settings.
- If you store your backups in object storage, and you do not configure Veeam Agent to perform active full backups periodically, Veeam Agent will create a full backup based on the last incremental backup and will assign a GFS flag to this full backup. If some data blocks required to create the full backup already reside in the object storage repository, the full backup will contain links to such data blocks. To avoid extra costs, Veeam Agent does not retrieve actual data blocks from the object storage repository.

If Veeam Agent assigns a GFS flag to a full backup file, this backup file can no longer be deleted or modified. Veeam Agent does not apply short-term retention policy settings to the full backup file. For example, Veeam Agent ignores the backup file when determining whether the number of allowed backup files is exceeded.

When the specified retention period ends, Veeam Agent unassigns the GFS flag from the full backup file. If the backup file does not have any other GFS flags assigned, it can be modified and deleted according to the short-term retention policy.

Veeam Agent assigns GFS flags in the similar way as Veeam Backup & Replication does for VM backup files. To learn about logic behind GFS flags, see the [Assignment of GFS Flags](#) and [Removal of GFS Flags](#) sections in the Veeam Backup & Replication User Guide.

## Limitations

When planning to use GFS retention policy, consider the following limitations:


- [Applicable to all backup targets except object storage] While applying the GFS retention policy, Veeam Agent does not create new full backup files. You must configure your backup jobs in a way you do not lose any essential data due to an insufficient number of full backup files. For example, if you configure monthly GFS retention, you need at least one full backup file per month.
- If a GFS flag is assigned to a full backup file in an active backup chain, the following applies:
  - Veeam Agent cannot transform the backup chain according to the short-term retention policy.
  - Veeam Agent is not able to merge data from incremental backup files into the full backup file.
- Veeam Agent assigns GFS flags only after you save GFS retention policy settings. This means that GFS flags are assigned only to those backup files created after the configuration, while backup files created earlier are not affected and previously assigned flags are not modified.
- You cannot store full backups to which GFS flags are assigned in backup repositories with rotated drives.
- Retention policy for deleted items does not apply to full backup files to which GFS flags are assigned.

## Active Full Backup

In some cases, you need to regularly create a full backup. For example, your corporate backup job may require that you create a full backup on weekend and run incremental backup on work days. To let you conform to these requirements, Veeam Agent lets you create active full backups.

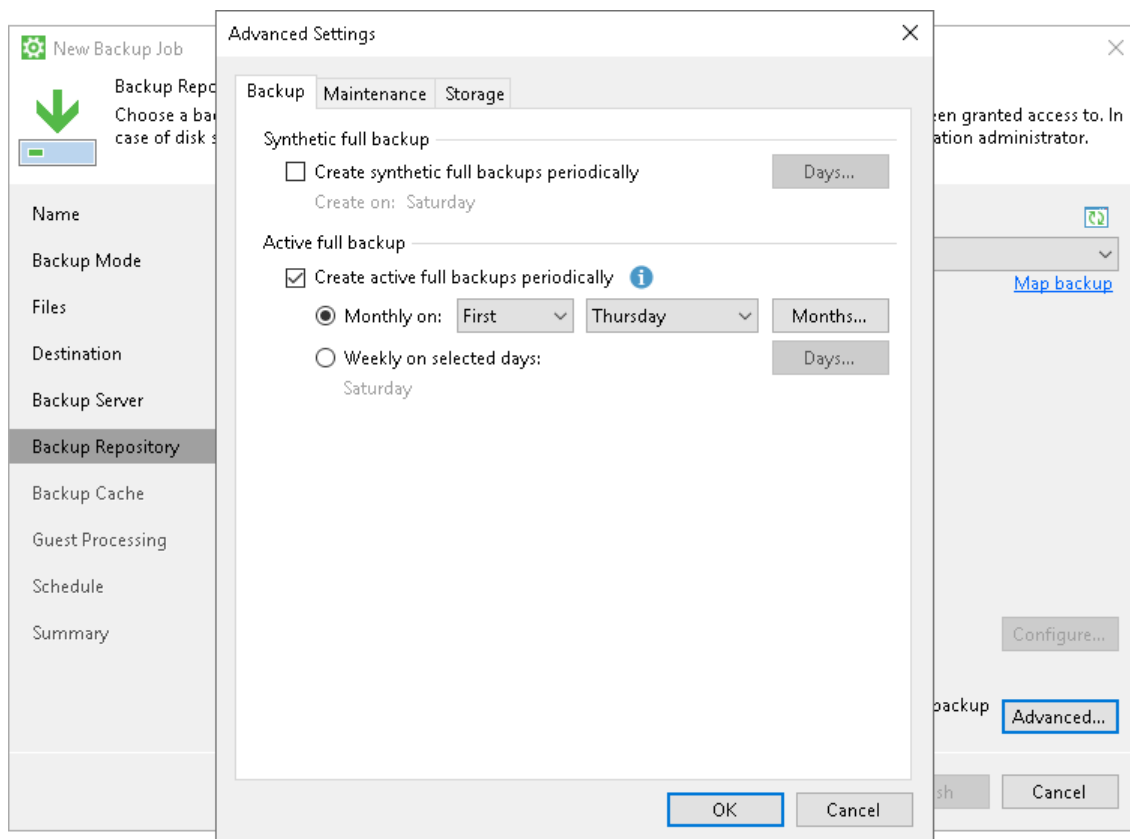
When Veeam Agent performs active full backup, it produces a full backup file and adds this file to the backup chain.

The active full backup resets the backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file and its subsequent incremental backup files remain on the disk. After the last incremental backup file created prior to the active full backup becomes outdated, Veeam Agent automatically deletes the previous backup chain. To learn more, see [Retention Job for Active Full Backups](#).



You can create active full backups manually or schedule a backup job to create active full backups periodically.


- To create an active full backup manually, use the **Active full backup** command from the Veeam Agent Tray menu. To learn more, see [Creating Active Full Backups](#).
- To schedule active full backups, specify scheduling settings in the **Advanced Settings** window of the **New Backup Job** wizard. You can schedule active full backups to run weekly, for example, every Saturday, or monthly, for example, every first Thursday of a month.




To be able to restore data from a Veeam Agent backup, you need to have a full backup file and a chain of subsequent incremental backup files on the disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you create an active full backup, in some days there will be more restore points on the disk than specified by retention job settings. Veeam Agent will remove the full backup chain only after the last incremental backup file in the chain becomes outdated.

For example, the retention job is set to 3 restore points. A full backup file is created on Sunday, incremental backup files are created on Monday and Tuesday, and an active full backup is created on Wednesday. Although the backup chain now contains 4 restore points, Veeam Agent will not delete the previous backup chain. Veeam Agent will wait for the next 2 incremental backup files to be created, and only then will delete the whole previous chain, which will happen on Friday. As a result, although the retention job is set to 3 restore points, the actual number of backup files on the disk will be greater for some time.



Veeam Agent treats the active full backup in the same way as a regular full backup. If some restore point becomes obsolete, Veeam Agent will re-build the full backup file to include in it data of the incremental backup file that follows the full backup file. After that, Veeam Agent will remove the earliest incremental backup file from the chain as redundant.



## Synthetic Full Backup

In some situations, running active full backups periodically may not be an option. Active full backups are resource-intensive and consume considerable amount of network bandwidth. As an alternative, you can create synthetic full backups.

### NOTE

Consider the following:

- Synthetic full backup is available only in the Workstation and Server editions of Veeam Agent for Microsoft Windows.
- Synthetic full backup is not available for backup jobs targeted at object storage.

In terms of data, the synthetic full backup is identical to a regular full backup. Synthetic full backup produces a VBK file that contains all data that you have chosen to back up. The difference between active and synthetic full backup lies in the way how backed-up data is retrieved:

- When you perform active full backup, Veeam Agent for Microsoft Windows reads backed-up data, compresses it and copies it to the target location.
- When you perform synthetic full backup, Veeam Agent for Microsoft Windows does not retrieve all backed-up data from the Veeam Agent computer. Instead, it creates a new incremental backup and then synthesizes a full backup from data you already have on the target location. Veeam Agent for Microsoft Windows accesses the previous full backup file and a chain of subsequent incremental backup files in the backup chain including a new incremental backup, consolidates data from these files and writes consolidated data into a new full backup file. As a result, the created synthetic full backup file contains the same data as an active full backup.

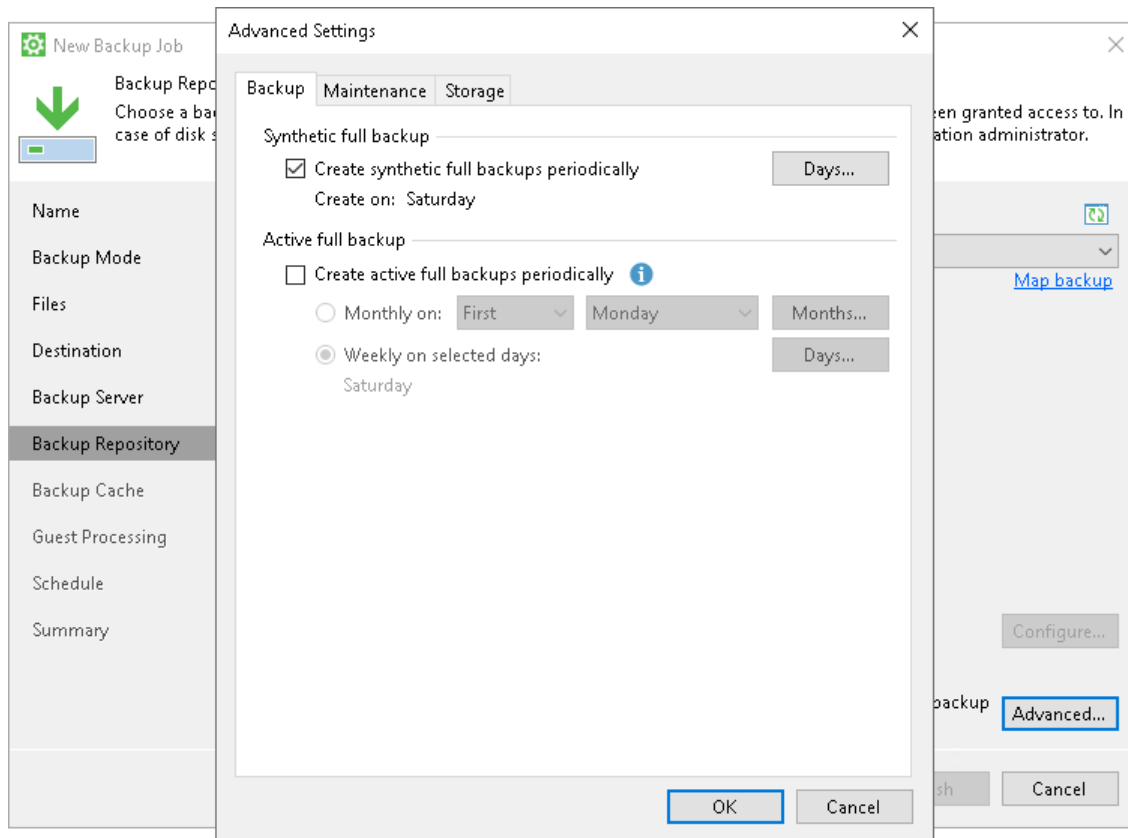
The synthetic full backup has a number of advantages:

- The synthetic full backup does not use network resources: it is created from backup files you already have on the target location.
- The synthetic full backup produces less load on the production environment: it is synthesized right on the target location.

Veeam Agent for Microsoft Windows treats synthetic full backups as regular full backups. As well as any other full backup file, the synthetic full backup file resets the backup chain. All subsequent incremental backup files use the synthetic full backup file as a new starting point.

A previously used full backup file and its subsequent incremental backup files remain on the disk. After the last incremental backup file created prior to the synthetic full backup becomes outdated, Veeam Agent for Microsoft Windows automatically deletes the previous backup chain. To learn more, see [Retention Policy for Synthetic Full Backups](#).

To create synthetic full backups, you must enable the **Create synthetic full backups periodically** option and schedule creation of synthetic full backups on specific days in the backup job settings.




## Retention Policy for Synthetic Full Backups

To be able to restore data from a Veeam Agent backup, you need to have a full backup file and a chain of subsequent incremental backup files on the disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you will not be able to restore data (since later incremental backup files depend on earlier incremental backup files).


For this reason, if you set up the backup job to create synthetic full backups, in some days there will be more restore points on the disk than specified by retention policy settings. Veeam Agent for Microsoft Windows will remove the full backup chain only after the last incremental backup file in the chain becomes outdated.

For example, the retention policy is set to 3 restore points. A full backup file is created on Sunday, incremental backup files are created Monday through Saturday, and synthetic full backup is scheduled on Thursday. Although the retention policy is already breached on Wednesday, the full backup is not deleted. Without the full backup, backup chain is useless, leaving you without any restore point at all. Veeam Agent for Microsoft Windows will wait for the next full backup file and 2 incremental backup files to be created, and only then will delete the whole previous chain, which will happen on Saturday.






Keep in mind that if the backup job is set up to create synthetic full backups, Veeam Agent for Microsoft Windows will never transform the backup chain. Instead, Veeam Agent for Microsoft Windows will always wait for the next full backup file and the necessary number of incremental backup files to be created, and only then will delete the whole previous chain. In the example above, Veeam Agent for Microsoft Windows will delete the previous chain every Saturday. As a result, although the retention policy is set to 3 restore points, the actual number of backup files on the disk will be greater most of the time.




## How Synthetic Full Backup Works

To create a synthetic full backup, Veeam Agent for Microsoft Windows performs the following steps:


1. On a day when synthetic full backup is scheduled, Veeam Agent for Microsoft Windows triggers a new backup job session. During this session, Veeam Agent for Microsoft Windows first performs incremental backup in a regular manner and adds a new incremental backup file to the backup chain. Incremental backup helps Veeam Agent for Microsoft Windows ensure that the synthetic full backup includes the latest changes of the backed-up data.



2. At the end of the backup job session, Veeam Agent for Microsoft Windows builds a new synthetic full backup using backup files that are already available in the backup chain, including the newly created incremental backup file.



- When the synthetic full backup is created, Veeam Agent for Microsoft Windows deletes the incremental backup file created at the beginning of the job session. As a result, you have a backup chain that consists of a full backup file, set of incremental backup files and synthetic full backup file.



- Every next job session creates a new incremental restore point starting from the synthetic full backup until the day on which synthetic full backup is scheduled. On this day, Veeam Agent for Microsoft Windows creates a new synthetic full backup.

## Compact of Full Backup File

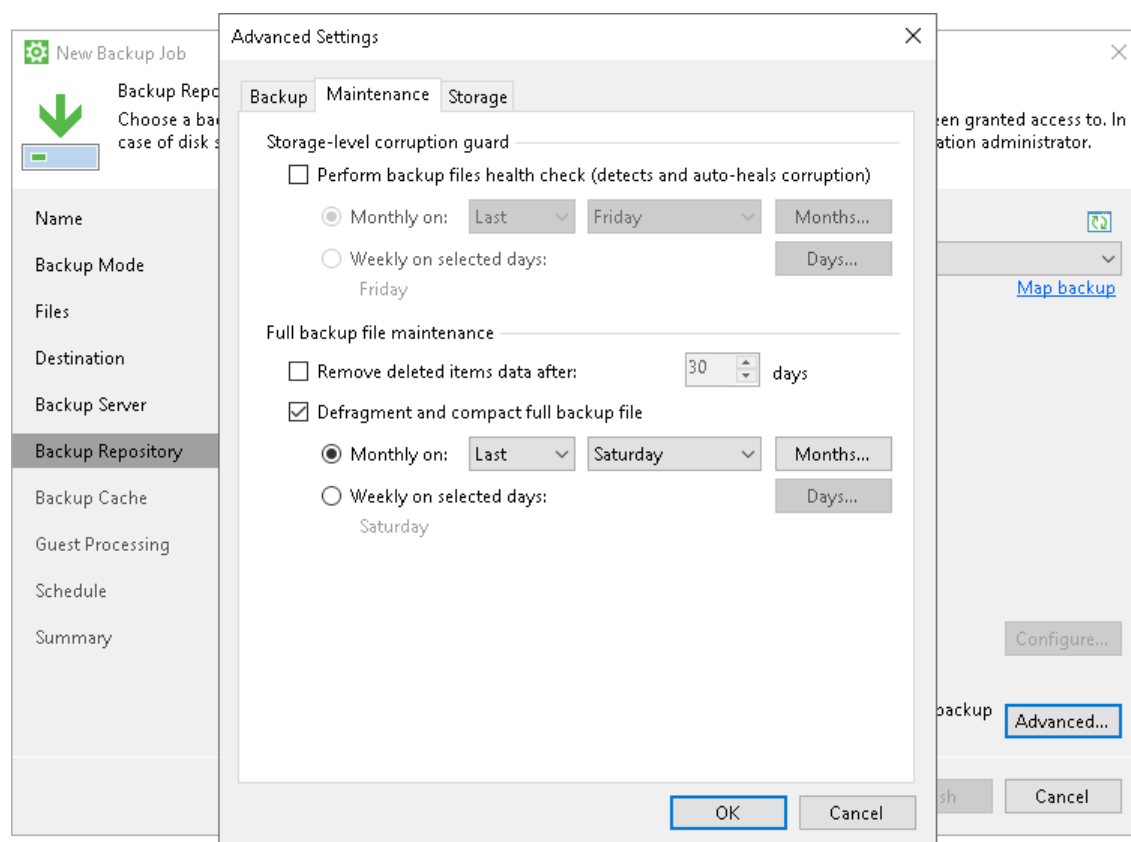
If you do not create periodic full backups, the backup job constantly transforms the full backup file in the backup chain to meet retention policy settings. The transformation process, however, has a side effect. In the long run, the full backup file grows large and gets badly fragmented. The file data occurs to be written to non-contiguous clusters on disk, and operations of reading and writing data from and to the backup file slow down.

To resolve the fragmentation problem, you can instruct Veeam Agent for Microsoft Windows to compact the full backup file periodically. During the file compact operation, Veeam Agent creates a new empty file and copies to this file data blocks from the original full backup file. As a result, the full backup file gets defragmented and the speed of reading and writing from and to the file increases.

### NOTE

Keep in mind that some data blocks belonging to data deleted from the backup source may not be deleted during the compact operation. As a result, the difference in size between the original and the new full backup files may not be equal to the size of the excluded data blocks. In this case, you can manually create an active full backup and reset the backup chain. To learn more about the active full backup, see [Active Full Backup](#).

To compact the full backup file periodically, you must enable the **Defragment and compact full backup file** option in the backup job settings and define the compact operation schedule. By default, the compact operation is performed on the last Saturday of every month. You can change the compact operation schedule and instruct Veeam Agent for Microsoft Windows to perform it weekly or monthly on specific days.



## Limitations for Full Backup File Compact

The full backup file compact has the following limitations:

- The **Defragment and compact full backup file** option can be enabled only for backup jobs for which active full and synthetic full backups are not scheduled.
- The **Defragment and compact full backup file** option is not available for backup jobs targeted at object storage.
- The compact full backup file operation is not performed during backup job sessions that produce active full backups. If the backup job starts again on the same day when the active full backup was created, Veeam Agent for Microsoft Windows does not perform the compact full backup operation. This limitation helps reduce the number of backup operations – Veeam Agent considers that the full backup is recent and does not need to be rebuilt.

If such situation occurs, Veeam Agent for Microsoft Windows triggers the full backup file compact operation during the next backup job session that produces an incremental backup file on another day.

- The target location must have enough space to store a file of the full backup size. During the compact process, Veeam Agent for Microsoft Windows creates auxiliary files that exist in the target location until the end of the compact operation.

- If you change the block size in backup job settings, Veeam Agent for Microsoft Windows does not change the block size in the compacted backup file till the next full backup. However, if you change compression settings in backup job settings, during the next compact file operation Veeam Agent for Microsoft Windows changes the compression level for the compacted backup file.
- The file compact operation is not performed during a backup job session that creates a restore point in the backup cache.
- The file compact operation is not performed during the backup cache synchronization process.

## Removal of Deleted Drives Data

During the compact operation, Veeam Agent for Microsoft Windows does not copy all data blocks from the VBK file to the newly created file. It copies only data blocks of Veeam Agent computer drives whose information is stored in the Veeam Agent for Microsoft Windows database. If a drive was removed from the Veeam Agent computer, its data is not copied to the new full backup file. This approach helps reduce the size of the full backup file and remove unnecessary data from it.

# Backup to Object Storage

You can store Veeam Agent backups in the following object storage types:

- S3 compatible (including WasabiCloud and IBM Cloud)
- Amazon S3
- Google Cloud Storage
- Microsoft Azure Blob Storage
- Veeam Data Cloud Vault added as a Veeam backup repository or Veeam Cloud Connect repository. To learn more, see [Backup Destinations](#).

Depending on your backup infrastructure, object storage can be available in different configurations. To learn more, see the following subsections:

- [Backup destinations](#)
- [Types of Connection to Object Storage in Veeam Backup & Replication](#)
- [Considerations and Limitations](#)

## Backup Destinations

You can back up Veeam Agent computer data to object storage in the following ways:

- Directly to object storage. In this case, Veeam Agent connects to the object storage and creates a backup repository in this storage.

Keep in mind that to connect to object storage, you need to specify an account with access permissions to read and write data.

To learn more, see [Object Storage Settings](#).

- To object storage added as a Veeam backup repository. In this case, an object storage repository is created in Veeam Backup & Replication. Veeam Agent connects to Veeam Backup & Replication and uses this Veeam backup repository.

To learn more, see [Veeam Backup Repository Settings](#).

- To object storage added as a Veeam Cloud Connect repository. In this case, a Veeam Cloud Connect service provider creates an object storage repository in Veeam Backup & Replication and exposes it as a cloud repository to tenants. Veeam Agent connects to the service provider and uses the Veeam Cloud Connect repository.

To learn more, see [Veeam Cloud Connect Repository Settings](#).

# Connection Types

If you back up data to object storage added as a Veeam backup repository or Veeam Cloud Connect repository, you must configure a repository beforehand on the Veeam Backup & Replication side. Depending on the repository configuration, Veeam Backup & Replication provides one of the following connection types to the repository in the object storage:

- Connection through a gateway server. With this connection type, Veeam Agent connects to the repository using a proxy component – a gateway server that is assigned in the Veeam Backup & Replication console. The backup data is transferred from the Veeam Agent computer to the gateway server, then it is transferred from the gateway server to the repository.
- Direct connection. With this connection type, Veeam Agent connects directly to the repository. The backup data is transferred from the Veeam Agent computer to the repository without proxy components. The access to this repository is managed by Application Programming Interface (API) that is provided by an external cloud service provider.

## Considerations and Limitations

Before you configure a backup job to store backups in object storage, consider the following:

- Veeam Agent does not support backup to object storage for which lifecycle rules are enabled. Enabling lifecycle rules may result in backup and restore failures.
- For backups located in object storage, synthetic full backup method is not supported.
- For backups located in object storage, compact full backup file option is not supported.
- Veeam Agent does not support direct backup to the Microsoft Azure Blob Storage under the general-purpose V1 storage account type.
- You can store backups only in those S3 compatible storage repositories that are accessible over the HTTPs protocol.
- [For object storage added as a Veeam backup repository or Veeam Cloud Connect repository] If you want to back up your data to the S3 compatible storage using a direct connection, you must additionally specify access permissions settings for the storage. To do so, enable the **Agents share credentials to object storage repository** or the **Provided by IAM/STS object storage capabilities** access control option. To learn more, see the [Managing Permissions for S3 Compatible Object Storage](#) section in the Veeam Backup & Replication User Guide.
- [For object storage added as a Veeam backup repository or Veeam Cloud Connect repository] Data recovery options are not available if you access the object storage repository using credentials with the read-only access permissions.
- If you have a backup job targeted at an object storage added as an extent of a scale-out backup repository in the direct connection mode and you put this extent to the Maintenance or Seal mode during the backup job session while there is no connection between Veeam Backup & Replication and the object storage, the current and subsequent backup job sessions will end successfully.

To learn about modes you can put extents of scale-out backup repositories to, see the [Service Actions with Scale-Out Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

# Backup Immutability

If you store your backup files in an object storage repository, Veeam Agent allows you to protect backup data from deletion or modification by making that data temporarily immutable. It is done for increased security: immutability protects data in your recent backups from loss as a result of attacks, malware activity or any other injurious actions.

## IMPORTANT

Backup immutability uses native object storage capabilities. You may incur additional API and storage charges from the storage provider.

## Supported Object Storage Types

Veeam Agent supports backup immutability for the following object storage types:

- Amazon S3
- S3 compatible storage that supports S3 Object Lock (including Wasabi)
- Microsoft Azure Blob Storage
- Veeam Data Cloud Vault

## NOTE

Veeam Agent does not support backup immutability for the Google Cloud storage.

## Before You Begin

Before you configure immutability for Veeam Agent backups, you must prepare the target storage account. Depending on the selected object storage type, perform the following actions:

- [S3 Compatible and Amazon S3 storage] When you create the S3 bucket, you must enable versioning and the S3 Object Lock feature for the bucket. For more information, see [AWS documentation](#).
- [S3 Compatible and Amazon S3 storage] After you create the S3 bucket with Object Lock enabled, make sure that the default retention is disabled to avoid unpredictable system behavior and data loss. To disable the default retention, edit the Object Lock retention settings as described in [AWS documentation](#).
- [Microsoft Azure Blob storage] You must enable blob versioning and version-level immutability support in the storage account. For more information, see [Microsoft documentation](#).

Consider the following about backup immutability:

- The effective immutability period consists of the user-defined immutability period and the block generation period automatically appended by Veeam Agent. For more information, see [How Backup Immutability Works](#) and [Block Generation](#).
- [S3 Compatible and Amazon S3 storage] Veeam Agent will use the *compliance* retention mode for each uploaded object. For more information on retention modes of S3 Object Lock, see [AWS documentation](#).
- [Microsoft Azure Blob storage] Do not enable immutability for already existing containers in the Microsoft Azure Portal. Otherwise, Veeam Agent will not be able to process these containers properly and it may result in data loss.

# Configuring Backup Immutability

Depending on how you create the backup job and configure [connection to object storage](#), you can define backup immutability settings in one of the following ways:

- [Backup Job wizard] You must specify the immutability period at the Bucket step of the wizard. For more information, see [Object Storage Settings](#).
- If you create the backup job that is targeted at an object storage repository configured as a Veeam backup repository or Veeam Cloud Connect repository, the immutability period in the settings of the repository must be specified in Veeam Backup & Replication. For details, see the [Adding Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

## Backup Immutability and Retention Policy

Backup immutability operates with backup data and related metadata (checkpoints) on the object storage side. Retention policy operates with logical representation of the stored data, or restore points, on the Veeam Agent side. These two mechanisms act independently from each other.

Veeam Agent will remove the irrelevant restore points per the defined backup retention policy. If the data associated with the removed restore point is still immutable, such data will remain in the repository until expiration of the immutability period. After that it will be automatically removed from the storage.

## Limitation of Backup Immutability

If you use Veeam Agent in the standalone mode, you can restore the immutable data that is associated with a restore point removed by retention policy only in Veeam Backup & Replication console. In Veeam Backup & Replication, you must perform the following actions:

1. Add the object storage repository that contains the necessary data to Veeam Backup & Replication. For more information, see the [Adding Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.
2. Roll back to the necessary checkpoint. For more information, see the [Immutability](#) section in the Veeam PowerShell Reference.
3. Remove the repository from the Veeam Backup & Replication infrastructure. For more information, see the [Removing Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

After that, you will be able to use Veeam Agent to restore data from the object repository in a regular manner.

## How Backup Immutability Works

After you specify the immutability period for a backup and run the backup job for the first time, Veeam Agent will append an additional period of 10 days to the specified immutability period. This additional period is called *block generation*. The resulting effective immutability period is the sum of the user-defined immutability period and the block generation period. All data blocks transferred to the target repository within the block generation period will have the same immutability expiration date. For example, data block *a* added on day 1 of the block generation period will have the same immutability expiration date as block *b* added on day 9. For more information, see [Block Generation](#).



During the effective immutability period, the following operations with backup data in the object storage repository will be prohibited:

- Manual removal of data from the backup repository.
- Removal of data by backup retention policy.
- Removal of data using any object storage provider tools.
- Removal of data by the technical support department of the object storage provider.

## Extension of Effective Immutability Period

During each transfer of data to the object storage repository, Veeam Agent creates a new checkpoint file with metadata that describes the latest state of the backup in the storage. The immutable blocks of data from a previous checkpoint may be reused in the newly created checkpoint. Veeam Agent keeps reused, or dependent, blocks of data locked by continuously assigning them to new generations and extending their effective immutability period. This guarantees that the effective immutability period is no less than the immutability period defined by user.

During data transfer, the effective immutability period for the backup is set as follows:

- [For new data blocks in the checkpoint] Immutability is set anew. The user-defined immutability period is appended with a 10-day block generation period.
- [For data blocks reused from the previous checkpoint] Immutability is extended to the immutability expiration date set for the new blocks.
- [For data blocks that are not reused in the checkpoint] Immutability is not extended. Such data blocks will remain in the repository until their immutability period is over. After that Veeam Agent will automatically remove them from the repository.

## Block Generation

When you specify an immutability period for the recent backups, Veeam Agent will automatically add 10 days to the immutability expiration date. This period is called *block generation*. The block generation period serves to reduce the number of requests to the object storage repository, which results in lower traffic and reduced storage costs. You do not have to configure it, the block generation period is applied automatically.

When the block generation period is appended to the user-defined immutability period, it means there is no need to extend the immutability period for old data blocks when adding new data blocks to the backup during that block generation period.

Consider this example. When you create a full backup to start a backup chain, all data blocks transferred to the object storage repository are new. For these new blocks of data, Veeam Agent will add the block generation period of 10 days to the specified immutability period. If the immutability period is set by user to the default period of 30 days, the effective immutability period with the added block generation period will become 40 days. The first full backup starts its generation that will last for 10 days. All new and reused data blocks within this block generation period will have the same immutability expiration date. For instance, a data block that was transferred to the target repository on day 9 will have the same immutability expiration date as a data block transferred on day 1. This mechanism guarantees that the effective immutability period for all the data blocks within a generation is no less than 30 days.

If a block generation period is over but data blocks from that generation are reused in the newly created checkpoint, their effective immutability period is automatically extended to ensure that the effective immutability period for all the data blocks in the new checkpoint is no less than the user-defined immutability period. For more information, see [How Backup Immutability Works](#).

# Backup to Deduplicating Storage Appliances

You can store backups in the deduplicating storage appliances added as backup repositories. To learn the full lists of supported appliances and their requirements and limitations, see the [Deduplicating Storage Appliances](#) section in the Veeam Backup & Replication User Guide.

If you want to use immutability with the deduplicating storage appliances, consider the limitations listed in the following subsections:

- [HPE StoreOnce Immutability and Veeam Agents](#)
- [Dell Data Domain Immutability and Veeam Agents](#)

## HPE StoreOnce Immutability and Veeam Agents

If you want to use immutability with HPE StoreOnce Catalyst repositories, make sure that the value for the Maximum ISV Controlled Data Retention setting in HPE StoreOnce exceeds or is equal to the value of each of the following settings configured in Veeam Backup & Replication:

- The immutability period specified in the backup repository settings.
- The long-term retention period configured in backup job settings.

If the value of at least one of these periods exceeds the maximum value set for immutability in HPE StoreOnce Catalyst Store, consider the following:

- If the immutability period specified in the backup repository settings exceeds the maximum value set for immutability in HPE StoreOnce Catalyst Store, Veeam Backup & Replication applies immutability to the created backups according to the maximum value set for immutability in HPE StoreOnce Catalyst Store.
- If the long-term retention period configured in backup job settings exceeds the maximum value set for immutability in HPE StoreOnce Catalyst Store, the immutability period for backups with GFS flags is set according to the maximum value set for immutability in HPE StoreOnce Catalyst Store. For example, if the backup is stored for 1 year, but the maximum value for immutability in HPE StoreOnce Catalyst Store is set to 6 months, the backups will be immutable for 6 months.

To learn more about immutability for the HPE StoreOnce Catalyst repositories, see the [HPE StoreOnce and Immutability](#) section in the Veeam Backup & Replication User Guide.

## Dell Data Domain Immutability and Veeam Agents

If you want to use immutability with the Dell Data Domain backup repository, make sure that the values of the following settings configured in Veeam Backup & Replication lie in the range between the minimum and maximum retention periods configured in Dell Data Domain:

- The immutability period specified in the backup repository settings.
- The long-term retention period configured in backup job settings.

The minimum and maximum values are included into the range.

If the value of at least one of these periods lies outside the established range, consider the following:

- If the immutability period specified in the backup repository settings lies outside the range between the minimum and maximum retention periods configured in Dell Data Domain, the immutability for backups is set equal to the nearest range value.

- If the long-term retention period configured in backup job settings lies outside the range between the minimum and maximum retention periods configured in Dell Data Domain, the immutability period for backups with GFS flags is set equal to the nearest range value. For example, if the backup is stored for 1 year, but the range for immutability in Dell Data Domain is set to from 2 to 6 months, the backups will be immutable for 6 months.

To learn more about immutability for the Dell Data Domain backup repositories, see the [Retention Lock](#) section in the Veeam Backup & Replication User Guide.

# Health Check for Backup Files

You can instruct Veeam Agent for Microsoft Windows to periodically perform the health check for your backups.

Veeam Agent offers two different health check mechanisms:

- In the first case, Veeam Agent performs a CRC check for metadata and a hash check for data blocks in the backup file to verify their integrity. The health check verifies the latest restore point in the backup chain to make sure that the restore point is consistent, and you will be able to restore data from this restore point. To learn more, see [Standard Health Check](#).
- In the second case, Veeam Agent does not perform a hash check for data blocks, but verifies metadata for the whole backup, not just the latest restore point. This mechanism uses less traffic than the standard health check as it does not read data from data blocks. This type of health check is intended and is the default option for object storage targets. To learn more, see [Health Check for Object Storage](#).

## Health Check Schedule

To run the health check periodically, enable the **Perform backup files health check** option in the backup job settings and define the health check schedule. By default, the health check is performed on the last Friday of every month. You can change the schedule and run the health check weekly or monthly on specific days.

When you configure the health check schedule, consider the following:

- The health check runs automatically during the first incremental backup job session on the days specified in the health check schedule. If the backup job runs several times on a specified day, the health check is performed only with the first run of the backup job on that day. For example, if the job is scheduled to run several times on Saturday, and the health check is scheduled on Saturday, the health check will only be performed during the first backup job session on Saturday.

The health check is not performed during the first full backup or subsequent active full backup job sessions.

- If Veeam Agent does not run any backup jobs on the day specified in the health check schedule, the health check will be performed during the first backup job session following that day.

For example, you scheduled to run the health check every last day of a month, while the backup job is scheduled to run every day and create an active full backup on Sundays. If the last day of a month falls on a Sunday, the health check will be performed on the following Monday with the first incremental backup job session on that day.

## Limitations for Health Check

- The health check is not performed during an active full backup job session started manually or automatically by schedule.
- The health check is not performed during a backup job session that creates a restore point in the backup cache.
- The health check is not performed during the backup cache synchronization process.

## Health Check Retries

The health check itself is started during the backup job session or the job retry session if the backup job session has failed. If the attempts are not successful, Veeam Agent for Microsoft Windows will perform the health check during the last job retry in any case.

If the health check detects corrupted data, Veeam Agent for Microsoft Windows completes the backup job with the *Error* status and starts the health check retry process. The health check retry starts as a separate backup job session. During the health check retry, Veeam Agent attempts to transport data blocks for the corrupted restore point from the Veeam Agent computer to the target location.

For a backup job started automatically upon schedule, the number of health check retries is equal to the number of job retries specified in the job settings. For a job started manually, Veeam Agent for Microsoft Windows performs 1 health check retry.

#### NOTE

Consider the following:

- If Veeam Agent for Microsoft Windows fails to fix the corrupted data during all health check retries, you must retry the job manually. In this case, Veeam Agent will transport the necessary data blocks from the Veeam Agent computer to fix the corrupted restore point.
- If you try to restore data from a corrupted restore point, Veeam Agent will display a warning message informing you that the restore operation may fail or the restored data may be corrupted.

## Standard Health Check

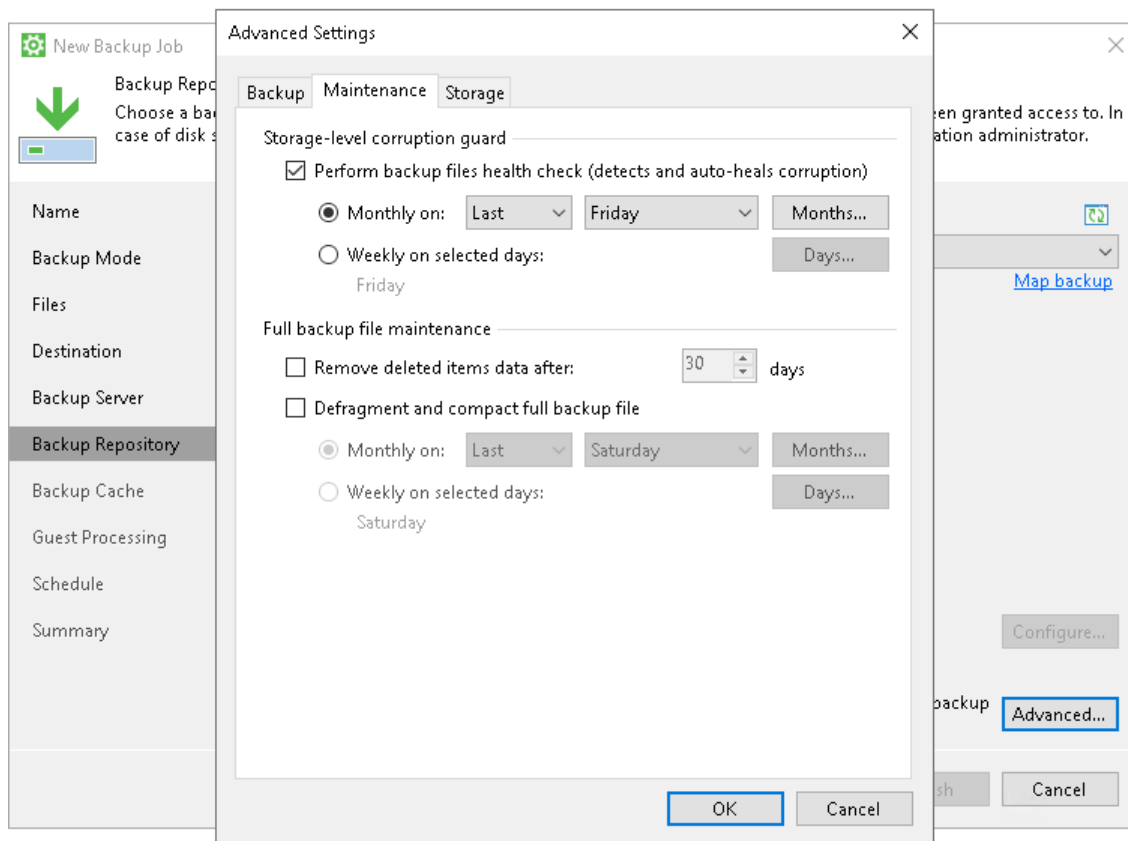
During the standard health check, Veeam Agent performs a CRC check for metadata and a hash check for data blocks in the backup file to verify their integrity. The health check helps make sure that the restore point is consistent, and you will be able to restore data from this restore point.

You can use the standard health check for all backup targets.

#### NOTE

Keep in mind that for the object storage targets, Veeam Agent offers another health check mechanism as default. To learn more, see [Health Check for Object Storage](#).

To run the standard health check periodically, enable the **Perform backup files health check** option in the backup job settings and define the health check schedule.



## How Standard Health Check Works

When Veeam Agent for Microsoft Windows saves a new restore point to the backup location, it calculates CRC values for backup metadata and hash values for data blocks in the backup file, and saves these values in the metadata of the backup file, together with the backed-up data. During the health check session, Veeam Agent uses these values to make sure that a verified restore point is consistent.

### NOTE

If the backup job is targeted at a Veeam backup repository or cloud repository, and you perform the health check for the encrypted backup files, Veeam Backup & Replication will pass encryption keys to the Veeam backup repository or cloud repository. For more information on encryption, see [Data Encryption](#).

Veeam Agent for Microsoft Windows performs the health check in the following way:

1. During the backup job session, Veeam Agent for Microsoft Windows creates a new restore point in the backup chain.
2. At the end of the backup job session, Veeam Agent for Microsoft Windows performs the health check. It calculates CRC values for backup metadata and hash values for data blocks in the backup file and compares them with the CRC and hash values that are already stored in the backup file.

During the health check, Veeam Agent for Microsoft Windows verifies the latest restore point in the backup chain (restore point created with the current backup job session – the session during which the health check is performed). If the latest restore point in the backup chain is incomplete (for example, if the current backup job session completed with an error), Veeam Agent checks the restore point preceding the latest one.

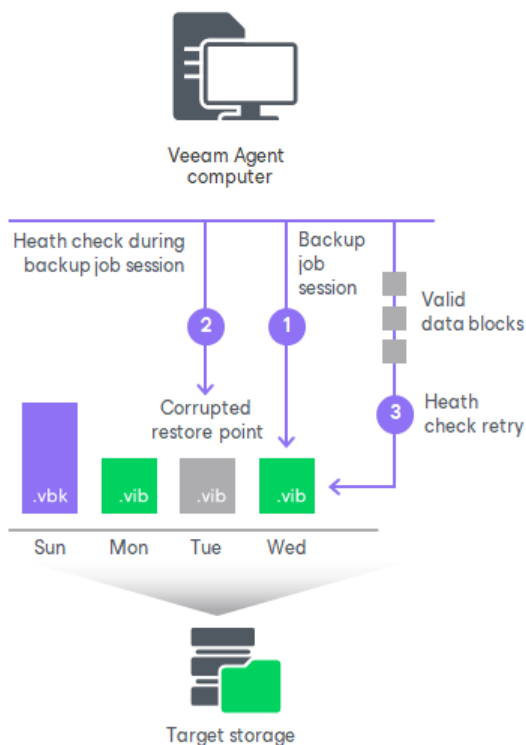
If the health check does not detect data corruption, the backup job session completes in a regular way.

3. If the health check detects corrupted data, Veeam Agent for Microsoft Windows completes the backup job with the *Error* status and starts the health check retry process. The health check retry starts as a separate backup job session.

Depending on the revealed data corruption, Veeam Agent performs the following actions:

- If the health check detects corrupted backup metadata in the full backup file, Veeam Agent marks the backup chain starting from this full restore point as corrupted in the Veeam Agent for Microsoft Windows database. During the health check retry, Veeam Agent transports the entire data selected for backup from the Veeam Agent Computer, creates a new full backup file in the target location and saves transported data blocks to this backup file.
- If the health check detects corrupted backup metadata in the incremental backup file, Veeam Agent removes information about this incremental restore point and subsequent incremental restore points from the Veeam Agent for Microsoft Windows database. During the health check retry, Veeam Agent transports incremental data relatively to the latest valid restore point in the backup chain from the Veeam Agent computer, creates a new incremental backup file in the target location and saves transported data blocks to this backup file.
- If the health check detects corrupted data blocks in the full or incremental backup file, Veeam Agent marks the restore point that includes the corrupted data blocks and subsequent incremental restore points as corrupted in the Veeam Agent for Microsoft Windows database. During the health check retry, Veeam Agent transports data blocks from the Veeam Agent computer. In addition, Veeam Agent transports data blocks that have changed since the backup job session that has triggered the health check. Veeam Agent stores these data blocks to the latest restore point that has been created with the current backup job session (session that has triggered the health check retry).

You can view the health check result in the restore point statistics. If the health check finds corrupted data, it will display information on where corrupt data has been detected — in metadata or blockstore, as well as list all restore points that share the corrupted data blocks.



# Health Check for Object Storage

For backup jobs targeted at object storage, Veeam Agent offers a special health check mechanism that differs from the standard health check in the following ways:

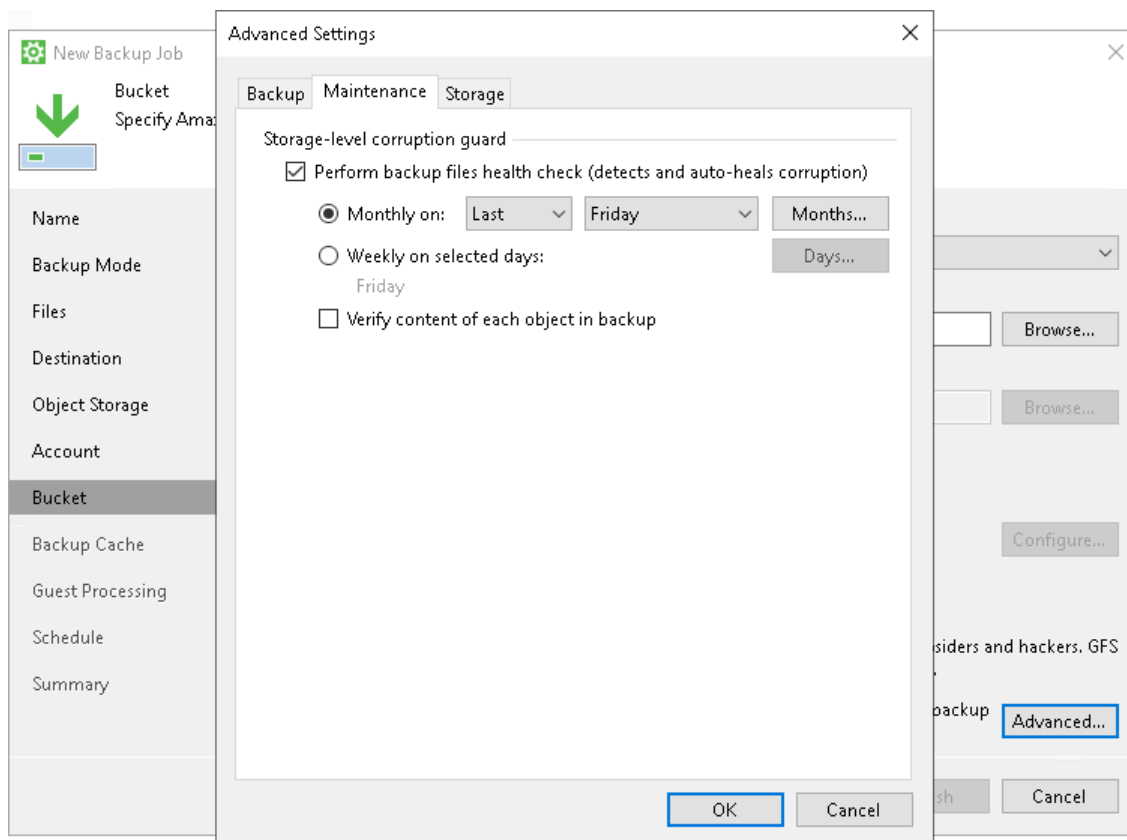
- The health check for object storage verifies metadata for the whole backup, not just the latest restore point.
- The health check for object storage does not read data from data blocks; it only lists data blocks to make sure all blocks in the storage are available for rebuilding every restore point in the backup chain.

The health check for object storage works faster than the standard health check and helps save traffic.

## NOTE

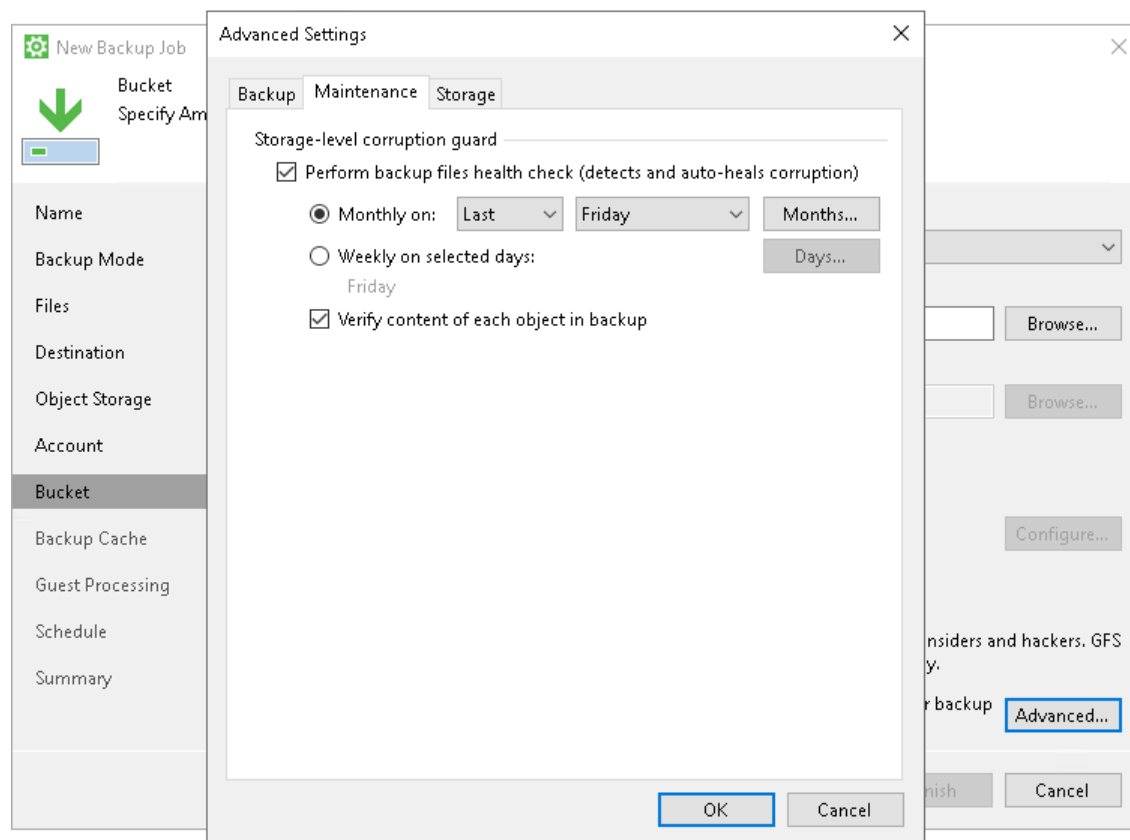
The health check mechanism described in this section is only available for object storage targets.

To run the special health check for object storage periodically, select the **Perform backup files health check** check box in the backup job settings and define the health check schedule. Make sure that the **Verify content of each object in backup** check box is not selected.





If necessary, you can enable the standard health check mechanism for object storage targets. To do so, select the **Verify content of each object in backup** check box in the backup job settings. Keep in mind that enabling this setting may result in additional charges from your object storage provider.



## How Health Check for Object Storage Works

Veeam Agent performs the health check of a backup in the following way:

1. During the backup job session after a new incremental backup file is created, Veeam Agent starts the health check of the whole backup. Veeam Agent checks if the metadata of the backup is consistent and no metadata is missing. Veeam Agent also checks if all data blocks for every restore point are available in the storage. Veeam Agent does not read data from data blocks.
2. If Veeam Agent does not find any corrupted data, the health check completes successfully.

If Veeam Agent detects corrupted data, the health check completes with an error. Depending on the detected data inconsistency, Veeam Agent behaves in one of the following ways:

- If the health check detects corrupted metadata, Veeam Agent marks the backup as corrupted in the Veeam Agent configuration database; the backup job session fails. During the next scheduled or manual backup job session, Veeam Agent will create a full backup and will start a new backup chain. The corrupted backup will become orphaned and will remain in the repository – you can keep or delete it.

- If the health check detects corrupted data blocks in the latest restore point in the backup chain, Veeam Agent launches the health check retry.

During the health check retry, Veeam Agent restarts the backup job to create a new restore point and transports data blocks from the Veeam Agent computer including the blocks that contain corrupted data and the blocks that have changed since the backup job session that has triggered the health check. Veeam Agent will not perform another health check after the job retry is finished; the next health check will run according to the defined schedule during another incremental backup job session.

- If the health check detects corrupted data blocks in an inactive backup chain, Veeam Agent does not launch the health check retry. Veeam Agent marks the backup and all related restore points as corrupted. The backup job session ends with a warning message.

You can view the health check result in the restore point statistics. If the health check finds corrupted data, it will display information on where corrupt data has been detected — in metadata or blockstore, as well as list all restore points that share the corrupted data blocks.

## NOTE

Consider the following:

In case of immutable backups, Veeam Agent performs the health check only for valid restore points according to the retention policy. Immutable data associated with removed restore points can still remain in the repository depending on the immutability period, but Veeam Agent will not perform the health check for such data.

For information about backup immutability settings, see [Backup Immutability](#).

# Changed Block Tracking

To perform incremental backup, Veeam Agent for Microsoft Windows needs to know what data blocks have changed since the previous job session. To get the list of changed data blocks, Veeam Agent for Microsoft Windows uses the changed block tracking mechanism, or CBT. CBT increases the speed and efficiency of incremental backups.

To keep track of changed data blocks, Veeam Agent for Microsoft Windows can use the following mechanisms:

- **Default CBT mechanism** – this mechanism is enabled by default in all installations of Veeam Agent for Microsoft Windows. To learn more, see [Default Changed Block Tracking Mechanism](#).
- **Veeam CBT driver** – this driver helps Veeam Agent for Microsoft Windows keep track of changed data blocks in a more efficient way. This functionality is available if the Veeam Agent computer meets a list of requirements. To learn more, see [Veeam Changed Block Tracking Driver](#).

## Default Changed Block Tracking Mechanism

Veeam Agent for Microsoft Windows performs the changed block tracking differently depending on the backup type:

- [CBT for volume-level backup](#)
- [CBT for file-level backup](#)

## CBT for Volume-Level Backup


In case of the volume-level backup, Veeam Agent for Microsoft Windows performs changed block tracking in the following way:

1. During the full backup job session, Veeam Agent reads the Master File Table (MFT) of the backed-up volume. Veeam Agent uses MFT records to create digests with file system metadata, transfers the created digests to the target location and stores them to the resulting backup file.
2. During subsequent incremental job sessions, Veeam Agent performs the following operations:
  - a. Reads the Master File Table (MFT) of the backed-up volume and creates the new digests with file system metadata.
  - b. Interacts with the target backup location to obtain digests from the backup file that was created during the previous job session.
  - c. Compares new and previous digests to detect files whose data blocks have changed on the volume since the previous job session.

During incremental backup, Veeam Agent for Microsoft Windows reads from the VSS snapshot only data blocks pertaining to files that have changed since the previous job session. If Veeam Agent cannot calculate information about the changed files, for example, if it fails to retrieve digests from the backup file, Veeam Agent will need to read all data blocks from the VSS snapshot. As a result, the backup may take significantly more time.

## NOTE

Veeam Agent for Microsoft Windows uses the default CBT mechanism for NTFS volumes only. As a result, for volumes that use other file systems, incremental backup will require greater time, because Veeam Agent for Microsoft Windows will read all data from the VSS snapshot to detect what blocks have changed since the last job session.



## CBT for File-Level Backup

In case of the file-level backup, Veeam Agent for Microsoft Windows performs changed block tracking in the following way:

1. During the full backup job session, Veeam Agent creates a new NTFS partition in the backup file. Veeam Agent uses this partition to store all backed-up files and folders.
2. During subsequent incremental job sessions, Veeam Agent performs the following operations:

- a. Compares the last modification time attribute of files on the Veeam Agent computer and files in the backup. This operation allows Veeam Agent to detect files that have changed since the previous job session.

- b. Transfers data blocks that contain metadata related to each changed file and its parent folders from the target backup location to the Veeam Agent computer.


If the backed-up file system has a complex folder structure with many hierarchy levels, and you back up data to a remote location, during incremental backup, the inbound network traffic on the Veeam Agent computer may exceed by far the outbound traffic. Significant amount of data can be transferred to the Veeam Agent computer even if few files are changed since the previous job session.

- c. Replaces file system metadata in the downloaded data blocks with the current file system metadata.
- d. Transfers data blocks with current metadata from the Veeam Agent computer to the target backup location along with changed files. In the target backup location, Veeam Agent stores data and metadata to the newly created backup file.

## NOTE

Consider the following:

- Veeam Agent for Microsoft Windows detects changed files using the last modification time attribute. If you change the file in any way, but the last modification time attribute remains unchanged, Veeam Agent will not back up this file during an incremental job session.
- If the backed-up file has changes, Veeam Agent will copy to the target location not only changed data blocks of the file but the entire file. To save space on the target storage, consider using the volume-level backup. To learn more, see [Backup Types](#).



## Veeam Changed Block Tracking Driver

You can set up Veeam Agent for Microsoft Windows to use the Veeam CBT driver instead of the default CBT mechanism. The Veeam CBT driver is a class filter driver for volume devices that helps Veeam Agent keep track of changed data blocks in a more efficient way. The driver is intended for computers running applications with large database files. The Veeam CBT driver is supported for volume-level and file-level backups.

## NOTE

In case of the file-level backup, consider the following:

- Veeam Agent does not use the Veeam CBT driver to back up deduplicated files.
- Veeam Agent uses the Veeam CBT driver to back up files of 50 MB and larger including files with unchanged last modification time attribute.
- When you run a backup job for the first time after switching to the Veeam CBT driver, the job may take a significant time to finish. This happens because the Veeam CBT driver cannot retrieve all necessary information about changed data blocks from backup files created with the default CBT mechanism. In this case, files for backup of which Veeam Agent uses the Veeam CBT driver are backed up as a whole.
- The Veeam CBT driver is designed to track data blocks that changed since the last run of the backup job. If you work with a file in an application that downloads the entire file into RAM, after you save the changes, all file blocks will be overwritten. Veeam Agent backs up such a file as a whole upon the next run of the backup job.

To use the Veeam CBT driver, the Veeam Agent computer must meet the following requirements:

- Run a 64-bit version of Microsoft Windows OS.
- Run one of the following OSes:
  - Microsoft Windows 11 (from version 21H2 to version 23H2).
  - Microsoft Windows 10 (from version 1803 to version 22H2).
  - Microsoft Windows Server OS that is supported by Veeam Agent. For more information, see [System Requirements](#).

- Run the Workstation or Server edition of Veeam Agent for Microsoft Windows.

## IMPORTANT

Consider the following:

- Prior to installing the Veeam CBT driver on a computer running Microsoft Windows Server 2008 R2 SP1, make sure that update [KB3033929](#) is installed in the OS.  
The update adds the SHA-2 code signing support that is required for verification of the Veeam CBT driver signature. Without this update installed, the OS running on a protected computer will fail to boot after you install the Veeam CBT driver. To learn more, see [this Microsoft KB article](#).
- Do not install the Veeam CBT driver on a computer running Microsoft Windows Server 2008 R2 SP1, 2012 or 2012 R2 if one or more volumes on this computer are encrypted with Microsoft BitLocker (or other encryption tool), or if you plan to use Microsoft BitLocker to encrypt volumes on this computer. Concurrent operation of Microsoft BitLocker and Veeam CBT driver may result in driver failures and may prevent the OS from starting.
- Do not install the Veeam CBT driver on a computer if you plan to use devices with hardware encryption made according to the TCG Opal Security Subsystem Class Specification. Operation of the driver on such devices may lead to a crash of the operating system. To learn more about the TCG Opal Security Subsystem Class Specification, see [this Trusted Computing Group webpage](#).

To enable the advanced CBT mechanism provided by the Veeam CBT driver, you need to install the driver in the Veeam Agent control panel. You can perform this operation at any time you need. To activate the driver after installation, Veeam Agent needs to reboot the computer. After computer reboot, the Veeam CBT driver will start keeping track of changed data blocks on computer volumes whose data you have selected for backup in the Veeam Agent backup job settings.

In contrast to the default CBT mechanism that supports NTFS volumes only, the Veeam CBT driver can keep track of changed data blocks on volumes that use the following file systems:

- FAT (only for volume-level backups)
- NTFS
- ReFS

Information about changed data blocks is registered in special VCT files. VCT files are stored in the `C:\ProgramData\Veeam\EndpointData\CtStore` folder on the Veeam Agent computer. When the backup job runs, Veeam Agent for Microsoft Windows uses VCT files to find out what data blocks have changed since the last run of the job, and copies only changed data blocks from the backed-up volume.

If you use the Veeam CBT driver, the `C:\ProgramData\Veeam\EndpointData\CtStore` folder accessibility is crucial for a successful backup. Consider the following:

- The folder and its content must not be compressed or encrypted.
- The folder and its parent folders must not be reparse points.

## NOTE

Consider the following:

- If the Veeam Agent computer shuts down unexpectedly, the Veeam CBT driver may fail to register information about changed data blocks in a VCT file. In this case, during the next backup job session, Veeam Agent for Microsoft Windows will need to read all data from the backed-up volume to create incremental backup. As a result, incremental backup will require greater time.
- If data blocks are changed on a volume while this volume is mounted on another Windows-based machine, during the next backup job session, Veeam Agent for Microsoft Windows will also read all data from the volume to create incremental backup.
- The Veeam CBT driver cannot detect data block changes made on a volume that is mounted in a non-Windows OS. For example, such changes can be made when you boot your computer using a Linux-based antivirus rescue disk. To continue the backup chain after such changes, you need to create active full backup instead of incremental backup. Alternatively, you can reset CBT. To learn more, see [Resetting CBT](#).

# Data Compression

Veeam Agent provides mechanisms of data compression and deduplication. Data compression and deduplication let you decrease traffic going over the network and disk space required for storing backup files.

## Data Compression

Data compression decreases the size of created backups but affects duration of the backup procedure. Veeam Agent allows you to select one of the following compression levels:

Compression Level	Compression Algorithm	Description
None	No compression	This compression level is recommended if you plan to store backup files on storage devices that support hardware compression and deduplication.
Dedupe-friendly	Rle	Optimized compression level for very low CPU usage. You can select this compression level if you want to decrease the load on the CPU of the Veeam Agent computer.
Optimal	Lz4	The default recommended compression level. It provides the best ratio between size of the backup file and time of the backup procedure.
High	Zstd 3	Provides up to 60% additional compression ratio over the Optimal level at the cost of 2x higher CPU usage and 2x slower restore.
Extreme	Zstd 9	Provides the smallest size of the backup file but reduces the backup performance. We recommend that you use the extreme compression level only on Veeam Agent computers with modern multi-core CPUs (6 cores recommended).

You can change data compression settings for existing backup jobs. New settings will not have any effect on previously created backup files in the backup chain. They will be applied to new backup files created after the settings were changed.

Compression settings are changed on the fly. You do not need to create a new full backup to use new settings — Veeam Agent will automatically apply the new compression level to newly created backup files.

## Storage Optimization

Depending on the type of storage you select as a backup target, Veeam Agent uses data blocks of different size, which optimizes the size of a backup file and job performance. You can choose one of the following storage optimization options:

- **4 MB** — select this option for backup jobs that can produce very large full backup files (larger than 16 TB). With this option selected, Veeam Agent will use data block size of 4096 KB.



- **1 MB** (default) – select this option for backup to SAN, DAS or local storage. With this option selected, Veeam Agent will use data block size of 1024 KB.

The SAN identifies larger blocks of data and therefore can process large amounts of data at a time. This option provides the fastest backup job performance.

- **512 KB** – select this option for backup to NAS and onsite backup. With this option selected, Veeam Agent will use data block size of 512 KB. This option reduces the size of an incremental backup file because of reduced data block sizes.
- **256 KB** – select this option if you plan to use WAN for offsite backup. With this option selected, Veeam Agent will use data block size of 256 KB. This results in the smallest size of backup files, allowing you to reduce the amount of traffic over WAN.

#### NOTE

If you change storage optimization settings, the new settings will be applied only after an active full backup is created. Veeam Agent will use the new block size for the active full backup and subsequent backup files in the backup chain. For more information on scheduling active full backups, see [Backup Settings](#).

# Guest Processing

For Server edition of Veeam Agent for Microsoft Windows, you can specify guest processing options. Veeam Agent for Microsoft Windows offers the following guest processing options:

- [Application-aware processing](#). You can create transactionally consistent backups of servers running applications that support Microsoft VSS. Application-aware processing guarantees that you can perform restore from Veeam Agent backups without data loss.
- [Pre-freeze and post-thaw scripts](#). You can use pre-freeze and post-thaw scripts to quiesce applications that do not support Microsoft VSS.
- [Transaction log truncation](#). You can set up the backup job to truncate transaction logs after the job successfully completes.
- [Transaction log backup for Microsoft SQL Server and Oracle](#). You can set up the backup job to back up transaction logs from servers running Microsoft SQL Server and archived logs of Oracle database systems.
- [File system indexing](#). You can set up the backup job to create a catalog of files and folders on the Veeam Agent computer OS. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you will be able to search for individual files in Veeam Agent backups and perform 1-click restore in Veeam Backup Enterprise Manager.

File system indexing is optional. If you do not enable this option in the backup job settings, you will still be able to perform 1-click restore from the Veeam Agent backup. For more information, see the [Veeam Backup Enterprise Manager User Guide](#).

## Supported Applications

Veeam Agent for Microsoft Windows supports VSS-aware processing for the following systems:

Specification	Requirement
<b>Microsoft Active Directory Domain Controllers</b>	<p>The following versions of Microsoft Active Directory Domain Services servers (domain controllers) are supported:</p> <ul style="list-style-type: none"><li>• Microsoft Windows Server 2022</li><li>• Microsoft Windows Server 2019</li><li>• Microsoft Windows Server 2016</li><li>• Microsoft Windows Server 2012 R2</li><li>• Microsoft Windows Server 2012</li><li>• Microsoft Windows Server 2008 R2 SP1</li></ul> <p>Minimum supported domain and forest functional level is Windows Server 2003.</p>
<b>Microsoft Exchange</b>	<p>The following versions of Microsoft Exchange are supported:</p> <ul style="list-style-type: none"><li>• Microsoft Exchange 2019</li><li>• Microsoft Exchange 2016</li><li>• Microsoft Exchange 2013 SP1</li><li>• Microsoft Exchange 2013</li></ul>

Specification	Requirement
<b>Microsoft SharePoint</b>	<p>The following versions of Microsoft SharePoint Server are supported:</p> <ul style="list-style-type: none"> <li>• Microsoft SharePoint Server Subscription Edition</li> <li>• Microsoft SharePoint Server 2019</li> <li>• Microsoft SharePoint Server 2016</li> <li>• Microsoft SharePoint Server 2013</li> </ul> <p>All editions are supported (Foundation, Standard, Enterprise).</p>
<b>Microsoft SQL Server</b>	<p>The following versions of Microsoft SQL Server are supported:</p> <ul style="list-style-type: none"> <li>• Microsoft SQL Server 2022</li> <li>• Microsoft SQL Server 2019</li> <li>• Microsoft SQL Server 2017</li> <li>• Microsoft SQL Server 2016</li> <li>• Microsoft SQL Server 2014</li> <li>• Microsoft SQL Server 2012</li> <li>• Microsoft SQL Server 2008 R2</li> <li>• Microsoft SQL Server 2008</li> </ul> <p>All editions of Microsoft SQL Server except LocalDB are supported.</p>
<b>Oracle</b>	<p>Oracle Database from version 11g to version 21c is supported for the following operating systems (32-bit and 64-bit architecture):</p> <ul style="list-style-type: none"> <li>• Microsoft Windows Server 2019</li> <li>• Microsoft Windows Server 2016</li> <li>• Microsoft Windows Server 2012 R2</li> <li>• Microsoft Windows Server 2012</li> <li>• Microsoft Windows Server 2008 R2 SP1</li> </ul> <p><b>Important notes:</b></p> <ul style="list-style-type: none"> <li>• Automatic Storage Management (ASM) is not supported.</li> <li>• Oracle Real Application Clusters (RAC) are not supported.</li> <li>• Oracle servers using Data Guard are not supported.</li> <li>• Oracle Database Express Edition is supported.</li> <li>• Configurations with different versions of Oracle Database deployed on the same server are not supported.</li> <li>• 32-bit Oracle running on 64-bit operating systems is not supported.</li> </ul>

## Application-Aware Processing

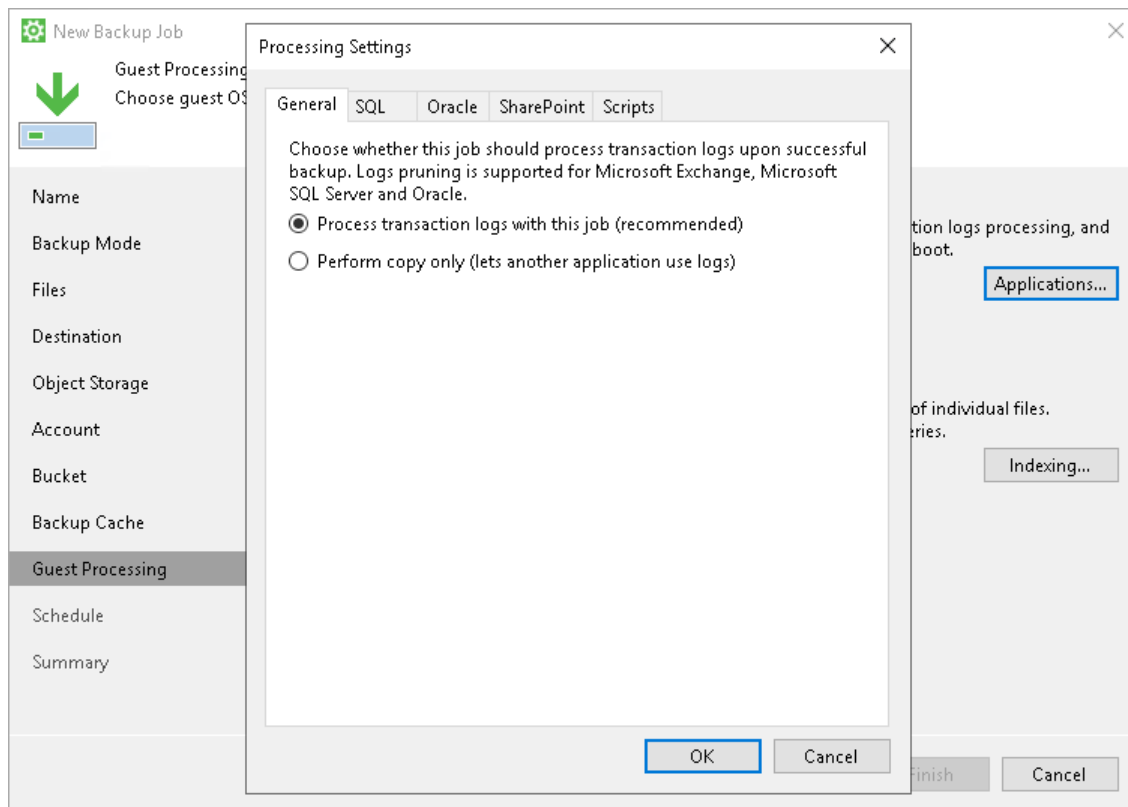
To create transactionally consistent backups of servers that run VSS-aware applications such as Microsoft SQL Server, Microsoft SharePoint, Microsoft Exchange or Oracle, you must enable application-aware processing for the backup job.

Application-aware processing is Veeam's proprietary technology based on Microsoft VSS. Microsoft VSS is responsible for quiescing applications and creating a consistent view of application data on the OS of the Veeam Agent computer. Use of Microsoft VSS ensures that there are no unfinished database transactions or incomplete application files when Veeam Agent requests the creation of a Microsoft VSS snapshot and starts copying backed-up data to the target location. For more information about Microsoft VSS, see [Microsoft documentation](#).

In the Server edition of Veeam Agent, the type of the VSS snapshot depends on application-aware processing settings:

- If application-aware processing is disabled for the backup job, Veeam Agent requests a copy-only VSS snapshot.
- If application-aware processing is enabled and the **Perform copy only** option is selected, Veeam Agent requests a copy-only VSS snapshot.
- If application-aware processing is enabled and the **Process transaction logs with this job** option is selected, Veeam Agent requests a full VSS snapshot.

To learn more, see [Specify Application-Aware Processing Settings](#).



Application-aware processing is supported for Microsoft Windows 2008 R2 SP1 and later. To use application-aware processing, you must have the latest updates installed on the Veeam Agent computer OS. To learn more, see [Supported Applications](#).

## IMPORTANT

If your computer OS runs an application that does not support Microsoft VSS (there is no VSS writer for this particular type of application, for example, MySQL), Veeam Agent for Microsoft Windows will not be able to utilize Microsoft VSS and application-aware processing for this application. To process such applications, you can use pre-freeze and post-thaw scripts. For more information, see [Pre-Freeze and Post-Thaw Scripts](#).

# Pre-Freeze and Post-Thaw Scripts

If Veeam Agent computer runs applications that do not support Microsoft VSS, you can instruct Veeam Agent for Microsoft Windows to run custom scripts during the backup job session. For example, the pre-freeze script may quiesce the file system and application data to bring the computer OS to a consistent state before Veeam Agent for Microsoft Windows requests the creation of a Microsoft VSS snapshot. After the VSS snapshot is created, the post-thaw script may bring the OS and applications to their initial state.

Veeam Agent for Microsoft Windows supports the following types of scripts:

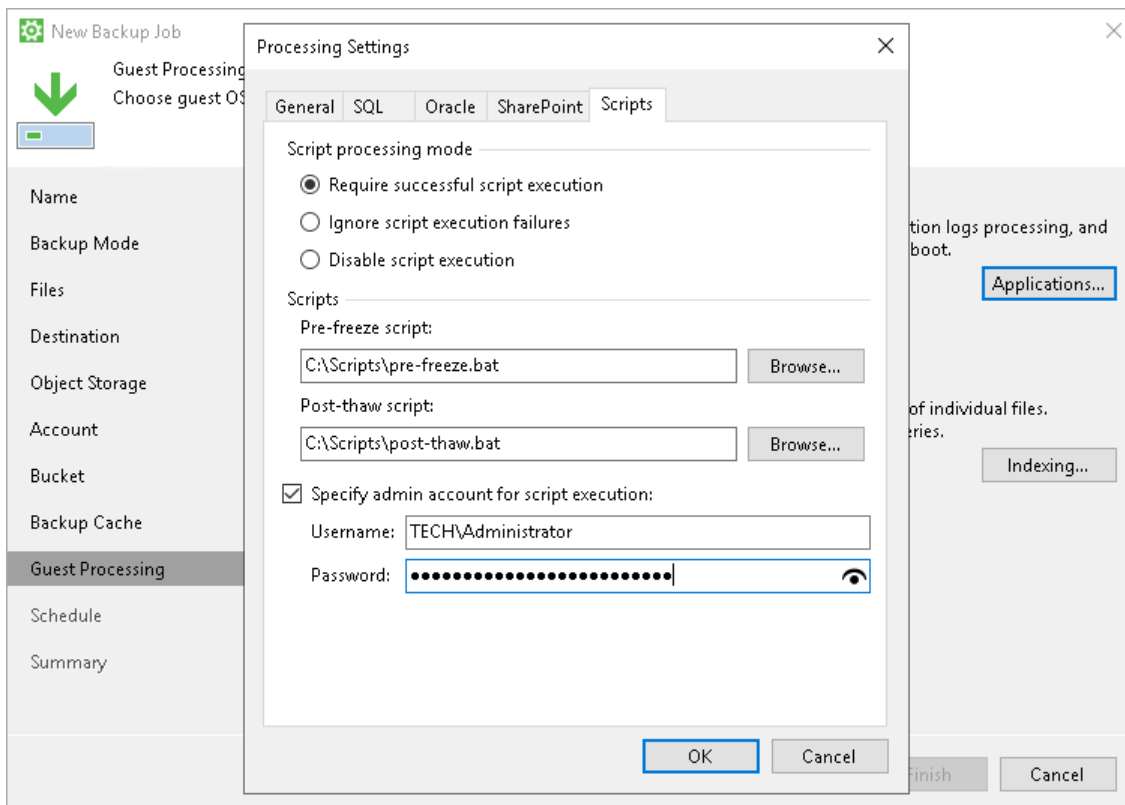
- Program files in the EXE, BAT and CMD format
- Windows script files in the JS, VBS and WSF format
- PowerShell script files in the PS1 format

You can use scripts of other formats as well, but we cannot guarantee correct processing of such scripts.

Scripts must be created beforehand. You must specify paths to them in the backup job settings. Scripts must reside on a local drive of the Veeam Agent computer.

When the backup job starts, Veeam Agent for Microsoft Windows executes scripts specified for the job. A script is considered to be executed successfully if "0" is returned.

The default time period for script execution is 10 minutes. If the script fails to execute before the timeout expires, Veeam Agent for Microsoft Windows displays an error message in the job session and error or warning messages issued during script execution.



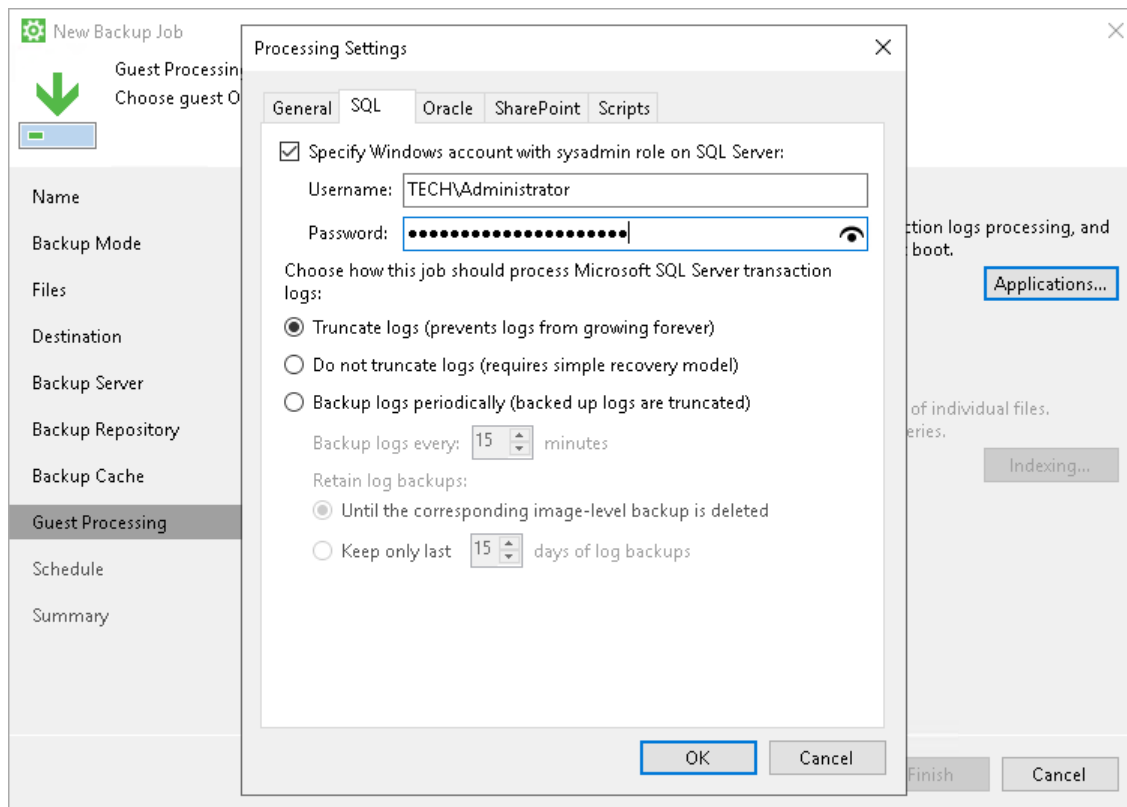
# Limitations for Pre-Freeze and Post-Thaw Scripts

Veeam Agent for Microsoft Windows has one limitation for pre-freeze and post-thaw scripts: you cannot stop a job when the pre-freeze or post-thaw script is executed. If the script hangs up, Veeam Agent for Microsoft Windows waits for 10 minutes and then continues running the job depending on the script processing mode.

## Transaction Log Truncation

If you back up database systems that use transaction logs, for example, Microsoft SQL Server, you can instruct Veeam Agent for Microsoft Windows to truncate transaction logs so that logs do not overflow the storage space. Veeam Agent for Microsoft Windows provides the following options of transaction logs handling:

- [Truncate logs](#)
- [Do not truncate logs](#)
- [Back up logs periodically](#)



## Truncate Logs

You can instruct Veeam Agent for Microsoft Windows to truncate logs after a backup is successfully created. With this option selected, Veeam Agent for Microsoft Windows behaves in the following way:

- If the backup job completes successfully, Veeam Agent for Microsoft Windows produces a backup file and truncates transaction logs on the Veeam Agent computer. As a result, you have the backup file that contains a computer image, image of a specific data volume or individual folders at a specific point in time.

In this scenario, you can recover a database to the point in time when the backup file was created. As transaction logs on the Veeam Agent computer are truncated, you cannot use them to get the restored database to some point in time between backup job sessions.

- If the backup job fails, Veeam Agent for Microsoft Windows does not truncate transaction logs on the Veeam Agent computer. In this scenario, you can restore computer data from the most recent point in the backup and use database system tools to apply transaction logs and get the database system to the necessary point in time after the restore point.

## Do not Truncate Logs

You can choose not to truncate transaction logs. We recommend this option if you use another backup tool together with Veeam Agent for Microsoft Windows.

For example, you can use Veeam Agent for Microsoft Windows to create a computer image backup and instruct the native Microsoft SQL Server log backup job to back up transaction logs. If you truncate transaction logs with Veeam Agent for Microsoft Windows, the chain of transaction logs will be broken, and the Microsoft SQL Server log backup job will not be able to produce a consistent log backup.

With this option selected, Veeam Agent for Microsoft Windows produces a backup file and does not trigger transaction log truncation. As a result, you have a backup file that contains a computer image, image of a specific data volume or individual folders captured at a specific point in time, and transaction logs. You can use transaction logs to restore the Veeam Agent computer to any point in time between job sessions. To do this, you must recover data from the backup file and use database system tools to apply transaction logs and get the database system to the necessary point in time.

## Back Up Logs Periodically

This option can be used if you back up Microsoft SQL Server or Oracle database system.

You can choose to back up database logs with Veeam Agent for Microsoft Windows. With this option selected, Veeam Agent for Microsoft Windows creates a backup and additionally copies Microsoft SQL Server transaction logs or Oracle archived logs and saves them to the backup location next to the backup files. To learn more, see [Microsoft SQL Server and Oracle Logs Backup](#).

In this scenario, you can use transaction logs to restore the Veeam Agent computer to any point in time between backup job sessions. To do that, you must recover data from the Veeam Agent backup and use Veeam Explorer for Microsoft SQL Server or Veeam Explorer for Oracle to perform transaction log replay and get the database system to a necessary point in time.

## Copy Only Backup

Some organizations prefer to back up Microsoft SQL Server databases and transaction logs with native Microsoft SQL Server tools or 3<sup>rd</sup> party backup tools. To restore database systems in a proper way, database administrators must be sure that they have database backups and a sequence of transaction log backups associated with these backups at hand.

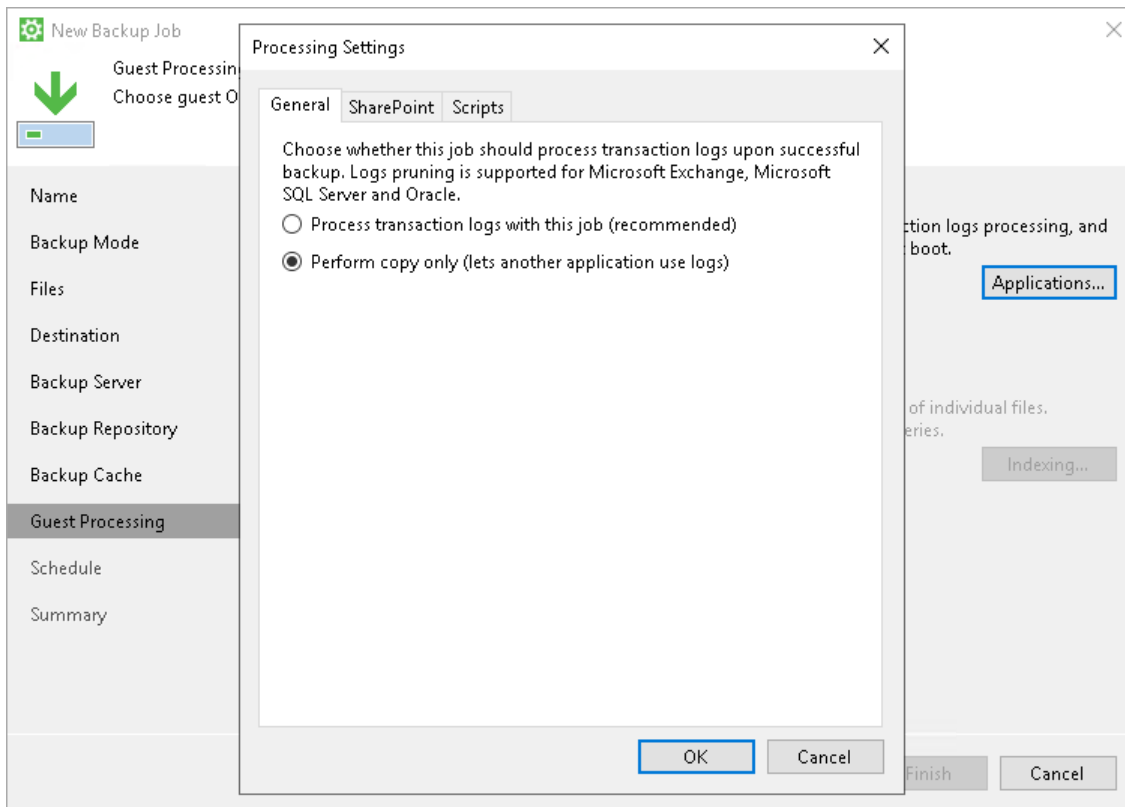
If you use native Microsoft SQL Server tools or 3<sup>rd</sup> party backup tools and also want to back up a machine that runs Microsoft SQL Server with Veeam Agent for Microsoft Windows, you must enable the **Perform copy only** option in the backup job settings.

The **Perform copy only** option indicates that a chain of database backups is created with native Microsoft SQL Server means or by a 3<sup>rd</sup> party tool, and instructs Veeam Agent to preserve this chain (backup history). If the **Perform copy only** option is enabled, Veeam Agent produces the backup of the Copy Backup type that is independent of the existing chain of database backups and does not contain transaction log data. As a result, the backup does not change the log sequence number and transaction log backup time. To learn more about Copy Backup and other VSS backup types, see [Microsoft documentation](#).

## IMPORTANT

Consider the following:

- Veeam Agent for Microsoft Windows does not truncate transaction logs after copy-only backup. For this reason, if you instruct the backup job to perform copy-only backup, you cannot specify transaction log handling settings for this job.
- Veeam Agent for Microsoft Windows supports one transaction log backup job per Veeam Agent computer. If you plan to configure several backup jobs using Veeam Agent for Microsoft Windows, you can enable database log backup settings for one job only.



## Microsoft SQL Server and Oracle Logs Backup

### NOTE

This and subsequent sections describe application-aware processing of Microsoft SQL Server and Oracle database systems in Veeam Agent for Microsoft Windows. You can perform item-level recovery of Microsoft SQL Server and Oracle systems if you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication. For more information, see [Veeam Backup & Replication Documentation](#).

## Microsoft SQL Server Log Backup

You can instruct the Veeam Agent backup job to create volume level or file-level backups and also periodically back up database transaction logs. If Microsoft SQL Server fails, you can restore Microsoft SQL Server from the necessary restore point of the Veeam Agent backup. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you can also use Veeam Explorer for Microsoft SQL Server to apply transaction logs and get databases on the Microsoft SQL Server to the necessary state between backups.



## Requirements for Microsoft SQL Server Transaction Log Backup

- Veeam Agent for Microsoft Windows supports transaction log backups for the following systems:
  - Microsoft SQL Server 2022
  - Microsoft SQL Server 2019
  - Microsoft SQL Server 2017
  - Microsoft SQL Server 2016
  - Microsoft SQL Server 2014 SP2
  - Microsoft SQL Server 2012 SP3
  - Microsoft SQL Server 2008 R2 SP3
  - Microsoft SQL Server 2008 SP4
- The database whose logs you want to back up must use the *Full* or *Bulk-logged* recovery model. In this case, all changes of the Microsoft SQL Server state will be written to transaction logs, and you will be able to replay transaction logs to restore the Microsoft SQL Server. You can use the Microsoft SQL Server Management Studio to switch to one of these models. For more information, see [Microsoft documentation](#).

## Oracle Log Backup

Veeam Agent for Microsoft Windows supports backup of Oracle database archived logs. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you can use Veeam Explorer for Oracle to apply archived logs and get Oracle databases to the necessary state between backups.

Database archived logs are created by the Oracle system. The Oracle database can run in one of the following logging modes:

- ARCHIVELOG turned on – logs are saved and can be used for recovery purposes.
- ARCHIVELOG turned off – no archived logs are saved. We do not recommend this mode, because it does not provide proper disaster recovery.

With ARCHIVELOG turned on, the Oracle system stores database archived logs in a certain location on the machine that runs the database system, as specified by the database administrator. Veeam Agent for Microsoft Windows allows you to set up the following ways of log handling:

- Instruct the backup job to collect log files from the Oracle system and ship them to the backup location where they are stored next to regular backup files created by Veeam Agent for Microsoft Windows.
- Skip log processing – log files remain untouched and are preserved within the Veeam Agent backup.

If you enable application-aware processing for Oracle, during the job session Veeam Agent for Microsoft Windows collects information about the database and processes archived logs according to job settings. Application-specific settings are configured at the **Guest Processing** step of the **New Backup Job** wizard – you can specify how logs should be managed for Oracle databases.

## Requirements for Oracle Archived Log Backup

- Veeam Agent for Microsoft Windows supports archived log backup and restore for Oracle database version 11.2 and later.
- Automatic Storage Management (ASM) is not supported.

- The database must run in the ARCHIVELOG mode.

## Database Log Backup Job

To back up database logs (Microsoft SQL Server transaction logs and Oracle archived logs), you must specify advanced settings for transaction log backup in the Veeam Agent backup job settings. The resulting job will comprise two jobs:

- Parent backup job — the backup job that creates a volume-level or file-level backup. The backup job becomes the parent job after you enable database log backup options at the **Guest Processing** step of the **New Backup Job** wizard.
- Child job — a transaction log backup job. Veeam Agent for Microsoft Windows automatically creates the child job if transaction log backup is enabled for the backup job. Session data of the transaction log backup job is stored in the Veeam Agent for Microsoft Windows database and displayed in the Veeam Agent control panel. To learn more, see [Transaction Log Backup Statistics](#).

The parent job runs in a regular manner — it starts by schedule or is started manually by the user. The transaction log backup job is triggered by the parent backup job. This sequence ensures that the restore point is present when it comes to transaction log replay.

### NOTE

Veeam Agent for Microsoft Windows supports one transaction log backup job per Veeam Agent computer. If you plan to configure several backup jobs using Veeam Agent for Microsoft Windows, you can enable database log backup settings for one job only.

## Sessions of Transaction Log Backup Jobs

The transaction log backup job runs permanently in the background, shipping transaction logs to the backup location at a specific time interval (by default, every 15 minutes). A sequence of time intervals between sessions of the parent backup job makes up a session of the transaction log backup job.

The transaction log backup session starts and stops in the following way:

- The initial session starts when the parent backup job schedule is enabled. After that, the session starts with every new session of the parent backup job.
- The session ends before the next session of the parent backup job or when this parent backup job is disabled.

## How Microsoft SQL Server Log Backup Works

The transaction log backup for Microsoft SQL Server is performed in the following way:

1. Veeam Agent for Microsoft Windows launches the parent backup job by schedule.
2. The parent backup job creates a volume-level or file-level backup and stores it to the backup location.
3. A new session of the transaction log backup job starts. Veeam Agent for Microsoft Windows copies transaction log files from the log archive destination (set by the Microsoft SQL Server administrator) to a temporary folder on the Veeam Agent computer file system.
4. Veeam Agent for Microsoft Windows detects what databases currently exist on the Microsoft SQL Server and maps this data with the information kept in the Veeam Agent for Microsoft Windows database. This periodic mapping reveals the databases for which Veeam Agent for Microsoft Windows must process transaction logs during this time interval.

5. Veeam Agent for Microsoft Windows transports transaction log backup copies from the temporary folder to the backup location and saves them as VLB files. As soon as copies of transaction log backups are saved to the backup location, transaction log backups in the temporary folder on the Veeam Agent computer are removed.

#### TIP

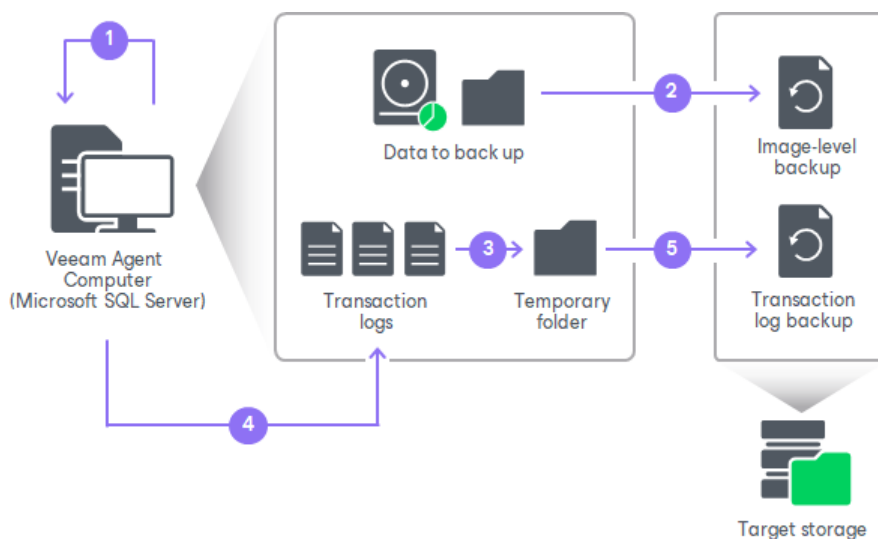
Default location of the temporary folder where Veeam Agent copies transaction log files to is `C:\ProgramData\Veeam\Endpoint\SqlLogBackup`. If it is necessary to change the location, you can use a registry value. To learn more, see [this Veeam KB article](#).

The session of the transaction log backup job remains working until the next start of the parent backup job. When a new session of the parent job starts, the transaction log backup job stops the current session and then starts a new session, performing steps 1-5.

Transaction logs that for some reason were not processed during the log backup interval remain in the temporary folder and are processed during the next log backup interval. To detect these remaining logs, Veeam Agent for Microsoft Windows enumerates log files in the temporary folder.

#### NOTE

If a new session of the transaction log backup starts and the parent backup job has not created a new restore point yet, the transaction log backup job will remain in the idle state, waiting for a new restore point to be created.



## How Oracle Archived Log Backup Works

The archived log backup for Oracle is performed in the following way:

1. Veeam Agent for Microsoft Windows launches the parent backup job by schedule.
2. The parent backup job creates a volume-level or file-level backup and stores it to the backup location.
3. A new session of the archived log backup starts. Veeam Agent for Microsoft Windows scans the Oracle system and collects information about databases whose logs must be processed, including:
  - List of all databases
  - Database state — a database is on or off, in which logging mode it runs

- Paths to all database files (configuration logs and so on) and other data required for backup

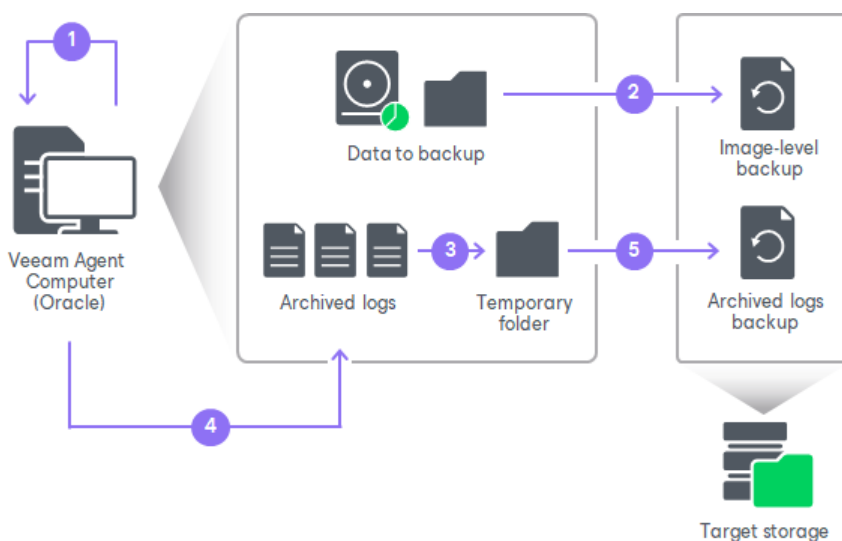
Veeam Agent for Microsoft Windows copies archived log files from the log archive destination (set by the Oracle administrator) to a temporary folder on the Veeam Agent computer.

4. Veeam Agent for Microsoft Windows maps information about the Oracle system collected at the step 3 with information kept in the Veeam Agent for Microsoft Windows database. This periodic mapping helps reveal databases for which Veeam Agent for Microsoft Windows must ship archived logs to the backup location during this time interval.
5. Archived log backup files are transferred from the temporary folder on the Veeam Agent computer to the backup location.
6. If you configure the backup job to delete archived log backup files, Veeam Agent deletes archived log backup files from the backup location according to selected [Oracle archived log settings](#).

Archived logs that for some reason were not processed during the log backup interval remain in the temporary folder and are processed during the next log backup interval. To detect these remaining logs, Veeam Agent for Microsoft Windows enumerates log files in the temporary folder.

## NOTE

If a new session of the archived log backup starts and the parent backup job has not created a new restore point yet, the transaction log backup job will remain in the idle state, waiting for a new restore point to be created.



## Retention for Database Log Backups

Transaction log backups are stored in files of the proprietary Veeam format – VLB. Veeam Agent for Microsoft Windows keeps transaction log backups together with the chain of backup files on the target location.


Veeam Agent for Microsoft Windows removes transaction log backups by retention. You can choose one of the following retention methods:

- [Retain logs according to the image-level backup](#)
- [Retain logs for the specified number of days](#)

## Retain Logs with Image-Level Backup

By default, Veeam Agent for Microsoft Windows retains transaction log backups together with the corresponding backup file. When Veeam Agent for Microsoft Windows removes a restore point from the backup chain, it also removes a chain of transaction logs relating to this restore point.

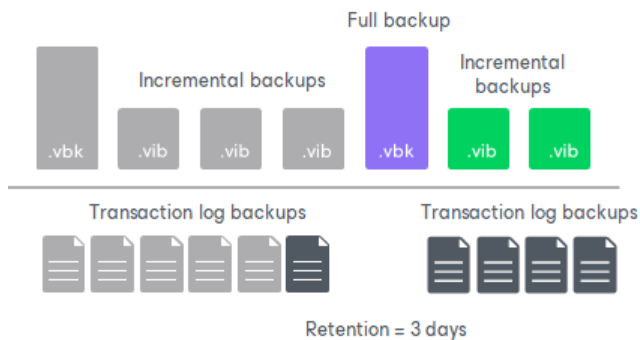
This method allows you to have both the file-level or volume-level backup and necessary transaction log backups at hand. If you need to recover a database to some state, you can restore a machine running Microsoft SQL Server or Oracle from the necessary restore point and perform transaction log replay to bring the database to the desired state.



## Retain Logs for a Number of Days

You can instruct Veeam Agent for Microsoft Windows to keep transaction logs only for a specific period of time. This retention setting can be used, for example, if you want to save on storage space and plan to retain transaction log backups for the last few days. In this case, you will be able to perform transaction log replay only to one of the most recent states.

If you select this retention method, you must make sure that retention policies for the Veeam Agent backup and transaction log backup are consistent. The restore point of the volume-level or file-level backup must always be preserved. If a backup of the database itself is missing, you will not be able to perform transaction log replay.



## File System Indexing

You can instruct Veeam Agent for Microsoft Windows to create an index of files and folders on the Veeam Agent computer OS during backup. If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you will be able to search for individual files inside Veeam Agent backups and perform 1-click restore in Veeam Backup Enterprise Manager.

## NOTE

File system indexing is optional. If you do not enable this option in the backup job settings, you will still be able to perform 1-click restore from the backup created with such backup job. For more information, see the [Preparing for File Browsing and Searching](#) section in the Veeam Backup Enterprise Manager User Guide.

The screenshot shows the 'New Backup Job' wizard in Veeam Backup Enterprise Manager, specifically the 'Guest Processing' step. The window has a title bar with a green gear icon and the text 'New Backup Job'. Below the title bar, there is a green arrow icon pointing down and the text 'Guest Processing' and 'Choose guest OS processing options.' A progress bar is shown below the text. On the left side, there is a list of steps: Name, Backup Mode, Files, Destination, Backup Server, Backup Repository, Backup Cache, Guest Processing (highlighted), Schedule, and Summary. The main area contains two options: 'Enable application-aware processing' (unchecked) and 'Enable file system indexing' (checked). The 'Enable application-aware processing' option has a description: 'Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.' and a button 'Applications...'. The 'Enable file system indexing' option has a description: 'Creates catalog of files to enable browsing, searching and 1-click restores of individual files. Indexing is optional, and is not required to perform instant file level recoveries.' and a button 'Indexing...'. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

**New Backup Job**

Guest Processing  
Choose guest OS processing options.

Name

Backup Mode

Files

Destination

Backup Server

Backup Repository

Backup Cache

**Guest Processing**

Schedule

Summary

☐ **Enable application-aware processing**  
Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.  
Customize application handling options for individual applications Applications...

☒ **Enable file system indexing**  
Creates catalog of files to enable browsing, searching and 1-click restores of individual files. Indexing is optional, and is not required to perform instant file level recoveries.  
Customize advanced file system indexing options Indexing...

< Previous Next > Finish Cancel

## How File Indexing Works

When you run a backup job with the file indexing option enabled, Veeam Agent for Microsoft Windows performs the following operations:

1. When the backup job starts, Veeam Agent for Microsoft Windows starts indexing the file system. The indexing procedure is carried out in parallel with the backup procedure. If indexing takes long, Veeam Agent for Microsoft Windows will not wait for the indexing procedure to complete. It will start copying data to the target location and continue file indexing.
2. When file indexing is complete, Veeam Agent for Microsoft Windows collects indexing data, writes it to an archive file and stores this archive file to the backup file along with the backed-up data.
3. If the backup job is set up to create backups in a Veeam backup repository, when the job completes, Veeam Guest Catalog Service running on the backup server also saves indexing data in the Veeam Catalog folder on the backup server.

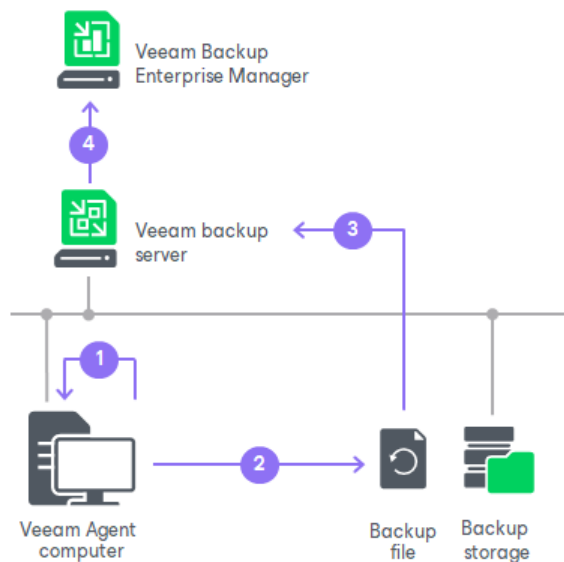
To learn more about the Veeam Guest Catalog Service, see the [Veeam Backup Catalog](#) section in the Veeam Backup & Replication User Guide.

4. During the next catalog replication session, the global Veeam Guest Catalog Service replicates data from the backup server to the Veeam Catalog folder on the Veeam Backup Enterprise Manager server.

## NOTE

Consider the following:

- If the backup job is set up to create backups in a Veeam Cloud Connect repository, Veeam Backup & Replication running on the SP backup server does not save indexing data in the Veeam Catalog folder.
- Due to the peculiarities of file indexing algorithms, file indexing takes more time if you backup volumes with a file system different from NTFS or volumes with data deduplication enabled.



# Data Encryption

Data security is an important part of the backup strategy. You must protect your information from unauthorized access, especially if you back up sensitive data to remote locations. To keep your data safe, you can use data encryption.

Data encryption transforms data to an unreadable, scrambled format with the help of a cryptographic algorithm and a secret key. If encrypted data is intercepted, it cannot be unlocked and read by the eavesdropper. Only intended recipients who know the secret key can reverse encrypted information back to a readable format.

In Veeam Agent, encryption works at the backup job level. Veeam Agent uses the block cipher encryption algorithm and stores data in the encrypted format to a backup file.

Encryption is performed on the trusted side depending on the backup target:

- Encryption is performed on the source side for all backup targets except the Veeam backup repository.
- Encryption is performed on the target side if you store backups in the Veeam backup repository.

Decryption is performed on the same side as encryption.

To create encrypted backups, you must enable the encryption option and specify a password that will be used for data encryption. To learn more, see [Storage Settings](#).

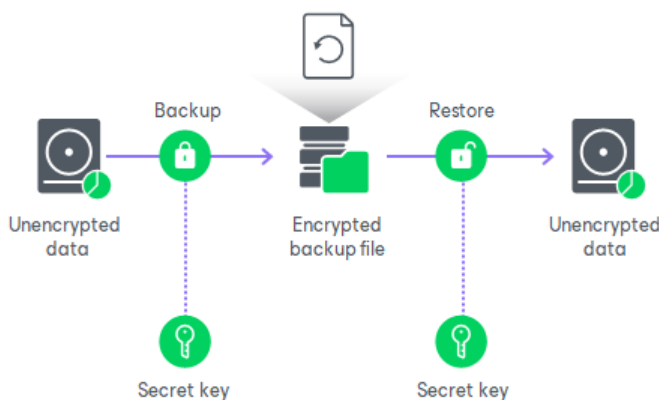
## NOTE

You cannot enable encryption options in the properties of the Veeam Agent backup job if you have chosen to create Veeam Agent backups in a Veeam backup repository. For such jobs, encryption options are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the [Data Encryption](#) section in the Veeam Backup & Replication User Guide.

## Encryption Algorithms

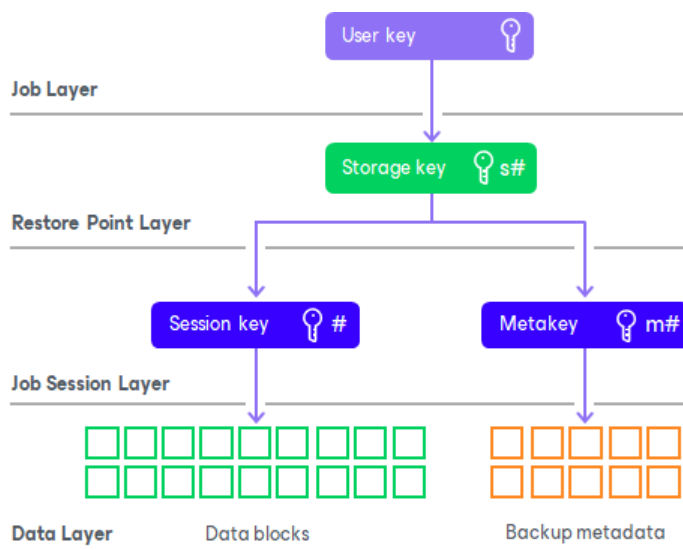
To encrypt data in backups and files, Veeam Agent employs a symmetric key encryption algorithm.

The symmetric, or single-key encryption algorithm, uses a single, common secret key to encrypt and decrypt data on the trusted side. To restore encrypted data, you must have the same secret key. Users who do not have the secret key cannot decrypt data and get access to it.





Veeam Agent relies on a hierarchical encryption scheme. Each layer in the hierarchy encrypts the layer below with a key of specific type.

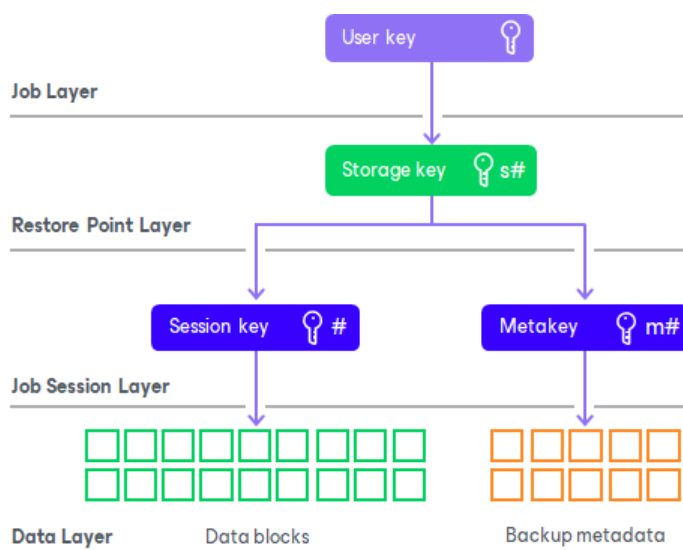


## Encryption Keys

An encryption key is a string of random characters that is used to bring data to a scrambled format and back to unscrambled. Encryption keys encode and decode initial data blocks or underlying keys in the key hierarchy.

Veeam Agent uses 4 types of keys:

- 3 service keys generated by Veeam Agent:
  - [Session Key](#)
  - [Metakey](#)
  - [Storage key](#)
- 1 key generated based on a user password: a [user key](#).




## Session Keys and Metakeys

The session key is the lowest layer in the encryption key hierarchy. When Veeam Agent encrypts data, it first encodes every data block in a file with a session key. For session keys, Veeam Agent uses the AES algorithm with a 256-bit key length in the CBC-mode.


Veeam Agent generates a new session key for every backup job session. For example, if you have created an encrypted backup job and run 3 job sessions, Veeam Agent will produce 3 backup files that will be encrypted with 3 different session keys:

- Full backup file encrypted with session key 1
- Incremental backup file encrypted with session key 2
- Incremental backup file encrypted with session key 3



The session key is used to encrypt only data blocks in backup files. To encrypt backup metadata, Veeam Agent applies a separate key – metakey. Use of a metakey for metadata raises the security level of encrypted backups.

For every job session, Veeam Agent generates a new metakey. For example, if you have run 3 job sessions, Veeam Agent will encrypt metadata with 3 metakeys.




In the encryption process, session keys and metakeys are encrypted with keys of a higher layer – storage keys. Cryptograms of session keys and metakeys are stored in the resulting file next to encrypted data blocks. Metakeys are additionally kept in the Veeam Agent database.

## Storage Keys

Backup files in the backup chain often need to be transformed, for example, when the earliest incremental backup file in the chain becomes obsolete and its data should be included into the full backup file. When Veeam Agent transforms a full backup file, it writes data blocks from several restore points to the full backup file. As a result, the full backup file contains data blocks that are encrypted in different job sessions with different session keys.

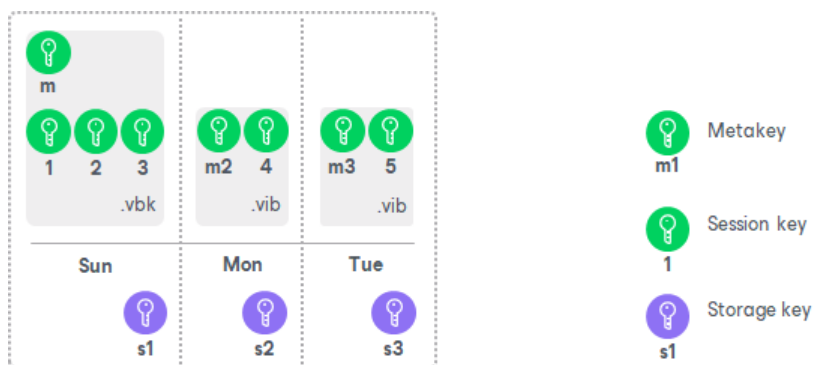
To restore data from such “composed” backup file, Veeam Agent would require a bunch of session keys. For example, if the backup chain contains restore points for 2 months, Veeam Agent would have to keep session keys for a 2-month period.



In such situation, storing and handling session keys would be resource consuming and complicated. To facilitate the encryption process, Veeam Agent uses another type of service key – a storage key.

For storage keys, Veeam Agent uses the AES algorithm with a 256-bit key length in the CBC-mode. A storage key is directly associated with one restore point in the backup chain. The storage key is used to encrypt the following keys in the encryption hierarchy:

- All session keys for all data blocks in one restore point
- Metafile encrypting backup metadata



During the restore process, Veeam Agent uses one storage key to decrypt all session keys for one restore point, no matter how many session keys were used to encrypt data blocks in this restore point. As a result, Veeam Agent does not need to keep the session keys history in the Veeam Agent database. Instead, it requires only one storage key to restore data from one file.


In the encryption process, storage keys are encrypted with a key of a higher layer – a user key. Cryptograms of storage keys are stored in the resulting file next to encrypted data blocks, and cryptograms of session keys and metafiles.

Storage keys are also kept in the Veeam Agent database. To maintain a set of valid storage keys in the database, Veeam Agent uses retention policy settings specified for the job. When some restore point is removed from the backup chain by retention, the storage key corresponding to this restore point is also removed from the Veeam Agent database.

## User Keys

When you enable encryption for a job, you must define a password to protect data processed by this job, and define a hint for the password. The password and the hint are saved in the job settings. Based on this password, Veeam Agent generates a user key.

The user key protects data at the job level. In the encryption hierarchy, the user key encrypts storage keys for all restore points in the backup chain.



Veeam Agent saves a hint for the password to its database and to the backup metadata file (VBM). When you decrypt a file, Veeam Agent displays a hint for the password that you must provide. After you enter a password, Veeam Agent derives a user key from the password and uses it to unlock the storage key for the encrypted file.

According to the security best practices, you should change passwords for encrypted jobs regularly. When you change a password for the job, Veeam Agent creates a new user key and uses it to encrypt new restore points in the backup chain. If you lose a password that was specified for encryption, you can change the password in the encryption settings. You can use the new password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.


## How Data Encryption Works

Data encryption is performed as part of the backup process. Encryption works at the source or target side depending on the backup target. As a result, encryption keys are not passed to the untrusted side, which helps to avoid data interception.

In Veeam Agent, the encryption process includes the following steps:

1. When you create a backup job, you enable the encryption option for the job and enter a password to protect data at the job level.
2. Veeam Agent generates a user key based on the entered password.
3. When you start an encrypted job, Veeam Agent creates a storage key and stores this key in its database.
4. Veeam Agent creates a session key and a metakey. The metakey is stored in the Veeam Agent database.
5. Veeam Agent processes job data in the following way:
  - a. The session key encrypts data blocks in the backup file. The metakey encrypts backup metadata.
  - b. The storage key encrypts the session key and the metakey.
  - c. The user key encrypts the storage key.

6. Encrypted data blocks are stored to the target location. The cryptograms of the user key, storage key, session key and metakey are stored in the resulting file next to encrypted data blocks.



## How Data Decryption Works

When you restore data from an encrypted backup file, Veeam Agent performs data decryption automatically in the background or requires you to provide a password.

- If encryption keys required to unlock the backup file are available in the Veeam Agent database, you do not need to enter the password. Veeam Agent uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore does not differ from that from an unencrypted one.

Automatic data decryption can be performed in one of the following situations:

- You encrypt and decrypt the backup file on the same Veeam Agent computer using the same Veeam Agent database.
  - You have included encryption keys into the Veeam Recovery Media and perform bare metal recovery after booting from this Veeam Recovery Media. To learn more, see [Specify Recovery Media Options](#).
- If encryption keys are not available in the Veeam Agent database, you need to provide a password to unlock the encrypted file.

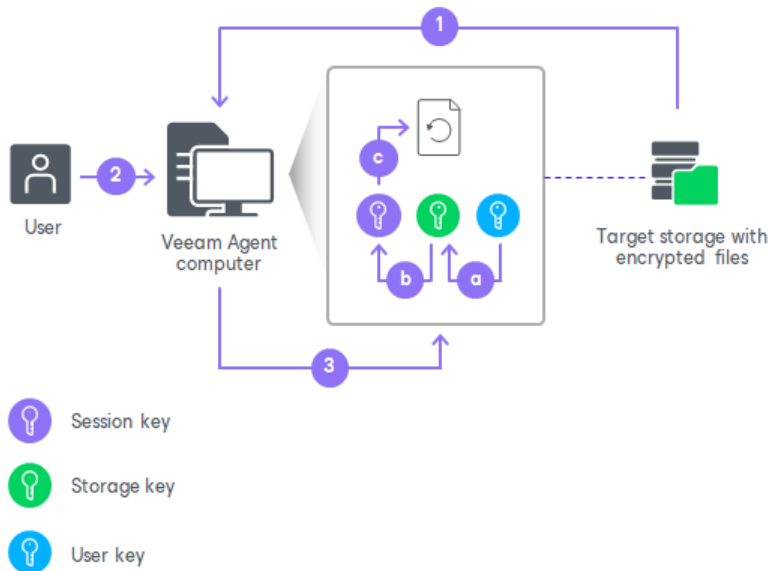
Data decryption is performed on the source or target side depending on the backup target. As a result, encryption keys are not passed to the untrusted side, which helps avoid data interception.

In Veeam Agent, the decryption process includes the following steps. Keep in mind that steps 1 and 2 are required only if you decrypt the file on the Veeam Agent computer other than the computer where the file was encrypted.

1. You select the backup from which you want to restore data. Veeam Agent notifies you that one or more files in the backup chain are encrypted and requires a password.
2. You specify a password for the imported file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

3. Veeam Agent reads the entered password and generates the user key based on this password. With the user key available, Veeam Agent performs decryption in the following way:
  - a. Veeam Agent applies the user key to decrypt the storage key.
  - b. The storage key, in its turn, unlocks underlying session keys and a metakey.
  - c. Session keys decrypt data blocks in the encrypted file.

After the encrypted file is unlocked, you can work with it as usual.

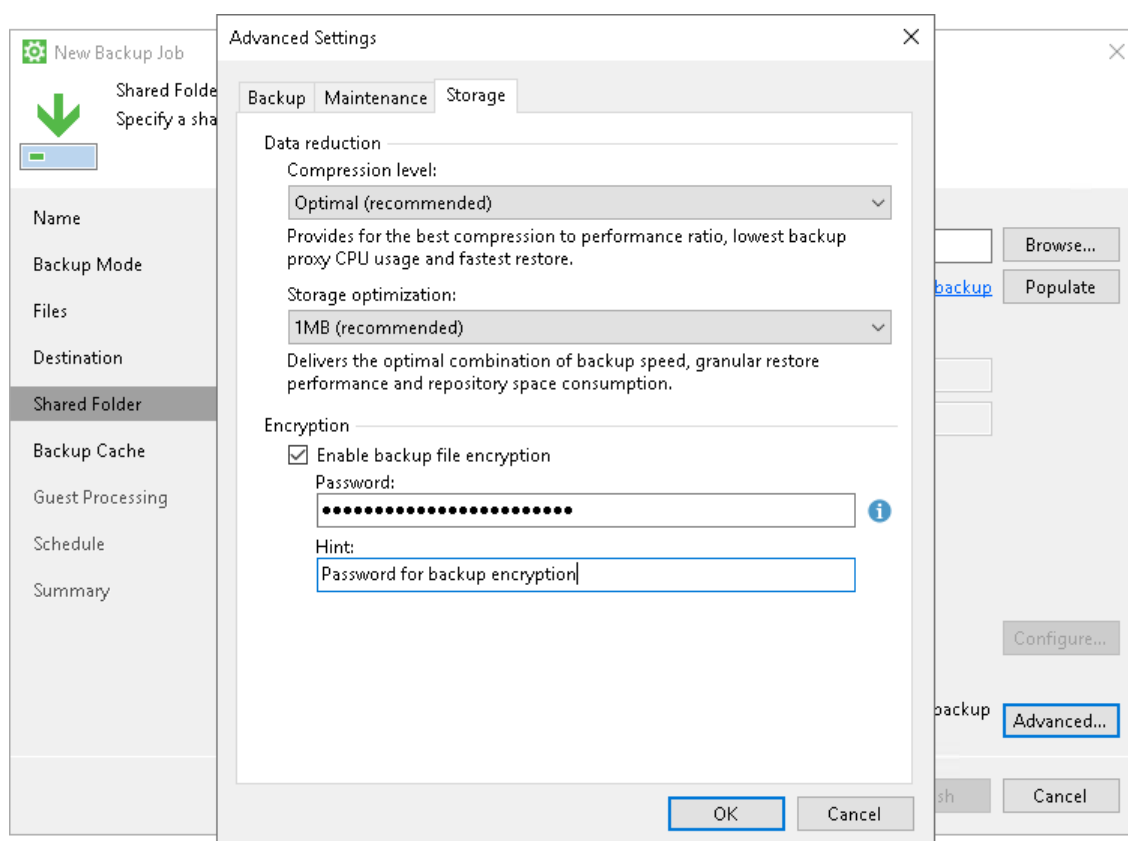


## Backup Job Encryption

Encryption for the backup job is configured in the advanced job settings. You should enable the encryption option and specify a password to protect data in backup files produced by the backup job.

## NOTE

You cannot specify encryption options for the backup job if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more, see the [Data Encryption](#) section of the Veeam Backup & Replication User Guide.



The backup job processing with encryption enabled includes the following steps:

1. You enable encryption for a backup job and specify a password.
2. Veeam Agent generates the necessary keys to protect backup data.
3. Veeam Agent encrypts data blocks and transfers them to the target location already encrypted.
4. On the target storage, encrypted data blocks are stored in a resulting backup file.



Restore of an encrypted backup file includes the following steps:

1. You select an encrypted backup file and define a password to decrypt the backup file. If the password has changed once or several times, you need to specify the latest password that was used to encrypt files in the backup chain.
2. Veeam Agent uses the provided password to generate user key and unlock the subsequent keys for backup file decryption.

3. Veeam Agent retrieves data blocks from the backup file, sends them to the target volume and decrypts them on the target volume.



## Resuming Encrypted Backup Chain

In some situations, encryption keys may be unavailable in the Veeam Agent for Microsoft Windows database, and Veeam Agent cannot create a new encrypted restore point in the backup chain. For example, this may happen after you change the password for encryption and then recover the entire Veeam Agent computer to a restore point that was created before you have changed the password. In this case, information about backup in the Veeam Agent database will become outdated and will not match backup metadata residing on the target location. To continue the existing encrypted backup chain, you need to provide the latest password in the Veeam Agent control panel.

When the backup job is started (either manually or upon the defined schedule), Veeam Agent detects the latest encrypted backup created by this job in the target location and displays a window in the Veeam Agent control panel offering to enter a password and continue the backup chain. You can choose to perform one of the following operations:

- To continue the existing encrypted backup chain, you can enter the password specified for encryption and click **OK**. If the password has changed more than once, you need to specify the latest password.

After you provide the correct password, Veeam Agent will use this password to decrypt backup metadata on the target location and update information about backup in its database. After that, Veeam Agent will create the new incremental backup file in the existing encrypted backup chain. To encrypt this backup file and subsequent backup files, Veeam Agent will use the password that is kept in its database at the time when continue the backup chain.

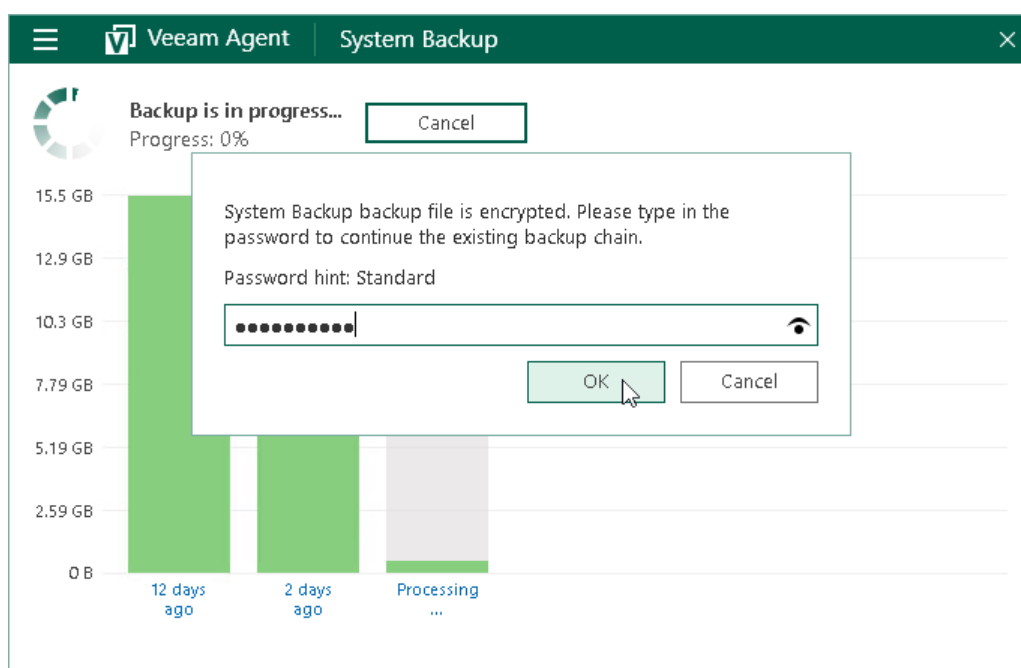
You will be able to use this password to restore data from any restore point in the backup chain, including restore points that were encrypted with an older password and restore points that were created before you have enabled the encryption option for the job.

- You can click **Cancel** to close the notification window and cancel the job. In this case, the next time the backup job is started, Veeam Agent will again prompt to enter the password.



## NOTE

Veeam Agent for Microsoft Windows displays the notification window for 1 hour. If you do not choose any option, after this time period expires, Veeam Agent will automatically close the window and cancel the job. If email notification settings are enabled for the backup job, Veeam Agent will also send an email report informing that the job was canceled.



## Encryption Best Practices

To guarantee the flawless process of data encryption and decryption, consider the following advice.

### Password

1. Use strong passwords that are hard to crack or guess. Consider the following recommendations:
  - a. The password must be at least 8 characters long.
  - b. The password must contain uppercase and lowercase characters.
  - c. The password must be a mixture of alphabetic, numeric and punctuation characters.
  - d. The password must significantly differ from the password you used previously.
  - e. The password must not contain any real information related to you, for example, date of birth, your pet's name, your logon name and so on.
2. Provide a meaningful hint for the password that will help you recall the password. The hint for the password must significantly differ from the password itself. The hint for the password is displayed when you select an encrypted backup server and attempt to unlock it.
3. Change passwords for encrypted jobs regularly. Use of different passwords helps increase the encryption security level.

## Encryption for Existing Job

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

Encryption is not retroactive. If you enable encryption for an existing backup job, Veeam Agent does not encrypt the previous backup chain created with this job. However, Veeam Agent encrypts backup metadata. As a result, you need to enter the password to restore data from unencrypted backup files in the backup chain as well as from encrypted backup files in this chain.

# Backup Cache

A remote storage specified as a target location for backup files may be unavailable at the time when the backup job must start. In this case, Veeam Agent cannot create a regular restore point upon the defined schedule. As a result, the backup chain on the remote storage will not contain a sequence of restore points that precisely complies with the backup schedule.

To overcome this limitation, Veeam Agent offers the concept of the backup cache. The backup cache is a temporary local storage in which Veeam Agent creates backup files in case the backed-up data cannot be transferred to a remote location. When the target location becomes available, Veeam Agent uploads backup files from the backup cache to the remote storage, adding regular restore points to the backup chain.

The backup cache lets you perform scheduled backup in due time ensuring that the resulting backup chain will contain "snapshots" of your data at desired points in time. This may be helpful, for example, for laptop users who go on business trips with no or limited access to the corporate network in which the backup location resides.

Technically, the backup cache is a local folder on the computer on which Veeam Agent is installed. A user can define a folder for the backup cache and the size of the backup cache in the backup job settings.

The backup cache is available if the following types of storage are chosen as a target location:

- Network shared folder
- Object storage
- Backup repository managed by a Veeam backup server
- Cloud repository managed by a Veeam Cloud Connect service provider

## How Backup Cache Works

When you create a backup job targeted at a remote storage, you can select to use the backup cache in its properties. The procedure of data backup with the backup cache enabled is performed in the following way:

1. The user enables the backup cache in the properties of the backup job targeted at a remote location.
2. During a regular backup job session, Veeam Agent for Microsoft Windows creates in the backup cache a map of data blocks on the remote location. Information about data blocks on the remote location is saved in a file with digests and stored in the folder

`C:\ProgramData\Veeam\EndpointData\CacheDigests.`

Veeam Agent for Microsoft Windows will use the created digests to create incremental backup files in the backup cache when the remote storage itself is unavailable at the time of scheduled backup.

3. If the target location is unavailable at the time when the scheduled backup job must start, Veeam Agent for Microsoft Windows creates the new restore point in the backup cache.

The target location is considered as unavailable in the following conditions:

- [For network shared folder] Veeam Agent for Microsoft Windows Service that runs on the protected computer cannot connect to the network shared folder. In this case, Veeam Agent for Microsoft Windows will immediately start creating the new restore point in the backup cache.
- [For object storage] Veeam Agent for Microsoft Windows Service that runs on the protected computer cannot connect to the object storage.

- [For Veeam backup repository] Veeam Agent for Microsoft Windows Service cannot connect to the Veeam Backup Service that runs on the backup server to which the backup repository specified as a target location for Veeam Agent backups is connected.
- [For Veeam Cloud Connect repository] Veeam Agent for Microsoft Windows Service cannot connect to one of the following services on the Veeam Cloud Connect provider side:
  - *Veeam Backup Service* that runs on the backup server used for managing the Veeam Cloud Connect infrastructure.
  - *Veeam Cloud Connect Service* that runs on the backup server used for managing the Veeam Cloud Connect infrastructure.
  - *Veeam Cloud Gateway Service* that runs on a cloud gateway deployed in the Veeam Cloud Connect infrastructure.

In case the connection to the target location is lost when the backup job is already running, Veeam Agent for Microsoft Windows performs backup based on the following rules:

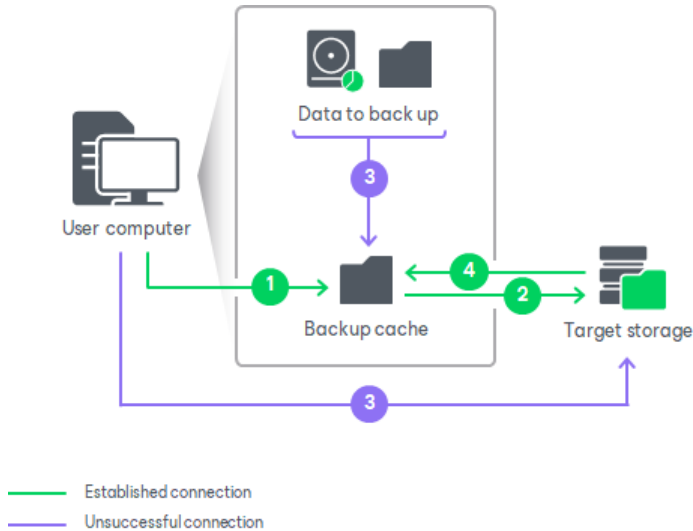
- [For network shared folder] Veeam Agent for Microsoft Windows immediately switches to the backup cache and writes to the backup cache all data that must be backed up within the current backup job session. If some data has already been transferred to the target location, Veeam Agent for Microsoft Windows starts the data transfer operation from the beginning and transfers all backed-up data to the backup cache.
- [For object storage, Veeam backup repository and Veeam Cloud Connect repository] Veeam Agent for Microsoft Windows tries to reconnect to the target location. The reconnection process may last 30 minutes or more. After the reconnection period expires, Veeam Agent for Microsoft Windows switches to the backup cache and writes to the backup cache only remaining data that has not been transferred to the target location yet. Data that has been already transferred to the target location remains in the backup repository.

The process of data backup to the backup cache practically does not differ from the regular one. The difference is that the resulting backup file is saved to the local folder instead of the remote storage.

If the target location becomes available after Veeam Agent for Microsoft Windows has started creating a restore point in the backup cache, Veeam Agent for Microsoft Windows will not switch back to the target location. Instead, Veeam Agent for Microsoft Windows will create a restore point in the backup cache and then upload this restore point to the target location.

4. After the remote storage becomes available, Veeam Agent for Microsoft Windows uploads backup files that were created in the backup cache to the target location. If more than one restore point was created in the backup cache, these restore points are uploaded to the target location sequentially, one by one.

Until all restore points are uploaded from the backup cache to the target location, Veeam Agent for Microsoft Windows will continue to create new restore points in the backup cache even if the target location is available at the time when the backup job is running.



## Limitations for Backup Cache

The backup cache has the following limitations:

- The backup cache is available only in Workstation and Server editions of Veeam Agent for Microsoft Windows.
- You cannot use the backup cache for a file-level backup job.
- You cannot restore data from backup files that reside in the backup cache. Until restore points are uploaded from the backup cache to the target location, these restore points are not considered as a fully valid part of the backup chain.
- Veeam Agent for Microsoft Windows does not support creating full backups (including active full backups and synthetic full backups) in the backup cache except for the very first full backup file in the backup chain.
- Veeam Agent for Microsoft Windows does not support creating encrypted backups in the backup cache. If encryption options are specified for the backup job, Veeam Agent will create unencrypted backup files in the backup cache. When the target location becomes available, Veeam Agent will encrypt data prior to uploading it to the remote storage.
- Veeam Agent for Microsoft Windows does not support creating transaction log backups in the backup cache. You cannot enable transaction log backup and the backup cache for the backup job simultaneously.
- Restore points created in the backup cache cannot be uploaded to the target location after you perform the following operations:
  - Change target location for backup files in the properties of the backup job
  - Change a folder defined for the backup cache
  - Move or delete backup files on the target location

- [For Veeam backup repository] Enable or disable data encryption settings in Veeam Backup & Replication

If the backup cache contains one or more restore points at the time when you perform one of these operations, you need to delete restore points from the backup cache.

- You cannot delete restore points from the backup cache while Veeam Agent for Microsoft Windows is performing the following operations:
  - Creating a new restore point in the backup cache.
  - Uploading a restore point from the backup cache to the target location.

# Backup to Rotated Drives

You can use rotated drives as a target location for backups. This scenario can be helpful if you want to store backups on several external hard drives (for example, USB or FireWire) and plan to swap these drives between different locations regularly.

Backup on rotated drives is performed in the following way:

1. Veeam Agent for Microsoft Windows creates a backup chain on an external drive that you use as a backup target. The backup chain consists of the first full backup and a set of subsequent incremental backups.
2. When you swap drives and attach a new external drive, Veeam Agent creates a separate backup chain on the new drive.
3. After you swap drives again, Veeam Agent detects if there is a backup chain on the currently attached drive. If the backup chain exists, Veeam Agent continues the existing chain: it creates a new incremental backup file and adds it to the existing backup files.

To use rotated drives for backup, you must perform the following actions:

1. Attach one of external drives from the set to your computer.
2. Configure the backup job to store backups on the currently connected external drive. To do this:
  - a. At the **Local Storage** step of the wizard, select the connected drive.
  - b. From the **Local drives** list, select the necessary volume on the connected drive and specify a folder where backups must be stored.
  - c. Save the job settings.

**New Backup Job**

Local Storage  
Choose a locally attached drive to back up to.

**Local Storage**

Storage device	Free space	Total space
Local Disk (C:)	67.2 GB	129.4 GB
Local Disk (E:)	89.4 GB	90.0 GB
Local Disk (F:)	89.9 GB	90.0 GB

Folder: F:\VeeamBackup\ [Browse...](#)

[Map backup](#)

Retention policy: 7 days

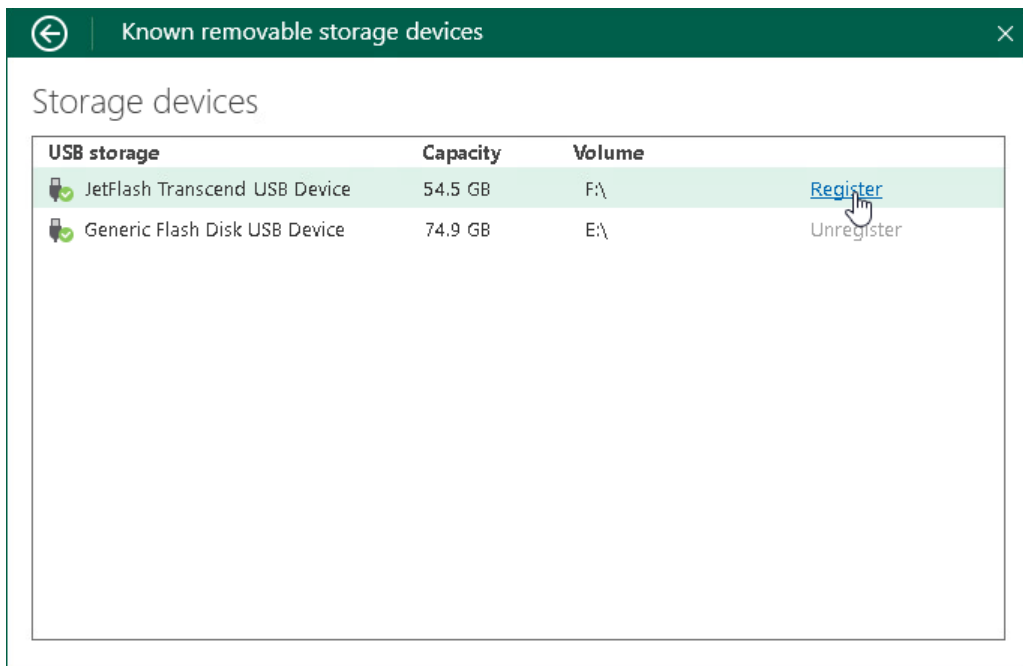
☐ Keep certain full backups longer for archival purposes  
Periodic fulls are not enabled in Advanced settings

Click Advanced to enable periodic full backups, configure encryption and other backup file settings

[Configure...](#) [Advanced...](#)

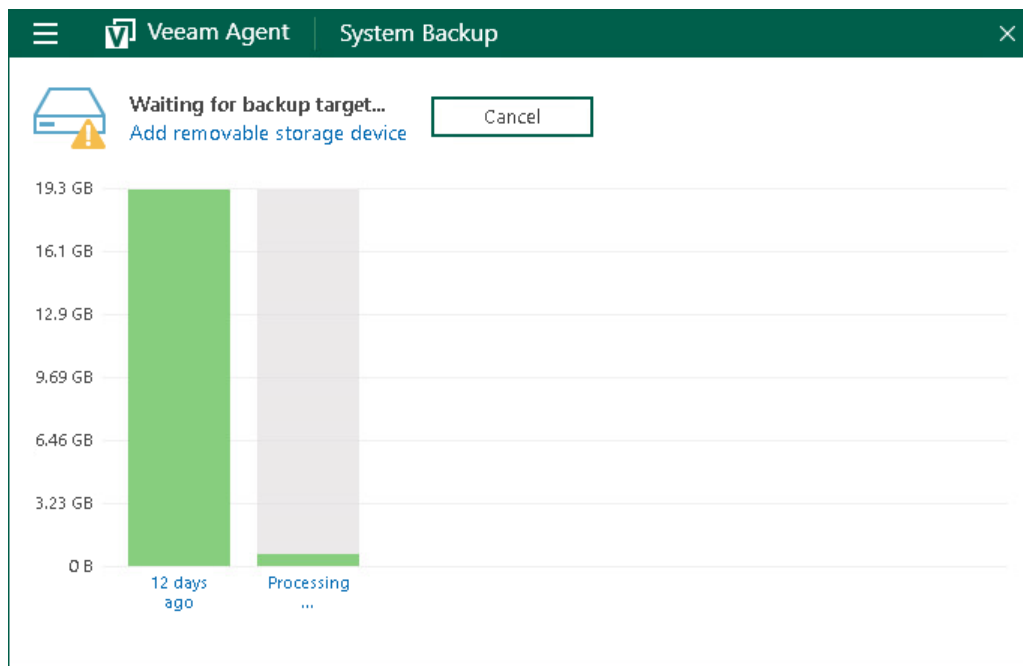
< Previous **Next >** Finish Cancel

3. When you need to swap files, disconnect the drive that was used previously and attach a new drive to your computer.
4. Register a newly connected drive as a known removable storage in Veeam Agent for Microsoft Windows. To do this:
  - a. Double-click the Veeam Agent icon in the system tray to open the control panel.
  - b. From the main menu, select **Settings**.
  - c. Click the **Manage registered storage devices** link.
  - d. Click **Register** next to the newly connected drive.





If you do not register the newly connected drive before the backup job starts, Veeam Agent will be unable to detect the backup target and launch the backup job. Veeam Agent will display a warning in the system tray and in the control panel. To register a new device, click the **Add removable storage device** link in the **Status** view of the control panel and register the newly connected drive as described above. To learn more, see [Managing Rotated Drives](#).



5. After you register the newly connected drive, you can start a new backup session manually or wait Veeam Agent to start a new session.

#### NOTE

Keep in mind that you cannot target a backup job at a certain set of rotated drives. If the backup job to rotated drive starts, and the drive that was used previously is not connected, Veeam Agent creates a backup file on any registered drive that is currently attached to the Veeam Agent computer.

# Backup of External Drives

You can use Veeam Agent for Microsoft Windows to back up data that resides on external USB/eSATA or FireWire drives connected to the Veeam Agent computer.

## IMPORTANT

Veeam Agent supports backup of external drives that support Microsoft VSS: HDD, SSD, and so on. USB flash drives (USB sticks) are not supported, because it is impossible to create a VSS snapshot on a device of this type. To learn how to check if your drive is supported, see [this Veeam KB article](#).

You can include an external drive in any type of backup: entire computer backup, volume-level backup or file-level backup.

- For entire computer backup, you can instruct Veeam Agent for Microsoft Windows to back up data of all external drives that are currently connected to the Veeam Agent computer. To do this, you must enable the **Include external USB drives** option at the **Backup Mode** step of the **New Backup Job** wizard.

An external drive whose data you want to back up must be connected to the Veeam Agent computer at the time when the backup job must start.

If an external drive that was backed up within a previous backup job session is not connected to the Veeam Agent computer when the backup job must start, Veeam Agent for Microsoft Windows will back up data of all existing drives and skip the absent one. The backup job session will complete successfully.

- For volume-level backup, you can include in the backup a specific volume that resides on an external drive. To do this, you must select the necessary volume at the **Volumes** step of the **New Backup Job** wizard.

If a USB drive that contains the volume included in the backup is not connected to the Veeam Agent computer when the backup job must start, the backup job will complete with an error.

- For file-level backup, you can include in the backup data of a specific volume that resides on an external drive (entire volume or specific folders with files). To do this, you must select the necessary object at the **Files** step of the **New Backup Job** wizard.

If a USB drive that contains a volume or folder included in the backup is not connected to the Veeam Agent computer when the backup job must start, the backup job will complete with an error.

## NOTE

Veeam Agent for Microsoft Windows does not back up external drives used as rotated drives. USB drives registered as a known removable storage in Veeam Agent for Microsoft Windows are excluded from the backup.

# Data Restore

Veeam Agent for Microsoft Windows offers two data restore scenarios:

- You can perform volume-level restore to recover the entire system image of your computer or specific computer volumes. To learn more, see [Volume-Level Restore](#).

When you perform volume-level restore, you can resize restored volumes to fit available space on target location. To learn more, see [Volume Resize](#).

- You can perform file-level restore to recover individual files and folders. To learn more, see [File-Level Restore](#).

If the backup is encrypted, Veeam Agent will decrypt it during the restore process. To learn more, see [Restore from Encrypted Backups](#).

# Volume-Level Restore

If data on a computer volume gets corrupted, you can restore this volume from the backup. For volume-level restore, you can use backups that were created at the volume level. File-level backups cannot be used for volume restore.

When you perform volume-level restore, Veeam Agent for Microsoft Windows restores the entire content of the volume. It retrieves from the backup data blocks pertaining to a specific volume and copies them to the necessary location.

Keep in mind that you cannot browse the volume in the backup and select individual application items, files and folders for restore. For granular file-level restore, you can use the file-level restore option. To learn more, see [File-Level Restore](#).

A volume can be restored to its original location or new location. If you restore the volume to its original location, Veeam Agent for Microsoft Windows overwrites data on the original volume. If you restore the volume to a new location, and the target disk contains any data, Veeam Agent for Microsoft Windows overwrites data in the target location with data retrieved from the backup.

A volume can be restored to a new location that has greater or less space than the size of the volume in the backup. Depending on the amount of free disk space on target location, you can select either to shrink or to extend the volume during restore. To learn more, see [Volume Resize](#).

## Limitations for Volume-Level Restore

Volume restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the swap file is currently hosted.
- You cannot restore a volume to the volume where the backup file used for restore is located.

To overcome the first two limitations, you can create a Veeam Recovery Media and use the **Veeam Bare Metal Recovery** wizard for volume-level restore. To learn more, see [Veeam Recovery Media](#).

## Volume Resize

With Veeam Agent for Microsoft Windows, you can resize backup volumes during [volume-level restore](#). When you select to resize a volume, Veeam Agent for Microsoft Windows restores data from the backup and resizes the restored volume to the specified size.

There are two ways to resize a volume depending on the amount of free disk space on the target location:

- **Volume shrink** – you can shrink a volume when you restore it to a new location that has less space than the size of the volume in the backup. You can also shrink a volume that is restored to its original location to free disk space on the target location. To learn more, see [How Volume Shrink Works](#).
- **Volume extend** – you can extend a backup volume when you restore it to a new location that has more available disk space than the size of the backup volume. To learn more, see [How Volume Extend Works](#).

Volume resize may be also helpful when you need to restore data after hardware upgrade. For example, you may want to resize volumes in the following situations:

- Shrink backup data to restore system volumes of your computer to a smaller disk after you replace an old HDD drive with a faster but less capacitive SSD drive.
- Extend the backup volume during volume-level restore to a new, more capacitive HDD drive.

You can restore and resize volumes:

- With the **Volume Restore** wizard when you restore volumes under Microsoft Windows system. For more information, see [Restoring Volumes](#).
- With the **Veeam Bare Metal Recovery** wizard when you perform restore from the Veeam Recovery Media. For more information, see [Restoring from Veeam Recovery Media](#).

The volume resize option is available only in the **Manual restore** mode at the **Disk Mapping** step of the wizard.

## Limitations for volume resize

Volume resize has the following limitations:

- You cannot restore a volume to the volume of the smaller size if the amount of data stored on the backup volume exceeds the free space on the target disk.
- You can only resize basic volumes that use the NTFS file system.
- If you resize a BitLocker encrypted volume during restore, the restored volume will be unencrypted.

## How Volume Shrink Works

When you restore a volume to a target location of the smaller size, Veeam Agent for Microsoft Windows performs the following operations:

1. When you select the **Resize** option to shrink a volume, Veeam Agent for Microsoft Windows mounts the backup volume to a temporary NTFS folder on the system drive, for example:  
`C:\Users\Username\AppData\Local\Temp.`

2. Veeam Agent for Microsoft Windows mounts the created NTFS folder as a VHD disk next to other disks that are present on the computer.

Mounting VBK file content as a VHD disk makes it possible for Veeam Agent for Microsoft Windows to use Microsoft Windows system's disk management tools to measure current size of the backup volume and maximum and minimum size for the restored volume.

3. Veeam Agent for Microsoft Windows sends a query request to the mounted VHD disk to calculate its size, amount of stored data and free disk space by which the volume can be shrunk.

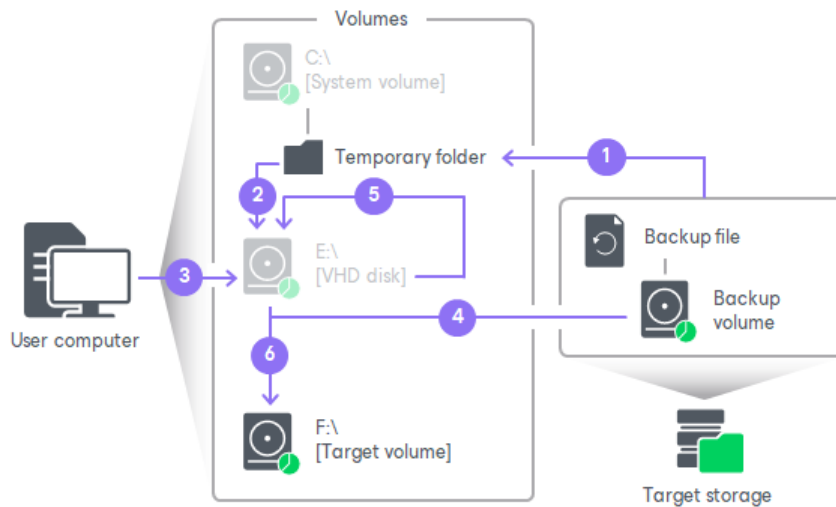
This step may take some time depending on the size of the backup volume and its data fragmentation ratio.

When the query is complete and you specify the desired size for the restored volume, Veeam Agent for Microsoft Windows unmounts the VHD disk.

4. When you start the restore process, Veeam Agent for Microsoft Windows creates on the target disk a volume of the specified size and restores to that volume the amount of backed-up data that fits the specified size.
5. Veeam Agent for Microsoft Windows mounts the backup volume as a VHD disk as described in steps 1 and 2 and starts to shrink it to the size of the target volume. During the process of volume shrink, empty data blocks from the part of the mounted VHD disk that does not fit the size of the target volume are moved to the part of the disk that contains actual data.

6. Veeam Agent for Microsoft Windows captures on the VHD disk data blocks that are moved during shrink and writes them to the target volume.

When all data blocks are written to the target volume, Veeam Agent for Microsoft Windows unmounts the VHD disk.



## How Volume Extend Works

When you restore a volume to a target location of the larger size, Veeam Agent for Microsoft Windows performs the following operations:

1. When you select the *Resize* option to extend a volume, Veeam Agent for Microsoft Windows mounts the backup volume to a temporary NTFS folder on the system drive, for example:

`C:\Users\Username\AppData\Local\Temp.`

2. Veeam Agent for Microsoft Windows mounts the created NTFS folder as a VHD disk next to other disks that are present on the computer.

Mounting VBK file content as a VHD disk makes it possible for Veeam Agent for Microsoft Windows to use Microsoft Windows system's disk management tools to measure current size of the backup volume and maximum and minimum size for the restored volume.

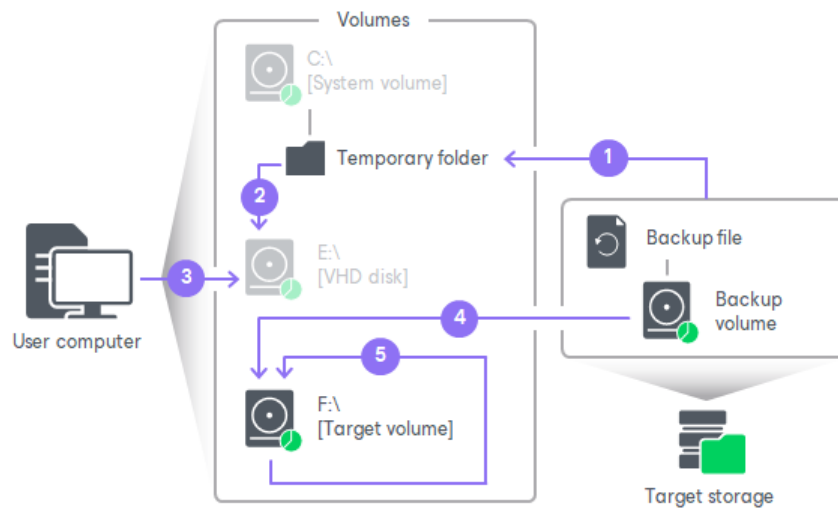
3. Veeam Agent for Microsoft Windows sends a query request to the mounted VHD disk to calculate its size, amount of stored data and free disk space by which the volume can be extended.

This step may take some time depending on the size of the backup volume and its data fragmentation ratio.

When the query is complete and you specify the desired size for the restored volume, Veeam Agent for Microsoft Windows unmounts the VHD disk.

4. When you start the restore process, Veeam Agent for Microsoft Windows creates on the target disk a volume of the same size as the backup volume and restores to that volume all data blocks from the backup volume.

- When all data blocks are written to the target location, Veeam Agent for Microsoft Windows extends the size of the target volume to the specified size.



# File-Level Restore

If you have lost or modified files and folders on your computer by mistake, you can restore a copy of the necessary objects from the backup. For file-level restore, you can use a backup of any type:

- Volume-level backup
- File-level backup

Veeam Agent for Microsoft Windows does not extract files and folders from the backup file. Instead, it uses Veeam's proprietary driver to publish the backup content directly into the computer file system, under `C:\VeeamFLR\<Volume N>`. For accessing the backup file content, Veeam Agent for Microsoft Windows uses a separate program – Virtual Disk Driver (VDK) that is provided with the product.

After the backup content is mounted, you can use a built-in Veeam Backup browser or Microsoft Windows Explorer to browse and copy necessary files and folders to your local machine drive, save them in a network shared folder or point applications to files and work with them in a regular way.



# Restore from Encrypted Backups

When you restore data from an encrypted backup, Veeam Agent for Microsoft Windows performs data decryption in the following ways:

- If encryption keys required to unlock the backup file are available in the Veeam Agent for Microsoft Windows database, you do not need to specify the password. Veeam Agent for Microsoft Windows uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore from the encrypted backup does not differ from that from an unencrypted one.

Automatic data decryption can be performed in one of the following situations:

- You encrypt and decrypt the backup file on the same Veeam Agent computer using the same Veeam Agent for Microsoft Windows database.
- You have included encryption keys into the Veeam Recovery Media and perform bare metal recovery after booting from this Veeam Recovery Media. To learn more, see [Specify Recovery Media Options](#).
- If encryption keys are not available in the Veeam Agent for Microsoft Windows database, you need to provide a password to unlock the encrypted file. The password must be the same as the password that was used to encrypt the backup file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent for Microsoft Windows, you can use the latest password to restore data from all restore points in the backup chain, including restore points that were encrypted with an old password and restore points that were created before you have enabled the encryption option for the job.

## NOTE

If you store backups in the Veeam backup repository, encryption keys are not available in the Veeam Agent for Microsoft Windows database but Veeam Agent does not require a password. Such backups are encrypted and decrypted on the Veeam Backup & Replication side, and Veeam Agent considers such backups unencrypted. To learn more, see [Data Encryption](#).

## Related Topic

[Data Encryption](#)

## Related Task

[Restoring Data from Encrypted Backups](#)

# Veeam Recovery Media

Veeam Agent for Microsoft Windows lets you create a Veeam Recovery Media — a recovery image of your computer.

The recovery image is a "copy" of your OS with the limited functionality — it contains all data required to run Microsoft Windows Recovery Environment (Windows RE), and provides an alternative way to boot your computer. If the OS installed on the computer fails to start for some reason, you can boot the Windows RE from the recovery image. After booting, you can do the following:

- You can use Veeam Agent for Microsoft Windows and Microsoft Windows tools to diagnose problems and fix errors on your computer.
- You can restore data from a backup to your computer. For this scenario, you must have a backup created with Veeam Agent for Microsoft Windows or system image created with Microsoft Windows.

The recovery image can be helpful if one of the following errors occur:

- The OS on the computer fails to start.
- The computer is blocked with malware and you cannot get access to your data.
- You want to perform bare metal restore from the backup on the computer without the OS and other software installed.
- You want to restore the system volume of the computer and so on.

You can create a recovery image on different kinds of media:

- Removable storage devices such as USB drives or SD cards
- CD/DVD/BD
- ISO images on local or external computer drives

When you boot from the Veeam Recovery Media, you can use the Veeam Agent for Microsoft Windows recovery environment to fix the OS system errors on your computer or restore data from the backup. Veeam Agent for Microsoft Windows offers a set of tools for the computer system image and data recovery:

- Bare Metal Recovery — the Veeam Agent for Microsoft Windows wizard to recover data on the original computer or a new computer.
- Windows Recovery Environment — a built-in Microsoft Windows tool to recover the computer system image.
- Tools — Veeam Agent for Microsoft Windows and Microsoft Windows utilities for advanced computer administration.

## Limitations for Veeam Recovery Media

- You cannot restore dynamic volumes using a Veeam Recovery Media. To restore dynamic volumes, you can recover data from the volume-level backup on a working computer system. To learn more, see [Restoring Volumes](#).
- The Veeam Recovery Media is based on the Microsoft Windows RE. Due to Microsoft limitations, Microsoft Windows RE automatically reboots after 72 hours of continuous use. All data that has not been saved before reboot will be lost.

# Drivers in Veeam Recovery Media

The Veeam Recovery Media created with Veeam Agent for Microsoft Windows contains the following data:

- Set of files required to start your computer OS from the recovery media.
- Diagnostic tools from Microsoft and Veeam.
- Drivers required to run hardware and devices on your computer in a regular way. When you boot your computer from the Veeam Recovery Media, drivers included into the Veeam Recovery Media are automatically loaded on the recovered OS.
- Network connection settings from your computer. When you boot your computer from the Veeam Recovery Media, network settings included into the Veeam Recovery Media are automatically applied and can be used to connect to the remote backup storage.
- If you have enabled data encryption options for the backup job, you can also include a decryption key into the Veeam Recovery Media. To learn more, see [Creating Veeam Recovery Media](#).

## NOTE

The Veeam Recovery Media contains all locales (languages) that are included in the OS of the Veeam Agent Computer.

You can include the following drivers in the Veeam Recovery Media:

- Drivers that are currently installed on your computer. Veeam Agent for Microsoft Windows detects hard disk controller drivers, network adapter drivers and USB controller drivers and includes them into the Veeam Recovery Media.
- Additional storage and network drivers. If you use non-standard drivers, you can include them in the created Veeam Recovery Media manually. For example, you can include drivers for a discrete network card, third-party USB 3.0 controllers and non-standard hard disk controllers.

## TIP

If you do not include some drivers in the Veeam Recovery Media, you can load them from the computer drive when you perform bare metal recovery. To learn more, see [Restoring from Veeam Recovery Media](#).

# BitLocker Encrypted Volumes Support

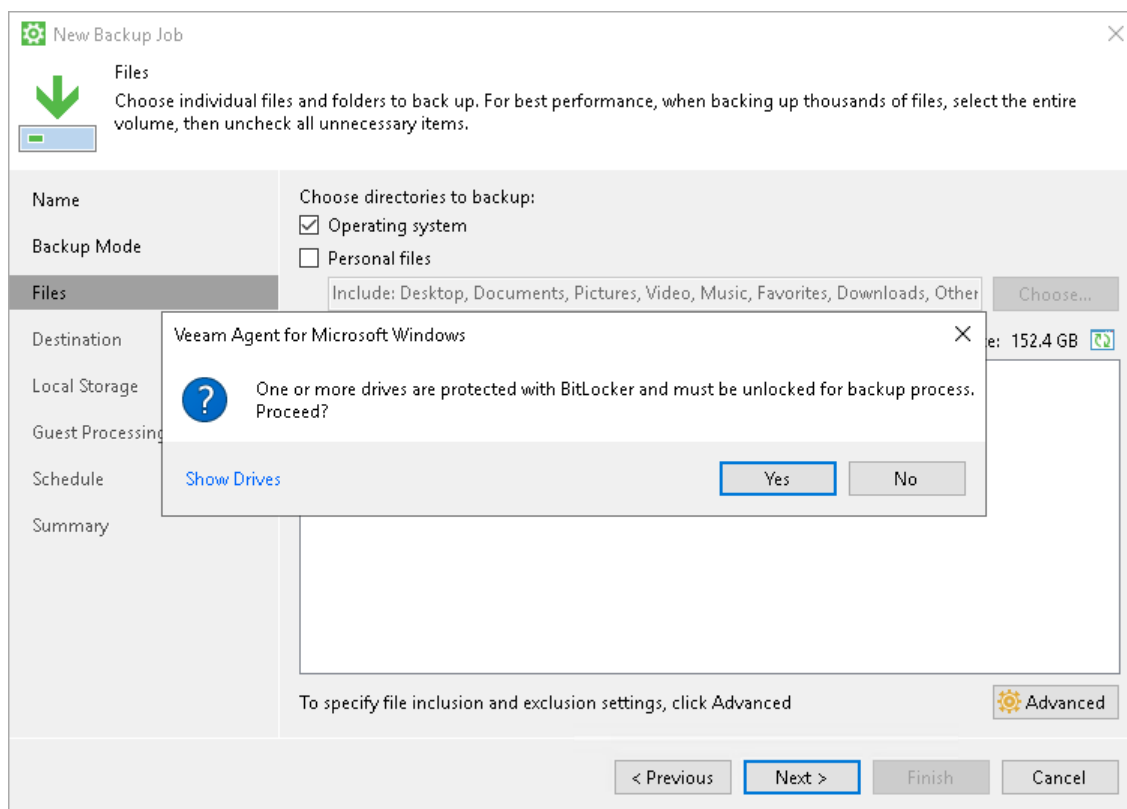
Veeam Agent for Microsoft Windows supports scenarios of data backup and restore to/from volumes encrypted with Microsoft Windows BitLocker.

## Data Backup

You can create backups of BitLocker encrypted volumes and store backups created with Veeam Agent for Microsoft Windows on BitLocker encrypted volumes.

BitLocker encrypted volumes (both source and target) must be unlocked at the moment when Veeam Agent for Microsoft Windows starts the backup operation.

- If the volume added to the backup scope is locked at the moment of backup, the backup job will be unable to process it and will fail.
- If the volume on which the backup file must be stored is locked at the moment of backup, the backup job will be unable to save the resulting file, and the job will fail.



## Data Restore

You can restore data from backups stored on BitLocker encrypted volumes and restore data to BitLocker encrypted volumes.

Veeam Agent for Microsoft Windows restores volumes in their initial state:

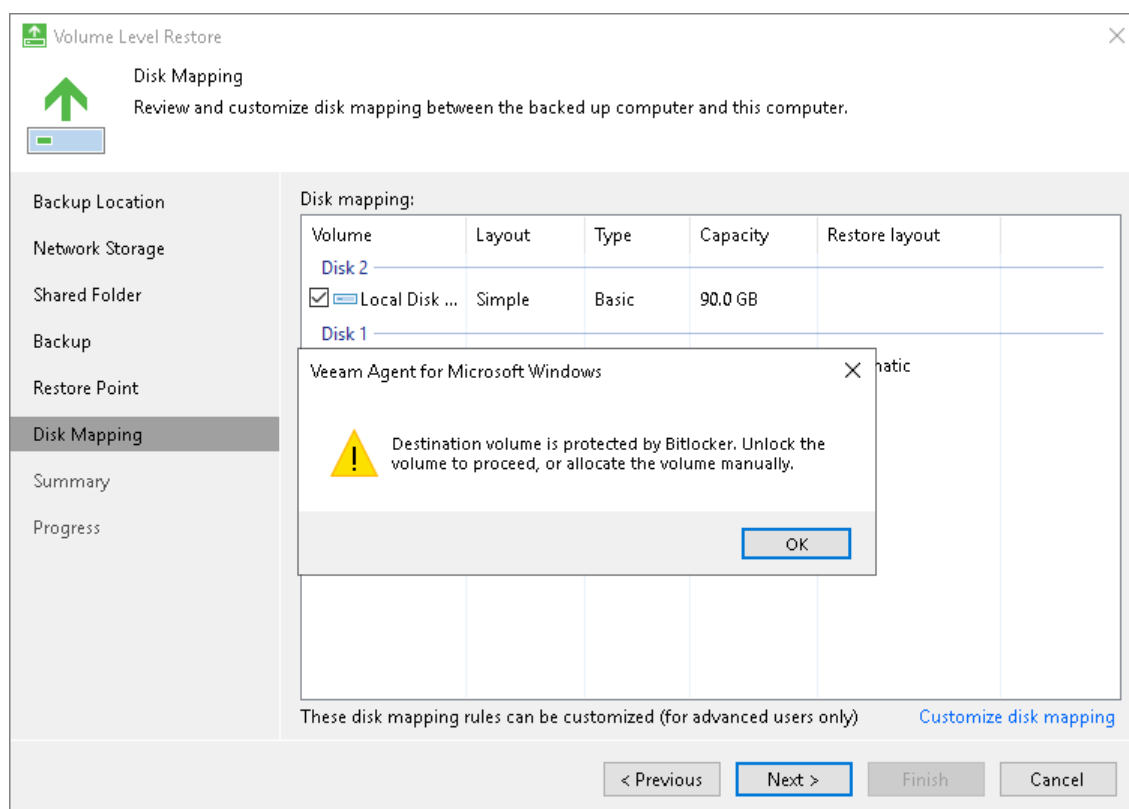
- If you restore an encrypted volume to its original location, the restored volume will be encrypted.
- If you restore an unencrypted volume to an encrypted volume, the restored volume will be unencrypted.

## IMPORTANT

If you resize a BitLocker encrypted volume during restore, the restored volume will be unencrypted. To learn more about volume resize, see [Volume Resize](#).

BitLocker encrypted volumes must be unlocked at the moment when you perform the restore operation.

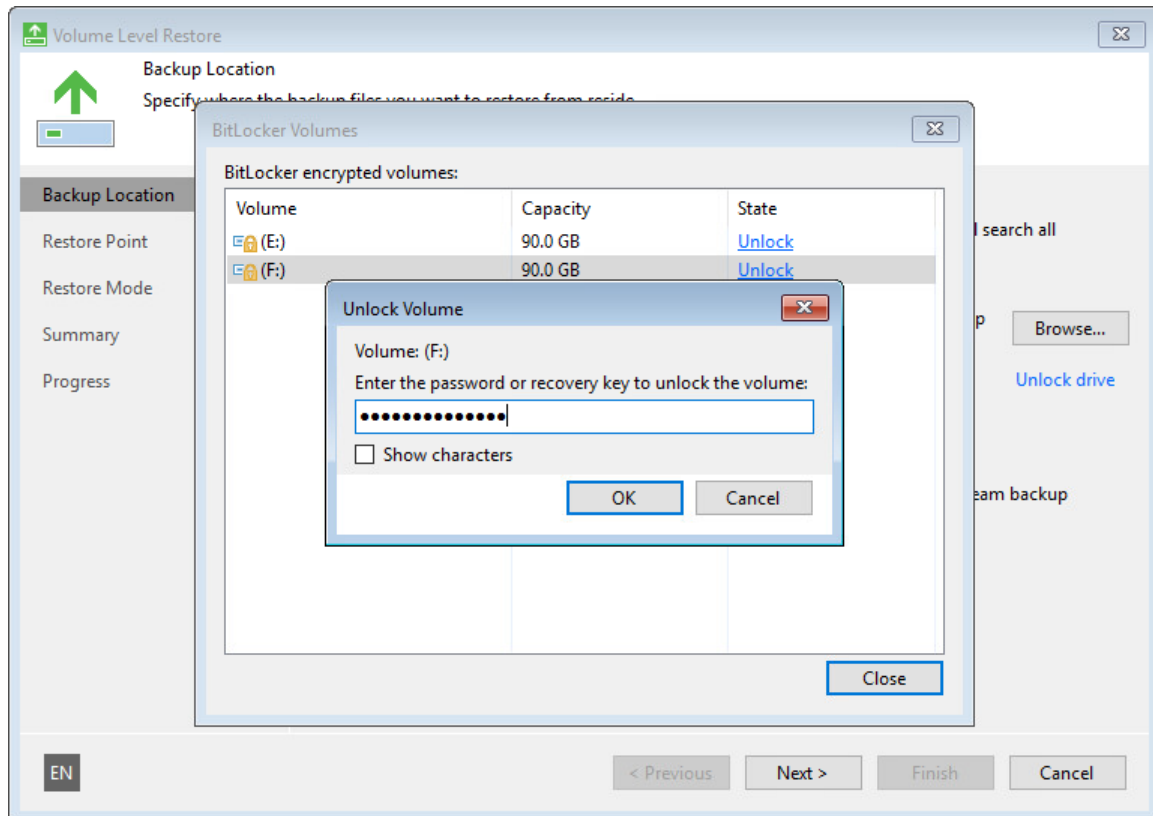
- If the backup file is stored on a locked volume, Veeam Agent for Microsoft Windows will fail to access it, and you will not be able to restore data from it.
- If you perform volume-level restore, and the target volume is locked, Veeam Agent for Microsoft Windows will display a warning and will ask you to unlock the volume. You can do this using the Microsoft Windows UI.



# Veeam Recovery Media

If you boot from the Veeam Recovery Media, you can restore data from backups stored on BitLocker encrypted volumes and restore data to BitLocker encrypted volumes.

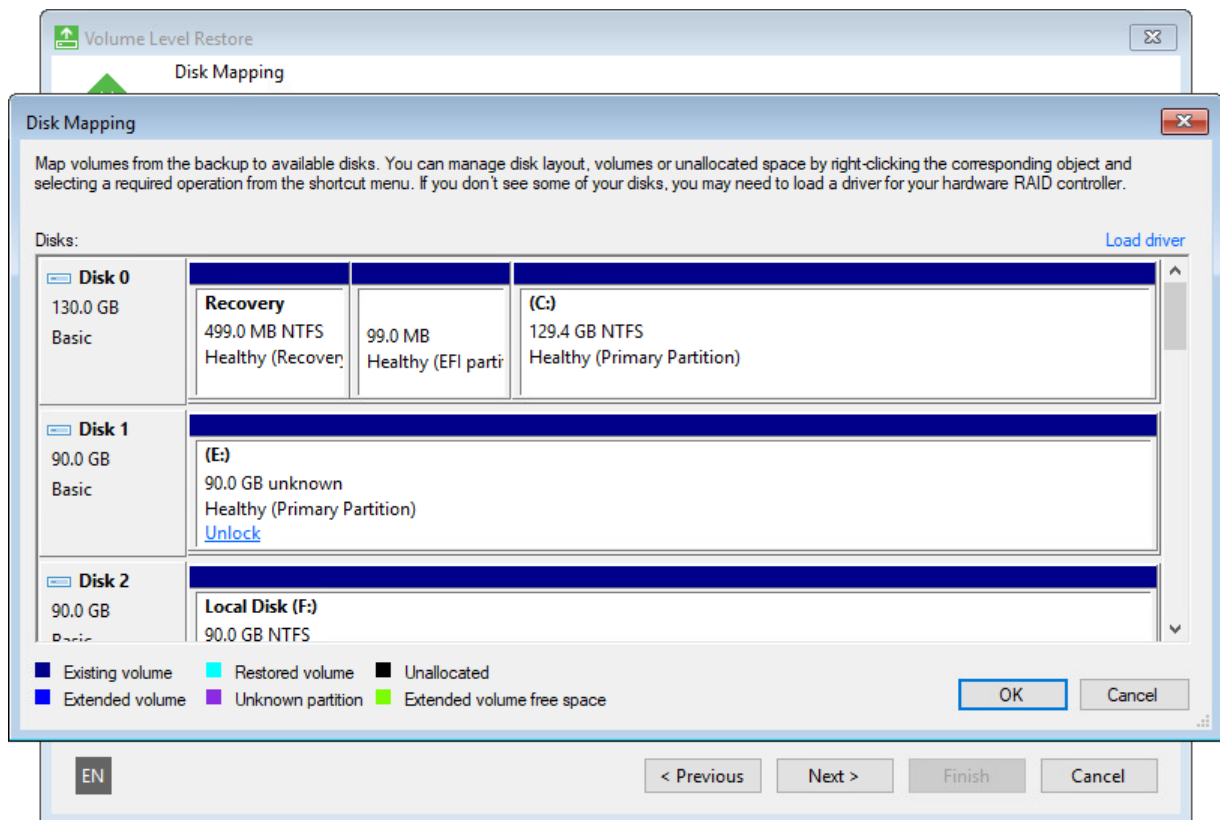
- If the backup file that you want to use for data restore resides on a locked volume, Veeam Agent for Microsoft Windows cannot access this backup file. To unlock the volume with the backup file, click **Unlock drive** under the **Backup file** field and enter a password for the volume.



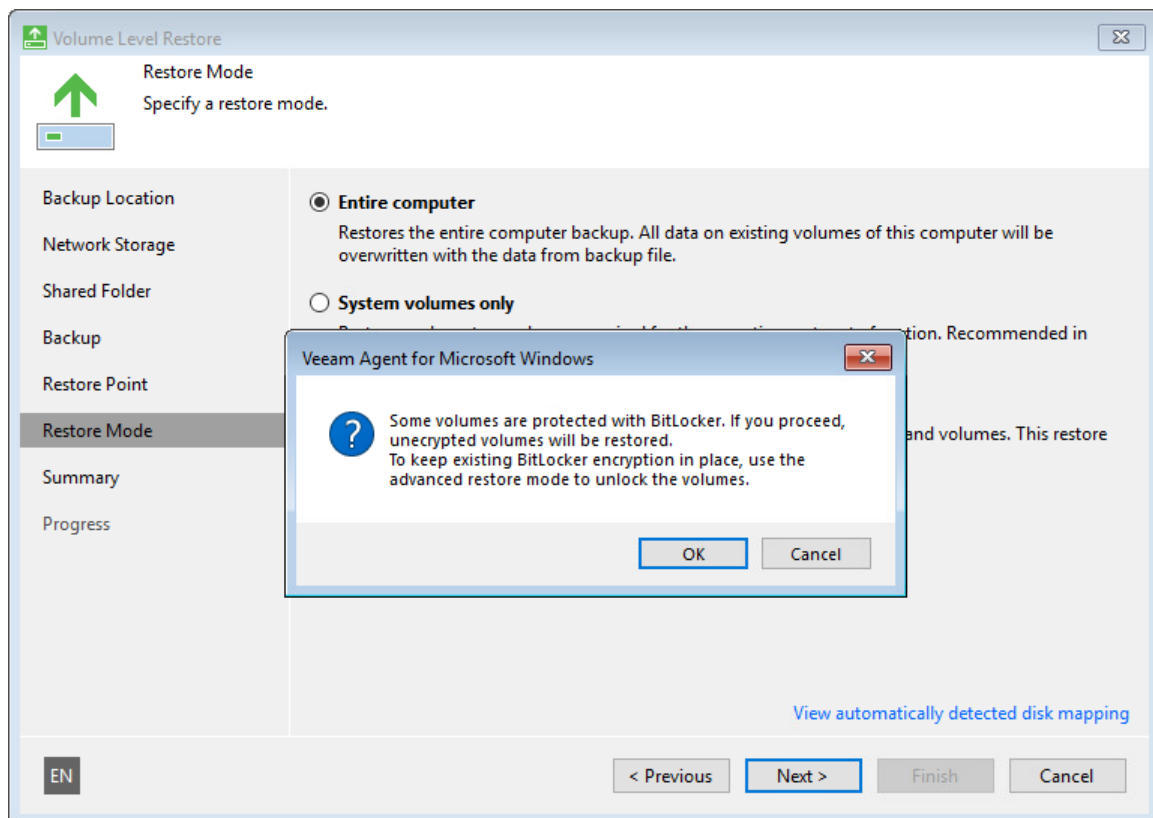
- If you restore a volume to its original location when it is BitLocker encrypted and locked, Veeam Agent tries to unlock the target volume automatically and keep BitLocker encryption enabled for the volume. If Veeam Agent cannot unlock the volume, it displays a warning message at the **Restore Mode** step of the wizard. In this case, you can use one of the following scenarios:

- You can restore data to the target volume and keep BitLocker encryption enabled for the volume. To do this, you must unlock the volume before you start data restore.

To unlock the volume, click **Cancel** in the warning window. At the **Restore Mode** step of the wizard, select **Manual Restore**. At the **Disk Mapping** step of the wizard, click **Customize disk mapping** and click **Unlock** under the necessary volume.



- You can restore data to the target volume and disable BitLocker encryption for the volume. To do this, click **OK** in the warning window. Veeam Agent will delete existing BitLocker encrypted partitions on the volume, format the disk and restore data from the backup as unencrypted.



- If you restore a volume to a new location that is BitLocker encrypted and locked, Veeam Agent will delete existing BitLocker encrypted partitions on the volume, format the disk and restore data from the backup as unencrypted.

## IMPORTANT

Veeam Agent for Microsoft Windows cannot back up volumes formatted as FAT32 and encrypted with BitLocker. In general, FAT32 does not allow storing VSS snapshots on the same volume. When Veeam Agent for Microsoft Windows triggers a VSS snapshot of a FAT32 formatted volume, the VSS snapshot is stored on another, non-FAT32 volume on the computer.

If BitLocker is enabled, the VSS cannot save the snapshot on another volume due to Microsoft limitations, and the backup process fails.



# Integration with Veeam Backup & Replication

If you plan to use Veeam Agent for Microsoft Windows 6.2 with Veeam Backup & Replication, you must install Veeam Backup & Replication 12.2 on the Veeam backup server.

## NOTE

You can also use Veeam Backup & Replication to manage Veeam Agent for Microsoft Windows on computers in your infrastructure. Within the Veeam Agent management scenario, you can remotely deploy Veeam Agent for Microsoft Windows to your computers, as well as configure and manage Veeam Agent backup jobs in Veeam Backup & Replication. To learn more, see the [Veeam Agent Management Guide](#).

The current guide covers subjects related to Veeam Agent for Microsoft Windows operating in the standalone mode.

You can store backup files created with Veeam Agent for Microsoft Windows in backup repositories managed by Veeam Backup & Replication. To do this, you must select a backup repository as a target location in the properties of the Veeam Agent backup job. To learn more about backup repositories, see the [Backup Repositories](#) and [Scale-Out Backup Repositories](#) sections in the Veeam Backup & Replication User Guide.

Veeam Agent for Microsoft Windows works with the backup repository as with any other target location. Backup files are stored in a separate folder; you can perform standard restore operations using these files.

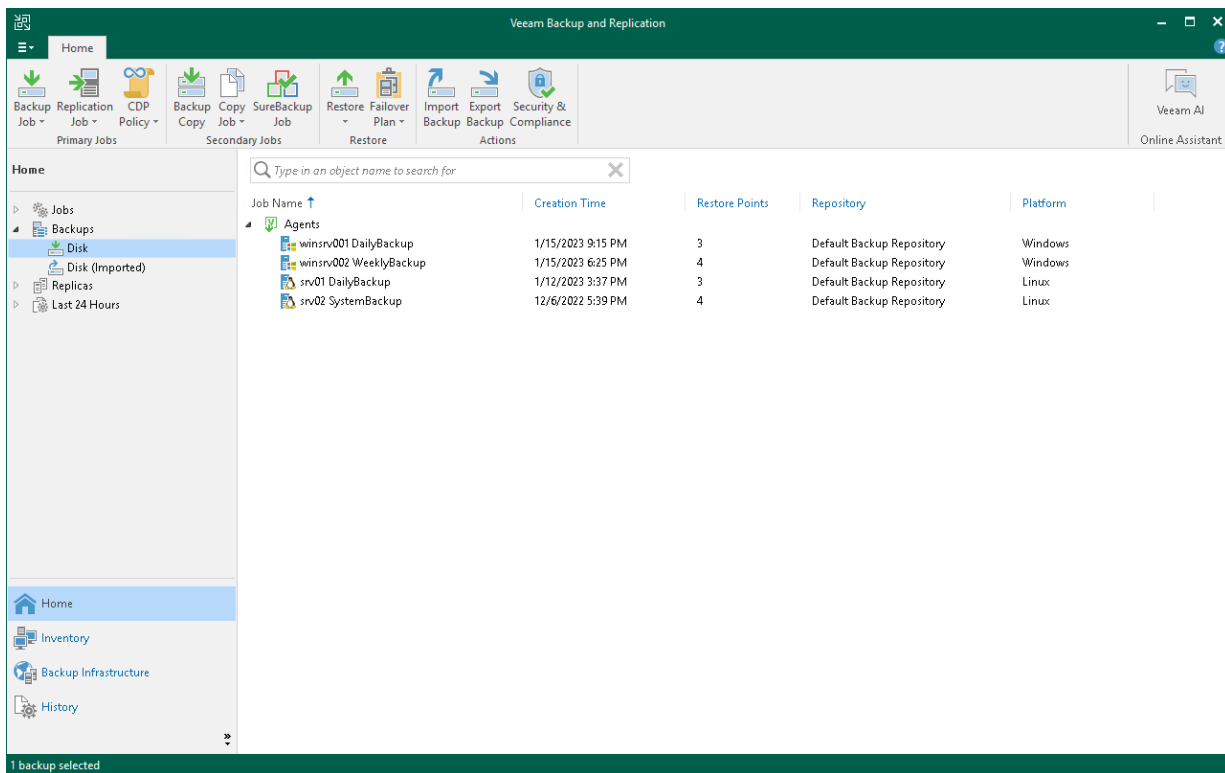
Information about Veeam Agent backups stored in the backup repositories, backup jobs and sessions becomes available in the Veeam Backup & Replication console:

- The Veeam Agent for Microsoft Windows backup job is displayed in the list of jobs in Veeam Backup & Replication.
- Backup files created with Veeam Agent for Microsoft Windows are displayed in the list of backups, under the **Backups > Disk** or **Backups > Object Storage** node.
- Performed job sessions are available in the **History** view of Veeam Backup & Replication.

Backup administrators working with Veeam Backup & Replication can perform a set of operations with Veeam Agent backups:

- Perform data protection operations: copy Veeam Agent backups to secondary backup repositories and archive these backups to tape; scan backups to detect malware or sensitive data.
- Perform restore operations: restore individual files and folders, application items; restore computer disks and convert them to the VMDK, VHD or VHDX format; restore Veeam Agent backups to Microsoft Azure, Amazon Elastic Compute Cloud, Google Compute Engine or to Hyper-V VMs; publish disks from backups.

- Perform administrative tasks: disable and delete Veeam Agent backup jobs, remove Veeam Agent backups and so on.



# Backup to Veeam Cloud Connect Repository

If you want to store your data in the cloud, you can connect to a Veeam Cloud Connect service provider (SP) and create Veeam Agent backups in a cloud repository.

## Limitations for Backup to Cloud Repository

Backup to a Veeam Cloud Connect repository has the following limitations:

- You can store backups in a Veeam Cloud Connect repository if you use the Workstation or Server edition of Veeam Agent for Microsoft Windows only.
- You cannot use a Veeam Cloud Connect repository as a target for standalone full backup. To learn more, see [Standalone Full Backup](#).
- Veeam Agent for Microsoft Windows does not support creating transaction log backups in the cloud repository. You cannot enable transaction log backup options in the properties of the backup job targeted at the cloud repository.

# Managing Veeam Agent in Veeam Backup & Replication

Veeam Backup & Replication lets you automate management of Veeam Agent for Microsoft Windows on multiple computers in your infrastructure. You can deploy Veeam Agent for Microsoft Windows, configure Veeam Agent backup jobs and perform other data protection and administration tasks on remote computers. To use the Veeam Agent management functionality in Veeam Backup & Replication, you must install Veeam Backup & Replication on the Veeam backup server.

To learn more, see the [Veeam Agent Management Guide](#).

# Planning and Preparation

Before you install Veeam Agent for Microsoft Windows, make sure that the target computer meets the system requirements and all required ports are open.

# System Requirements

The protected computer must meet requirements listed in the table below.

NOTE

The following system requirements apply to Veeam Agent for Microsoft Windows operating in the standalone mode.

To learn about system requirements for Veeam Agent for Microsoft Windows managed by Veeam Backup & Replication, see the [System Requirements](#) section in the Veeam Agent Management Guide.

Specification	Requirement
Hardware	<p>CPU: x86-64 processor.</p> <p>Memory: 2 GB RAM or more. Memory consumption varies depending on number and size of processed disks.</p> <p>Disk Space: 200 MB for product installation.</p> <p>Network: 1 Mbps or faster. High latency and reasonably unstable WAN links are supported.</p> <p>System firmware: BIOS or UEFI.</p> <p>Drive encryption: Microsoft BitLocker (optional). BitLocker encrypted volumes must be unlocked at the moment when Veeam Agent starts the backup or restore operation. Only Microsoft BitLocker is supported for drive encryption. Other drive encryption products are not supported.</p>

Specification	Requirement
OS	<p>Both 64-bit and 32-bit (where applicable) versions of the following operating systems are supported<sup>1</sup>:</p> <ul style="list-style-type: none"> <li>• Microsoft Windows Server 2022</li> <li>• Microsoft Windows Server 2019</li> <li>• Microsoft Windows Server 2016</li> <li>• Microsoft Windows Server General Availability Channel (from version 1803 to version 20H2)</li> <li>• Microsoft Windows Server 2012 R2</li> <li>• Microsoft Windows Server 2012</li> <li>• Microsoft Windows Server 2008 R2 SP1<sup>2</sup></li> <li>• Microsoft Windows 11 (from version 21H2 to version 23H2)</li> <li>• Microsoft Windows 10 (from version 1909 to version 22H2)</li> <li>• Microsoft Windows 10 Long-Term Servicing Channel (versions 2015, 2016, 2019)</li> <li>• Microsoft Windows 8.1</li> <li>• Microsoft Windows 7 SP1</li> </ul> <p><sup>1</sup> Consider the following:</p> <ul style="list-style-type: none"> <li>• Small Business Server, Server Essentials, and Server Storage editions of Microsoft Windows Server OSes are supported.</li> <li>• Server Core installations of Microsoft Windows Server OSes are supported for Veeam Agent for Microsoft Windows managed by Veeam Backup &amp; Replication only. To learn more, see the <a href="#">Veeam Agent Management Guide</a>. Veeam Agent operating in the standalone mode does not support Server Core installations.</li> <li>• Microsoft Windows OSes installed on ReFS boot partitions are not supported.</li> <li>• Windows Embedded / Windows IoT OSes are supported (except for custom builds by certain vendors that do not have components required for Veeam Agent operation).</li> <li>• Microsoft Failover Clusters are supported for Veeam Agent for Microsoft Windows managed by Veeam Backup &amp; Replication only. To learn more, see the <a href="#">Veeam Agent Management Guide</a>. Veeam Agent operating in the standalone mode does not support Microsoft Failover Clusters.</li> </ul> <p><sup>2</sup>Veeam CBT driver is supported only if Microsoft Windows update <a href="#">KB3033929</a> is installed on the Veeam Agent computer.</p>

Specification	Requirement
<b>File System</b>	<p>Microsoft Windows FAT32/exFAT, NTFS, ReFS file systems are supported.</p> <p>The supported file system must reside on a volume that is 64 TB or smaller, because Veeam Agent uses the Microsoft Software Shadow Copy Provider to create a volume shadow copy during the backup. To learn more about the limitation, see <a href="#">Microsoft documentation</a>.</p>
<b>Software</b>	<p>The following required 3rd party software is included in the setup program:</p> <ul style="list-style-type: none"> <li>• Microsoft .NET Framework 4.5.2</li> <li>• Windows Universal C Runtime Library</li> </ul> <p>When installing the product, the setup program checks whether all prerequisite software is available on the target computer. If some of the required software components are missing, the missing software is installed automatically.</p>
<b>Database</b>	SQLite database engine (installed with the product).

Consider the following:

- Veeam Agent for Microsoft Windows works with only those hard drive types that are supported by the Microsoft Windows OS. Thus, Veeam Agent supports the 512 bytes and 4 KB sector hard drives only. Other hard drive types are not supported. To learn more, see [Microsoft documentation](#).
- Supported culture settings depend on the version of Microsoft Windows OS installed on your computer. For more information, see [Microsoft documentation](#).

## Backup Source

Veeam Agent for Microsoft Windows supports backup of data that resides in the following types of storage:

- Local (internal) storage of the protected computer.
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives (USB sticks and SD cards are not supported).
- Storage Area Network (SAN), such as iSCSI connected volumes.

Keep in mind that Veeam Agent cannot work with third-party volume managers installed on the protected computer. Such managers may not allow Veeam Agent to interact with necessary interfaces and services properly.

## Backup Target

Backup can be performed to the following types of storage:

### *Disk-based storage*

- Local (internal) storage of the protected computer (not recommended).



- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives, and raw device mapping (RDM) volumes.

#### **IMPORTANT**

Storage devices with the exFAT file system are not supported as a backup target.

- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) share.
- Storage Area Network (SAN), such as iSCSI connected volumes.
- On-premises object storage.
- Veeam Backup & Replication 12.2 backup repository.

#### *Cloud storage*

- Cloud-based object storage.
- Veeam Cloud Connect 12.0 or later cloud repository.

## Network

Consider the following:

- If you back up to a repository managed by a Veeam backup server, Veeam Agent for Microsoft Windows must be able to establish a direct IP connection to the Veeam Backup & Replication server. Veeam Agent cannot work with Veeam Backup & Replication that is located behind a NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 or IPv6 addresses.

# Permissions

Depending on the scenario, the user accounts must have the permissions listed in the following subsections:

- [Permissions for Backup to Object Storage](#)
- [Permissions for Guest Processing](#)

If you plan to back up data to object storage, make sure that the user account that you use to connect to the object storage has the required permissions. The list of required permissions differs depending on the selected object storage:

- [Amazon S3 or S3 compatible](#)
- [Google Cloud Storage](#)

## Amazon S3 or S3 compatible

If you plan to back up data to the Amazon S3 or S3 compatible storage, make sure the user account that you plan to use has the following permissions:

Identity-based permission:

```
{
  "s3:ListAllMyBuckets"
}
```

Resource-based permissions:

```
{
  "s3:DeleteObject",
  "s3:GetBucketLocation",
  "s3:GetBucketObjectLockConfiguration",
  "s3:GetBucketVersioning",
  "s3:GetObject",
  "s3:ListBucket",
  "s3:PutObject"
}
```

### TIP

For information about required permissions for Amazon S3 storage with immutability enabled, see the [Using Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

## Google Cloud Storage

If you plan to back up data to the Google Cloud storage, make sure the user account that you plan to use has the following permissions:

```
{
  "storage.buckets.get",
  "storage.buckets.list",
  "storage.objects.create",
  "storage.objects.delete",
  "storage.objects.get",
  "storage.objects.list"
}
```

## Permissions for Guest Processing

To use guest processing, make sure to configure user accounts according to the requirements listed in this section.

Consider the following general requirements when choosing a user account:

- [For file indexing] Choose a user account that has administrator privileges.
- When using Active Directory accounts, make sure to provide a user account in the *DOMAIN\Username* format.
- When using local user accounts, make sure to provide a user account in the *Username or HOST\Username* format.
- To process a Domain Controller server, make sure that you are using a user account that is a member of the *DOMAIN\Administrators* group.
- To back up a Read-Only Domain controller, a delegated RODC administrator account is sufficient. For more information, see [Microsoft documentation](#).

Depending on the application you need to back up, the user must have the permissions listed in the table below:

Application	Required Permission
<b>Microsoft SQL Server</b>	<p>To back up Microsoft SQL Server data, the user whose account you plan to use must be:</p> <ul style="list-style-type: none"> <li>• Local Administrator on the Veeam Agent computer.</li> <li>• System administrator (has the <i>Sysadmin</i> role) on the target Microsoft SQL Server.</li> </ul> <p>If you need to provide minimal permissions, the user account must be assigned the following roles and permissions:</p> <ul style="list-style-type: none"> <li>• SQL Server instance-level role: <i>public</i> and <i>dbcreator</i>.</li> <li>• Database-level roles and roles for the model system database: <i>db_backupoperator</i>, <i>db_denydatareader</i>, <i>public</i>; for the master system database – <i>db_backupoperator</i>, <i>db_datareader</i>, <i>public</i>; for the msdb system database – <i>db_backupoperator</i>, <i>db_datareader</i>, <i>public</i>, <i>db_datawriter</i>.</li> <li>• Securables: <i>view any definition</i>, <i>view server state</i>, <i>connect SQL</i>.</li> </ul>
<b>Microsoft Active Directory</b>	To back up Microsoft Active Directory data, the user account must be a member of the built-in <i>Administrators</i> group.
<b>Microsoft Exchange</b>	To back up Microsoft Exchange data, the user account must have the Local Administrator permissions in Microsoft Exchange.
<b>Oracle</b>	<p>To back up Oracle data, the user account must be configured as follows:</p> <ul style="list-style-type: none"> <li>• The user account must be a member of both the <i>Local Administrators</i> group and the <i>ORA_DBA</i> group (if OS authentication is used).</li> <li>• The user account must be granted <i>SYSDBA</i> privileges.</li> </ul>
<b>Microsoft SharePoint</b>	<p>To back up Microsoft SharePoint server, the user account must have the <i>Farm Administrator</i> role.</p> <p>To back up Microsoft SQL databases of the Microsoft SharePoint Server, the user account must have the same privileges as for the <a href="#">Microsoft SQL Server</a>.</p>

# Ports

## NOTE

The following tables describe network ports that must be opened to ensure proper communication of Veeam Agent operating in the standalone mode with other infrastructure components.

To learn about ports required to enable proper work of Veeam Agent for Microsoft Windows managed by Veeam Backup & Replication, see the [Ports](#) section in the Veeam Agent Management Guide.

## IMPORTANT

The list of ports required for computers booted from the Veeam Recovery Media is the same as the list of ports required for Veeam Agent computers.

## Communication Between Veeam Agent Components

The following table describes network ports that must be opened to enable proper communication between Veeam Agent for Microsoft Windows components.

From	To	Protocol	Port	Notes
Veeam Agent Computer	Veeam Agent Computer	TCP	9395+, 6183+	Ports used locally on the Veeam Agent computer for communication between Veeam Agent components and Veeam Agent for Microsoft Windows Service.  If the default port number is already in use, Veeam Agent for Microsoft Windows Service will try to use the next port number.
	Veeam Update Notification Server (agents.butler.veeam.com)	TCP	443	Default port used to download information about available updates from the Veeam Update Notification Server over the Internet.

# Communication with Veeam Backup & Replication Repositories

The following table describes network ports that must be opened to ensure proper communication with Veeam backup repositories.

From	To	Protocol	Port	Notes
Veeam Agent Computer	Veeam Backup Server	TCP	10001	<p>Default port used by Veeam Agent for Microsoft Windows operating in the standalone mode for communication with the Veeam Backup server.</p> <p>Data between the Veeam Agent computer and backup repositories is transferred directly, bypassing Veeam backup servers.</p>
	Linux server performing the role of a backup repository	TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
	Microsoft Windows server performing the role of a backup repository	TCP	49152-65535	Dynamic RPC port range. For more information, see <a href="#">Microsoft documentation</a> .
		TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
	Shared folder SMB (CIFS) share	TCP UDP	137 to 139, 443, 445	<p>Ports used as a transmission channel from the Veeam Agent computer to the target SMB (CIFS) share.</p> <p>Ports 137 to 139 are used by backup infrastructure components to communicate using NetBIOS.</p> <p>Port 443 is used to connect the target SMB (CIFS) share over QUIC. For more information, see <a href="#">Microsoft documentation</a>.</p>

From	To	Protocol	Port	Notes
	Gateway Microsoft Windows server	TCP UDP	137 to 139, 445	<p>If an SMB (CIFS) share is used as a backup repository and a Microsoft Windows server is selected as a gateway server for this CIFS share, these ports must be opened on the gateway Microsoft Windows server.</p> <p>Ports 137 to 139 are used by backup infrastructure components to communicate using NetBIOS.</p>
		TCP	49152-65535	Dynamic RPC port range. For more information, see <a href="#">Microsoft documentation</a> .
		TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.

## Communication with Veeam Cloud Connect Repositories

The following table describes network ports that must be opened to ensure proper communication with Veeam Cloud Connect repositories.

From	To	Protocol	Port	Notes
Veeam Agent Computer	Cloud gateway	TCP	6180	Port on the cloud gateway used to transport Veeam Agent data to the Veeam Cloud Connect repository.
	Certificate Revocation Lists	TCP	80 or 443 (most popular)	<p>Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider.</p> <p>Generally, information about CRL locations can be found on the CA website.</p>

## Communication with Object Storage

The following table describes network ports that must be opened to ensure proper communication with object storage if you back up data to object storage directly or to object storage added as a Veeam backup repository with the direct connection mode. For more information about object storage connection modes, see [Connection Types](#).

From	To	Protocol	Port	Notes
Veeam Agent Computer	Amazon S3 object storage	TCP	443	Used to communicate with the Amazon S3 object storage through the following endpoints: <ul style="list-style-type: none"><li>• *.amazonaws.com (for both <i>Global</i> and <i>Government</i> regions)</li><li>• *.amazonaws.com.cn (for <i>China</i> region)</li></ul> All AWS service endpoints are specified in the <a href="#">AWS documentation</a> .
			80	Used to verify the certificate status through the following endpoints: <ul style="list-style-type: none"><li>• *.amazontrust.com</li><li>• *.cloudfront.net</li></ul> Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.
	Microsoft Azure object storage	TCP	443	Used to communicate with the Microsoft Azure object storage through the following endpoints: <ul style="list-style-type: none"><li>• xxx.blob.core.windows.net (for <i>Global</i> region)</li><li>• xxx.blob.core.chinacloudapi.cn (for <i>China</i> region)</li><li>• xxx.blob.core.usgovcloudapi.net (for <i>Government</i> region)</li></ul> Consider that the <xxx> part of the address must be replaced with your actual storage account URL that can be found in the Azure management portal.



From	To	Protocol	Port	Notes
			80	<p>Used to verify the certificate status through the following endpoints:</p> <ul style="list-style-type: none"> <li>ocsp.digicert.com</li> <li>ocsp.msocsp.com</li> </ul> <p>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself. For more details, see also <a href="#">Microsoft documentation</a>.</p>
	Google Cloud storage	TCP	443	<p>Used to communicate with Google Cloud storage through the following endpoints:</p> <ul style="list-style-type: none"> <li>storage.googleapis.com</li> </ul> <p>All cloud endpoints are specified in <a href="#">this Google article</a>.</p>
			80	<p>Used to verify the certificate status through the following endpoints:</p> <ul style="list-style-type: none"> <li>ocsp.pki.goog</li> <li>pki.goog</li> <li>crl.pki.goog</li> </ul> <p>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.</p>
	IBM Cloud object storage	TCP	Depends on device configuration	Used to communicate with IBM Cloud object storage.
	S3 compatible object storage	TCP	Depends on device configuration	Used to communicate with S3 compatible object storage.
	Veeam Data Cloud Vault storage	TCP	443	Used to communicate with the Veeam Data Cloud Vault storage through the <code>xxx.blob.core.windows.net</code> endpoint.

## Communication with Mail Servers

The following table describes network ports that must be opened to ensure proper communication with mail servers.

From	To	Protocol	Port	Notes
Veeam Agent Computer	SMTP server	TCP	25	Default port used by the SMTP server.
		TCP	587	Port used by the SMTP server if SSL is enabled.
	Gmail REST API ( <code>gmail.googleapis.com</code> )	TCP	443	Port used for communication with Google Mail services.
	Microsoft Graph REST API ( <code>graph.microsoft.com</code> , <code>login.microsoftonline.com</code> )	TCP	443	Port used for communication with Microsoft Exchange Online organizations.

# Installation and Configuration

You can install Veeam Agent for Microsoft Windows on any computer whose data you plan to protect — server, desktop, laptop or tablet.

# Before You Begin

Before you start the installation process, check the following prerequisites:

- The computer on which you plan to install Veeam Agent must satisfy system requirements specified in this document. To learn more, see [System Requirements](#).
- Make sure to install all available updates for the operating system on the computer you want to protect. Otherwise, the correct functioning of Veeam Agent is not guaranteed.
- If your computer runs Microsoft Windows 11 version 22H2 or 23H2, make sure that the Smart App Control feature is disabled. Otherwise, it may cause installation issues.

For information about Smart App Control, see [this Microsoft KB article](#).

- You must run the Veeam Agent setup file under the Administrator account or any user account that has Administrator privileges on the computer where you plan to install the product.
- Veeam Agent for Microsoft Windows requires the following components:
  - SQLite database engine
  - Microsoft .NET Framework 4.5.2 or later
  - Windows Universal C Runtime Library

If these components are not pre-installed on the computer, the setup will install them during the product installation process.

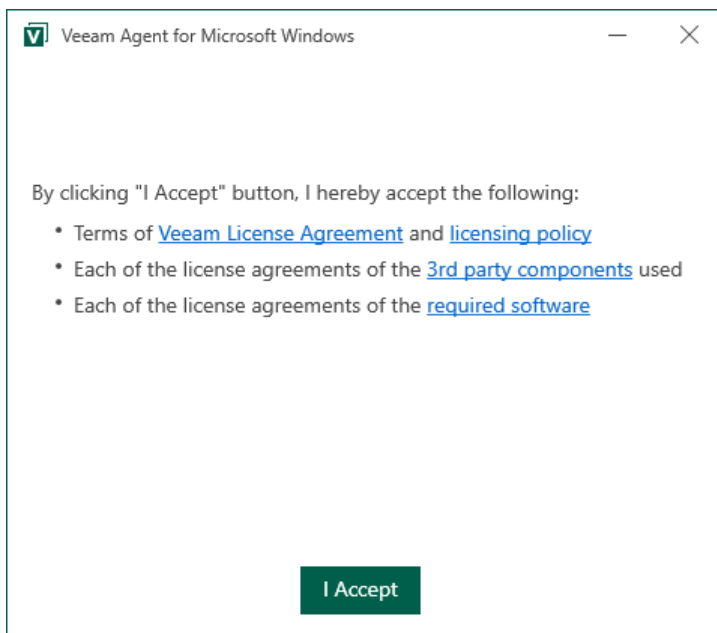
- The product program files are placed to the `%ProgramFiles%\Veeam\Endpoint Backup` folder on the system volume. Make sure that you have enough free space on the system volume to install the product. Veeam Agent requires at least 200 MB.
- [Recommended] If you want to create a recovery image of your computer on a USB storage device, CD/DVD/BD or make an ISO image, prepare the necessary device/media or make sure that you have enough free disk space in the target location. On average, the size of the created recovery image is 500 MB.

During the recovery image creation, Veeam Agent formats the removable storage device. If you have important information on the device, create a copy of this data in some other location.

# Installing Veeam Agent for Microsoft Windows

To install Veeam Agent for Microsoft Windows:

1. Download the setup archive for Standalone Veeam Agent for Microsoft Windows from [this Veeam webpage](#) and save it on the computer where you plan to install the product.
2. Double-click the downloaded setup archive. In the open archive, double-click the setup file to launch the installation wizard and click **Next**.
3. Click **I Accept** to accept license agreements and install Veeam Agent for Microsoft Windows.



## NOTE

In some cases, computer reboot is required to complete the installation.

After the installation process is complete, you can instruct Veeam Agent for Microsoft Windows to create a recovery image for your computer. To learn more, see [Creating Veeam Recovery Media](#).

# Installing Veeam Agent for Microsoft Windows in Unattended Mode

You can install Veeam Agent for Microsoft Windows in the unattended mode using the command line interface. The unattended installation mode does not require user interaction – the installation runs automatically in the background, and you do not have to respond to the installation wizard prompts. You can use the unattended installation mode to automate the Veeam Agent installation process in large-scale environments.

## Prerequisite Software

During the product installation, Veeam Agent for Microsoft Windows automatically sets up the following required prerequisite components:

- SQLite database engine

Veeam Agent for Microsoft Windows 6.2 uses SQLite database instead of Microsoft SQL Server 2012 Express LocalDB. If you install Veeam Agent 6.2 on a computer with an existing SQL Server database, migration to the new database will be performed during the upgrade of Veeam Agent. If necessary, you can skip migration to the new database. To learn more, see [Installation Without Database Migration](#).

- Windows Universal C Runtime Library

Veeam Agent for Microsoft Windows also sets up Microsoft .NET Framework 4.5.2 if it does not detect this component on the computer during the product installation.

In some cases, installation of prerequisite software requires computer reboot. This can happen, for example, if you have an earlier version of a prerequisite component installed on the computer and during the installation process this component is used by third-party software.

In this situation, unattended setup will install Veeam Agent for Microsoft Windows but will not start the Veeam Agent for Microsoft Windows service. After you reboot the computer, the Veeam Agent for Microsoft Windows service will be started and Veeam Agent for Microsoft Windows will be fully functioning.

## Installation

To install Veeam Agent for Microsoft Windows version 6.2, use a command with the following syntax:

```
<path_to_exe> /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware
```

where `<path_to_exe>` – path to the Veeam Agent for Microsoft Windows installation file.

Veeam Agent for Microsoft Windows uses the following codes to report about the installation results:

- 1000 – Veeam Agent for Microsoft Windows has been successfully installed.
- 1001 – prerequisite components required for Veeam Agent for Microsoft Windows have been installed on the machine. Veeam Agent for Microsoft Windows has not been installed. The machine needs to be rebooted.
- 1002 – Veeam Agent for Microsoft Windows installation has failed.
- 1004 – migration to SQLite database has failed.

- 1101 — Veeam Agent for Microsoft Windows has been installed. The machine needs to be rebooted.

## Installation Without Database Migration

If you do not want to migrate the existing SQL database during the installation of Veeam Agent, you can skip migration to the new database. In this case, you will need to create backup jobs from scratch.

To install Veeam Agent for Microsoft Windows version 6.2 and skip migration, use a command with the following syntax:

```
<path_to_exe> /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware /skipmigration
```

To learn more about database migration, see [Upgrading Veeam Agent for Microsoft Windows](#).

# Using Sysprep and Veeam Agent for Microsoft Windows

You can pre-install Veeam Agent in a custom Microsoft Windows system image that will be used for deployment on different computers. To do this, you should perform a set of configuration steps in the reference Microsoft Windows system installation that will be included in a deployment image.

To configure a custom Microsoft Windows system image with Veeam Agent:

1. Install Veeam Agent in a Microsoft Windows system image. To learn more, see [Installing Veeam Agent for Microsoft Windows](#).
2. Configure the backup job in the way you want it to work on computers with pre-installed Veeam Agent. To learn more, see [Creating Backup Jobs](#).

## NOTE

It is advised to configure the backup job for the entire computer backup. In case of volume-level backup, it may be necessary to reconfigure the backup job after Microsoft Windows is deployed to the target computer and include the necessary volumes in the backup once again. This may happen if volumes' GUIDs were changed at the stage of Microsoft Windows generalization with Sysprep.

3. Create a registry value: `HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Endpoint Backup\SysprepMode (DWORD)=1`.

This registry value is used to regenerate the job ID when Veeam Agent starts for the first time on the new computer. If you do not create the registry value, the backup job may fail as soon as it is started on the new computer.

4. Run the Sysprep tool in the Generalize mode to remove any system-specific data. If you need to run the Sysprep tool in the Audit mode, do not forget to re-create the registry value afterwards.
5. Deploy the image on the necessary computers in any convenient way. To learn more about deployment of Microsoft Windows system to new computers, see [Microsoft documentation](#).

When you deploy the created image on the computer, Veeam Agent will re-generate internal IDs of the backup jobs. As a result, the backup job will be fully functional.

## NOTE

After Veeam Agent for Microsoft Windows starts on the new computer for the first time, the registry value is changed to 0 so that the job ID is not changed during subsequent starts of Veeam Agent for Microsoft Windows.



# Upgrading Veeam Agent for Microsoft Windows

For Veeam Agent for Microsoft Windows, upgrade to newer versions is supported. You can start the upgrade process from the Veeam Agent control panel when the new version becomes available. To learn how to check for product updates, see [Checking for New Product Versions and Updates](#).

## NOTE

Upgrade to version 6.2 is supported for Veeam Agent for Microsoft Windows version 4.0 and later.

## Before You Begin

Before you upgrade Veeam Agent for Microsoft Windows to version 6.2, check the following prerequisites:

- Make sure that you do not have backup jobs targeted at Microsoft OneDrive.  
Starting from Veeam Agent for Microsoft Windows 6.0, backup to Microsoft OneDrive is deprecated. To continue working with backups created with an earlier version of Veeam Agent and stored in a Microsoft OneDrive storage, download these backups and move them to another storage. To learn more, see [Downloading Backups from Microsoft OneDrive](#).
- [For upgrade of Veeam Agent for Microsoft Windows 4.0 and 5.0] Make sure that your local Microsoft SQL Server instance and other SQL components that were installed with the previous version of Veeam Agent work correctly. Veeam Agent for Microsoft Windows 6.2 uses SQLite database instead of Microsoft SQL Server 2012 Express LocalDB. Migration to the new database will be performed during the upgrade of Veeam Agent. If any required SQL components are missing, migration will fail.
- Make sure that there are no running jobs.  
If you have any running jobs, let them complete successfully before you start the upgrade. Disable any periodic jobs temporarily to prevent them from starting during the upgrade. If the protected computer runs VSS-aware applications and backup of database logs (Microsoft SQL Server transaction logs or Oracle archived logs) is enabled in the backup job for the computer, disable this backup job too.

During the upgrade process, configuration and backup files that were created with the previous version of Veeam Agent are not impacted in any way.

## Upgrade Process

### NOTE

Consider the following:

- In some cases, upgrade to the new version of Veeam Agent may require computer reboot.
- You can also download the Veeam Agent setup archive from [this Veeam webpage](#). Save the downloaded archive on the computer where you plan to install the new version of the product and double-click the setup archive to start the upgrade.

To upgrade Veeam Agent for Microsoft Windows:

1. Double-click the Veeam Agent icon in the system tray, or right-click the Veeam Agent icon in the system tray and select **Control Panel**.

2. Open the **About** tab.
3. If the new version of Veeam Agent is available, click **Download**.
4. When the download is complete, click **Install** to run the setup archive.
5. To upgrade Veeam Agent, you must accept the terms of the license agreements. Read the license agreements and click **Accept and Update**.
6. To preserve settings from the previous product installation, click **Yes** in the opened window.

If the migration fails for any reason, the upgrade process will stop. In this case, you will be able to continue using the earlier version of Veeam Agent as the existing Microsoft SQL Server database and preinstalled SQL components are not deleted during the migration.

#### NOTE

Migration will not start if at least one of the following conditions is met:

- You use Veeam Agent version 3.0 or earlier. In this case, a new database will be created during the upgrade.
- Migration to SQLite database was performed earlier.

7. After the installation is complete, click **Finish**.

## Unattended Upgrade

You can upgrade Veeam Agent to a newer version in the unattended mode using the same command that is used for unattended installation. To learn more, see [Installing Veeam Agent for Microsoft Windows in Unattended Mode](#).

If necessary, you can skip the migration of the SQL database during the unattended upgrade. To learn more, see [Installation Without Database Migration](#).

# Configuring Advanced Settings

## Connection Settings for Veeam Backup Server

If you want to connect Veeam Agent computer to Veeam backup server as a member of the protection group for pre-installed Veeam Agents, you must apply connection settings from the configuration file. The configuration file is one of the Veeam Agent for Microsoft Windows setup files that you must obtain from your System Administrator. To learn more about protection group for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

To connect Veeam Agent for Microsoft Windows to Veeam backup server, use Veeam Agent Configurator. To learn more, see the [SetVBRSettings](#) section in the Veeam Agent Configurator Reference.

# Uninstalling Veeam Agent for Microsoft Windows

To uninstall Veeam Agent for Microsoft Windows:

1. From the **Start** menu, select **Control Panel > Programs and Features**.
2. In the programs list, right-click Veeam Agent and select **Uninstall**. Wait for the process to complete.

The prerequisite components installed and used by Veeam Agent are not removed during the uninstall process. To remove each of the remaining components, right-click it in the programs list and select **Uninstall**.

# What You Do Next

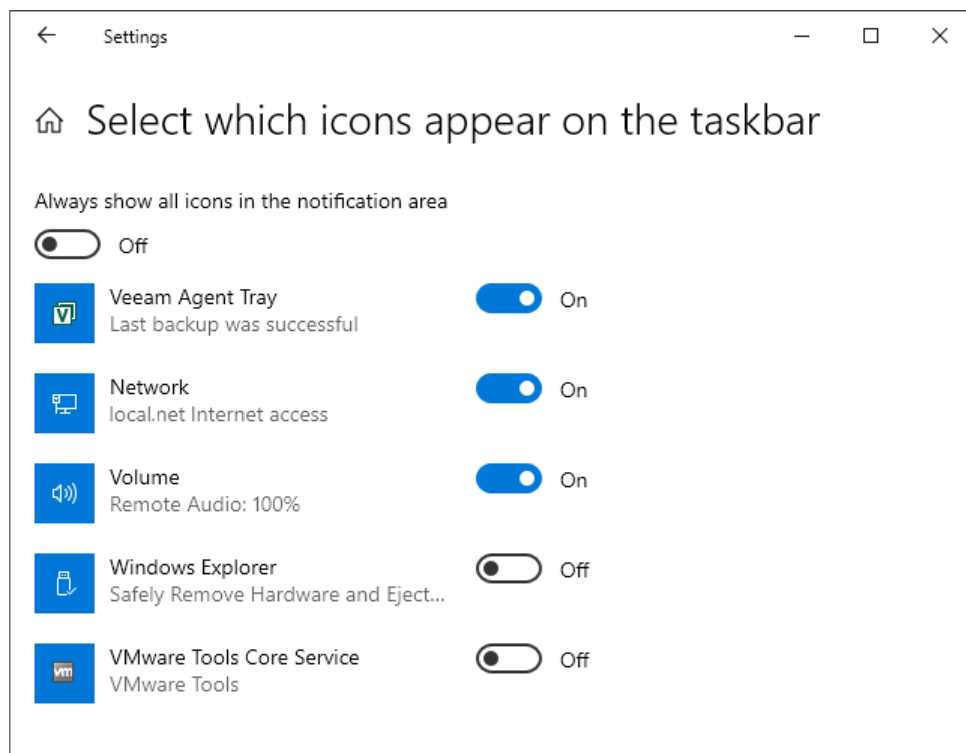
After the product installation, Veeam Agent for Microsoft Windows displays its icon in the system tray. You can use the system tray icon to perform main operations in Veeam Agent for Microsoft Windows:

- Configure the backup job and start ad-hoc backup operations
- Launch restore wizards
- Open the Veeam Agent control panel
- Monitor the state of backup tasks and so on

Depending on the current settings of your Microsoft Windows OS, the Veeam Agent for Microsoft Windows icon may not be displayed in the system tray.

To bring the icon to the system tray:

1. In Microsoft Windows, open the **Notification Area Settings** window. To do this, do either of the following:
  - Click the arrow in the system tray and click **Taskbar Settings**.
  - From the Microsoft Windows main menu, select **Control Panel** and navigate to **Taskbar and Navigation**. In the **Taskbar** section, select **Notification Area**.
2. In the **Notification Area Settings** window, select **On** for the Veeam Agent Tray.
3. In the **Behaviors** column, set the **Show icon and notification** setting for it.



# Getting Started

To protect your computer from a disaster of any kind, you must perform the following operations in Veeam Agent for Microsoft Windows:

1. Create a Veeam Recovery Media.

The Veeam Recovery Media provides an alternate way to boot the Microsoft Windows RE. If your computer fails to start or the hard disk gets corrupted, you can boot the Windows RE from the Veeam Recovery Media and restore your data.

To learn more, see [Creating Veeam Recovery Media](#).

2. Define what data you want to back up and configure the backup job.

Before you configure a backup job, you should decide on the following backup details:

- Backup scope: entire computer image, individual computer volumes or specific computer folders.
- Backup destination: where you want to store created backups.
- Backup schedule: how often you want to back up your data.

After that, you can configure the backup job. The backup job runs automatically by the defined schedule, captures the data that you have added to the backup scope and creates a chain of restore points in the target location. If your data gets lost or corrupted, you can restore it from the required restore point.

To learn more, see [Creating Backup Jobs](#).

3. Specify Veeam Agent settings.

You can define resource usage settings during backup, instruct Veeam Agent to automatically check for new product versions and so on. To learn more, see [Specifying Settings](#).

4. Monitor backup task performance.

You can use the Veeam Agent control panel to check how backup tasks are being performed, what errors have occurred during backup job sessions and so on. To learn more, see [Reporting](#).

5. In case of a disaster, you can restore the entire computer image or specific data on the computer. To learn more, see [Performing Restore](#).

# Licensing

You can use Veeam Agent for Microsoft Windows as a free product. In this case, you do not need to obtain and install any license.

To work with a commercial version of Veeam Agent for Microsoft Windows, you must obtain a license and install it on the protected computer. If you do not install a license, you will be able to use the Free edition of the product only.

If you plan to use a commercial version of the product with Veeam Backup & Replication, you must install and manage the license in the Veeam Backup & Replication console or in Veeam Backup Enterprise Manager. To learn more, see [Managing License](#).

# Product Editions

Veeam Agent for Microsoft Windows offers three product editions that define product functionality and operation modes:

- *Server* – a commercial edition that provides access to all product functions. The Server edition is intended for performing data protection tasks on servers that run the Microsoft Windows OS. To use the Server edition of Veeam Agent for Microsoft Windows, you must obtain and install a license on the protected computer. The license must have a number of instances that is enough to protect a machine with the Server product edition.
- *Workstation* – a commercial edition that offers capabilities for performing data protection tasks on desktop computers and laptops that run the Microsoft Windows OS. To use the Workstation edition of Veeam Agent for Microsoft Windows, you must obtain and install a license on the protected computer. The license must have a number of instances that is enough to protect a machine with the Workstation product edition.
- *Free* – a free edition that offers limited capabilities. In contrast to the Workstation and Server editions, the Free edition does not require a license.

For more information about product editions, pricing and features available for them, see [this Veeam webpage](#).

To learn more about instance licensing, see [this Veeam webpage](#).

When you install a license on the protected computer, you can select the product edition of Veeam Agent for Microsoft Windows: Workstation or Server (if both editions are supported by the license). To learn more, see [Selecting Product Edition](#).

If you use Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you should manage product licenses and editions from the Veeam Backup & Replication console. To learn more, see [Managing License](#).

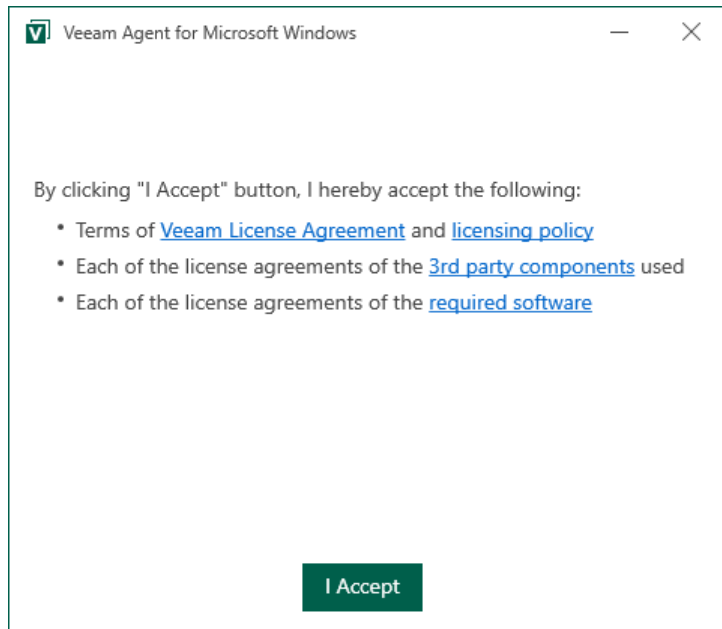
If the license has expired and you have used backup job options available for the Workstation and Server editions, you must disable these options in the properties of the backup job. Otherwise, the backup job will fail.



# License Agreements

When you install Veeam Agent for Microsoft Windows, you must accept the terms of license agreements.

To view the license agreements, click the corresponding links in the installation window.



# Installing License

When you launch the Veeam Agent control panel for the first time, Veeam Agent displays a notification window offering to install a license. You can choose to install the license immediately or postpone this operation.

- If you choose to install the license, you can immediately browse for the license key on your computer and complete the license installation process.
- If you choose to postpone the license installation process, you will be able to install a license later at any time you need.

Until you install a license, you can use the Free edition of the product. To switch to a commercial version of Veeam Agent for Microsoft Windows, you need to obtain and install a license.

## NOTE

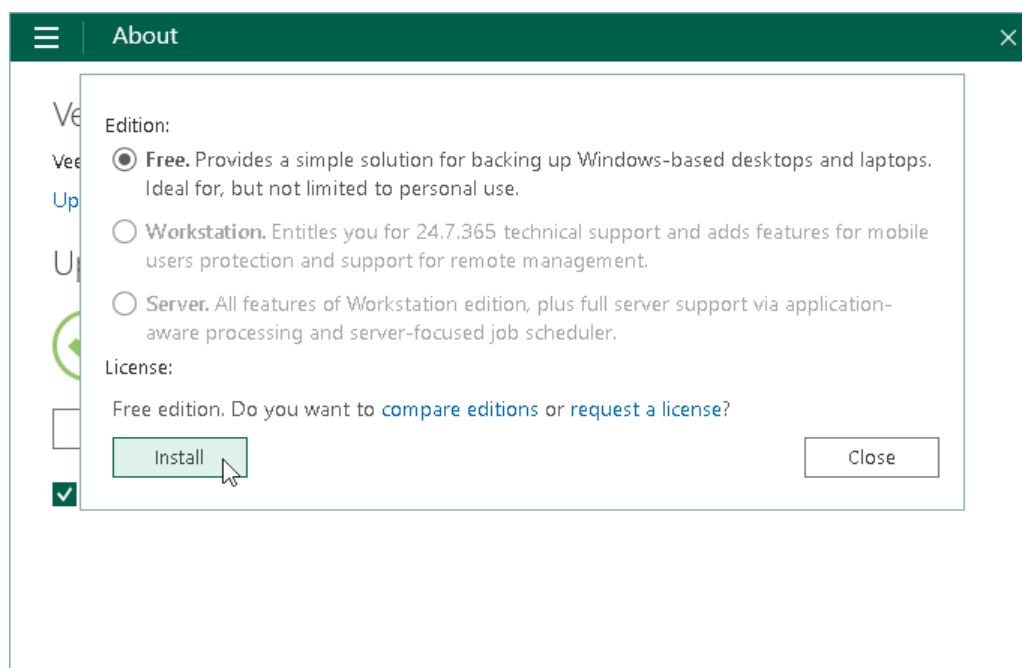
If you plan to use a Veeam Backup & Replication repository as a target location for Veeam Agent backups, you must install a license in Veeam Backup & Replication. The license must have enough instances to protect machines with Veeam Agents that back up data to the Veeam Backup & Replication repository. To learn more, see [Managing License](#).

To install a license:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **About**.
3. In the **Version** section, click the following link:
  - a. **Update license to get additional features** – if the license is not installed yet, and you run the Free edition of Veeam Agent for Microsoft Windows.
  - b. **Manage license and edition** – if the license is already installed on the Veeam Agent computer, and you want to change the license or select the product edition.
4. In the dialog window, click **Install** and browse for the LIC file.

Veeam Agent for Microsoft Windows will install the license and select the product edition that is allowed by the license. If a license supports both the Workstation and Server editions, Veeam Agent for Microsoft Windows will select the product edition based on the type of the Microsoft Windows OS installed on the Veeam Agent computer.

You can change the product edition manually if needed. To learn more, see [Selecting Product Edition](#).



# Selecting Product Edition

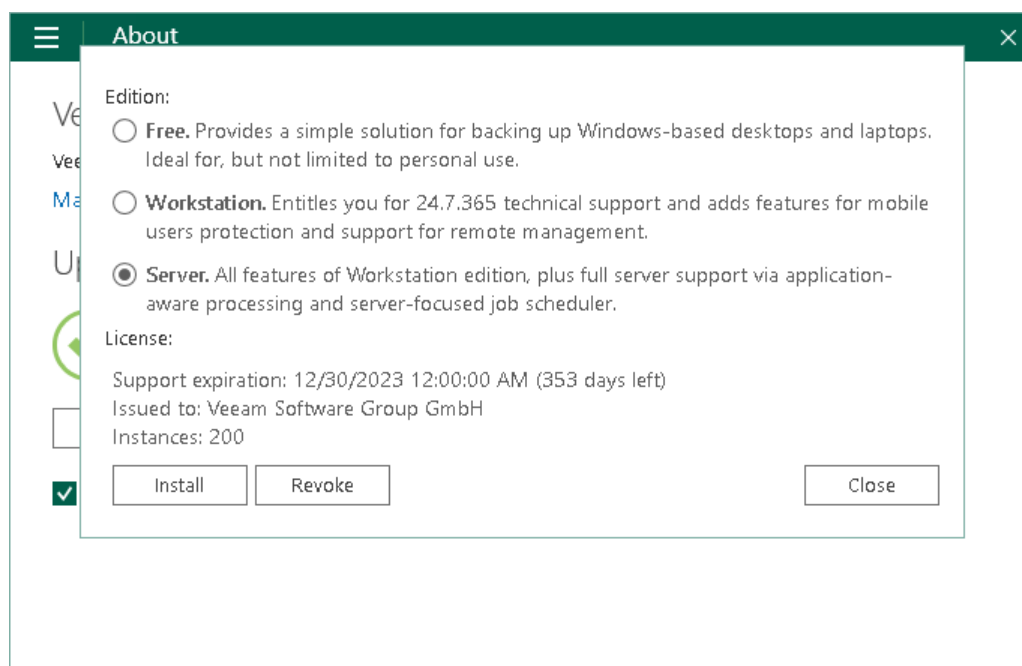
When you install a license, Veeam Agent for Microsoft Windows automatically selects the product edition that is allowed by the license. If a license supports both the Workstation and Server editions, Veeam Agent for Microsoft Windows will select the product edition based on the type of the Microsoft Windows OS installed on the Veeam Agent computer.

You can change the product edition manually if needed. To select the product edition:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **About**.
3. In the **Version** section, click **Manage license and edition**.
4. In the dialog window, in the **Edition** section, select the desired product edition. To learn more about editions of Veeam Agent for Microsoft Windows, see [Product Editions](#).

## NOTE

After you switch from the Server edition to the Workstation edition, or vice versa, Veeam Agent for Microsoft Windows will disable the backup job. This operation is required, because backup retention policies and available backup job options differ in Workstation and Server editions. To enable the job, you must edit the backup job settings in accordance with the selected edition.

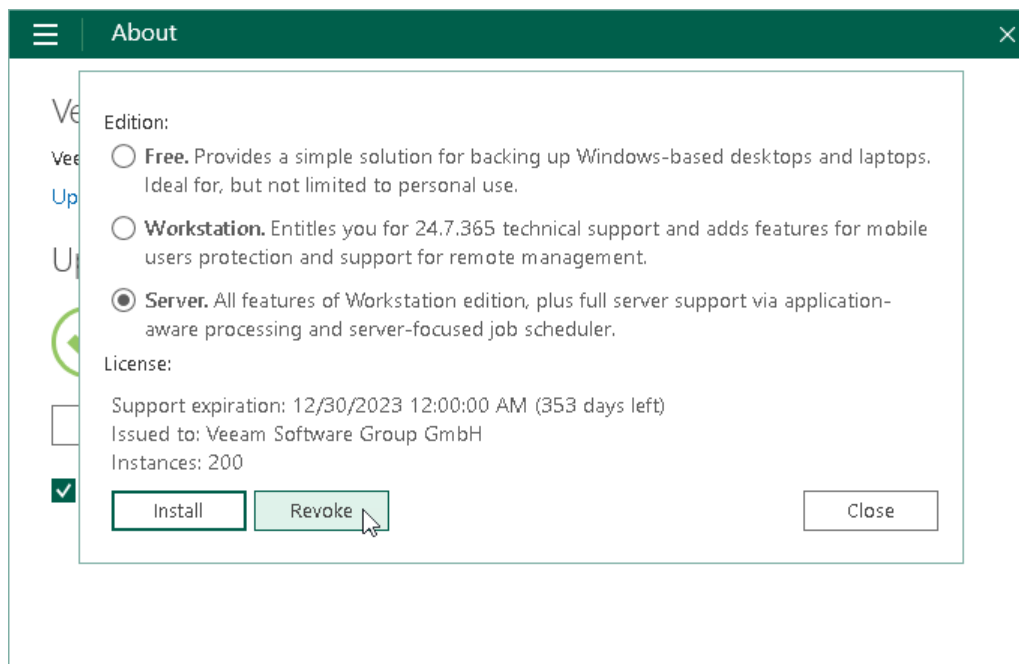


# Revoking License

You can revoke a license at any time if needed, for example, after the license is expired, and you want to continue using the Free edition of Veeam Agent for Microsoft Windows.

To revoke a license:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **About**.
3. In the **Version** section, click **Manage license and edition**.
4. In the dialog window, click **Revoke**.



# Performing Backup

To protect your computer and data, you can perform the following operations:

- Create a Veeam Recovery Media. You can use the Veeam Recovery Media to boot the Microsoft Windows RE from the recovery media in case the OS on your computer fails to start. To learn more, see [Creating Veeam Recovery Media](#) and [Creating Veeam Recovery Media with Command Line Interface](#).
- Back up your data. You can back up your data to protect the entire computer image, individual volumes or folders on your computer. Veeam Agent lets you configure a scheduled backup job with the default settings right after the product installation, configure a backup job with custom settings or create ad-hoc backups at any time you need. You can use data backup to restore necessary information if data on your computer gets corrupted or you delete some files and folders by mistake. To learn more, see [Creating Backup Jobs](#).

# Creating Veeam Recovery Media

You can create a Veeam Recovery Media — a recovery media for your computer. The Veeam Recovery Media contains all data that is required to run the Microsoft Windows RE. If your computer stops working or the hard disk fails, you can boot from the Veeam Recovery Media, instead of booting from the hard drive. After booting, you can use Veeam and Microsoft tools to fix errors, recover the system image of your computer and your data.

## NOTE

In some cases, Windows Recovery Environment components may be missing on the system, and Veeam Agent for Microsoft Windows will not find them during the Veeam Recovery Media creation. In such case you will be prompted to do one of the following:

- Insert the Windows Installation Media so that Veeam Agent for Microsoft Windows can load the necessary components from it.
- Download and install Windows Assessment and Deployment Kit (MS ADK).

Keep in mind that the Veeam Recovery Media created with MS ADK components will not contain the following tools:

- Windows Recovery Environment
- Memory Diagnostic
- Startup Repair

# Before You Begin

Before you create a Veeam Recovery Media, check the following prerequisites:

## Removable Storage Device Scenario (USB, SD Card and Other)

- The removable storage device must be inserted into a corresponding slot on the computer or connected to the computer.
- The removable storage device must have enough capacity to store the created recovery image. On average, the size of the created recovery image without manually loaded drivers is 500 MB.
- During the recovery image creation, Veeam Agent for Microsoft Windows formats the removable storage device. If you have important information on the device, create a copy of this data in some other location.
- If you want to include specific storage and network drivers into the recovery image, place them to a local folder on your computer or in a network shared folder to which you have access and read permissions. During the recovery image creation process, you will be able to define a path to this folder, and Veeam Agent for Microsoft Windows will include the drivers into the recovery image.
- [For Microsoft Windows 2008 R2 and later] If you want your computer to detect a Wi-Fi network and connect to it after you boot from the recovery image, enable the Wireless LAN Service feature on your computer. In this case, Veeam Agent for Microsoft Windows will add wireless networking support files to the Veeam Recovery Media. To learn more about the Wireless LAN Service, see [Microsoft documentation](#).

## CD/DVD/BD Scenario

- An empty or re-writable CD/DVD/BD must be inserted into a CD/DVD/BD drive on the computer.
- The CD/DVD/BD must have enough capacity to store the created recovery image. On average, the size of the created recovery image without manually loaded drivers is 500 MB.
- If you want to include specific storage and network drivers into the recovery image, place them to a local folder on your computer or in a network shared folder on which you have read permissions. During the recovery image creation, you will be able to define a path to this folder, and Veeam Agent for Microsoft Windows will include the drivers into the recovery image.
- [For RW CD/DVD/BD] During the recovery image creation, Veeam Agent for Microsoft Windows erases information on the CD/DVD/BD. If you have important information on the CD/DVD/BD, create a copy of this data in some other location.
- [For Microsoft Windows 2008 R2 and later] If you want your computer to detect a Wi-Fi network and connect to it after you boot from the recovery image, enable the Wireless LAN Service feature on your computer. In this case, Veeam Agent for Microsoft Windows will add wireless networking support files to the Veeam Recovery Media. To learn more about the Wireless LAN Service, see [Microsoft documentation](#).

## Local Target and Shared Folder Scenario (ISO)

- If you want to include specific storage and network drivers into the recovery image, place them to a local folder on your computer or in a network shared folder on which you have read permissions. During the recovery image creation, you will be able to define a path to this folder, and Veeam Agent for Microsoft Windows will include the drivers into the recovery image.
- [For shared folders] If you plan to save the created ISO file in a network shared folder, make sure that you have access to this folder and write permissions on it.



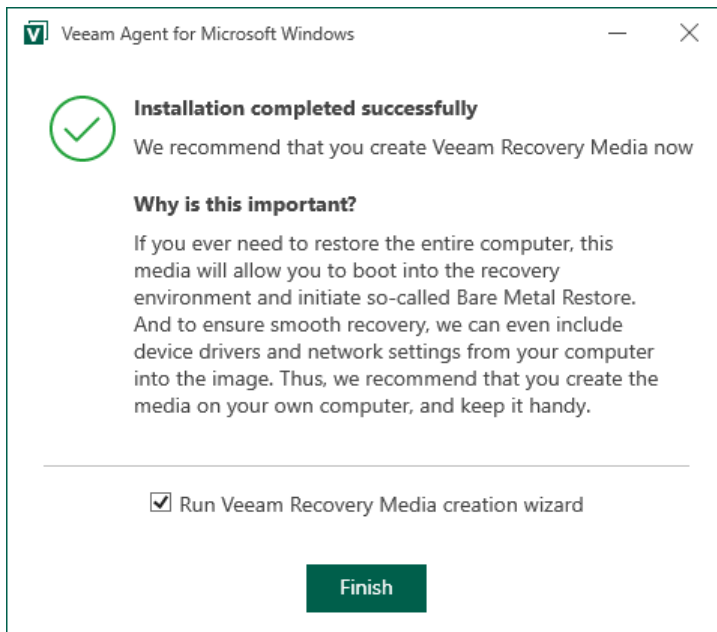
- [For Microsoft Windows 2008 R2 and later] If you want your computer to detect a Wi-Fi network and connect to it after you boot from the recovery image, enable the Wireless LAN Service feature on your computer. In this case, Veeam Agent for Microsoft Windows will add wireless networking support files to the Veeam Recovery Media. To learn more about the Wireless LAN Service, see [Microsoft documentation](#).

# Step 1. Launch Create Recovery Media Wizard

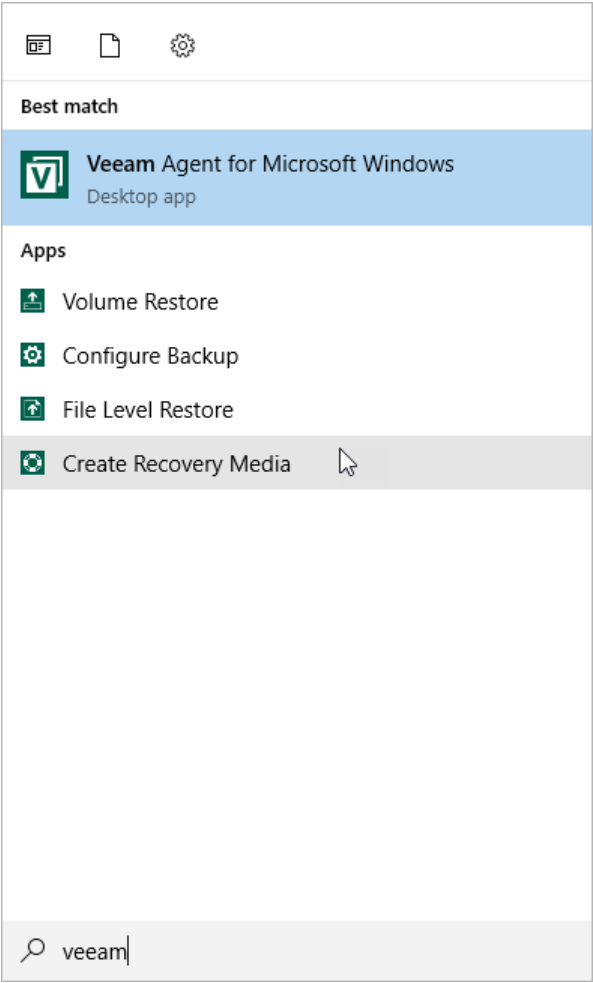
You can launch the **Create Recovery Media** wizard right after the product installation process or at any time later.

To launch the **Create Recovery Media** wizard after installation:

1. At the last step of the installation wizard, select the **Run Veeam Recovery Media creation wizard** check box.
2. Click **Finish**. Veeam Agent for Microsoft Windows will automatically launch the **Create Recovery Media** wizard.



To launch the **Create Recovery Media** wizard at any time, from the Microsoft Windows Start menu, select **All Programs > Veeam > Tools > Create Recovery Media** or use the Microsoft Windows search to find the **Create Recovery Media** option on your computer.



## Step 2. Specify Recovery Media Options

At the **Recovery Media** step of the wizard, specify on which type of media you want to create a recovery image and what drivers you want to include in the recovery image.

1. In the **Available bootable media types** list, select a media for the recovery image. You can create the following types of recovery images:
  - Recovery image on a removable storage device. You can create a recovery image on a USB drive, SD card and so on. Veeam Agent for Microsoft Windows displays all removable storage devices currently attached to your computer. Select the necessary one in the list.
  - Recovery image on an optical disk. You can create a recovery image on a CD, DVD or BD. Veeam Agent for Microsoft Windows displays all CD, DVD and BD drives available on your computer. Select the necessary one in the list.
  - ISO file with the recovery image. You can create a recovery image in the ISO file format and save the resulting file locally on your computer or in a network shared folder.
2. If you have enabled data encryption options for the backup job and want to include the decryption key in the recovery image, select the **Include decryption key for seamless restore from encrypted backup** check box. In this case, when you use the created recovery image to perform bare metal recovery, you will not have to enter the password used for encryption.
3. If you want to include in the recovery image current network settings, make sure that the **Include network connections settings from this computer** check box is selected. When you use the created Veeam Recovery Media to boot your computer, these settings will be automatically applied and will be used to connect to the remote backup storage.
4. If you want to include in the recovery image storage and network drivers that are currently installed on your computer, make sure that the **Include hardware drivers from this computer** check box is selected. Veeam Agent for Microsoft Windows will detect hard disk controller drivers, network adapter drivers and USB controller drivers and include them into the recovery image. When you use the created Veeam Recovery Media to boot your computer, these drivers will be automatically injected into Windows RE.
5. If you want to include in the recovery image additional storage and network drivers that you may need when booting from the recovery image, select the **Include the following additional storage and network hardware drivers** check box, click **Add** and select a folder containing necessary drivers. The folder that you select must contain all files of the driver package (files in CAT, INF and SYS formats).

We strongly recommend that you enable this option if you use drivers that are not included into the Microsoft Windows installation DVD. For example, you can include drivers for a discrete network card, third party USB 3.0 controllers and non-standard hard disk controllers.

## IMPORTANT

Consider the following:

- When you boot your computer from the Veeam Recovery Media, Veeam Agent for Microsoft Windows does not automatically install additional drivers included in the recovery image. You need to install such drivers manually using the **Load Driver** tool. To learn more, see [Using Veeam Agent and Microsoft Windows Tools](#).
- We do not recommend to include large amount of additional drivers (1 GB and more) in the Veeam Recovery Media. When you boot your computer from the Veeam Recovery Media, Veeam Agent for Microsoft Windows loads all additional drivers stored in the Veeam Recovery Media into your computer RAM. If the total size of the recovery environment is approximately equal to or greater than the amount of RAM, Windows RE will fail to load.

**Create Recovery Media**

Recovery Media  
Specify bare metal recovery media options.

**Recovery Media**

Image Path  
Ready to Apply  
Progress

Available bootable media types:

Name	Type	Capacity
E:\	LOCAL DISK (E:)	414.6 GB
Image	ISO image file	

☒ Include decryption key for seamless restore from encrypted backup (protects from password loss)

☒ Include network connections settings from this computer (recommended)

☒ Include hardware drivers from this computer (recommended)

☒ Include the following additional storage and network hardware drivers:

Folder
C:\Drivers

Add...  
Remove

< Previous   **Next >**   Finish   Cancel

## Step 3. Specify Path to ISO

The **Image Path** step of the wizard is available if you have selected to create an ISO file with the recovery image.

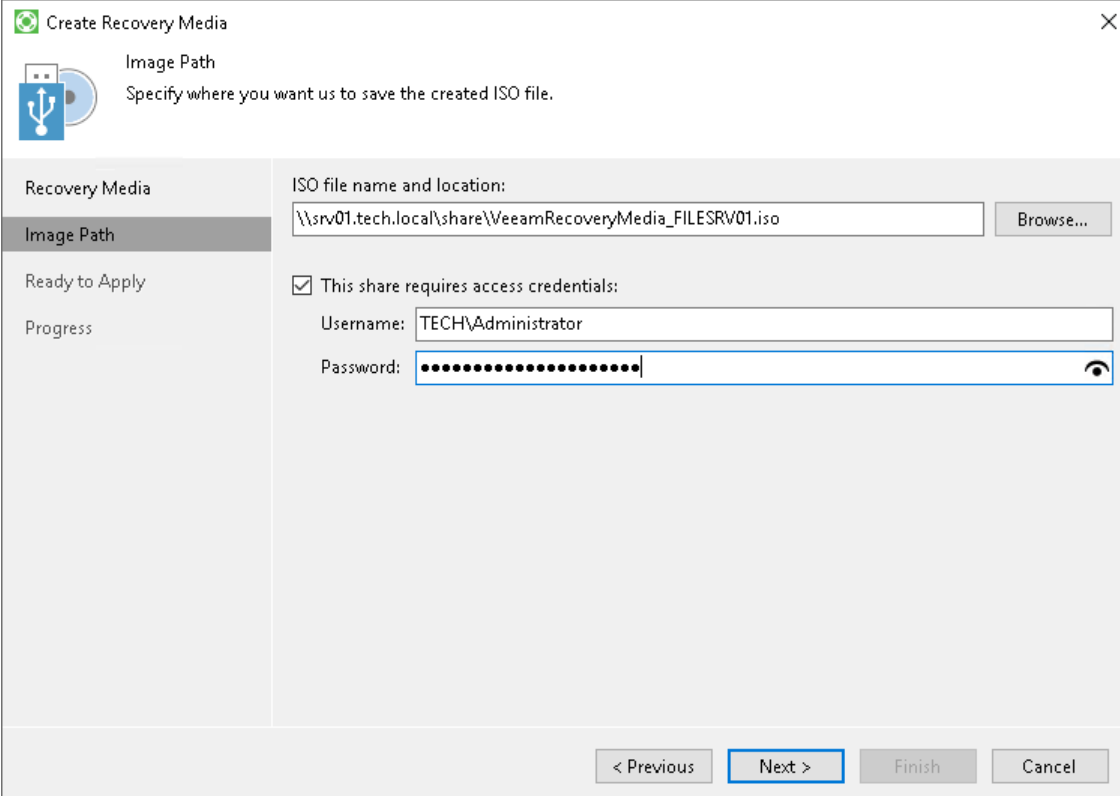
Select a location where you want to save the ISO file.

1. In the **ISO file name and location** field, specify a real path to the folder where you want to save the created recovery image and the ISO file name. You can save the ISO file in the following locations:
  - Local folder: select the necessary folder on your computer.
  - Network shared folder: specify a UNC path to a network shared folder. Keep in mind that a UNC path always starts with two back slashes (\\).

We strongly recommend that you store the recovery image in a location other than a local computer drive. If you choose to save the recovery image in a local folder on your computer, you can copy it to an external location afterwards. In this case, the recovery image will always be available should computer volumes get corrupted or the computer fail to start.

2. If you chose to save the ISO file in a network shared folder and this folder requires authentication, select the **This share requires access credentials** check box and enter the user name and password in the **Username** and **Password** fields. The user name must be specified in the *DOMAIN\UserName* format.

To view the entered password, click and hold the eye icon on the right of the **Password** field.

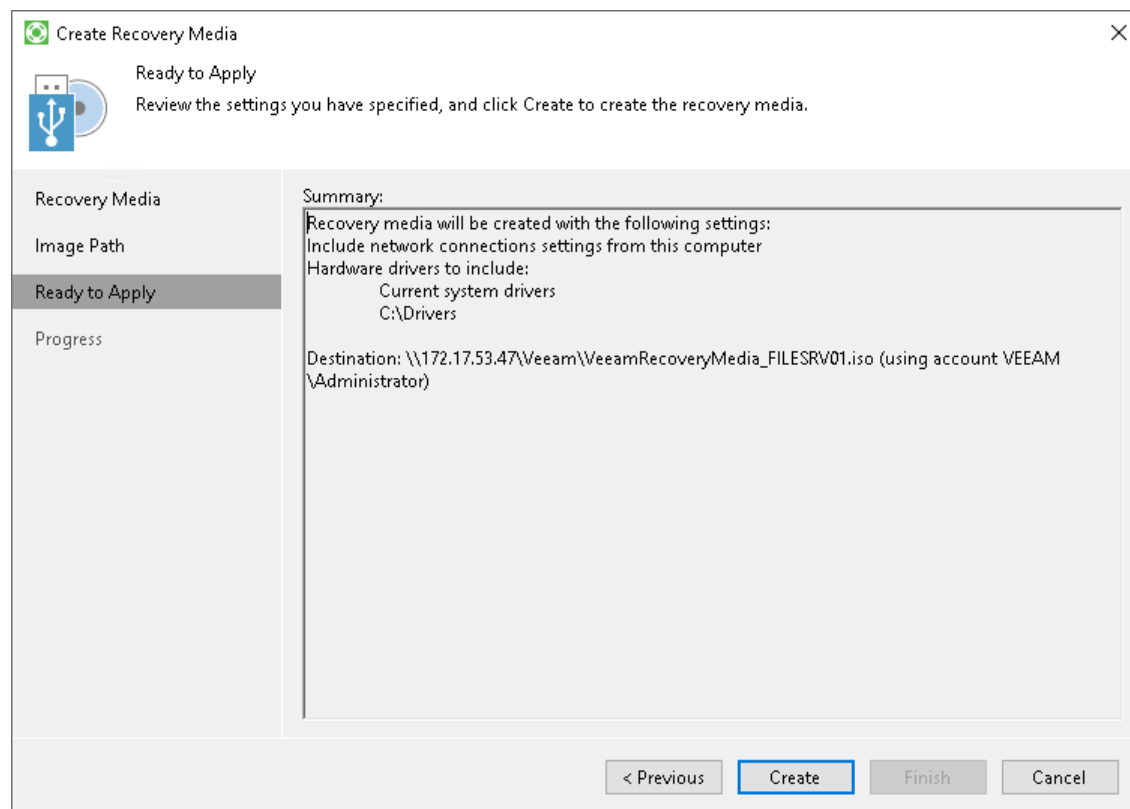


The screenshot shows the 'Create Recovery Media' wizard window. The title bar says 'Create Recovery Media' with a close button. The left sidebar has four items: 'Recovery Media', 'Image Path' (selected), 'Ready to Apply', and 'Progress'. The main area is titled 'Image Path' with the instruction 'Specify where you want us to save the created ISO file.' Below this, there is a section 'ISO file name and location:' with a text box containing '\\srv01.tech.local\share\VeeamRecoveryMedia\_FILESrv01.iso' and a 'Browse...' button. Below that is a checkbox labeled 'This share requires access credentials:' which is checked. Under the checkbox are two text boxes: 'Username:' with 'TECH\Administrator' and 'Password:' with a masked password '.....'. To the right of the password box is an eye icon. At the bottom of the window are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 4. Review Recovery Image Settings

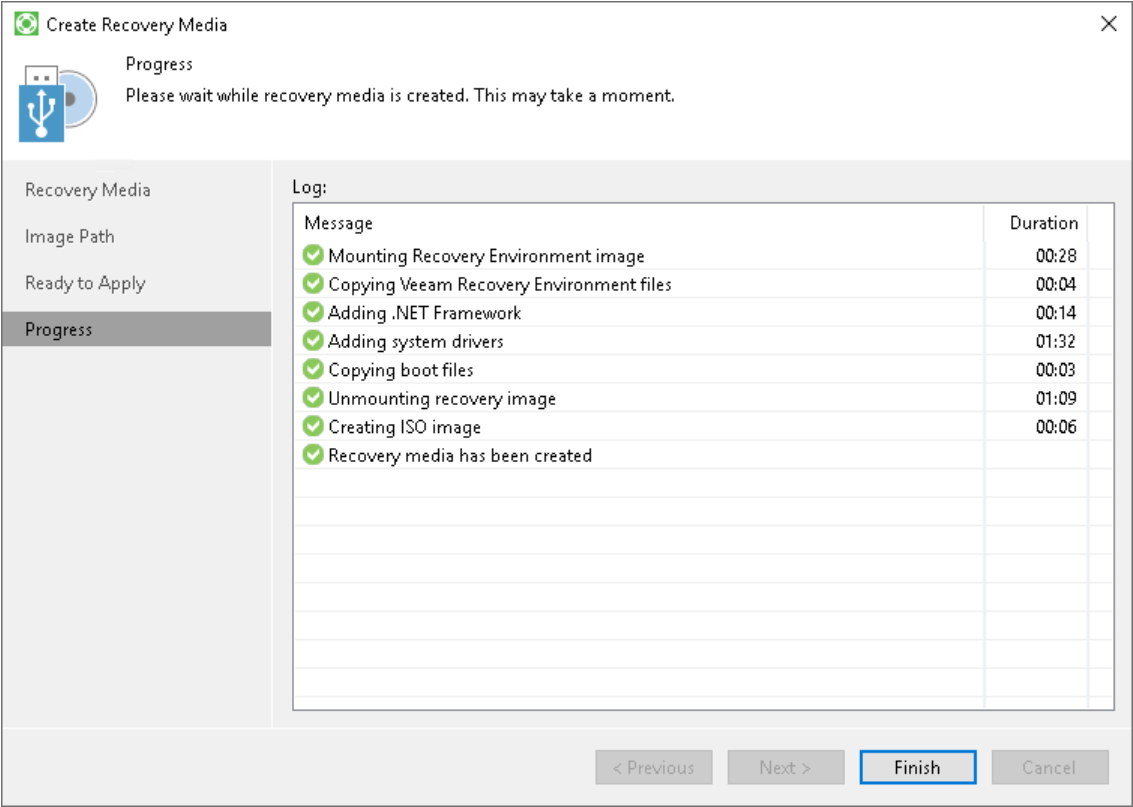
At the **Ready to Apply** step of the wizard, review settings of the recovery image that you plan to create and click **Create**.

Veeam Agent for Microsoft Windows will collect files necessary for recovery image creation and write the resulting recovery image to the specified target or burn it to CD/DVD/BD.



The process of recovery image creation may take some time. Wait for the process to complete and click **Finish** to exit the wizard.

If you want to interrupt the process of recovery image creation, click **Cancel** or close the wizard window.





# What You Do Next

[For ISO] After the recovery image is created, you can burn the created ISO to a CD/DVD/BD. To do this, you can use native Microsoft Windows tools or third-party software.

# Creating Veeam Recovery Media with Command Line Interface

In addition to creating the Veeam Recovery Media with the **Create Recovery Media** wizard, you can also use the command line interface to create the recovery image and save it as an ISO file. To create the Veeam Recovery Media, use a command with the following syntax:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /createrecoverymediaiso /f:<location>\<name>.iso
```

where:

- <location> – path to the folder in which the recovery image ISO file will be created. All folders that are mentioned in the path must be created in advance. Names of folders in the path must not contain spaces.
- <name> – name of the resulted ISO file. The file name must not contain spaces.

## NOTE

Consider the following:

- When you create the Veeam Recovery Media using the command line interface, Veeam Agent for Microsoft Windows includes in the recovery image storage and network drivers that are currently installed on the Veeam Agent computer, and current network settings of this computer. If you want to include additional drivers or data decryption key in the recovery image, you must create the Veeam Recovery Media using the **Create Recovery Media** wizard. To learn more, see [Creating Veeam Recovery Media](#).
- If you do not use the /f: option to specify a path to the folder in which to save the resulting ISO file, Veeam Agent for Microsoft Windows will save the ISO file to the C:\Users\%UserProfile%\Documents folder with the VeeamRecoveryMedia\_<machine>.iso name, where <machine> – name of your Veeam Agent computer.
- You cannot save the ISO file to a network shared folder using the command line interface.

You can use the last exit code to verify if the backup job has completed successfully. To check the last exit code, use the %ERRORLEVEL% variable in cmd.exe. Veeam Agent can provide the following exit codes:

- 0 – recovery image successfully created
- 1 – the create recovery image operation failed

## Example of Use

The example below creates an ISO file with the VeeamRE name in the E:\Veeam\RecoveryImage folder.

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /createrecoverymediaiso /f:E:\Veeam\RecoveryImage\VeeamRE.iso
```

# Creating Backup Jobs

To back up your data, you must configure a backup job. The backup job defines how, where and when to back up data. You can choose one of the following backup types:

- Backup of an entire computer image
- Backup of specific computer volumes, for example, a system volume or secondary volume
- Backup of individual folders, for example, documents folder or folder with music

# Before You Begin

Before you configure a backup job, consider the following:

- The target location where you plan to store backup files must have enough free space. The required space depends on the size of backed-up data and backup job settings.
- Available backup job options depend on the edition of Veeam Agent for Microsoft Windows. Make sure that you have obtained and installed a license that has the desired number of instances. To learn more, see [Licensing](#).
- [For Veeam backup repository] You can store created backups in a backup repository only if the backup server runs Veeam Backup & Replication 12.2. If you plan to use a commercial version of Veeam Agent for Microsoft Windows with Veeam Backup & Replication, you must install a license in Veeam Backup & Replication that has a sufficient number of instances to protect machines with Veeam Agent. To learn more, see [Managing License](#).
- [For Veeam backup repository] If you plan to use a Veeam backup repository as a target for backups, you must preconfigure user access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported. To learn more about KMS encryption for Veeam backup repositories, see the [Key Management System Keys](#) section in the Veeam Backup & Replication User Guide.
- [For Veeam Cloud Connect repository] If you plan to use a cloud repository as a target for backups, make sure that the service provider has communicated to you the necessary data: cloud gateway settings, user account settings and certificate thumbprint.
- A user account under which you launch the **New Backup Job** wizard must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

Backup has the following limitations:

- You cannot save the backup of entire computer on the local computer disk. Use an external hard drive or USB drive, network shared folder or backup repository as a target location.
- Veeam Agent for Microsoft Windows does not back up data to which symbolic links are targeted. It only backs up the path information that the symbolic links contain. After restore, identical symbolic links are created in the restore destination.
- Keep in mind that Veeam Agent stops running the backup job after 21 days (504 hours).

# Step 1. Launch New Backup Job Wizard

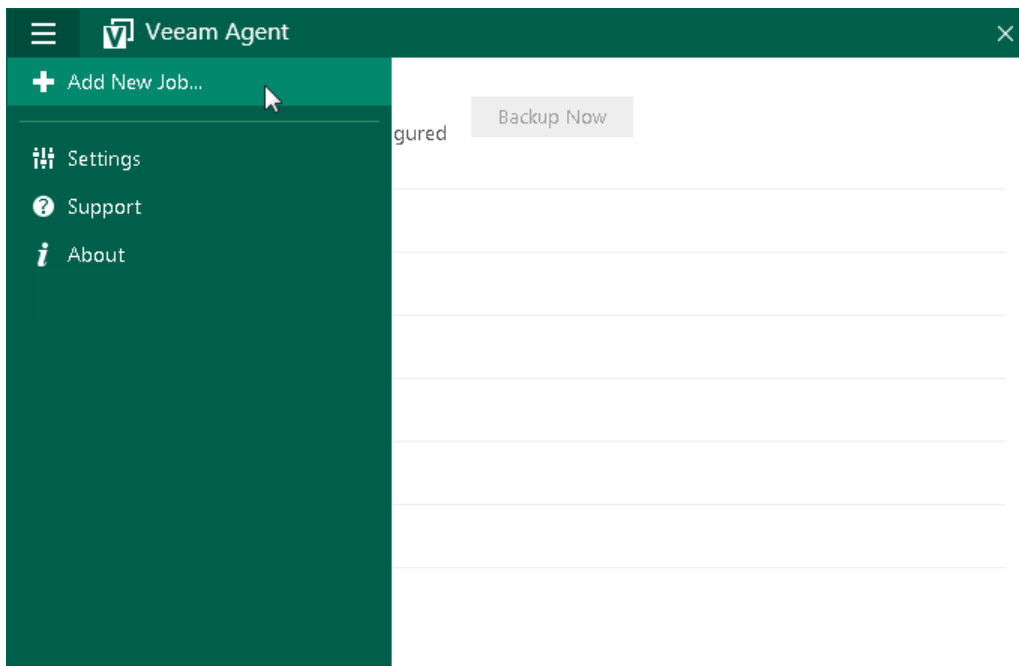
To launch the **New Backup Job** wizard, do either of the following:

- [If no backup jobs are configured] Double-click the Veeam Agent for Microsoft Windows icon in the system tray.
- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Backup > Configure backup**.

## NOTE

The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**. Then, from the main menu, select **Add New Job**.
- From the Microsoft Windows **Start** menu, select **All Programs > Veeam > Tools > Configure Backup** or use the Microsoft Windows search to find the *Configure Backup* option on your computer.



## Step 2. Specify Job Name and Description

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

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

New Backup Job

Name

Type in a name and description for this backup job.

Name

System Backup

Description:

Created by FILESR01\Administrator at 12/28/2022 10:47 AM.

< Previous Next > Finish Cancel

## Step 3. Select Backup Mode

At the **Backup Mode** step of the wizard, select the mode in which you want to create a backup:

1. In the **Backup mode** section, select the backup mode. You can select one of the following options:
  - **Entire computer** – select this option if you want to create a backup of the entire computer image. When you restore data from such backup, you will be able to recover the entire computer image as well as data on specific computer volumes: files, folders, application data and so on. With this option selected, you will pass to the [Destination](#) step of the wizard.
  - **Volume level backup** – select this option if you want to create a backup of specific computer volumes, for example, all volumes except the system one. When you restore data from such backup, you will be able to recover data on these volumes only: files, folders, application data and so on. With this option selected, you will pass to the [Volumes](#) step of the wizard.
  - **File level backup** – select this option if you want to create a backup of individual folders on your computer. When you backup data in this mode, you can configure additional rules to include/exclude files of the specific type in/from your backup. When you restore data from such backup, you will be able to recover backed-up files and folders. With this option selected, you will pass to the [Files](#) step of the wizard.
2. [For entire computer backup] If you want to include in the backup one or more external USB drives, select the **Include external USB drives** check box. With this option selected, Veeam Agent for Microsoft Windows will include in the backup all supported external drives that are connected to the Veeam Agent computer at the time when the backup job starts. Veeam Agent for Microsoft Windows supports backup of external drives that support Microsoft VSS: HDD, SSD, and so on. USB flash drives (USB sticks) are not supported. To learn more, see [Backup of External Drives](#).

### TIP

File-level backup is typically slower and requires more network traffic than volume-level backup. Depending on the performance capabilities of your computer and backup environment, the difference between file-level and volume-level backup job performance may increase significantly. If you plan to back up all folders with files on a specific volume or back up large amount of data, we recommend that you configure volume-level backup instead of file-level backup.

New Backup Job

Backup Mode

Choose what data you want to back up from this computer.

Name

Backup Mode

Destination

Local Storage

Guest Processing

Schedule

Summary

Entire computer (recommended)

Back up your entire computer image for fast recovery on any level. Deleted, temporary and page files are automatically excluded from the image to reduce the backup size.

☒ Include external USB drives

Volume level backup

Back up images of selected volumes, for example only data volumes. Deleted, temporary and page files are automatically excluded from the image to reduce the backup size.

File level backup (slower)

Back up individual files and folders by mask. This mode produces an image-based backup with only selected files included in the image.

< Previous

Next >

Finish

Cancel



## Step 4. Specify Backup Scope Settings

Specify backup scope for the backup job:

- [Select volumes to back up](#) — if you have selected the **Volume level backup** option at the [Backup Mode](#) step of the wizard.
- [Select folders to back up](#) — if you have selected the **File level backup** option at the [Backup Mode](#) step of the wizard.

### Selecting Volumes to Back Up

The **Volumes** step of the wizard is available if you have chosen to create a volume-level backup.

At this step of the wizard, you must specify the backup scope — define what volumes you want to include in the backup.

To specify the backup scope, choose volumes and objects that you want to include in the backup. You can include the following data in the backup:

- Computer volumes (located on local drives or external USB drives). To include individual volumes of your computer to the backup scope, select check boxes next to necessary volumes.

If you select all volumes at this step of the wizard, Veeam Agent will create a backup of the entire computer image. However, if a new volume appears on a Veeam Agent computer, Veeam Agent will not include this volume in the backup scope automatically. Veeam Agent will back up only those volumes that are explicitly selected for backup.

#### NOTE

You cannot include in the backup volumes located on virtual hard disks (VHD or VHDX).

- Mount points. To include individual volumes mounted on your computer to the backup scope, select check boxes next to necessary mount points.

In the list of objects to back up, Veeam Agent displays available mount points after volumes.

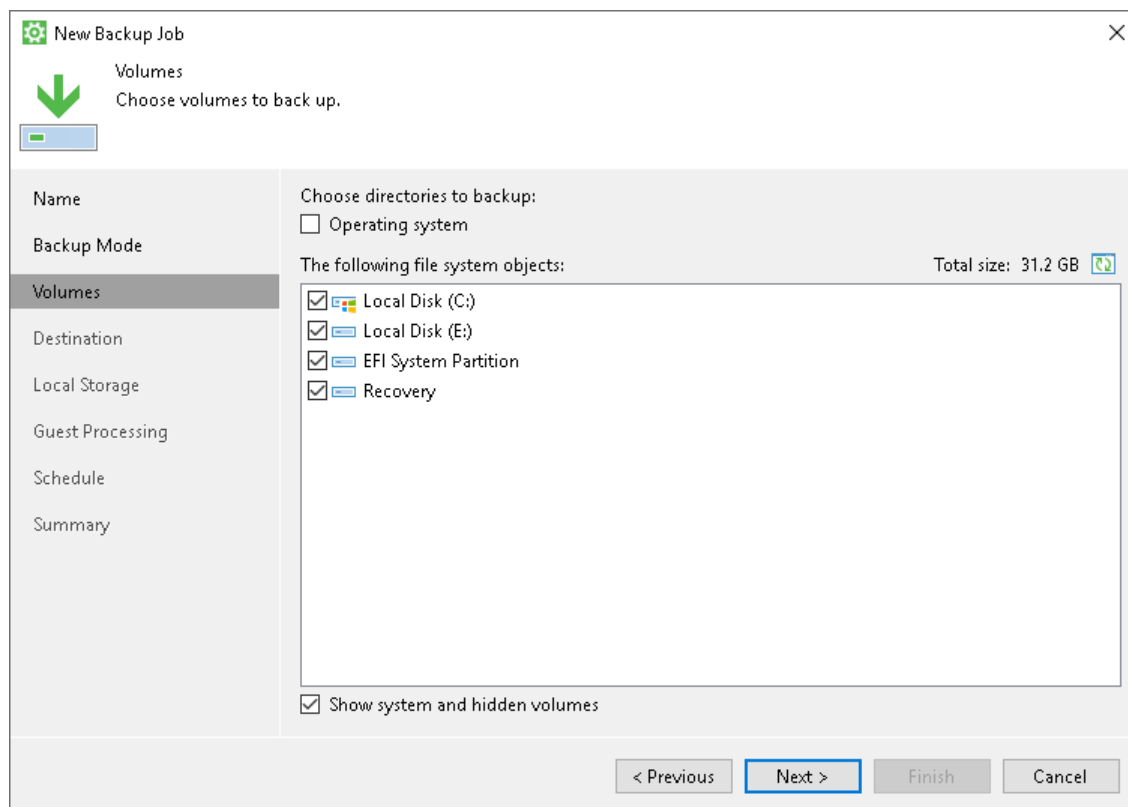
- System state data. To include system state data into the backup, do one of the following:
  - In the **Choose directories to backup** list, select the **Operating system** check box.
  - In the **The following file system objects** list, select the check box next to the system volume. Typically, the system volume is the **Local Disk (C:)** volume.
  - Select the Show system and hidden volumes check box at the bottom of the window. In the **The following file system objects** list, select the **System Reserved** check box and the check box next to the system volume. Typically, the system volume is the Local Disk (C:) volume.

Veeam Agent for Microsoft Windows will include in the backup scope the Microsoft Windows system partition and boot partition of your computer. For GPT disks on Microsoft Windows 8.1, 10, 11, 2012, 2012 R2, 2016, 2019 and 2022, Veeam Agent will additionally back up the recovery partition. To learn more, see [System State Data Backup](#).

When you include a system volume in the backup, Veeam Agent automatically includes the System Reserved/UEFI or other system partitions in the backup too. If you do not want to back up the system state data, you can clear the **System Reserved/EFI System Partition** check box. However, in this case Veeam Agent does not guarantee that the OS will boot properly when you attempt to recover from such backup. To learn more, see [System State Data Backup](#).

## NOTE

Veeam Agent automatically adds to the list of exclusions the following Microsoft Windows objects for all computer users: temporary files folder, Recycle Bin, Microsoft Windows pagefile, hibernate file and VSS snapshot files from the System Volume Information folder.



## Selecting Folders to Back Up

The **Files** step of the wizard is available if you have chosen to create a file-level backup.

In the file-level backup mode, you can create two types of backups:

- File-level backup that includes individual folders on your computer.
- Hybrid backup that contains individual folders and specific volumes of your computer.

At this step of the wizard, you must specify the backup scope – define what folders with files or entire volumes you want to include in the backup.

To specify the backup scope, select check boxes next to necessary objects. You can include the following data in the backup:

- Operating system data – data related to the OS installed on your computer.
- Personal files – data related to user profiles. With this option enabled, Veeam Agent for Microsoft Windows will include in the backup scope settings and data related to Veeam Agent computer user profiles.

You can specify what personal data to include in the backup and choose whether to exclude roaming profiles from the backup. To do this, click **Choose** next to the **Personal files** field and select the necessary options in the **Personal Folders** window. To learn more, see [Personal Data Backup](#).

- System reserved data – system data required to boot the OS installed on your computer. With this option enabled, Veeam Agent for Microsoft Windows will include in the backup scope Microsoft Windows system partition and boot partition of your computer. For GPT disks on Microsoft Windows 8.1, 10, 11, 2012, 2012 R2, 2016, 2019 and 2022, Veeam Agent for Microsoft Windows will additionally back up the recovery partition. To learn more, see [System State Data Backup](#).
- Individual folders. To learn more, see [Selecting Folders](#).
- Individual computer volumes and mount points (located on local drives or external USB drives). To learn more, see [Selecting Volumes](#).

## IMPORTANT

Consider the following:

- Veeam Agent for Microsoft Windows does not include the following Microsoft Windows objects in the backup: temporary files folder, Recycle Bin, Microsoft Windows pagefile, hibernate file and VSS snapshot files from the System Volume Information folder.  
Also, Veeam Agent does not include in the file-level backup files encrypted with Windows Encrypting File System (EFS).
- If you have a file-level backup job configured and need to extend the volume where backed-up files reside, we strongly recommend to create an active full backup after the volume is extended. Otherwise, Veeam Agent may skip files during the job run even if these files are added to the backup scope. To learn more, see [Creating Active Full Backups](#).

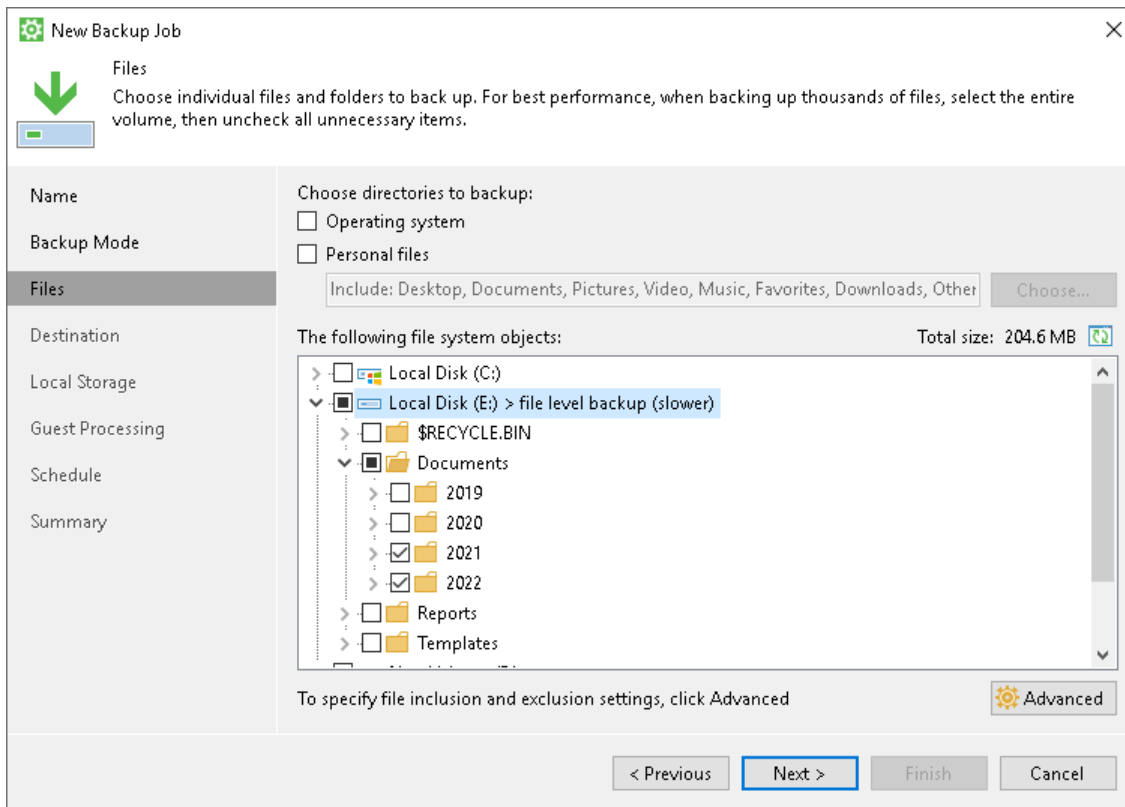
## Selecting Folders

To include a folder in the file-level backup:

1. In the **The following file system objects** list, browse to the folder that you want to back up and select the check box next to this folder.

After you select a folder, Veeam Agent will display the *file level backup* mark next to the volume that contains this folder.

2. If you want to include or exclude files of a specific type in/from the backup, click **Advanced**. To learn more, see [Configuring Filters](#).

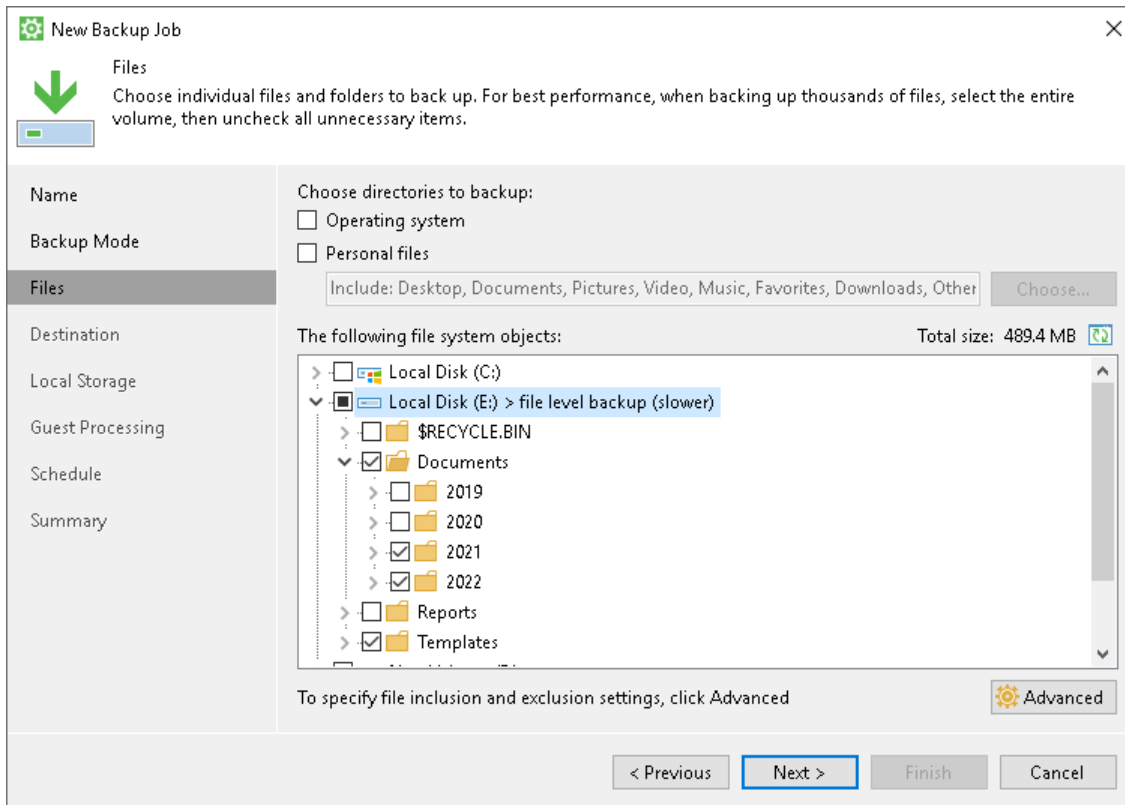


Alternatively, if you want to quickly include a large number of folders in the backup, you can select folders in the following way:

1. Select the check box next to the volume that contains folders you want to back up.  
Veeam Agent will display the *volume level backup* mark next to the selected volume.
2. Expand the volume and clear the check boxes next to one or more folders that you want to exclude from the backup.

3. Select the check box next to the volume once again.

Veeam Agent will display the *file level backup* mark next to the volume. You can include or exclude files that reside on this volume in/from the backup. To learn more, see [Configuring Filters](#).



## Selecting Volumes

You can include individual volumes in the backup at the **Files** step of the wizard. This may be useful if you want to have a volume-level backup of the volume with specific folders excluded. Volume-level backup also offers better performance in case you backup a complex folder structure with many files.

You can also include individual mount points in the backup at the **Files** step of the wizard, but you cannot exclude folders from mount points.

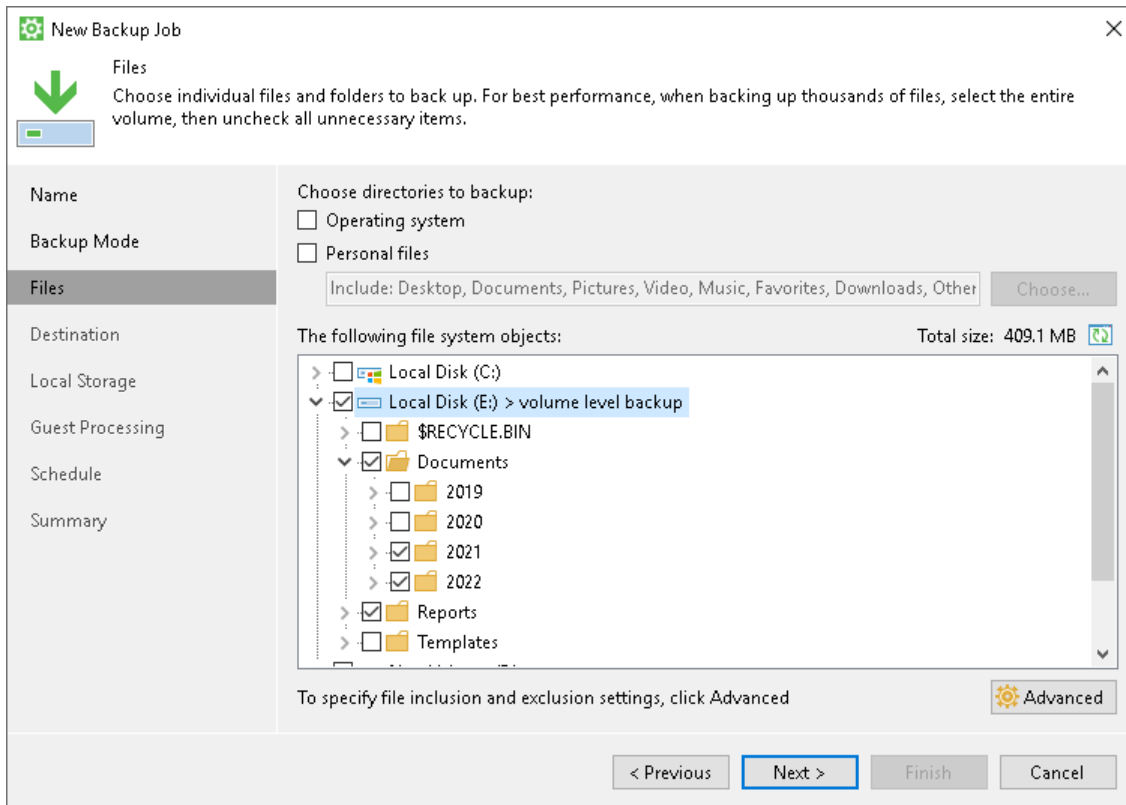
To include a volume or a mount point in the file-level backup:

1. In the **The following file system objects** list, select the check box next to the item that you want to back up.
2. If you want to exclude a folder on the selected volume from the backup, expand the volume and clear the check box next to the necessary folder.

### IMPORTANT

Consider the following:

- If you select a volume for backup, you cannot use filters to include or exclude files of a specific type in/from the backup. You can only exclude specific folders that reside on the volume. If you want to include or exclude files, you must select individual folders for backup. To learn more, see [Selecting Folders](#).
- If you select a deduplicated volume for backup, make sure that you do not exclude any folders from the backup. Otherwise, the backup files may get corrupted.



## Configuring Filters

To include or exclude files of a specific type in/from the file-level backup, you can configure filters.

### NOTE

Veeam Agent for Microsoft Windows applies filters to files in specific folders that you include in the backup. Filters are not applied to computer volumes and mount points selected for backup. If you plan to create a hybrid backup that will contain volumes, mount points, and folders, filters will be applied to files in folders only.

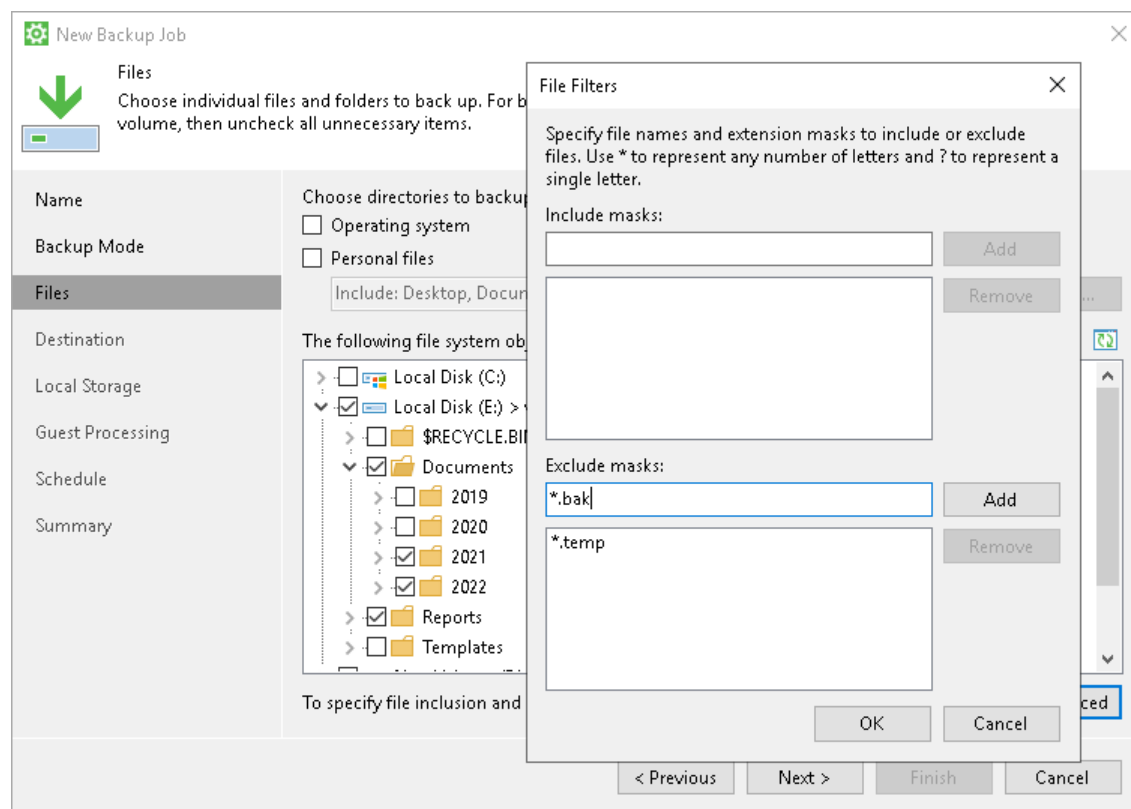
To configure a filter:

1. At the **Files** step of the wizard, click **Advanced**.
2. Specify what files you want to back up:
  - In the **Include masks** field, specify file names and masks for file types that you want to back up, for example, `MyMovie.avi`, `*filename*`, `*.docx`, `*.mp3`. Veeam Agent will create a backup only for selected files. Other files will not be backed up.
  - In the **Exclude masks** field, specify file names and masks for file types that you do not want to back up, for example, `OldPhotos.rar`, `*.temp`, `*.tmp`, `*.back`. Veeam Agent will back up all files except files of the specified type.
3. Click **Add**.
4. Repeat steps 2–3 for each mask that you want to add.

You can use a combination of include and exclude masks. Keep in mind that exclude masks have a higher priority than include masks. For example, you can specify masks in the following way:

- Include mask: \*.avi
- Exclude mask: \*movie\*

Veeam Agent will include in the backup all files of the AVI format that do not contain *movie* in their names.



## Step 5. Select Backup Destination

At the **Destination** step of the wizard, select a target location for the created backup.

You can store backup files in one of the following locations:

- **Local storage** – select this option if you want to save the backup on a removable storage device attached to the computer or on a local computer drive. With this option selected, you will pass to the [Local Storage](#) step of the wizard.
- **Object storage** – select this option if you want to save the backup in object storage. With this option selected, you will pass to the [Object storage](#) step of the wizard.
- **Shared folder** – select this option if you want to save the backup in a network shared folder. With this option selected, you will pass to the [Shared folder](#) step of the wizard.
- **Veeam backup repository** – select this option if you want to save the backup in a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Backup Server](#) step of the wizard.
- **Veeam Cloud Connect repository** – select this option if you want to save the backup in a cloud repository exposed to you by the Veeam Cloud Connect service provider. With this option selected, you will pass to the [Service Provider](#) step of the wizard.



## IMPORTANT

Consider the following:

- We strongly recommend that you store backups in the external location like USB storage device or shared network folder. You can also keep your backup files on the separate non-system local drive.
- If you select to store the backup in a local folder included in the backup scope, Veeam Agent will automatically exclude this folder from the backup.
- For the Free product edition, the **Object storage** and the **Veeam Cloud Connect repository** options are not available.

The screenshot shows the 'New Backup Job' wizard in Veeam Agent. The 'Destination' step is selected in the left sidebar. The main area displays instructions and five radio button options for backup destinations. The 'Next >' button is highlighted with a blue border.

**New Backup Job** [Close]

**Destination**  
Choose a target location for your backup. We highly recommend that you do not store your backups on the same computer that you are protecting.

**Left Sidebar:**

- Name
- Backup Mode
- Files
- Destination**
- Local Storage
- Guest Processing
- Schedule
- Summary

**Options:**

- ☒ **Local storage**  
Back up to a locally attached storage device such as USB, Firewire or eSATA external hard drive. Backing up to internal hard drives is not recommended.
- ☐ **Object storage**  
Back up to an on-prem or cloud object storage.
- ☐ **Shared folder**  
Back up to an SMB (CIFS) share on a Network Attached Storage (NAS) device, or on a regular file server.
- ☐ **Veeam backup repository**  
Back up to a repository managed by Veeam Backup & Replication 12 or later server.
- ☐ **Veeam Cloud Connect repository**  
Back up to a cloud repository managed by Veeam Cloud Connect service provider.

**Buttons:** < Previous | **Next >** | Finish | Cancel

# Step 6. Specify Backup Storage Settings

Specify backup storage settings for the backup job:

- [Local storage settings](#) – if you have selected the **Local storage** option at the [Destination](#) step of the wizard.
- [Object storage settings](#) – if you have selected the **Object storage** option at the [Destination](#) step of the wizard.
- [Shared folder settings](#) – if you have selected the **Shared folder** option at the [Destination](#) step of the wizard.
- [Veeam backup repository settings](#) – if you have selected the **Veeam backup repository** option at the [Destination](#) step of the wizard.
- [Veeam Cloud Connect repository settings](#) – if you have selected the **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.

## Local Storage Settings

The **Local Storage** step of the wizard is available if you have chosen to save the backup on a local drive of your computer.

Specify local storage settings:

1. In the **Local drives** list, select a drive where you want to store the backup.
2. In the **Folder** field, specify a path to the folder where backup files must be saved. By default, Veeam Agent saves files in the `VeeamBackup` folder.
3. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

4. Specify short-term backup retention policy settings:
  - [For Free and Workstation product editions] In the **Keep backups for <N> days (excluding days with no backup)** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days. After this period is over, Veeam Agent will remove the earliest restore points from the backup chain.  
To learn more, see [Backup Retention Policy in Free and Workstation Editions](#).
  - [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.
    - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
    - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

To learn more, see [Backup Retention Policy in Server Edition](#).

5. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the [Long-Term Retention Policy](#) section in the Veeam Backup & Replication User Guide.
6. Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

## IMPORTANT

USB storage devices formatted as FAT32 do not allow storing files larger than 4 GB in size. For this reason, we recommend that you do not use such USB storage devices as a backup target.

**New Backup Job**

Local Storage  
Choose a locally attached drive to back up to.

**Local drives:**

Storage device	Free space	Total space
Local Disk (C:)	67.4 GB	129.4 GB
Local Disk (E:)	89.4 GB	90.0 GB
Local Disk (F:)	89.9 GB	90.0 GB

**Folder:**  
E:\VeeamBackup\ [Browse...](#) [Map backup](#)

Retention policy: 7 days

☒ Keep certain full backups longer for archival purposes  
1 weekly, 1 monthly [Configure...](#)

Click Advanced to enable periodic full backups, configure encryption and other backup file settings [Advanced...](#)

< Previous **Next >** Finish Cancel

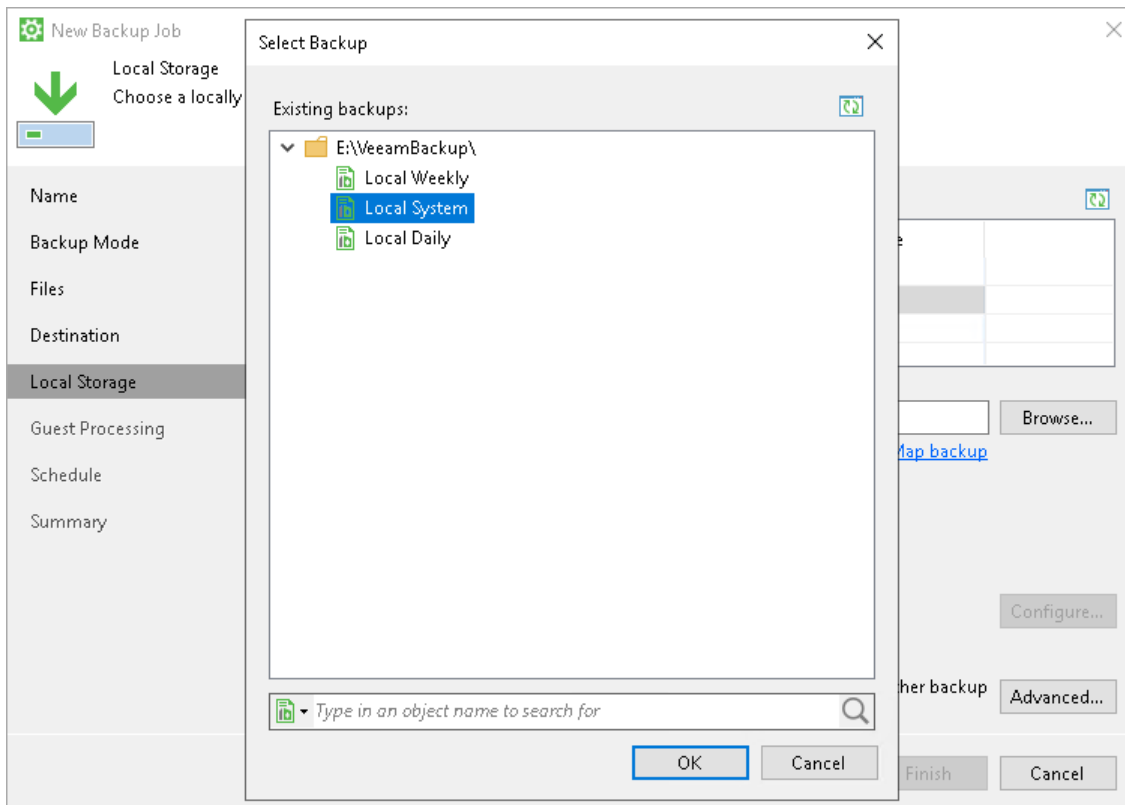
## Mapping Backup Job

If you have selected to map the job to the already created backup that is stored on a local drive, perform the following steps:

1. In the **Folder** field, specify the same path as was specified in the **Folder** field for the job that was used to create the backup.
2. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.

Keep in mind that Veeam Agent displays backups stored in the folder that is specified in the **Folder** field and its first-level subfolders.

3. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.



## Object Storage Settings

The **Object Storage** step of the wizard is available if you have selected the **Object storage** option at the [Destination](#) step of the wizard.

At the **Object Storage** step of the wizard, select the object storage. You can select one of the following options:

- **S3 compatible** – select this option if you want to create a backup in the S3 compatible storage. With this option selected, you will pass to the [Account](#) step of the wizard.

### NOTE

If you plan to store backups in an IBM or Wasabcloud storage, use the **S3 compatible** option.

- **Amazon S3** – select this option if you want to create a backup in the Amazon S3 storage. With this option selected, you will pass to the [Account](#) step of the wizard.
- **Google Cloud Storage** – select this option if you want to create a backup in the Google Cloud storage. With this option selected, you will pass to the [Account](#) step of the wizard.

- **Microsoft Azure Blob Storage** — select this option if you want to create a backup in the Microsoft Azure storage. With this option selected, you will pass to the [Account](#) step of the wizard.

**New Backup Job**

Object Storage  
Select an object storage type to back up to.

**Object Storage**

- ☒ **S3 Compatible**  
Select this option if you want to store backup files on on-premises object storage system or on a cloud storage provider.
- ☐ **Amazon S3**  
Select this option if you want to store backup files on the Amazon S3 cloud object storage.
- ☐ **Google Cloud Storage**  
Select this option if you want to store backup files on the Google Cloud Storage.
- ☐ **Microsoft Azure Blob Storage**  
Select this option if you want to store backup files on the Azure Blob Storage.

< Previous   **Next >**   Finish   Cancel

## S3 Compatible Settings

If you have selected to store backup files in the S3 compatible storage, specify the following settings:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in object storage.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.

### NOTE

If you want to connect to the repository using an IPv6 address and port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is the `ipv6-literal.net` notation of the IPv6 address.
- `port` is the number of the port that Veeam Agent will use to connect to the cloud storage.

For example: `https://2001-db8-1--bb8-c0b8-112.ipv6-literal.net:9000`

2. In the **Region** field, specify the storage region based on your regulatory and compliance requirements.
3. In the **Access key** field, enter the access key ID.

4. In the **Secret key** field, enter the secret access key.

The screenshot shows the 'New Backup Job' wizard in the 'Account' step. The left sidebar contains a list of steps: Name, Backup Mode, Files, Destination, Object Storage, Account (selected), Bucket, Backup Cache, Guest Processing, Schedule, and Summary. The main area is titled 'Account' with the instruction 'Specify an access and secret key for connecting to the S3 compatible storage system.' Below this, there are several input fields: 'Service point:' with the value 'https://myservicepoint.com', 'Region:' with the value 'global', 'Storage account:', 'Access key:' with the value 'F314BDAC823E13D427D692D49C129', and 'Secret key:' which is currently filled with dots and has a visibility icon (an eye) to its right. At the bottom right, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

### NOTE

You cannot select a folder that is managed by the Veeam Backup & Replication server.

If necessary, you can create a new folder. To do this, click the **New Folder** option in the **Select Folder** window.

3. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

4. In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep certain full backups for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see [Long-Term Retention Policy](#).

5. To protect recent backups from modifications, select the **Make recent backups immutable for <N> days** check box and specify the number of days for which you want to keep your backups immutable. By default, Veeam Agent keeps backups immutable for 30 days. To learn more, see [Backup Immutability](#).

#### NOTE

If you use the GFS retention scheme and enable immutability for a backup, the restore points with GFS flags become immutable for the whole GFS retention period. You will not be able to delete such restore points till the end of the GFS retention period.

6. Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

After that, Veeam Agent will create a new repository in the cloud storage where you can store backups.

The screenshot shows the 'New Backup Job' wizard window, specifically the 'Bucket' step. The window has a sidebar on the left with tabs: Name, Backup Mode, Files, Destination, Object Storage, Account, **Bucket**, Backup Cache, Guest Processing, Schedule, and Summary. The 'Bucket' tab is selected. The main area contains the following fields and options:

- Bucket:** A text field containing 'buck-1' and a 'Browse...' button.
- Folder:** A text field containing 'veeam\_backup' and a 'Browse...' button.
- Retention policy:** A field with a spinner set to '7' and a dropdown menu set to 'days'.
- ☒ **Keep certain full backups longer for archival purposes** (1 weekly, 1 monthly) with a 'Configure...' button.
- ☒ **Make recent backups immutable for:** 30 days. Below this is a descriptive text: 'Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy. Object storage must support S3 Object Lock feature.'
- A link labeled [Map backup](#).
- A button labeled **Advanced...** with the text 'Click Advanced to enable periodic full backups, configure encryption and other backup file settings' above it.

At the bottom of the window are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

# Mapping Backup Job

You can map the job to the already created backup that is stored in the S3 compatible storage. To map the backup job, perform the following steps:

1. In the **Bucket** field, specify a bucket where the created backup is stored:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
2. In the **Folder** field, specify a folder where the created backup is stored:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

## NOTE

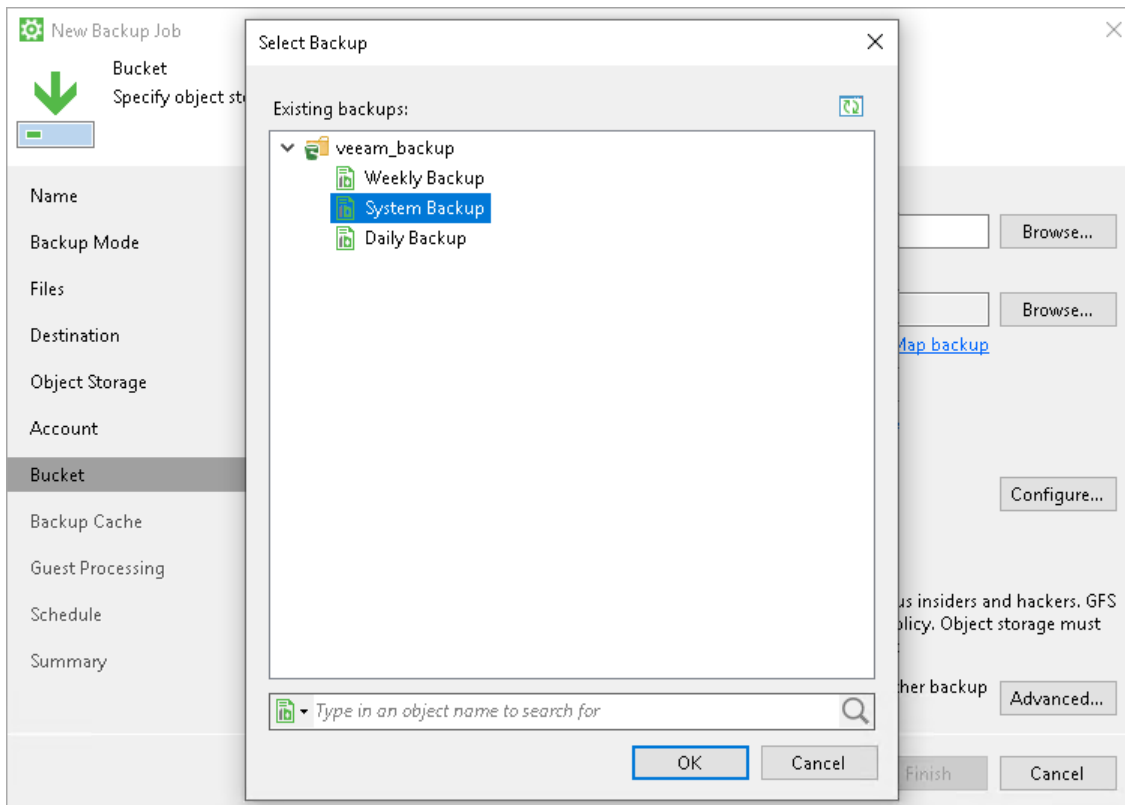
You cannot select a folder that is managed by the Veeam Backup & Replication server.

3. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.

Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same account used to connect to the S3 compatible storage.



4. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.



## Amazon S3 Settings

If you have selected to store backup files in the Amazon S3 storage, specify the following settings:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in object storage.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter the access key ID.
2. In the **Secret key** field, enter the secret access key.

3. From the **AWS region** drop-down list, select the AWS region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Global** region.

The screenshot shows the 'New Backup Job' wizard in the 'Account' step. The left sidebar contains a list of steps: Name, Backup Mode, Files, Destination, Object Storage, Account (highlighted), Bucket, Backup Cache, Guest Processing, Schedule, and Summary. The main area is titled 'Account' and contains the instruction 'Specify an Amazon access and secret key for connecting to the Amazon S3 storage.' Below this, there are four fields: 'Amazon AWS account:' (empty), 'Access key:' (containing 'RTWGHGT6BDRHKYT9SE1'), 'Secret key:' (masked with dots and a toggle icon), and 'AWS region:' (a dropdown menu set to 'Global'). At the bottom right, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in object storage and specified account settings to connect to the storage.

### IMPORTANT

You must create the bucket where you want to store your backup data beforehand. When you create a bucket, consider Amazon bucket naming rules. For example, it is not recommended that you use dots (.) in the bucket name. For more information on bucket naming rules, see [this AWS article](#).

Specify settings for the bucket in the storage:

1. From the **Data center** drop-down list, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.

b. In the **Select Folder** window, do the following:

- i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
- ii. Select the necessary folder and click **OK**.

#### NOTE

You cannot select a folder that is managed by the Veeam Backup & Replication server.

If necessary, you can create a new folder. To do this, click the **New Folder** option in the **Select Folder** window.

4. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

5. In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep certain full backups for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see [Long-Term Retention Policy](#).

6. To protect recent backups from modifications, select the **Make recent backups immutable for <N> days** check box and specify the number of days for which you want to keep your backups immutable. By default, Veeam Agent keeps backups immutable for 30 days. To learn more, see [Backup Immutability](#).

#### NOTE

If you use the GFS retention scheme and enable immutability for a backup, the restore points with GFS flags become immutable for the whole GFS retention period. You will not be able to delete such restore points till the end of the GFS retention period.

7. Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

After that, Veeam Agent will create a new repository in the cloud storage where you can store backups.

## Mapping Backup Job

You can map the job to the already created backup that is stored in the Amazon S3 storage. To map the backup job, perform the following steps:

1. From the **Data center** drop-down list, select the geographic region where the created backup is stored.
2. In the **Bucket** field, specify a bucket where the created backup is stored:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder where the created backup is stored:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

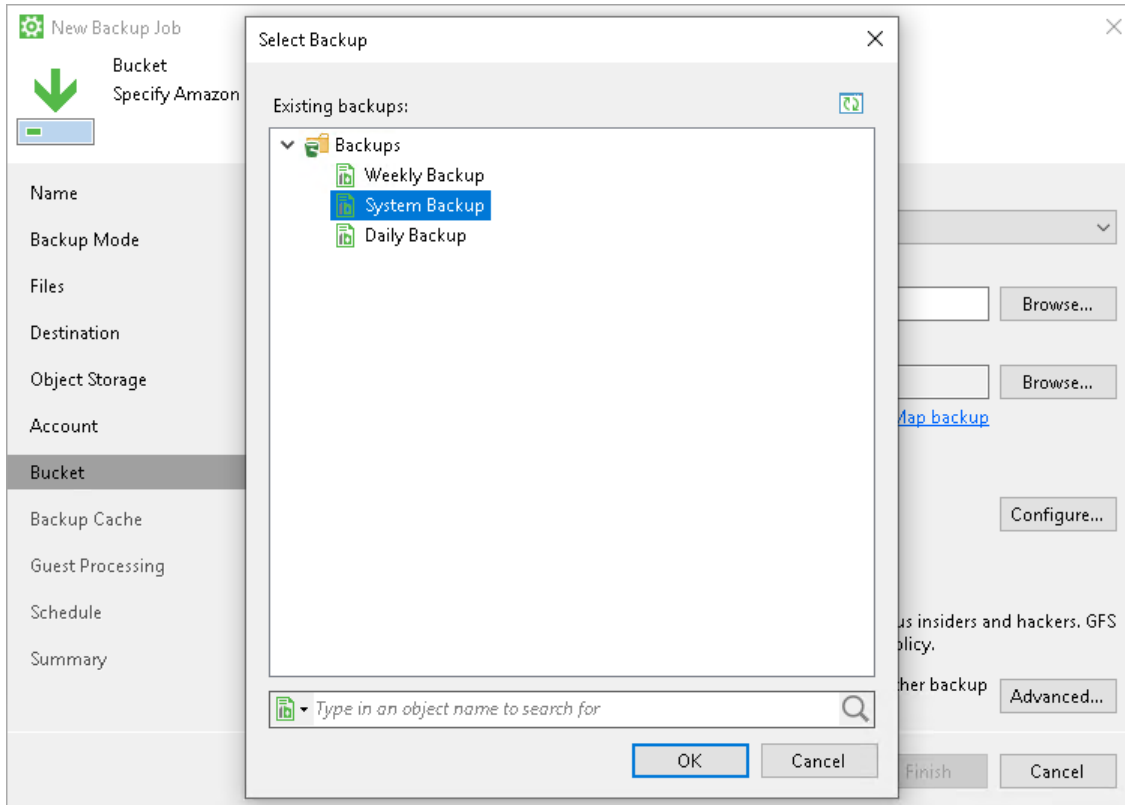
### NOTE

You cannot select a folder that is managed by the Veeam Backup & Replication server.

4. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.

Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same account used to connect to the Amazon S3 storage.

5. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.



## Google Cloud Storage Settings

If you have selected to store backup files in the Google Cloud storage repository, specify the following settings:

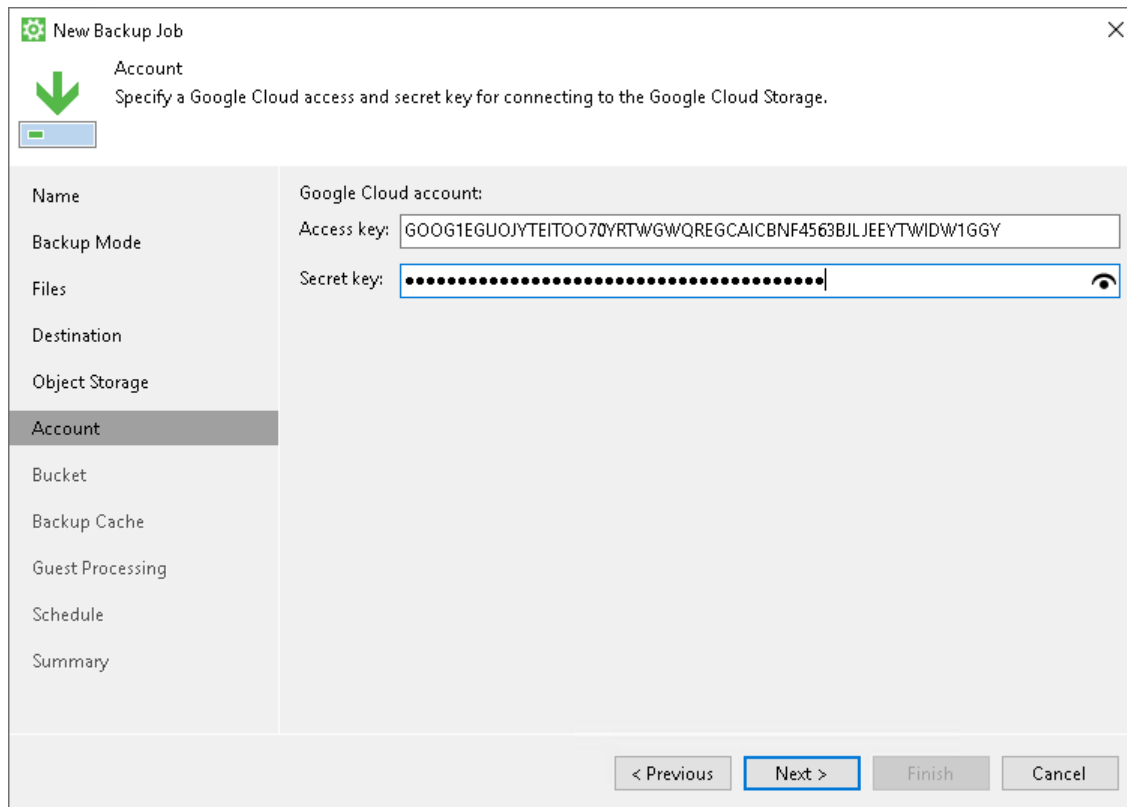
1. [Specify account settings](#).
2. [Specify bucket settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in object storage.

To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) keys associated with the Google Cloud account. Veeam Agent will use the HMAC keys to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see [Google Cloud documentation](#).

If you have not created the HMAC keys beforehand, you can create the keys in the Google Cloud console, as described in [Google Cloud documentation](#).



The screenshot shows the 'New Backup Job' wizard in the Veeam Agent for Microsoft Windows. The 'Account' step is selected in the left-hand navigation pane. The main area contains the following fields:

- Google Cloud account:** A text field with the value 'GOOG1EGUOJYTEITOO70YRTWGWQREGCAICBNF4563BJLJEEYTWDW1GGY'.
- Access key:** A text field with the value 'GOOG1EGUOJYTEITOO70YRTWGWQREGCAICBNF4563BJLJEEYTWDW1GGY'.
- Secret key:** A text field with a masked value (dots) and a toggle icon on the right.

At the bottom of the wizard, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. From the **Data center region** drop-down list, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

## NOTE

You cannot select a folder that is managed by the Veeam Backup & Replication server.

If necessary, you can create a new folder. To do this, click the **New Folder** option in the **Select Folder** window.

4. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

5. In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep certain full backups for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see [Long-Term Retention Policy](#).

5. Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

After that, Veeam Agent will create a new repository in the cloud storage where you can store backups.

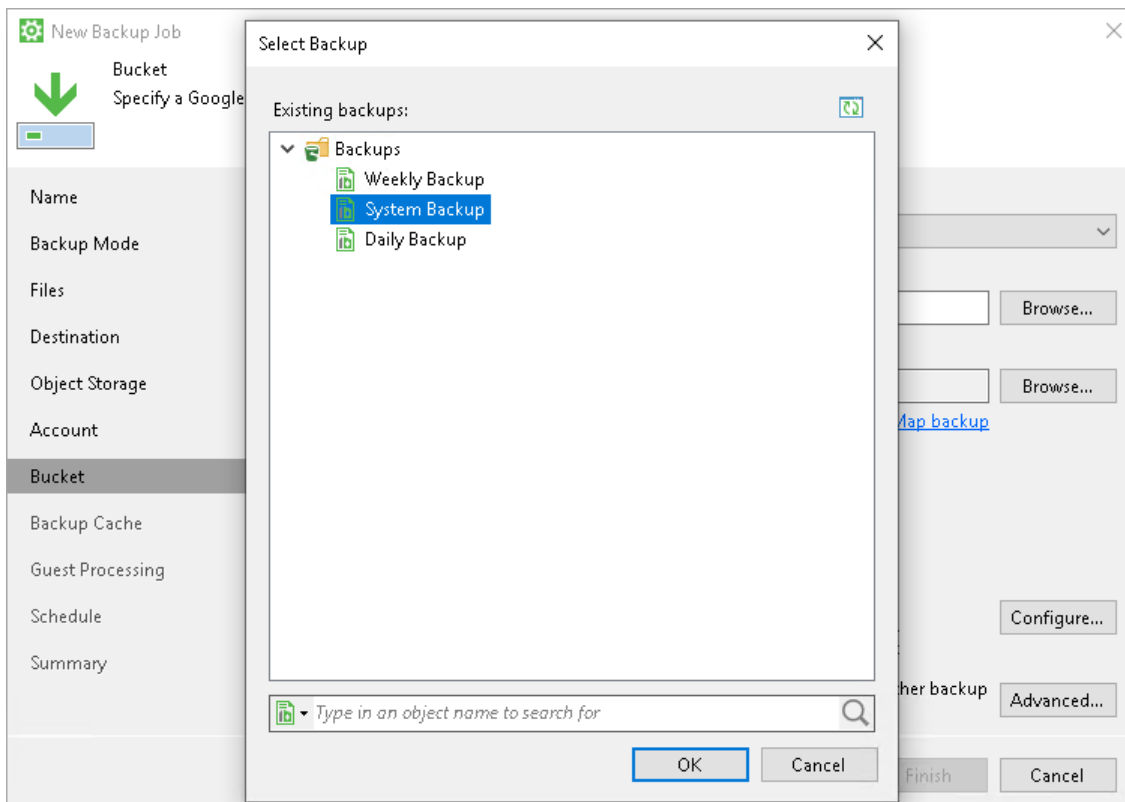
The screenshot shows the 'New Backup Job' wizard window. The 'Bucket' step is selected in the left sidebar. The main area displays configuration options for Google Cloud storage. The 'Data center region' is set to 'eu-west1 (Belgium)'. The 'Bucket' field contains 'eu-west1' with a 'Browse...' button. The 'Folder' field contains 'Backups' with a 'Browse...' button. A 'Map backup' link is visible. The 'Retention policy' is set to '7 days'. The 'Keep certain full backups longer for archival purposes' checkbox is checked, with a 'Configure...' button. An 'Advanced...' button is also present. At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

## Mapping Backup Job

You can map the job to the already created backup that is stored in the Google Cloud storage. To map the backup job, perform the following steps:

1. From the **Data center region** drop-down list, select the geographic region where the created backup is stored.

2. In the **Bucket** field, specify a bucket where the created backup is stored:
    - a. Select the **Browse** option.
    - b. In the **Select Bucket** window, do the following:
      - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
      - ii. Select the necessary bucket and click **OK**.
  3. In the **Folder** field, specify a folder where the created backup is stored:
    - a. Select the **Browse** option.
    - b. In the **Select Folder** window, do the following:
      - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
      - ii. Select the necessary folder and click **OK**.
- NOTE**
- You cannot select a folder that is managed by the Veeam Backup & Replication server.
4. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.
- Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same account used to connect to the Google Cloud storage.
5. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.





## Microsoft Azure Blob Storage Settings

If you have selected to store backup files in the Microsoft Azure Blob storage, specify the following settings:

1. [Specify account settings.](#)
2. [Specify container settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in object storage.

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.
3. From the **Region** drop-down list, select the Microsoft Azure region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Azure Global (Standard)** region.

### NOTE

The **Allow storage account key access** option for **Shared Key authorization** must be enabled in the storage account. For information on how to find this option, see [Microsoft documentation](#).

The screenshot shows the 'New Backup Job' wizard window. The 'Account' step is selected in the left sidebar. The main area displays the 'Microsoft Azure Blob Storage account:' section with three fields: 'Account' (containing 'newstorage'), 'Shared key' (masked with dots and a toggle icon), and 'Region' (a dropdown menu set to 'Azure Global (Standard)'). At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

New Backup Job	
Account Specify a Microsoft Azure account and shared key for connecting to the Microsoft Azure Blob Storage.	
Name	Microsoft Azure Blob Storage account:
Backup Mode	Account: newstorage
Files	Shared key: [masked]
Destination	Region: Azure Global (Standard)
Object Storage	
Account	
Container	
Backup Cache	
Guest Processing	
Schedule	
Summary	

## Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to save backup files in the Microsoft Azure storage and specified account settings to connect to the storage.

Specify settings for the container in the storage:

1. From the **Container** drop-down list, select a container in the storage.
2. In the **Folder** field, specify a folder in the container:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the container name or click the arrow to the left of the container name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

#### NOTE

You cannot select a folder that is managed by the Veeam Backup & Replication server.

If necessary, you can create a new folder. To do this, click the **New Folder** option in the **Select Folder** window.

3. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

4. In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.

To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep certain full backups for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see [Long-Term Retention Policy](#).

5. To protect recent backups from modifications, select the **Make recent backups immutable for <N> days** check box and specify the number of days for which you want to keep your backups immutable. By default, Veeam Agent keeps backups immutable for 30 days. To learn more, see [Backup Immutability](#).

#### NOTE

If you use the GFS retention scheme and enable immutability for a backup, the restore points with GFS flags become immutable for the whole GFS retention period. You will not be able to delete such restore points till the end of the GFS retention period.

6. Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

After that, Veeam Agent will create a new repository in the cloud storage where you can store backups.

**New Backup Job**

Container  
Specify Microsoft Azure Blob Storage container to use.

**Name**  
**Backup Mode**  
**Files**  
**Destination**  
**Object Storage**  
**Account**  
**Container**  
**Backup Cache**  
**Guest Processing**  
**Schedule**  
**Summary**

Container: veeam

Folder: Backups [Browse...](#)

[Map backup](#)

Retention policy: 7 days

☒ Keep certain full backups longer for archival purposes  
1 weekly, 1 monthly [Configure...](#)

☒ Make recent backups immutable for: 30 days  
Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.

Click Advanced to enable periodic full backups, configure encryption and other backup file settings [Advanced...](#)

< Previous Next > Finish Cancel

## Mapping Backup Job

You can map the job to the already created backup that is stored in the Microsoft Azure Blob storage. To map the backup job, perform the following steps:

1. From the **Container** drop-down list, select a container where the created backup is stored.
2. In the **Folder** field, specify a folder where the created backup is stored:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the container name or click the arrow to the left of the container name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

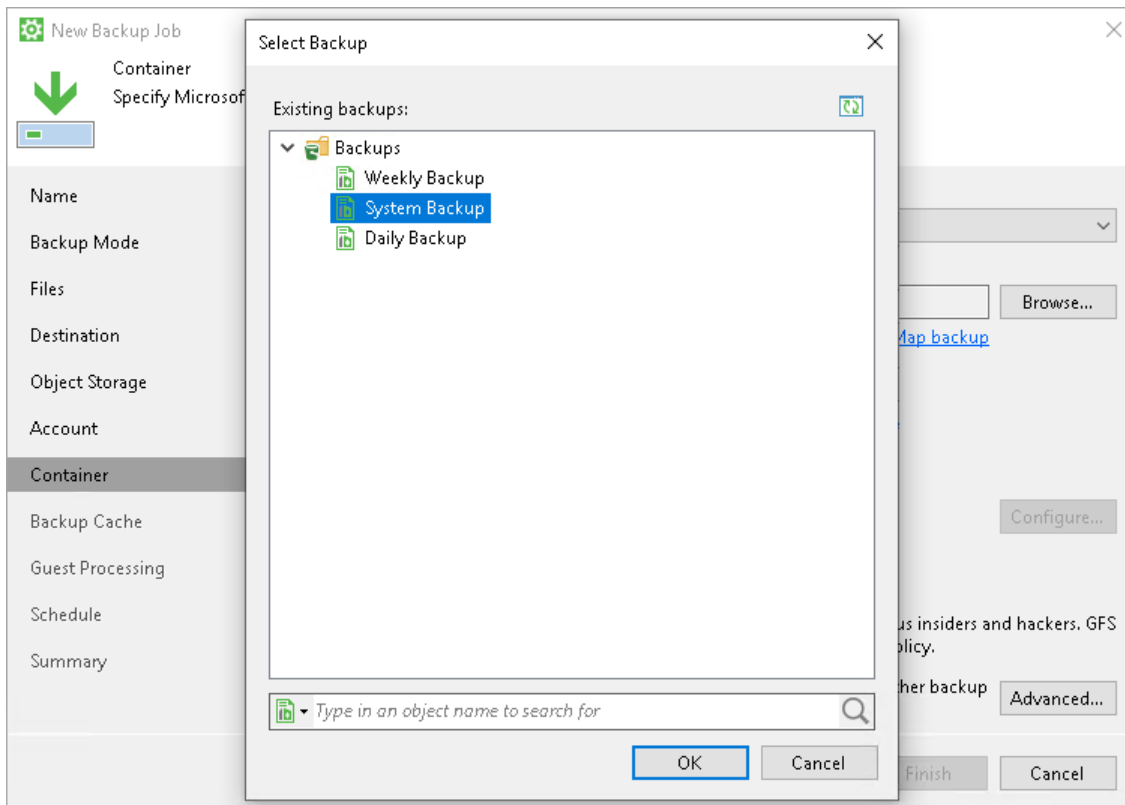
### NOTE

You cannot select a folder that is managed by the Veeam Backup & Replication server.

3. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.

Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same account used to connect to the Microsoft Azure Blob storage.

4. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.



## Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have chosen to save the backup in a network shared folder.

Specify shared folder settings:

1. In the **Shared folder** field, type a UNC name of the network shared folder in which you want to store backup files. Keep in mind that the UNC name always starts with two backslashes (\\).
2. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

3. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the [down-level logon name](#) format. For example, `DOMAIN\UserName` or `HOSTNAME\UserName`.

To view the specified password, click and hold the eye icon on the right of the **Password** field.

If you do not select the **This share requires access credentials** check box, Veeam Agent for Microsoft Windows will connect to the shared folder using the `NT AUTHORITY\SYSTEM` account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can grant *Full Control* access permissions on the shared folder and underlying file system to the computer account (`DOMAIN\COMPUTERNAME$`).

4. To view how much free space is available in the selected shared folder, click **Populate**.

5. Specify short-term backup retention policy settings:

- [For Free and Workstation product editions] In the **Keep backups for <N> days (excluding days with no backup)** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent for Microsoft Windows keeps backup files for 7 days. After this period is over, Veeam Agent for Microsoft Windows will remove the earliest restore points from the backup chain.

To learn more, see [Backup Retention Policy in Free and Workstation Editions](#).

- [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.
  - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
  - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

To learn more, see [Backup Retention Policy in Server Edition](#).

6. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the [GFS Retention Policy](#) section in the Veeam Backup & Replication User Guide.

7. Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

The screenshot shows the 'New Backup Job' wizard in Veeam Agent for Microsoft Windows, specifically the 'Shared Folder' step. The window has a title bar with a green gear icon and a close button. On the left is a sidebar with navigation links: Name, Backup Mode, Files, Destination, Shared Folder (selected), Backup Cache, Guest Processing, Schedule, and Summary. The main area contains the following fields and controls:

- Shared Folder:** A green downward arrow icon and the text 'Specify a shared folder and an account to connect to this shared folder.'
- Shared folder:** A text box containing '\\WINSRV001\SharedBackups' and a 'Browse...' button.
- Permissions:** A list box showing 'n/a' and a 'Map backup' link.
- Access Credentials:** A checked checkbox 'This share requires access credentials:' with fields for 'Username: TECH\administrator' and 'Password: [masked]'.
- Retention policy:** A spinner box set to '7' and a dropdown menu set to 'days'.
- Archival Purposes:** A checked checkbox 'Keep certain full backups longer for archival purposes' with a sub-label '1 weekly, 1 monthly' and a 'Configure...' button.
- Advanced Settings:** A text box with the message 'Click Advanced to enable periodic full backups, configure encryption and other backup file settings' and an 'Advanced...' button.
- Navigation:** At the bottom are buttons for '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

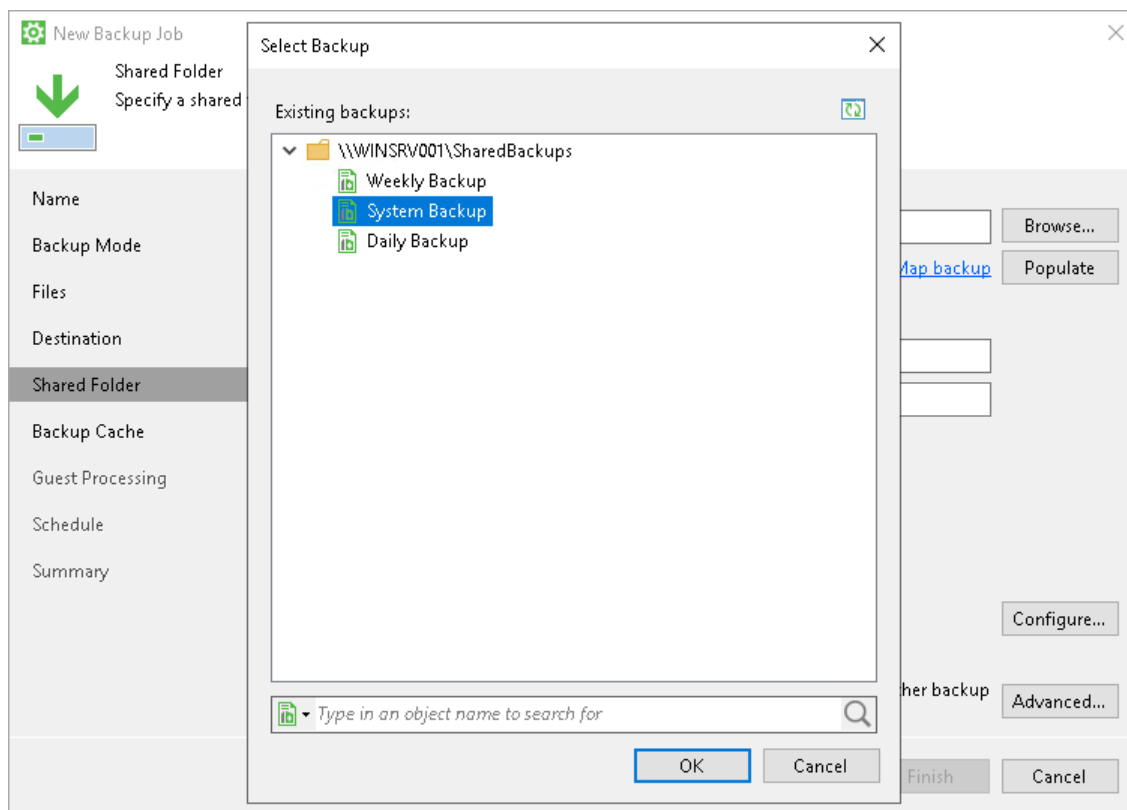
# Mapping Backup Job

If you have selected to map the job to the already created backup that is stored in a network shared folder, perform the following steps:

1. In the **Shared folder** field, specify the same path as was specified in the **Shared folder** field for the job that was used to create the backup.
2. Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window.

Keep in mind that Veeam Agent displays backups stored in the folder that is specified in the **Shared folder** field and its first-level subfolders.

3. If you map the job to an encrypted backup, and encryption keys are not available in the Veeam Agent database, you must provide the password specified for encryption. In the displayed window, enter the password and click **OK**.



## Veeam Backup Repository Settings

If you have selected to store backup files in a Veeam backup repository, specify settings to connect to the backup repository:

1. [At the Backup Server step of the wizard, specify backup server settings.](#)
2. [At the Backup Repository step of the wizard, select the Veeam backup repository.](#)

## Specifying Backup Server Settings

The **Backup Server** step of the wizard is available if you have chosen to store backup files in a Veeam backup repository.

Specify settings for the Veeam backup server that manages the target backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup server. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. Select the **Specify your personal credentials** check box. In the **Username** and **Password** fields, specify a user name and password of the account that has access to this backup repository. Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

If you do not select the **Specify your personal credentials** check box, Veeam Agent for Microsoft Windows will connect to the backup repository using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can add the computer account (*DOMAIN\COMPUTERNAME\$*) to an AD group and grant access rights on the backup repository to this group.

Setting access permissions on the backup repository to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). If you have set such permissions on the backup repository, you can omit specifying credentials. However, we recommend this scenario for demo environments only.

## IMPORTANT

Consider the following:

- If you plan to use a commercial version of Veeam Agent with Veeam Backup & Replication, you must install a license in Veeam Backup & Replication in advance, before connecting to the backup server.
- If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain in the backup repository.

The screenshot shows the 'New Backup Job' wizard in Veeam Agent for Microsoft Windows. The 'Backup Server' step is selected in the left-hand navigation pane. The main area contains the following fields and options:

- Name:** A label for the backup job.
- Backup Mode:** A dropdown menu.
- Files:** A section for file selection.
- Destination:** A section for destination selection.
- Backup Server:** The current step, containing:
  - Veeam backup server name or IP address:** A text field with the value '172.24.30.15'.
  - Port:** A spinner box with the value '10001'.
  - Specify your personal credentials:** A checked checkbox.
  - Username:** A text field with the value 'TECH\Administrator'.
  - Password:** A password field with masked characters and a visibility toggle icon.
- Backup Repository:** A section for repository selection.
- Backup Cache:** A section for cache selection.
- Guest Processing:** A section for guest processing settings.
- Schedule:** A section for scheduling.
- Summary:** A section for job summary.

At the bottom of the wizard, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted.

# Selecting Backup Repository

The **Backup Repository** step of the wizard is available if you have chosen to save backup files in a Veeam backup repository.

Specify settings for the target backup repository:

1. From the **Backup repository** list, select a backup repository where you want to store created backups. The **Backup repository** list displays only those backup repositories on which you have permissions to store data.

To refresh the list of backup repositories, click the **Refresh** button at the top right corner of the **Backup repository** field. Backup repositories list refresh may be required if you change permission settings for a specific backup repository on the Veeam backup server and want to display this backup repository in the **New Backup Job** wizard. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

2. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

3. Specify short-term backup retention policy settings:

- [For Free and Workstation product editions] In the **Keep backups for <N> days (excluding days with no backup)** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days. After this period is over, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Backup Retention Policy in Free and Workstation Editions](#).

- [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.
  - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
  - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

To learn more, see [Backup Retention Policy in Server Edition](#).

4. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the [GFS Retention Policy](#) section in the Veeam Backup & Replication User Guide.



5. Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

The screenshot shows the 'New Backup Job' wizard in Veeam Backup & Replication. The 'Backup Repository' step is active, indicated by a green arrow and a green bar in the left sidebar. The sidebar lists the following steps: Name, Backup Mode, Files, Destination, Backup Server, Backup Repository (selected), Backup Cache, Guest Processing, Schedule, and Summary. The main area contains the following information:

- Backup repository:** A dropdown menu showing 'Default Backup Repository (Created by Veeam Backup)' with a 'Map backup' link to its right.
- Files:** A bar chart showing '64.8 GB free of 129.4 GB'.
- Retention policy:** A dropdown menu set to '7' days.
- ☒ **Keep certain full backups longer for archival purposes** (1 weekly, 1 monthly). A 'Configure...' button is to the right.
- Click Advanced to enable periodic full backups, configure encryption and other backup file settings.** An 'Advanced...' button is to the right.

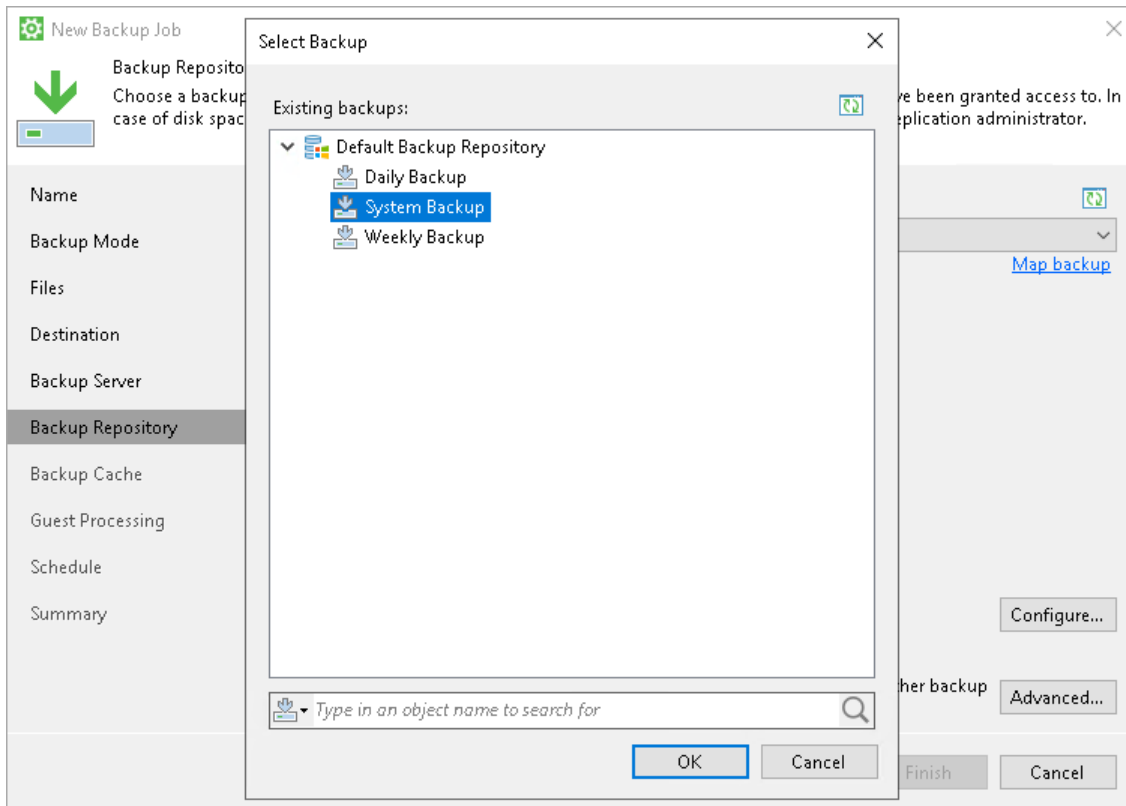
At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Mapping Backup Job

If you have selected to map the job to the already created backup that is stored in a Veeam backup repository, perform the following steps:

1. If you map the job to an encrypted backup, you must decrypt the backup on the Veeam Backup & Replication side before mapping. To learn more, see the [Restoring Data from Encrypted Backups](#) section in the Veeam Backup & Replication User Guide.

2. Click the **Map backup** link and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window. Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same account used to connect to the repository.



The backup job mapping can be helpful in case of the backup seeding. If you want to seed backup files, you must place these files in a specific folder on the target location. In the Veeam backup repository, Veeam Agent stores backups in the folders with paths of the following format:

```
<path_to_folder>\<domain_name>_<user_name>\<backup_job_name>
```

where:

- <path\_to\_folder> – path to the folder in the Veeam backup repository.
- <domain\_name> – domain name specified to connect to the repository.
- <user\_name> – user name specified to connect to the repository.
- <backup\_job\_name> – name of the backup job.

For example:

```
C:\Backup\TECH_Administrator\System Backup
```

To learn more about backup seeding scenarios for Veeam backup repositories, see [Moving Veeam Agent Backups to Veeam Backup Repository](#).

## Veeam Cloud Connect Repository Settings

If you have selected to store backup files in a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings](#).

2. [Verify the TLS certificate and specify user account settings.](#)
3. [Select the cloud repository.](#)

## Specifying Service Provider Settings

The **Service Provider** step of the wizard is available if you have chosen to save backup files in a Veeam Cloud Connect repository.

Specify settings for the cloud gateway that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.

### TIP

You can look for service providers who offer Repository as a Service using Veeam Backup & Replication. The list of service providers is published on the Veeam website and constantly updated. You can select the necessary service provider from the list and contact this service provider to get the cloud repository service.

To find a service provider, click the *Click here to open the directory* link. Veeam Agent for Microsoft Windows will open a webpage on the Veeam website. Use the filter on the webpage to find the necessary service provider by the type of provided cloud services, service provider datacenter location or service area.

The screenshot shows the 'New Backup Job' wizard in Veeam Agent for Microsoft Windows, specifically the 'Service Provider' step. The window has a title bar with a gear icon and the text 'New Backup Job'. Below the title bar, there is a green downward arrow icon and the text 'Service Provider' and 'Specify a DNS name or IP address and a port number received from the service provider.' A progress bar shows the 'Service Provider' step is active. On the left, a sidebar lists the steps: Name, Backup Mode, Files, Destination, Service Provider (highlighted), Credentials, Backup Resources, Backup Cache, Guest Processing, Schedule, and Summary. The main area contains a 'DNS name or IP address:' text box with '172.24.30.114' entered, and a 'Port:' spinner box set to '6180'. Below these fields, a note states: 'Default service provider's port is 6180. If connection cannot be established, contact with your service provider to make sure the settings are correct.' At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'. At the very bottom of the window, there is a link: 'Search for resellers and service providers that offer cloud repositories and cloud hosts for off-site backup and disaster recovery. [Click here to open the directory](#)'.

## Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to save backup files in a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the cloud repository.

1. At the top of the wizard window, Veeam Agent for Microsoft Windows displays information about the TLS certificate obtained from the SP side. You can view the certificate settings and verify the TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

- To view the TLS certificate, click the certificate link.
  - To verify the TLS certificate with a thumbprint, copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Thumbprint for certificate verification** field. Click **Verify**. Veeam Agent for Microsoft Windows will check if the thumbprint you enter matches the thumbprint of the obtained TLS certificate.
2. In the **Username** field, enter the user name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The user name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.
  3. In the **Password** field, provide a password for the tenant or subtenant account.

The screenshot shows the 'New Backup Job' wizard window, specifically the 'Credentials' step. The window has a title bar with a green gear icon and a close button. Below the title bar, there's a green downward arrow icon and the text 'Credentials' and 'Specify credentials that you have received from the service provider and validate the certificate.' A progress bar shows the 'Credentials' step is active. On the left, a sidebar lists the wizard steps: Name, Backup Mode, Files, Destination, Service Provider, Credentials (selected), Backup Resources, Backup Cache, Guest Processing, Schedule, and Summary. The main area displays a warning message: 'Warning: Warning: Remote certificate name mismatch: CN=Veeam Software, O=Veeam Software, OU=Veeam SoftwareRemote certificate chain errors: WarningUntrustedRoot (A certifi...'. Below the warning, it shows the 'Certificate: CN=Veeam Software, O=Veeam Software, OU=Veeam Software' and the 'Thumbprint for certificate verification: DE513021D51DE609FE9DBFEF0773126B54BB7D8C'. There is a 'Verify' button next to the thumbprint field. Below the thumbprint field, there are 'Username:' and 'Password:' fields. The 'Username' field contains 'Tech\User1' and the 'Password' field is masked with dots. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Selecting Cloud Repository

The **Backup Resources** step of the wizard is available if you have chosen to save backup files in a cloud repository and specified settings to connect to the SP.

Specify settings for the cloud repository:

1. From the **Available cloud repositories** list, select a cloud repository where you want to store created backups. The **Available cloud repositories** list displays only those backup repositories on which you have permissions to store data.

To refresh the list of cloud repositories, click the **Refresh** button at the top right corner of the **Available cloud repositories** field. Cloud repositories list refresh may be required if you change permission settings for a specific cloud repository and want to display this cloud repository in the **New Backup Job** wizard.

2. If you want to map the job to a specific backup that was previously created on the same Veeam Agent computer, click the **Map Backup** link and select the backup.

To learn more, see [Mapping Backup Job](#).

3. Specify short-term backup retention policy settings:

- [For Free and Workstation product editions] In the **Keep restore points for the last <N> days when computer was used** field, specify the number of days for which you want to store backup files in the target location. By default, Veeam Agent for Microsoft Windows keeps backup files for 7 days. After this period is over, Veeam Agent for Microsoft Windows will remove the earliest restore points from the backup chain.

To learn more, see [Backup Retention Policy in Free and Workstation Editions](#).

- [For Server product edition] In the **Retention policy** field, specify the number of restore points or days for which you want to store backup files in the target location. By default, Veeam Agent keeps backup files for 7 days.
  - Select the **restore points** option if you want Veeam Agent to remove restore points from the backup chain when the number of restore points exceeds the specified limit.
  - Select the **days** option if you want Veeam Agent to remove the earliest restore points from the backup chain when the specified period is over.

To learn more, see [Backup Retention Policy in Server Mode](#).

4. [For Workstation and Server product editions] To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep some periodic full backups longer for archival purposes** check box and click **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see the [GFS Retention Policy](#) section in the Veeam Backup & Replication User Guide.

- Click **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

**New Backup Job**

**Backup Resources**  
Choose a cloud backup repository. You can only select between repositories your service provider has assigned to you.

**Available cloud repositories:**

Repository	Free space	Capacity	
User 02 Backups	100.0 GB	100.0 GB	

[Map backup](#)

Retention policy: 7 days

☒ Keep certain full backups longer for archival purposes  
1 weekly, 1 monthly

[Click Advanced to enable periodic full backups, configure encryption and other backup file settings](#)

[Configure...](#) [Advanced...](#)

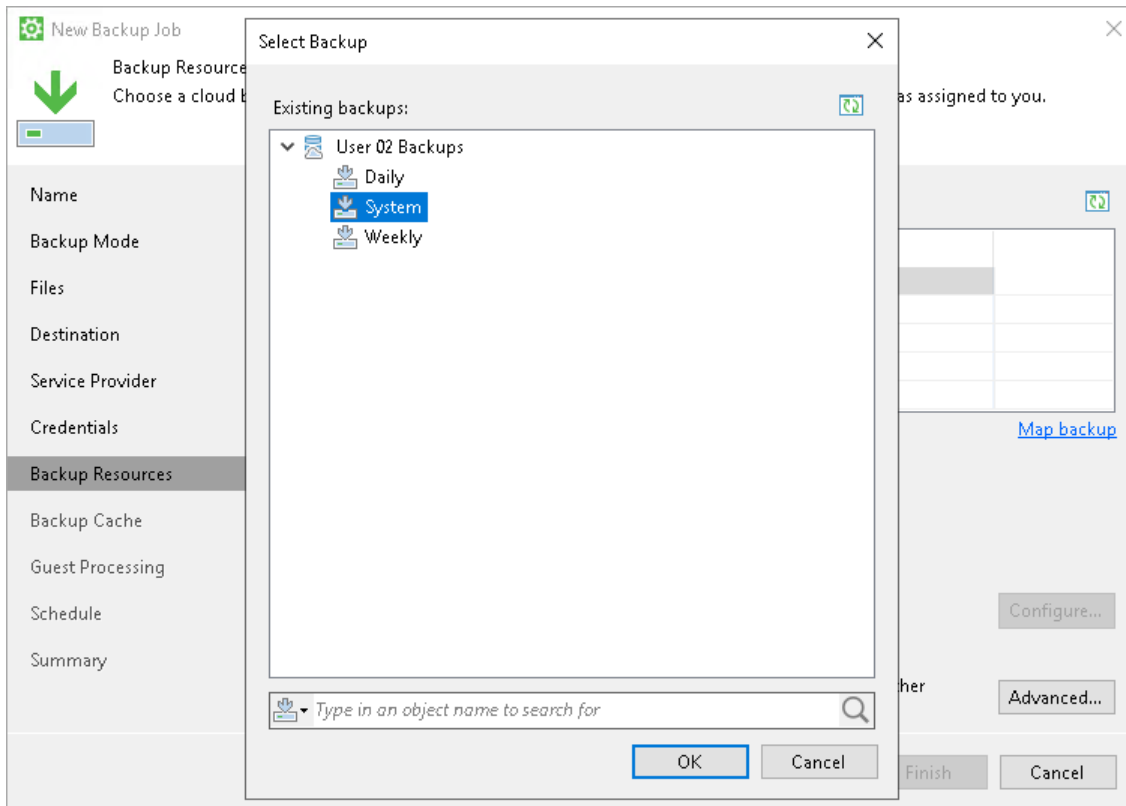
< Previous Next > Finish Cancel

## Mapping Backup Job

If you have selected to map the job to the already created backup, perform the following steps:

- Click **Map backup** and select the desired backup. To find the backup, you can use the search field at the bottom of the **Select Backup** window. Keep in mind that Veeam Agent displays only those backups that were created on the same Veeam Agent computer with the same tenant or subtenant account used to connect to the repository.

2. If you map the job to an encrypted backup, and encryption keys in the cloud repository database do not coincide with the encryption keys in the Veeam Agent database, you must provide the password specified for encryption after the job start. In the dialog window, enter the password and click **OK**. To learn more, see [Resuming Encrypted Backup Chain](#).



The backup job mapping can also be helpful in case of the backup seeding. If you want to seed backup files, you must place these files in a specific folder on the target location. On the cloud repository, Veeam Agent stores backups in the folders with paths of the following format:

- If a tenant account is specified in the **Username** field at the **Credentials** step of the wizard:  
`<path_to_repository>\<tenant_name>\<backup_job_name>`
- If a subtenant account is specified in the **Username** field at the **Credentials** step of the wizard:  
`<path_to_repository>\<tenant_name>\Users\<subtenant_name>\<backup_job_name>`

where:

- `<path_to_repository>` – path to the folder on the cloud repository.
- `<tenant_name>` – name of the tenant account.
- `<subtenant_name>` – name of the subtenant account.
- `<backup_job_name>` – name of the backup job.

For example:

- `C:\Backup\TechCompany\System Backup`
- or
- `C:\Backup\TechCompany\Users\User01\System Backup`

To learn more about backup seeding scenarios for cloud repositories, see [Moving Veeam Agent Backups to Veeam Cloud Connect Repository](#).



# Step 7. Specify GFS Retention Policy

This step of the wizard is available if you use the Workstation or Server edition of Veeam Agent for Microsoft Windows.

## NOTE

If you want to store your backups in object storage, and you do not configure Veeam Agent to perform active full backups periodically, Veeam Agent still requires a full backup to assign a GFS flag to a full restore point. In this case, Veeam Agent creates a full backup based on the last incremental backup. If some data blocks required to create the full backup reside in object storage, the full backup contains links that point at such data blocks. Veeam Agent does not retrieve actual data blocks from the object storage to avoid extra costs.

To configure GFS retention policy settings for the backup job:

1. Click the **Configure** button at one of the following steps of the wizard:
  - **Local Storage** — if you have selected the **Local storage** option at the **Destination** step of the wizard.
  - **Shared Folder** — if you have selected the **Shared folder** option at the **Destination** step of the wizard.
  - **Bucket/Container** — if you have selected the **Object storage** option at the **Destination** step of the wizard.
  - **Backup Repository** — if you have selected the **Veeam backup repository** option at the **Destination** step of the wizard.
  - **Backup Resources** — if you have selected the **Veeam Cloud Connect repository** option at the **Destination** step of the wizard.
2. In the **Configure GFS** window, do the following:
  - a. If you want to create weekly restore points for archival purposes, select the **Keep weekly full backups for** check box. Then specify the number of weeks during which you want to prevent restore points from being modified and deleted.
    - In the **If multiple full backups exist, use the one from** list, select a week day when Veeam Agent must assign the weekly GFS flag to a full restore point.
    - [For object storage target if active full backup schedule is not enabled] In the **Create weekly full on this day** list, select a week day when Veeam Agent must create a full backup and assign the weekly GFS flag to a full restore point.
  - b. If you want to create monthly restore points for archival purposes, select the **Keep monthly full backups for** check box. Then specify the number of months during which you want to prevent restore points from being modified and deleted.

In the **Use weekly full backup for the following week of a month** list, select a week when Veeam Agent must assign the monthly GFS flag to a full restore point. A week equals to 7 calendar days; for example, the first week of May is days 1–7, and the last week of May is days 25–31.

- c. If you want to create yearly restore points for archival purposes, select the **Keep yearly full backups** for check box. Then specify the number of years during which you want to prevent restore points from being modified and deleted.

In the **Use monthly full backup for the following month** list, select a month when Veeam Agent must assign the yearly GFS flag to a full restore point.

**New Backup Job**

**Backup Resources**  
Choose a cloud backup repository. You can only select between repositories your service provider has assigned to you.

**Available cloud repositories:**

**Configure GFS**

- ☐ Keep weekly full backups for: 1 weeks  
If multiple full backups exist, use the one from: Sunday
- ☐ Keep monthly full backups for: 1 months  
Use weekly full backup for the following week of a month: First
- ☒ Keep yearly full backups for: 1 years  
Use monthly full backup for the following month: January

OK Cancel

Map backup

Configure...

Advanced...

< Previous Next > Finish Cancel

## NOTE

If you select to assign multiple types of GFS flags, the flags begin to depend on each other. For more information on this dependency, see the [Assignment of GFS Flags](#) section in the Veeam Backup & Replication User Guide.

# Step 8. Specify Advanced Backup Settings

In the **Advanced Settings** window, specify advanced settings for the backup job:

- [Backup settings](#)
- [Maintenance settings](#)
- [Storage settings](#)

## Backup Settings

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

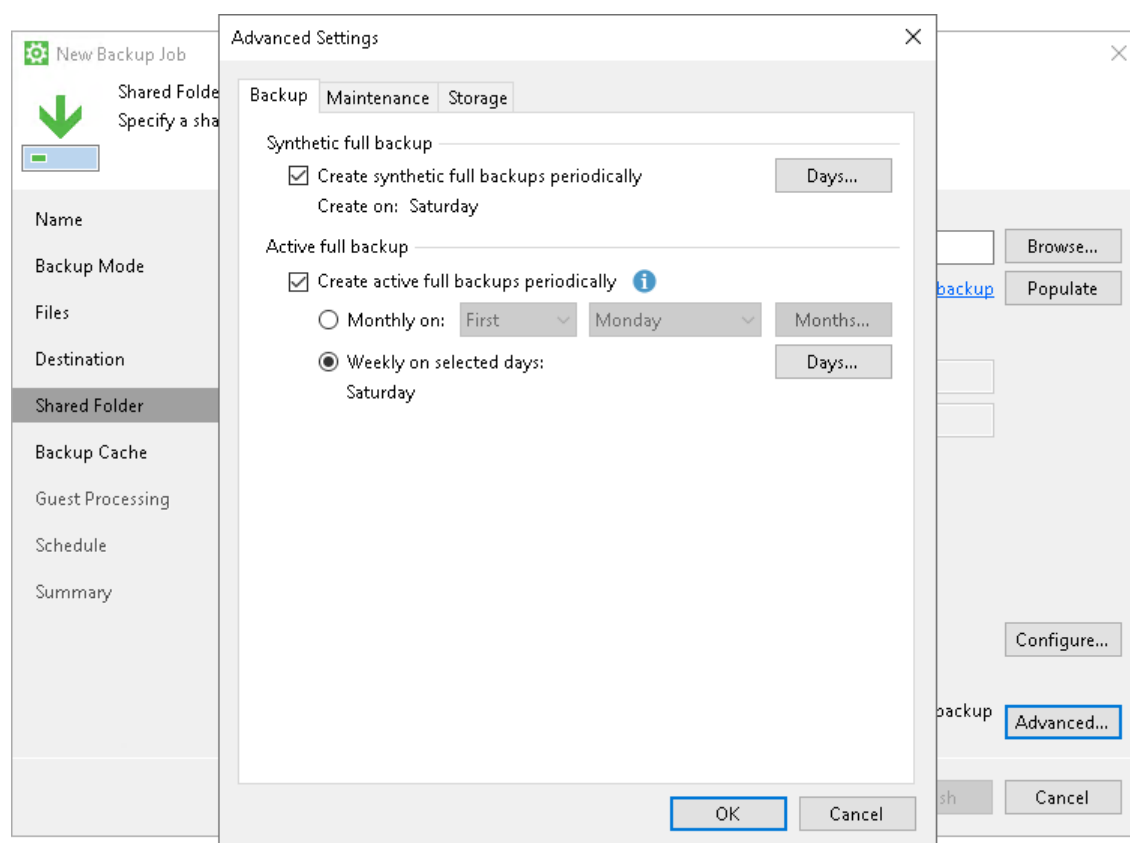
1. Click **Advanced** at one of the following steps of the wizard:
  - [Local Storage](#) — if you have selected the **Local storage** option at the [Destination](#) step of the wizard.
  - **Bucket/Container** — if you have selected the **Object storage** option at the [Destination](#) step of the wizard.
  - [Shared Folder](#) — if you have selected the **Shared folder** option at the [Destination](#) step of the wizard.
  - [Backup Repository](#) — if you have selected the **Veeam backup repository** option at the [Destination](#) step of the wizard.
  - [Backup Resources](#) — if you have selected the **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.
2. If you want to periodically create synthetic full backups, on the **Backup** tab, select the **Create synthetic full backups periodically** check box and click **Days** to schedule synthetic full backups on the necessary week days.
3. If you want to periodically create active full backups, select the **Create active full backups periodically** check box. Use the **Monthly on** or **Weekly on selected days** options to define scheduling settings.

Keep in mind that to create active and synthetic full backups on selected days, Veeam Agent must run the backup job on these days. To learn more about backup job schedule, see [Scheduling Settings](#).

## NOTE

Consider the following:

- Synthetic full backup is available only in the Workstation and Server editions of Veeam Agent for Microsoft Windows.
- Synthetic full backup is not available for backup jobs targeted at object storage.
- Before scheduling periodic full backups, you must make sure that you have scheduled the backup job to run on the enough free space on the target location. As an alternative, you can create active full backups manually when needed. For more information, see [Creating Active Full Backups](#).
- If you schedule the active full backup and synthetic full backup on the same day, Veeam Agent for Microsoft Windows will perform only active full backup. Synthetic full backup will be skipped.



## Maintenance Settings

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

1. Click **Advanced** at one of the following steps of the wizard:
  - [Local Storage](#) — if you have selected the **Local storage** option at the [Destination](#) step of the wizard.
  - [Bucket/Container](#) — if you have selected the **Object storage** option at the [Destination](#) step of the wizard.
  - [Shared Folder](#) — if you have selected the **Shared folder** option at the [Destination](#) step of the wizard.
  - [Backup Repository](#) — if you have selected the **Veeam backup repository** option at the [Destination](#) step of the wizard.

- [Backup Resources](#) — if you have selected the **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.
2. Click the **Maintenance** tab.
  3. To periodically perform a health check for the latest restore point in the backup chain, enable the **Perform backup files health check** option in the **Storage-level corruption guard** section and specify the health check schedule. For more information, see [Health Check for Backup Files](#).
  4. [For object storage target] Veeam Agent offers a special health check mechanism for object storage targets. To run the health check for object storage, enable the **Perform backup files health check** option in the **Storage-level corruption guard** section and specify the health check schedule. For more information about the health check for object storage, see [Enabling Health Check for Object Storage Target](#).
  5. [For Veeam backup repository and cloud repository targets] Select the **Remove deleted items data after** check box and specify the number of days for which you want to keep the backup created with the backup job in the target location. If Veeam Agent for Microsoft Windows does not create new restore points for the backup, the backup will remain in the target location for the period that you have specified. When this period is over, the backup will be removed from the target location.

By default, the retention period for outdated backups is 30 days. Do not set this retention period to 1 day or a similar short interval. In the opposite case, the backup job may work not as expected and remove data that you still require. For more information, see [Retention Policy for Outdated Backups](#).

#### NOTE

The **Remove deleted items data after** option is available only if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.

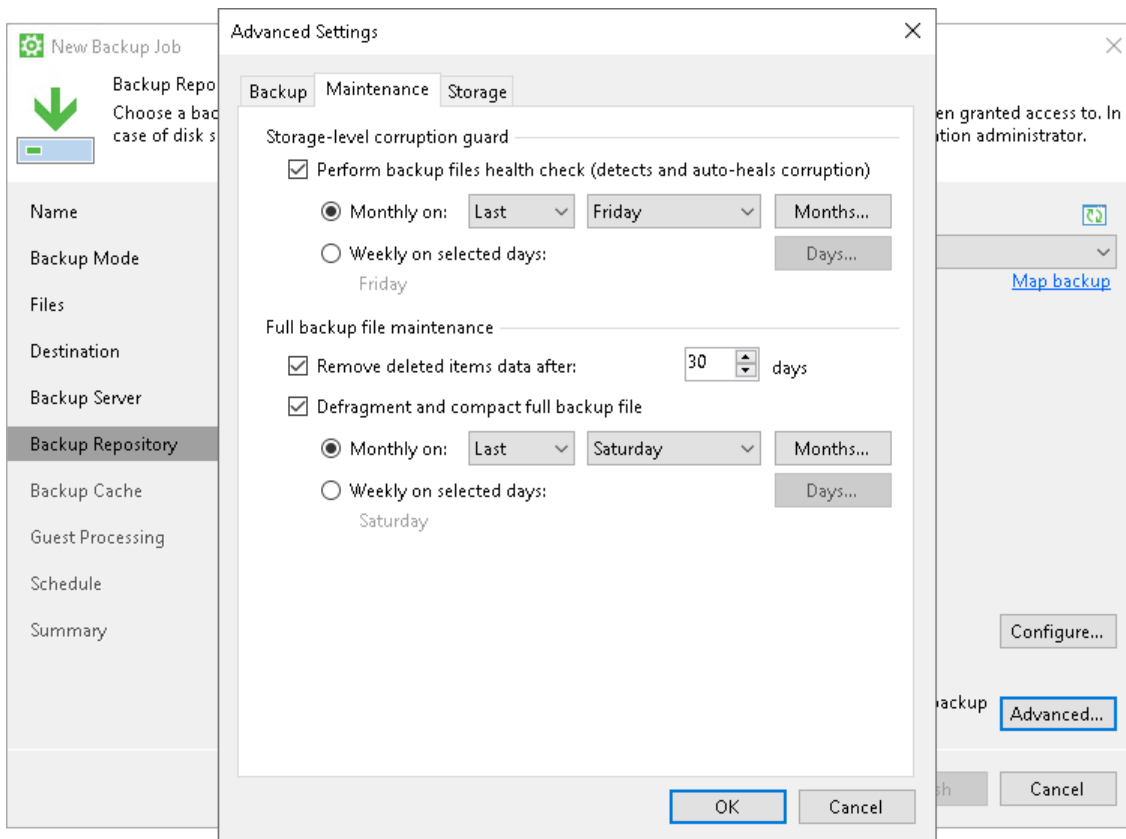
6. To periodically compact a full backup, select the **Defragment and compact full backup file** check box and specify the schedule for the compact operation.

During the compact operation, Veeam Agent for Microsoft Windows creates a new empty file and copies to this file data blocks from the full backup file. As a result, the full backup file gets defragmented, and the speed of reading from and writing to the backup file increases.

If the full backup file contains data blocks for deleted drives, Veeam Agent for Microsoft Windows will remove these data blocks. For more information, see [Compact of Full Backup File](#).

## NOTE

The **Defragment and compact full backup file** option is not available for backup jobs targeted at object storage.

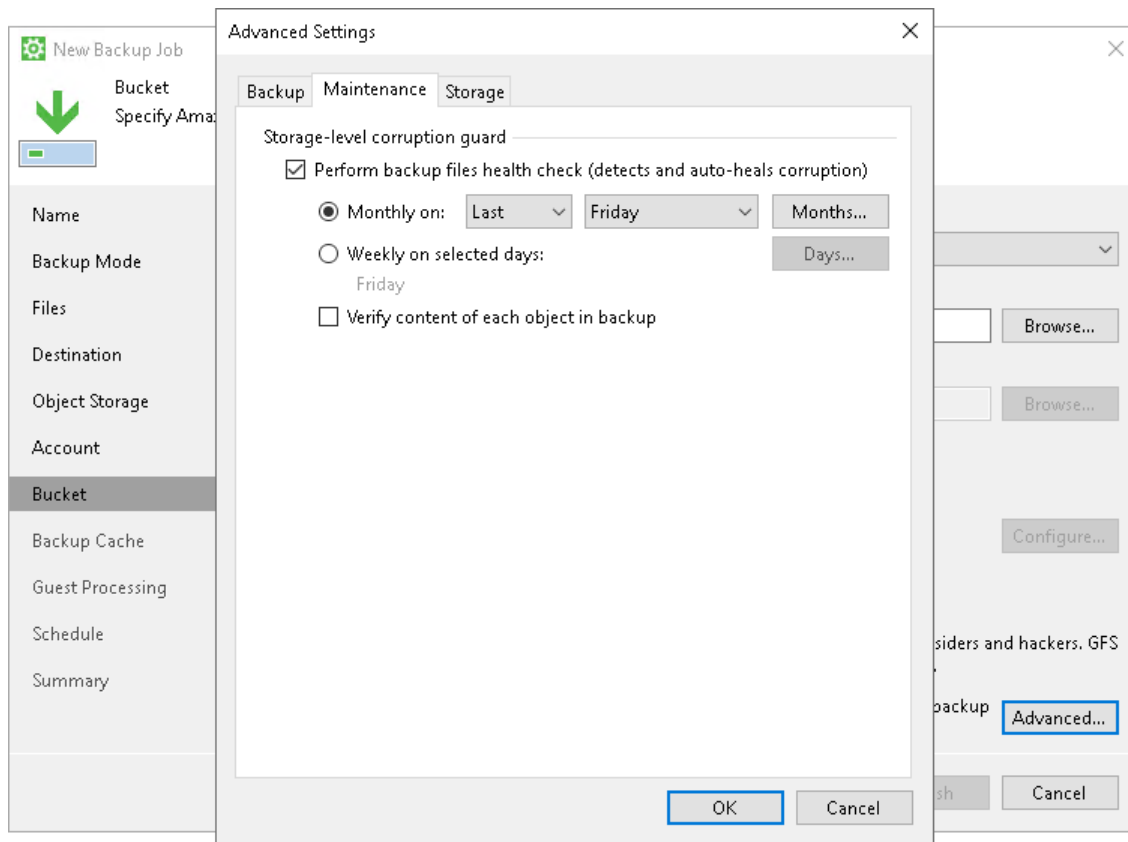


## Enabling Health Check for Object Storage Target

To periodically perform a health check for backups that reside in object storage, enable the Perform backup files health check option in the Storage-level corruption guard section and specify the health check schedule.

Veeam Agent offers a special health check mechanism for object storage targets as default. Unlike the standard health check, the health check for object storage verifies metadata for the whole backup, not just the latest restore point, and does not read data from data blocks. For more information, see [Health Check for Object Storage](#).

If necessary, you can switch from the health check for object storage to the standard health check. To do so, select the **Verify content of each object in backup** check box in the backup job settings. Keep in mind that enabling this setting may result in additional charges from your object storage provider. For information on how the standard health check works, see [Standard Health Check](#).



## Storage Settings

To specify storage settings for the backup job:

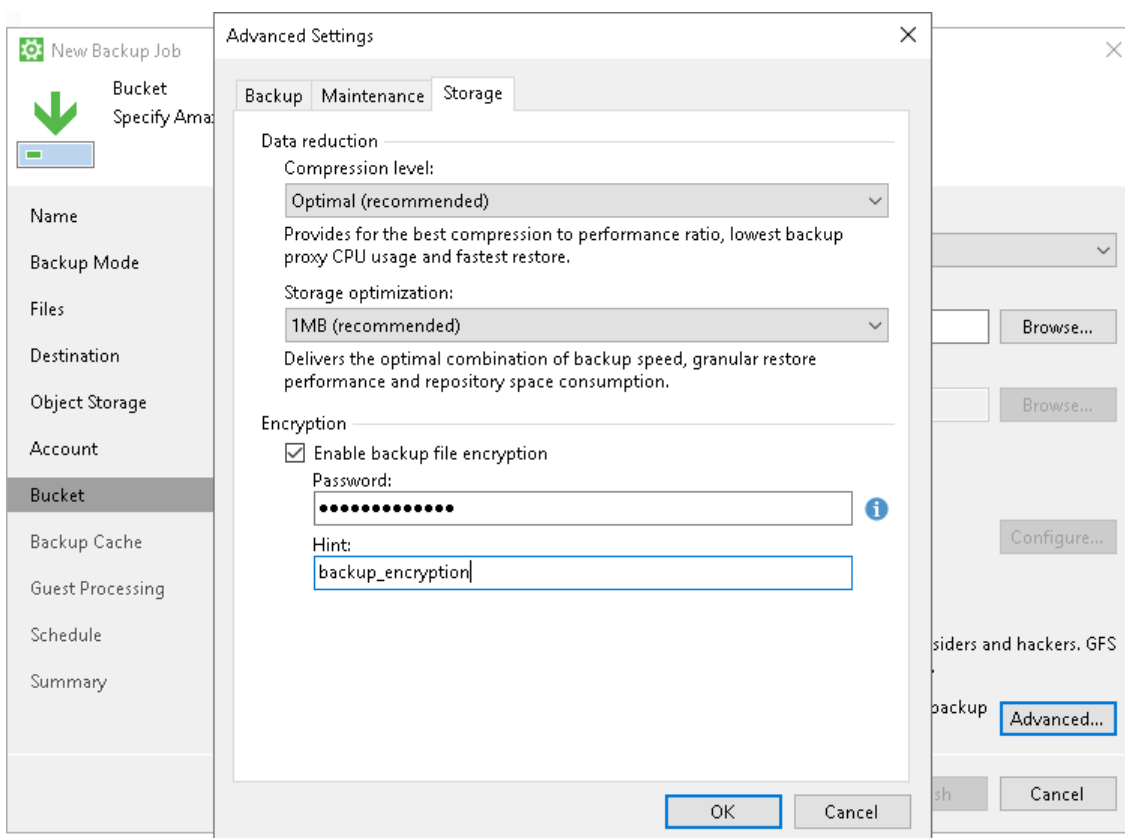
1. Click **Advanced** at one of the following steps of the wizard:
  - [Local Storage](#) — if you have selected the **Local storage** option at the [Destination](#) step of the wizard.
  - **Bucket/Container** — if you have selected the **Object storage** option at the [Destination](#) step of the wizard.
  - [Shared Folder](#) — if you have selected the **Shared folder** option at the [Destination](#) step of the wizard.
  - [Backup Repository](#) — if you have selected the **Veeam backup repository** option at the [Destination](#) step of the wizard.
  - [Backup Resources](#) — if you have selected the **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.
2. Click the **Storage** tab.
3. From the **Compression level** list, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.
4. In the **Storage optimization** section, select what size of data blocks you plan to use: *4 MB*, *1 MB*, *512 KB*, *256 KB*. Veeam Agent for Microsoft Windows will use data blocks of the chosen size to optimize the size of backup files and job performance.

5. If you want to encrypt the content of backup files, in the **Encryption** section, specify encryption settings for the backup job:
  - a. Select the **Enable backup file encryption** check box.
  - b. In the **Password** field, type a password that you want to use for encryption.
  - c. In the **Hint** field, type a hint for the password. In case you lose the password, the specified hint will help you to remember the lost password.

## NOTE

Consider the following:

- You cannot specify encryption options for the backup job if you have chosen to save backup files in a Veeam backup repository. Encryption of Veeam Agent backups stored in the backup repository are managed per repository by a backup administrator working with Veeam Backup & Replication. To learn more, see the [Data Encryption](#) and [Access Permissions](#) sections in the Veeam Backup & Replication User Guide.
- If you lose a password that was specified for encryption, you can change the password in the encryption settings. After the backup job creates a new restore point encrypted with the new password, you will be able to use this password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.
- If you enable encryption for the existing backup job that has already created one or more restore points, during the next job session, Veeam Agent for Microsoft Windows will create active full backup. The created full backup file and subsequent incremental backup files in the backup chain will be encrypted with the specified password.
- Encryption is not retroactive. If you enable encryption for the existing backup job, Veeam Agent for Microsoft Windows does not encrypt the previous backup chain created with this job.





## Step 9. Specify Backup Cache Settings

The **Backup Cache** step of the wizard is available if you have chosen to save backup files in a remote storage: in object storage, in a network shared folder, in a Veeam backup repository or in a Veeam Cloud Connect repository.

Specify backup cache settings:

1. Select the **Enable backup cache** check box.
2. In the **Location** field, specify a path to the folder on your computer in which backup files must be stored.
3. In the **Maximum size** field, specify the size for the backup cache.

When defining the size of the backup cache, assume the following:

- Each full backup file may consume about 50% of the backed-up data size.
- Each incremental backup file may consume about 10% of the backed-up data size.

### TIP

For the backup cache, you can use a dedicated removable storage device, for example, a USB key or an SD card. In this case, the backup cache will not consume disk space on the local drive of the Veeam Agent computer.

The screenshot shows the 'New Backup Job' wizard in Veeam Agent for Microsoft Windows, specifically the 'Backup Cache' step. The window title is 'New Backup Job' with a close button (X) in the top right corner. On the left, there is a sidebar with a tree view containing the following items: 'Name', 'Backup Mode', 'Files', 'Destination', 'Object Storage', 'Account', 'Bucket', 'Backup Cache' (which is highlighted with a dark background), 'Guest Processing', 'Schedule', and 'Summary'. Above the sidebar, there is a green downward arrow icon and a small progress bar. The main area of the wizard contains the following elements: A checkbox labeled 'Enable backup cache' which is checked. Below it, a text description: 'Backups remain in the cache until a connection to the backup target can be established. Once that happens, cached backups are automatically uploaded to the backup target and then deleted from the cache.' A 'Location:' label followed by a text input field containing 'E:\BackupCache' and a 'Browse...' button to its right. Below the input field, a status bar shows a blue icon and the text '89.9 GB free of 90.0 GB'. At the bottom, there is a 'Maximum size:' label followed by a numeric input field set to '10' and a unit dropdown menu set to 'GB'. At the very bottom of the window, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

# Step 10. Specify Guest Processing Settings

The **Guest Processing** step of the wizard is available in the Server edition of Veeam Agent for Microsoft Windows.

You can enable the following settings for guest OS processing:

- [Application-aware processing](#)
- [Transaction log handling for Microsoft SQL Server](#)
- [Archived log handling for Oracle databases](#)
- [SharePoint account settings](#)
- [Use of pre-freeze and post-thaw scripts](#)
- [File indexing](#)

The screenshot shows the 'New Backup Job' wizard in the 'Guest Processing' step. The window title is 'New Backup Job' with a close button (X) in the top right corner. Below the title bar, there is a green arrow icon pointing down and the text 'Guest Processing' and 'Choose guest OS processing options.' A vertical list of steps is on the left: Name, Backup Mode, Files, Destination, Object Storage, Account, Bucket, Backup Cache, Guest Processing (highlighted), Schedule, and Summary. The main area contains two checked options: 'Enable application-aware processing' and 'Enable file system indexing'. Each option has a description and a button to open a configuration dialog ('Applications...' and 'Indexing...'). At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Step	Setting	Description	Action
Guest Processing	<input checked="" type="checkbox"/> <b>Enable application-aware processing</b>	Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.	<a href="#">Applications...</a>
	<input checked="" type="checkbox"/> <b>Enable file system indexing</b>	Creates catalog of files to enable browsing, searching and 1-click restores of individual files. Indexing is optional, and is not required to perform instant file level recoveries.	<a href="#">Indexing...</a>

## Application-Aware Processing

If your computer runs VSS-aware applications, you can enable application-aware processing to create a transactionally consistent backup. The transactionally consistent backup guarantees proper recovery of applications without data loss.

To enable application-aware processing:

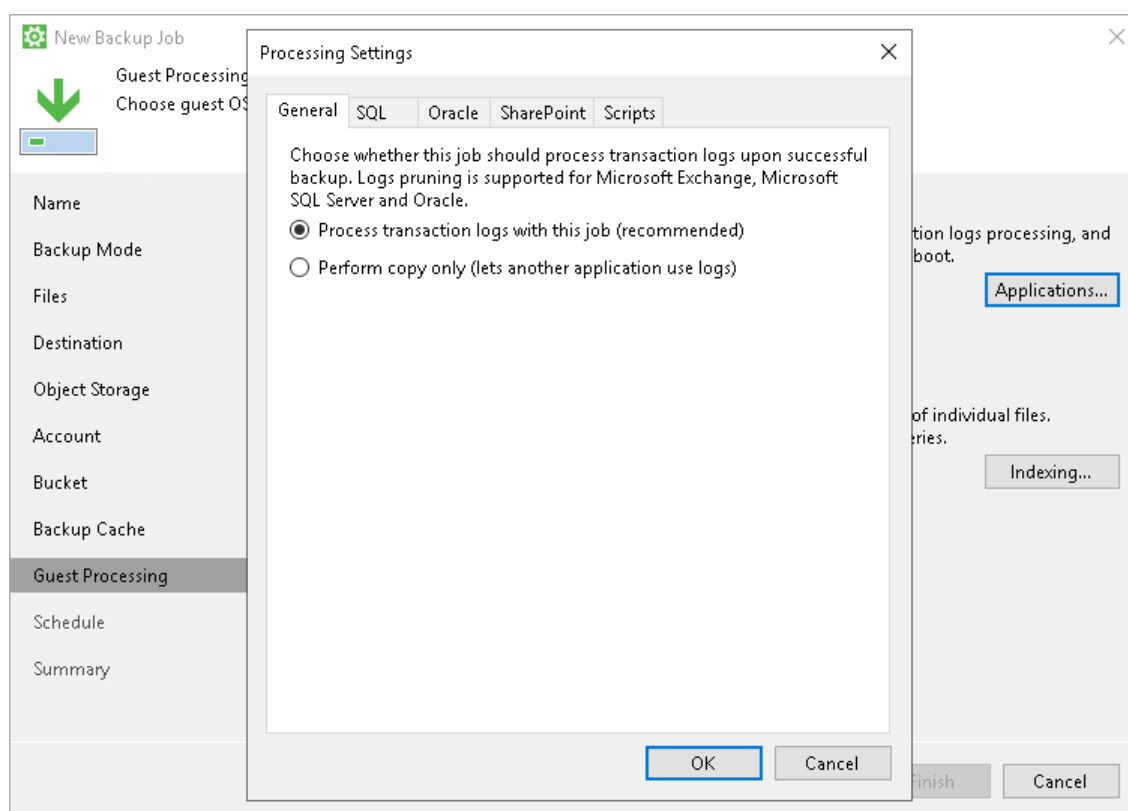
1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.

3. In the **Processing Settings** window, on the **General** tab, specify if Veeam Agent must process transaction logs [For Microsoft Exchange, Microsoft SQL and Oracle] or copy-only backups must be created.
  - a. Select **Process transaction logs with this job** if you want Veeam Agent to process transaction logs.
    - [For Microsoft Exchange] With this option selected, Veeam Agent will wait for backup to complete successfully and then trigger truncation of transaction logs. If the backup job fails, the logs will remain untouched until the next backup job session.
    - [For Microsoft SQL Server and Oracle] You will have to specify settings for transaction log handling on the **SQL** and **Oracle** tabs of the **Processing Settings** window. For more information, see [Microsoft SQL Server Transaction Log Settings](#) and [Oracle Archived Log Settings](#).
  - b. Select **Perform copy only** if you use another tool to maintain consistency of the database state. Veeam Agent will create a copy-only backup. The copy only backup preserves the chain of full/differential backup files and transaction logs. After a copy-only backup, Veeam Agent does not trigger truncation of transaction logs. For more information, see [Microsoft documentation](#).

## IMPORTANT

Consider the following:

- [For Microsoft Exchange] Veeam Agent for Microsoft Windows performs truncation of Microsoft Exchange transaction logs only if all disks that contain Microsoft Exchange databases and logs are included in a volume-level backup job.
- [For Microsoft SQL Server and Oracle] If both Microsoft SQL Server and Oracle Server are installed on one guest OS, you can enable log backup settings for one application only: either Microsoft SQL Server or Oracle.



# Microsoft SQL Server Transaction Log Settings

If you back up Microsoft SQL Server, you can specify how Veeam Agent for Microsoft Windows must process database transaction logs:

1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.
3. In the **Processing Settings** window, on the **General** tab, select **Process transaction logs with this job**.
4. In the **Processing Settings** window, click the **SQL** tab.
5. To specify a user account that Veeam Agent will use to connect to the Microsoft SQL Server, select the **Specify Windows account with sysadmin role on SQL Server** check box and enter a user name and password for the Microsoft Windows user account. This account must have the roles and permissions as specified in section [Permissions for Guest Processing](#).

## NOTE

You cannot use Microsoft SQL Server accounts (for example, the SA account) to connect to the database.

6. Specify how transaction logs must be processed. You can select one of the following options:
  - Select **Truncate logs** to truncate transaction logs after successful backup. Veeam Agent will wait for the backup to complete successfully and then truncate transaction logs. If the backup job fails, the logs will remain untouched until the next backup job session.
  - Select **Do not truncate logs** to preserve transaction logs. When the backup job completes, Veeam Agent will not truncate transaction logs.

We recommend that you enable this option only for databases with log truncation managed by a database administrator or databases that use the *Simple* recovery model. If you enable this option for databases that use the *Full* or *Bulk-logged* recovery model, transaction logs may grow large and consume all disk space. In this case, the database administrator must take care of transaction logs.

- Select **Backup logs periodically** to back up transaction logs with Veeam Agent. Veeam Agent will periodically copy transaction logs to the backup location and store them together with the image-level backup. During the backup job session, transaction logs will be truncated.

For more information, see [Microsoft SQL Server and Oracle Logs Backup and Restore](#).

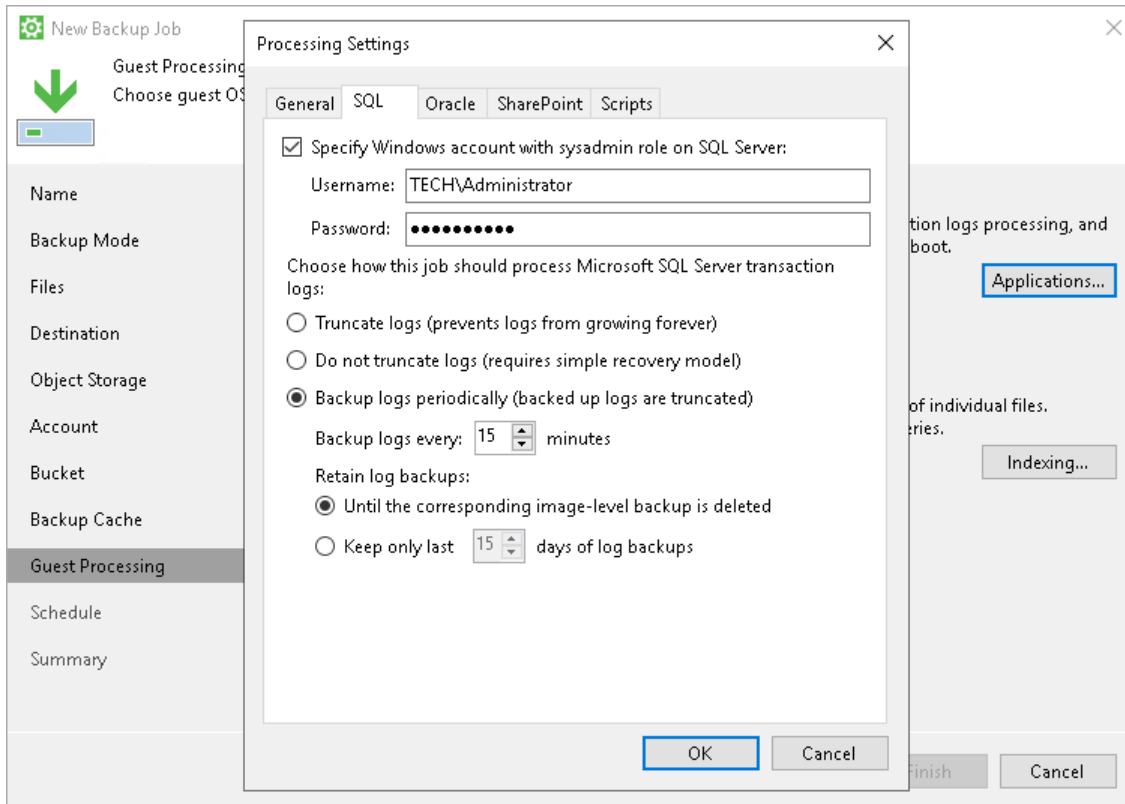
## NOTE

You can enable transaction log backup in one backup job only. If you have already configured a job that is set up to back up transaction logs, you will not be able to select the **Backup logs periodically** option in the properties of another backup job.

If you have selected to back up transaction logs with Veeam Agent, you must specify settings for transaction log backup:

1. In the **Backup logs every <N> minutes** field, specify the frequency for transaction log backup. By default, transaction logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.

2. In the **Retain log backups** section, specify retention policy for transaction logs stored in the backup location.
  - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for image-level backups and transaction log backups.
  - Select **Keep only last <N> days of log backups** to keep transaction logs for a specific number of days. By default, transaction logs are kept for 15 days. If you select this option, you must make sure that retention for transaction logs is not greater than retention for the image-level backup. For more information, see [Retention for Transaction Log Backups](#).



## Oracle Archived Log Settings

If you back up an Oracle database, you can specify how Veeam Agent for Microsoft Windows must process archived logs:

1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.
3. In the **Processing Settings** window, on the **General** tab, select **Process transaction logs with this job**.
4. In the **Processing Settings** window, click the **Oracle** tab.
5. Specify a user account that Veeam Agent for Microsoft Windows will use to connect to the Oracle database:
  - a. Select the **Specify Oracle database account with SYSDBA privileges** check box.
  - b. In the **Account** field, select what type of user account you plan to use: *Oracle account* or *Windows account*.

- c. In the **Username** and **Password** fields, enter a username and password for the account.

The specified account must have SYSDBA rights on the Oracle database.

6. In the **Archived logs** section, specify if Veeam Agent for Microsoft Windows must delete archived logs on the Oracle database:
  - Select **Do not delete archived logs** if you want Veeam Agent for Microsoft Windows to preserve archived logs. When the backup job completes, Veeam Agent for Microsoft Windows will not delete archived logs.

We recommend that you select this option for databases for which the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space. In this case, the database administrator must take care of archived logs him-/herself.
  - Select **Delete logs older than <N> hours** or **Delete logs over <N> GB** if you want Veeam Agent for Microsoft Windows to delete archived logs that are older than <N> hours or larger than <N> GB. Veeam Agent for Microsoft Windows will wait for the backup to complete successfully and then trigger archived logs deletion through Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next backup job session.

#### TIP

If you configure backup job to back up archived logs, Veeam Agent also triggers archived logs deletion after each log backup job session.

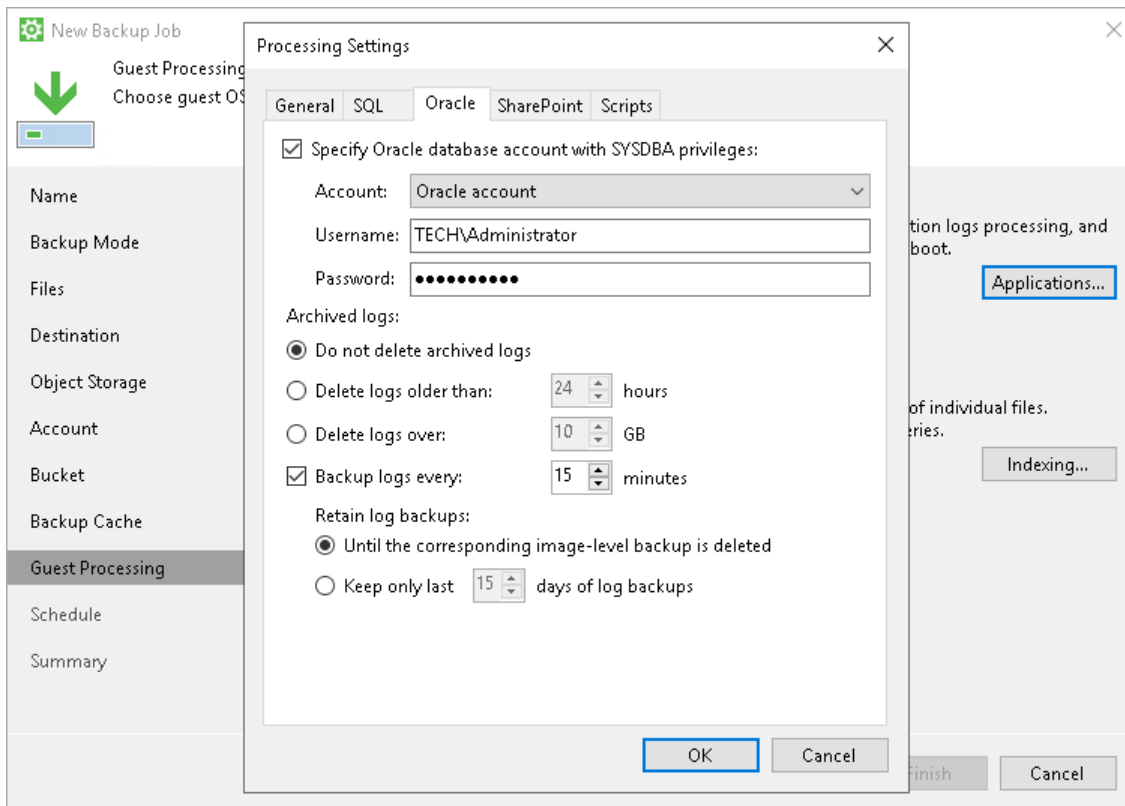
7. To back up Oracle archived logs with Veeam Agent for Microsoft Windows, select the **Backup log every <N> minutes** check box and specify the frequency for archived log backup. By default, archived logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.

#### NOTE

You can enable archived log backup in one backup job only. If you have already configured a job that is set up to back up archived logs, you will not be able to select the **Backup logs every <N> minutes** option in the properties of another backup job.

8. In the **Retain log backups** section, specify retention policy for archived logs stored in the backup location:
  - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for Veeam Agent backups and archived log backups.

- Select **Keep only last <n> days** to keep archived logs for a specific number of days. By default, archived logs are kept for 15 days. If you select this option, you must make sure that retention for archived logs is not greater than retention for the Veeam Agent backups. For more information, see [Retention for Archived Log Backups](#).

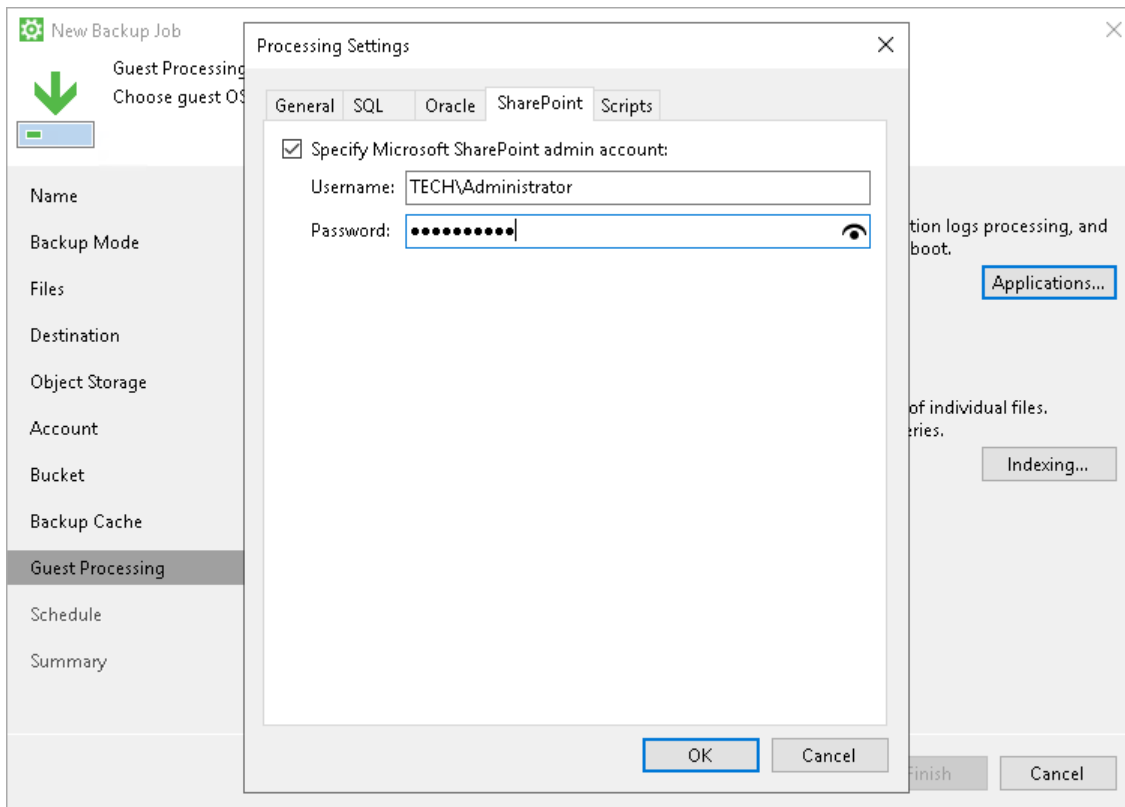


## Microsoft SharePoint Account Settings

If you back up Microsoft SharePoint, you must specify a user account that has enough permissions on the application:

1. At the **Guest Processing** step of the wizard, make sure that the **Enable application-aware processing** check box is selected.
2. Click **Applications**.
3. In the **Processing Settings** window, click the **SharePoint** tab.

4. Select the **Specify Microsoft SharePoint admin account** check box and enter a user name and password for the user account.



## Pre-Freeze and Post-Thaw Scripts

If you plan to back up data of applications that do not support VSS, you can specify what scripts Veeam Agent for Microsoft Windows must use to quiesce the OS on your computer. The pre-freeze script quiesces the file system and application data to bring the OS to a consistent state before Veeam Agent for Microsoft Windows requests the creation of a VSS snapshot. After the VSS snapshot is created, the post-thaw script brings the file system and applications to their initial state.

To specify pre-freeze and post-thaw scripts for the job:

1. At the **Guest Processing** step, click **Applications**.
2. In the **Processing Settings** window, click the **Scripts** tab.
3. In the **Script processing mode** section, specify the scenario for scripts execution:
  - Select **Require successful script execution** if you want Veeam Agent for Microsoft Windows to stop the backup process if the script fails.
  - Select **Ignore script execution failures** if you want to continue the backup process even if script errors occur.
  - Select **Disable script execution** if you do not want to run scripts.
4. In the **Scripts** section, specify paths to pre-freeze and post-thaw scripts. Scripts must reside on a local drive of the Veeam Agent computer.

Veeam Agent for Microsoft Windows supports the following types of scripts:

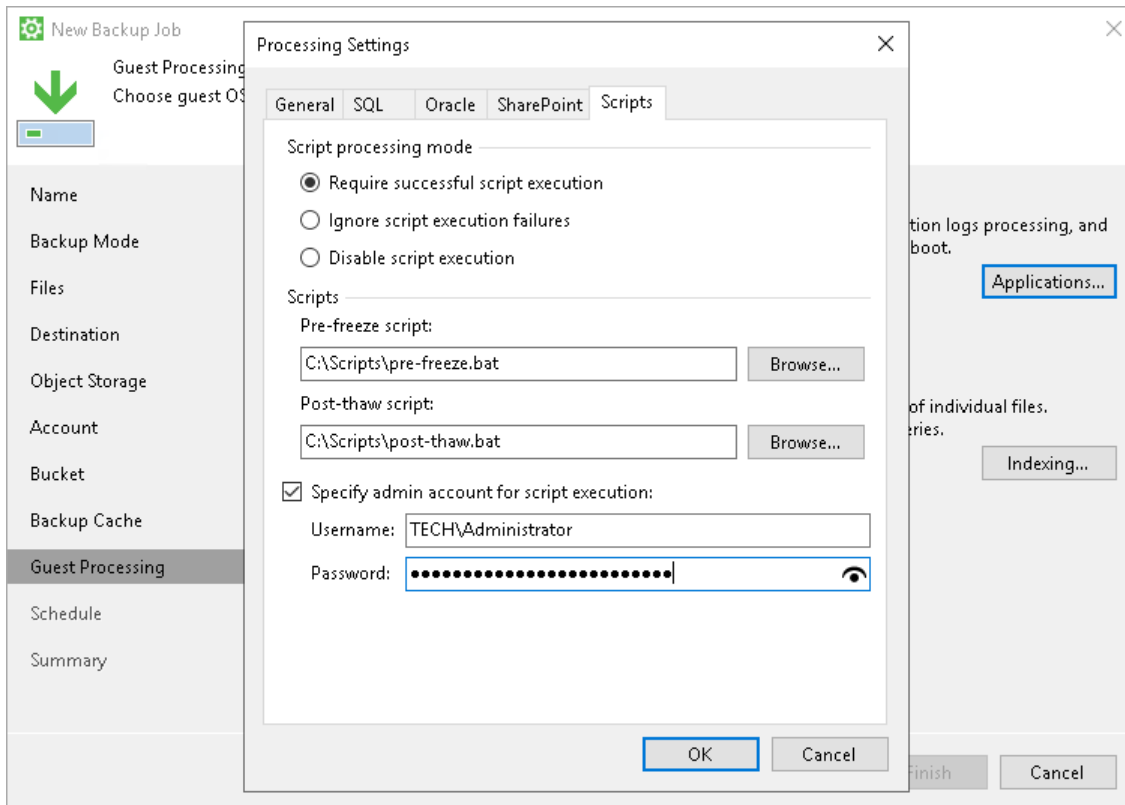
- Program files in the EXE, BAT and CMD format



- Windows script files in the JS, VBS and WSF format
- PowerShell script files in the PS1 format

You can use scripts of other formats as well, but we cannot guarantee correct processing of such scripts.

5. By default, Veeam Agent for Microsoft Windows performs guest processing activities under the Local System account. To specify a user account that Veeam Agent for Microsoft Windows will use to run pre-freeze and post-thaw scripts, select the **Specify admin account for script execution** check box and enter a user name and password for the user account.



## File Indexing

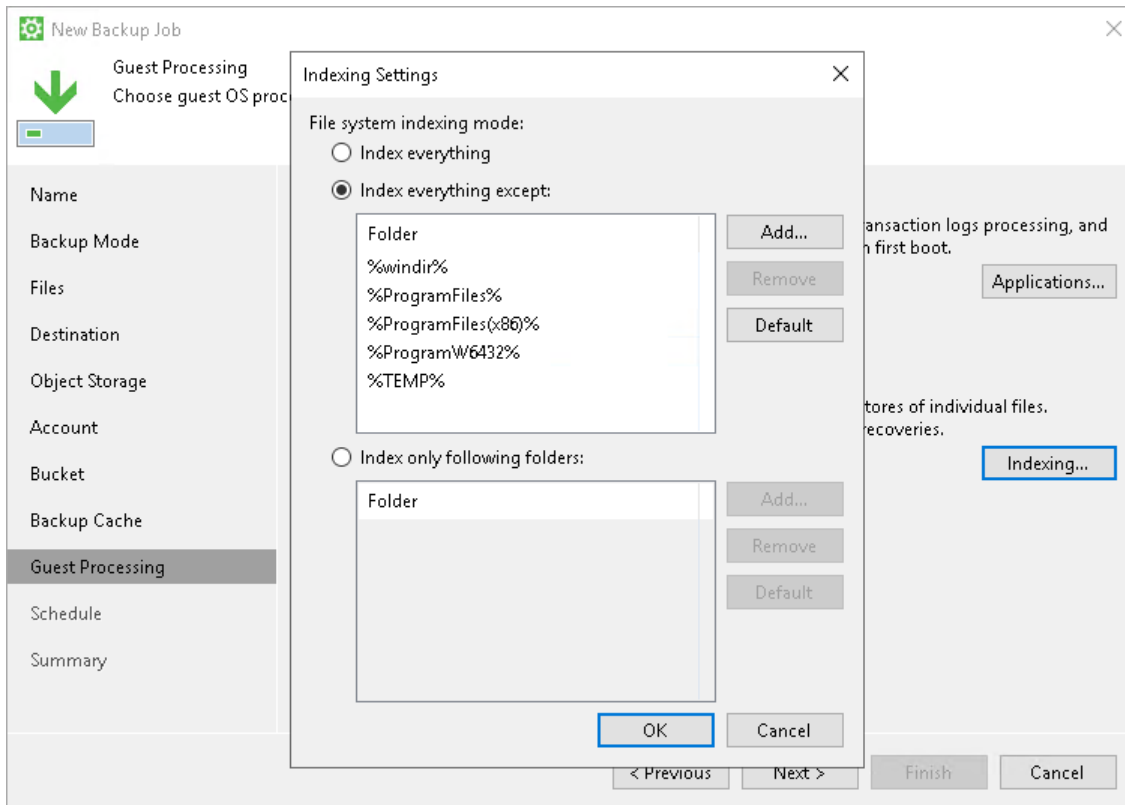
To specify file indexing options:

1. At the **Guest Processing** step of the wizard, click **Indexing**.
2. In the **Indexing Settings** window, specify the indexing scope:
  - Select **Index everything** if you want to index all files within the backup scope that you have specified at the **Backup mode** step of the wizard. Veeam Agent for Microsoft Windows will index all files that reside:
    - on your computer OS (for entire computer backup)
    - on the volumes that you have selected for backup (for volume-level backup)
    - in the directories that you have selected for backup (for file-level backup)

- Select **Index everything except** if you want to index all files on your computer OS except those defined in the list. By default, system folders are excluded from indexing. You can add or delete folders using the **Add** and **Remove** buttons on the right. You can also use system environment variables to form the list, for example: `%windir%`, `%ProgramFiles%` and `%Temp%`.

To reset the list of folders to its initial state, click **Default**.

- Select **Index only following folders** to define folders that you want to index. You can add or delete folders to index using the **Add** and **Remove** buttons on the right. You can also use system environment variables to form the list, for example: `%windir%`, `%ProgramFiles%` and `%Temp%`.



# Step 11. Specify Backup Schedule

At the **Schedule** step of the wizard, specify the schedule according to which you want to perform backup. Backup job scheduling options differ depending on the edition of Veeam Agent for Microsoft Windows:

- [Scheduling Settings in Free and Workstation Editions](#)
- [Scheduling Settings in Server Edition](#)

## Scheduling Settings in Free and Workstation Editions

At the **Schedule** step of the wizard, set up Veeam Agent for Microsoft Windows to run the backup job periodically at specific time or after specific computer events:

- [Select the time to launch the backup job.](#)
- [Select events that trigger the backup job launch.](#)

### IMPORTANT

If the power scheme on your computer does not allow using wake up timers, Veeam Agent for Microsoft Windows will ask you to change the power scheme settings. Click **Yes** to allow Veeam Agent for Microsoft Windows to wake your computer from sleep for backup.

You can manually change the power scheme settings on your computer. To do this, navigate to **Control Panel > All Control Panel Items > Power Options > Edit Plan Settings**.

## Selecting Time to Launch Backup Job

In the **Periodically** section, specify time to launch the backup job:

1. Select the **Daily at** check box and use the fields on the right to specify time and days when the backup job must start:
  - *Everyday*— select this option to start the job at specific time daily.
  - *On week-days*— select this option to start the job at specific time on week-days.
  - *On these days*— select this option to start the job at specific time on selected days.

You can leave the **Daily at** check box unchecked to configure the backup job without daily schedule. In this case, you will still be able to use the configured backup job to perform backup automatically [at specific events](#). You can also use the configured backup job to create ad-hoc incremental and standalone full backups. To learn more, see [Performing Ad-Hoc Backups](#).

2. If you have selected the *On these days* option, click the **Days** button and clear check boxes for the days when the job must not start.
3. Select the action that Veeam Agent for Microsoft Windows must perform in case your computer is powered off at the time when the scheduled backup job must start.
  - *Backup once powered on*— select this option if you want Veeam Agent for Microsoft Windows to start the scheduled backup job when you power on the computer.

- *Skip backup* – select this option if you want Veeam Agent for Microsoft Windows not to start the scheduled backup job when the computer is powered on. Veeam Agent for Microsoft Windows will perform backup at the next scheduled time.
4. If you want Veeam Agent for Microsoft Windows to perform a finalizing action after the backup job completes successfully, select the necessary action:
- *Keep running* – select this option if the computer must keep on working.
  - *Sleep* – select this option if you want Veeam Agent for Microsoft Windows to bring your computer to the standby mode.
  - *Shutdown* – select this option if you want Veeam Agent for Microsoft Windows to shut down your computer.
  - *Hibernate* – select this option if you want Veeam Agent for Microsoft Windows to bring your computer to the hibernate mode. This option is available if the hibernate mode is enabled on your computer. To learn more, see [Microsoft documentation](#).

Veeam Agent for Microsoft Windows applies this setting only to scheduled backups. If you start standalone full backup or incremental backup manually, Veeam Agent for Microsoft Windows will ignore this setting, and the computer will not be shut down or brought to the standby mode when the backup job completes.

When the backup job completes, Veeam Agent for Microsoft Windows will prompt a dialog with a countdown to the selected post-job action. You can select to proceed to the action immediately or to cancel the action. To learn more, see [Controlling Backup Post-Job Action](#).

## Selecting Events that Trigger Backup Job Launch

In the **At the following events** section, specify computer events that trigger the backup job launch:

- Select the **Lock** check box if you want to start the scheduled backup job when the user locks the computer.

### NOTE

The backup job does not always start right after the lock event is created. For example, when you put the Veeam Agent computer to sleep, the operating system creates the lock event before the entering sleep event but the backup job starts after the computer resumes from the sleep mode. The backup job does not start if the computer stayed in the sleep mode longer than 15 minutes.

- Select the **Log off** check box if you want to start the scheduled backup job when the user working with the computer performs a logout operation.
- Select the **When backup target is connected** check box if you want to start the scheduled backup job when the backup storage becomes available. With this option selected, Veeam Agent will detect state changes for networks and drives. When such change is detected, Veeam Agent will check if the target backup repository is accessible. If so, it will start the backup job.
- Select the **Eject removable storage once backup is completed** check box if you want Veeam Agent for Microsoft Windows to unmount the storage device after the backup job completes successfully. With this option selected, backup files on the removable storage will be protected from encrypting ransomware, such as CryptoLocker.

Veeam Agent applies this setting only to backup jobs triggered by computer events. In case of backup jobs scheduled to run at specific time, Veeam Agent will ignore this setting, and the storage device will not be unmounted after the backup job completes successfully.

## IMPORTANT

The *Eject removable storage once backup is completed* option does not guarantee a bulletproof protection against ransomware. To ensure your backups are safe, keep the OS up to date and regularly scan your backup repository for virus threats using modern antivirus software.

- Use the **Back up no more often than every <N> <time units>** field to restrict the frequency of backup job sessions. Specify a minutely, hourly or daily interval between the backup job sessions.

The *Back up no more often than every <N> <time units>* option is applied only to job sessions started at specific events. Daily backups are performed according to defined schedule regardless of the time interval specified for this setting.

The screenshot shows the 'New Backup Job' wizard in the 'Schedule' step. The left sidebar has a green arrow icon and a list of steps: Name, Backup Mode, Files, Destination, Shared Folder, Schedule (selected), and Summary. The main area is titled 'Schedule' with the instruction 'Specify the schedule settings to run your backup job regularly.' Below this, there are two sections: 'Periodically' and 'At the following events'. The 'Periodically' section includes a note about waking the computer from sleep, a 'Daily at' checkbox (checked) with a time dropdown set to '12:30 AM' and a frequency dropdown set to 'Everyday', and two more dropdowns for 'If computer is powered off at this time' (set to 'Skip backup') and 'Once backup is taken, computer should' (set to 'Keep running'). The 'At the following events' section has checkboxes for 'Lock', 'Log off', and 'When backup target is connected' (checked), with an information icon next to the last one. Below these is a checkbox for 'Eject removable storage once backup is completed (ransomware protection)' which is also checked. At the bottom of this section is a field 'Back up no more often than every' with a value of '2' and a unit dropdown set to 'Hour'. At the very bottom of the wizard are four buttons: '< Previous', 'Apply' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Scheduling Settings in Server Edition

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

### IMPORTANT

If the power scheme on your computer does not allow using wake up timers, Veeam Agent for Microsoft Windows will ask you to change the power scheme settings. Click **Yes** to allow Veeam Agent for Microsoft Windows to wake your computer from sleep for backup.

You can manually change the power scheme settings on your computer. To do this, navigate to **Control Panel > All Control Panel Items > Power Options > Edit Plan Settings**.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the backup job manually to create backup.
2. Define scheduling settings for the job:
  - To run the job at specific time daily, on defined week days or with specific periodicity, select **Daily at this time**. Use the fields on the right to configure the necessary schedule.
  - To run the job once a month on specific days, select **Monthly at this time**. Use the fields on the right to configure the necessary schedule.
  - To run the job repeatedly throughout a day with a specific time interval, select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

A repeatedly run job is started by the following rules:

- Veeam Agent for Microsoft Windows always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
- If you define permitted hours for the job, after the denied interval is over, Veeam Agent for Microsoft Windows will immediately start the job and then run the job by the defined schedule.

For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

- To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right. A new backup job session will start as soon as the previous backup job session finishes.
3. In the **Automatic retry** section, define whether Veeam Agent for Microsoft Windows must attempt to run the backup job again if the job fails for some reason. Enter the number of attempts to run the job and define time intervals between them. If you select continuous backup, Veeam Agent for Microsoft Windows will retry the job for the defined number of times without any time intervals between the job runs.
  4. In the **Backup window** section, define the time interval within which the backup job must complete. The backup window prevents the job from overlapping with production hours and ensures that the job does not impact performance of your server. To set up a backup window for the job:
    - a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

- b. In the **Time Periods** window, define the allowed hours and prohibited hours for backup. If the job exceeds the allowed window, it will be automatically terminated.

The screenshot shows the 'New Backup Job' dialog box with the 'Schedule' tab selected. The left sidebar contains a list of tabs: Name, Backup Mode, Files, Destination, Shared Folder, Backup Cache, Guest Processing, Schedule (highlighted), and Summary. The main area is titled 'Schedule' and contains the following settings:

- Run the job automatically:** ☒ (checked)
- Daily at this time:** ☒ (selected). Time: 12:30 AM. Frequency: Everyday. Button: Days...
- Monthly at this time:** ☐ (unselected). Time: 10:00 PM. Frequency: Fourth. Day: Saturday. Button: Months...
- Periodically every:** ☐ (unselected). Frequency: 1. Unit: Hours. Button: Schedule...
- Automatic retry:**
  - ☒ Retry failed job. Count: 3 times.
  - Wait before each retry attempt for: 10 minutes.
- Backup window:**
  - ☐ Terminate job if it exceeds allowed backup window. Button: Window...
  - If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

At the bottom of the dialog are four buttons: '< Previous', 'Apply' (highlighted with a blue border), 'Finish', and 'Cancel'.

# Step 12. Review Backup Job Settings

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

1. Review settings of the configured backup job.

In addition to backup job settings, Veeam Agent for Microsoft Windows displays a command to start the job using the command-line interface.

2. To start the job after you close the wizard, select the **Run the job when I click Finish** check box.
3. Click **Finish**.

**New Backup Job**

**Summary**  
You have successfully created the new backup job.

**Summary:**

Guest file system indexing: [enabled]

Index everything except:

- %windir%
- %ProgramFiles%
- %ProgramFiles(x86)%
- %ProgramW6432%
- %TEMP%

Script processing: [disabled]

Schedule -----

Schedule: server

Daily at [12:30 AM]

Retry failed job 3 times

Wait before each retry attempt for: 10 minutes

Command line to start the job:

"C:\Program Files\Veeam\Endpoint Backup\Veeam.Endpoint.Manager.exe" "backup" "9b821634-6cce-495d-a971-24a3f2392d74"

☒ Run the job when I click Finish

< Previous   Next >   **Finish**   Cancel



# What You Do Next

After you configure the scheduled backup job, Veeam Agent for Microsoft Windows displays a clock over its icon in the system tray. The clock identifies that your computer is protected with the scheduled backup job. Veeam Agent for Microsoft Windows will periodically start the scheduled backup job to back up selected data and add a new restore point to the backup chain in the target location.

If necessary, you can also perform the following backup operations when you need it:

- [Create a standalone full backup](#)
- [Create an incremental backup](#)
- [Create an active full backup](#)

If some of your data gets lost or corrupted, you can do the following:

- [Recover all computer volumes or specific volumes from the backup](#)
- [Recover individual files and folders from the backup](#)

# Managing Backup Jobs

After you configure the scheduled backup job, you can perform the following actions with it:

- [Edit the backup job settings](#)
- [Disable and enable the backup job](#)
- [Stop the backup job](#)
- [Remove the backup job](#)

# Editing Backup Job Settings

If you want to change settings of a backup job, you can edit it at any time. For example, you may want to edit the backup job to add a new folder to the backup scope, change the target location or job scheduling settings.

To access backup job settings, do one of the following:

- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**. Then hover over the name of the backup job whose settings you want to change, and select **Edit job**.
- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Backup > Configure backup**.

## NOTE

The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

Then edit the job settings as required. To learn more about available job settings, see [Creating Backup Jobs](#).

If you change the target location for the backup job, during the next backup job session Veeam Agent for Microsoft Windows will perform full data backup. All subsequent backup sessions will produce incremental backups – Veeam Agent for Microsoft Windows will copy only changed data to the target location and add a new incremental backup file to the backup chain.

If you change the backup scope for the backup job, during the next backup job session Veeam Agent for Microsoft Windows will create a new incremental backup. The backup will contain the following data:

- All data blocks pertaining to new data added to the backup scope.
- Changed data blocks pertaining to original data in the backup scope (data that was processed by the job at the time before you changed the backup scope).

## TIP

Full backup takes much more time than incremental backup. If you change the target location, you can copy an existing backup chain to the new location manually. In this case, the new backup job session will produce an incremental backup file and add it to the backup chain.

# Editing Encryption Settings

If you change encryption settings for the backup job, during the next backup job session Veeam Agent for Microsoft Windows will create active full backup – encrypted (if encryption was enabled) or unencrypted (if encryption was disabled). All subsequent backup sessions will produce incremental backups.

Enabling or disabling encryption does not affect backup files that were created before you have changed encryption settings.

If the backup chain contains encrypted and unencrypted backup files, you need to provide a password to restore data from any restore point in this chain. After all encrypted backup files are removed from the backup chain according to retention policy, you will be able to restore data from remaining unencrypted restore points without providing a password.

# Disabling and Enabling Scheduled Backups

You can disable scheduled backups if you do not want to run automatic backups for some period of time. For example, you may want to put backup activities on hold if you plan to perform resource consuming operations on your computer at the time when the backup job is scheduled. After the operations are completed, you can enable scheduled backups again.

The disabling option is applicable to backup job sessions started upon schedule. You can create standalone full backups and perform ad-hoc incremental backup even if the scheduled backups are disabled.

The disabling option is applicable to all backup jobs configured in Veeam Agent for Microsoft Windows. If you enable this option, all jobs that you configured will not start automatically upon the defined schedule. If you want to prevent a specific job from starting automatically, disable scheduling options in the properties of this job.

If you configured a backup job that is set up to create database log backups, after you disable scheduled backups, the database log backup job will be disabled, too.

The disabling option does not put on hold the backup cache synchronization process. If Veeam Agent has created one or more backup files in the backup cache, and then the backup target becomes available, Veeam Agent will immediately upload backup files to the target location.

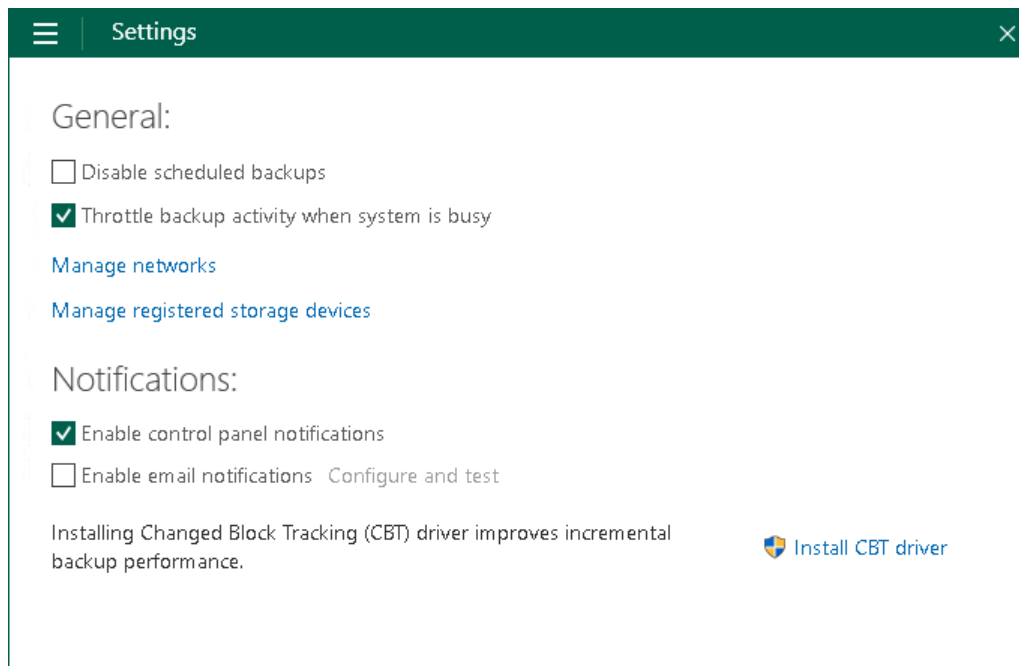
To disable scheduled backups:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Select the **Disable scheduled backups** check box.

To enable scheduled backups:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. From the main menu, select **Settings**.

3. Clear the **Disable scheduled backups** check box.




# Stopping Backup Job

You can stop the running backup job before the job session completes, for example, if the backup process is about to take long, and you do not want the job to produce workload on the production environment during business hours.

When you stop a backup job, the job session will finish immediately. Veeam Agent will not produce a new restore point during the backup job session.

To stop a backup job:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. In the main menu, do either of the following:
  - Hover over the name of the running backup job that you want to stop and select **Stop job**.
  - Hover over the name of the running backup job that you want to stop and select **Open**. Then click **Cancel** in the job statistics window.



# Removing Backup Job

You can remove a backup job configured in Veeam Agent for Microsoft Windows. When you remove a backup job, Veeam Agent for Microsoft Windows also removes information about backup job sessions from the Veeam Agent database. After you remove the job, information about sessions of this job is not displayed in the **Status** view of the control panel any longer.

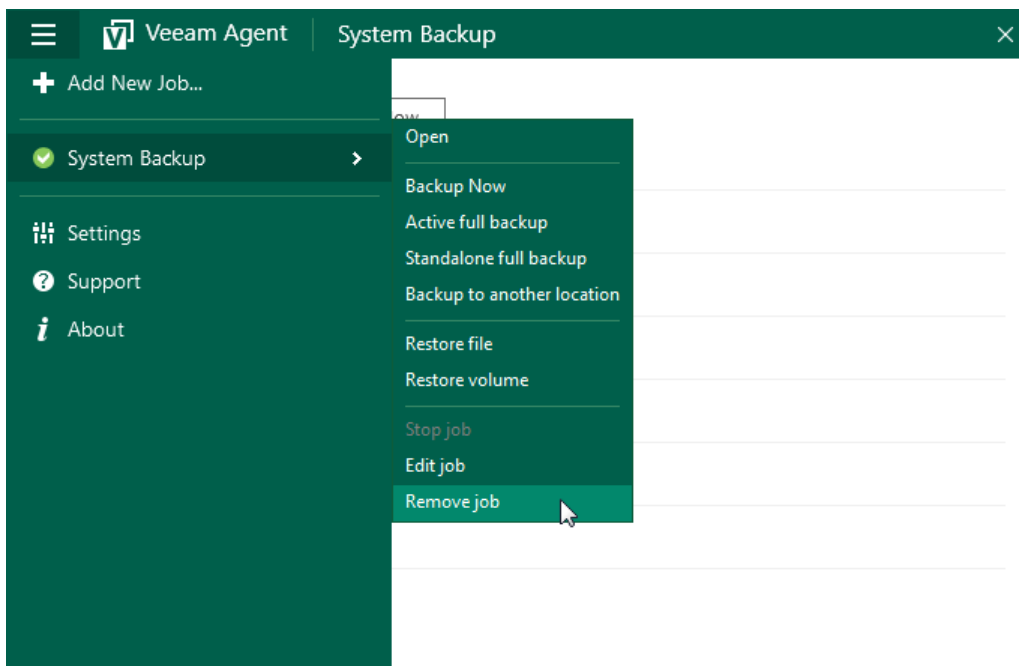
## NOTE

You cannot remove a backup job that is currently running. You must stop the job first.

Backup files created by the deleted backup job remain intact in the target location.

To remove a backup job:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to remove, and select **Remove job**.
3. In the displayed window, click **Yes**.



# Controlling Backup Post-Job Action

In the Free and Workstation product editions, you can set up Veeam Agent to perform a finalizing action after the backup job completes successfully:

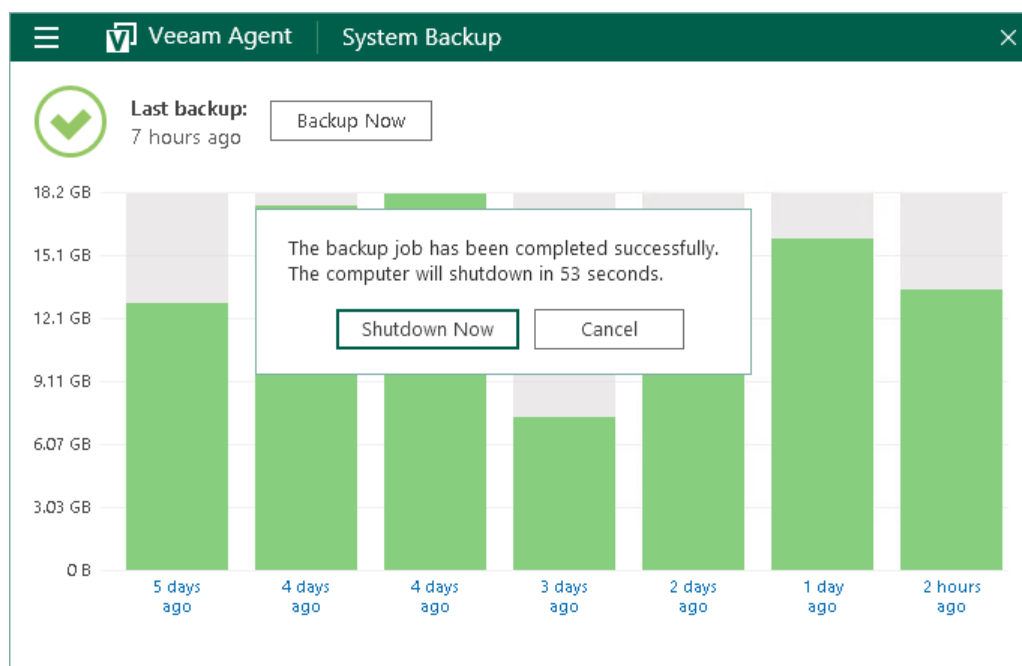
- *Sleep* – bring your computer to the standby mode.
- *Hibernate* – bring your computer to the hibernate mode.
- *Shutdown* – shut down your computer.

To learn more, see [Specify Backup Schedule](#).

When the backup job completes, Veeam Agent opens control panel and prompts a dialog with a countdown to the specified action. Timeout between the backup job completion and the backup post-job action is 60 seconds.

- To proceed to the backup post-job action immediately, click **Sleep/Hibernate/Shutdown Now**.
- To cancel the action (for example, if you want to continue working or to save your data before turning off the computer), click **Cancel**.

If you do not select any option, Veeam Agent will perform the specified action when timeout expires.





# Performing Ad-Hoc Backups

In addition to running scheduled backups, you can create ad-hoc backups of your data at any time you need. Veeam Agent for Microsoft Windows lets you perform the following types of ad-hoc backups:

- [Incremental ad-hoc backup](#)
- [Active full backup](#)
- [Standalone full backup](#)
- [Backup to another location](#)

# Creating Incremental Backups

You can create an ad-hoc incremental backup of your data in addition to the scheduled backup. Ad-hoc incremental backup may be necessary if you want to capture your data at a specific point in time, for example, before you install new software on your computer. Ad-hoc incremental backup lets you produce an additional restore point in the backup chain at any time and does not require you to reconfigure the scheduling settings in the backup job.

Before you perform ad-hoc incremental backup, check the following prerequisites:

- The backup job that you want to use to perform ad-hoc incremental backup must be configured and successfully run at least once.
- You cannot perform ad-hoc incremental backup if a backup task of any type is currently running. These include a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.

To perform ad-hoc incremental backup:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. Do either of the following:
  - If the Veeam Agent control panel currently displays statistics of the job that you want to use to perform ad-hoc incremental backup, in the job statistics window, click **Backup Now**.
  - If you want to perform ad-hoc incremental backup using another backup job configured in Veeam Agent, in the main menu, hover over the name of the necessary job and select **Backup now**.

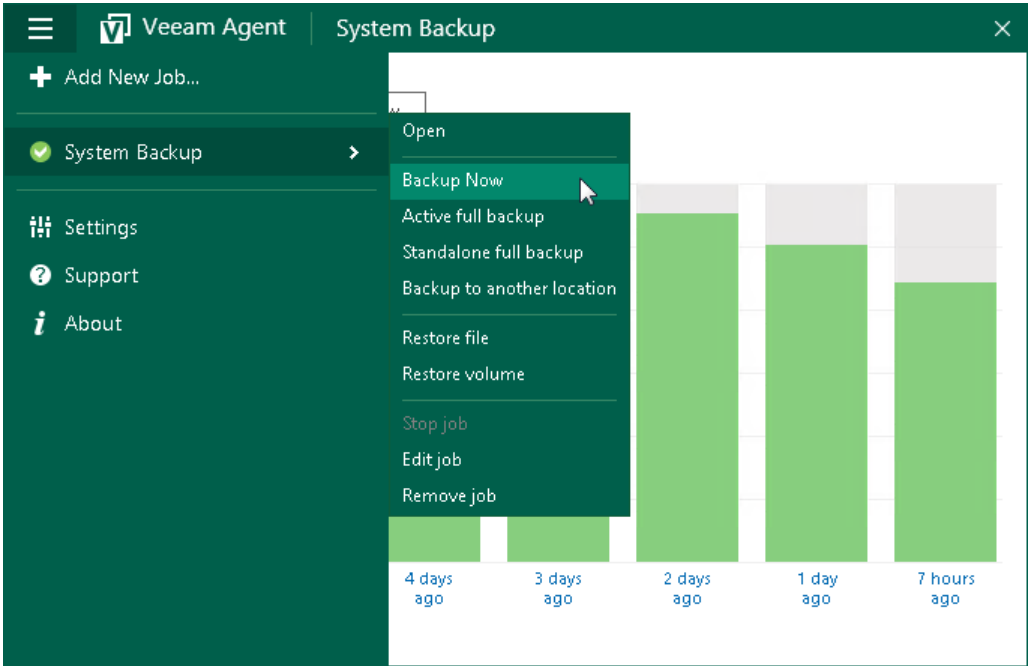
If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the ad-hoc incremental backup task from the system tray menu:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select **Backup > Backup now**.

## NOTE

The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

Veeam Agent for Microsoft Windows will perform incremental backup using settings of the scheduled backup job and add a new restore point to the backup chain in the target location.



# Creating Active Full Backups

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

## IMPORTANT

If you have a file-level backup job configured in Veeam Agent and you need to extend the volume where backed-up files reside, we strongly recommend to create an active full backup after the volume is extended. Otherwise, Veeam Agent may skip files during the job run even if these files are added to the backup scope.

Before you create an active full backup, check the following prerequisites:

- The backup job must be configured.
- You cannot create an active full backup if a backup task of any type is currently running. This includes a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.
- A user account under which you start the **Active full backup** operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

To perform active full backup:

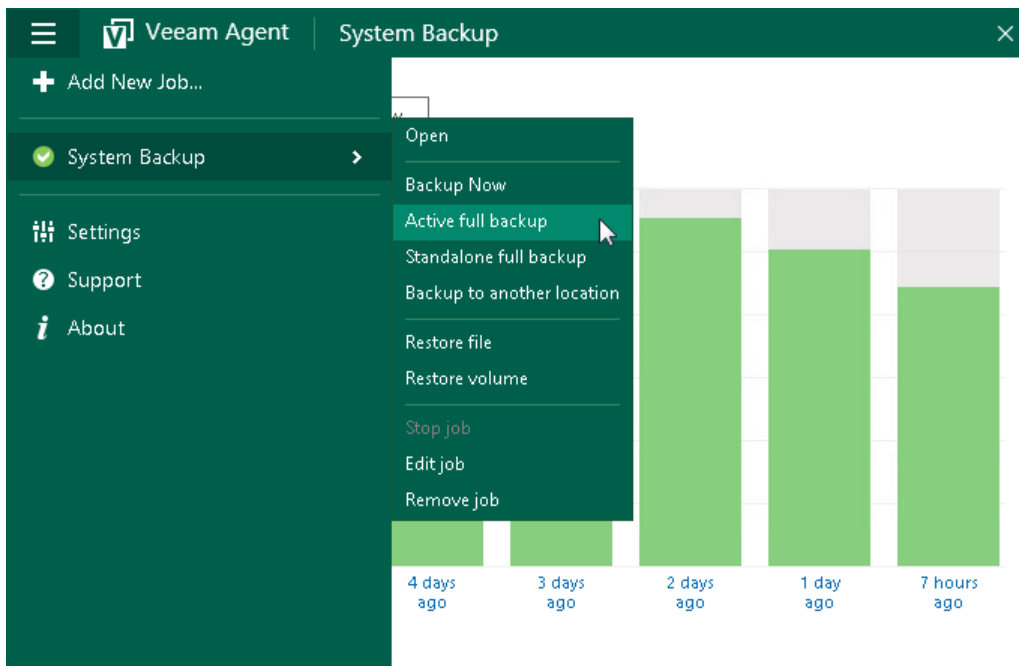
1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to use to create an active full backup, and select **Active full backup**. Veeam Agent for Microsoft Windows will create a full backup file using settings of the backup job and add this backup file to the backup chain.

If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the active full backup task from the system tray menu:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select **Backup > Active full backup**.

## NOTE

The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.



# Creating Standalone Full Backups

If you want to back up your data at a specific point in time, you can create a standalone full backup. The standalone full backup is independent: it is not followed by subsequent incremental backups and is not removed by retention. You can use the standalone full backup to create an additional restore point from which you can recover your data.

Before you create a standalone full backup, check the following prerequisites and limitations:

- The backup job must be configured.
- A user account under which you start the Standalone full backup operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.
- You cannot create a standalone full backup if a backup task of any type is currently running. This includes a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.
- You cannot create a standalone full backup if you have chosen to store backup files in one of the following locations:
  - Veeam backup repository
  - Veeam Cloud Connect repository
  - Object storage

If you want to create a full backup file not associated with the backup chain, you can perform standalone full backup to another location. To learn more, see [Performing Backup to Another Location](#).

To create a standalone full backup:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to use to create a standalone full backup, and select **Standalone full backup**. Veeam Agent for Microsoft Windows will create a full backup file using settings of the backup job. The resulting full backup file will be saved in the target location specified in the job settings, and placed to a separate folder. The folder is named in the following way: `<BackupJobName>.adhoc.<DateandTime>`.

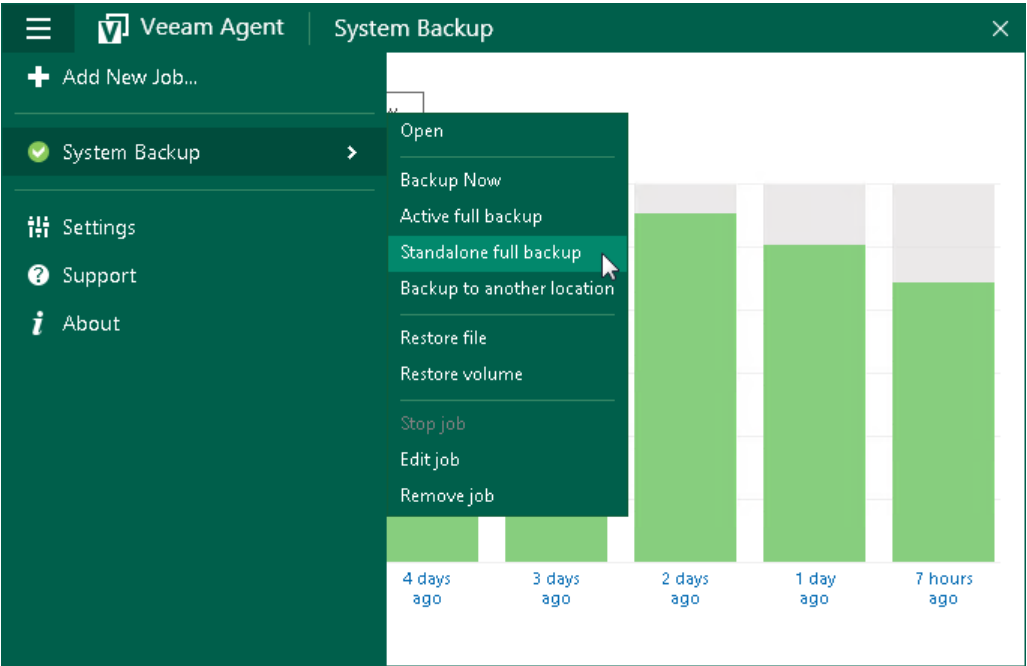
If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the standalone full backup task from the system tray menu:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select **Backup > Standalone full backup**.

## NOTE

The **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent for Microsoft Windows.

You can also create a standalone full backup in a location that is not specified in the backup job settings. To learn more, see [Performing Backup to Another Location](#).



# Performing Backup to Another Location

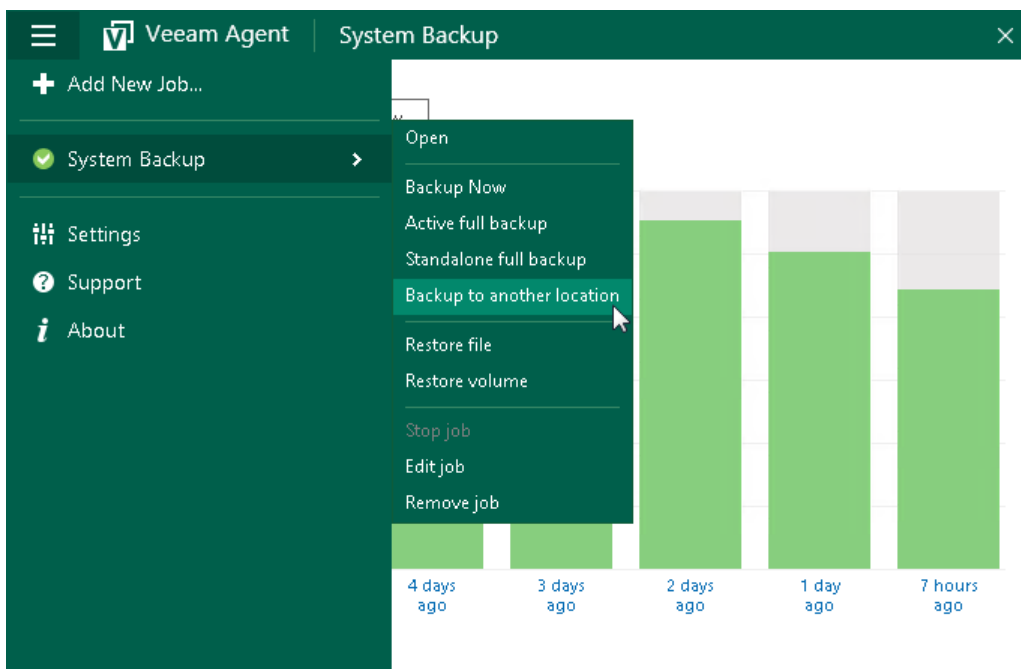
You can create a standalone full backup in a separate location that is not specified as a target location in the backup job settings. Performing backup to another location is similar to creating regular standalone full backups. The main difference is that you must manually select a target location in which Veeam Agent for Microsoft Windows will save the backup file.

Before you perform backup to another location, check the following prerequisites:

- The backup job must be configured.
- You cannot perform backup to another location if a backup task of any type is currently running. This includes a scheduled backup, standalone full backup, active full backup or ad-hoc incremental backup.
- A user account under which you start the **Backup to another location** operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

To perform backup to another location:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click it and select **Control Panel**.
2. In the main menu, hover over the name of the backup job that you want to use to perform backup, and select **Backup to another location**.



## NOTE

If only one job is configured in Veeam Agent for Microsoft Windows, you can also start the standalone full backup task from the system tray menu:

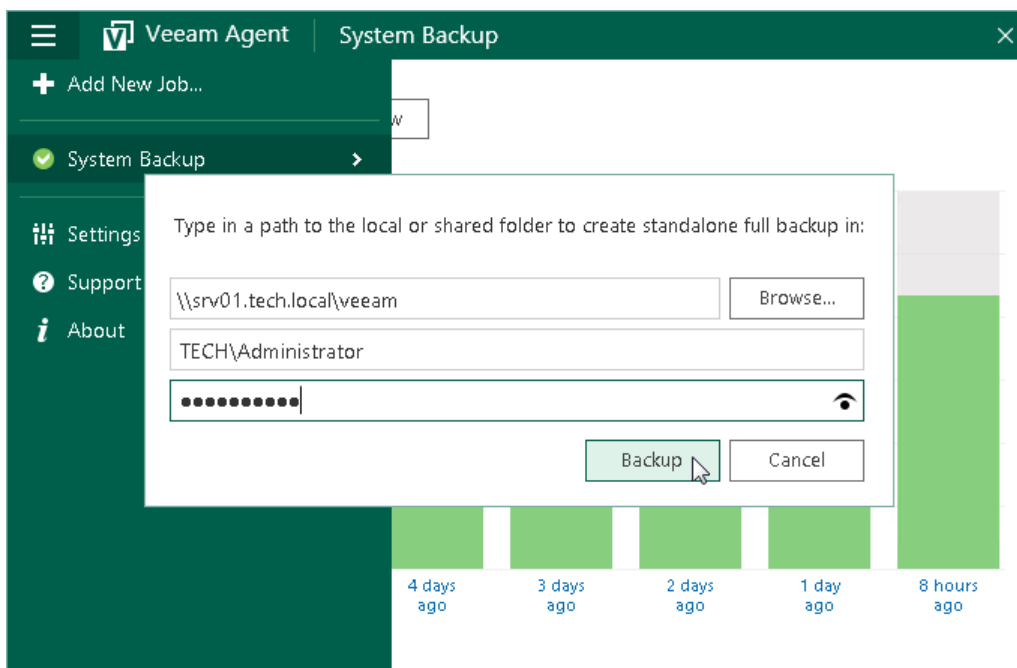
1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
2. Select Backup > Backup to another location.

Note that the **Backup** option is not available in the system tray menu if multiple backup jobs are configured in Veeam Agent.



3. In the standalone full backup dialog window, specify the target location for the backup file:
  - If you want to save the backup file in a folder on a local drive or a removable storage device, click **Browse** and select the necessary folder or type a path to the folder where backup file must be saved.
  - If you want to save the backup file in a network shared folder, type a UNC name of the network shared folder. Keep in mind that the UNC name always starts with two back slashes (\\). If the network shared folder requires authentication, specify a user name and password of the account that has access permissions on this shared folder. The user name must be specified in the *DOMAIN\UserName* format.
4. Click **Backup**.

Veeam Agent for Microsoft Windows will create a full backup file using settings of the selected backup job. The resulting full backup file will be saved in a separate folder in the specified location. The folder is named in the following way: <BackupJobName>.adhoc.<DateandTime>.



# Performing Backup with Command Line Interface

In addition to running scheduled backup jobs and performing ad-hoc backups from the Veeam Agent tray agent or control panel, you can create backups with the command line interface. For example, you can use commands for running a backup job in custom scripts to set up more detailed backup schedule than the daily schedule configured with the control panel.

You can run a backup job from the command line interface to create the following types of backups:

- Full or incremental backup (regular restore point in the backup chain)
- Active full backup
- Standalone full backup
- Backup to another location

## IMPORTANT

If you create backups with the command line interface, Veeam Agent will ignore the following network connection restrictions:

- Disabled backup over metered connections
- Disabled backup over VPN connections
- Backup over selected wireless networks

To learn more about network usage settings, see [Restricting Network Connections Usage](#).

Before you create a backup from the command line interface, check the following prerequisites:

- The backup job must be configured.
- You cannot run a backup job from the command line interface if a backup task for this job is currently running. This includes a scheduled backup, standalone full backup or ad-hoc incremental backup.

If more than one backup job is configured, you can run multiple backup jobs at the same time.

- To create a standalone full backup or backup to another location, you must run the command line interface with administrative privileges.

- If more than one backup job is configured, you must specify the job ID in the "00000000-0000-0000-0000-000000000000" format. You can find the job ID at the **Summary** step of the backup job wizard.

**Edit Backup Job [System Backup]**

Summary  
You have successfully edited the backup job.

Name	Summary:
Backup Mode	Compression level: Optimal (recommended)
Files	Storage optimization: 1MB (recommended)
Destination	Storage encryption: disabled
Local Storage	Verify the backup monthly on the last Friday of January, February, March, April, May, June, July, August, September, October, November, December
Schedule	GFS: disabled
Summary	Schedule ----- Schedule: workstation Daily at [3:42 PM] If computer is shutdown at this time: [skip backup] Once backup is taken, computer should: [shutdown] Create active full backups weekly on Saturday
	Command line to start the job: "C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" "backup" "979cce02-1028-4eea-860e-86eea237d646"

☐ Run the job when I click Finish

< Previous   Next >   **Finish**   Cancel

## Creating Backups

To perform backup, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /backup
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /backup "<ID>"
```

where <ID> – unique job ID generated by Veeam Agent for Microsoft Windows.

## Creating Active Full Backups

To create an active full backup, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /activefull
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /activefull "<ID>"
```

where <ID> – unique job ID generated by Veeam Agent for Microsoft Windows.

## Creating Standalone Full Backups

To create a standalone full backup, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone "<ID>"
```

where <ID> – unique job ID generated by Veeam Agent for Microsoft Windows.

## Performing Backup to Another Location

To create a standalone full backup to a different location than a location that is specified in the backup job settings, use a command with the following syntax:

- If one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone <location>
```

- If more than one backup job is configured:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /standalone <location> "<ID>"
```

where:

- <location> – path to a folder in which the backup should be created.
- <ID> – unique job ID generated by Veeam Agent for Microsoft Windows.

### IMPORTANT

You can specify a network shared folder as a target location for standalone full backup only if read and write permissions on this folder are granted to *Everyone* or to the *LocalSystem* account of the Veeam Agent computer. You cannot specify credentials to access the network shared folder in the command.

# Monitoring Backup Job Status

When you start a backup job from the command line interface, it runs automatically in the background. You can view information about the backup job session or the created restore point in the control panel. To learn more, see [Viewing Statistics in Control Panel](#).

You can also use the last exit code to verify if the backup job has completed successfully. To check the last exit code, use the `%ERRORLEVEL%` variable in `cmd.exe`.

Veeam Agent for Microsoft Windows can provide the following exit codes:

- 0 – backup successfully created
- -1 – backup job failed to start or completed with error
- 5 – backup job is currently running and cannot be started from the command line interface

# Deleting Backups

Backup files created with Veeam Agent for Microsoft Windows are removed automatically according to the retention policy settings.

If necessary, you can remove the backup files manually with a file manager, for example, Microsoft Windows Explorer. In this case, delete the whole backup chain from the target location. After you remove the whole backup chain from the target location, during the next backup job session, Veeam Agent for Microsoft Windows will produce a new full backup. All subsequent backups will be incremental.

## IMPORTANT

Do not delete individual backup files from the backup chain. If you delete a full or incremental backup file, the chain will be broken. In this case, Veeam Agent for Microsoft Windows may fail to perform the scheduled backup next time.

## Deleting Backups with Command Line Interface

You can use the command line interface to delete Veeam Agent backups from a Veeam Cloud Connect repository or object storage.

When you delete a backup from a Veeam Cloud Connect repository, Veeam Agent for Microsoft Windows deletes actual backup files from the repository and removes records about the backup from the Veeam Backup & Replication database on the SP backup server. After information about the backup is removed from the SP backup server, Veeam Backup & Replication removes this information from its database and console on the tenant backup server, too.

Before you delete a backup, check the following prerequisites:

- To perform the delete backup operation, you must run the command line interface with administrative privileges.
- The delete backup operation cannot be performed if a backup or a restore task is currently running.
- Check the following backup job settings:
  - The target location from which you want to delete a backup must be specified as a target location for backup files in the backup job settings.
  - Credentials of the user account whose backup you want to delete must be specified in the backup job settings.

To learn more, see [Veeam Cloud Connect Repository Settings](#) and [Object Storage Settings](#).

To delete a Veeam Agent backup, use a command with the following syntax:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" /deletebackup
```

## TIP

If more than one backup job is configured, Veeam Agent for Microsoft Windows will provide a list of backup jobs. To delete a backup, type the number that is displayed next to the backup job whose backup you want to delete.

Veeam Agent for Microsoft Windows can provide the following exit codes:

- 0 – backup job was successfully deleted
- -1 – the delete backup job operation failed

# Managing Backup Cache

You can perform the following operations with the backup cache:

- [Monitor backup cache activity](#)
- [Pause backup cache synchronization](#)
- [Delete restore points from the backup cache](#)



# Monitoring Backup Cache Activity

You can use the Veeam Agent control panel to view information about backup cache activity. In the **Backup Cache** view, Veeam Agent for Microsoft Windows displays a list of restore points that were created in the backup cache, their status and size of the resulting backup file. For restore points that are being uploaded or have been already uploaded to the target location, Veeam Agent for Microsoft Windows also displays the upload speed.

## Viewing Restore Points in Backup Cache

To view information about backup cache activity, open the **Backup Cache** view in one of the following ways:

- If one backup job is configured:

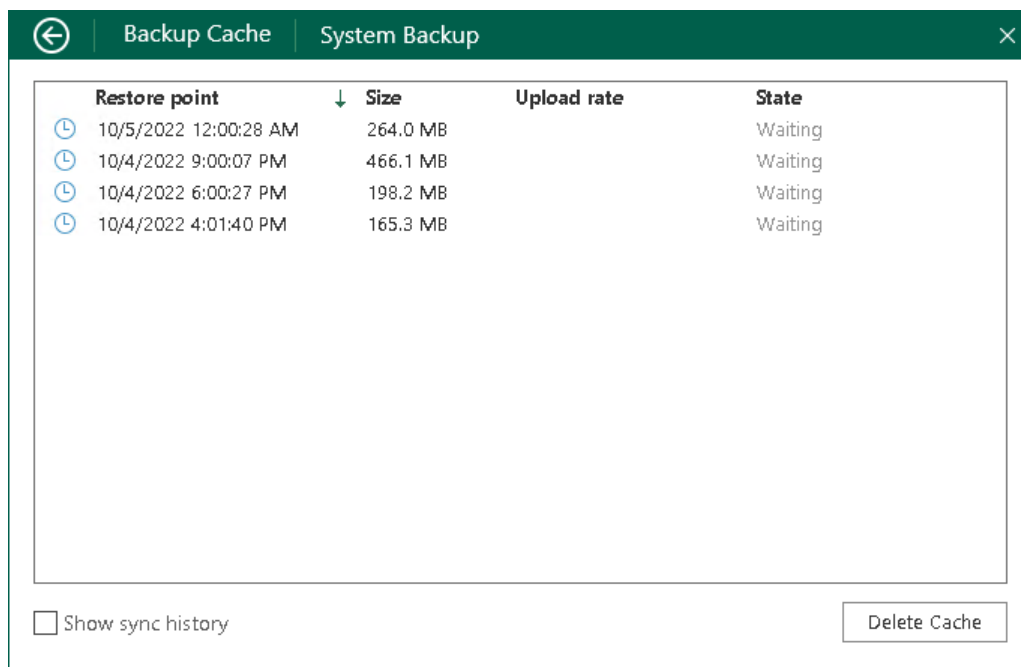
Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **View cache**. The Veeam Agent control panel will open, and you will pass immediately to the **Backup Cache** view.

### NOTE

The **View cache** option is not available in the system tray menu if more than one backup job is configured in Veeam Agent for Microsoft Windows.

- If more than one backup job is configured:

Double-click the Veeam Agent for Microsoft Windows icon, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**. In the main menu, hover over the name of the job that created restore points in the backup cache and select **Open backup cache**.

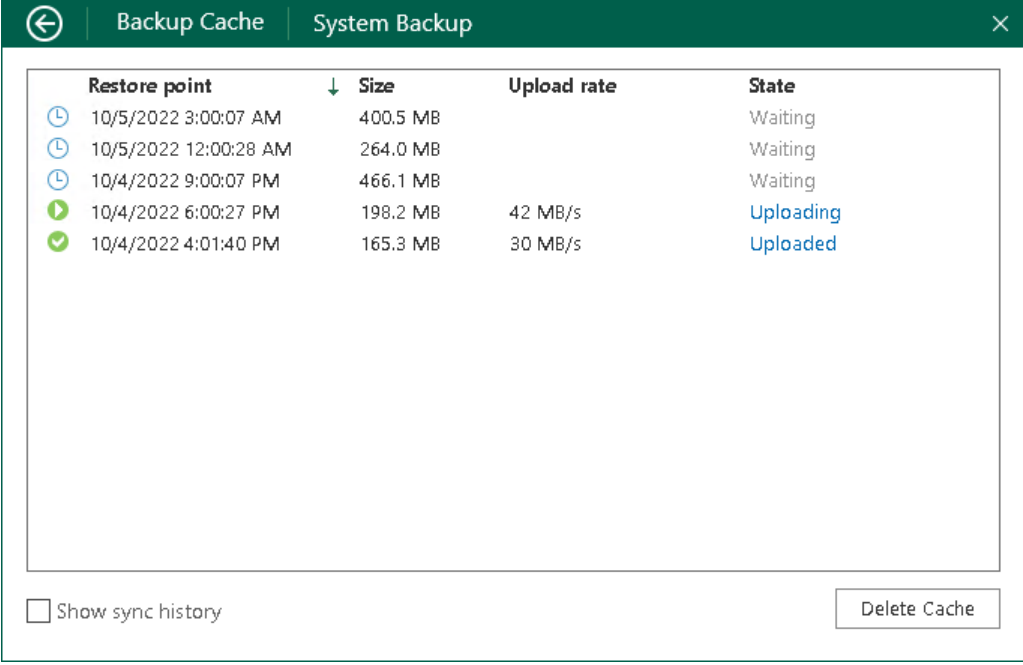


Restore point	Size	Upload rate	State
10/5/2022 12:00:28 AM	264.0 MB		Waiting
10/4/2022 9:00:07 PM	466.1 MB		Waiting
10/4/2022 6:00:27 PM	198.2 MB		Waiting
10/4/2022 4:01:40 PM	165.3 MB		Waiting

☐ Show sync history Delete Cache

## Viewing Backup Cache History

By default, the **Backup Cache** view contains a list of restore points that are waiting for upload or currently being uploaded to the target location. Restore points that have already been uploaded to the target location are not displayed in the list. To view such restore points, in the **Backup Cache** view, select the **Show sync history** check box.



	Restore point	↓ Size	Upload rate	State
🕒	10/5/2022 3:00:07 AM	400.5 MB		Waiting
🕒	10/5/2022 12:00:28 AM	264.0 MB		Waiting
🕒	10/4/2022 9:00:07 PM	466.1 MB		Waiting
🟢	10/4/2022 6:00:27 PM	198.2 MB	42 MB/s	Uploading
✅	10/4/2022 4:01:40 PM	165.3 MB	30 MB/s	Uploaded

☐ Show sync history Delete Cache

## Viewing Upload Details for Restore Points

For every restore point that is being uploaded or has been uploaded to the target location, you can also view detailed information on the upload process:

1. Open the **Backup Cache** view in one of the following ways:
  - [If the backup cache is enabled in the properties of one job] Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **View cache**.
  - [If the backup cache is enabled in the properties of multiple jobs] Double-click the Veeam Agent for Microsoft Windows icon, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**. In the main menu, hover over the name of the job that created restore points in the backup cache and select **Open backup cache**.
2. In the **Backup Cache** view, click one of the following links next to the necessary restore point:
  - **Uploading** – for a restore point that is currently being uploaded to the target location.
  - **Uploaded** – for a restore point that has been already uploaded to the target location.

In the **Upload details** view, Veeam Agent for Microsoft Windows will provide detailed information about operations performed as part of the restore point upload process.

Upload details		System Backup	
Action	Duration		
✓ Preparing for upload: Increment 1:33 PM Thursday 1/5/2023			
✓ Required backup infrastructure resources have been assigned	0:00:20		
✓ Hard disk 1 (90.0 GB) 6.0 MB read at 478 KB/s	0:00:14		
✓ Hard disk 2 (90.0 GB) 6.0 MB read at 485 KB/s	0:00:14		
✓ Finalizing	0:00:03		
✓ Processing finished at 1/10/2023 9:23:15 AM			

# Pausing Backup Cache Synchronization

After at least one restore point is created in the backup cache, Veeam Agent for Microsoft Windows starts monitoring availability of the target location. To perform this operation, Veeam Agent for Microsoft Windows starts the backup cache synchronization job that runs in the background.

The backup cache synchronization job can be paused automatically by Veeam Agent for Microsoft Windows or manually.

Veeam Agent pauses the backup cache synchronization job automatically in case Veeam Agent computer is put into sleep or hibernate mode during the job session, and the backup job is targeted at one of the following locations:


- Local computer drive
- Network shared folder

If the connection was lost during data transfer, Veeam Agent does not transfer all the data again. Instead, Veeam Agent continues data transfer that was started before the connection loss. After the connection is restored, Veeam Agent transfers only those data blocks that were not transferred before.

You can pause the backup cache synchronization job manually, for example, if you know that the target location will not become available for a while and want to reduce impact of Veeam Agent for Microsoft Windows on your OS performance.

To pause backup cache synchronization manually:

- Right-click the Veeam Agent for Microsoft Windows icon in the system tray.
- Select one of the following options:
  - **Pause sync > 10 minutes** – to pause backup cache synchronization for 10 minutes.
  - **Pause sync > 1 hour** – to pause backup cache synchronization for 1 hour.
  - **Pause sync > Until network change** – to pause backup cache synchronization until new network settings are applied to the network adapter of the Veeam Agent computer.



# Deleting Restore Points from Backup Cache

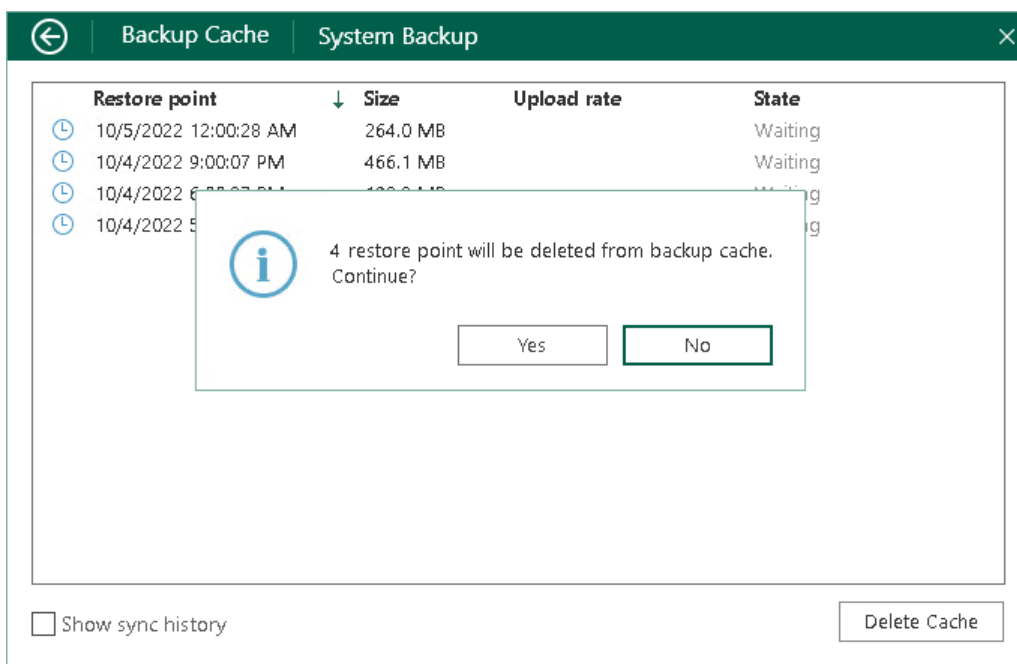
You can delete restore points from the backup cache manually if needed. For example, it is required to delete restore points after Veeam Agent for Microsoft Windows creates one or more restore points at the backup cache, and then you change the target location for backup files in the backup job settings.

To delete restore points from the backup cache:

1. Open the **Backup Cache** view in one of the following ways:
  - [If the backup cache is enabled in the properties of one job] Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **View cache**.
2. In the **Backup Cache** view, click **Delete Cache**.
3. In the window notifying that restore points will be deleted from the backup cache, click **Yes**.

## NOTE

The first backup job session following the deletion of restore points from the backup cache must complete successfully and create backup files on the target location. During this session, Veeam Agent for Microsoft Windows will create a new map of target location data blocks in the backup cache. If you delete restore points from the backup cache, and then run the backup job when the target location is unavailable, the backup job will fail.



# Performing Restore

If you experience a problem with your computer, your data gets lost or corrupted, you can use one of the following options to recover your data or bring the computer back to work:

- [Restoring from Veeam Recovery Media](#)
- [Using Veeam Agent and Microsoft Windows Tools](#)
- [Using Microsoft Windows Recovery Environment](#)
- [Restoring Volumes](#)
- [Restoring Files and Folders](#)
- [Restoring Data from Encrypted Backups](#)

# Restoring from Veeam Recovery Media

If the OS on your computer fails to start, you can use the Veeam Recovery Media to recover your computer. The Veeam Recovery Media will help you boot the computer in the limited mode. After booting, you can use Veeam Agent for Microsoft Windows or standard Microsoft Windows tools to diagnose problems and fix errors. You can also use a backup created with Veeam Agent for Microsoft Windows to restore the whole system image of your computer or specific volumes on your computer.

# Before You Begin

Before you boot from the recovery image and recover your data, check the following prerequisites:

- You must have a successfully created recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For data recovery, you can use a volume-level backup created with Veeam Agent for Microsoft Windows or system image created with Microsoft Windows. Make sure that the backup or system image is available on the computer drive (local or external), in object storage, in a network shared folder, in the backup repository managed by a Veeam backup server or in the cloud repository.

## NOTE

For a backup created by a backup copy job, you can perform bare metal restore if the backup is stored in the Veeam backup repository. For more information about backup copy jobs, see [Performing Backup Copy for Veeam Agent Backups](#).

- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- Recovery images for Microsoft Windows 32-bit OSes can be booted in the BIOS system only. Recovery images for Microsoft Windows 64-bit OSes can be booted in the BIOS and UEFI systems.

Consider the following:

- When you create a Veeam Recovery Media, Veeam Agent for Microsoft Windows stores settings for languages added to the list of input languages on your computer. If necessary, you can switch between languages using a hotkey combination when working with the **Veeam Recovery Media** wizard. The default key combination is typically **[Shift] + [Alt]**.
- NIC Teaming is not supported by Veeam Recovery Media due to Windows Recovery Environment limitations. To avoid network connection issues in Veeam Recovery Media, we recommend to disable NIC Teaming in the network switch configuration.
- You can open the Command Prompt at any moment. To do this, press **[Shift] + [F10]** on the keyboard.
- If you perform restore on a tablet, you can use a virtual keyboard to enter necessary restore settings in the **Veeam Recovery Media** wizard.



# Step 1. Boot from Veeam Recovery Media

To boot from the Veeam Recovery Media:

1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.  
  
[For removable storage device] Attach the removable storage device with the recovery image to your computer.
2. Start your computer.
3. You will be offered to boot the OS from the CD/DVD/BD or attached removable storage. Press any key on the keyboard to continue.
4. Wait for Veeam Agent for Microsoft Windows to load files from the Veeam Recovery Media. Loading the OS from the Veeam Recovery Media usually takes more time than loading the OS from the local computer drive.
5. After the OS has loaded, make sure network settings are specified correctly and configure network if necessary. To learn more, see [Select Network Adapter or Wireless Network](#).
6. Choose the necessary recovery tool to use. Veeam Agent for Microsoft Windows offers the following tools:
  - [Bare Metal Recovery](#) — the Veeam Recovery Media wizard to recover data on the original computer or perform bare metal recovery.
  - [Windows Recovery Environment](#) — built-in Microsoft Windows tools to recover the computer system image.
  - [Tools](#) — Veeam Agent for Microsoft Windows and Microsoft Windows utilities for advanced computer administration.

## TIP

Consider the following:

- To shut down or restart your computer, click the **Power Options** button at the bottom right corner of the Veeam Recovery Media screen and select the necessary option: **Shut down** or **Restart**.
- You can use the Veeam Recovery Media to start the Microsoft Windows OS in the safe mode. To do this, click the **Power Options** button at the bottom right corner of the Veeam Recovery Media screen and select **Boot in safe mode**. In the displayed window, select the necessary safe boot mode (*Minimal*, *Network* or *Repair*), and click **OK**.



## Step 2. Select Network Adapter or Wireless Network

To open the **Network settings** window, click the **Network Settings** button at the bottom right corner of the **Veeam Recovery Media** screen.

### TIP

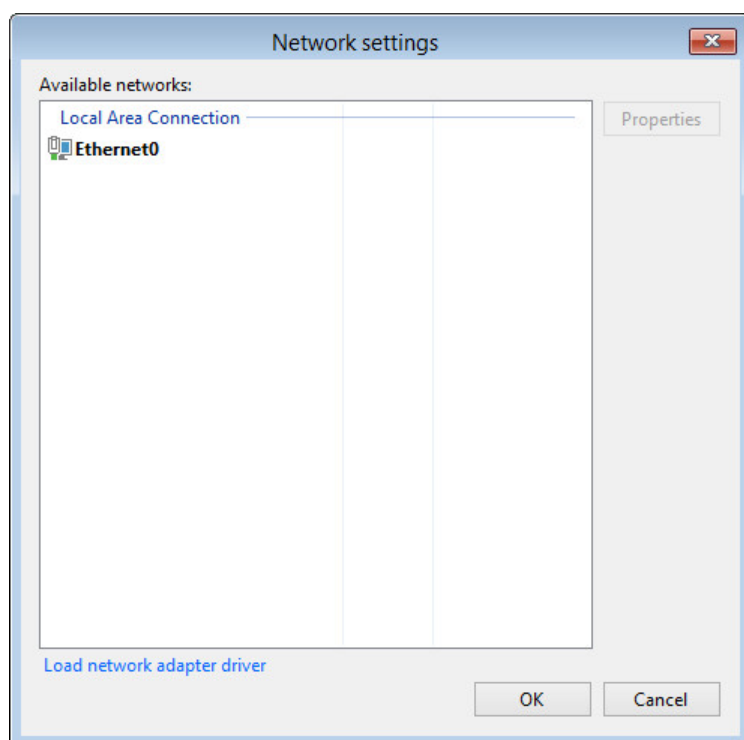
The *Network Settings* button appearance may vary depending on the detected network connection: Ethernet or wireless. If your computer is connected to a wireless network, the *Network Settings* button will indicate Wi-Fi signal strength.

Select a network adapter or wireless network that you want to use to connect to the network shared folder or Veeam backup repository where the backup resides.

- If network connection settings are included in the Veeam Recovery Media, or if there is a DHCP server in your network, Veeam Agent for Microsoft Windows will configure the network settings automatically and display available network adapters in the list.
- If you want to access the network shared folder or Veeam backup repository using a wireless network, select the necessary network in the list and click **Next**. If the wireless network is password protected, you will be prompted to specify a password for this network.
- You can manually configure TCP/IP v4 settings for adapters if necessary. To do this, select an adapter in the list and click **Properties**.

### NOTE

You will be prompted to configure network settings manually if Veeam Agent for Microsoft Windows does not detect available networks and there are no network settings included in the Veeam Recovery Media.



# Installing Network Adapter Drivers

The list of networks can be empty. This can happen in two situations:

- The driver for the network adapter is included in the Veeam Recovery Media but failed to be installed automatically for some reason.
- The driver for the network adapter is not included in the Veeam Recovery Media.

To install drivers that were included in the Veeam Recovery Media:

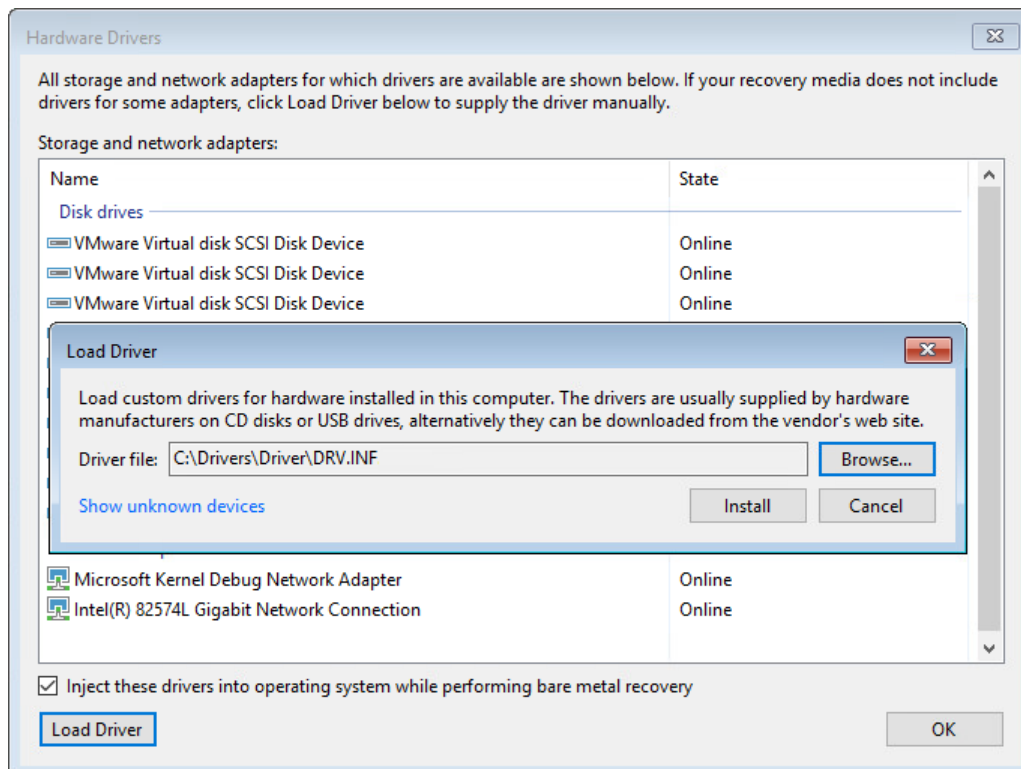
1. At the **Network settings** window, click Load network adapter driver.
2. In the **Hardware Drivers** window, select the necessary device.

If you want to include in the restored operating system all the drivers that were saved to the Veeam Recovery Media, select the **Inject these drivers into operating system while performing bare metal recovery** option. In case the option is not selected, the restored operating system will include only default Windows hardware drivers.

3. Click the **Install** link next to the selected device.

To install drivers that were not included in the Veeam Recovery Media:

1. At the **Network settings** window, click **Load network adapter driver**.
2. At the bottom of the **Hardware Drivers** window, click the **Load Driver** button and select the INF file in the driver package folder. You can also click the **Show unknown devices** link to see a list of all existing devices without drivers. This information may help you to identify the exact device for which you need to install the driver.
3. Click **Install**.



## Step 3. Launch Veeam Recovery Media Wizard

To launch the **Veeam Recovery Media** wizard, on the **Veeam Recovery Media** screen, click **Bare Metal Recovery**.

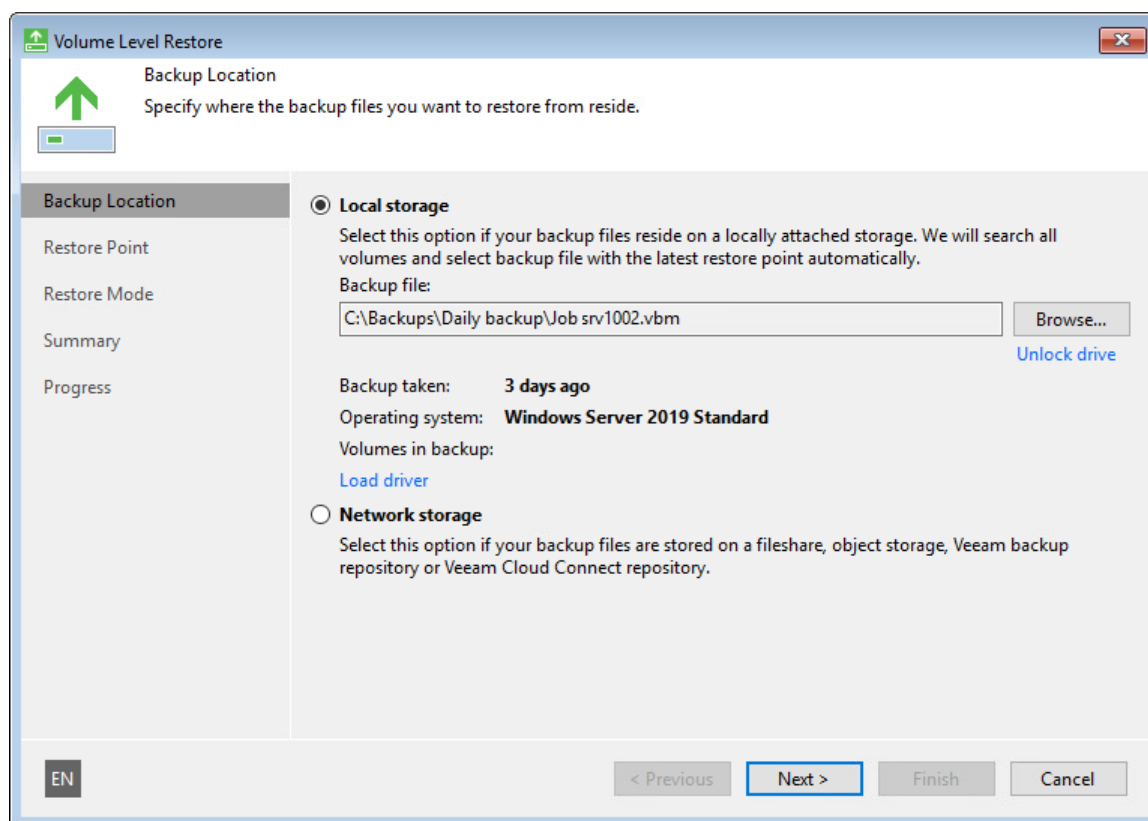


## Step 4. Specify Backup File Location

At the **Backup Location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

By default, Veeam Agent for Microsoft Windows automatically locates the latest backup on the computer drive and you pass immediately to the [Restore Point](#) step of the wizard. If Veeam Agent for Microsoft Windows fails to locate the backup on the local computer drive for some reason, or the backup file is located in object storage, in a network shared folder, in a backup repository or cloud repository, select where the backup file resides:

- **Local storage** – select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. Click **Browse** and select a backup metadata file (VBM).
- **Network storage** – select this option if the backup file resides in object storage, in a network shared folder, in a backup repository managed by a Veeam backup server or in a cloud repository exposed to you by a Veeam Cloud Connect service provider. In this case, the Veeam Recovery Media wizard will include additional steps for specifying the backup file location settings.



## Installing Drivers for Remote Storage Devices

A removable storage device with the backup file may not be displayed in the list of devices. This can happen in two situations:

- The driver for the remote storage device is included in the Veeam Recovery Media but failed to be installed automatically for some reason.
- The driver for the remote storage device is not included in the Veeam Recovery Media.

To install drivers that were included in the Veeam Recovery Media:

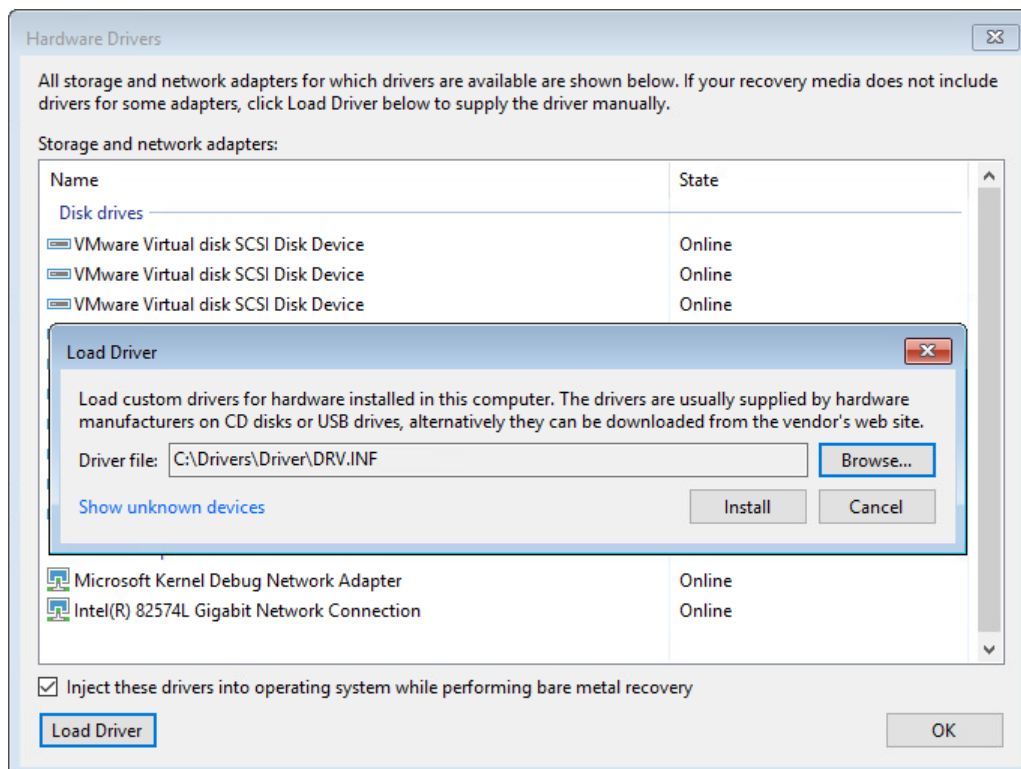
1. At the **Backup Location** step of the wizard, click **Load driver**.
2. In the **Hardware Drivers** window, select the necessary device.

If you want to include in the restored operating system all the drivers that were saved to the Veeam Recovery Media, select the **Inject these drivers into operating system while performing bare metal recovery** option. In case the option is not selected, the restored operating system will include only default Windows hardware drivers.

3. Click the **Install** link next to the selected device.

To install drivers that were not included in the Veeam Recovery Media:

1. At the **Backup Location** step of the wizard, click **Load driver**.
2. At the bottom of the **Hardware Drivers** window, click the **Load Driver** button and select the INF file in the driver package folder. You can also click the **Show unknown devices** link to see a list of all existing devices without drivers. This information may help you to identify the exact device for which you need to install the driver.
3. Click **Install**.

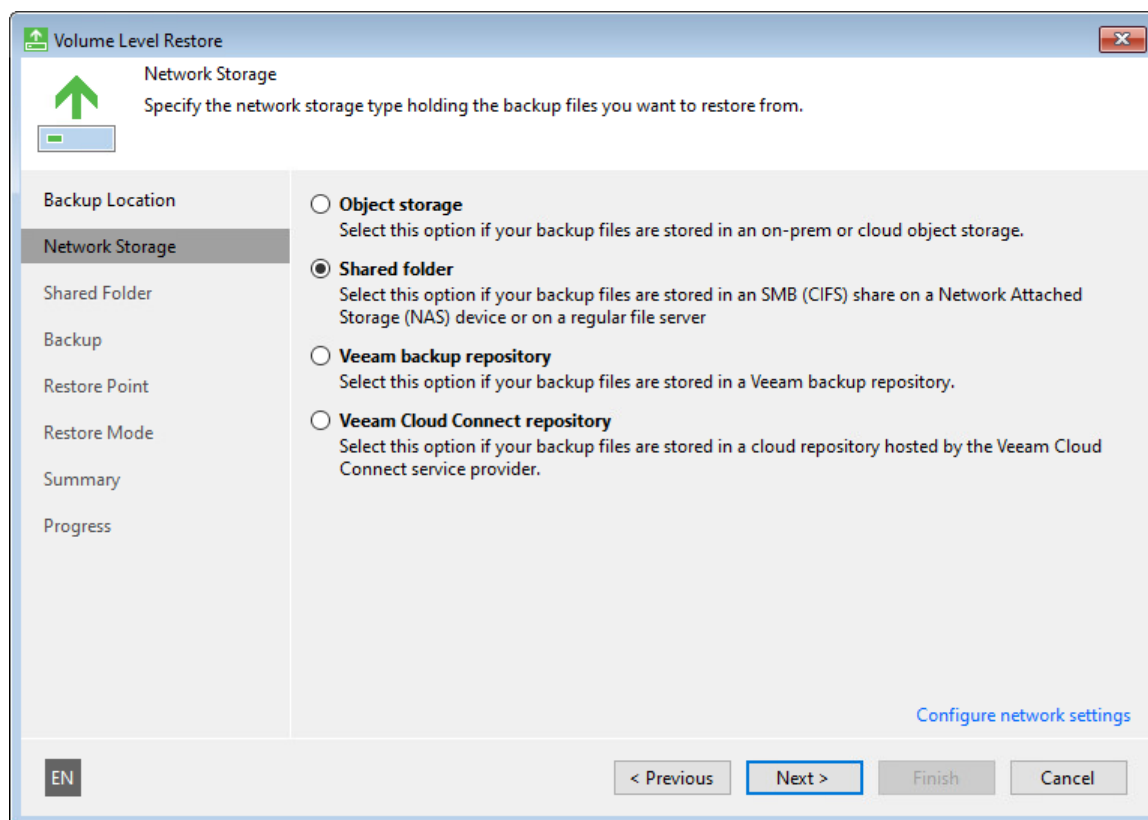


## Step 5. Select Network Storage Type

The **Network Storage** step of the wizard is available if you have selected to restore data from a backup file that resides in a remote location — in object storage, in a network shared folder, in a backup repository or in a cloud repository.

Select where the backup file resides:

- **Object storage** — select this option if the backup file resides in object storage. With this option selected, you will pass to the [Object Storage](#) step of the wizard.
- **Shared folder** — select this option if the backup file resides in a network shared folder. With this option selected, you will pass to the [Shared Folder](#) step of the wizard.
- **Veeam backup repository** — select this option if the backup file resides in a backup repository managed by a Veeam backup server. With this option selected, you will pass to the [Backup Server](#) step of the wizard.
- **Veeam Cloud Connect repository** — select this option if the backup file resides in a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the [Service Provider](#) step of the wizard.





# Step 6. Specify Network Storage Settings

Specify settings for the remote storage that contains a backup file from which you plan to restore data:

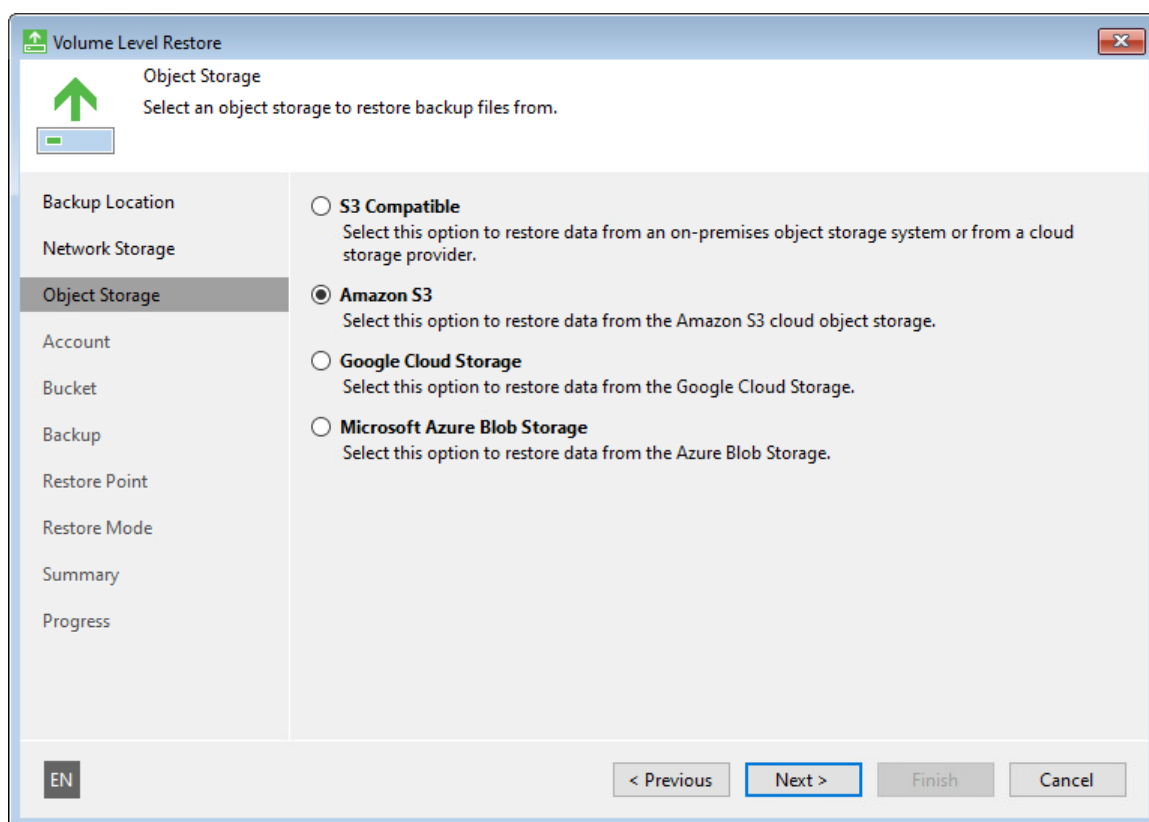
- [Shared folder settings](#) – if you have selected the **Shared folder** option at the [Network Storage](#) step of the wizard.
- [Veeam backup repository settings](#) – if you have selected the **Veeam backup repository** option at the [Network Storage](#) step of the wizard.
- [Veeam Cloud Connect repository settings](#) – if you have selected the **Veeam Cloud Connect repository** option at the [Network Storage](#) step of the wizard.

## Object Storage Settings

The **Object Storage** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

Specify settings for the object storage:

- [Specify S3 compatible settings](#).
- [Specify Amazon S3 settings](#).
- [Specify Google Cloud Storage settings](#).
- [Specify Microsoft Azure Blob Storage settings](#).



## S3 Compatible Settings

If you have selected to restore data from a backup file located in the S3 compatible storage, specify the following settings:

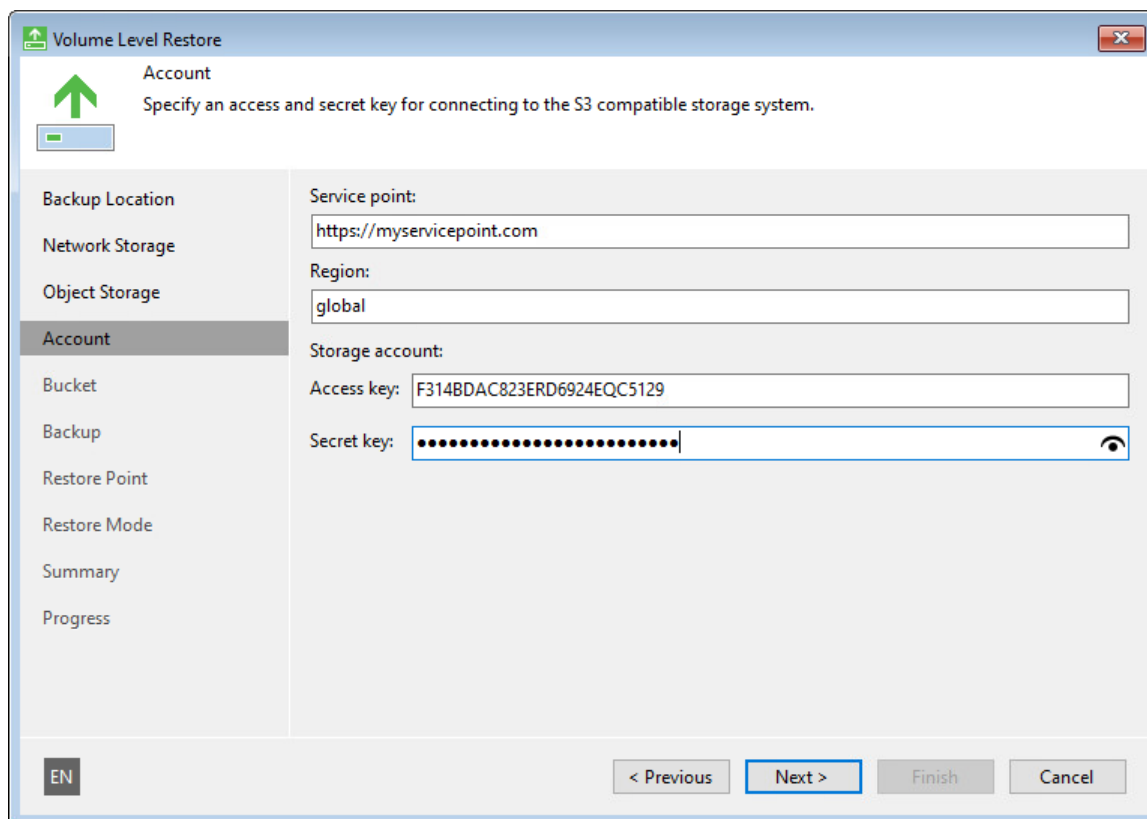
1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.
2. In the **Region** field, specify the storage region.
3. In the **Access key** field, enter the access key ID.
4. In the **Secret key** field, enter the secret access key.



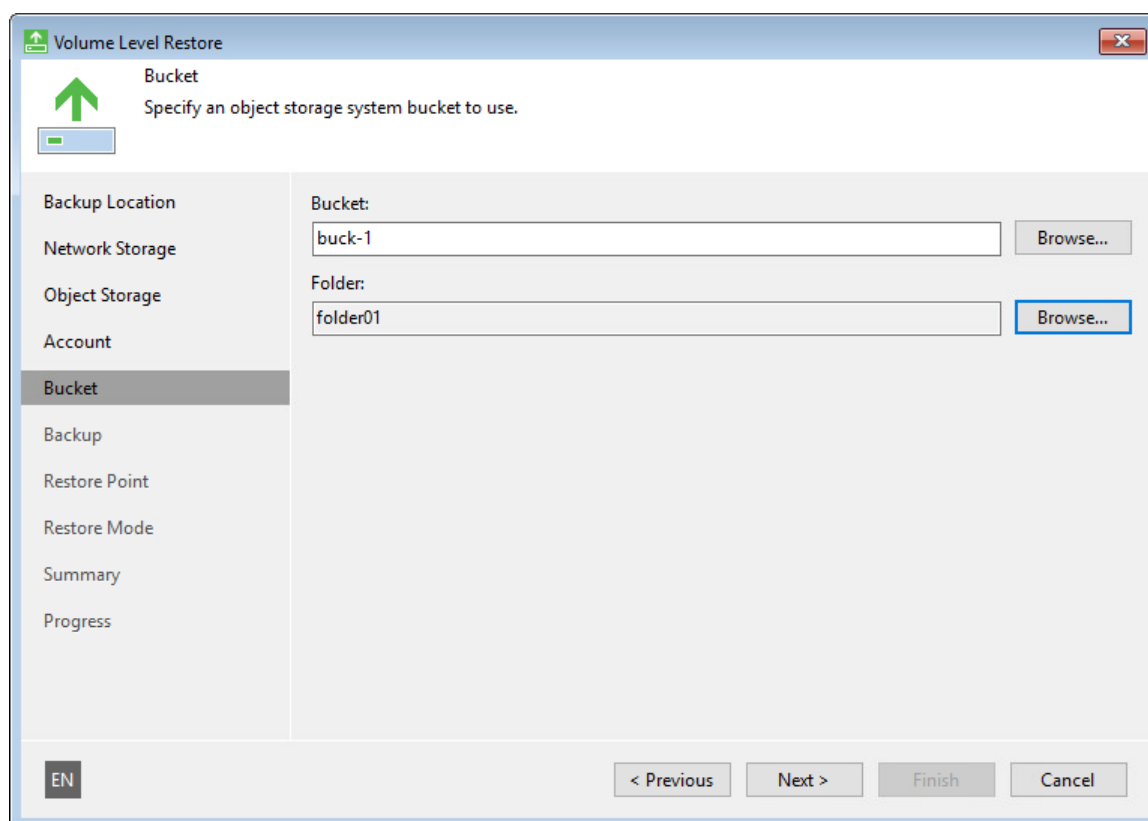
The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Inside, there's a green arrow icon and the text 'Account' and 'Specify an access and secret key for connecting to the S3 compatible storage system.' On the left is a sidebar with a tree view containing: Backup Location, Network Storage, Object Storage, Account (selected), Bucket, Backup, Restore Point, Restore Mode, Summary, and Progress. The main area has four input fields: 'Service point:' with 'https://myservicepoint.com', 'Region:' with 'global', 'Access key:' with 'F314BDAC823ERD6924EQC5129', and 'Secret key:' with a masked password '.....'. At the bottom are buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'. There's also an 'EN' button in the bottom left corner.

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.



## Amazon S3 Settings

If you have selected to restore data from a backup file located in the Amazon S3 storage, specify the following settings:

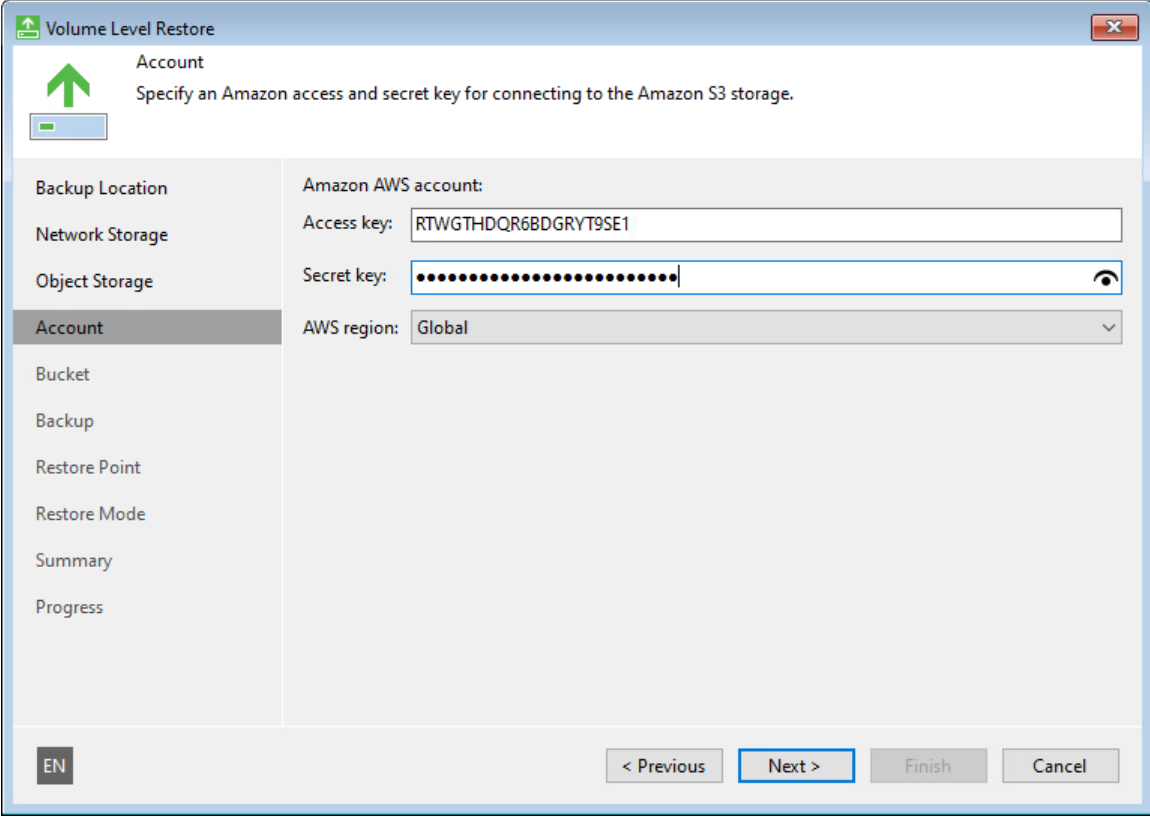
1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter the access key ID.
2. In the **Secret key** field, enter the secret access key.
3. From the **AWS region** drop-down list, select the AWS region. By default, Veeam Agent uses the **Global** region.



The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Inside, there's a green arrow icon and the text 'Account' and 'Specify an Amazon access and secret key for connecting to the Amazon S3 storage.' On the left is a sidebar with a tree view containing: Backup Location, Network Storage, Object Storage, **Account** (selected), Bucket, Backup, Restore Point, Restore Mode, Summary, and Progress. The main area has the following fields: 'Amazon AWS account:' (empty), 'Access key:' (text box with 'RTWGTQHDQR6BDGRYT9SE1'), 'Secret key:' (password box with dots and a visibility icon), and 'AWS region:' (dropdown menu showing 'Global'). At the bottom left is an 'EN' button. At the bottom right are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. From the **Data center** drop-down list, select the geographic region where the backup file is stored.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.

3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

The screenshot shows the 'Volume Level Restore' wizard window. The 'Bucket' step is selected in the left sidebar. The main area is titled 'Bucket' with the instruction 'Specify an Amazon S3 bucket to use.' Below this, there is a 'Data center:' dropdown menu set to 'EU (Paris)'. Underneath, there is a 'Bucket:' text field containing 'buck-1' and a 'Browse...' button. Below that, there is a 'Folder:' text field containing 'veeam\_backup' and another 'Browse...' button. At the bottom of the window, there are navigation buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'. A language dropdown 'EN' is located at the bottom left.

## Google Cloud Storage Settings

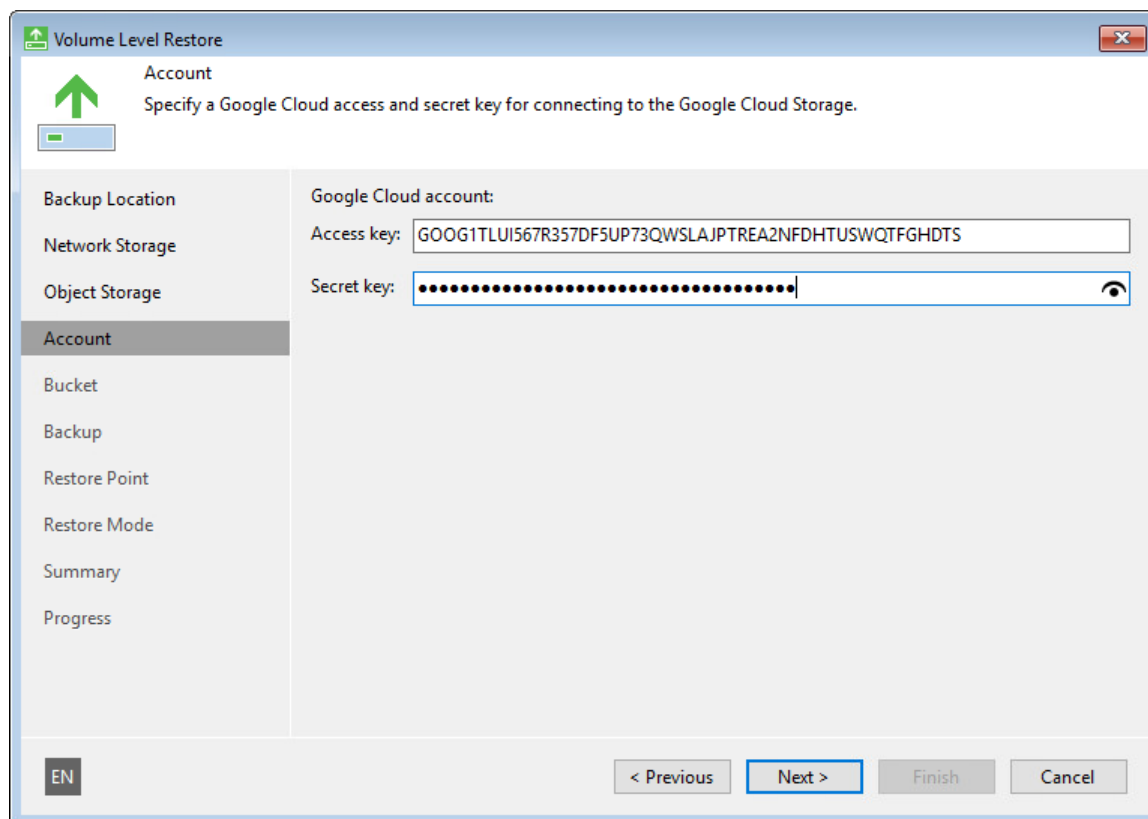
If you have selected to restore data from a backup file located in the Google Cloud storage, specify the following settings:

1. [Specify account settings](#).
2. [Specify bucket settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) keys associated with the Google Cloud account. Veeam Agent will use the HMAC keys to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see [Google Cloud documentation](#).



## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. From the **Data center region** drop-down list, select the geographic region where the backup file is stored.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.

- ii. Select the necessary folder and click **OK**.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Inside, there's a green arrow icon and the text 'Bucket' and 'Specify a Google Cloud storage bucket to use.' Below this is a progress bar. On the left is a sidebar with a list of steps: Backup Location, Network Storage, Object Storage, Account, **Bucket** (highlighted), Backup, Restore Point, Restore Mode, Summary, and Progress. The main area contains the following fields: 'Data center region:' with a dropdown menu showing 'europe-west1 (Belgium)'; 'Bucket:' with a text field containing 'eu-west1' and a 'Browse...' button; and 'Folder:' with a text field containing 'veeam\_backup' and a 'Browse...' button. At the bottom left is an 'EN' button, and at the bottom right are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

## Microsoft Azure Blob Storage Settings

If you have selected to restore data from a backup file located in the Microsoft Azure Blob storage, specify the following settings:

1. [Specify account settings](#).
2. [Specify container settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Microsoft Azure Blob storage, specify the following:

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.

3. From the **Region** drop-down list, select the Microsoft Azure region. By default, Veeam Agent uses the **Azure Global (Standard)** region.

**Volume Level Restore**

**Account**  
Specify a Microsoft Azure account and shared key for connecting to the Microsoft Azure Blob Storage.

**Backup Location**  
**Network Storage**  
**Object Storage**  
**Account**  
Container  
Backup  
Restore Point  
Restore Mode  
Summary  
Progress

Microsoft Azure Blob Storage account:

Account: newstorage

Shared key: .....

Region: Azure Global (Standard)

EN < Previous Next > Finish Cancel

## Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to restore data from a backup file located in the Microsoft Azure Blob storage and specified account settings to connect to the storage.

Specify settings for the container in the storage:

1. From the **Container** drop-down list, select a container in the storage.
2. In the **Folder** field, specify a folder in the container:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the container name or click the arrow to the left of the container name to view the list of available folders.



- ii. Select the necessary folder and click **OK**.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Inside, there's a green upward arrow icon and the text 'Container' and 'Specify a Microsoft Azure Blob Storage container to use.' Below this is a sidebar with a list of steps: Backup Location, Network Storage, Object Storage, Account, Container (highlighted), Backup, Restore Point, Restore Mode, Summary, and Progress. The main area has a 'Container:' dropdown menu with 'veeam' selected, and a 'Folder:' text box with 'veeam\_backup' entered. A 'Browse...' button is next to the folder text box. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'. There is also a small 'EN' button in the bottom left corner.

## Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. In the **Shared folder** field, enter a UNC name of the network shared folder with a backup file. Keep in mind that the UNC name always starts with two back slashes (\\).

2. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the [down-level logon name](#) format. For example, `DOMAIN\UserName` or `HOSTNAME\UserName`.

To view the specified password, click and hold the eye icon on the right of the **Password** field.

The screenshot shows the 'Volume Level Restore' wizard window. The 'Shared Folder' step is active, indicated by a green arrow icon and a highlighted sidebar item. The main area contains the following fields and controls:

- Shared folder:** A text box containing '\\srv01.tech.local\Veeam' and a 'Browse...' button.
- Authentication:** A checked checkbox labeled 'This share requires access credentials:'.
- Username:** A text box containing 'TECH\Administrator'.
- Password:** A text box filled with dots, with an eye icon on the right to toggle visibility.

The sidebar on the left lists the following steps: Backup Location, Network Storage, Shared Folder (selected), Backup, Restore Point, Restore Mode, Summary, and Progress. At the bottom, there are navigation buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'. An 'EN' button is also present in the bottom left corner.

## Backup Server Settings

The **Backup Server** step of the wizard is available if you have chosen to restore data from a backup file located in a backup repository.

Specify settings for the Veeam backup server that manages the backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup server. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. Select one of the following authentication methods to access the Veeam backup repository:
  - **Credentials** – in the **Username** and **Password** fields, enter a user name and password of the account that has access to this backup repository.

Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

## NOTE

If you want to perform bare metal restore from a backup created by Veeam Agent operating in the managed mode and stored in the Veeam backup repository, you must use an account that has the Veeam Backup Administrator or Veeam Restore Operator role on the Veeam backup server. For more information about user roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

The screenshot shows the 'Volume Level Restore' wizard window, specifically the 'Backup Server' step. The window has a title bar with a green icon and a close button. Below the title bar, there's a green upward arrow icon and the text 'Backup Server'. A instruction box says: 'Specify a Veeam Backup & Replication server name and authentication method. You can use Windows credentials in the DOMAIN\USERNAME format or a recovery token from your backup administrator.' On the left, a sidebar lists steps: 'Backup Location', 'Network Storage', 'Backup Server' (selected), 'Backup', 'Restore Point', 'Restore Mode', 'Summary', and 'Progress'. The main area contains fields for 'Veeam backup server name or IP address' (172.24.30.118) and 'Port' (10001). There are two radio buttons: 'Credentials' (selected) and 'Recovery token'. Under 'Credentials', there are fields for 'Username' (TECH\Administrator) and 'Password' (masked with dots). Under 'Recovery token', there is a 'Token' field. A note at the bottom says: 'Recovery tokens can be created using the Veeam Backup & Replication console.' At the bottom of the window, there are buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. A small 'EN' button is also visible on the left.

- **Recovery token** – in the **Token** field, enter the recovery token created on the Veeam Backup & Replication side. In this case, you will be able to recover data only from the backup for which the recovery token was generated.

The recovery token must be created beforehand. To learn more, see the [Creating Recovery Token](#) section in the Veeam Agent Management Guide.

After you click **Next**, you will be prompted to check the Veeam backup server certificate. Click **View Certificate** to check the certificate. Click **Continue** to connect to the Veeam backup server.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. On the left is a sidebar with a green upward arrow icon and a progress bar. The sidebar contains the following items: 'Backup Location', 'Network Storage', 'Backup Server' (which is highlighted), 'Backup', 'Restore Point', 'Restore Mode', 'Summary', and 'Progress'. The main area of the window is titled 'Backup Server' and contains the following text: 'Specify a Veeam Backup & Replication server name and authentication method. You can use Windows credentials in the DOMAIN\USERNAME format or a recovery token from your backup administrator.' Below this text are two sections. The first section is 'Veeam backup server name or IP address:' with a text box containing '172.24.30.118' and a 'Port:' dropdown menu set to '10005'. The second section has two radio buttons: 'Credentials' (unselected) and 'Recovery token' (selected). Below the 'Recovery token' radio button is a 'Token:' text box containing '44aE-ccA8-A73d-aC3b|'. Below the token text box is a note: 'Recovery tokens can be created using the Veeam Backup & Replication console.' At the bottom of the window are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'. There is also a small 'EN' button in the bottom left corner.

After you click **Next**, you will be prompted to check the Veeam backup server certificate. Click **View Certificate** to check the certificate. Click **Continue** to connect to the Veeam backup server.

## Service Provider Settings

If you have selected to restore data from a backup file located in a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings.](#)
2. [Verify the TLS certificate and specify user account settings.](#)

## Specifying Service Provider Settings

The **Service provider** step of the wizard is available if you have chosen to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.

2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. The main window has a green arrow icon and the text 'Service Provider' and 'Specify a DNS name or IP address and a port number received from the service provider.' On the left is a sidebar with a list of steps: Backup Location, Network Storage, Service Provider (highlighted), Credentials, Backup, Restore Point, Restore Mode, Summary, and Progress. The main area contains a 'DNS name or IP address:' text box with '172.24.30.114' entered, and a 'Port:' spinner box with '6180' selected. Below these is a note: 'Default service provider's port is 6180. If connection cannot be established, contact your service provider to make sure the settings are correct.' At the bottom are buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'. There is also a small 'EN' button in the bottom left corner.

## Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the service provider.

1. At the top of the wizard window, Veeam Agent for Microsoft Windows displays information about the TLS certificate obtained from the SP side. You can view the certificate settings and verify the TLS certificate.  
  
TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.
  - To view the TLS certificate, click the certificate link.
  - To verify if the TLS certificate with a thumbprint, copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Fingerprint for certificate verification** field. Click **Verify**. Veeam Agent for Microsoft Windows will check if the thumbprint you enter matches the thumbprint of the obtained TLS certificate.
2. In the **Username** field, enter the user name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The user name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar reads 'Volume Level Restore'. The main window has a green arrow icon and the text 'Credentials' and 'Specify credentials that you have received from the service provider and validate the certificate.' On the left is a navigation pane with the following items: 'Backup Location', 'Network Storage', 'Service Provider', 'Credentials' (highlighted), 'Backup', 'Restore Point', 'Restore Mode', 'Summary', and 'Progress'. The main area displays a message: 'This certificate has been validated.' with a certificate icon. Below this, it says 'Verified by: CN=Veeam Software, O=Veeam Software, OU=Veeam Software'. There are two input fields: 'Username:' with the text 'TechCompany\User01' and 'Password:' with a masked password of 15 dots and a toggle eye icon. At the bottom left is a language dropdown set to 'EN'. At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Volume Level Restore

Credentials

Specify credentials that you have received from the service provider and validate the certificate.

Backup Location

Network Storage

Service Provider

Credentials

Backup

Restore Point

Restore Mode

Summary

Progress

This certificate has been validated.

Verified by: [CN=Veeam Software, O=Veeam Software, OU=Veeam Software](#)

Username: TechCompany\User01

Password: •••••••••••••••

EN

< Previous Next > Finish Cancel

## Step 7. Select Backup

The **Backup** step of the wizard is available if you have chosen to restore data from a backup file that resides in a remote location — in object storage, in a network shared folder, in a Veeam backup repository or Veeam Cloud Connect repository.

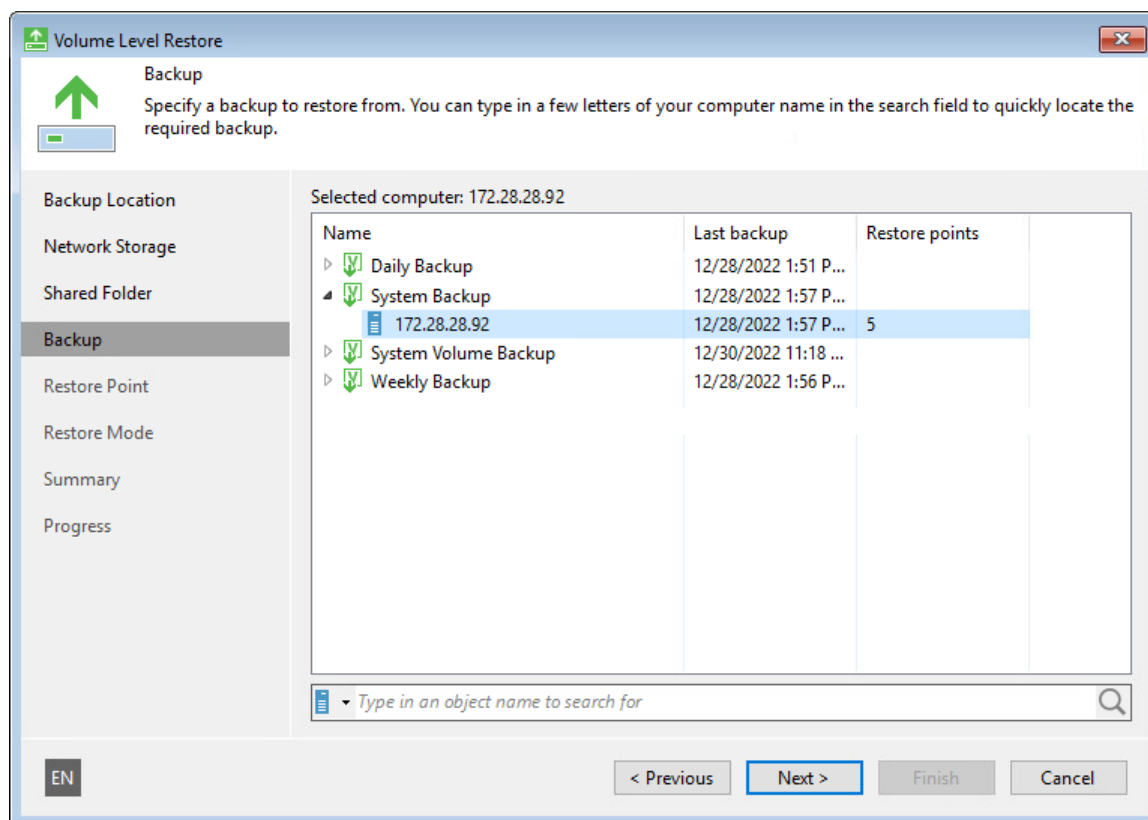
From the list of backups, select a backup from which you want to recover data. To quickly find the necessary backup, use the search field at the bottom of the window: enter a backup name or a part of it in the search field and click the **Start search** button on the right or press **[ENTER]**.

In the list of backups, Veeam Agent for Microsoft Windows displays only those backups that meet the following criteria:

1. Backups created at the volume level. File-level backups are not displayed.
2. [For backup repository target] Backups accessible by the user whose credentials are specified at the [Backup Server](#) step of the wizard:
  - If you specify credentials for the user who has access to the backup repository, the list of backups will include only backups created by this user.
  - If you specify credentials for the Backup Administrator on the backup server, the list of backups will include all Veeam Agent backups stored in the backup repository.
3. [For cloud repository target] Backups accessible by the user whose credentials are specified at the [Credentials](#) step of the wizard:
  - If you specify credentials for the tenant account, the list of backups will include backups created by all users who create backups under this tenant account and its subtenant accounts.
  - If you specify credentials for the subtenant account, the list of backups will include only those Veeam Agent backups that were created under this subtenant account.

## NOTE

If you want to restore data from an encrypted backup, and the Veeam Recovery Media from which you booted your computer does not contain encryption keys required to unlock the backup file, you need to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).



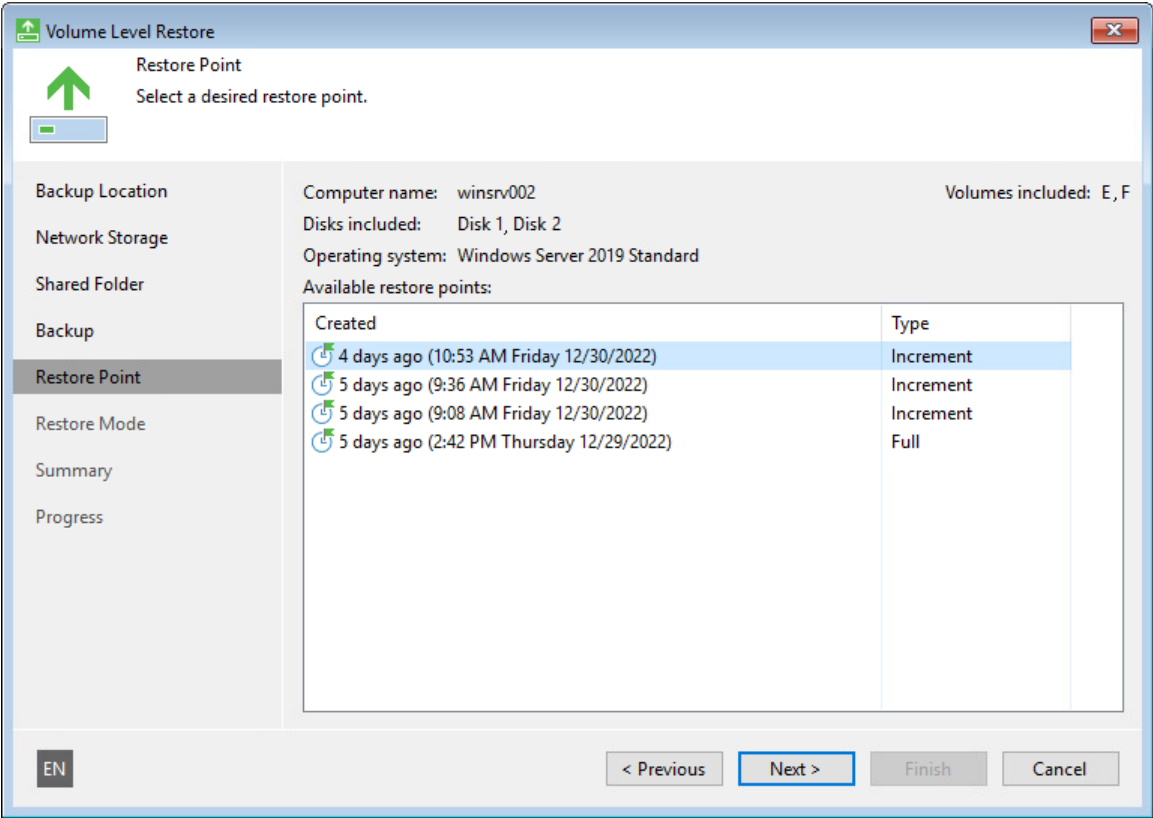


# Step 8. Select Restore Point

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

By default, Veeam Agent for Microsoft Windows uses the latest restore point. However, you can select any valid restore point to recover files and folders to a specific point in time.

Veeam Agent displays only restore points of volume-level backups. For example, if you have run 2 job sessions to create a backup of all computer volumes and then changed the backup scope to file-level backup, Veeam Agent will display only 2 restore points in the list.



## Step 9. Select Data Restore Mode

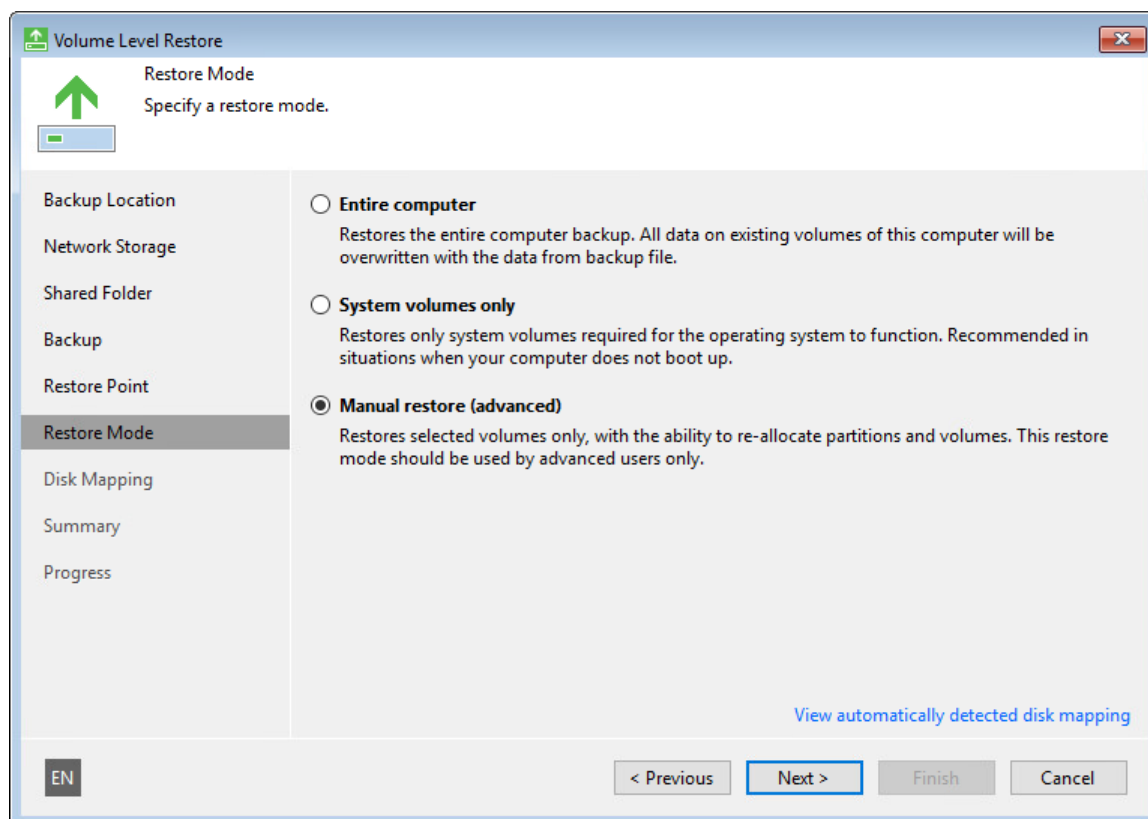
At the **Restore Mode** step of the wizard, select the data restore mode:

- **Entire computer** – select this option if you want to restore the whole system image of your computer. In this case, Veeam Agent for Microsoft Windows will attempt to map volumes from the backup to existing computer volumes and will overwrite existing data with data restored from the backup.
- **System volumes only** – select this option if you want to restore only system state data and the system volume (volume on which the Microsoft OS is installed). In this case, Veeam Agent for Microsoft Windows will restore the Microsoft Windows system partition and boot partition from the backup to your computer. For GPT disks on Microsoft Windows 8.1, 10, 11, 2012, 2012 R2, 2016, 2019 and 2022 Veeam Agent for Microsoft Windows will additionally restore the recovery partition.
- **Manual restore** – select this option if you want to choose what computer volumes you want to restore and manually allocate disk space on restored volumes. This option is recommended for users who have experience in working with Microsoft Windows disks and partitions.

To view the current disk allocations settings on your computer, at the bottom of the wizard click **View automatically detected disk mapping**.

### IMPORTANT

You will not be able to restore data in the *Entire computer* or *System volumes only* mode, if disks on a computer have not enough space to embed volume data from the backup. In this situation, you will be prompted to use the *Manual restore* mode.



# Step 10. Map Restored Disks

The **Disk Mapping** step of the wizard is available if at the [Restore Mode](#) step of the wizard you have chosen to restore data in the *Manual* mode.

You can map volumes that you want to restore from the backup to disks on the target computer.

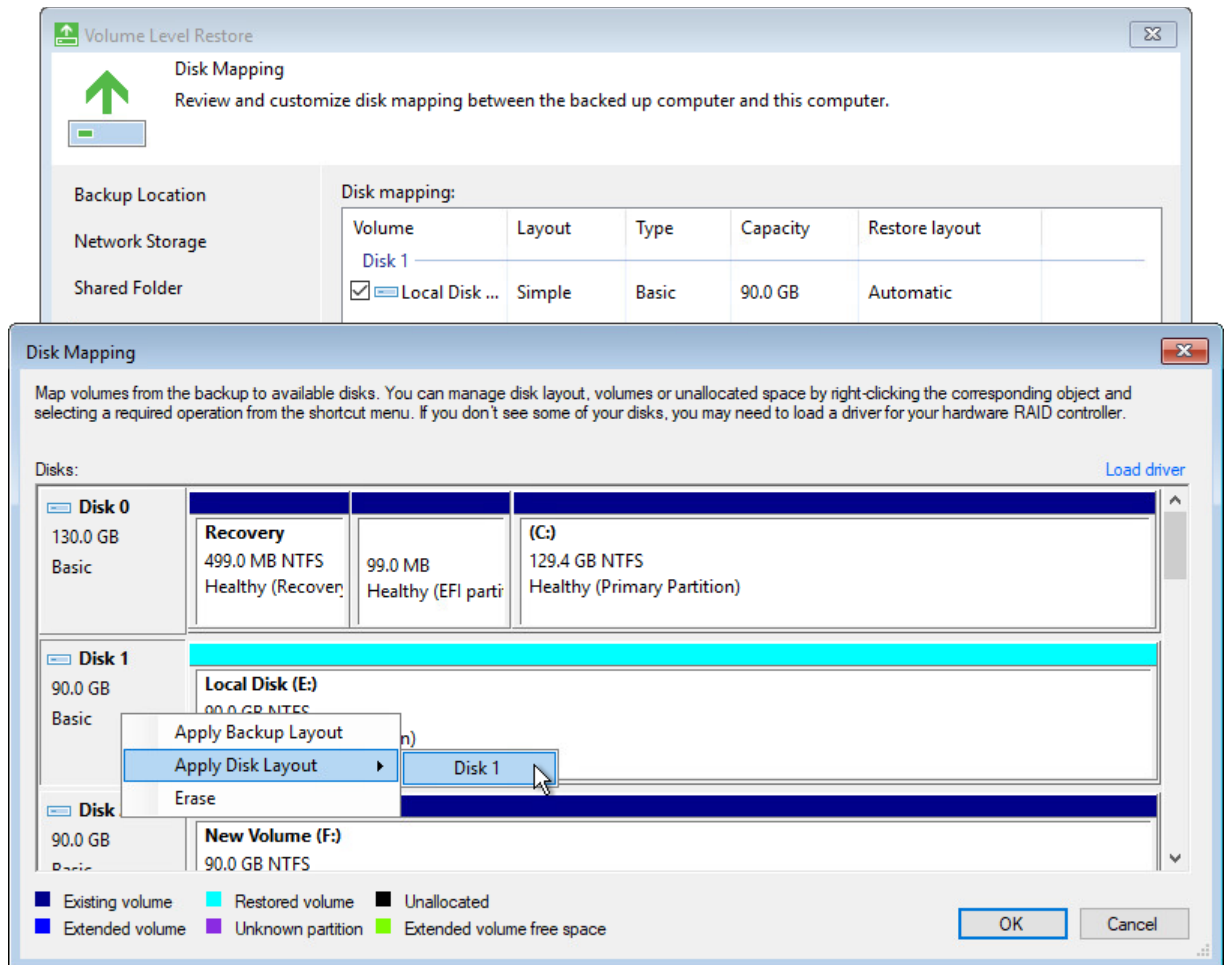
## IMPORTANT

We strongly recommend that you change disk mapping settings only if you have experience in working with Microsoft Windows disks and partitions. If you make a mistake, your computer data may get corrupted.

To map volumes:

1. Select check boxes next to volumes that you want to restore from the backup.
2. [For restore to a new location] By default, Veeam Agent for Microsoft Windows restores all volumes to their initial location. If the initial location is unavailable, a volume is restored to a disk of the same or larger size. To map the restored volume to another computer disk, at the bottom of the wizard click **Customize disk mapping**. In the **Disk Mapping** window, specify how volumes must be restored:
  - Right-click the target disk on the left and select the necessary disk layout:
    - **Apply Backup Layout** – select this option if you want to apply to disk the settings that were used on your computer at the moment when you performed backup.
    - **Apply Disk Layout** – select this option if you want to apply to the current disk settings of another disk.

- **Erase** – select this option if you want to discard the current disk settings.



- Right-click unallocated disk space in the disk area on the right and select what volume from the backup you want to place on this computer disk.

If you restore to a dynamic disk, after you select the volume to place on the disk, you will pass to the **Allocate Volume** window. To close the window, click **OK**.

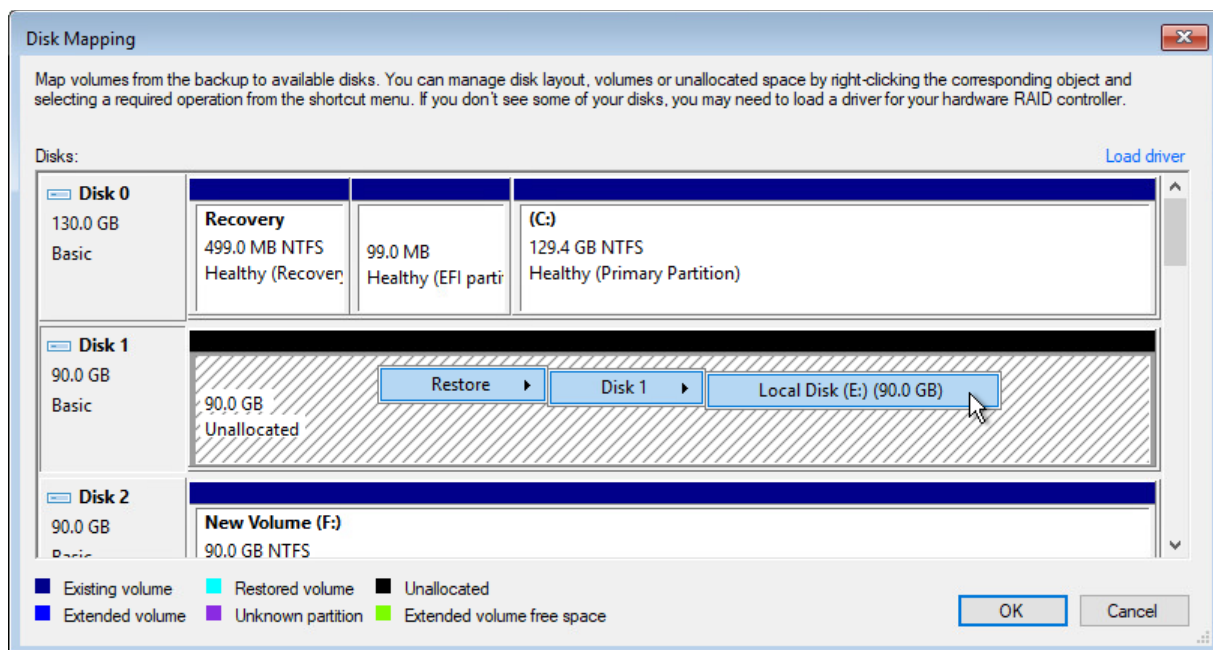
## NOTE

When you use the Veeam Recovery Media, you can recover volumes from dynamic disks only as simple volumes. If you want to restore such volumes as spanned, striped, mirrored, or parity volumes, you need to perform the following operations:

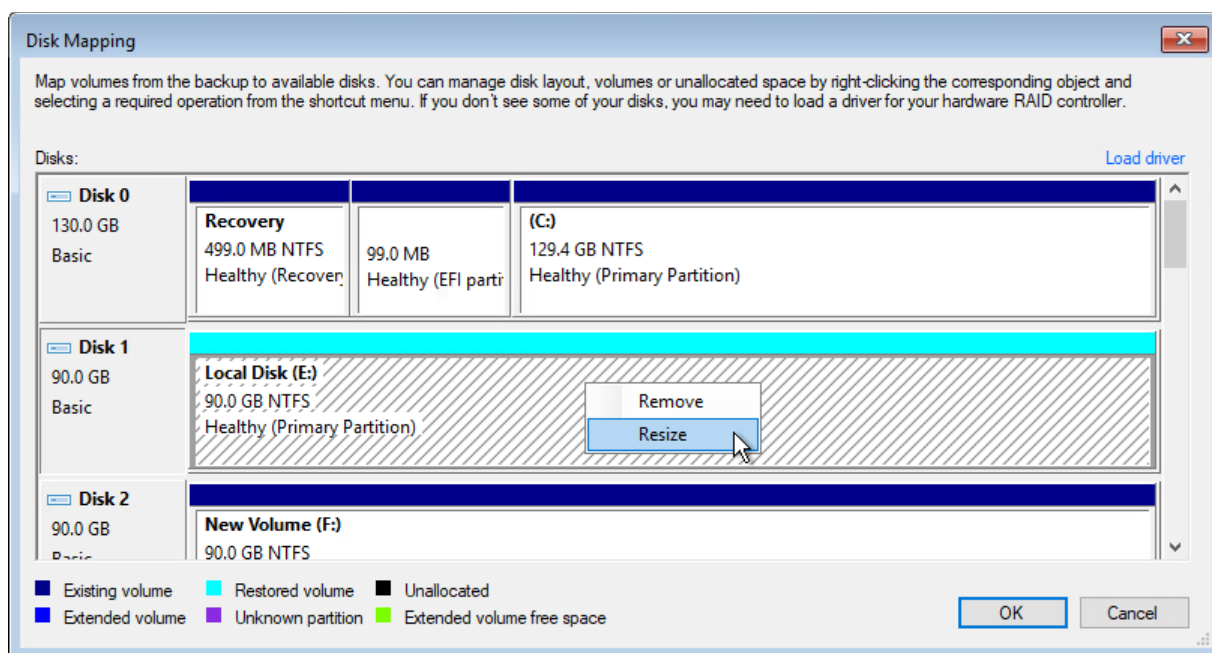
1. Using the Veeam Recovery Media, restore the system volume from any disk and data volumes from basic disks.
2. Using Veeam Agent, restore volumes from dynamic disks. To learn more, see [Restoring Volumes](#).

To learn more about volume types that you can create on dynamic disks, see [Microsoft documentation](#).

If you want to change disk layout configured by Veeam Agent for Microsoft Windows, right-click an automatically mapped volume and select **Remove**. You will be able to use the released space for mapping volumes in your own order.



3. [For restore with volume resize] You can resize a volume mapped by Veeam Agent for Microsoft Windows to a target computer disk. To resize a volume, right-click it in the **Disk Mapping** window and select **Resize**. With this option selected, you will pass to the [Volume Resize](#) window.



## NOTE

If you map a backup volume that is larger than the amount of available space on the target disk, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume. After you agree and click **OK**, Veeam Agent for Microsoft Windows will prepare to shrink the volume to the size of available disk space.

# Installing Storage Adapter Drivers

A computer disk may not be available in the list of disks. This can happen in two situations:

- The driver for the storage adapter is included in the Veeam Recovery Media but failed to be installed automatically for some reason.
- The driver for the storage adapter is not included in the Veeam Recovery Media.

To install drivers that were included in the Veeam Recovery Media:

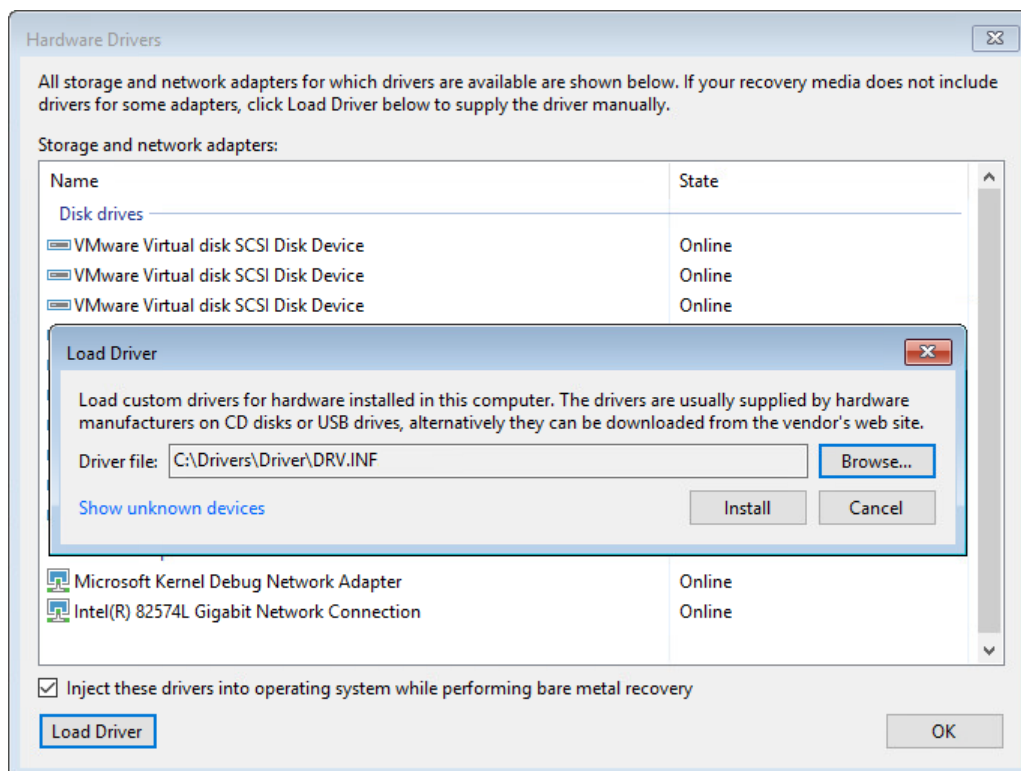
1. At the **Disk Mapping** step of the wizard, click **Load driver**.
2. In the **Hardware Drivers** window, select the necessary device.

If you do not want to save drivers for listed devices to the restored operating system, clear the **Inject these drivers into operating system while performing bare metal recovery** check box.

3. Click the **Install** link next to the selected device.

To install drivers that were not included in the Veeam Recovery Media:

1. At the Disk Mapping step of the wizard, click **Load driver**.
2. At the bottom of the **Hardware Drivers** window, click the **Load Driver** button and select the INF file in the driver package folder. You can also click the **Show unknown devices** link to see a list of all existing devices without drivers. This information may help you to identify the exact device for which you need to install the driver.
3. Click **Install**.



# Step 11. Resize Restored Volumes

At the **Disk Mapping** step of the wizard you can set the necessary size for the restored volumes. You can resize a volume if you have chosen to restore data in the *Manual* mode and customize disk layout. A volume will be shrunk or extended to the specified size during the process of data restore.


## NOTE

By default, Veeam Agent for Microsoft Windows displays volume size in megabytes (MB). This allows you to specify the desired size for the volume precisely. You can also choose to display volume size in gigabytes (GB). This may be helpful when you need to resize volumes on larger computer disks and want to simplify disk size calculations.

When you use GB as a volume size unit, you can specify volume size with integral numbers, for example, 1 GB, 60 GB or 200 GB, but not 0,8 GB, 60,5 GB or 200,7 GB. However, if the maximum volume size is in fact greater than the displayed value for less than 1 GB, Veeam Agent for Microsoft Windows will automatically add the exceeding amount of disk space to the extended volume. For example, if the maximum volume size is 60,2 GB, Veeam Agent for Microsoft Windows will display this size as 60 GB. When you specify 60 GB as a desired volume size, Veeam Agent for Microsoft Windows will extend the volume to 60,2 GB.

To resize a volume:

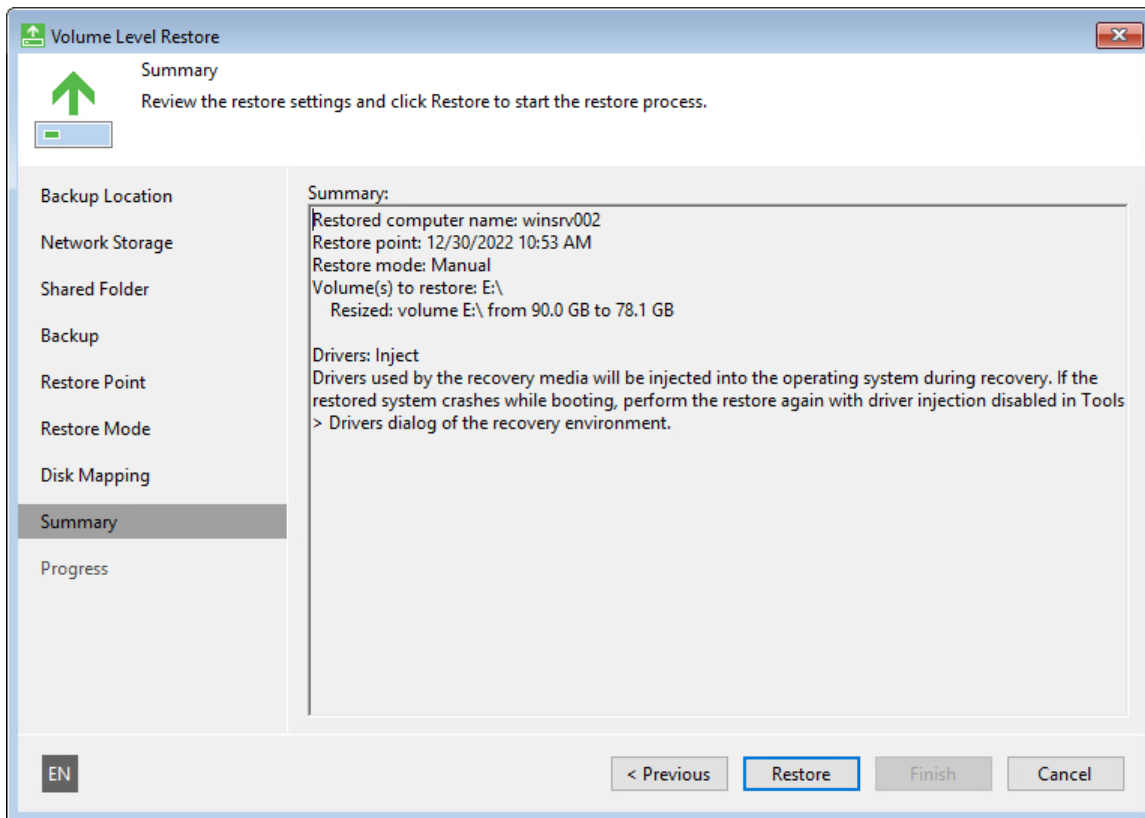
1. Specify a volume you want to resize:
  - a. Right-click a restored volume mapped to a target disk and select **Resize**.
  - b. [For volume shrink] Right-click unallocated disk space and select what volume from the backup you want to place on the computer disk. If the selected volume is larger than the amount of unallocated disk space, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume.
2. In the **Volume Resize** window, select the volume size unit and specify the desired size for the restored volume.



## Step 12. Start Restore Process

At the **Summary** step of the wizard, finalize the recovery process.

1. Review the specified recovery settings.
2. Click **Restore** to start the recovery process. Veeam Agent for Microsoft Windows will perform partition re-allocation operations if necessary, restore the necessary data from the backup and overwrite data on your computer with it.





# Using Veeam Agent and Microsoft Windows Tools

When you boot from the Veeam Recovery Media, you can use a set of tools to repair typical causes of unbootable OS, diagnose your computer and perform advanced administration tasks. Veeam Agent for Microsoft Windows offers its native tools and standard Microsoft Windows recovery tools.

## IMPORTANT

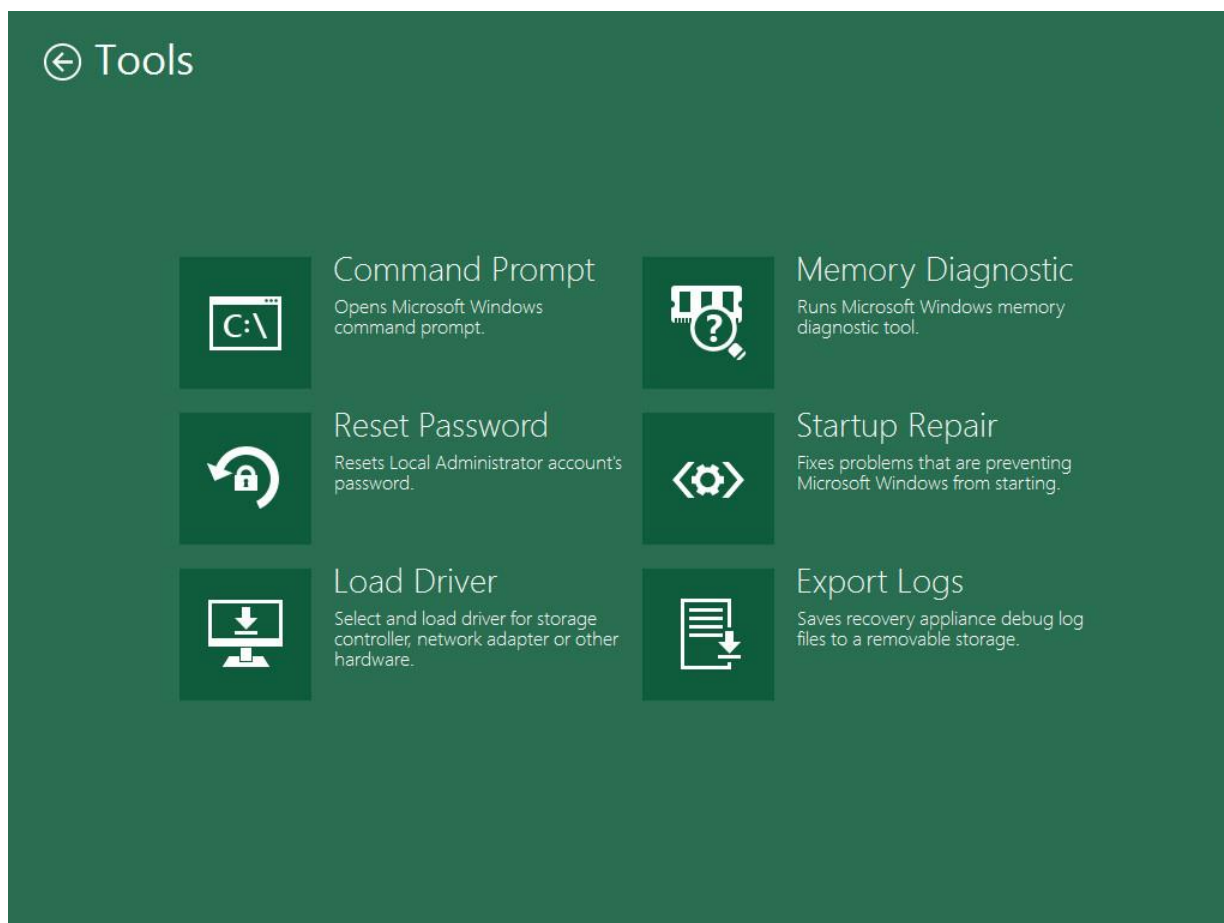
Veeam Agent for Microsoft Windows includes Microsoft Windows Tools in the Veeam Recovery Media. If some of Microsoft Windows Tools components are missing on the computer, some of Microsoft Windows Tools may not be available when you boot from the Veeam Recovery Media.

To open the tools view, on the **Veeam Recovery Media** screen, click **Tools**. Then choose the necessary tool from the list:

- **Command Prompt** – use this option to start the Microsoft Windows command prompt (`cmd.exe`).
- **Reset Password** – use this option to reset a password for the built-in Administrator account to none. The next time you boot your computer from the hard disk under the Administrator account, you will not have to specify any password. Consider the following:
  - The password reset option does not function on domain controller machines.
  - If the built-in Administrator account is disabled, this account will be enabled by the password reset option.
- **Load Driver** – use this option to load from external sources drivers that are not available on the Veeam Recovery Media. Drivers can be loaded from the computer drive or from a network shared folder.
- **Memory Diagnostic** (Microsoft utility) – use this option to check the system memory of your computer and detect potential problems. The utility can be started during the current work session or when you boot your computer the next time. To learn more, see [Microsoft documentation](#).
- **Startup Repair** (Microsoft utility) – use this option to fix system problems that may prevent Microsoft Windows from starting, for example, missing and damaged system files or the corrupted boot sector. To learn more, see [this Microsoft KB article](#).
- **Export Logs** – use this option to export the Veeam Agent for Microsoft Windows debug logs to a ZIP file and save this file on a removable storage appliance attached to your computer.

## NOTE

Do not save the archive file with debug logs on the local disk x: of the recovery image OS. This local disk is a temporary storage that will be automatically deleted after you finish working with the Veeam Recovery Media.

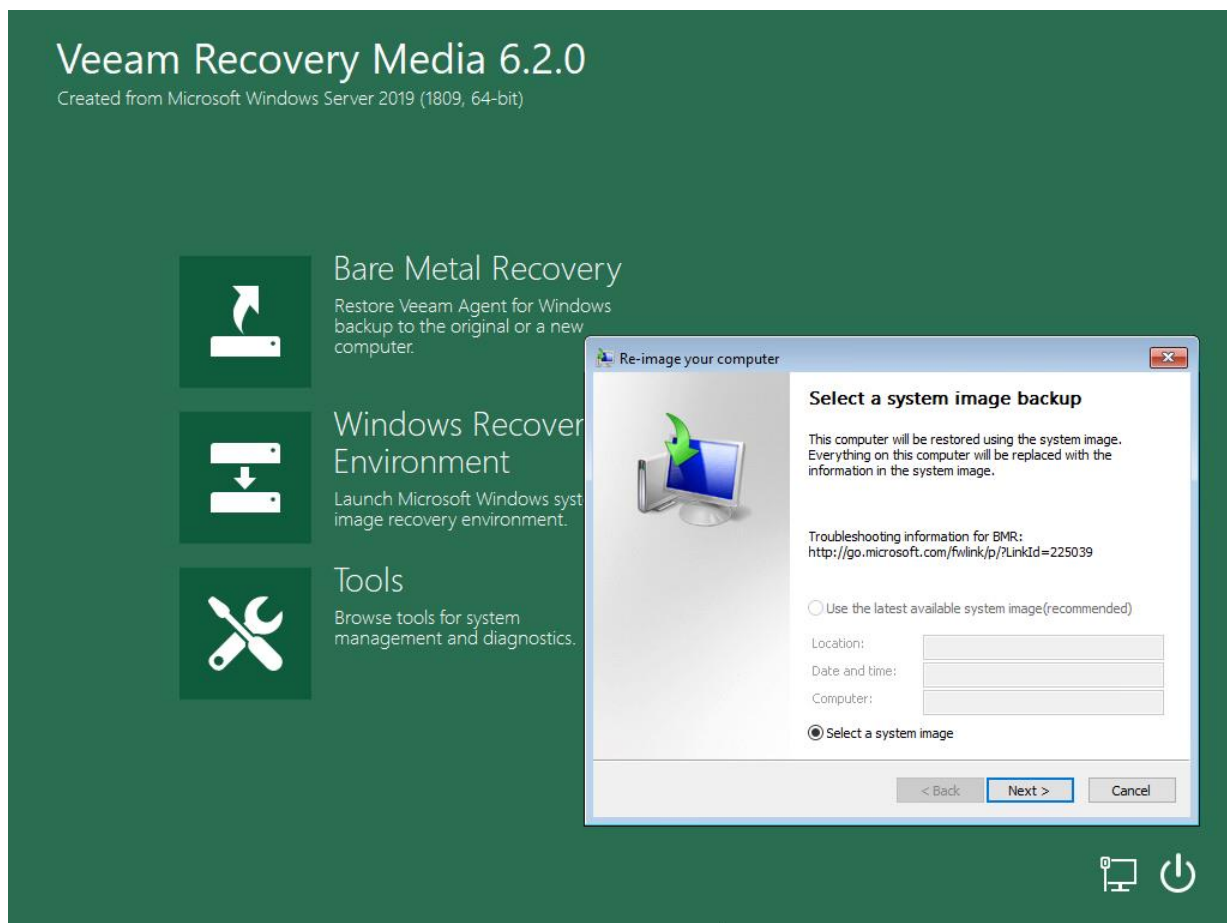


# Using Microsoft Windows Recovery Environment

If you have a Microsoft system image on the computer drive or a DVD archive with Microsoft system images, you can recover your computer using the Microsoft Windows System Image Recovery tool.

To access the Microsoft Windows System Image Recovery tool, on the Veeam Recovery Media screen, click **Windows Recovery Environment**.

The process of recovery does not differ from the process performed in Microsoft Windows. To learn more, see [this Microsoft KB article](#).



# Restoring Volumes

You can restore a specific computer volume or all volumes from the volume-level backup.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent for Microsoft Windows will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent for Microsoft Windows will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

# Before You Begin

Before you begin the volume-level restore process, check the following prerequisites and limitations.

## General

- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders and on backup repositories] You must have access to the target location where the backup file resides.
- [For backup repository targets] If you plan to restore data from a backup stored in a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- A user account under which you start the restore operation must have administrative privileges on the Veeam Agent computer. If the account under which you are currently logged on to Microsoft Windows does not have administrative privileges, you will be prompted to enter administrator credentials.

## ReFS Restore

If you restore a ReFS volume from a backup that was made on another machine, the Veeam Agent computer where you perform restore must run the OS that supports the specific ReFS version. For example, you can restore a ReFS 3.x volume only to a computer that runs one of the following OSes:

- Microsoft Windows 10 version 1803 or later
- Microsoft Windows Server 2016 or later

## NTFS Restore

- If you restore an NTFS volume from a backup that was made on another machine, and data deduplication was enabled for the volume, the Veeam Agent computer where you perform restore must run the same OS version or later as the backed-up machine OS.
- If you restore an NTFS volume from a backup that was made on another machine, and the allocation unit size of the volume was set to 128K or greater, the Veeam Agent computer where you perform restore must run one of the following OSes:
  - Microsoft Windows 10 version 1803 or later
  - Microsoft Windows Server 2019 or later

## Volume-Level Restore Limitations

Volume-level restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Microsoft Windows swap file is hosted.
- You cannot restore a volume to the volume where the backup file that you use for restore is located.


To overcome the first two limitations, you can boot from the recovery image and use the **Volume Level Restore** wizard for volume-level restore. To learn more, see [Restoring from Veeam Recovery Media](#).

# Step 1. Launch Volume Level Restore Wizard

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

- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Restore > Entire volumes**.
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the control panel, click a bar of the necessary backup job session. Click **Restore Volumes** at the bottom of the window.
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the main menu, hover over the name of the job that created the backup from which you want to restore data, and select **Restore volume**.
- From the Microsoft Windows **Start** menu, select **All Programs > Veeam > Tools > Volume Restore**.

If Veeam Agent for Microsoft Windows automatically detects backups of your computer in the target location, you will pass immediately to the [Restore Point](#) step of the wizard.

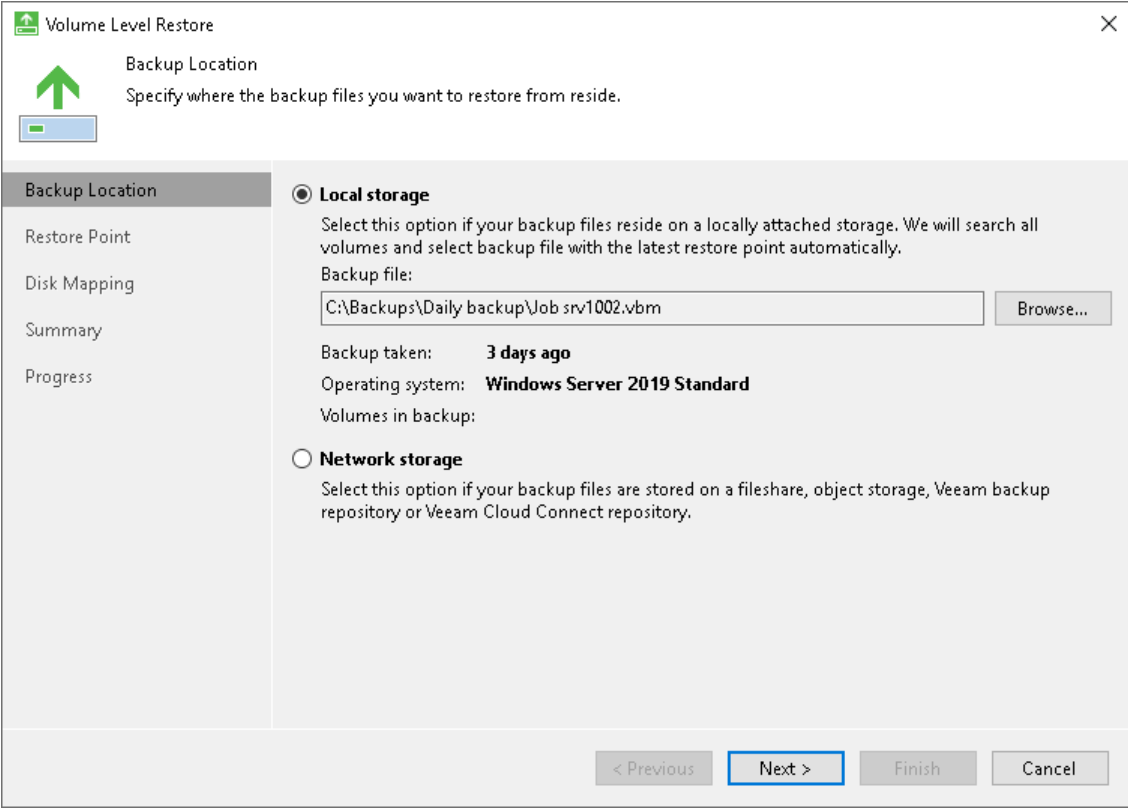


## Step 2. Specify Backup File Location

At the **Backup Location** step of the wizard, specify where the backup file that you plan to use for restore resides.

If you configured only one backup job, Veeam Agent automatically locates the latest backup of this backup job on the target storage, and you pass immediately to the [Restore Point](#) step of the wizard. If there are several backup jobs configured or Veeam Agent fails to locate the backup for some reason, specify where the backup file resides:

- **Local storage** – select this option if the backup file resides on the computer drive, external drive or removable storage device that is currently connected to your computer. Click **Browse** and select a backup metadata file (VBM).
- **Network storage** – select this option if the backup file resides in object storage, in a network shared folder, in a backup repository managed by a Veeam backup server or in a cloud repository exposed to you by a Veeam Cloud Connect service provider. In this case, the Volume Level Restore wizard will include additional steps for specifying file location settings.



The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Inside, there's a green arrow icon and the text 'Backup Location' and 'Specify where the backup files you want to restore from reside.' Below this is a sidebar with a list of steps: 'Backup Location' (selected), 'Restore Point', 'Disk Mapping', 'Summary', and 'Progress'. The main area has two radio button options: 'Local storage' (selected) and 'Network storage'. Under 'Local storage', there's a description: 'Select this option if your backup files reside on a locally attached storage. We will search all volumes and select backup file with the latest restore point automatically.' Below this is a 'Backup file:' label and a text box containing 'C:\Backups\Daily backup\Job srv1002.vbm', followed by a 'Browse...' button. Further down, it shows 'Backup taken: 3 days ago', 'Operating system: Windows Server 2019 Standard', and 'Volumes in backup:'. The 'Network storage' option has a description: 'Select this option if your backup files are stored on a fileshare, object storage, Veeam backup repository or Veeam Cloud Connect repository.' At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 3. Select Network Storage Type

The **Network Storage** step of the wizard is available if you have chosen to restore data from a backup file that resides in a remote location — in object storage, in a network shared folder, in a backup repository or in a cloud repository.

Select where the backup file is located:

- **Object storage** — select this option if the backup file resides in object storage. With this option selected, you will pass to the [Object Storage](#) step of the wizard.
- **Shared folder** — select this option if the backup file resides in a network shared folder. With this option selected, you will pass to the [Shared Folder](#) step of the wizard.
- **Veeam backup repository** — select this option if the backup file resides in a backup repository managed by a Veeam backup server. With this option selected, you will pass to the [Backup Server](#) step of the wizard.
- **Veeam Cloud Connect repository** — select this option if the backup file resides in a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the [Service Provider](#) step of the wizard.

**Volume Level Restore**

**Network Storage**  
Specify the network storage type holding the backup files you want to restore from.

**Backup Location**

- Network Storage**
- Shared Folder
- Backup
- Restore Point
- Disk Mapping
- Summary
- Progress

☐ **Object storage**  
Select this option if your backup files are stored in an on-prem or cloud object storage.

☒ **Shared folder**  
Select this option if your backup files are stored in an SMB (CIFS) share on a Network Attached Storage (NAS) device or on a regular file server

☐ **Veeam backup repository**  
Select this option if your backup files are stored in a Veeam backup repository.

☐ **Veeam Cloud Connect repository**  
Select this option if your backup files are stored in a cloud repository hosted by the Veeam Cloud Connect service provider.

< Previous   **Next >**   Finish   Cancel



## Step 4. Specify Network Storage Settings

Specify settings for the network storage that contains a backup file from which you plan to restore data:

- [Object storage settings](#) — if you have selected the **Object storage** option at the [Network Storage](#) step of the wizard.
- [Shared folder settings](#) — if you have selected the **Shared folder** option at the [Network Storage](#) step of the wizard.
- [Veeam backup repository settings](#) — if you have selected the **Veeam backup repository** option at the [Network Storage](#) step of the wizard.
- [Veeam Cloud Connect repository settings](#) — if you have selected the **Veeam Cloud Connect repository** option at the [Network Storage](#) step of the wizard.

### Object Storage Settings

The **Object Storage** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

Specify settings for the object storage:

- [Specify S3 compatible settings.](#)
- [Specify Amazon S3 settings.](#)
- [Specify Google Cloud Storage settings.](#)
- [Specify Microsoft Azure Blob Storage settings.](#)

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Inside, there's a green arrow icon and the text 'Object Storage' and 'Select an object storage to restore backup files from.' Below this is a list of steps: 'Backup Location', 'Network Storage', 'Object Storage' (which is selected and highlighted), 'Account', 'Bucket', 'Backup', 'Restore Point', 'Disk Mapping', 'Summary', and 'Progress'. To the right of the list, there are four radio button options: 'S3 Compatible' (unselected), 'Amazon S3' (selected), 'Google Cloud Storage' (unselected), and 'Microsoft Azure Blob Storage' (unselected). Each option has a brief description. At the bottom right, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Volume Level Restore

Object Storage  
Select an object storage to restore backup files from.

Backup Location

Network Storage

Object Storage

Account

Bucket

Backup

Restore Point

Disk Mapping

Summary

Progress

☐ **S3 Compatible**  
Select this option to restore data from an on-premises object storage system or from a cloud storage provider.

☒ **Amazon S3**  
Select this option to restore data from the Amazon S3 cloud object storage.

☐ **Google Cloud Storage**  
Select this option to restore data from the Google Cloud Storage.

☐ **Microsoft Azure Blob Storage**  
Select this option to restore data from the Azure Blob Storage.

< Previous   Next >   Finish   Cancel

## S3 Compatible Settings

If you have selected to restore data from a backup file located in the S3 compatible storage, specify the following settings:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.
2. In the **Region** field, specify the storage region.
3. In the **Access key** field, enter the access key ID.
4. In the **Secret key** field, enter the secret access key.

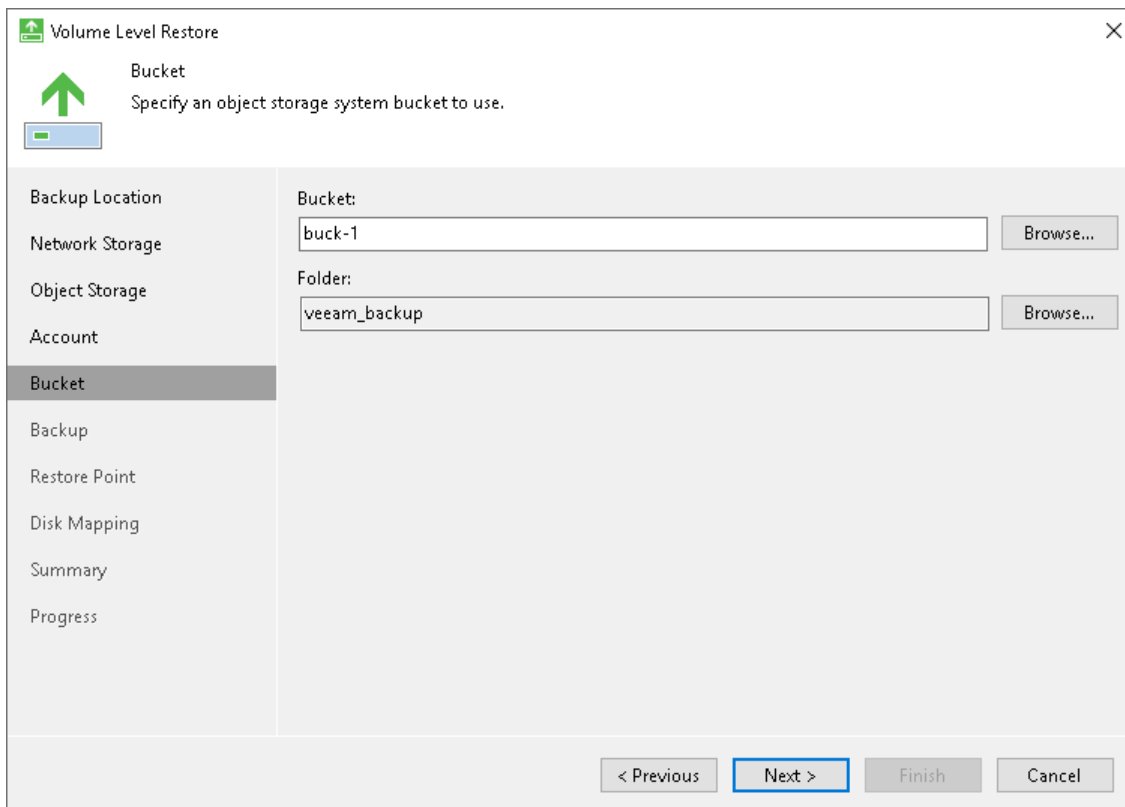
The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore' with a close button. Inside, there's a green arrow icon and the text 'Account' and 'Specify an access and secret key for connecting to the S3 compatible storage system.' Below this is a sidebar with a list of steps: Backup Location, Network Storage, Object Storage, Account (highlighted), Bucket, Backup, Restore Point, Disk Mapping, Summary, and Progress. The main area contains four fields: 'Service point:' with the value 'https://myservicepoint.com', 'Region:' with the value 'global', 'Storage account:' with the value 'F314BDAC823E13D427D692D49C129', and 'Secret key:' with a masked value represented by dots. At the bottom right, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **SelectBucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **SelectFolder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.



## Amazon S3 Settings

If you have selected to restore data from a backup file located in the Amazon S3 storage, specify the following settings:

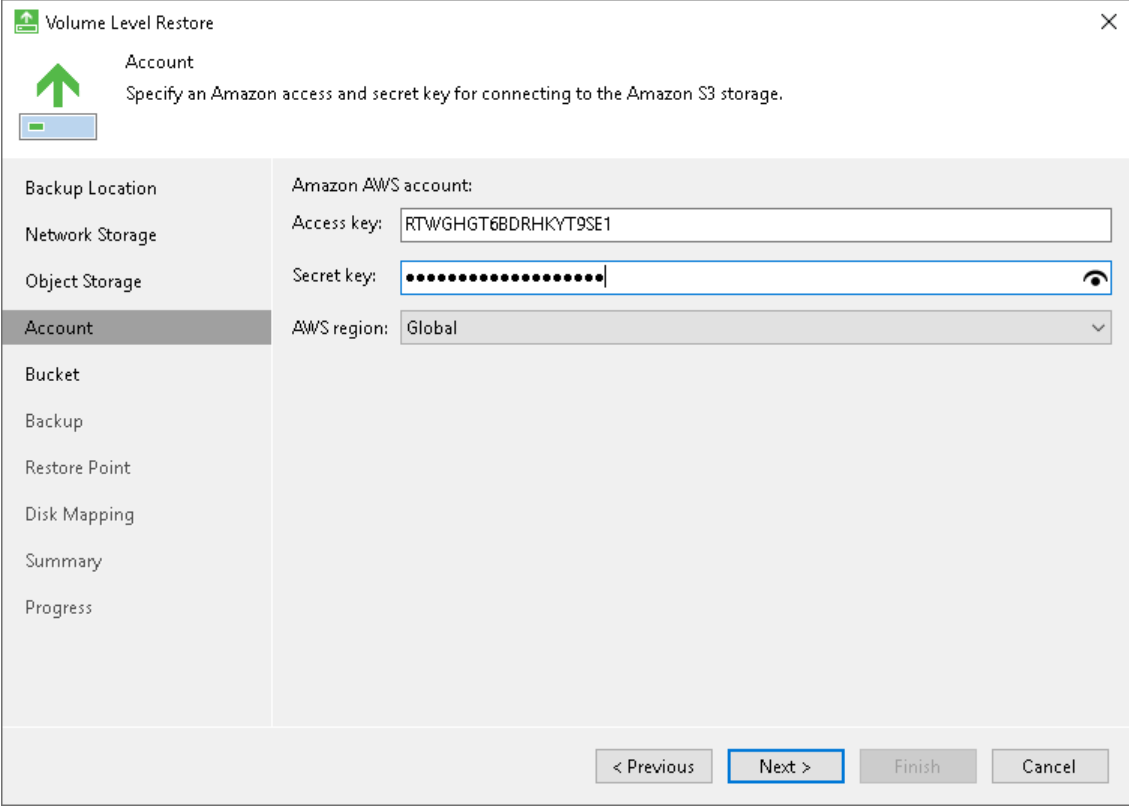
1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter the access key ID.
2. In the **Secret key** field, enter the secret access key.
3. From the **AWS region** drop-down list, select the AWS region. By default, Veeam Agent uses the **Global** region.



The screenshot shows the 'Volume Level Restore' wizard window. The 'Account' step is selected in the left sidebar, which also lists 'Backup Location', 'Network Storage', 'Object Storage', 'Bucket', 'Backup', 'Restore Point', 'Disk Mapping', 'Summary', and 'Progress'. The main area contains the following fields:

- Amazon AWS account:** (Label)
- Access key:** (Text field containing 'RTWGHGT6BDRHKYT9SE1')
- Secret key:** (Text field with masked characters and a toggle icon)
- AWS region:** (Drop-down menu showing 'Global')

At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. From the **Data center** drop-down list, select the geographic region where the backup file is stored.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.

3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
    - ii. Select the necessary folder and click **OK**.

**Volume Level Restore**

**Bucket**  
Specify an Amazon S3 bucket to use.

**Backup Location**  
**Network Storage**  
**Object Storage**  
**Account**  
**Bucket**  
Backup  
Restore Point  
Disk Mapping  
Summary  
Progress

**Data center:**  
US East (N. Virginia)

**Bucket:**  
bucket\_3 **Browse...**

**Folder:**  
veeam\_backup **Browse...**

< Previous **Next >** Finish Cancel

## Google Cloud Storage Settings

If you have selected to restore data from a backup file located in the Google Cloud storage, specify the following settings:

1. [Specify account settings](#).
2. [Specify bucket settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) keys associated with the Google Cloud account. Veeam Agent will use the HMAC keys to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see [Google Cloud documentation](#).

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. From the **Data center region** drop-down list, select the geographic region where the backup file is stored.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.

- ii. Select the necessary folder and click **OK**.

**Volume Level Restore**

**Bucket**  
Specify a Google Cloud storage bucket to use.

**Backup Location**  
Network Storage  
Object Storage  
Account  
**Bucket**  
Backup  
Restore Point  
Disk Mapping  
Summary  
Progress

**Data center region:**  
Belgium

**Bucket:**  
west1 **Browse...**

**Folder:**  
veeam\_backup **Browse...**

< Previous **Next >** Finish Cancel

## Microsoft Azure Blob Storage Settings

If you have selected to restore data from a backup file located in the Microsoft Azure Blob storage, specify the following settings:

1. [Specify account settings.](#)
2. [Specify container settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Microsoft Azure Blob storage, specify the following:

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.

3. From the **Region** drop-down list, select the Microsoft Azure region. By default, Veeam Agent uses the **Azure Global (Standard)** region.

**Volume Level Restore**

**Account**  
Specify a Microsoft Azure account and shared key for connecting to the Microsoft Azure Blob Storage.

**Backup Location**  
**Network Storage**  
**Object Storage**  
**Account**  
Container  
Backup  
Restore Point  
Disk Mapping  
Summary  
Progress

**Microsoft Azure Blob storage account:**

Account:

Shared key:

Region:

< Previous   **Next >**   Finish   Cancel

## Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to restore data from a backup file located in the Microsoft Azure Blob storage and specified account settings to connect to the storage.

Specify settings for the container in the storage:

1. From the **Container** drop-down list, select a container in the storage.
2. In the **Folder** field, specify a folder in the container:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the container name or click the arrow to the left of the container name to view the list of available folders.



ii. Select the necessary folder and click **OK**.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Below the title bar, there is a green upward arrow icon and the text 'Container' and 'Specify a Microsoft Azure Blob Storage container to use.' On the left side, there is a vertical list of steps: 'Backup Location', 'Network Storage', 'Object Storage', 'Account', 'Container' (which is highlighted), 'Backup', 'Restore Point', 'Disk Mapping', 'Summary', and 'Progress'. The main area of the wizard is divided into two sections. The top section is labeled 'Container:' and has a dropdown menu with 'veeam' selected. The bottom section is labeled 'Folder:' and has a text box containing 'veeam\_backup' and a 'Browse...' button. At the bottom of the wizard, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

## Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have chosen to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. In the Shared folder field, enter a UNC name of the network shared folder with the backup file. Keep in mind that the UNC name always starts with two back slashes (\\).
2. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the [down-level logon name](#) format. For example, *DOMAIN\UserName* or *HOSTNAME\UserName*.

If you do not select the **This share requires access credentials** check box, Veeam Agent for Microsoft Windows will connect to the shared folder using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed.

3. To view the entered password, click and hold the eye icon on the right of the **Password** field.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Below the title bar, there is a green arrow icon and the text 'Shared Folder' and 'Specify an SMB (CIFS) shared folder UNC path and credentials.' On the left side, there is a vertical list of steps: 'Backup Location', 'Network Storage', 'Shared Folder' (which is highlighted), 'Backup', 'Restore Point', 'Disk Mapping', 'Summary', and 'Progress'. The main area of the window contains the following fields and controls:

- 'Shared folder:' text box with the value '\\filesrv01\sharedbackups' and a 'Browse...' button to its right.
- A checked checkbox labeled 'This share requires access credentials:'.
- 'Username:' text box with the value 'TECH\Administrator'.
- 'Password:' text box with masked characters (dots) and an eye icon to its right.

At the bottom of the window, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

## Backup Server Settings

The **Backup Server** step of the wizard is available if you have chosen to restore data from a backup file located in a backup repository.

Specify settings for the Veeam backup server that manages the backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup server. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. Select the **Specify your personal credentials** check box. In the **Username** and **Password** fields, enter a user name and password of the account that has access to this backup repository. Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

If you do not select the **Specify your personal credentials** check box, Veeam Agent for Microsoft Windows will connect to the backup repository using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can add the computer account (*DOMAIN\COMPUTERNAME\$*) to an AD group and grant access rights on the backup repository to this group.

Setting access permissions on the backup repository to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). If you have set such permissions on the backup repository, you can omit specifying credentials. However, we recommend this scenario for demo environments only.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore'. Inside, there's a green arrow icon and the text 'Backup Server'. Below that, it says 'Specify a Veeam Backup & Replication server name and authentication method. You can use Windows credentials in the DOMAIN\USERNAME format or a recovery token from your backup administrator.' On the left, there's a sidebar with options: 'Backup Location', 'Network Storage', 'Backup Server' (selected), 'Backup', 'Restore Point', 'Disk Mapping', 'Summary', and 'Progress'. The main area has a 'Veeam backup server name or IP address:' field with '172.24.30.118' and a 'Port:' dropdown set to '10001'. Below that, there's a checkbox 'Specify your personal credentials:' which is checked. Under this checkbox, there are 'Username:' and 'Password:' fields. The 'Username' field contains 'TECH\Administrator' and the 'Password' field is filled with dots. At the bottom right, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Service Provider Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings.](#)
2. [Verify the TLS certificate and specify user account settings.](#)

## Specifying Service Provider Settings

The **Service provider** step of the wizard is available if you have chosen to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.

2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.

The screenshot shows the 'Volume Level Restore' wizard window. The 'Service Provider' step is selected in the left-hand navigation pane. The main area contains a text box for 'DNS name or IP address' with the value '172.24.30.114' and a 'Port' dropdown menu set to '6180'. Below these fields, a note states: 'Default service provider's port is 6180. If connection cannot be established, contact your service provider to make sure the settings are correct.' At the bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step	Field	Value
Service Provider	DNS name or IP address	172.24.30.114
	Port	6180

## Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the service provider.

1. At the top of the wizard window, Veeam Agent for Microsoft Windows displays information about the TLS certificate obtained from the SP side. You can view the certificate settings and verify the TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

- To view the TLS certificate, click the certificate link.
- To verify if the TLS certificate with a thumbprint, copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Fingerprint for certificate verification** field. Click **Verify**. Veeam Agent for Microsoft Windows will check if the thumbprint you enter matches the thumbprint of the obtained TLS certificate.

2. In the **Username** field, enter the user name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The user name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.

The screenshot shows the 'Volume Level Restore' wizard window. The title bar says 'Volume Level Restore' with a close button. The main area is titled 'Credentials' with a subtitle 'Specify credentials that you have received from the service provider and validate the certificate.' On the left is a navigation pane with the following items: 'Backup Location', 'Network Storage', 'Service Provider', 'Credentials' (selected), 'Backup', 'Restore Point', 'Disk Mapping', 'Summary', and 'Progress'. The main content area shows a message: 'This certificate has been validated.' with a verified by string: 'CN=Veeam Software, O=Veeam Software, OU=Veeam Software'. Below this are two input fields: 'Username:' with the value 'TechCompany\User01' and 'Password:' with masked characters '.....'. At the bottom are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Volume Level Restore

**Credentials**  
Specify credentials that you have received from the service provider and validate the certificate.

Backup Location  
Network Storage  
Service Provider  
**Credentials**  
Backup  
Restore Point  
Disk Mapping  
Summary  
Progress

This certificate has been validated.  
Verified by: [CN=Veeam Software, O=Veeam Software, OU=Veeam Software](#)

Username: TechCompany\User01

Password: .....

< Previous   Next >   Finish   Cancel

## Step 5. Select Backup

The **Backup** step of the wizard is available if you have chosen to restore data from a backup file located in object storage, in a network shared folder or in a backup repository.

From the list of backups, select a backup from which you want to recover data. To quickly find the necessary backup, use the search field at the bottom of the window: enter a backup name or a part of it in the search field and click the **Start search** button on the right or press **[ENTER]**.

In the list of backups, Veeam Agent for Microsoft Windows displays only those backups that meet the following criteria:

1. Backups created at the volume level. File-level backups are not displayed.
2. [For backup repository target] Backups accessible by the user whose credentials are specified at the [Backup Server](#) step of the wizard:
  - If you specify credentials for the user who has access to the backup repository, the list of backups will include only backups created by this user.
  - If you specify credentials for the user who is assigned the *Backup Administrator* or *Restore Operator* role on the backup server, the list of backups will include all Veeam Agent backups stored in the backup repository.
3. [For cloud repository target] Backups accessible by the user whose credentials are specified at the [Credentials](#) step of the wizard:
  - If you specify credentials for the tenant account, the list of backups will include backups created by all users who create backups under this account.
  - If you specify credentials for the subtenant account, the list of backups will include only those Veeam Agent backups that were created under this subtenant account.

## NOTE

If you restore data from an encrypted backup that was created on another Veeam Agent computer, you need to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).

The screenshot shows the 'Volume Level Restore' wizard in the 'Backup' step. The left sidebar contains a list of steps: Backup Location, Network Storage, Shared Folder, Backup (selected), Restore Point, Disk Mapping, Summary, and Progress. The main area displays a table of backups for the selected computer 172.24.25.94. The table has three columns: Name, Last backup, and Restore points. The 'System Backup' folder is expanded, showing a sub-entry '172.24.25.94' which is highlighted. Below the table is a search bar with the placeholder text 'Type in an object name to search for'. At the bottom are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

**Volume Level Restore**

**Backup**  
Specify a backup to restore from. You can type in a few letters of your computer name in the search field to quickly locate the required backup.

Backup Location  
Network Storage  
Shared Folder  
**Backup**  
Restore Point  
Disk Mapping  
Summary  
Progress

Selected computer: 172.24.25.94

Name	Last backup	Restore points
▶ [✓] Daily Backup	12/28/2022 1:51 P...	
▶ [✓] System Backup	12/28/2022 1:57 P...	
172.24.25.94	12/28/2022 1:57 P...	4
▶ [✓] Weekly Backup	12/28/2022 1:56 P...	

Type in an object name to search for

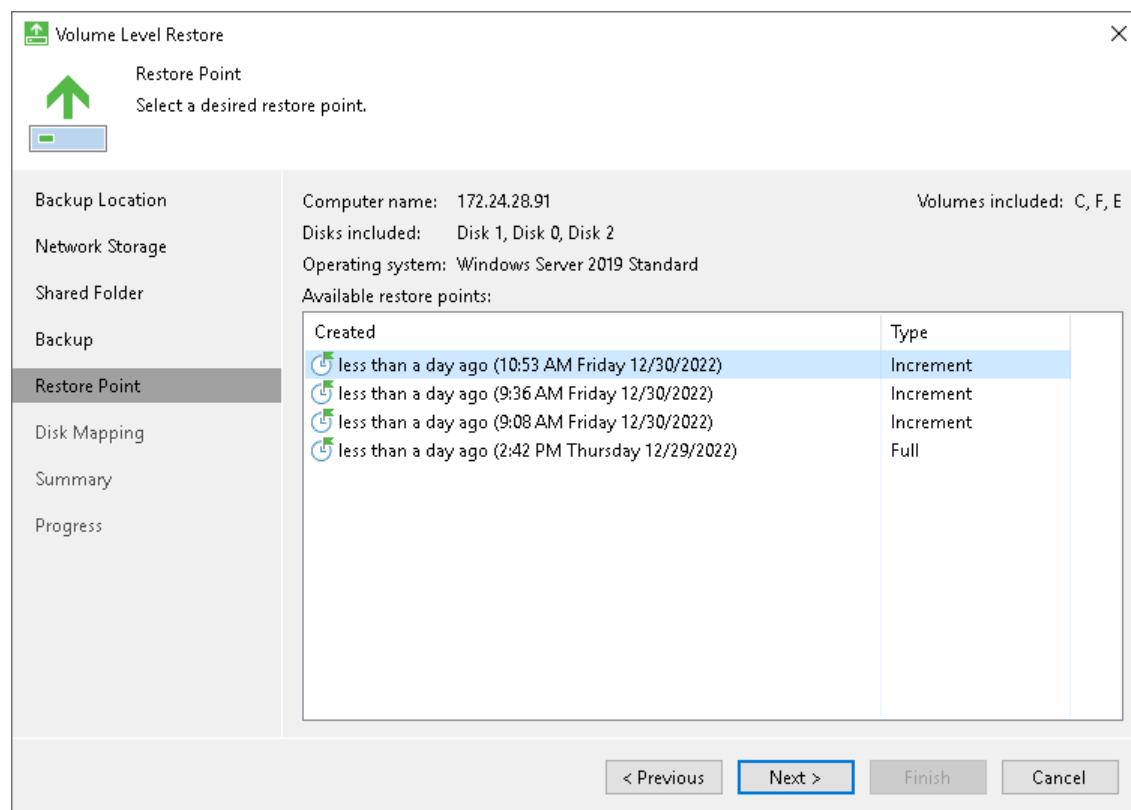
< Previous   **Next >**   Finish   Cancel

## Step 6. Select Restore Point

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

By default, Veeam Agent for Microsoft Windows uses the latest restore point. However, you can select any valid restore point to recover volumes to a specific point in time.

Veeam Agent for Microsoft Windows displays restore points for volume-level backups only. For example, if you have run 3 job sessions to create a backup of all computer volumes and then changed the backup scope to file-level backup, Veeam Agent for Microsoft Windows will display only 3 restore points in the list.





# Step 7. Map Restored Disks

At the **Disk Mapping** step of the wizard, select what volumes you want to restore and map volumes from the backup to volumes on your computer.

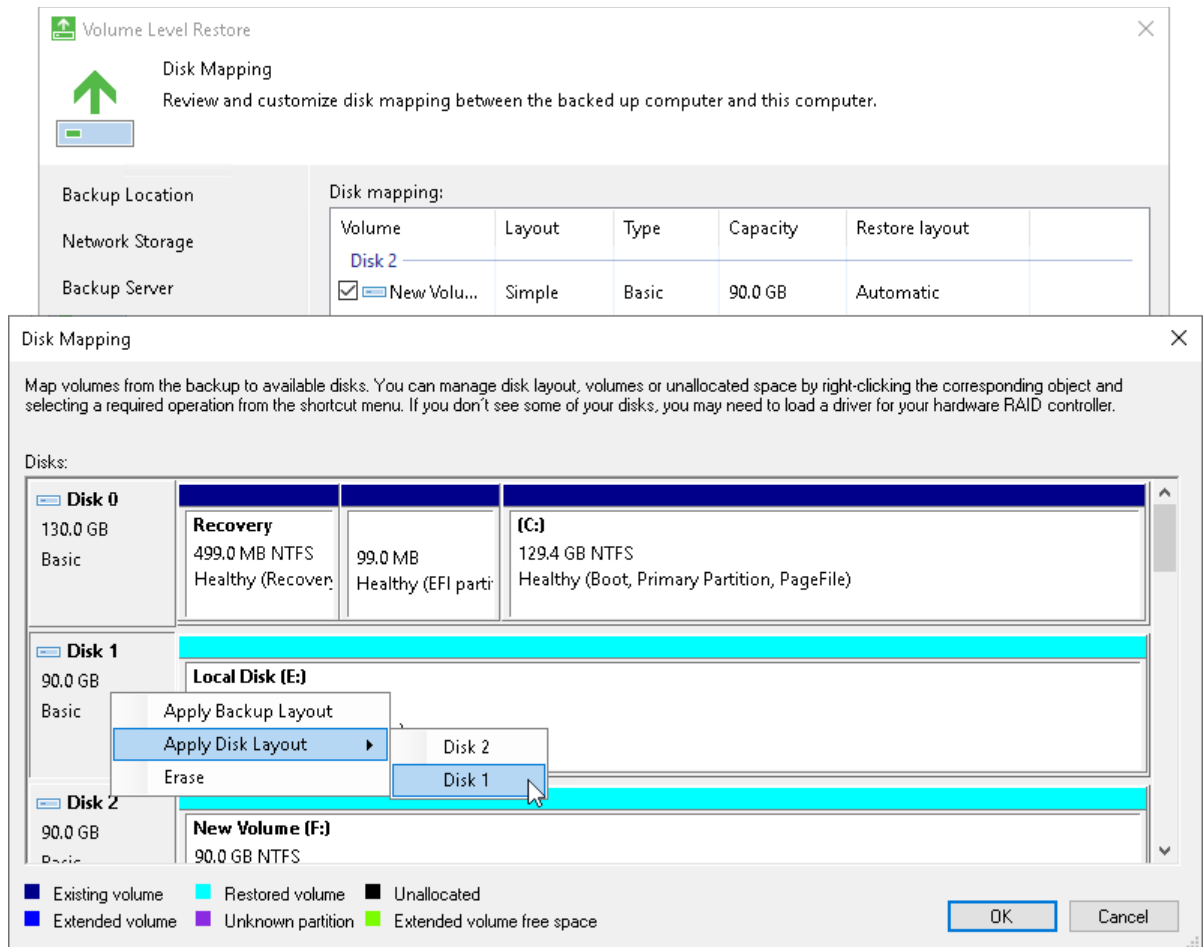
## IMPORTANT

We strongly recommend that you change disk mapping settings only if you have experience in working with Microsoft Windows disks and partitions. If you make a mistake, your computer data may get corrupted.

To select volumes to restore:

1. Select check boxes next to volumes that you want to restore from the backup.
2. [For restore to a new location] By default, Veeam Agent for Microsoft Windows restores volumes to their initial location. If the initial location is unavailable, a volume is restored to a disk of the same or larger size. To map the restored volume to another computer disk, at the bottom of the wizard click **Customize disk mapping**. In the **Disk Mapping** window, specify how volumes must be restored:
  - Right-click the target disk on the left and select the necessary disk layout:
    - **Apply Backup Layout** – select this option if you want to apply to disk the settings that were used on your computer at the moment when you performed backup.
    - **Apply Disk Layout** – select this option if you want to apply to the current disk settings of another disk.

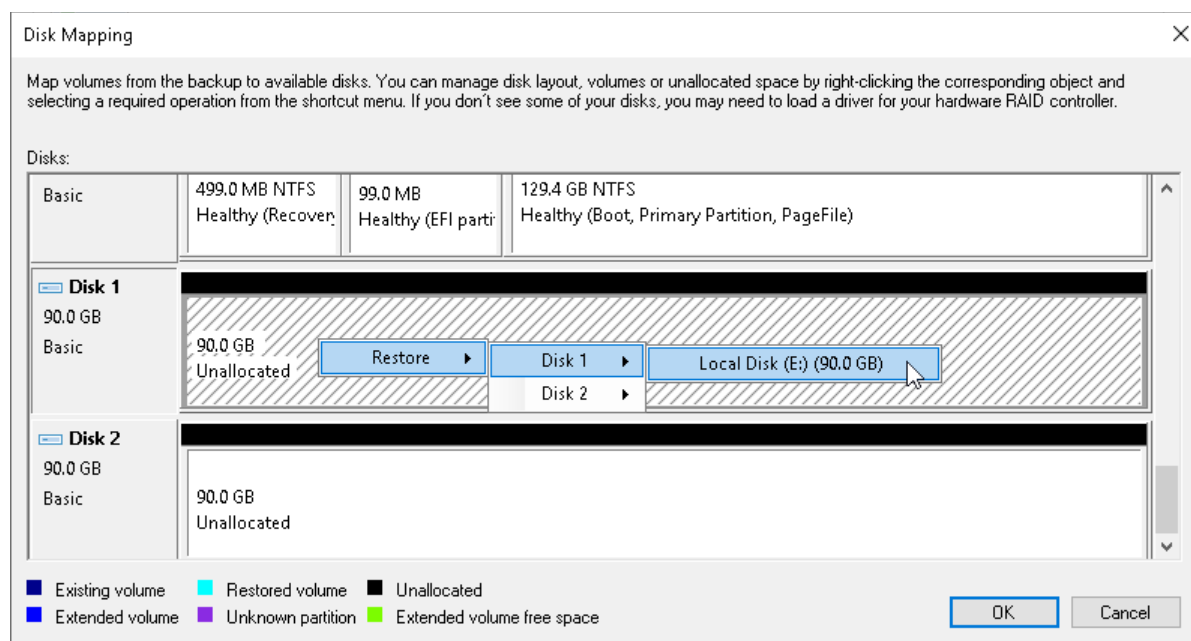
- **Erase** – select this option if you want to discard the current disk settings.



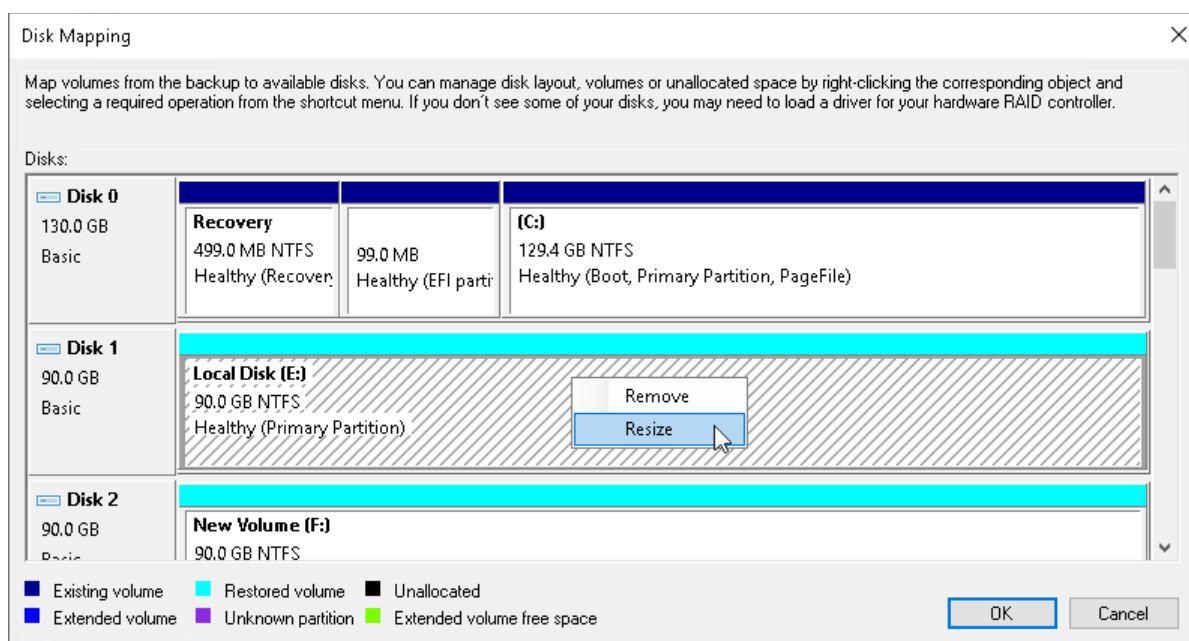
- Right-click unallocated disk space in the disk area on the right and select what volume from the backup you want to place on this computer disk.

If you restore to a dynamic disk, you can specify the type and size of the target volume before mapping. After you select the volume to place on the disk, you will pass to the [Allocate Volume](#) window.

If you want to change disk layout configured by Veeam Agent for Microsoft Windows, right-click an automatically mapped volume and select **Remove**. You will be able to use the released space for mapping volumes in your own order.



3. [For restore with volume resize] You can resize a volume mapped by Veeam Agent for Microsoft Windows to a target computer disk. To resize a volume, right-click it in the **Disk Mapping** window and select **Resize**. With this option selected, you will pass to the [Volume Resize](#) window.



## NOTE

If you map a backup volume that is larger than the amount of available space on the target disk, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume. After you agree and click **OK**, Veeam Agent for Microsoft Windows will prepare to shrink the volume to the size of available disk space.

# Allocating Restored Volumes to Dynamic Disks

At the **Disk Mapping** step of the wizard, you can set the necessary type and size for the restored volumes. A volume will get the specified type and size during the process of data restore.

You can do this only if all the following conditions are met:

- You restore data to a dynamic disk. To learn more about dynamic volume types, see [Microsoft documentation](#).
- At the **Restore Mode** step of the wizard, you have selected the **Manual restore** option.
- At the **Disk Mapping** step of the wizard, you have clicked the **Customize disk mapping** link.

To set volume type and size:

1. Right-click unallocated space of a dynamic disk and select what volume from the backup you want to place on this disk. Veeam Agent for Microsoft Windows will prompt you to set the restored volume type and size.
2. In the **Allocate** window, specify the restored volume type:
  - If you want to restore the volume as a simple volume, from the **Restore as** drop-down list, select the **Simple volume** option and click **OK**.
  - If you want to restore the volume as a spanned, striped, mirrored, or parity volume, from the **Restore as** drop-down list, select the desired option and specify the following settings:
    - i. In the **Available disks** list, select volume extents you want to restore. You can add or delete volume extents using the **Add**, **Remove**, and **Remove All** buttons.
    - ii. In the **Specify the amount of disk space to allocate** field, specify the total size of restored extents.
- iii. Click **OK**.

Allocate Volume Local Disk (E:)

Restore as: Span volume

Available disks:

- Disk 2
- Disk 1

Selected disks:

Add

Remove

Remove All

Total volume Local Disk (E:) size (MB): 92142

Total size of selected disks (MB): 0

Specify the amount of disk space to allocate (MB): 0

OK Cancel

## Step 8. Resize Restored Volumes

At the **Disk Mapping** step of the wizard you can set the necessary size for the restored volumes. You can resize a volume if you have chosen to restore data in the *Manual* mode and customize disk layout. A volume will be shrunk or extended to the specified size during the process of data restore.


### NOTE

By default, Veeam Agent for Microsoft Windows displays volume size in megabytes (MB). This allows you to specify the desired size for the volume precisely. You can also choose to display volume size in gigabytes (GB). This may be helpful when you need to resize volumes on larger computer disks and want to simplify disk size calculations.

When you use GB as a volume size unit, you can specify volume size with integral numbers, for example, 1 GB, 60 GB or 200 GB, but not 0,8 GB, 60,5 GB or 200,7 GB. However, if the maximum volume size is in fact greater than the displayed value for less than 1 GB, Veeam Agent for Microsoft Windows will automatically add the exceeding amount of disk space to the extended volume. For example, if the maximum volume size is 60,2 GB, Veeam Agent for Microsoft Windows will display this size as 60 GB. When you specify 60 GB as a desired volume size, Veeam Agent for Microsoft Windows will extend the volume to 60,2 GB.

To resize a volume:

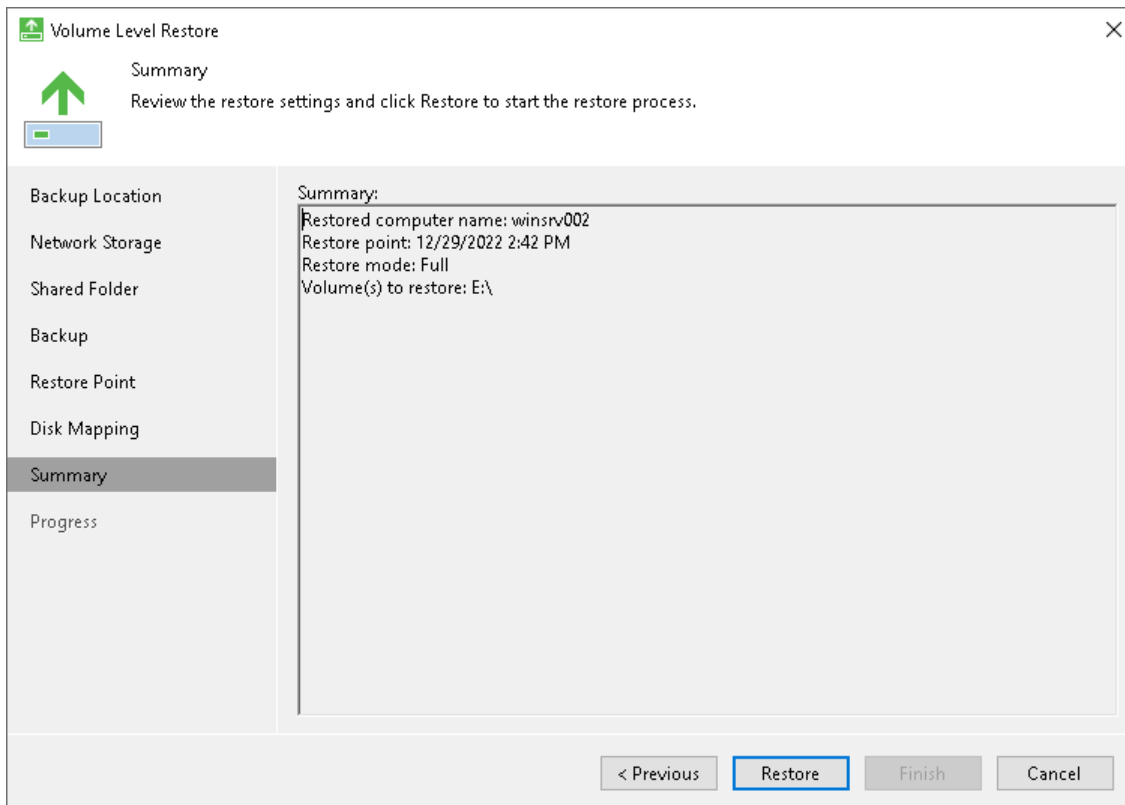
1. Specify a volume you want to resize:
  - a. Right-click a restored volume mapped to a target disk and select **Resize**.
  - b. [For volume shrink] Right-click unallocated disk space and select what volume from the backup you want to place on the computer disk. If the selected volume is larger than the amount of unallocated disk space, Veeam Agent for Microsoft Windows will prompt you to shrink the restored volume.
2. In the **Volume Resize** window, select the volume size unit and specify the desired size for the restored volume.



## Step 9. Complete Restore Process

At the **Summary** step of the wizard, complete the procedure of volume-level restore.

1. Review settings of the restore process.
2. Click **Restore** to start the recovery process. Veeam Agent for Microsoft Windows will perform partition re-allocation operations if necessary, restore the necessary volume data from the backup and overwrite volume data on your computer with the restored data.



# Restoring Files and Folders

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent for Microsoft Windows publishes the backup content directly into the computer file system and displays it in the Veeam Backup browser. You can restore files and folders to their initial location, copy files and folders to a new location or target applications to restored files and work with them as usual.

# Before You Begin

Before you begin the file-level restore process, check the following prerequisites.

## General

- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders and on backup repositories] You must have access to the target location where the backup file resides.
- [For backup repository targets] If you plan to restore data from a backup stored in a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

### NOTE

Consider the following:

- If you perform restore under an account that does not have administrative privileges, you might not be able to access certain files or folders due to the user's security context settings. To learn more about the security context, see [Microsoft documentation](#).
- When you restore files or folders, target EFS settings are applied to the restored objects. For example, if you restore files or folders to a folder encrypted with EFS, the restored objects are encrypted.

## ReFS Restore

If you restore files or folders from a ReFS volume-level backup that was made on another machine, the Veeam Agent computer where you perform restore must run the OS that supports the specific ReFS version. For example, you can restore files from a ReFS 3.x volume only to a computer that runs one of the following OSes:

- Microsoft Windows 10 version 1803 or later
- Microsoft Windows Server 2016 or later

## NTFS Restore

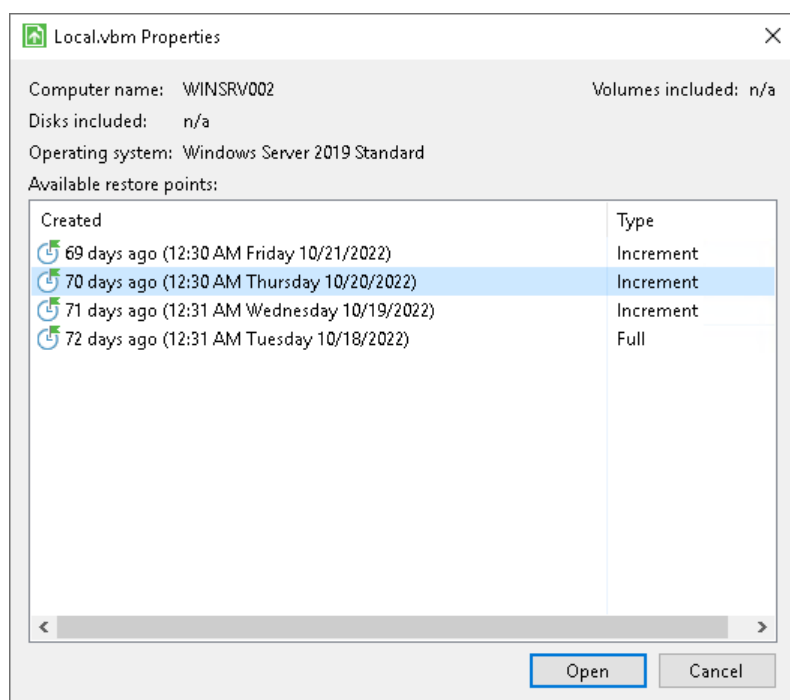
- If you restore files or folders from an NTFS volume-level backup that was made on another machine, and data deduplication is enabled for some volumes of the backup, the Veeam Agent computer where you perform restore must run the same OS version or later as the backed-up machine OS.
- If you restore files or folders from an NTFS volume-level backup that was made on another machine, and the allocation unit size of volumes is set to 128K or greater, the Veeam Agent computer where you perform restore must run one of the following OSes:
  - Microsoft Windows 10 version 1803 or later
  - Microsoft Windows Server 2019 or later



# Step 1. Launch File Level Restore Wizard

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

- Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Restore > Individual files**.
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the control panel, click a bar of the necessary backup job session. Click **Restore Files** at the bottom of the window. Veeam Agent for Microsoft Windows will automatically publish the backup content into the computer file system and [open the Veeam Backup browser](#).
- Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**. In the main menu, hover over the name of the job that created the backup from which you want to restore data, and select **Restore file**.
- From the Microsoft Windows **Start** menu, select **All Programs > Veeam > File Level Restore**.
- In Microsoft Windows Explorer, double-click the necessary VBK or VBM file or right-click the file and select **Extract**. In the displayed window, select the restore point from which you want to recover files and click **Open**. Veeam Agent for Microsoft Windows will automatically publish the backup content into the computer file system and [open the Veeam Backup browser](#).

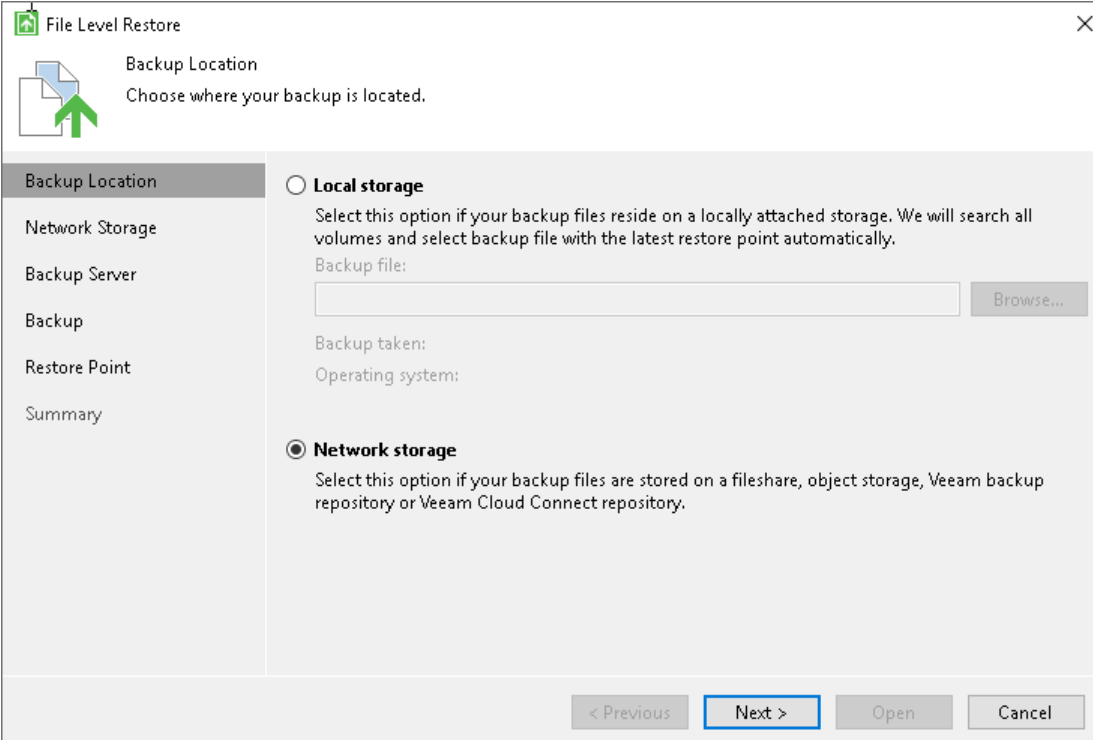


## Step 2. Specify Backup File Location

At the **Backup Location** step of the wizard, specify where the backup file that you plan to use for restore resides.

If you configured only one backup job, Veeam Agent automatically locates the latest backup of this backup job on the target storage, and you pass immediately to the [Restore Point](#) step of the wizard. If there are several backup jobs configured or Veeam Agent fails to locate the backup for some reason, specify where the backup file resides:

- **Local storage** – select this option if the backup file resides on the computer drive, external drive or removable storage device that is currently connected to your computer. Click **Browse** and select a backup metadata file (VBM).
- **Network storage** – select this option if the backup file resides in object storage, in a network shared folder, in a backup repository managed by a Veeam backup server or in a cloud repository exposed to you by a Veeam Cloud Connect service provider. In this case, the Veeam Recovery Media wizard will include additional steps for specifying the backup file location settings.



The screenshot shows the 'File Level Restore' wizard window. The title bar says 'File Level Restore'. Inside, there's a 'Backup Location' section with a document icon and a green arrow pointing up. Below this is a list of steps: 'Backup Location', 'Network Storage', 'Backup Server', 'Backup', 'Restore Point', and 'Summary'. The 'Backup Location' step is selected. The main area has two radio buttons: 'Local storage' (unselected) and 'Network storage' (selected). The 'Local storage' option has a description: 'Select this option if your backup files reside on a locally attached storage. We will search all volumes and select backup file with the latest restore point automatically.' It also has a 'Backup file:' label, a text input field, and a 'Browse...' button. Below this are labels for 'Backup taken:' and 'Operating system:'. The 'Network storage' option has a description: 'Select this option if your backup files are stored on a fileshare, object storage, Veeam backup repository or Veeam Cloud Connect repository.' At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Open', and 'Cancel'.

## Step 3. Select Remote Storage Type

The **Remote Storage** step of the wizard is available if you have chosen to restore data from a backup file that resides in a remote location — in object storage, in a network shared folder, in a backup repository or in a cloud repository.

Specify where the backup file resides:

- **Object storage** — select this option if the backup file is located in object storage. With this option selected, you will pass to the [Object Storage](#) step of the wizard.
- **Shared folder** — select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the [Shared Folder](#) step of the wizard.
- **Veeam backup repository** — select this option if the backup file resides in a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Backup Server](#) step of the wizard.
- **Veeam Cloud Connect repository** — select this option if the backup file resides in a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the [Service Provider](#) step of the wizard.

## Step 4. Specify Remote Storage Settings

Specify settings for the remote storage that contains a backup file from which you plan to restore data:

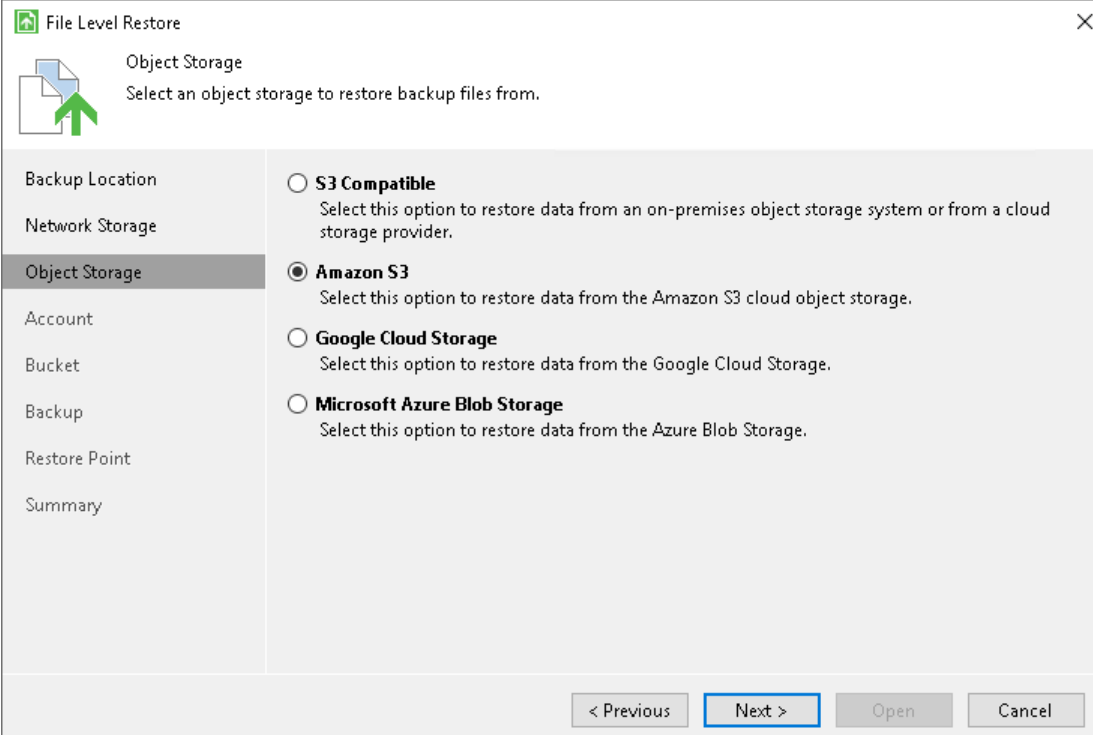
- [Object storage settings](#) – if you have selected the **Object storage** option at the [Remote Storage](#) step of the wizard.
- [Shared folder settings](#) – if you have selected the **Shared folder** option at the [Remote Storage](#) step of the wizard.
- [Veeam backup repository settings](#) – if you have selected the **Veeam backup repository** option at the [Remote Storage](#) step of the wizard.
- [Veeam Cloud Connect repository settings](#) – if you have selected the **Veeam Cloud Connect repository** option at the [Remote Storage](#) step of the wizard.

### Object Storage Settings

The **Object Storage** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

Specify settings for the object storage:

- [Specify S3 compatible settings.](#)
- [Specify Amazon S3 settings.](#)
- [Specify Google Cloud Storage settings.](#)
- [Specify Microsoft Azure Blob Storage settings.](#)



The screenshot shows the 'File Level Restore' wizard window. The title bar says 'File Level Restore'. Inside, there's a sub-header 'Object Storage' with a description: 'Select an object storage to restore backup files from.' Below this is a list of options on the right and a sidebar on the left. The sidebar has a tree view with 'Object Storage' selected. The options on the right are: 'S3 Compatible' (unselected), 'Amazon S3' (selected), 'Google Cloud Storage' (unselected), and 'Microsoft Azure Blob Storage' (unselected). At the bottom, there are four buttons: '< Previous', 'Next >', 'Open', and 'Cancel'.

**File Level Restore**

**Object Storage**  
Select an object storage to restore backup files from.

**Backup Location**

- Network Storage
- Object Storage**
- Account
- Bucket
- Backup
- Restore Point
- Summary

☐ **S3 Compatible**  
Select this option to restore data from an on-premises object storage system or from a cloud storage provider.

☒ **Amazon S3**  
Select this option to restore data from the Amazon S3 cloud object storage.

☐ **Google Cloud Storage**  
Select this option to restore data from the Google Cloud Storage.

☐ **Microsoft Azure Blob Storage**  
Select this option to restore data from the Azure Blob Storage.

< Previous   **Next >**   Open   Cancel

## S3 Compatible Settings

If you have selected to restore data from a backup file located in the S3 compatible storage, specify the following settings:

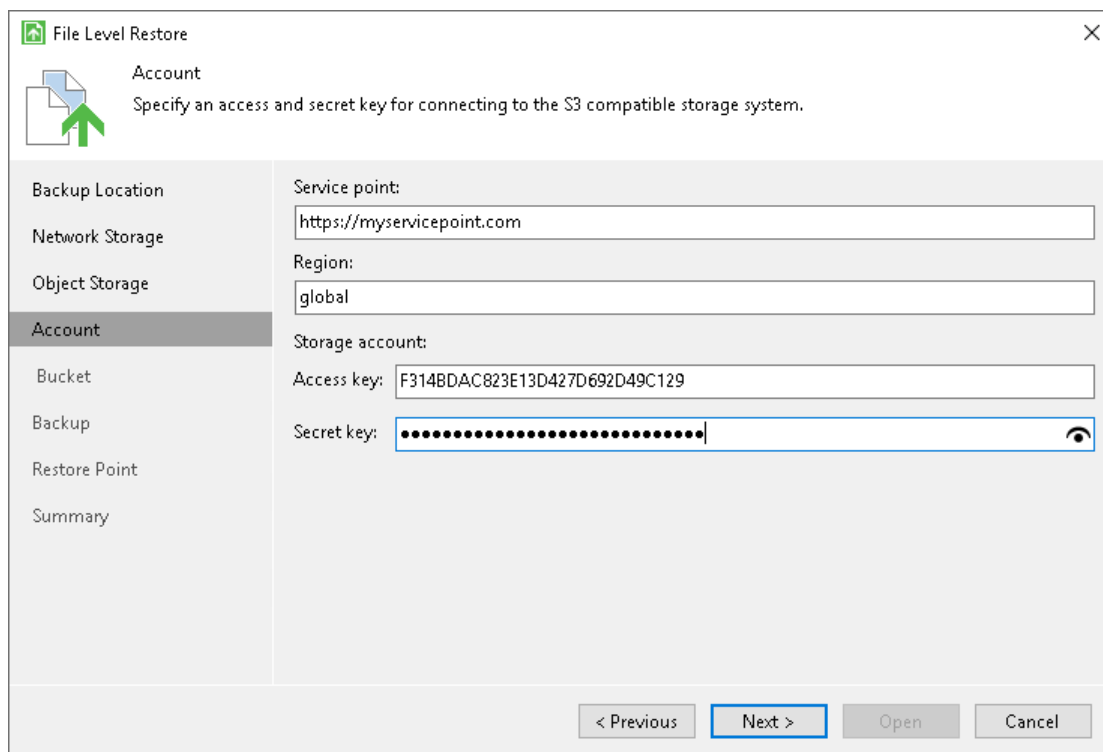
1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.
2. In the **Region** field, specify the storage region.
3. In the **Access key** field, enter the access key ID.
4. In the **Secret key** field, enter the secret access key.



The screenshot shows the 'File Level Restore' wizard window. The title bar says 'File Level Restore' with a close button. The main window has a left sidebar with a tree view containing: Backup Location, Network Storage, Object Storage, Account (selected), Bucket, Backup, Restore Point, and Summary. The main area is titled 'Account' with a subtitle 'Specify an access and secret key for connecting to the S3 compatible storage system.' Below this, there are four input fields: 'Service point' with the value 'https://myservicepoint.com', 'Region' with the value 'global', 'Storage account:' (empty), 'Access key:' with the value 'F314BDAC823E13D427D692D49C129', and 'Secret key:' with a masked value represented by dots. At the bottom right, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Open', and 'Cancel'.

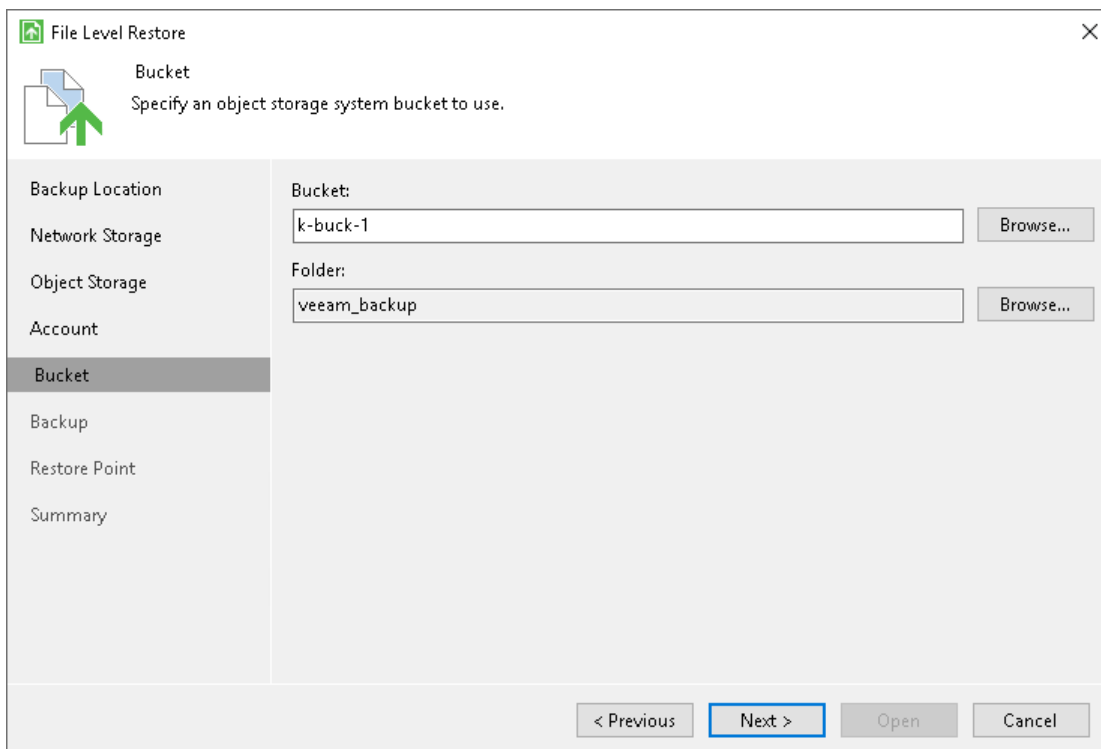
## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.

- b. In the **Select Bucket** window, do the following:
      - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
      - ii. Select the necessary bucket and click **OK**.
  2. In the **Folder** field, specify a folder in the bucket:
    - a. Select the **Browse** option.
    - b. In the **Select Folder** window, do the following:
      - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.
      - ii. Select the necessary folder and click **OK**.



## Amazon S3 Settings

If you have selected to restore data from a backup file located in the Amazon S3 storage, specify the following settings:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter the access key ID.

2. In the **Secret key** field, enter the secret access key.
3. From the **AWS region** drop-down list, select the AWS region. By default, Veeam Agent uses the **Global** region.

**File Level Restore**

**Account**  
Specify an Amazon access and secret key for connecting to the Amazon S3 storage.

**Backup Location**

**Network Storage**

**Object Storage**

**Account**

**Bucket**

**Backup**

**Restore Point**

**Summary**

Amazon AWS account:

Access key: RTWGHGT6BDRHKYT9SE1

Secret key: [Masked]

AWS region: Global

< Previous   Next >   Open   Cancel

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. From the **Data center** drop-down list, select the geographic region where the backup file is stored.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.

ii. Select the necessary folder and click **OK**.

**File Level Restore**

Bucket  
Specify an Amazon S3 bucket to use.

Backup Location  
Network Storage  
Object Storage  
Account  
**Bucket**  
Backup  
Restore Point  
Summary

Data center:  
EU (Paris)

Bucket:  
buck-1 **Browse...**

Folder:  
system\_backup **Browse...**

< Previous **Next >** Open Cancel

## Google Cloud Storage Settings

If you have selected to restore data from a backup file located in the Google Cloud storage, specify the following settings:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.



To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) keys associated with the Google Cloud account. Veeam Agent will use the HMAC keys to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see [Google Cloud documentation](#).

**File Level Restore**

**Account**  
Specify a Google Cloud access and secret key for connecting to the Google Cloud Storage.

**Backup Location**  
**Network Storage**  
**Object Storage**  
**Account**  
 Bucket  
 Backup  
 Restore Point  
 Summary

**Google Cloud account:**  
 Access key: GOOG1TVPOQJ6KLUQ1LCBDJKSQYUISLAJPQTA2PNDHTEUISFIWQKZr  
 Secret key: [Masked] [Eye icon]

< Previous   **Next >**   Open   Cancel

## Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to restore data from a backup file located in object storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. From the **Data center region** drop-down list, select the geographic region where the backup file is stored.
2. In the **Bucket** field, specify a bucket in the storage:
  - a. Select the **Browse** option.
  - b. In the **Select Bucket** window, do the following:
    - i. Double-click the region name or click the arrow to the left of the region name to view the list of available buckets.
    - ii. Select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the bucket name or click the arrow to the left of the bucket name to view the list of available folders.

ii. Select the necessary folder and click **OK**.

**File Level Restore**

Bucket  
Specify a Google Cloud storage bucket to use.

Backup Location  
Network Storage  
Object Storage  
Account  
**Bucket**  
Backup  
Restore Point  
Summary

Data center region:  
Belgium

Bucket:  
west1 **Browse...**

Folder:  
veeam\_backup **Browse...**

< Previous **Next >** Open Cancel

## Microsoft Azure Blob Storage Settings

If you have selected to restore data from a backup file located in the Microsoft Azure Blob storage, specify the following settings:

1. [Specify account settings](#).
2. [Specify container settings](#).

## Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to restore data from a backup file located in object storage.

To connect to the Microsoft Azure Blob storage, specify the following:

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.

3. From the **Region** drop-down list, select the Microsoft Azure region. By default, Veeam Agent uses the **Azure Global (Standard)** region.

**File Level Restore**

**Account**  
Specify a Microsoft Azure account and shared key for connecting to the Microsoft Azure Blob Storage.

**Backup Location**  
**Network Storage**  
**Object Storage**  
**Account**  
**Container**  
**Backup**  
**Restore Point**  
**Summary**

Microsoft Azure Blob storage account:

Account: newstorage

Shared key: ..... 🔍

Region: Azure Global (Standard) ▼

< Previous   **Next >**   Open   Cancel

## Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to restore data from a backup file located in the Microsoft Azure Blob storage and specified account settings to connect to the storage.

Specify settings for the container in the storage:

1. From the **Container** drop-down list, select a container in the storage.
2. In the **Folder** field, specify a folder in the container:
  - a. Select the **Browse** option.
  - b. In the **Select Folder** window, do the following:
    - i. Double-click the container name or click the arrow to the left of the container name to view the list of available folders.

ii. Select the necessary folder and click **OK**.

The screenshot shows the 'File Level Restore' wizard window. The title bar says 'File Level Restore'. The main heading is 'Container' with the instruction 'Specify a Microsoft Azure Blob Storage container to use.' Below this, there is a sidebar on the left with a tree view containing: Backup Location, Network Storage, Object Storage, Account, Container (selected), Backup, Restore Point, and Summary. The main area has two input fields: 'Container:' with a dropdown menu showing 'veeam' and 'Folder:' with a text box containing 'veeam\_backup' and a 'Browse...' button. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Open', and 'Cancel'.

## Shared Folder Settings

The **Shared Folder** step of the wizard is available if you have chosen to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. In the **Shared folder** field, type in a UNC name of the network shared folder with the backup file. Keep in mind that the UNC name always starts with two back slashes (\\).
2. If the network shared folder requires authentication, select the **This share requires access credentials** check box and specify a user name and password of the account that has *Full Control* access permissions on this shared folder. The user name must be specified in the [down-level logon name](#) format. For example, *DOMAIN\UserName* or *HOSTNAME\UserName*.

If you do not select the **This share requires access credentials** check box, Veeam Agent for Microsoft Windows will connect to the shared folder using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed.

3. To view the entered password, click and hold the eye icon on the right of the **Password** field.

The screenshot shows the 'File Level Restore' wizard window. On the left is a sidebar with navigation options: 'Backup Location', 'Network Storage', 'Shared Folder' (which is selected and highlighted), 'Backup', 'Restore Point', and 'Summary'. The main area of the window is titled 'Shared Folder' and contains the instruction: 'Specify a UNC path to a shared folder, and credentials (if required for share access)'. Below this, there is a 'Shared folder:' text box containing the path '\\WINSRV001\SharedBackups' and a 'Browse...' button to its right. Underneath, a checkbox labeled 'This share requires access credentials:' is checked. Below the checkbox are two text boxes: 'Username:' containing 'TECH\Administrator' and 'Password:' containing a series of dots. An eye icon is positioned to the right of the password field. At the bottom of the window are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Open', and 'Cancel'.

## Backup Server Settings

The **Backup Server** step of the wizard is available if you have chosen to restore data from a backup file located in a backup repository.

Specify settings for the Veeam backup server that manages the backup repository:

1. In the **Veeam backup server name or IP address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent for Microsoft Windows must communicate with the backup server. By default, Veeam Agent for Microsoft Windows uses port 10001.
3. Select the **Specify your personal credentials** check box. In the **Username** and **Password** fields, enter a user name and password of the account that has access to this backup repository. Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

If you do not select the **Specify your personal credentials** check box, Veeam Agent for Microsoft Windows will connect to the backup repository using the *NT AUTHORITY\SYSTEM* account of the computer where the product is installed. You can use this scenario if the Veeam Agent computer is joined to the Active Directory domain. In this case, you can add the computer account (*DOMAIN\COMPUTERNAME\$*) to an AD group and grant access rights on the backup repository to this group.

Setting access permissions on the backup repository to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). If you have set such permissions on the backup repository, you can omit specifying credentials. However, we recommend this scenario for demo environments only.

The screenshot shows the 'File Level Restore' wizard window. The 'Backup Server' step is selected in the left sidebar. The main area contains the following fields and options:

- Backup Location:** Network Storage
- Backup Server:** Veeam backup server name or IP address: 172.24.33.10, Port: 10001
- ☒ Specify your personal credentials:
  - Username: TECH\Administrator
  - Password: [masked]

At the bottom, there are four buttons: '< Previous', 'Next >', 'Open', and 'Cancel'.

## Service Provider Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings.](#)
2. [Verify the TLS certificate and specify user account settings.](#)

## Specifying Service Provider Settings

The **Service provider** step of the wizard is available if you have chosen to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.

2. In the **Port** field, specify the port over which Veeam Agent for Microsoft Windows will communicate with the cloud gateway. By default, port 6180 is used.

The screenshot shows the 'File Level Restore' wizard window. The 'Service Provider' step is selected in the left sidebar. The main area contains a text box for 'DNS name or IP address' with the value '172.24.38.39' and a 'Port' spinner box set to '6180'. A note below states: 'Default service provider's port is 6180. If connection cannot be established, contact your service provider to make sure the settings are correct.' The bottom of the window has buttons for '< Previous', 'Next >', 'Open', and 'Cancel'.

File Level Restore

Service Provider  
Specify a DNS name or IP address and a port number received from the service provider.

Backup Location  
Network Storage  
**Service Provider**  
Credentials  
Backup  
Restore Point  
Summary

DNS name or IP address: 172.24.38.39 Port: 6180

Default service provider's port is 6180. If connection cannot be established, contact your service provider to make sure the settings are correct.

< Previous Next > Open Cancel

## Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the service provider.

1. At the top of the wizard window, Veeam Agent for Microsoft Windows displays information about the TLS certificate obtained from the SP side. You can view the certificate settings and verify the TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

- To view the TLS certificate, click the certificate link.
- To verify if the TLS certificate with a thumbprint, copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Fingerprint for certificate verification** field. Click **Verify**. Veeam Agent for Microsoft Windows will check if the thumbprint you enter matches the thumbprint of the obtained TLS certificate.

2. In the **Username** field, enter the user name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The user name of the subtenant account must be specified in the *TENANT\SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.

**File Level Restore**

**Credentials**  
Specify credentials that you have received from the service provider and validate certificate.

Backup Location  
Network Storage  
Service Provider  
**Credentials**  
Backup  
Restore Point  
Summary

This certificate has been validated.  
Verified by: [CN=Veeam Software, O=Veeam Software, OU=Veeam Software](#)

Username:

Password:

< Previous   **Next >**   Open   Cancel



## Step 5. Select Backup

The **Backup** step of the wizard is available if you have chosen to restore data from a backup file that resides in a remote location — in object storage, in a network shared folder, in a backup repository or in a cloud repository.

From the list of backups, select a backup from which you want to recover data. To quickly find the necessary backup, use the search field at the bottom of the window: enter a backup name or a part of it in the search field and click the **Start search** button on the right or press **[ENTER]**.

If you restore data from a backup stored in the backup repository, Veeam Agent for Microsoft Windows displays only those backups that are accessible by the user whose credentials are specified at the [Backup Server](#) step of the wizard:

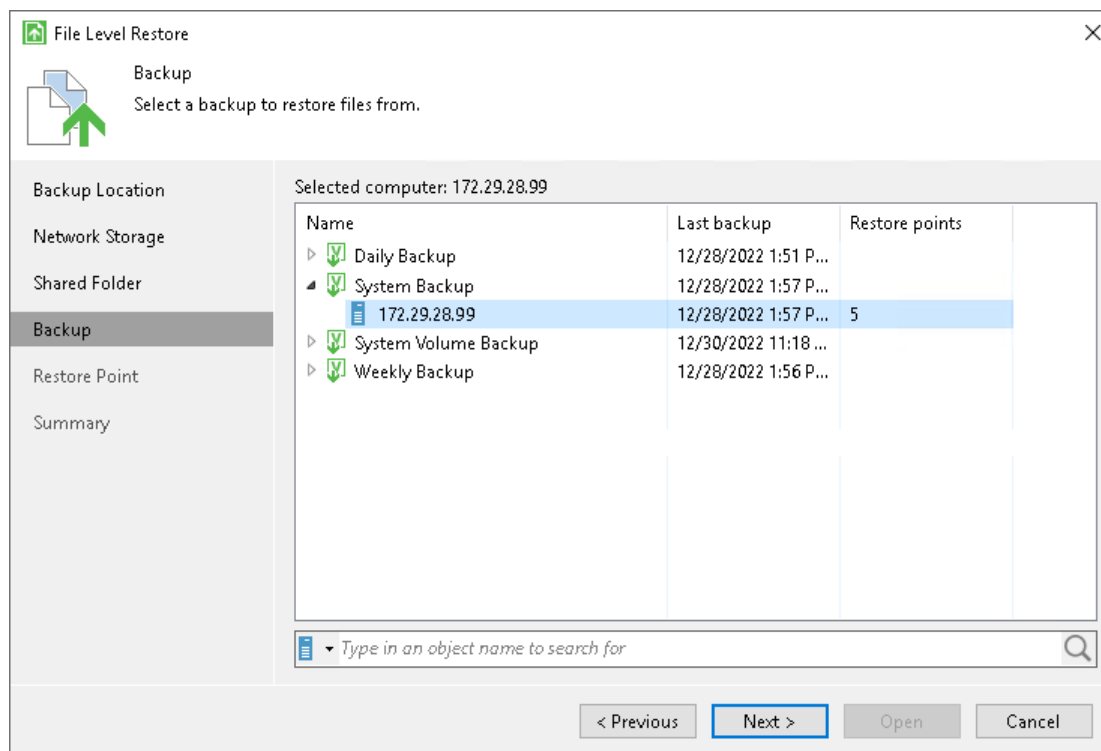
- If you specify credentials for the user who has access to the backup repository, the list of backups will include only backups created by this user.
- If you specify credentials for the user who is assigned the *Backup Administrator* or *Restore Operator* role on the backup server, the list of backups will include all Veeam Agent backups stored in the backup repository.

If you restore data from a backup stored in the cloud repository, Veeam Agent for Microsoft Windows displays only those backups that are accessible by the user whose credentials are specified at the [Credentials](#) step of the wizard:

- If you specify credentials for the tenant account, the list of backups will include backups created by all users who create backups under this account.
- If you specify credentials for the subtenant account, the list of backups will include only those Veeam Agent backups that were created under this subtenant account.

## NOTE

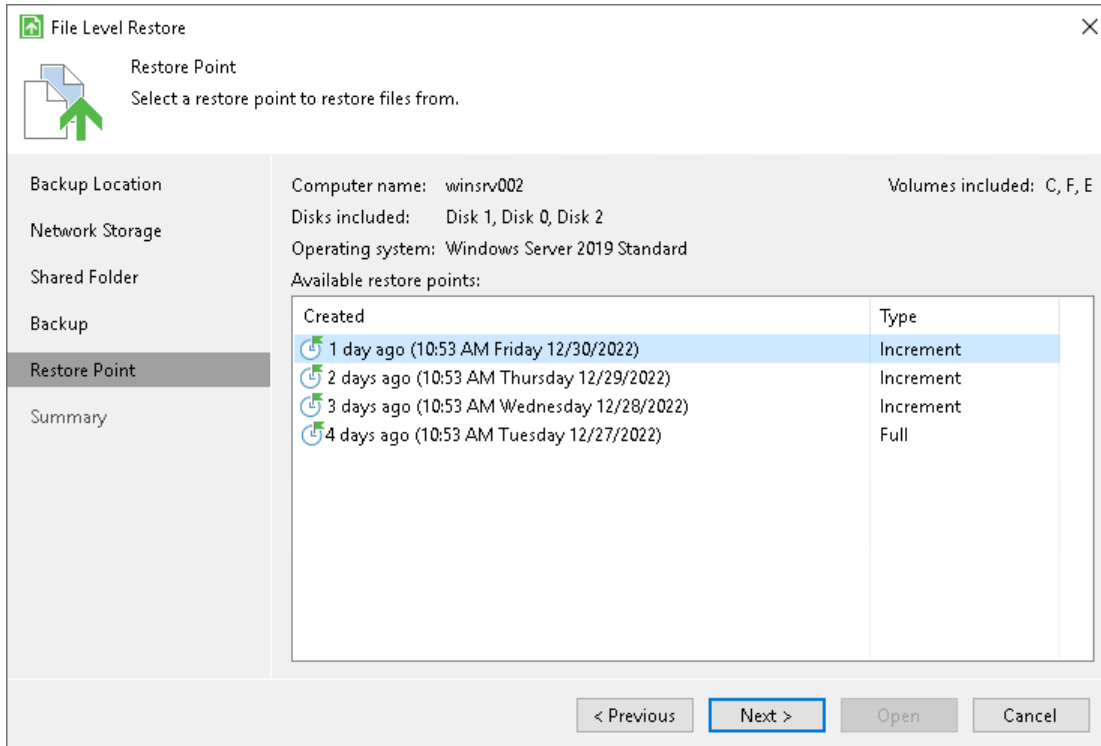
If you restore data from an encrypted backup that was created on another Veeam Agent computer, you need to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).



## Step 6. Select Restore Point

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

By default, Veeam Agent for Microsoft Windows uses the latest restore point. However, you can select any valid restore point to recover files and folders to a specific point in time.



**File Level Restore**

**Restore Point**  
Select a restore point to restore files from.

**Backup Location**  
Network Storage  
Shared Folder  
Backup  
**Restore Point**  
Summary

Computer name: winsrv002  
Disks included: Disk 1, Disk 0, Disk 2  
Operating system: Windows Server 2019 Standard  
Volumes included: C, F, E

Available restore points:

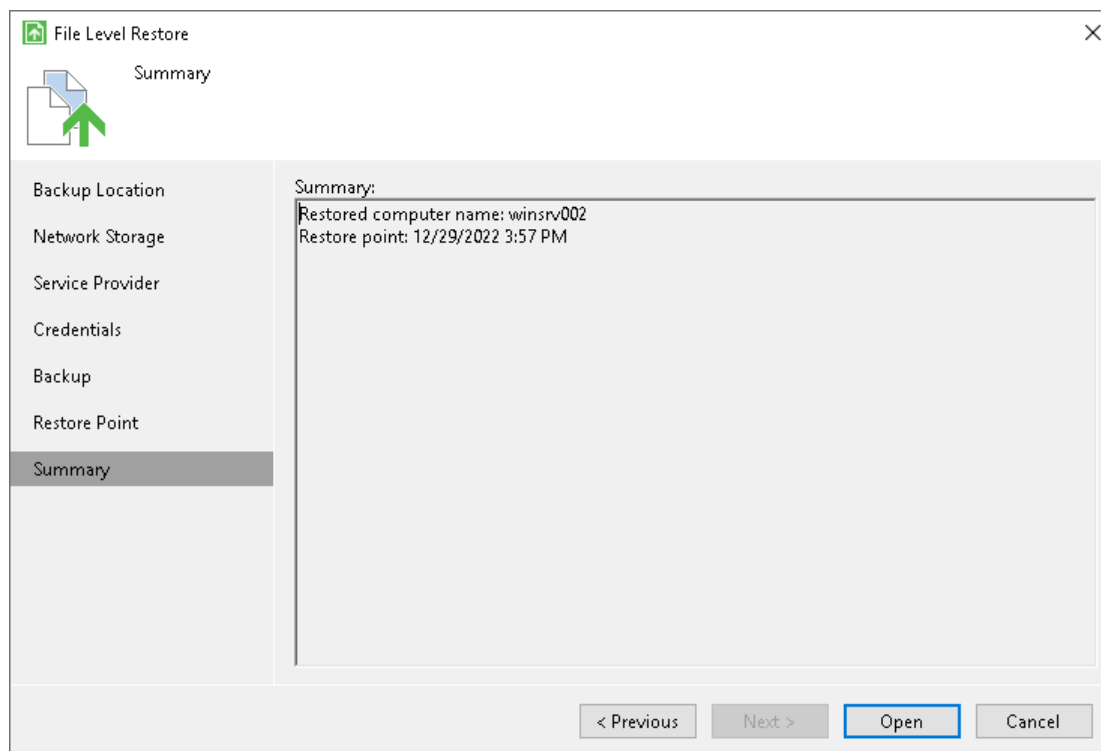
Created	Type
1 day ago (10:53 AM Friday 12/30/2022)	Increment
2 days ago (10:53 AM Thursday 12/29/2022)	Increment
3 days ago (10:53 AM Wednesday 12/28/2022)	Increment
4 days ago (10:53 AM Tuesday 12/27/2022)	Full

< Previous   **Next >**   Open   Cancel

## Step 7. Complete Restore Process

At the **Summary** step of the wizard, complete the procedure of file-level restore.

1. Review settings of the restore process.
2. Click **Open**. Veeam Agent for Microsoft Windows will retrieve the content of the backup file, publish it directly into the file system of your computer and display it in the Veeam Backup browser.



## Step 8. Save Restored Files

When the restore process is complete, Veeam Agent for Microsoft Windows opens the Veeam Backup browser displaying the content of the backup file.

You can perform the following operations with restored files and folders:

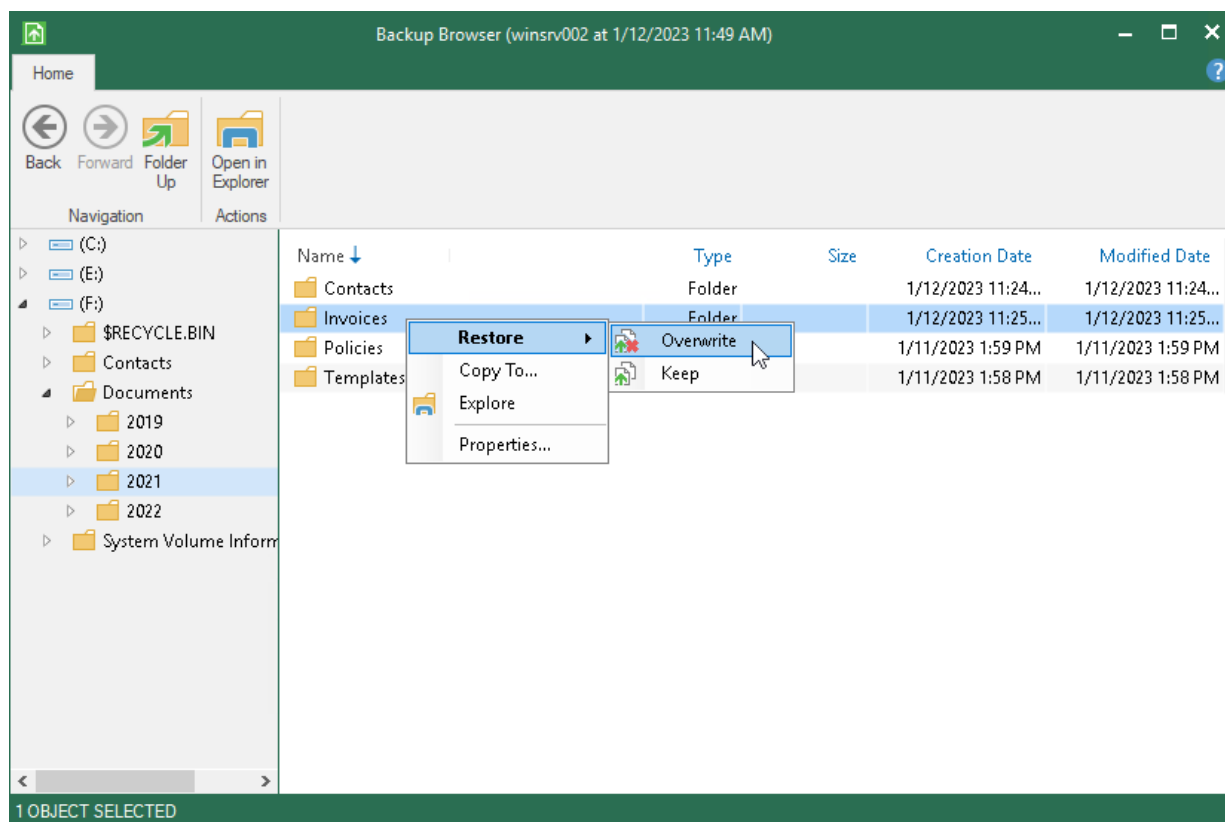
- [Save files and folders to their initial location](#)
- [Save files and folders to a new location](#)
- [Open files in Microsoft Windows Explorer](#)

After you finish working with files and folders, [close the Veeam Backup browser](#).

### Saving Files to Initial Location

To save restored files or folders to their initial location, right-click the necessary item in the file system tree or in the details pane on the right and select one of the following commands:

- To overwrite the original item on your computer with the item restored from the backup, select **Restore > Overwrite**.
- To save the item restored from the backup next to the original item on your computer, select **Restore > Keep**. Veeam Agent for Microsoft Windows will add the *RESTORED-* prefix to the restored file or folder name and save it in the same location where the original file resides.



# Saving Files to New Location

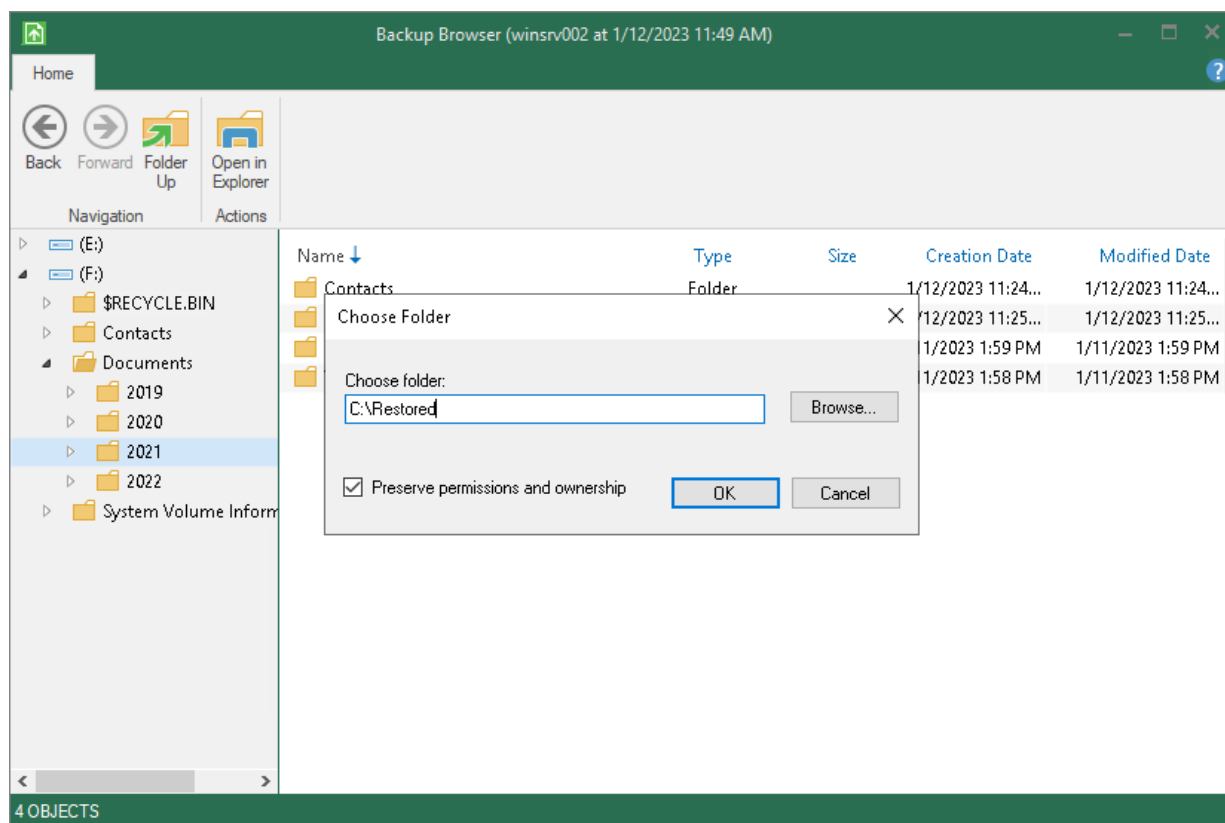
To save restored files or folders on your computer or to a network shared folder, right-click the necessary item in the file system tree or in the details pane on the right and select **Copy To**.

When restoring file objects, you can choose to preserve their original NTFS permissions:

- Select the **Preserve permissions and ownership** check box to keep the original ownership and security permissions for restored items. Veeam Agent for Microsoft Windows will copy selected files or folders with associated Access Control Lists, preserving granular access settings.

If access settings of a file or folder that you want to restore are inherited from a parent folder, when you restore this file or folder without the parent folder, its access settings will not be preserved.

- Leave the **Preserve permissions and ownership** check box not selected if you do not want to preserve the original ownership and access settings for restored items. Veeam Agent for Microsoft Windows will change security settings: the user who launched the Veeam Agent for Microsoft Windows will be set as the owner of the restored items. Access permissions will be inherited from the folder to which the restored items are copied.

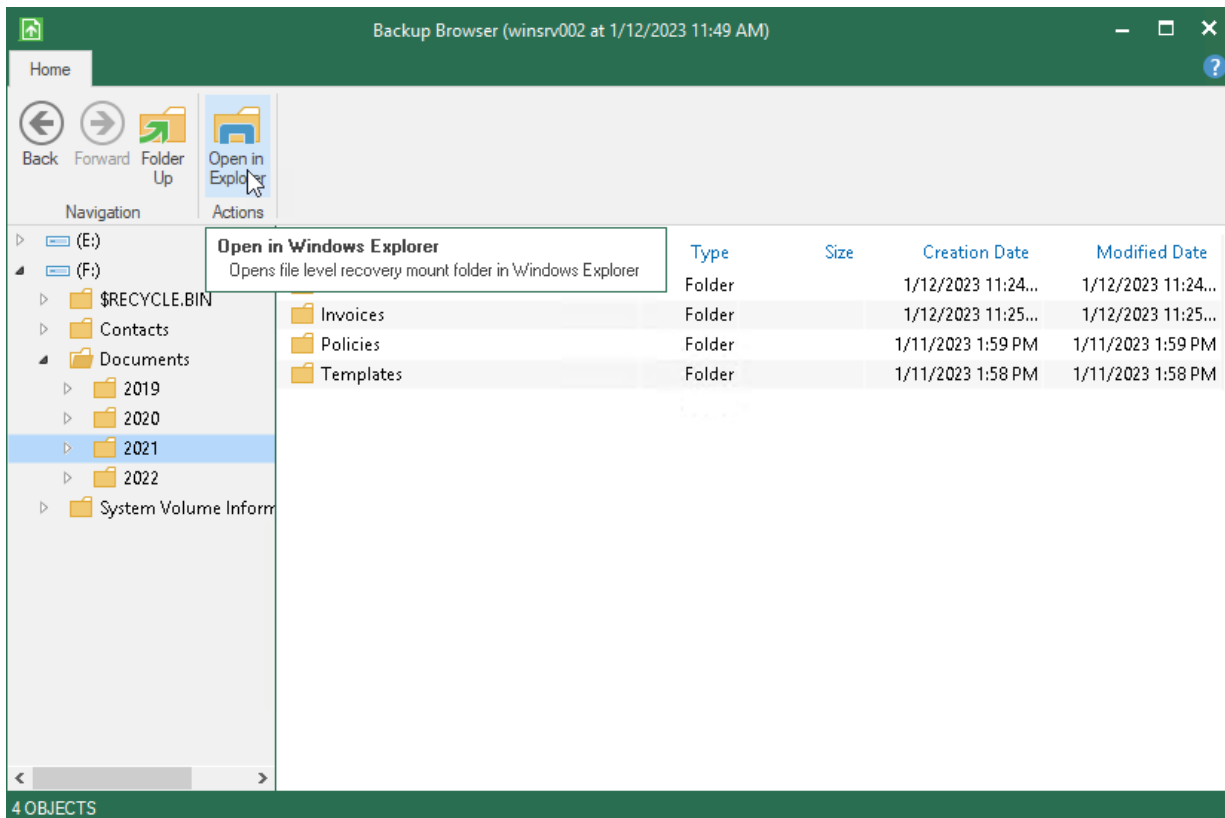


## Working with Windows Explorer

You can use Microsoft Windows Explorer to work with restored files and folders. To do this, do either of the following:

- In Veeam Backup browser, select the necessary file or folder and click **Open in Explorer** on the toolbar. Veeam Agent for Microsoft Windows will open the selected folder or file in Microsoft Windows Explorer.

- Open Microsoft Windows Explorer and browse for restored files and folders. The backup content is mounted under the `C:\VeeamFLR\ServerName` folder.



We recommend that you use the Veeam Backup browser instead of Microsoft Windows Explorer for file-level restore. Use of the Veeam Backup browser has the following advantages:

1. You can browse the guest OS file system ignoring the file system ACL settings.
2. You can preserve permissions and ownership during file-level restore.

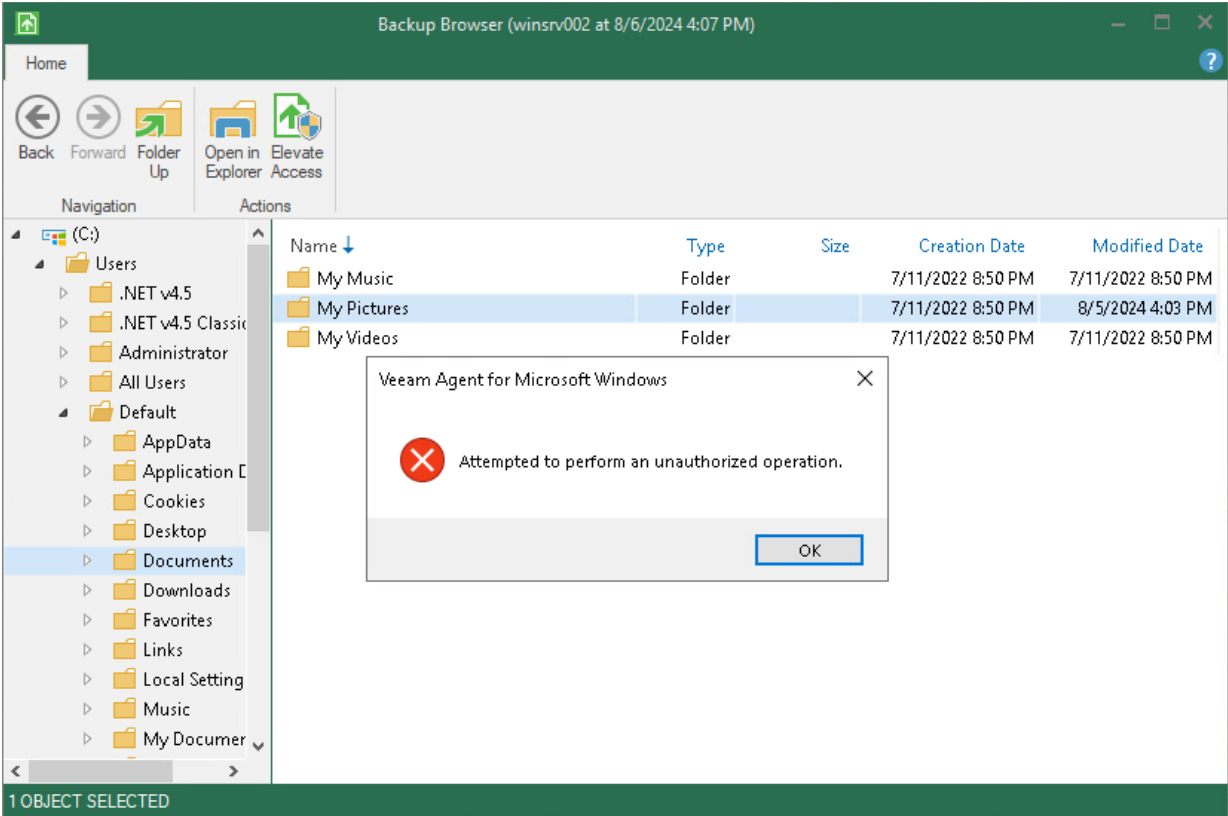
If you open the file system through the Microsoft Windows Explorer, these capabilities will not be available.

To learn more, see *SeBackupPrivilege* and *SeRestorePrivilege* in [Microsoft documentation](#).

# Elevating Access Rights

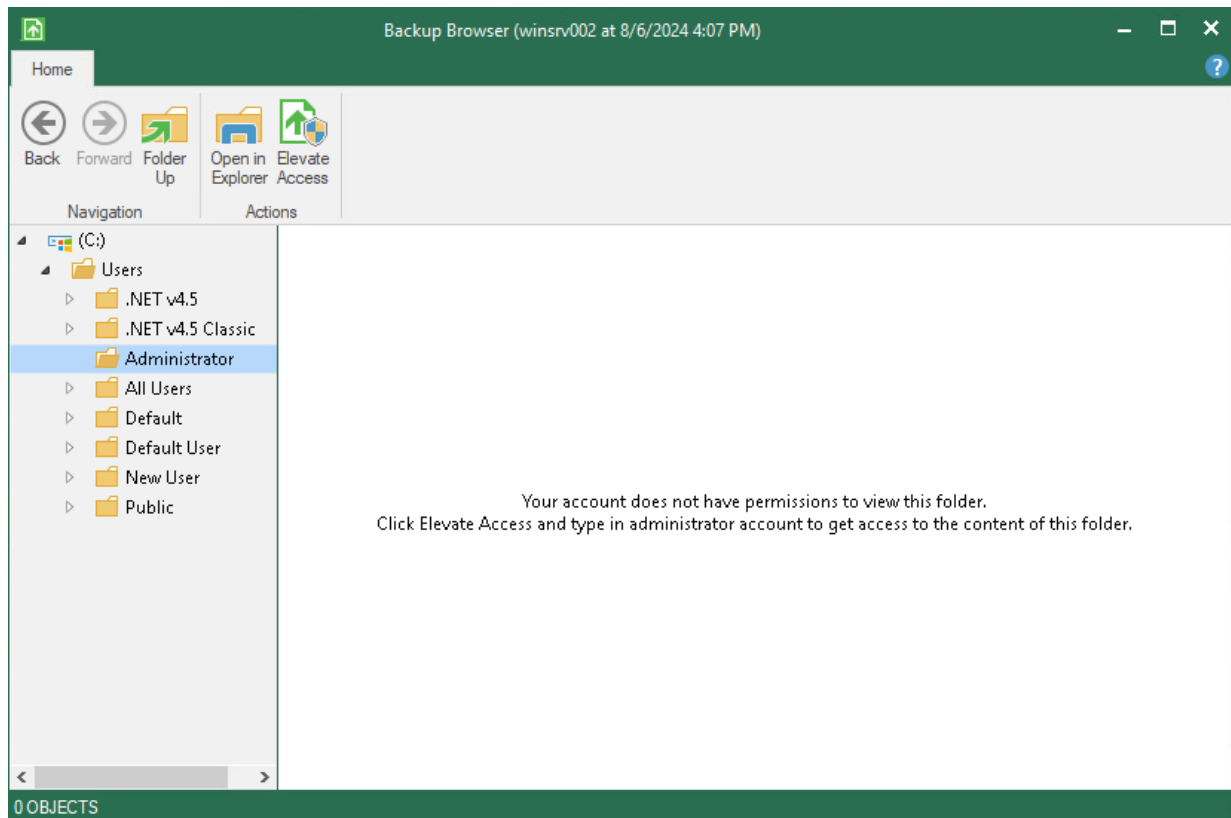
If you work under an account that does not have administrative privileges, you might not have access rights to restore the files you need. For example, this can happen in the following situations:

- You are trying to restore files to a folder you do not have access to.





- You are trying to select a file or a folder to restore from a location you do not have access to.



## NOTE

You can open the Veeam Backup browser under an account without administrative privileges only if the backup is stored in a Veeam backup repository or Veeam Cloud Connect repository.

In these cases, you can switch to the Administrator account with the Veeam Backup browser. To change the user, do the following:

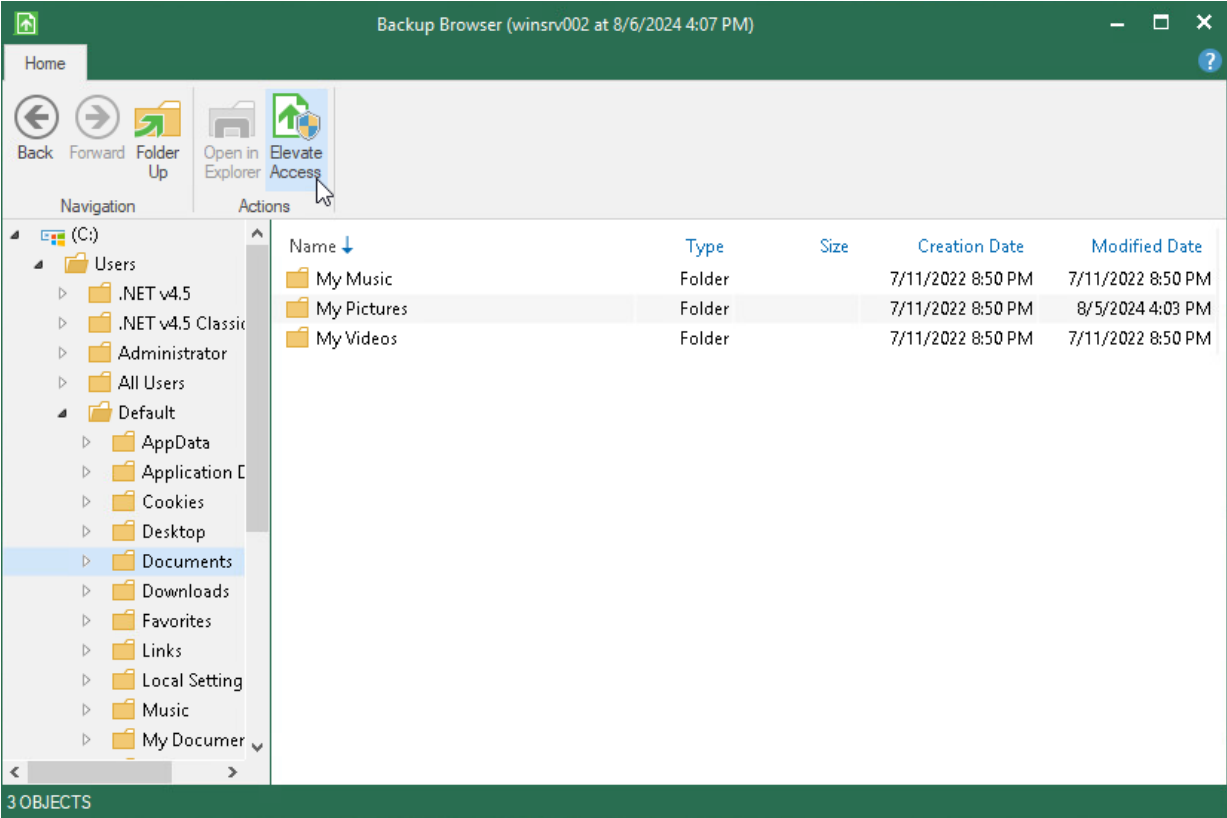
1. Open the **File Level Restore** wizard from the Veeam Agent control panel.
2. In the Veeam Backup browser, click **Elevate Access** on the toolbar.

Keep in mind that Veeam Agent displays the **Elevate Access** button only if you work under an account without administrative privileges. To allow users without administrative privileges to open the Veeam Backup browser and restore from the file-level backup, do one of the following:

- If Veeam Agent operates in the standalone mode, use a registry value. For more information, [contact Veeam Customer Support](#).
- If Veeam Agent operates in the managed mode, use the Veeam Backup & Replication console. For more information, see the [Veeam Agent for Microsoft Windows Settings](#) section in the Veeam Agent Management Guide.

3. In the dialog window, provide administrator credentials.

After that, Veeam Agent will restart the Veeam Backup browser with the updated access rights.



## Closing Veeam Backup Browser

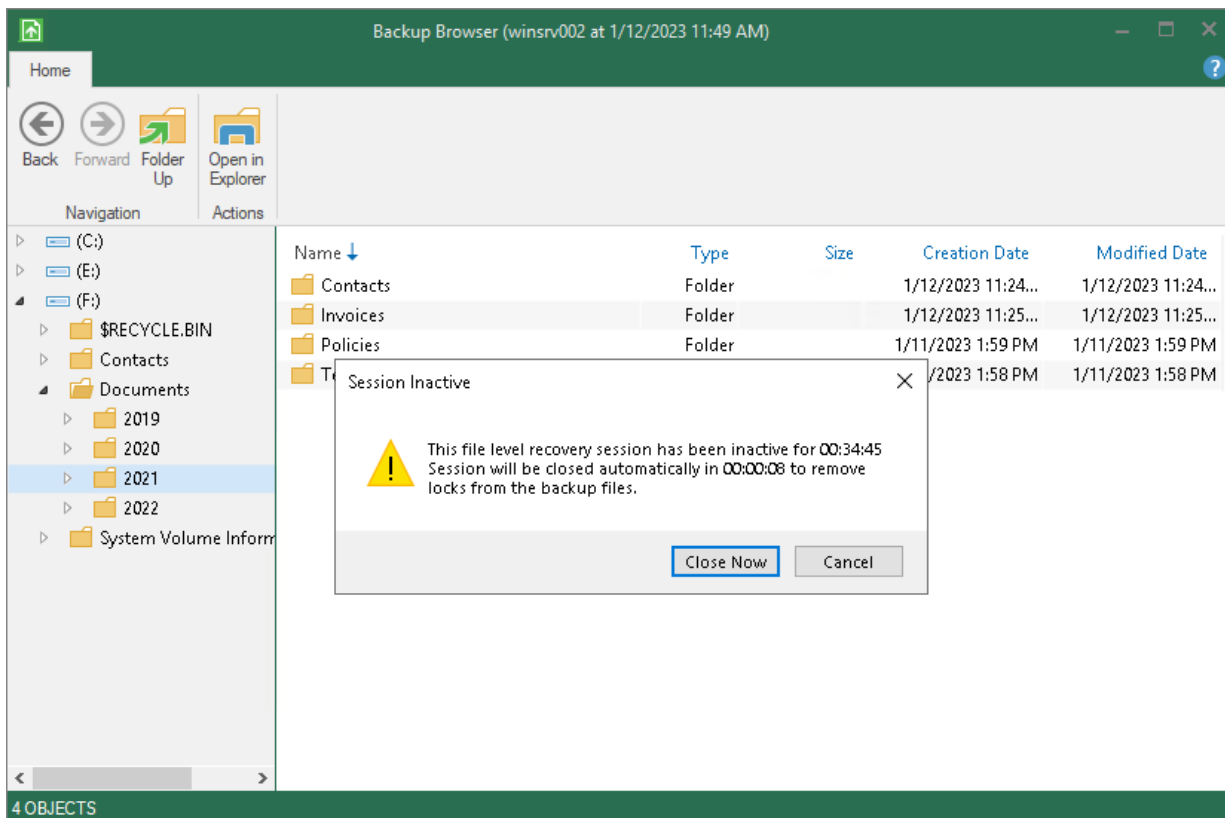
You can browse restored files and folders only while the Veeam Backup browser is open. After the Veeam Backup browser is closed, Veeam Agent for Microsoft Windows unmounts the backup content from your computer.

We recommend that you close the Veeam Backup browser after you finish restoring files and folders. Every 5 minutes, Veeam Agent for Microsoft Windows checks if there is any activity in the Veeam Backup browser. If the user or product components and services have not performed any actions for 30 minutes, Veeam Agent for Microsoft Windows displays a warning that the Veeam Backup browser is to be closed within 5 minutes.

After the warning is displayed, you can perform one of the following actions:

- You can close the Veeam Backup browser manually.
- You can click **Cancel** to postpone the close operation. In this case, the Veeam Backup browser will remain open for 30 minutes. After this period expires, Veeam Agent for Microsoft Windows will display the warning again.

- You can perform no action at all. In this case, the Veeam backup browser will be automatically closed in 5 minutes.

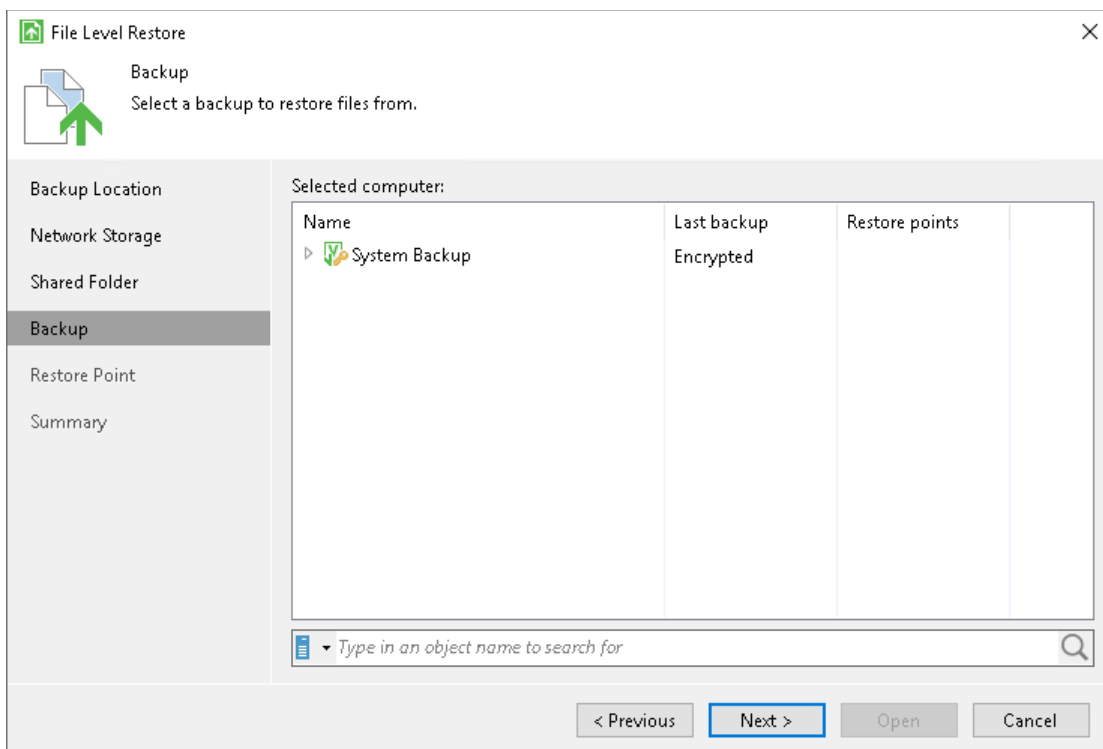


# Restoring Data from Encrypted Backups

You can perform restore from encrypted backups. In this case, backup files must be decrypted during the restore process.

To decrypt backup files and perform restore:

1. Launch the necessary data restore wizard:
  - If you want to perform file-level or volume-level restore from an encrypted backup that was created on a Veeam Agent computer, right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Restore > Individual files** or **Restore > Entire volumes**. To learn more, see [Restoring Files and Folders](#) and [Restoring Volumes](#).
  - If you want to perform bare metal recovery from an encrypted backup, boot from the Veeam Recovery Media and launch the **Veeam Recovery Media** wizard. To learn more, see [Restoring from Veeam Recovery Media](#).
2. At the **Backup Location** step of the wizard, specify where the encrypted backup file that you plan to use for restore resides. If the backup file resides in a remote location, at subsequent steps of the wizards, select the backup location type and specify settings to connect to the backup location.
3. At the **Backup** step of the wizard, select the encrypted backup.



4. In the **Specify Password** window, in the **Password** field, enter the password for the backup file.  
In the **Hint** field of the **Specify Password** window, Veeam Agent for Microsoft Windows displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.

## NOTE

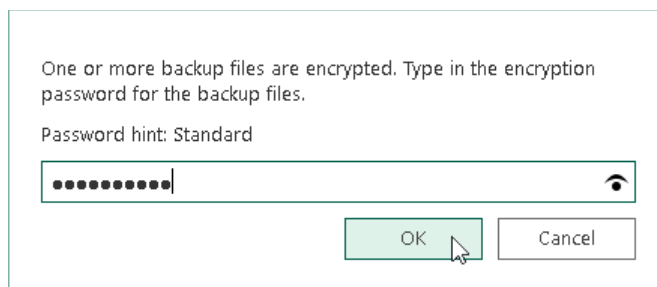
Veeam Agent does not require a password in the following cases:

- You perform restore from a backup file on the same Veeam Agent computer where the backup file was created using the same Veeam Agent for Microsoft Windows database.
- You perform restore from a backup file that resides in the Veeam backup repository. Such backups are encrypted and decrypted on the Veeam Backup & Replication side, and Veeam Agent considers such backups unencrypted. To learn more, see [Data Encryption](#).  
The restore process for backups stored in the Veeam backup repository is the same as for unencrypted backups. To learn more, see [Performing Restore](#).
- You have included encryption keys into the Veeam Recovery Media and perform bare metal recovery after booting from this Veeam Recovery Media. To learn more, see [Specify Recovery Media Options](#).

Veeam Agent requires a password if at least one restore point in the backup chain is encrypted. To learn more, see [Encryption for Existing Job](#).

If you changed the password one or several times while the backup chain was created, you need to specify the latest password. In Veeam Agent for Microsoft Windows, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

If you enter correct password, Veeam Agent for Microsoft Windows will decrypt the backup metadata. You will be able to pass to the Restore Point step of the wizard and continue the restore operation in a regular manner.



One or more backup files are encrypted. Type in the encryption password for the backup files.

Password hint: Standard

.....

OK Cancel

# Reporting

Veeam Agent for Microsoft Windows provides several ways to get information about performed backups:

- You can view information about performed backups in the Control Panel.
- You can get information about the backup state using the Veeam Agent for Microsoft Windows tray agent.
- You can get information about the backup progress using the Veeam Agent for Microsoft Windows taskbar button.
- You can get information about Veeam Agent for Microsoft Windows events using the events bar in the Control Panel.
- You can get information about Veeam Agent for Microsoft Windows events using Windows Notification Center.
- You can get information about performed backups in email reports.

# Viewing Statistics in Control Panel


You can use the Veeam Agent control panel to view statistics about performed backups. To open the Control Panel, do either of the following:

- Double-click the Veeam Agent icon in the system tray.
- Right-click the Veeam Agent icon in the system tray and select **Control Panel**.

For every configured backup job, Veeam Agent displays statistics in a separate view in the control panel. To identify the job whose statistics is currently displayed in the Control Panel, check the name of the job at the top of the Control Panel window.

If you configured multiple backup jobs in Veeam Agent, after you open the Control Panel, the Control Panel displays statistics for the first job that you configured. To switch to another job, in the main menu, hover over the name of the necessary job and select **Open**.

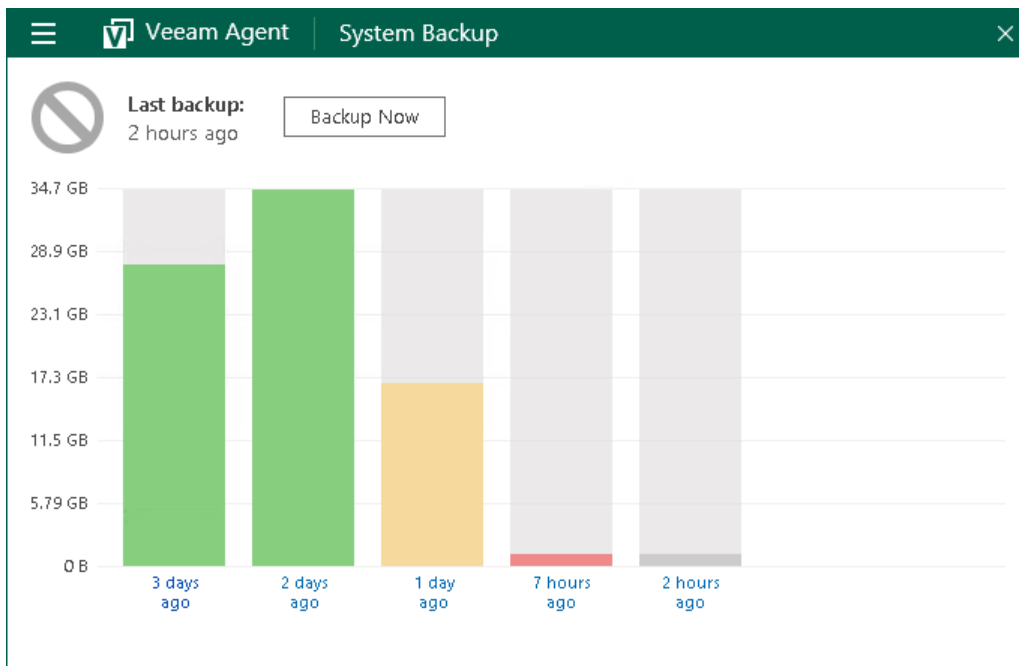
The Control Panel displays information about backup job sessions that run previously and a backup job session that is currently running. Every bar represents a separate backup job session. To view general information about a specific job session, hover the mouse over the necessary bar in the chart. Veeam Agent will provide the following details: backup status, backup time and size of the resulting backup file.



The bar color identifies the status of the backup job session. The backup job session can complete with one of the following statuses:


- *Success* (green color) – the backup job is currently running or has completed successfully.
- *Warning* (yellow color) – the backup job has completed with a warning. Veeam Agent for Microsoft Windows has managed to create the resulting backup file but you need to pay your attention to some alerts, for example: the target location is running low on disk space.
- *Error* (red color) – the backup job has completed with an error. The resulting backup file has not been created.

- *Canceled* (gray color) – the user has canceled the backup job session. The resulting backup file has not been created.




By default, Veeam Agent displays the size of created backup files. To display the duration of backup job sessions:

1. Double-click the Veeam Agent icon in the system tray or right-click the icon and select **Control Panel**.
2. In the backup statistics view, right-click the backup job sessions chart.
3. In the **Backup history** menu, select the **Duration** option.





If the backup cache is enabled for the job, Veeam Agent for Microsoft Windows also displays status of the restore point created within the job session. To learn more, see [Viewing Status of Restore Points in Backup Cache](#).



# Viewing Statistics for Separate Restore Points

You can view the following information about separate restore points in the backup chain:

- [View general statistics](#) — for any separate restore point.
- [View transaction log backup statistics](#) — for restore points created by the backup job with transaction log backup enabled.

## General Statistics

Veeam Agent for Microsoft Windows provides the following information about separate restore points in the backup chain:

- Backed-up items: items that you have chosen to back up.
- Backup duration: duration of the backup job session.
- Restore point size: size of the resulting backup file.
- Total backup size: total size of all backup files created by the backup job in the target location.

Keep in mind that if you store your backups in a scale-out backup repository with a performance tier located in object storage, the total backup size may differ from the size of all restore points in the backup chain. In this case, Veeam Agent displays only the size of the restore points stored in the performance tier.

For more information about the scale-out backup repository, see the [Scale-Out Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

- Average backup time: average time of all successful backup job sessions displayed in the chart.
- Free disk space: amount of free disk space remaining in the target location.
- Details on operations performed during the backup job session.

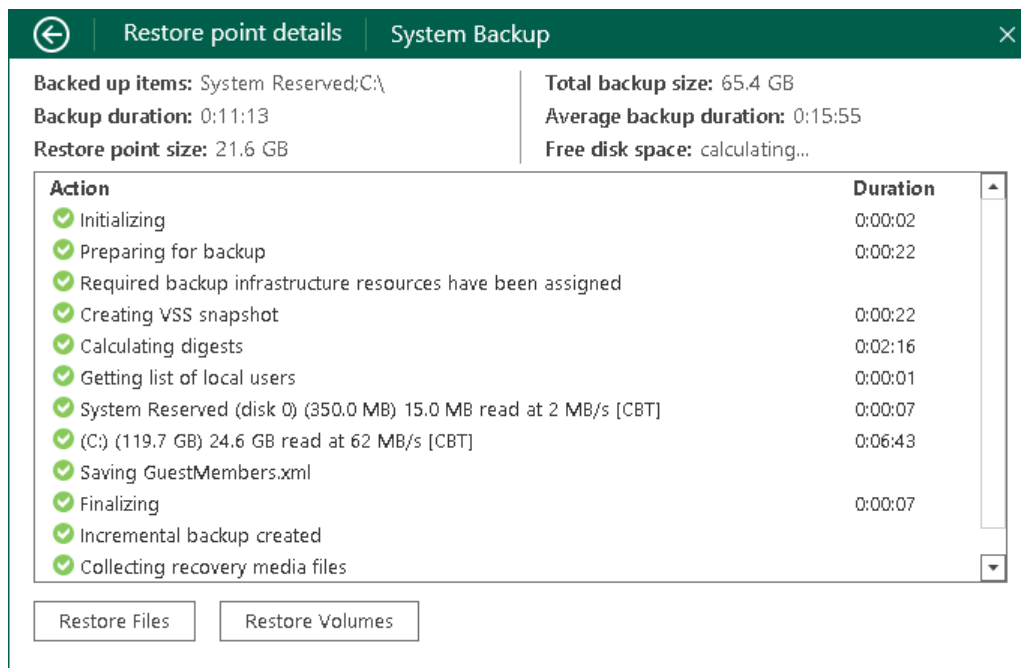
To view the restore point statistics:

1. Double-click the Veeam Agent icon in the system tray or right-click the icon and select **Control Panel**.
2. At the top of the control panel window, check the name of the backup job.

If multiple backup jobs are configured in Veeam Agent, and you want to view statistics of another job, in the main menu, hover over the name of the necessary job and select **Open**.

3. Click the necessary bar in the chart.
4. Veeam Agent will display detailed statistics on the selected backup job session. To get back to a chart view, click the arrow icon at the top left corner of the window.

If transaction log backup is enabled for the job, you can also view transaction log backup statistics. To learn more, see [Transaction Log Backup Statistics](#).



## Transaction Log Backup Statistics

If transaction log backup is enabled for the job, you can use the Veeam Agent control panel to view transaction log backup statistics.

Veeam Agent for Microsoft Windows provides the following information about transaction log processing:

- Protected databases: number of databases that were backed up at least once during the last session.
- Unprotected databases: number of databases that failed to be backed up during the last session.
- Excluded databases: databases excluded from processing. Databases may be excluded for the following reasons: database status is *Offline*, database recovery model is set to *Simple*, database is read-only, database was deleted after the latest full backup, database is added to the list of exclusions.
- Average log size: average amount of data read from the OS through all intervals.
- Max log size: maximal amount of data read from the OS over all 15-min intervals.
- Total log size: total amount of data written to the target location.
- SLA: how many log backup intervals completed in time with successful log backup (calculated as percentage of total number of intervals).
- Misses: how many intervals were missed (number of intervals).
- Max delay: difference between the configured log backup interval and time actually required for log backup. If exceeded, a warning is issued.
- Details on operations performed during the transaction log backup job session.

To view statistics on the transaction log backup processing:



1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray or right-click the icon and select **Control Panel**.

2. At the top of the control panel window, check the name of the backup job.

If multiple backup jobs are configured in Veeam Agent for Microsoft Windows, and you want to view statistics of another job, in the main menu, hover over the name of the necessary job and select **Open**.


3. Click the necessary bar in the chart.
4. In the **Restore point details** window, click the **Change to database view** link at the bottom right corner of the window. Veeam Agent for Microsoft Windows will display detailed statistics on the transaction log backup. To get back to the general statistics for the selected restore point, click the **Change to backup view** link.


To get back to a chart view, click the arrow icon at the top left corner of the window.


 Restore point details System Backup 


Protected databases: 2	Average log size: 330 KB	SLA: 100%
Unprotected databases: 0	Maximum log size: 364 KB	Misses: 0
Excluded databases: 3	Total log size: 660 KB	Max delay: 0:00:00

Action	Duration
✓ Preparing for backup	0:00:09
✓ Transaction log backup interval 1 hour 15 minutes	
✓ Backed up 296.0 KB of transaction logs for 2 databases: VEEAMSQL2019\CrmDB;VEEA...	
✓ New transaction log backup interval started at 10/4/2022 7:39:11 AM	
✓ Enumerating SQL Server databases	
✓ Performing SQL Server transaction log backup for VEEAMSQL2019\VeeamBackupRepo	
✓ Saving 364.0 KB of transaction logs to backup repository	0:00:05
✓ Transaction log backup completed at 11.4 KB/s with bottleneck: Target (Network)	
▶ Waiting for transaction log backup interval to expire	0:52:15

 Errors


 Warnings

 Success

 [Change to backup view](#)

# Viewing Information About Job Retries

If the backup job started by schedule has failed for some reason, Veeam Agent for Microsoft Windows retries the job. All backup job retries are performed within one backup job session. For this reason, Veeam Agent for Microsoft Windows displays them as one bar in the chart.



## NOTE

For portable devices, Veeam Agent for Microsoft Windows does not automatically retry the backup job if a device is working on battery and the battery level is below 20%.

To view detailed information about the backup job retries:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. At the top of the control panel window, check the name of the backup job.  
If multiple backup jobs are configured in Veeam Agent for Microsoft Windows, and you want to view statistics of another job, in the main menu, hover over the name of the necessary job and select **Open**.
3. Click the necessary bar in the chart.
4. At the bottom right corner of the window, click the **Show retries** link.

5. After you view details, you can hide them. To do this, at the bottom right corner of the window, click the **Hide retries** link.

Restore point details

Weekly Backup

Backed up items: n/a

Backup duration: 0:00:06

Restore point size: n/a

Total backup size: 0 B

Average backup duration: 0:01:51

Free disk space: n/a

Action	Duration
Retry: 3	
Initializing	
Error: Failed to connect to backup repository Default Backup Repository. It may be offl	
Processing finished with errors at 1/12/2023 4:56:41 PM	
Retry: 4	
Initializing	
Error: Failed to connect to backup repository Default Backup Repository. It may be offl	
Processing finished with errors at 1/12/2023 5:06:55 PM	
Retry: 5	
Initializing	
Error: Failed to connect to backup repository Default Backup Repository. It may be offl	
Processing finished with errors at 1/12/2023 5:17:08 PM	

Restore Files




Restore Volumes

Hide retries

# Viewing Status of Restore Points in Backup Cache

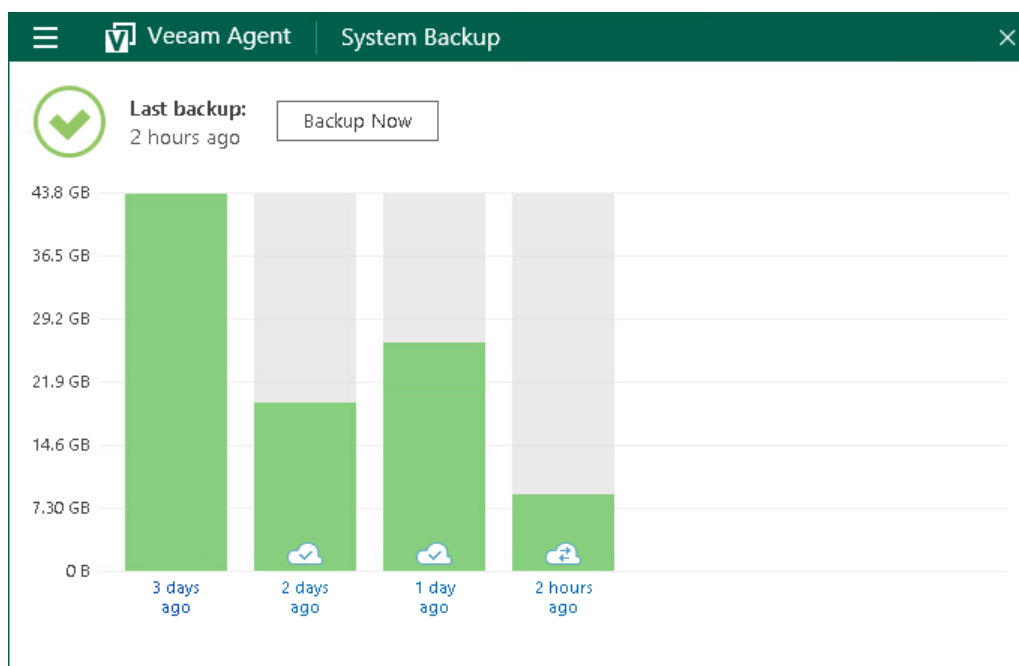
If the backup cache is enabled for the job, you can monitor status of restore points in the backup cache. Veeam Agent for Microsoft Windows displays the restore point status through the icon on a bar in the chart.

The icon can be in one of the following states:

Icon	Description	Backup state
	Check mark over the icon.	The backup file created within the backup session is saved on the target storage.
—	No icon.	The backup file created within the backup session is saved in the backup cache.
	Sync sign over the icon.	The backup file is being uploaded from the backup cache to the target storage.
	Error sign over the icon.	The backup file was not uploaded to the target storage or has been deleted from the backup cache.

## TIP










You can also monitor the backup cache activity and view detailed statistics on the restore point upload process. To learn more, see [Monitoring Backup Cache Activity](#).




# Monitoring Backup State with Tray Agent

The Veeam Agent for Microsoft Windows icon displayed in the system tray lets you monitor the state of your backups and get informed about the computer protection status.

The icon can be in one of the following states:


Icon	Description	Backup state
	Question mark over the icon	The backup job is not configured.
	Veeam Agent for Microsoft Windows icon	The backup job is set up but scheduling settings for the job are not configured.
	Animated icon	The backup task is being performed. To view the backup task progress, hover the mouse over the icon.
	Clock over the icon	<ul style="list-style-type: none"><li>• The backup job has been scheduled; waiting for the first backup job session.</li><li>• The latest session of the scheduled backup job has completed successfully; waiting for the next backup job session.</li></ul>
	Sync sign over the icon	Veeam Agent is uploading backup files from the backup cache to the target storage.
	Cancel sign over the icon	The latest session of the scheduled backup job has been canceled.
	Error sign over the icon	An error occurred during the latest backup job session, and the session was terminated.
	Minus sign over the icon	The scheduled backup job is disabled.
	Grey icon	The tray agent is not connected to the Veeam Agent for Microsoft Windows service.



Icon	Description	Backup state
	Warning sign over the icon	<ul style="list-style-type: none"> <li>The backup job has completed with a warning, for example, the target location is running low on space.</li> <li>[If you have selected a removable storage device as a target destination in the backup job settings] The target removable storage device is not connected to the computer. In this case, Veeam Agent also displays a warning on the notifications bar in the control panel. You can attach the target removable storage device to the computer within 10 minutes, and Veeam Agent will automatically start the scheduled backup job.</li> </ul>

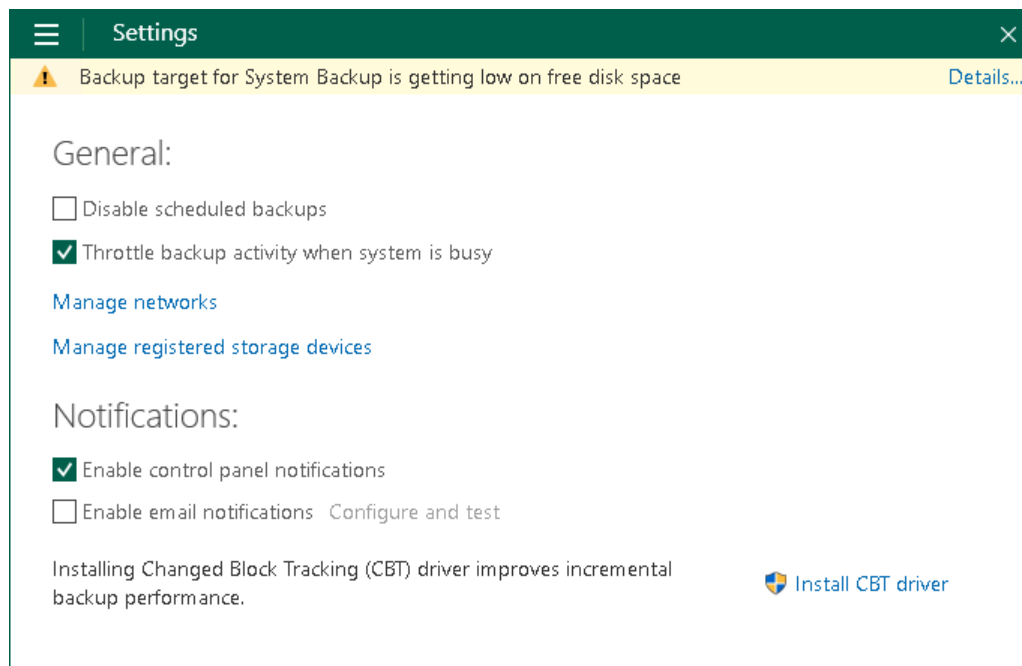
# Monitoring Backup Process in Taskbar Button

You can monitor the backup process with the Veeam Agent taskbar button. Veeam Agent displays on the taskbar button a green progress bar that reflects the bar for the currently running job session in the control panel. As a result, you can track the process of the backup file creation while working with another application without having to switch to the control panel.



# Viewing and Dismissing Veeam Agent Events

If a warning event occurs, Veeam Agent displays a notification bar with the event description in the control panel window. You can get detailed information about events and dismiss events not to get alerted of them in future. For the full list of notifications, see [Control Panel Notifications](#).



Veeam Agent displays only the latest event in the notification bar. To view detailed information about all events:

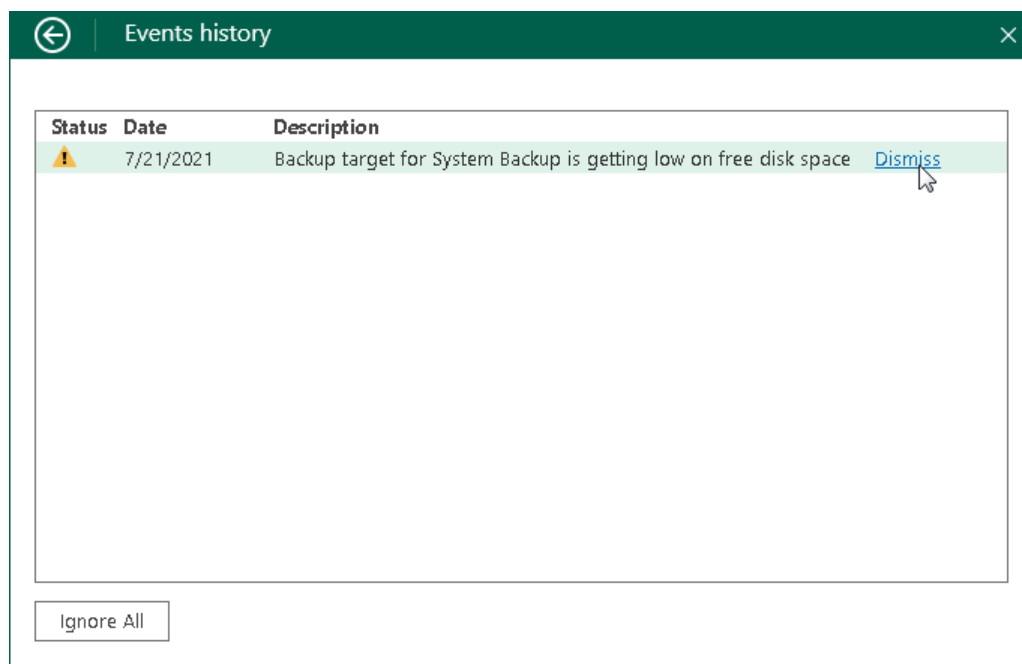
1. Double-click the Veeam Agent icon in the system tray, or right-click the Veeam Agent icon in the system tray and select **Control Panel**.
2. Click **Details** on the notification bar at the top of the control panel window.

To dismiss events:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. Click **Details** on the notification bar at the top of the control panel window.
3. Click **Dismiss** next to the necessary event. To dismiss all events at once, click **Ignore All** at the bottom left corner of the window.

## TIP

You can disable notifications at all. To learn more, see [Disabling Control Panel Notifications](#).



## Control Panel Notifications


Veeam Agent for Microsoft Windows can inform you about the following events using the notification bar:

Notification	Description
New software update is available	A newer version of Veeam Agent for Microsoft Windows is available for download.
Veeam Recovery Media must be recreated due to a recent Veeam Agent for Microsoft Windows upgrade	You must recreate the Veeam Recovery Media after you have upgraded Veeam Agent for Microsoft Windows to the newer version.
Recovery media OS version is different from the current OS version	You must recreate the Veeam Recovery Media after you have updated the Microsoft Windows OS.
Recovery media has not been created	You have not created the Veeam Recovery Media yet. If the OS on the computer fails to boot, you will not be able to boot the OS using the Veeam Recovery Media to diagnose and fix problems, or restore data from the backup.

Notification	Description
Backup target is not available, backup is postponed	Veeam Agent for Microsoft Windows is unable to detect the backup target at the time when the scheduled backup job must start (for example, the backup target is a removable storage device that is not currently connected to the Veeam Agent computer). Veeam Agent for Microsoft Windows displays this notification during the job retry sessions.
Backup target is not available	Veeam Agent for Microsoft Windows is unable to detect the backup target by the moment when the scheduled backup job must start (for example, backup target is a removable storage device that is not currently connected to the Veeam Agent computer). Veeam Agent for Microsoft Windows displays this notification after the job has failed.
Backup target has not been seen for <N> days	The backup target has not been seen for <N> days. This notification is displayed if scheduled backups have not been created for 2 days or more.
Backup target is getting low on free disk space	Free space on the target storage disk is below 10%. If Veeam Agent for Microsoft Windows does not have enough disk space for backup operation, the job session will fail.
Skipping scheduled backup, because battery level is too low	[For laptops and tablets] The battery level is below 20%. Veeam Agent for Microsoft Windows does not start a new backup session in this case.
This application is managed by your system administrator	This notification is displayed if Veeam Agent for Microsoft Windows is operating in the read-only mode. Veeam Agent for Microsoft Windows operates in the read-only mode if backup operations are managed by Veeam Backup & Replication or Veeam Service Provider Console.
Your license will expire in <N> days	The Veeam Agent for Microsoft Windows license will expire, and the grace period will start in <N> days.
Your license has expired and needs to be renewed	The Veeam Agent for Microsoft Windows license has expired, and the grace period has started. During the grace period, you can perform all types of backup and restore operations. You must obtain a new license before the end of the grace period.
Your license has expired, and your grace period is over	The Veeam Agent for Microsoft Windows license has expired, and the grace period is over. In this case, Veeam Agent for Microsoft Windows does not perform backup operations. However, you are able to restore data from existing backups. To use Veeam Agent for Microsoft Windows in the full functionality mode, you must obtain a new license.

# Viewing Events with Windows Notification Center

In addition to displaying notifications about Veeam Agent events in the control panel window, Veeam Agent can inform you about its events using the Windows Notification Center. Once an event occurs, a message notifying about this event will appear on the Microsoft Windows desktop. This lets you view the notification without the need to open the Veeam Agent Control Panel and take the necessary actions immediately. For the full list of notifications that appears on the Microsoft Windows desktop, see [Windows Desktop Notifications](#).



The notification closes automatically in a few seconds. To close the notification manually, click the message.

To disable desktop notifications, right-click the message and select **Turn off notifications for Veeam Agent for Microsoft Windows**.

## Windows Desktop Notifications

Veeam Agent for Microsoft Windows can inform you about the following events using the Windows Notification Center:

Notification	Description
New software update is available	A newer version of Veeam Agent for Microsoft Windows is available for download.
Veeam Recovery Media must be recreated due to a recent Veeam Agent for Microsoft Windows upgrade	You must recreate the Veeam Recovery Media after you have upgraded Veeam Agent for Microsoft Windows to the newer version.
Recovery media OS version is different from the current OS version	You must recreate the Veeam Recovery Media after you have updated the Microsoft Windows OS.
Recovery media has not been created	You have not created the Veeam Recovery Media yet. If the OS on the computer fails to boot, you will not be able to boot the OS using the Veeam Recovery Media to diagnose and fix problems, or restore data from the backup.
Backup target is not available	Veeam Agent for Microsoft Windows is unable to detect the backup target by the moment when the scheduled backup job must start (for example, backup target is a removable storage device that is not currently connected to the Veeam Agent computer). Veeam Agent for Microsoft Windows displays this notification after the job has failed.

Notification	Description
Your license will expire in <N> days	The Veeam Agent for Microsoft Windows license will expire, and the grace period will start in <N> days.
Your license has expired and needs to be renewed	The Veeam Agent for Microsoft Windows license has expired, and the grace period has started. During the grace period, you can perform all types of backup and restore operations. You must obtain a new license before the end of the grace period.
Your license has expired, and your grace period is over	The Veeam Agent for Microsoft Windows license has expired, and the grace period is over. In this case, Veeam Agent for Microsoft Windows does not perform backup operations. However, you are able to restore data from existing backups. To use Veeam Agent for Microsoft Windows in the full functionality mode, you must obtain a new license.

# Viewing Job Session Results in Email Reports

You can receive email notifications with Veeam Agent for Microsoft Windows job results. When the backup job session completes, Veeam Agent for Microsoft Windows will send a report containing data on the job session to the specified email address.

The report contains the following data:

- Cumulative job session statistics: session duration details, details of the session performance, amount of read, processed and transferred data, backup size, compression and deduplication ratios.
- Detailed statistics for the computer processed within the session: processing duration details, backup data size, amount of read and transferred data, list of warnings and errors (if any).
- If the backup job is set up to create database log backups, the report contains statistics for the database log backup job: a list of databases that were backed up at least once during the last session and information for the latest log processing intervals.
- If you use the backup cache, the report also contains statistics on the backup cache activity: a list of restore points created in the backup cache, their status and upload details. To learn more about the backup cache, see [Backup Cache](#).

To receive email reports, you must enable and configure email notifications in the Veeam Agent for Microsoft Windows Control Panel. To learn more, see [Enabling Email Notifications](#). Once email notifications are configured, Veeam Agent for Microsoft Windows will send email report for every backup job session that is started by schedule, manually or when you perform standalone full or incremental ad-hoc backup.

If the scheduled backup job fails, Veeam Agent for Microsoft Windows does not send a report after every job retry. Instead, Veeam Agent for Microsoft Windows sends one report on the first error within the job session and another report on the last job session result — success or error.

Agent Backup job: System Backup (Full)							
Veeam Agent for Microsoft Windows							
Success							
Friday, 28 October 2022 10:00:00							
Success	1	Start time	10:00:00	Total size	120.0 GB	Backup size	15.6 GB
Warning	0	End time	10:14:07	Data read	22.2 GB	Dedupe	5.4x
Error	0	Duration	0:14:07	Transferred	15.5 GB	Compression	1.4x
Details							
Name	Status	Start time	End time	Size	Read	Transferred	Duration
FILESRV01	Success	10:00:00	10:14:07	120.0 GB	22.2 GB	15.5 GB	0:14:06



# Specifying Settings

You can use global settings of Veeam Agent to accomplish the following tasks:

- [Throttle backup activities.](#)
- [Restrict network connections usage.](#)
- [Manage backup storage devices.](#)
- [Disable Control Panel notifications.](#)
- [Enable email notifications.](#)
- [Check for new product versions and updates.](#)

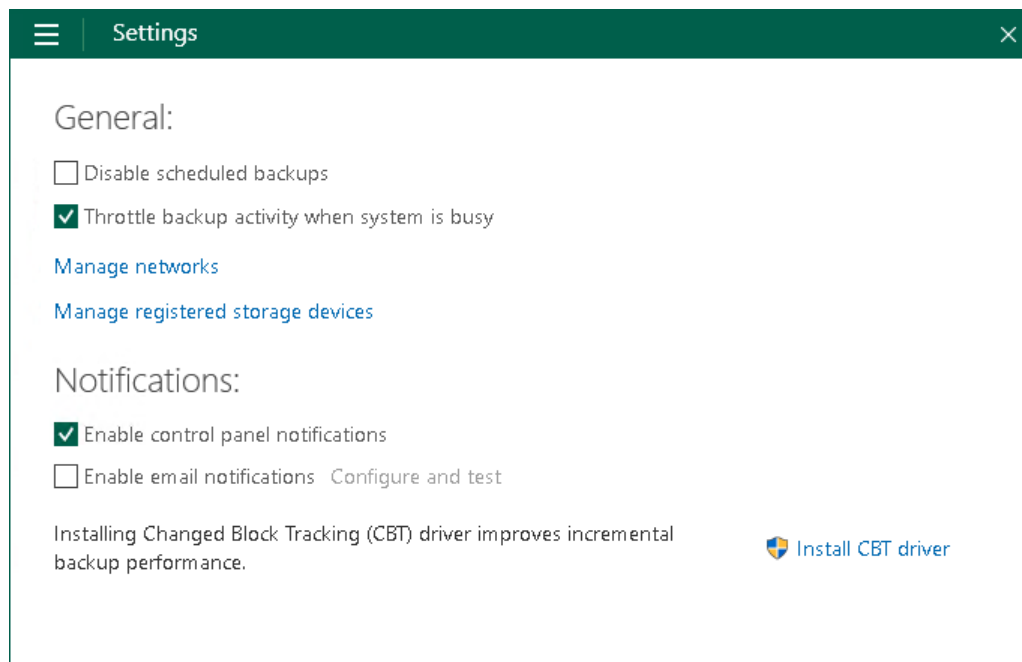
# Throttling Backup Activities

You can instruct Veeam Agent for Microsoft Windows to throttle its activities during backup. The throttling option can help you avoid situations when backup tasks consume all available hard disk resources and hinder work of other applications and services.

With throttling enabled, Veeam Agent for Microsoft Windows sets low priority for Veeam Agent for Microsoft Windows components engaged in the backup process (in particular, the *VeeamAgent.exe* process). If this option is not enabled, Veeam Agent for Microsoft Windows components have normal priority.

To enable the throttling option for backup activities:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Select the **Throttle backup activity when system is busy** check box.



# Restricting Network Connections Usage

You can instruct Veeam Agent for Microsoft Windows to throttle network traffic for backup jobs and disable backup over Internet connections of certain types. This helps limit the impact of Veeam Agent operations on network performance and avoid extra costs.

Network usage settings apply to all types of backups: scheduled and ad-hoc. To restrict network connections usage for Veeam Agent, you can perform the following tasks:

- [Limit bandwidth consumption for Veeam Agent.](#)
- [Disable backup over metered connections.](#)
- [Disable backup over VPN connections.](#)
- [Allow Veeam Agent to use specific Wi-Fi networks only.](#)

# Limiting Bandwidth Consumption

To reduce the impact of Veeam Agent operations on network performance, you can limit bandwidth consumption for Veeam Agent backup jobs. Limiting bandwidth consumption prevents jobs from utilizing the entire bandwidth available in your environment to ensure that enough traffic is provided for other network operations. We strongly recommend to limit bandwidth consumption if you perform backup to a remote location over a slow network connection.

## NOTE

Bandwidth consumption limit does not apply to restore operations.

By default, Veeam Agent is set up to use the entire bandwidth available in your environment. To limit bandwidth consumption for Veeam Agent backup jobs:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Click the **Manage networks** link.
4. In the **Throttling** section, select the **Limit bandwidth consumption to** check box and specify the maximum speed for transferring backed-up data from the Veeam Agent computer to the target location.

Network

Throttling:

☒ Limit bandwidth consumption to: 1 MB/s

Networks:

☒ Restrict metered connections usage

☐ Restrict VPN connections usage

☐ Restrict Wi-Fi usage to these networks only:

SSID
------

Add...

Remove

# Disabling Backup over Metered Connections

Veeam Agent can disable backup over metered Internet connections to help you avoid extra costs. If you use a metered Internet connection, your service provider charges you by the amount of data sent and received by your computer. Veeam Agent automatically detects metered connections and will not perform backup when your computer is on such connection.

Consider the following limitations and requirements:

- Veeam Agent disables backup over metered Internet connections only on computers that run Microsoft Windows 8.1 and later. If the computer runs an earlier version of Microsoft Windows, this option is not applicable.
- You must specify which connections are metered in Microsoft Windows. To learn more, see [this Microsoft KB article](#).

By default, backup over metered connections is disabled. To enable backup over metered connections:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Click the **Manage networks** link.
4. In the **Networks** section, clear the **Restrict metered connections usage** check box.

## NOTE

For the cases when backup over metered connections is disabled, consider the following:

- If you start the backup job manually when only a metered connection is available, Veeam Agent for Microsoft Windows will display a warning and ask you to confirm that you want to use this connection for backup.
- If you start the backup job with the command line interface, Veeam Agent will ignore the setting and will use metered connections for backup.
- If you enable the backup cache for the backup job, and this backup job starts upon schedule when only a metered connection is available, Veeam Agent for Microsoft Windows will save the backup file to the backup cache. If the backup cache is not enabled, the backup job will fail.

The screenshot shows the 'Network' settings window in Veeam Agent. The window has a dark green header with a back arrow icon and a close 'X' icon. The title 'Network' is centered in the header. Below the header, the 'Throttling:' section contains a checked checkbox 'Limit bandwidth consumption to:' followed by a text input field containing '1' and a dropdown menu showing 'MB/s'. The 'Networks:' section has three unchecked checkboxes: 'Restrict metered connections usage', 'Restrict VPN connections usage', and 'Restrict Wi-Fi usage to these networks only:'. Below these checkboxes is a large, empty rectangular area with a light gray background, intended for listing SSIDs. To the right of this area are two buttons: 'Add...' and 'Remove'. A small upward-pointing arrow is visible in the top right corner of the SSID list area.

# Disabling Backup over VPN Connections

Veeam Agent can disable backup over VPN connections. This helps you avoid extra costs if your VPN service provider charges you by the amount of data sent and received by your computer. Veeam Agent automatically detects VPN connections and will not perform backup when your computer is on such connection.

By default, backup over VPN connections is enabled. To disable backup over VPN connections:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Click the **Manage networks** link.
4. In the **Networks** section, select the **Restrict VPN connections usage** check box.

Keep in mind that Veeam Agent detects VPN connections created with Microsoft Windows in-built tools and the following 3rd party VPN providers:

- Cisco AnyConnect
- Fortinet SSL VPN
- Open VPN
- Palo Alto Networks GlobalProtect VPN

## NOTE

For the cases when backup over VPN connections is disabled, consider the following:

- If you start the backup job manually when only a VPN connection is available, Veeam Agent will display a warning and ask you to confirm that you want to use this connection for backup.
- If you start the backup job with the command line interface, Veeam Agent will ignore the setting and will use VPN connections for backup.
- If you enable the backup cache for the backup job, and this backup job starts upon schedule when only a VPN connection is available, Veeam Agent will save the backup file to the backup cache. If the backup cache is not enabled, the backup job will fail.

The screenshot shows the 'Network' settings window in Veeam Agent. The window has a dark green header with a back arrow icon and a close 'X' icon. The title 'Network' is centered in the header. Below the header, the 'Throttling:' section contains a checked checkbox 'Limit bandwidth consumption to:', followed by a text input field containing '1' and a dropdown menu showing 'MB/s'. The 'Networks:' section contains three checkboxes: 'Restrict metered connections usage' (unchecked), 'Restrict VPN connections usage' (checked), and 'Restrict Wi-Fi usage to these networks only:' (unchecked). Below these checkboxes is a large, empty rectangular area labeled 'SSID' in the top-left corner, with an upward-pointing arrow icon in the top-right corner. To the right of this area are two buttons: 'Add...' and 'Remove'.



# Selecting Wireless Networks for Backup

You can restrict usage of Wi-Fi networks for Veeam Agent for Microsoft Windows. This may be useful, for example, to avoid transmission of backed-up data over insecure or slow wireless network connections.

By default, Veeam Agent for Microsoft Windows is set up to back up data over any Wi-Fi network to which the Veeam Agent computer is connected. You can select one or more specific Wi-Fi networks and instruct Veeam Agent for Microsoft Windows to perform backup over these networks only. Backup over other wireless networks will be disabled.

To restrict usage of wireless networks:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Click the **Manage networks** link.
4. Select the **Restrict Wi-Fi usage to these networks only** check box and click **Add**.
5. Select one or more Wi-Fi networks in the list and click **Save**.

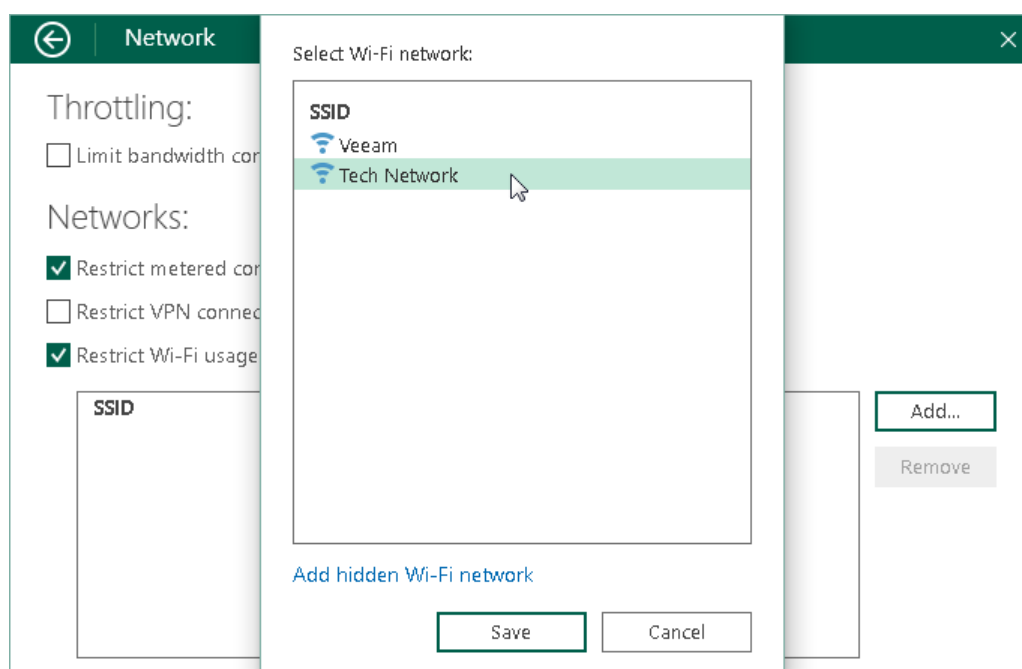
If you want to allow Veeam Agent to back up data over a hidden wireless network, click the **Add hidden Wi-Fi network** link. In the displayed window, specify the SSID of the necessary network and click **Save**.

If you want to restrict usage over any wireless networks, select the **Restrict Wi-Fi usage to these networks only** check box and do not add any networks to the list.

## NOTE

For the cases when backup over Wi-Fi networks is restricted, consider the following:

- If you start the backup job manually when selected Wi-Fi networks are not available, Veeam Agent for Microsoft Windows will display a warning and ask you to confirm that you want to use this connection for backup.
- If you start the backup job with the command line interface, Veeam Agent will ignore the setting and will use all available Wi-Fi networks for backup.
- If you enable the backup cache for the backup job, and this backup job starts upon schedule when selected Wi-Fi networks are not available, Veeam Agent for Microsoft Windows will save the backup file to the backup cache. If the backup cache is not enabled, the backup job will fail.



# Managing Rotated Drives

You can use a rotated drives scheme for storing backups. To do this, you can create backups on several external drives (for example, USB or FireWire) and swap these drives when needed.

The drive on which you plan to store a backup must be registered in Veeam Agent for Microsoft Windows. If the drive is not registered, Veeam Agent for Microsoft Windows will not be able to detect the drive and store a backup on it.

## TIP

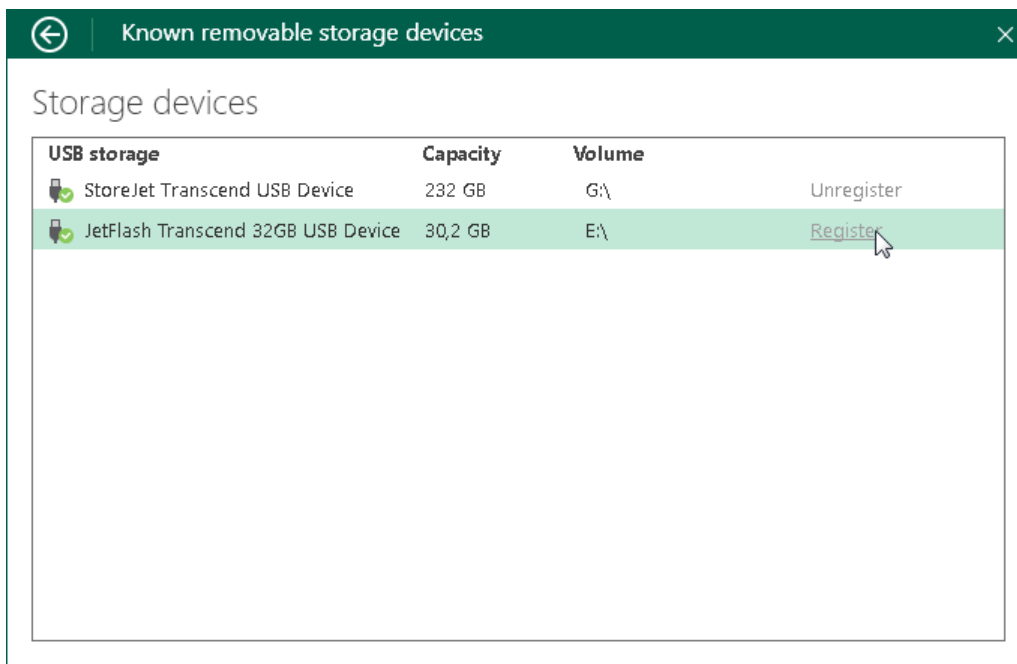
When you target a backup job at an external drive for the first time, this drive is registered automatically.

Consider the following limitations:

- You can register and unregister drives if you have selected to store backups on an external drive connected to the computer. If you have selected to store backups on a local computer drive, in a network shared folder or in a backup repository, registering options will be disabled.
- You cannot unregister all drives at once. At least one drive will remain registered in Veeam Agent for Microsoft Windows.

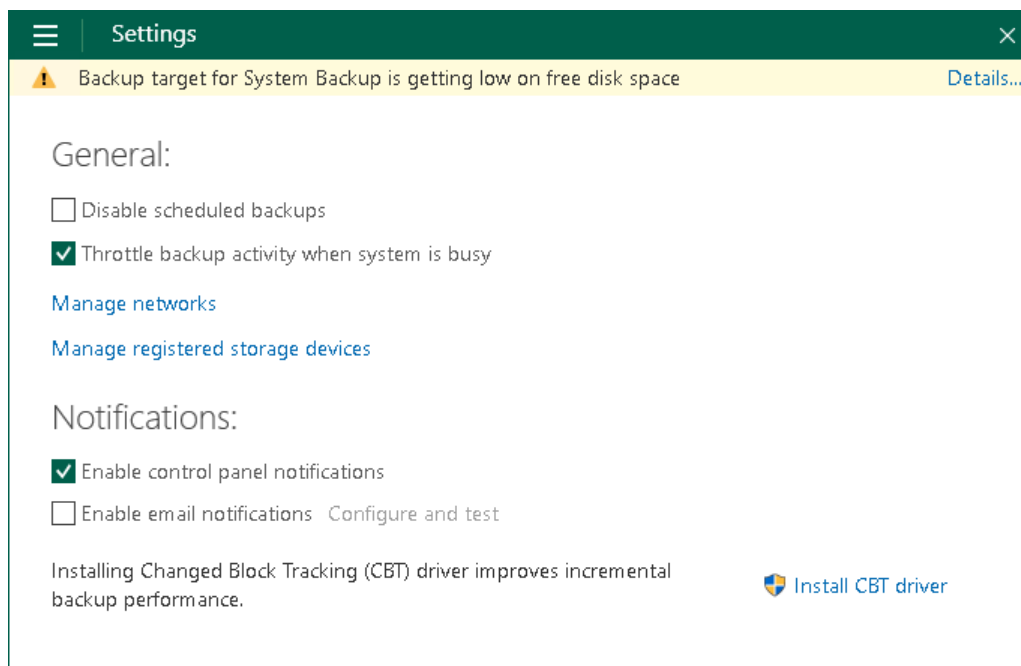
To register and unregister a drive in Veeam Agent for Microsoft Windows:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Click the **Manage registered storage devices** link.
4. In the list of devices, click **Register/Unregister** next to the necessary backup storage device.



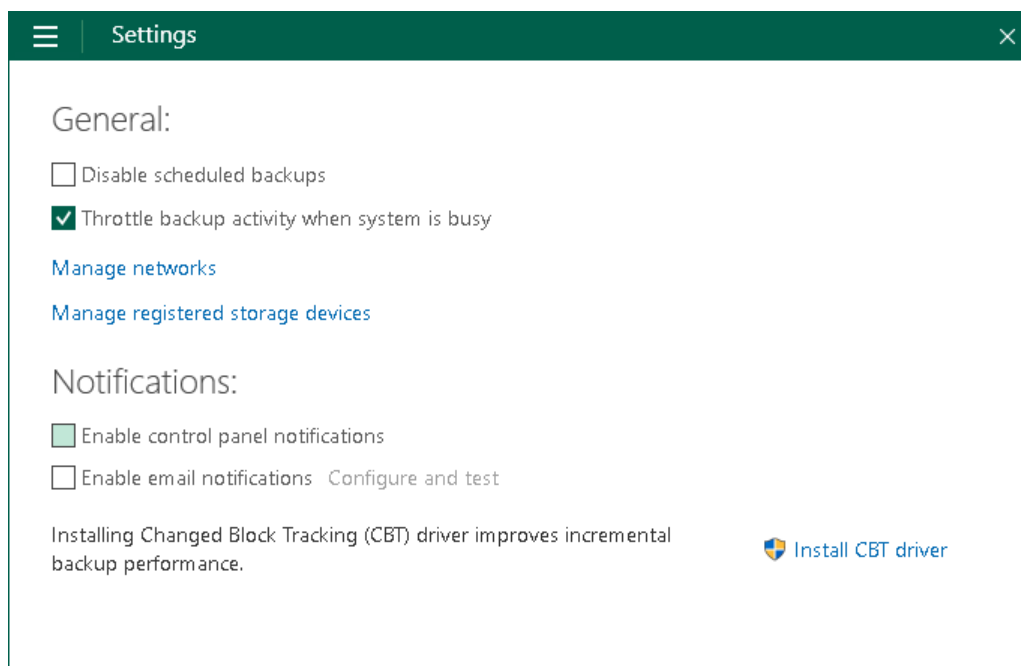
# Disabling Control Panel Notifications

Veeam Agent for Microsoft Windows displays warning and information messages on the notification bar in the Control Panel. If necessary, you can disable Veeam Agent notifications.



To disable notifications:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. In the **Notifications** section, clear the **Enable Control Panel notifications** check box.

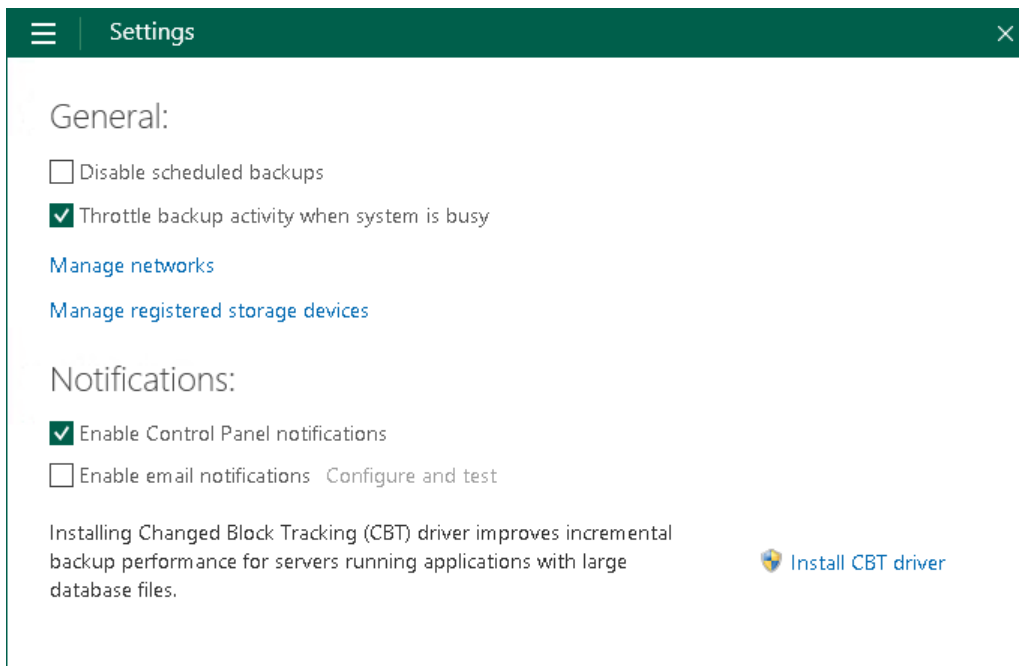


# Enabling Email Notifications

You can enable Veeam Agent for Microsoft Windows email notifications to receive reports containing data on the latest backup job session statistics and result.

To enable email notifications:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. In the **Notifications** section, select the **Enable email notifications** check box and click the **Configure and test** link.
4. In the **Configure and test email notifications** window, select one of the **Mail Server** options and specify the required settings:
  - o [Custom SMTP server settings](#).
  - o [Gmail server settings](#).
  - o [Microsoft 365 server settings](#).



## Disabling Email Notifications

To disable email notifications, clear the **Enable email notifications** check box in the **Settings** tab of the Control Panel. Current email notifications configuration will remain saved in the Veeam Agent database.

# Custom SMTP Server Settings

To connect to the SMTP server, specify the following settings in the **Configure and test email notifications** window:



1. In the **Mail server** section, specify the following:
  - a. Select the **Custom SMTP Server (Basic authentication)** option from the drop-down list.
  - b. In the **SMTP Server DNS Name or IP address** field, enter a full DNS name or IP address of the SMTP server that will be used for sending email notifications.
  - c. In the **Port** field, specify the port number for the SMTP server.

**NOTE**

Sending email notifications using Implicit TLS (over port 465) is not supported. For more information about Implicit TLS, see [this RFC section](#).

  - d. In the **Username** field, specify a user name for the account that has rights to access the SMTP server.
  - e. In the **Password** field, enter a password for the account that has rights to access the SMTP server.
  - f. To use a secure SSL/TLS connection for email operations, select the **Use secure connection (SSL/TLS)** check box.
2. In the **Email settings** section specify the following:
  - a. Specify the sender email address in the **From email address** field. You can specify your account email address or its alias.
  - b. Specify the recipient email address in the **To email address** field. You can specify several recipient email addresses separated with a comma or semicolon.
  - c. In the **Email subject** field, specify a subject for the message. You can use the following variables in the subject:
    - i. *%JobResult%*
    - ii. *%ComputerName%*
    - iii. *%JobName%*
    - iv. *%CompletionTime%*
3. In the **Notify on** section, select the **Success**, **Warning** or **Error** check boxes to receive email notification if a job is run successfully, not successfully or with a warning.

4. [Optional] Click **Test now** to validate the SMTP server settings and send a test email.

 **Configure and test email notifications** 

Mail server:

Custom SMTP server (Basic authentication) ▼

smtp.tech.local

587

notifications@tech.local


☒ Use secure connection (SSL/TLS)

●●●●●●●●●●●●●●●●

Email settings:

notifications@tech.local

administrators@tech.local

[%JobResult%] %ComputerName% - %JobName% - %CompletionTime% 

Notify on: ☒ Success ☒ Warning ☒ Error

Test Now

Click to test specified server settings

# Gmail Server Settings

To connect to the Gmail server, specify the following settings in the **Configure and test email notifications** window:

1. In the **Mail server** section, specify the following:
  - a. Select the **Google Gmail (Modern authentication)** option from the drop-down list.
  - b. Click the **Sign in with Google** button, enter your Google account credentials and complete Google authorization process.

## IMPORTANT

When you are prompted to allow Veeam Agent sending emails on behalf of your Google account, grant the permission and click **Continue**.

2. In the **Email settings** section, specify the following:
  - a. If necessary, change the email address in the **From email address** field to its alias.
  - b. If necessary, change the recipient email address in the **To email address** field. You can specify several recipient email addresses separated with a comma or semicolon.
  - c. In the **Email subject** field, specify a subject for the message. You can use the following variables in the subject:
    - i. *%JobResult%*
    - ii. *%ComputerName%*
    - iii. *%JobName%*
    - iv. *%CompletionTime%*
3. In the **Notify on** section, select the **Success**, **Warning** or **Error** check boxes to receive email notification if a job is run successfully, not successfully or with a warning.
4. [Optional] Click **Test now** to validate the Gmail server settings and send a test email.

To specify custom authentication options, do the following:

1. Click the **Server settings** button.
2. Select the **Use custom settings** check box.



3. Specify the application client ID and the client secret.

To learn how to register your custom application, see [Registering Application in Google Cloud Console](#).

Configure and test email notifications

Mail server:

Google Gmail (Modern authentication) Server settings

Application is authorized. Sign in with Google

Email settings:

notifications@tech.local

administrators@tech.local

[%JobResult%] %ComputerName% - %JobName% - %CompletionTime% i

Notify on: ☒ Success ☒ Warning ☒ Error

Test Now [Click to test specified server settings](#)

## Registering Application in Google Cloud Console

If you want to use your own web application for email notifications, you need to configure it in the Google Cloud console. To do this, perform the following steps:

1. Log in to the [Google Cloud console](#) under a Google account that has permissions to create applications.
2. Create a new project and enable Gmail API for the project. To do this, open **APIs and services > Library > Gmail API > Manage** and click **Enable API**.
3. Create OAuth credentials. To do this, perform the following steps:
  - a. Open **APIs and services > Credentials**. Click **Create credentials** and select **OAuth client ID**.
  - b. In the **Application type** field, select **Desktop app**.
  - c. In the **Name** field, specify the name of your OAuth 2.0 client.
  - d. Click **Create** to generate the application client ID and the client secret. In the opened window, you can copy credentials or download them in the JSON format. You can also find them later in the **APIs and services > Credentials** section when editing your OAuth 2.0 client ID.

Client ID	587339997875-63b3pht11492fo0qfssprhd9up759.apps.googleusercontent.com
Client secret	GOCSPK-2WwLwB2Gh1dyWw59W2G-1LRkQ
Creation date	10 January 2023 at 16:39:15 GMT+4

4. Open **APIs and services > OAuth consent screen** and click **Edit App**. Specify your application name and the user support email and click **Save and continue**.

5. If your application is in the **Testing** status, you must specify test users. To do this, at the **Test users** step of the **Edit App** wizard, click **Add users**. To apply changes, click **Save and continue**. Note that only test users will have access to the app.

API

APIs and services

Enabled APIs and services

Library

Credentials

OAuth consent screen

Page usage agreements

Edit app registration

OAuth consent screen

Scopes

3 Test users

4 Summary

Test users

While publishing status is set to 'Testing,' only test users are able to access the app. Allowed user cap prior to app verification is 100, and is counted over the entire lifetime of the app. [Learn more](#)

+ ADD USERS

Filter

Enter property name or value

User information	
backup_notifications@tech.com	
test_email_notifications1@tech.com	
test_email_notifications2@tech.com	

SAVE AND CONTINUE

CANCEL

After you finish the registration, specify custom application registration settings when configuring the mail server for Google Gmail OAuth 2.0 authentication. For more information, see [Gmail Server Settings](#).

#### NOTE

You can leave your application in the **Testing** status and do not publish it. In that case, you will receive a warning message *Google hasn't verified this app* when connecting to your application. If you want to verify it, see [this Google article](#).

# Microsoft 365 Server Settings

## NOTE

Before connecting to the Microsoft 365 server, check the version of Microsoft .NET Framework installed on your computer. Microsoft 365 server modern authentication requires Microsoft .NET Framework 4.6.2 or later.

If an older version of Microsoft .NET Framework is installed, perform the upgrade before proceeding to authorization.

To connect to the Microsoft 365 server, specify the following settings in the **Configure and test email notifications** window:

1. In the **Mail server** section, specify the following:
  - a. Select the **Microsoft 365 (Modern authentication)** option from the drop-down list.
  - b. Click the **Authorize now** button and enter your Microsoft 365 account credentials or choose an account from the suggested list.

## IMPORTANT

When you are prompted to grant Veeam Agent access to your Microsoft 365 profile, accept the permissions listed in the opened window.

2. In the **Email settings** section, specify the following:
  - a. If necessary, change the email address in the **From email address** field to its alias.
  - b. If necessary, change the recipient email address in the **To email address** field. You can specify several recipient email addresses separated with a comma or semicolon.
  - c. In the **Email subject** field, specify a subject for the message. You can use the following variables in the subject:
    - i. *%JobResult%*
    - ii. *%ComputerName%*
    - iii. *%JobName%*
    - iv. *%CompletionTime%*
3. In the **Notify on** section, select the **Success**, **Warning** or **Error** check boxes to receive email notification if a job is run successfully, not successfully or with a warning.
4. [Optional] Click **Test now** to validate the Microsoft 365 server settings and send a test email.

To specify custom authentication options, do the following:

1. Click the **Server settings** button.
2. Select the **Use custom settings** check box.

3. Specify the application client ID and the tenant ID.

To learn how to register your custom application, see [Registering Application in Microsoft Azure Portal](#).

Configure and test email notifications

Mail server:

Microsoft 365 (Modern authentication) Server settings

✓ Application is authorized. [Re-authorize now.](#)

Email settings:

notifications@tech.local

administrators@tech.local

[%JobResult%] %ComputerName% - %JobName% - %CompletionTime% ⓘ

Notify on: ☒ Success ☒ Warning ☒ Error

Test Now Click to test specified server settings

## Registering Application in Microsoft Azure Portal

If you want to use your own web application for email notifications, you need to configure it in the Microsoft Azure portal. To do this, perform the following steps:

1. Log in to the [Microsoft Azure portal](#) under Exchange Online credentials that has permissions to register Azure AD applications.
2. Register the application. To do this, open **Azure Active Directory > App registrations** and click **New registration**:
  - a. In the **Name** field, specify the name of your application.
  - b. In the **Supported account types** section, select the **Accounts in this organizational directory only** option.

c. Click **Register**.

Home > My Directory Name | App registrations >

## Register an application

**\* Name**

The user-facing display name for this application (this can be changed later).

✓

### Supported account types

Who can use this application or access this API?

- ☒ Accounts in this organizational directory only (My Directory Name only - Single tenant)
- ☐ Accounts in any organizational directory (Any Azure AD directory - Multitenant)
- ☐ Accounts in any organizational directory (Any Azure AD directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)
- ☐ Personal Microsoft accounts only

[Help me choose...](#)

### Redirect URI (optional)

We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

Select a platform ▼

e.g. https://example.com/auth

By proceeding, you agree to the [Microsoft Platform Policies](#)

**Register**

After registration, you can copy application (client) ID and directory (tenant) ID. You can also find these credentials later in the **Overview** section of you application properties.

Home > My Directory Name | App registrations >

## email-app01

Search << Delete Endpoints Preview features

Overview

Quickstart

Integration assistant

Manage

Branding & properties

Authentication

### Essentials

Display name : [email-app01](#)

Application (client) ID : 83k1a03z-eb07-457a-b2fb-99n6095s4c1c

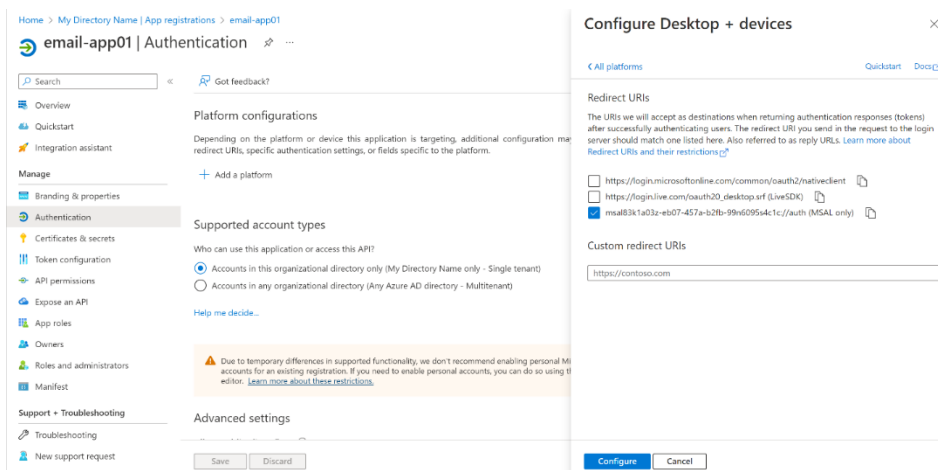
Object ID : 97p48830-3849-2671-77ku-35f465f22000

Directory (tenant) ID : 5582cvp7-02h5-a056-d8f4-ce8a3722da6g

Supported account types : [My organization only](#)

3. Add a platform configuration for your application. To do this, open **Authentication > Platform configurations** and click **Add a platform**:
  - a. Select **Mobile and desktop applications**.
  - b. Select the MSAL redirect URI generated in the following format: `msal<applicationid>://auth`.

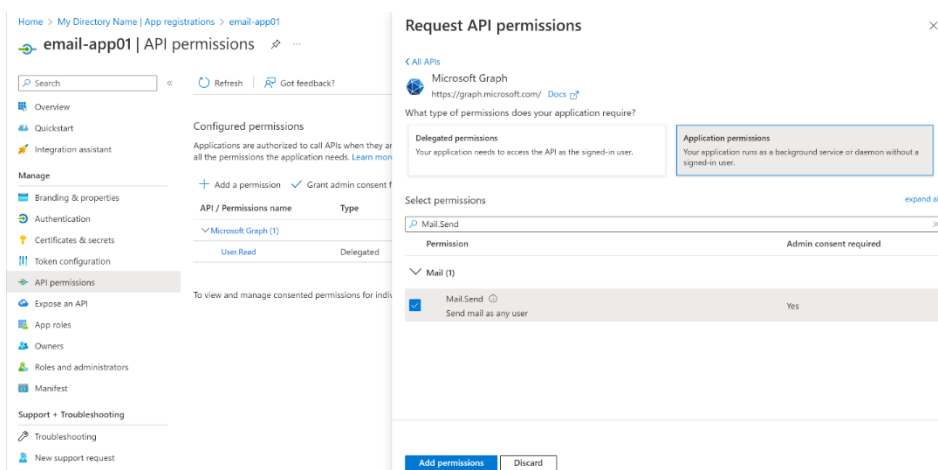
### c. Click **Configure**.



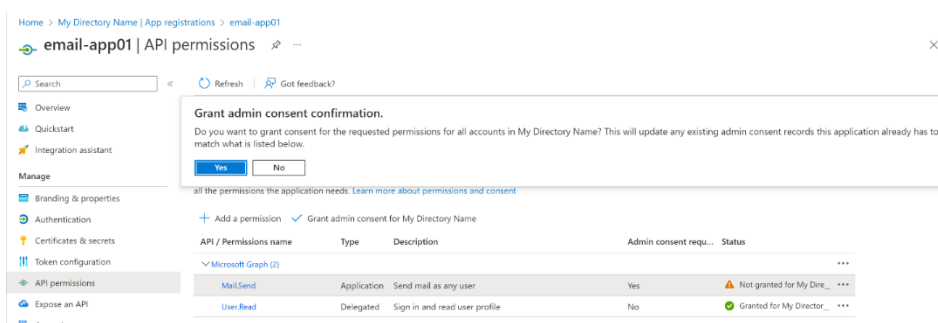
4. Grant the application the *Mail.Send* permission of Microsoft Graph. This will allow Veeam Agent for Microsoft Windows to call the Microsoft Graph API for sending email notifications. To do this, open **API permissions** and click **Add a permission**:

a. Select **Microsoft Graph > Application Permissions**.

b. Select the *Mail.Send* permission from the list and click **Add permissions**.



5. Click **Grant admin consent for <Your Directory Name>**. In the displayed window, click **Yes** to confirm the operation.



After you finish the registration, specify custom application registration settings when configuring the mail server for Microsoft 365 OAuth 2.0 authentication. For more information, see [Microsoft 365 Server Settings](#).

# Checking for New Product Versions and Updates

You can set up Veeam Agent for Microsoft Windows to automatically notify you about new product versions and updates. When a new version or patch becomes available, Veeam Agent for Microsoft Windows displays a notification in the notification bar. You can download the setup file and update Veeam Agent for Microsoft Windows. To learn more, see [Upgrading Veeam Agent for Microsoft Windows](#).

By default, automatic notifications are enabled. To disable notifications:

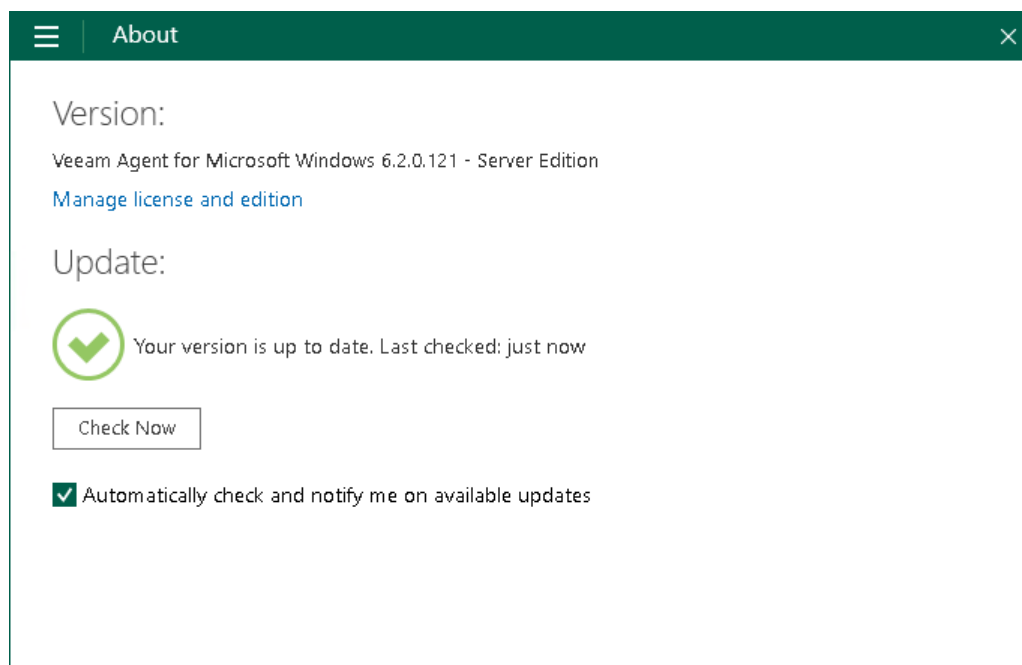
1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **About**.
3. In the **Update** section, clear the **Automatically check and notify me on available updates** check box.

To manually check if product updates are available, click **Check Now**.

## NOTE

Consider the following:

- To get information about available product updates, Veeam Agent for Microsoft Windows sends request to the Veeam Update Notification Server (`agents.butler.veeam.com`).
- For downloading setup files, Veeam Agent uses the Background Intelligent Transfer Service (BITS). If this service is disabled on the Veeam Agent computer, Veeam Agent for Microsoft Windows will not be able to download a setup file.



# Managing Veeam CBT Driver

You can set up Veeam Agent to use the Veeam CBT driver instead of the default CBT mechanism. The Veeam CBT driver offers more efficient changed block tracking mechanism that will be useful for machines running applications with large database files. To learn more about how the Veeam CBT driver works, see [Veeam Changed Block Tracking Driver](#).

You can perform the following operations with the Veeam CBT driver in Veeam Agent:

- [Install the Veeam CBT driver](#).
- [Remove the Veeam CBT driver](#).
- [Remove the Veeam CBT driver with Veeam Recovery Media](#).
- [Reset the Veeam CBT driver](#).



# Installing Veeam CBT Driver

You can install the Veeam CBT driver at any time you need. To use the Veeam CBT driver, the Veeam Agent computer must meet the following requirements:

- Run a 64-bit version of Microsoft Windows OS.
- Run one of the following OSes:
  - Microsoft Windows 11 (from version 21H2 to version 23H2).
  - Microsoft Windows 10 (from version 1803 to version 22H2).
  - Microsoft Windows Server OS that is supported by Veeam Agent. For more information, see [System Requirements](#).
- Run the Workstation or Server edition of Veeam Agent for Microsoft Windows.

## IMPORTANT

Consider the following:

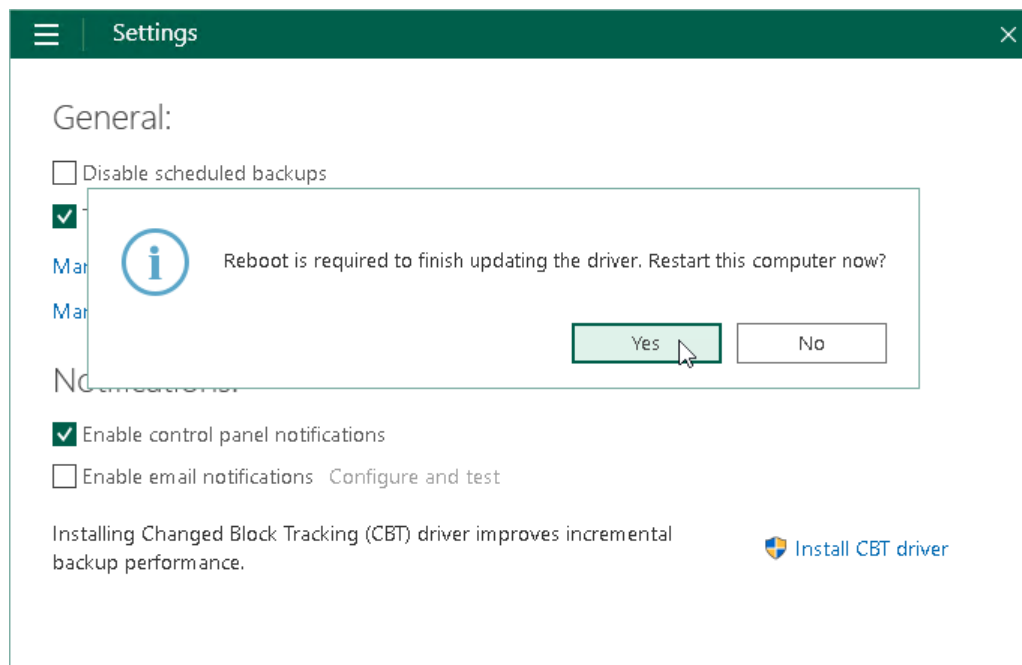
- Prior to installing the Veeam CBT driver on a computer running Microsoft Windows Server 2008 R2 SP1, make sure that update [KB3033929](#) is installed in the OS.  
The update adds the SHA-2 code signing support that is required for verification of the Veeam CBT driver signature. Without this update installed, the OS running on a protected computer will fail to boot after you install the Veeam CBT driver. To learn more, see [this Microsoft KB article](#).
- Do not install the Veeam CBT driver on a computer running Microsoft Windows Server 2008 R2 SP1, 2012 or 2012 R2 if one or more volumes on this computer are encrypted with Microsoft BitLocker (or other encryption tool), or if you plan to use Microsoft BitLocker to encrypt volumes on this computer. Concurrent operation of Microsoft BitLocker and Veeam CBT driver may result in driver failures and may prevent the OS from starting.
- Do not install the Veeam CBT driver on a computer if you plan to use devices with hardware encryption made according to the TCG Opal Security Subsystem Class Specification. Operation of the driver on such devices may lead to a crash of the operating system. To learn more about the TCG Opal Security Subsystem Class Specification, see [this Trusted Computing Group webpage](#).

To install the Veeam CBT driver:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Click **Install CBT driver**.

4. To complete the installation process, Veeam Agent for Microsoft Windows needs to reboot the computer. To reboot the computer immediately, in the displayed window, click **Yes**. After Veeam Agent for Microsoft Windows reboots the computer, the driver will start tracking blocks that are changing on the volumes whose data you chose to back up in the backup job settings.

If you choose not to reboot the computer immediately, Veeam Agent for Microsoft Windows will continue to use the default CBT mechanism until the next computer reboot.



# Removing Veeam CBT Driver

You can quickly remove the Veeam CBT driver, for example, if your Veeam Agent computer does not run applications with large database files any more, and you do not need to perform advanced change block tracking on this computer.

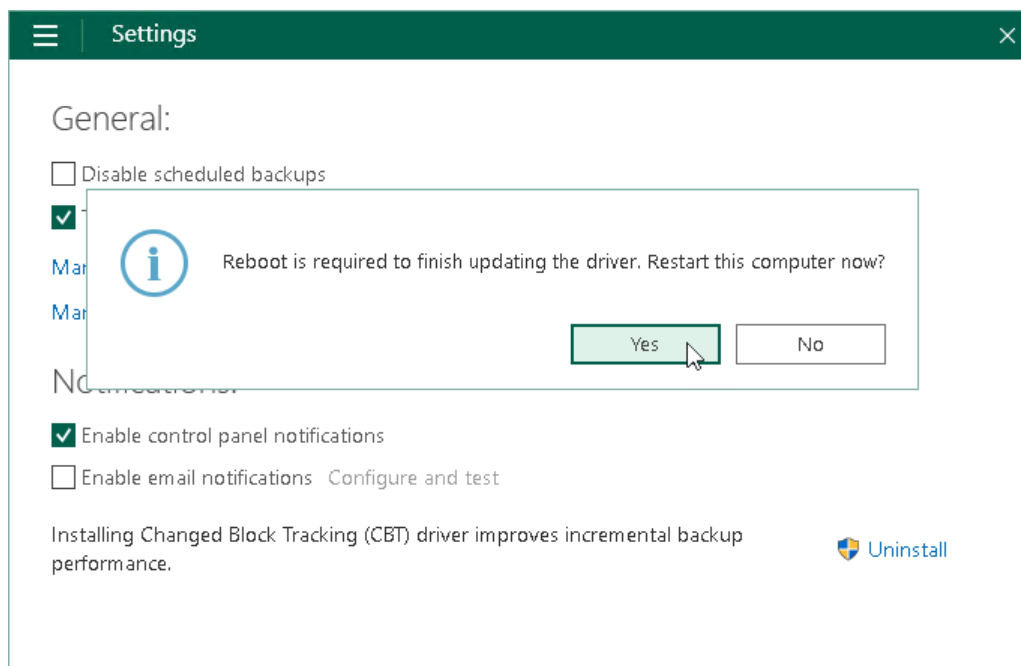
To remove the Veeam CBT driver:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Settings**.
3. Click **Uninstall**.
4. To complete the uninstallation process, Veeam Agent for Microsoft Windows needs to reboot the computer. To reboot the computer immediately, in the displayed window, click **Yes**. After computer reboot, Veeam Agent for Microsoft Windows will use the default CBT mechanism to get the list of changed data blocks.

## TIP

You can also uninstall the Veeam CBT driver in the Microsoft Windows control panel:

1. From the **Start** menu, select **Control Panel > Programs and Features**.
2. In the programs list, right-click **Veeam CBT Driver** and select **Uninstall**.
3. In the displayed window, click **Yes**.



# Removing CBT Driver with Veeam Recovery Media

You can use the Veeam Recovery Media to remove the Veeam CBT driver from your Veeam Agent computer. This operation may be required, for example, if the OS on your computer fails to start after you have installed the Veeam CBT driver in Veeam Agent for Microsoft Windows.

To remove the Veeam CBT driver:

1. [Boot from the Veeam Recovery Media](#).
2. On the Veeam Recovery Media screen, click **Tools > Command Prompt** or press **[Shift] + [F10]**.
3. Use a command with the following syntax:

```
X:\VeeamRecovery\Veeam.Endpoint.Recovery.exe -RemoveVeeamCBTDriver
```

4. Reboot the Veeam Agent computer.

## NOTE

Veeam Agent for Microsoft Windows will remove the Veeam CBT Driver from the Veeam Agent computer. However, a record about the driver will remain in the Microsoft Windows control panel. To remove the record, from the **Start** menu, select **Control Panel > Programs and Features**. Then right-click **Veeam CBT Driver** in the programs list and select **Uninstall**.

# Resetting CBT

In some cases, it may be required to reset CBT data collected by the Veeam CBT driver. For example, this may be necessary if you want to avoid performing active full backup after a volume was changed in a non-Windows OS.

To reset CBT, run the command line interface with administrative privileges and use one of the following commands:

- To reset CBT for all volumes of the Veeam Agent computer, use a command with the following syntax:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" RESETCBT all
```

- To reset CBT for a specific volume, use a command with the following syntax:

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" RESETCBT %volumeMountPoint%
```

or

```
"C:\Program Files\Veeam\Endpoint Backup\Veeam.EndPoint.Manager.exe" RESETCBT %volumeUUID%
```

where:

- %volumeMountPoint% — mount point of the volume, for example: C:\.
- %volumeUUID% — ID of the volume, for example: \\?\Volume{1214be80-1165-41e5-8244-8fbf79d85c31}.

After CBT reset, during the next backup job session, Veeam Agent for Microsoft Windows will create incremental backup. The backup job session will require greater time, because Veeam Agent for Microsoft Windows will need to read all data from the backed-up volumes.

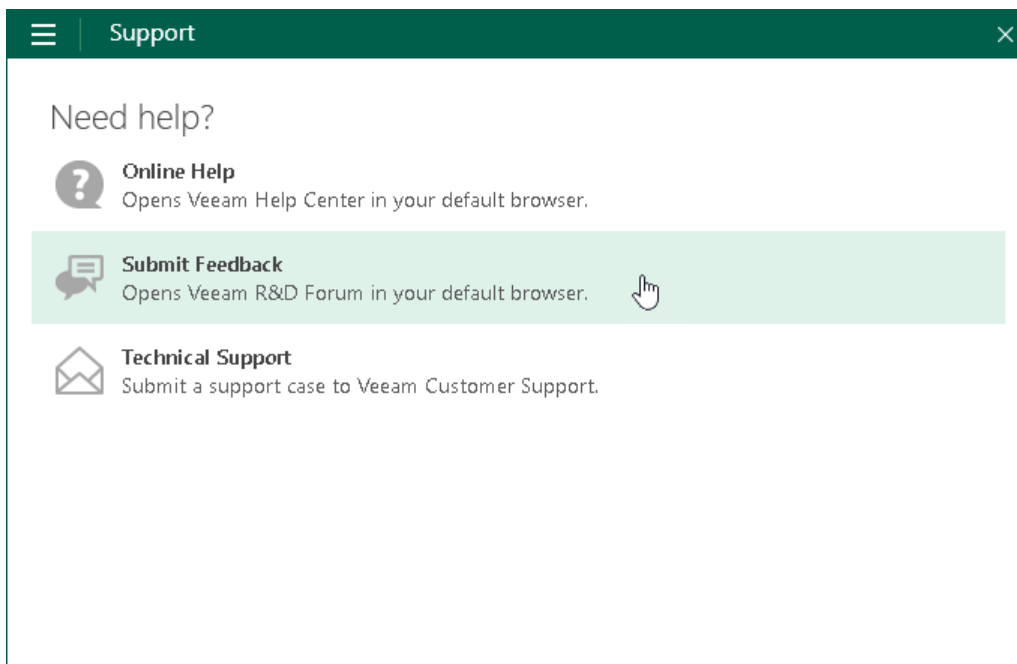
# Getting Support

If you have any questions or want to share your feedback about Veeam Agent, you can use one of the following options:

- You can search for the information on the necessary subject in the current Veeam Agent for Microsoft Windows User Guide.
- You can visit [Veeam R&D Forums](#) and share your opinion or ask a question.
- You can submit a support case to the Veeam Customer Support directly from the product. To learn more, see [Reporting Issues](#).

To access help and support options in Veeam Agent:

1. Right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Support**.
3. Click one of available options to get support on the product.



# Reporting Issues

You can submit a support case in the Control Panel. To submit a support case:

1. Double-click the Veeam Agent for Microsoft Windows icon in the system tray, or right-click the Veeam Agent for Microsoft Windows icon in the system tray and select **Control Panel**.
2. From the main menu, select **Support**.
3. Click **Technical Support**.
4. In the email field of the **Report an issue** window, enter a valid email address.

If the email address that you have entered is not registered at the [Veeam Customer Support Portal](#), click **Register** on the right of the email field. Veeam Software will register your email address and send you a verification email to the specified address. When you receive a verification email, open it and click a link provided in the email to complete the verification procedure. After the verification procedure is complete, you will be able to submit a support case.



5. In the description fields, enter a short and detailed description of your problem.
6. Select the **I agree that debug logs will be uploaded to Veeam servers automatically** check box and click **Submit Case**.

Veeam Agent will automatically collect logs from your computer (without additional warnings) and open a support case at the [Veeam Customer Support Portal](#).

## IMPORTANT

Consider the following:

- If you have any questions about the product functionality, do not submit a support case through the [Veeam Customer Support Portal](#) and do not send an email to the Veeam Customer Support directly. To submit a support case, use the Control Panel in Veeam Agent for Microsoft Windows.
- You can submit a support case only in the Control Panel of the current version of Veeam Agent. If you use an older version of Veeam Agent, upgrade Veeam Agent and check whether the problem still exists in the current version. If the problem exists, use the Control Panel to submit a support case.

 Report an issue 

Paid edition users may open support cases using the form below or by using Customer Support Portal. Free edition users are encouraged to report issues using the form below – response will be provided to the specified email address on best effort basis depending on support staff availability, but is not guaranteed.

Hi, I have a problem. The tray icon in the System Tray is gray, and Veeam Agent reports that the tray agent is not connected to the Veeam Agent for Microsoft Windows service.

What can I do to fix this problem?

Thank you,  
John

☒ I agree that debug logs will be uploaded to Veeam servers automatically

Submit Case



# Using with Veeam Backup & Replication

If you have the Veeam backup infrastructure deployed in the production environment, you can use Veeam Agent together with Veeam Backup & Replication.

## IMPORTANT

If you plan to use Veeam Agent for Microsoft Windows 6.2 with Veeam Backup & Replication, you must install Veeam Backup & Replication 12.2 on the Veeam backup server.

## NOTE

This and subsequent sections describe tasks with Veeam Backup & Replication available for Veeam Agent operating in the standalone mode. For information about tasks available in Veeam Backup & Replication within the Veeam Agent management scenario, see the [Veeam Agent Management Guide](#).

## Tasks with Veeam Backup & Replication

Veeam Backup & Replication lets you perform a number of additional data protection and disaster recovery tasks, as well as administrative actions with Veeam Agent backups. You can:

- [Grant access permissions on backup repositories.](#)
- [Manage Veeam Agent licenses.](#)

### *Data protection tasks*

- [Create Veeam Agent backups on backup repositories.](#)
- [Create Veeam Agent backups on Veeam Cloud Connect repositories.](#)
- [Copy Veeam Agent backups to secondary backup repositories.](#)
- [Archive Veeam Agent backups to tape.](#)

### *Restore tasks*

- [Restore Veeam Agent backups to Hyper-V VMs.](#)
- [Restore Veeam Agent backups to VMware vSphere VMs.](#)
- [Restore Veeam Agent backups to Nutanix VMs.](#)
- [Restore volumes from Veeam Agent backups.](#)
- [Restore files and folders from Veeam Agent backups..](#)
- [Restore application items from Veeam Agent backups.](#)
- [Restore disks from Veeam Agent backups.](#)
- [Publish disks to analyze backup content.](#)
- [Restore data from Veeam Agent backups to Amazon EC2.](#)
- [Restore data from Veeam Agent backups to Microsoft Azure.](#)
- [Restore data from Veeam Agent backups to Google Compute Engine.](#)

- [Export restore points of Veeam Agent backups to standalone full backup files.](#)

#### *Administrative tasks*

- [Import Veeam Agent backups.](#)
- [Enable and disable Veeam Agent backup jobs.](#)
- [View Veeam Agent backup job statistics.](#)
- [Delete Veeam Agent backup jobs.](#)
- [View Veeam Agent backup properties.](#)
- [Create recovery tokens.](#)
- [Copy Veeam Agent backups.](#)
- [Remove Veeam Agent backups.](#)
- [Delete Veeam Agent backups.](#)
- [Configure global settings.](#)
- [Assign roles to users.](#)

# Setting Up User Permissions on Backup Repositories

To be able to store backups in a backup repository managed by a Veeam backup server, the user must have access permissions on this backup repository.

## IMPORTANT

Veeam Agent for Microsoft Windows does not support Veeam backup repositories with enabled KMS encryption. To learn more about KMS encryption for Veeam backup repositories, see the [Key Management System Keys](#) section in the Veeam Backup & Replication User Guide.

## NOTE

If you plan to create backups in a Veeam backup repository with Veeam Agent backup jobs configured in Veeam Backup & Replication, you do not need to grant access permissions on the backup repository to users. In the Veeam Agent management scenario, to establish a connection between the backup server and protected computers, Veeam Backup & Replication uses a TLS certificate. To learn more, see the [Configuring Security Settings](#) section in the Veeam Agent Management Guide.

Access permissions are granted to security principals such as users and AD groups by the backup administrator working with Veeam Backup & Replication. Users with granted access permissions can target Veeam Agent backup jobs at this backup repository and perform restore from backups located in this backup repository.

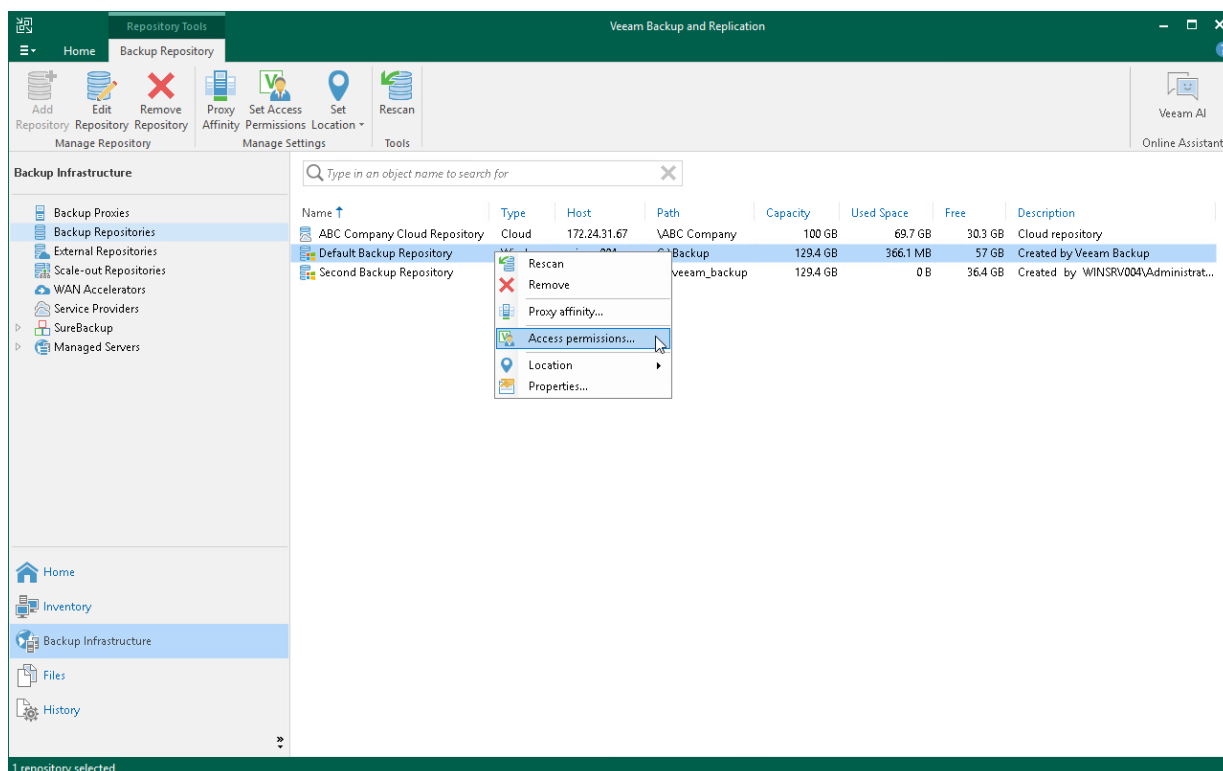
Right after installation, access permissions on the default backup repository are set to *Allow to everyone* for testing and evaluation purposes. If necessary, you can change these settings.

After you create a new backup repository, access permissions on this repository are set to *Deny to everyone*. To allow users to store backups in the backup repository, you must grant users with access permissions to this repository.

To grant access permissions to a security principal:

1. In Veeam Backup & Replication, open the **Backup Infrastructure** view.
2. In the inventory pane, click one of the following nodes:
  - The **Backup Repositories** node – if you want to grant access permissions on a regular backup repository to Veeam Agent users.
  - The **Scale-out Repositories** node – if you want to grant access permissions on a scale-out backup repository to Veeam Agent users.


3. In the working area, select the necessary backup repository and click **Set Access Permissions** on the ribbon, or right-click the backup repository and select **Access permissions**. If you do not see the **Set Access Permissions** button on the ribbon or the **Access permissions** command is not available in the shortcut menu, press and hold the [Ctrl] key, right-click the backup repository and select **Access permissions**.



4. In the **Access Permissions** window, in the **Standalone applications** tab, specify to whom you want to grant access permissions on this backup repository:
  - **Allow to everyone** – select this option if you want all users to be able to store backups on this backup repository. Setting access permissions to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). However, we recommend this scenario for demo environments only.
  - **Allow to the following accounts or groups only** – select this option if you want only specific users to be able to store backups on this backup repository. Click **Add** to add the necessary users and groups to the list.
5. If you want to encrypt Veeam Agent backup files stored in the backup repository, select the **Encrypt backups stored in this repository** check box and choose the necessary password from the field below. If you have not specified a password beforehand, click **Add** on the right or the **Manage passwords** link to add a new password. Veeam Backup & Replication will encrypt files at the backup repository side using its built-in encryption mechanism. To learn more, see [Veeam Backup & Replication Documentation](#).

## IMPORTANT

If Veeam Agent is set up to use the backup cache, and the backup cache contains one or more restore points, Veeam Agent will automatically remove these restore points from the backup cache after you enable or disable the encryption option for the backup repository.



# Managing License

If you plan to use Veeam Agent with Veeam Backup & Replication, you must install a license in Veeam Backup & Replication or Veeam Backup Enterprise Manager. The license must have a total number of instances that is sufficient to protect machines (servers and workstations) on which you plan to install Veeam Agent. For more information, see [Veeam Licensing Policy](#).

After Veeam Agent connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming instances in the license. The product edition for Veeam Agent is selected depending on the type of the OS running on the Veeam Agent computer. You can switch to another commercial edition of Veeam Agent manually if needed. If you do not want Veeam Agents to consume instances, you can restrict instance consumption. For more information, see [Managing Instance Consumption by Veeam Agents](#).

The number of backup jobs configured in Veeam Agent does not impact instance consumption. For example, if 2 backup jobs are configured in Veeam Agent that operates in the Server edition, this Veeam Agent will consume instances required for 1 server.

Veeam Agent obtains information about the license from Veeam Backup & Replication and keeps it locally on the Veeam Agent computer. Information about the license is valid for 32 days. If Veeam Agent does not connect to Veeam Backup & Replication during this period, Veeam Backup & Replication will revoke its license.

## NOTE

In addition to managing Veeam Agent licenses, you can use the Veeam Backup & Replication console to manage Veeam Agent backup jobs and perform operations with backups created by these jobs.

If your backup server is connected to Veeam Backup Enterprise Manager, you can use Veeam Backup Enterprise Manager to manage licenses and perform restore tasks with Veeam Agent backups. You cannot manage Veeam Agent backup jobs with Veeam Backup Enterprise Manager.

For more information on Veeam Backup & Replication licensing, see the [Licensing](#) section in the Veeam Backup & Replication User Guide.

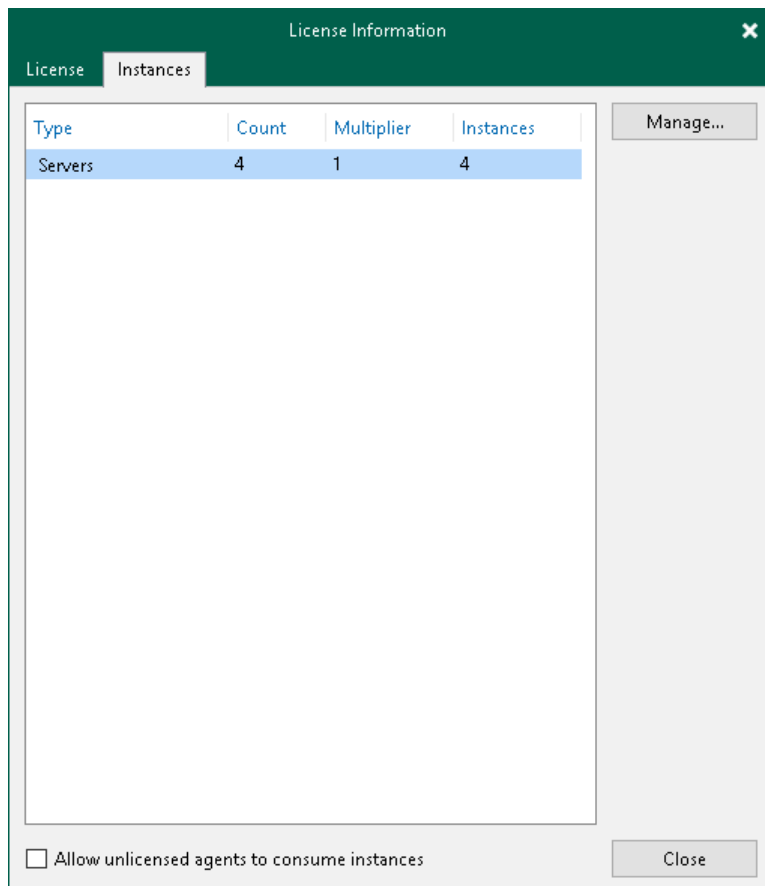
# Managing Instance Consumption by Veeam Agents

By default, Veeam Backup & Replication allows Veeam Agents to connect to the Veeam backup server and consume instances in the license. If you do not want Veeam Agents to consume instances, you can restrict instance consumption.

If you restrict instance consumption, Veeam Backup & Replication will switch all Veeam Agents connected to this Veeam backup server to the free edition that offers limited capabilities. For information about Veeam Agent editions, see [Product Editions](#).

To restrict instance consumption by Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, click the **Instances** tab.
3. On the **Instances** tab, clear the **Allow unlicensed agents to consume instances** check box.
4. Click **Close**.



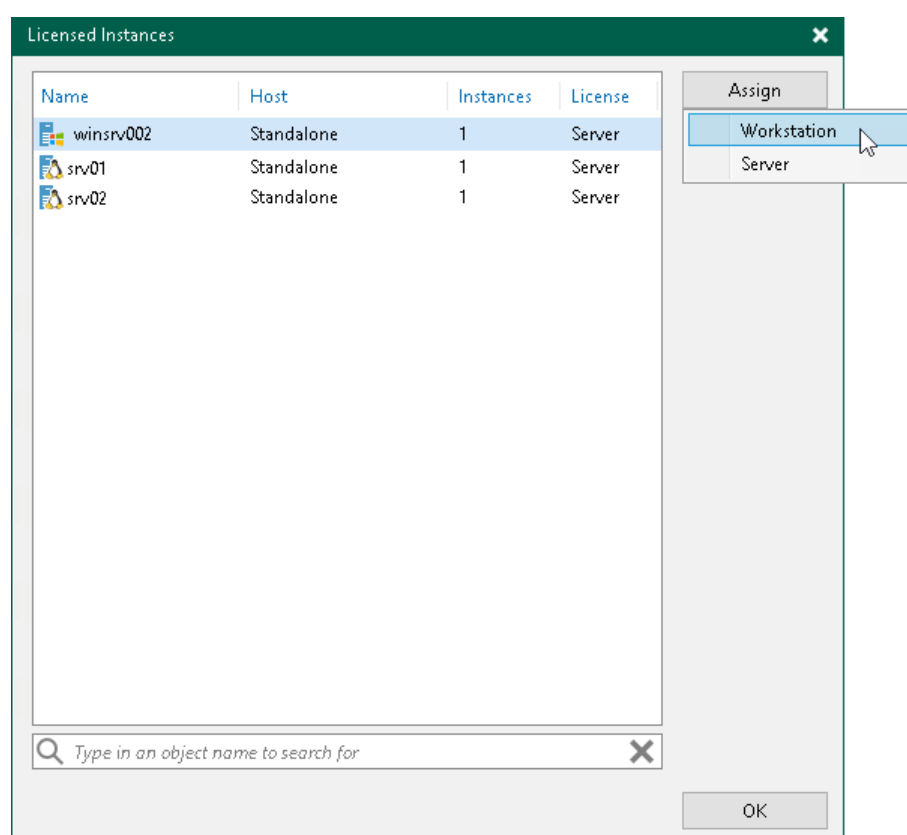
# Assigning License to Veeam Agent

After Veeam Agent connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming the license. The product edition for Veeam Agent is selected depending on the type of the OS running on the protected computer.

You can also assign a license to Veeam Agent manually if needed. When you assign a license, you can select the product edition, too.

To assign a license:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the **Licensed Instances** window, select the Veeam Agent to which you want to assign the license, click **Assign** and select the desired product edition: *Workstation* or *Server*.





# Viewing Licensed Veeam Agents and Revoking License

When Veeam Agent connects to the backup server, Veeam Backup & Replication applies a license to the Veeam Agent. You can view to which Veeam Agents the license is currently applied.

To view a list of licensed Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.

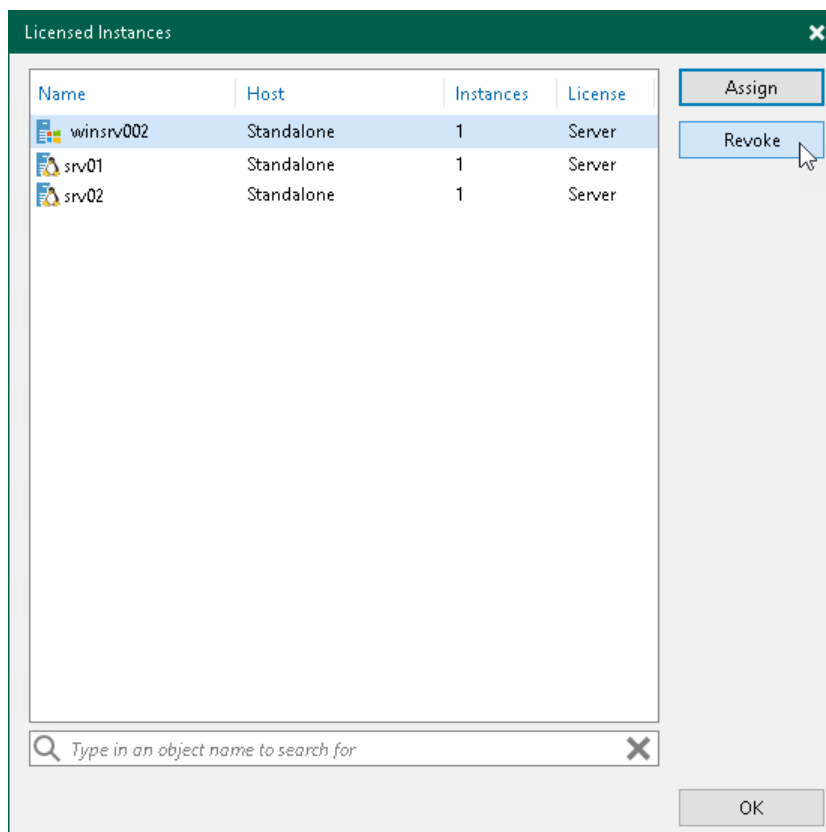
In the list of licensed instances, Veeam Backup & Replication displays Veeam Agents that have established a connection with the backup server when you created the backup job.

## Revoking License from Veeam Agents

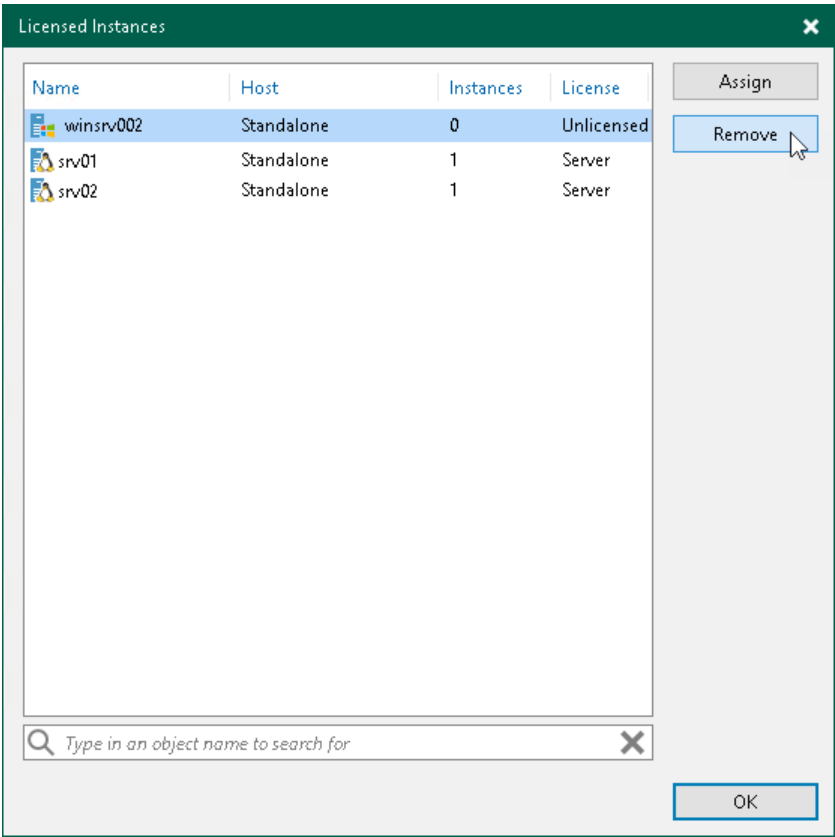
You can revoke the license from some Veeam Agents and re-apply it to other protected workloads. License revoking can be helpful, for example, if you do not want to use some Veeam Agents with Veeam Backup & Replication anymore.

To revoke a license from the Veeam Agent:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the Licensed Instances window, select a Veeam Agent and click **Revoke**. Veeam Backup & Replication will revoke the license from the Veeam Agent, and the license will be freed for other workloads that you want to protect with Veeam products.



The Veeam Agent from which you have revoked the license will become unable to connect to the Veeam backup server but will remain in the **Licensed Instances** list. To allow this Veeam Agent to create backups in the Veeam backup repository, select the Veeam Agent and click **Remove**. During the next backup job session, the Veeam Agent will connect to the Veeam backup server and start consuming the license.



# Performing Data Protection Tasks

You can perform the following data protection tasks:

- Back up your data and store the resulting backup files in one of the following types of Veeam backup repositories:
  - [In a backup repository managed by a Veeam backup server](#)
  - [In a Veeam Cloud Connect repository](#)
- [Copy Veeam Agent backups from the backup repository to a secondary backup repository with backup copy jobs.](#)
- [Use SureBackup.](#)
- [Archive Veeam Agent backups to tapes with backup to tape jobs.](#)
- [Scan Veeam Agent backups.](#)

# Backing Up to Backup Repositories

You can store backups created with Veeam Agent in backup repositories connected to Veeam backup servers. To do this, you must perform the following actions:

1. [Set up user permissions at the backup repository side.](#)
2. [Point the Veeam Agent backup job to the backup repository.](#)


## NOTE

Consider the following:

- A Veeam Agent backup job can be started automatically upon the defined schedule or manually from the Veeam Agent computer. You cannot start, stop, retry or edit Veeam Agent backup jobs in the Veeam Backup & Replication console.
- If the user is granted restore permissions on the Veeam backup server, the user will be able to see all backups in the backup repository.
- The user who creates a Veeam Agent backup in the backup repository is set as the owner of the backup file. The backup file owner can access this file and restore data from it. If the user who is not the backup file owner needs to perform operations with the backup file, the user must have the Veeam Backup & Replication role that allows to perform these operations. To learn more about roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

Backup jobs targeted at the backup repository become visible in Veeam Backup & Replication under the **Jobs > Backup** node in the **Home** view. Backups created with Veeam Agent are available under the **Backups > Disk** node in the **Home** view.

The Veeam Backup Administrator working with Veeam Backup & Replication can manage Veeam Agent backup jobs and restore data from Veeam Agent backups. To learn more, see [Restoring Data from Veeam Agent Backups](#) and [Performing Administration Tasks](#).



# Backing Up to Cloud Repositories

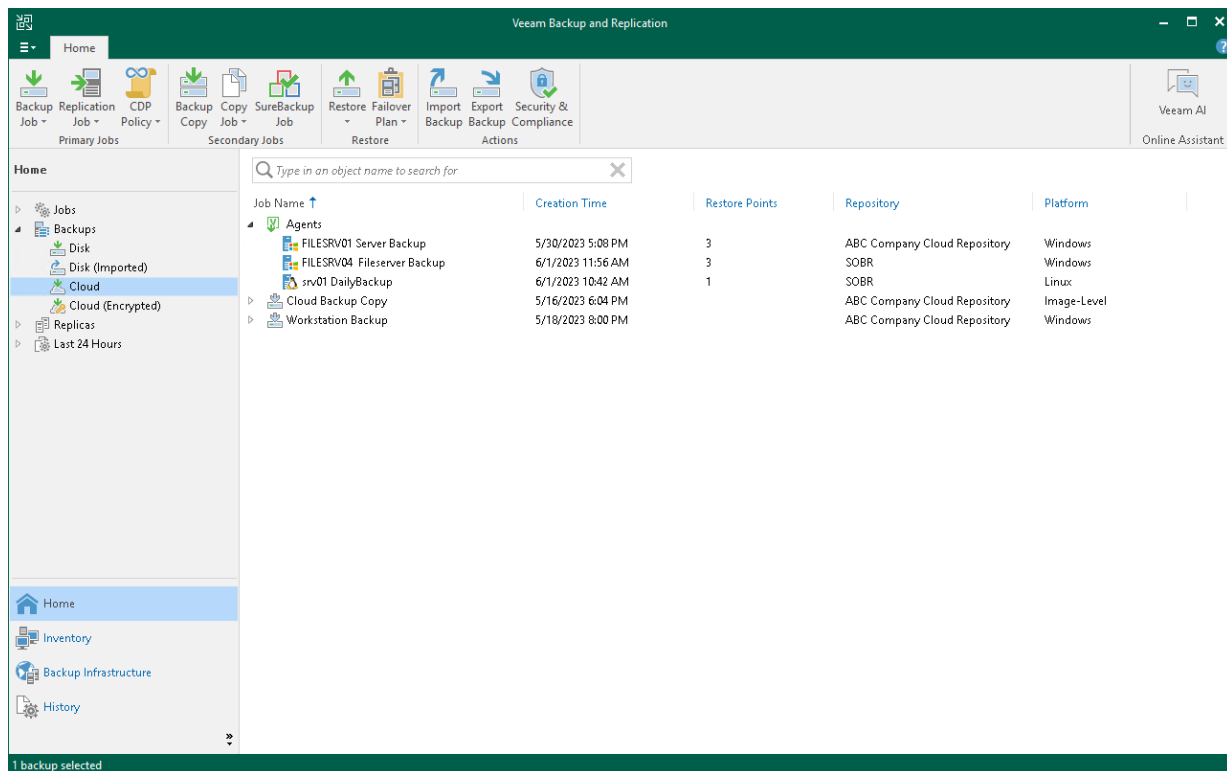
You can store backups created with Veeam Agent in cloud repositories provided to you by a Veeam Cloud Connect service provider. To do this, you must connect to the service provider and point the backup job to the cloud repository. To learn more, see [Specify Service Provider Settings](#).

## Veeam Agent Backups on Tenant Side

Backups created with Veeam Agent are available under the **Cloud** node in the **Home** view of the Veeam Backup & Replication console deployed on the tenant side.

The backup administrator working with Veeam Backup & Replication on the tenant side can manage Veeam Agent backups created in the cloud repository and restore data from such backups. To recover data from a Veeam Agent backup, you can perform the following operations:

- [Export computer disks as virtual disks](#).
- [Restore guest OS files](#).
- [Export restore points to standalone full backup files](#).



## Veeam Agent Backups on Service Provider Side

The service provider can view information about backup and restore sessions performed by Veeam Agent users. The full list of sessions is available in the **History** view of the Veeam backup console deployed on the service provider side. The list of sessions performed within the last 24 hours is available under the **Last 24 hours** node in the **Cloud Connect** view of the Veeam backup console on the service provider side. The service provider cannot view detailed statistics about individual sessions in the list.

The service provider cannot perform restore tasks with Veeam Agent backups that are stored in the cloud repository. The service provider can perform the following restore tasks with unencrypted Veeam Agent backups stored in the cloud repository:

- Instant recovery
- Disk restore
- Disk publish

To learn more, see the [Restoring Data from Tenant Backups](#) section in the Veeam Cloud Connect Guide.

HomeView

Backup Job

Replication Job

CDP Policy

Backup Copy

SureBackup Job

Restore

Import Backup

Export Backup

Security & Compliance

Primary JobsSecondary JobsRestoreActions

Veeam AI  
Online Assistant

Cloud Connect

Cloud Connect

Cloud Gateways

Gateway Pools

Tenants

Backup Storage

Replica Resources

Last 24 Hours

Home

Inventory

Backup Infrastructure

Storage Infrastructure

Cloud Connect

Type in an object name to search for

Job Name	Session Type	Status	Start Time	End Time	Tenant	Data Sent	Data Received
srv001_daily_backup	Cloud Backup	Success	11/30/2023 11:24 AM	11/30/2023 11:25 AM	ABC Company	54.1 KB	655 MB
srv001_volume_backup	Cloud Backup	Success	11/30/2023 11:26 AM	11/30/2023 11:27 AM	ABC Company	15.6 KB	15.8 KB
srv002_system_backup	Cloud Backup	Success	11/30/2023 11:25 AM	11/30/2023 11:26 AM	ABC Company	97.8 KB	2.2 MB
srv003_system_backup	Cloud Backup	Success	11/30/2023 11:19 AM	11/30/2023 11:20 AM	ABC Company	40.6 KB	20.5 MB

5 sessions

# Performing Backup Copy for Veeam Agent Backups

You can configure backup copy jobs that will copy backups created with Veeam Agent to a secondary backup repository.

Backup copy jobs treat Veeam Agent backups as usual backup files. The backup copy job setup and processing procedures practically do not differ from the same procedures for a backup copy job that processes VM backups. To learn more about backup copy jobs, see the [Backup Copy](#) section in the Veeam Backup & Replication User Guide.

When mapping a backup copy job to a Veeam Agent backup, consider the limitations listed in the [Map Backup File](#) section in the Veeam Backup & Replication User Guide.

**New Backup Copy Job**

**Objects**  
Add objects which backups should be mirrored to the target repository. Immediate backup copy job will process image-level and transaction log backups.

**Job**

**Objects**

**Target**

**Data Transfer**

**Schedule**

**Summary**

**Objects to process:**

Name	Type	Size
System Backup	Agent for Windows	1.43 GB
Default Backup Repository	Repository	1.44 GB
Second Backup Repository	Repository	1.44 GB

**Add...**  
**Remove**  
**Exclusions...**  
**Recalculate**

**Total size: 4.32 GB**

☒ Include database transaction log backups (increases bandwidth usage)

**< Previous** **Next >** **Finish** **Cancel**

## Restoring Data from Copies of Veeam Agent Backups

Backups copied to the secondary backup repository do not preserve user access permissions. At the same time, users who created backups do not have access permissions on these secondary repositories. For this reason, users cannot restore data from their backups residing in the secondary site.

To overcome this limitation, you can delegate the restore task to backup administrators who work with Veeam Backup & Replication. Backup administrators can use Veeam Backup & Replication options to recover data from such backups: for example, perform file-level restore or retrieve necessary application items with Veeam Explorers.

You can also restore data from the copied backup stored in the target repository using Veeam Agent.

To do this:

1. In Veeam Agent, launch the **Volume Level Restore** wizard to restore volumes or **File Level Restore** wizard to restore files and folders. You can also boot from the Veeam Recovery Media and launch the **Veeam Recovery Media** wizard for data restore.
2. At the **Backup Location** step of the wizard, select **Network storage**.
3. At the **Network Storage** step of the wizard, select to restore data from the backup repository.
4. At the **Backup Server** wizard, specify settings for the Veeam backup server that manages the target backup repository where the copied backup is located.
5. Select the **Specify your personal credentials** check box and provide credentials for the user who has the *Veeam Backup Administrator* or *Veeam Restore Operator* role on the Veeam backup server.
6. Pass through the next steps of the wizard and select a backup and restore point from which you want to restore data.

The screenshot shows the 'Volume Level Restore' wizard window. The left sidebar contains the following steps: Backup Location, Network Storage, Backup Server (highlighted), Backup, Restore Point, Disk Mapping, Summary, and Progress. The main area is titled 'Backup Server' and includes the instruction: 'Specify a Veeam Backup & Replication server name and authentication method. You can use Windows credentials in the DOMAIN\USERNAME format or a recovery token from your backup administrator.' Below this, there is a text field for 'Veeam backup server name or IP address' containing '172.17.53.13' and a 'Port' spinner set to '10001'. A checkbox labeled 'Specify your personal credentials:' is checked. Below it, the 'Username' field contains 'VEEAM\Administrator' and the 'Password' field is masked with dots. At the bottom right, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.



# Using SureBackup

Veeam Backup & Replication offers the SureBackup technology to test backups and check if you can recover data from them. You can verify any restore point of a backed-up computer protected with Veeam Agent for Microsoft Windows.

To learn more about the logic behind SureBackup, see the [How SureBackup Works](#) section in the Veeam Backup & Replication User Guide.

Before creating the SureBackup job, check limitations for Veeam Agent backups below. Then launch the **New SureBackup Job** wizard to create the SureBackup job. To learn more, see the [Creating SureBackup Job](#) section in the Veeam Backup & Replication User Guide.

## Limitations

For backups created with Veeam Agent for Microsoft Windows, SureBackup has the following limitations:

- SureBackup is not supported for backups stored in the Veeam Cloud Connect repository.
- SureBackup is not supported for backups stored in the archive tier of the the scale-out backup repository.
- [For full recoverability testing mode] SureBackup is not supported for backups containing drives greater than 64 TB.
- [For full recoverability testing mode] If you plan to verify computer recovery with VMware vSphere, consider the following:
  - SureBackup is not supported for backups of 4 KB sector drives.
  - SureBackup is not supported for backups of storage spaces.
  - SureBackup is not supported for backups containing more than 54 drives.
- [For full recoverability testing mode] When Veeam Backup & Replication publishes virtual machines based on backed-up Veeam Agent computers in the isolated virtual environment, all these virtual machines are included in the first isolated network added during the virtual lab configuration. To learn more, see the [Create Isolated Networks](#) section in the Veeam Backup & Replication User Guide.
- [For full recoverability testing mode] SureBackup is not supported for file-level backups. You must use entire machine or volume-level backup of the protected computer. The backup must include the computer system volume. To learn more about backup types, see [Backup Types](#).
- [For full recoverability testing mode] SureBackup is not supported if the Microsoft Windows system partition and boot partition of the backed-up computer are located on different drives.
- [For full recoverability testing mode] SureBackup is not supported for failover clusters.
- [For full recoverability testing mode] If you plan to verify computer recovery with Microsoft Hyper-V, SureBackup is not supported for application groups with computers connected to different networks.
- [For full recoverability testing mode] If you plan to verify computer recovery with Microsoft Hyper-V, SureBackup is not supported for EFI-based Veeam Agent computers that run Windows 7, Windows Server 2008 or Windows Server 2008 R2.

# Archiving Veeam Agent Backups to Tape

You can configure backup to tape jobs to archive Veeam Agent backups to tape.

Backup to tape jobs treat Veeam Agent backups as usual backup files. The archiving job setup and processing procedures practically do not differ from the regular ones. To learn more about backup to tape jobs, see the [Backup to Tape](#) section in the Veeam Backup & Replication User Guide.

## NOTE

For the **After this job** option in the backup to tape job schedule settings, you cannot select a backup job managed by Veeam Agent or a standalone Veeam Agent backup job as the preceding backup job.

[illegible]

# Scanning Backup

If you want to scan restore points of a backup after a malware attack or to look for some sensitive data in a backup, you can run the scan backup session:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. In the working area, expand the Veeam Agent backup, select the necessary computer in the backup and click **Scan Backup** on the ribbon or right-click the computer and select **Scan Backup**.
4. Specify the scan mode you want to use:
  - **Find the last clean restore point**
  - **Find the last clean restore point in range**
  - **Scan content of all restore points in range**
5. If you want to use antivirus software as a scan engine, select the **Scan restore points with an antivirus engine** check box. For more information, see the [Secure Restore](#) section in the Veeam Backup & Replication User Guide.
6. If you want to use a YARA rule as a scan engine, select the **Scan restore points with the following YARA rule** check box and specify the YARA file located in the Veeam Backup & Replication product folder:  
C:\Program Files\Veeam\Backup and Replication\Backup\YaraRules.

If you do not want to create a malware detection event for a YARA rule, you can add a `SuppressMalwareDetectionNotification` tag to the name of the rule. For example:

```
rule SearchFileHash : SuppressMalwareDetectionNotification
```

In this case, the malware detection event will not be created but the scan backup session will be finished with the *Warning* status.

7. Configure the scan range. You can specify the following options:
  - Scan all restore points, from most recent restore point to the oldest available restore point.
  - Scan restore points created during a specific time period.

Veeam Backup & Replication selects the order in which to scan restore points depending on the selected scan mode:

- In the **Find the last clean restore point** mode, Veeam Backup & Replication scans restore points from the most recent to the oldest.
- In the **Find the last clean restore point in range** mode, Veeam Backup & Replication scans restore points in the optimal order.
- In the **Scan content of all restore points in range** mode, Veeam Backup & Replication scans restore points from the oldest to the most recent.

If you want to continue the scan backup session after the first malware or the first piece of specific information is found, select the **Continue scanning all remaining files after the first occurrence** check box.

8. Click **OK**.

To learn more, see the [Scan Backup](#) section in the Veeam Backup & Replication User Guide.

Scan Backup

✕

Performs an ad-hoc scan of your backups with an antivirus or YARA engine to find the latest malware-free restore point, or to detect a presence of specific data, such as personal information.

Scan mode:

☒ Find the last clean restore point  
Restore points will be scanned sequentially starting from the most recent one until the first malware-free backup is found. Use this option when a cyber-attack is known to have started recently.

☐ Find the last clean restore point in range  
Restore points will be scanned in the optimal order to identify the last clean backup with least scans possible. Use this option if you are unsure when an attack started or when dealing with a sleeping malware.

☐ Scan content of all restore points in range  
All restore points in range will be scanned sequentially. Use this option for backup content analysis with an applicable YARA rule, for example when looking for personally identifiable information (PII), personal health information (PHI) or payment card industry (PCI) data.

Scan engine:

☐ Scan restore points with an antivirus engine

☒ Scan restore points with the following YARA rule:

FindFileByHash.yara

▼

[Copy YARA rules location to clipboard](#)

Scan range:

From:	To:
<input checked="" type="radio"/> Most recent restore point	<input checked="" type="radio"/> Oldest available restore point
<input type="radio"/> Start date: <input type="text"/>	<input type="radio"/> End date: <input type="text"/>

☒ Continue scanning all remaining files after the first occurrence

Hide Scan Range

OK

Cancel

# Restoring Data from Veeam Agent Backups

You can perform the following restore operations:

- [Restore Veeam Agent backups to VMware vSphere VMs](#)
- [Restore Veeam Agent backups to Hyper-V VMs](#)
- [Restore Veeam Agent backups to Nutanix AHV VMs](#)
- [Restore Veeam Agent backups to Proxmox VE VMs](#)
- [Restore data from Veeam Agent backups to Microsoft Azure](#)
- [Restore data from Veeam Agent backups to Amazon EC2](#)
- [Restore data from Veeam Agent backups to Google Compute Engine](#)
- [Restore computer volumes from Veeam Agent backups](#)
- [Restore individual files and folders from Veeam Agent backups](#)
- [Restore application items from Veeam Agent backups with Veeam Explorers](#)
- [Export computer disks as VMDK, VHD or VHDX disks](#)
- [Publish disks to analyze backup content](#)
- [Export restore points of Veeam Agent backups to standalone full backup files](#)

# Restoring Veeam Agent Backup to vSphere VM

In the Veeam Backup & Replication console, you can use Instant Recovery to restore a Veeam Agent computer as a VMware vSphere VM in your virtualization environment.

A restored VMware vSphere VM will have the same settings as the backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves the settings of the Veeam Agent computer from the backup and applies them to the target VM. These settings include:

- Amount of RAM.
- Number of CPU cores.
- Number of network adapters.
- Network adapter settings.
- BIOS UUID.

If you do not want to preserve the backed-up machine UUID for a VMware vSphere VM, you can create a new UUID during the Instant Recovery configuration process.

- Number of disks and volumes.
- Size of volumes.

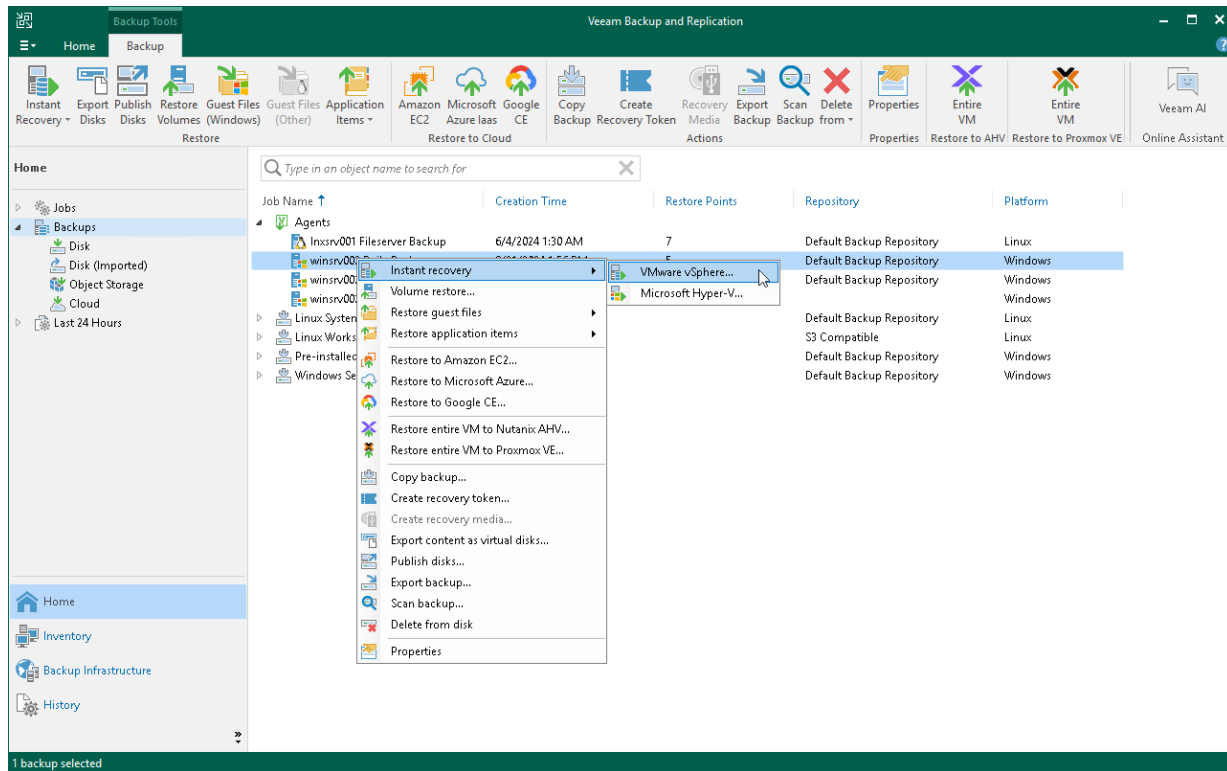
## Considerations and Limitations

If you restore a Veeam Agent computer to a VMware vSphere VM, consider the following:

- You can use entire machine or volume-level backups of Microsoft Windows computers. Volume-level backups must include the computer system drive.
- You can use backups of Microsoft Windows computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.
- If you restore a workload to the production network, make sure that the original workload is powered off.
-

# Restore to vSphere VM

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the [Performing Instant Recovery of Workloads to VMware vSphere](#) section in the Veeam Backup & Replication User Guide.



# Restoring Veeam Agent Backup to Hyper-V VM

In the Veeam Backup & Replication console, you can use Instant Recovery to restore a Veeam Agent computer as a Hyper-V VM in your virtualization environment.

A restored Hyper-V VM will have the same settings as the backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves settings of the Veeam Agent computer from the backup and applies them to the target VM.

## Considerations and Limitations

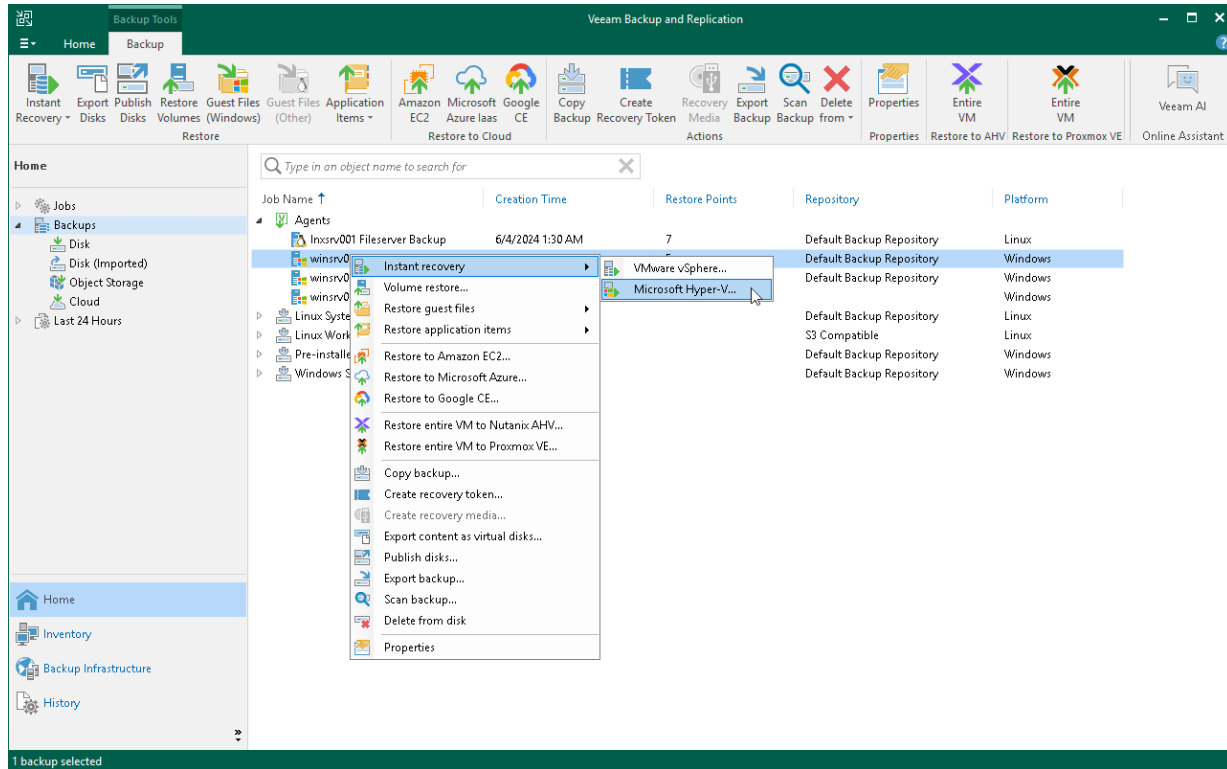
If you restore a Veeam Agent computer to a Hyper-V VM, consider the following:

- You can use backups of Microsoft Windows computers stored in a Veeam backup repository only. You cannot use backups stored in a Veeam Cloud Connect repository for this operation.
- You cannot recover an EFI-based Veeam Agent computer that runs Windows 7, Windows Server 2008 or Windows Server 2008 R2 to a Hyper-V VM. These OSes can be restored only to a Generation 1 VM that does not support EFI. To learn more, see [this Microsoft article](#).
- Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.
- Veeam Agent computer disks are recovered as dynamically expanding virtual disks.
- By default, Veeam Backup & Replication automatically powers on a VM after restore. If you do not want to power on a VM after restore, you can change this setting during the Instant Recovery configuration process.
-



# Restore to Hyper-V VM

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the [Performing Instant Recovery of Workloads to Hyper-V](#) section in the Veeam Backup & Replication User Guide.



# Restoring Veeam Agent Backup to Nutanix VM

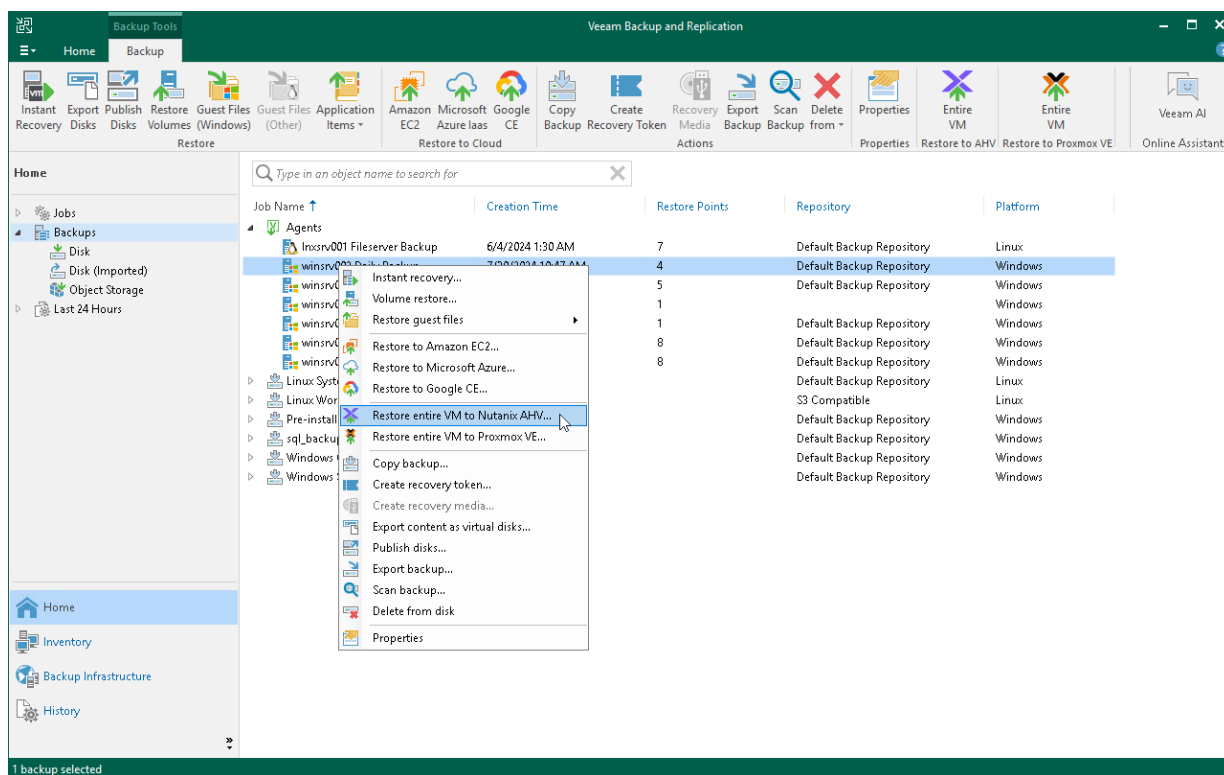
You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a Nutanix AHV VM in your virtualization environment.

## Considerations and Limitations

If you restore a Veeam Agent computer to a Nutanix AHV VM, keep in mind that you can use backups of Microsoft Windows computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.

## Restore to Nutanix AHV

The procedure of restore to Nutanix AHV for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more about restore to Nutanix AHV, see the [Restoring VMs Using Veeam Backup & Replication Console](#) section in the Veeam Backup for Nutanix AHV User Guide.



# Restoring Veeam Agent Backup to Proxmox VM

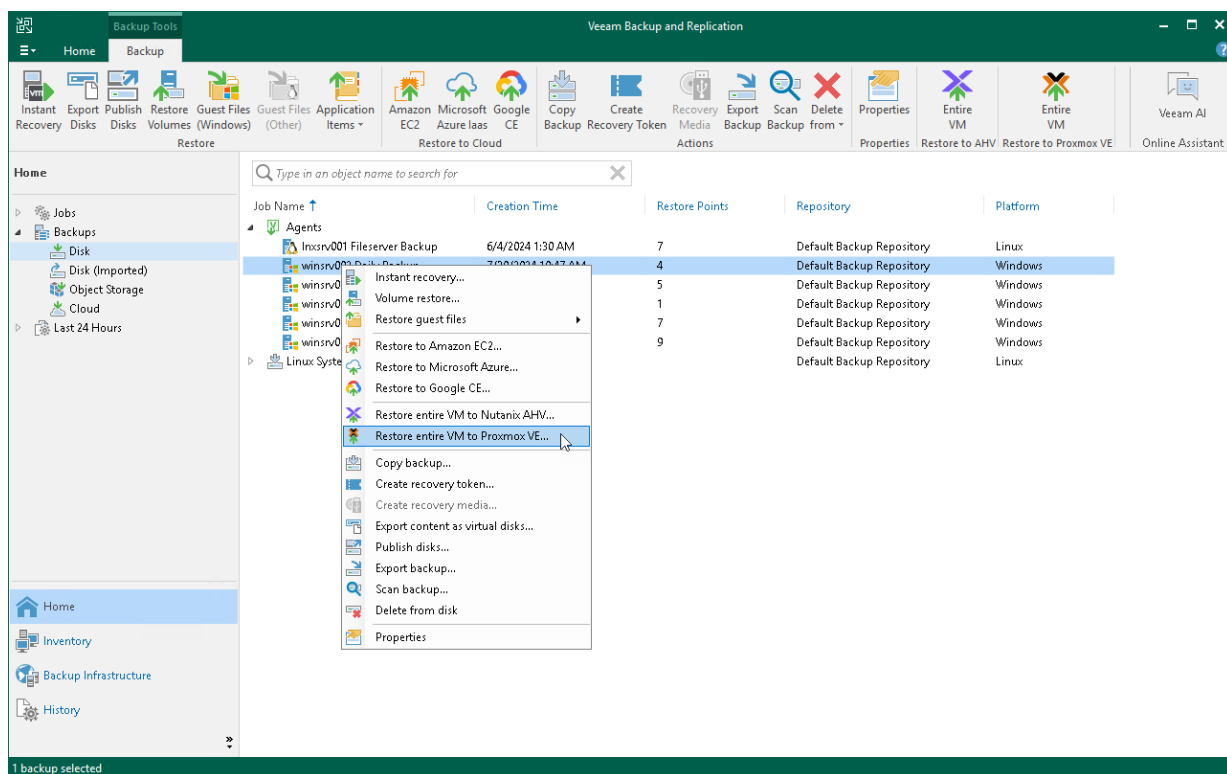
You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a Proxmox VE VM in your virtualization environment.

## Considerations and Limitations

If you restore a Veeam Agent computer to a Proxmox VE VM, keep in mind that you can use backups of Microsoft Windows computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.

## Restore to Proxmox VE

The procedure of restore to Proxmox VE for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more about restore to Proxmox VE, see the [Performing VM Restore](#) section in the Veeam Backup for Proxmox VE User Guide.



# Restoring to Microsoft Azure

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Microsoft Azure.

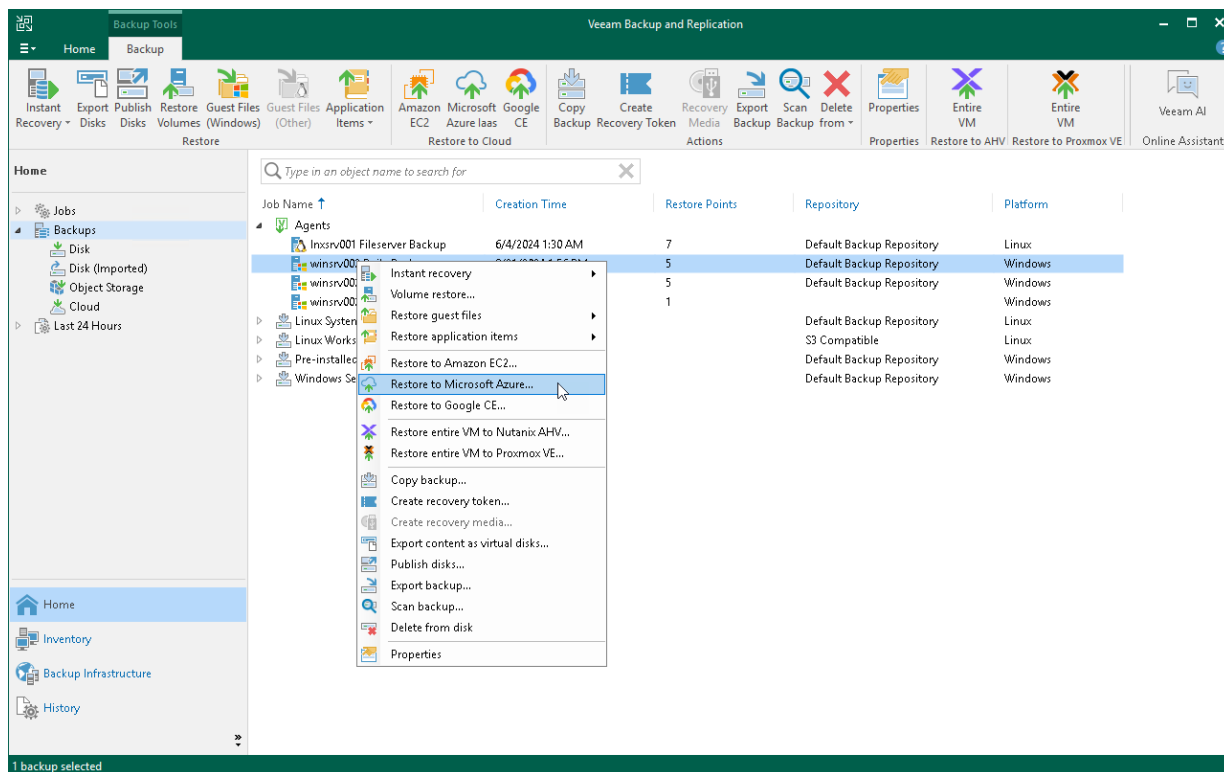
## Considerations and Limitations

If you restore a Veeam Agent computer to Microsoft Azure, consider the following:

- You can use backups of Microsoft Windows computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level.
- If you recover an EFI-based system to Microsoft Azure, Veeam Agent will restore a BIOS-based Generation 1 VM.
- Veeam Backup & Replication offers experimental support for generation 2 VMs within restore to Microsoft Azure feature. To learn more, see the [Generation 2 VM Support](#) section in the Veeam Backup & Replication User Guide.

## Restore to Microsoft Azure

The procedure of restore to Microsoft Azure from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Microsoft Azure, see the [Restoring to Microsoft Azure](#) section in the Veeam Backup & Replication User Guide.



# Restoring to Amazon EC2

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Amazon EC2.

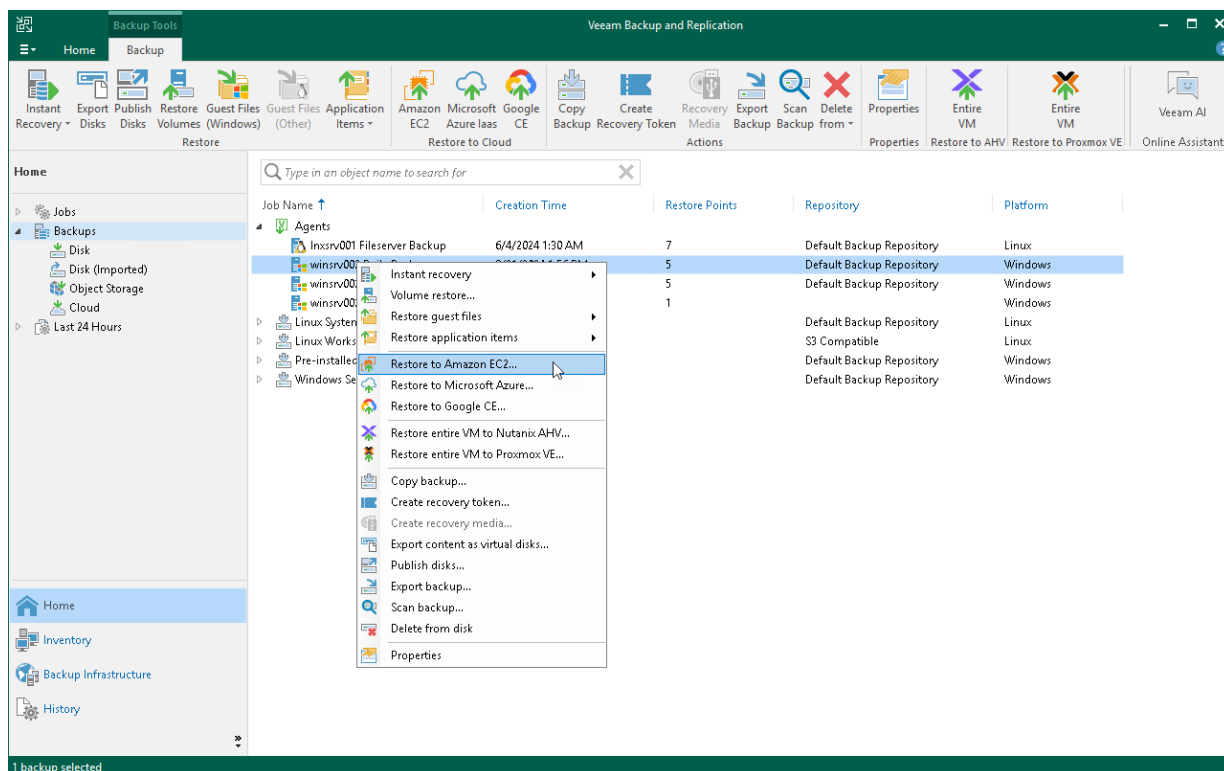
## Considerations and Limitations

If you restore a Veeam Agent computer to Amazon EC2, consider the following:

- You can use backups of Microsoft Windows computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level.

## Restore to Amazon EC2

The procedure of restore to Amazon EC2 from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Amazon EC2, see the [Restoring to Amazon EC2](#) section in the Veeam Backup & Replication User Guide.



# Restoring to Google Compute Engine

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Google Compute Engine.

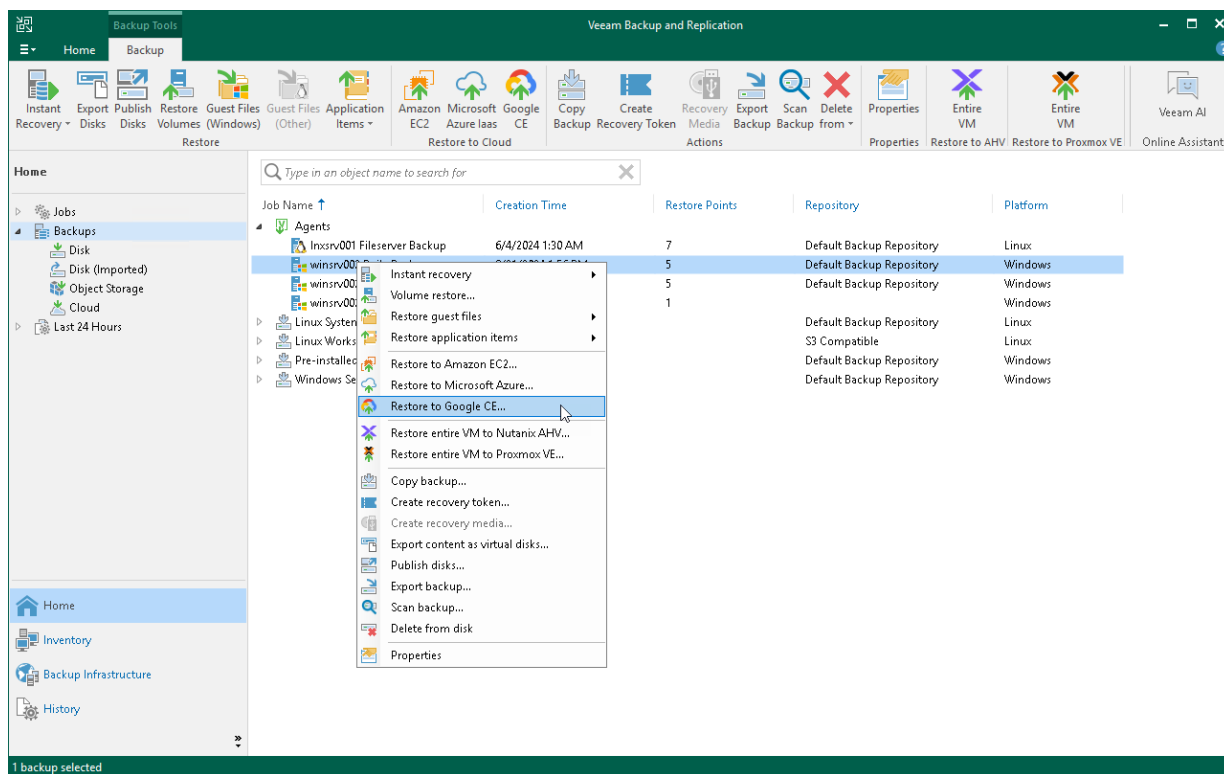
## Considerations and Limitations

If you restore a Veeam Agent computer to Google Compute Engine, consider the following:

- You can use backups of Microsoft Windows computers stored in a Veeam backup repository. You cannot perform this operation with Veeam Agent backups created on the Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level.

## Restore to Google Compute Engine

The procedure of restore to Google Compute Engine from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Google Compute Engine, see the [Restoring to Google Compute Engine](#) section in the Veeam Backup & Replication User Guide.



# Restoring Volumes

You can use Veeam Backup & Replication to restore a specific computer volume or all volumes from a volume-level backup created with Veeam Agent for Microsoft Windows.

If data on a computer volume gets corrupted, you can restore this volume from the backup. For volume-level restore, you can use backups that were created at the volume level. File-level backups cannot be used for volume restore.

When you perform volume-level restore, Veeam Backup & Replication restores the entire content of the volume. It retrieves from the backup data blocks pertaining to a specific volume and copies them to the necessary location. Keep in mind that you cannot browse the volume in the backup and select individual application items, files and folders for restore. For granular file-level restore, you can use the [restore guest files](#) option.

A volume can be restored to its original location or a new location. If you restore the volume to its original location, Veeam Backup & Replication overwrites data on the original volume. If you restore the volume to a new location, and the target disk contains any data, Veeam Backup & Replication overwrites data in the target location with data retrieved from the backup.

A volume can be restored to a new location that has greater or less space than the size of the volume in the backup. Depending on the amount of free disk space on target location, you can select either to shrink or to extend the volume during restore. To learn more, see [Volume Resize](#).

- [Complete the restore process](#).

## Before You Begin

Before you begin the volume-level restore process, check the following prerequisites:

- The volume-level backup from which you plan to restore data must be successfully created at least once.
- A computer on which you want to restore a volume must be added to the Veeam Backup & Replication inventory and run Veeam Agent for Microsoft Windows operating in the standalone mode.

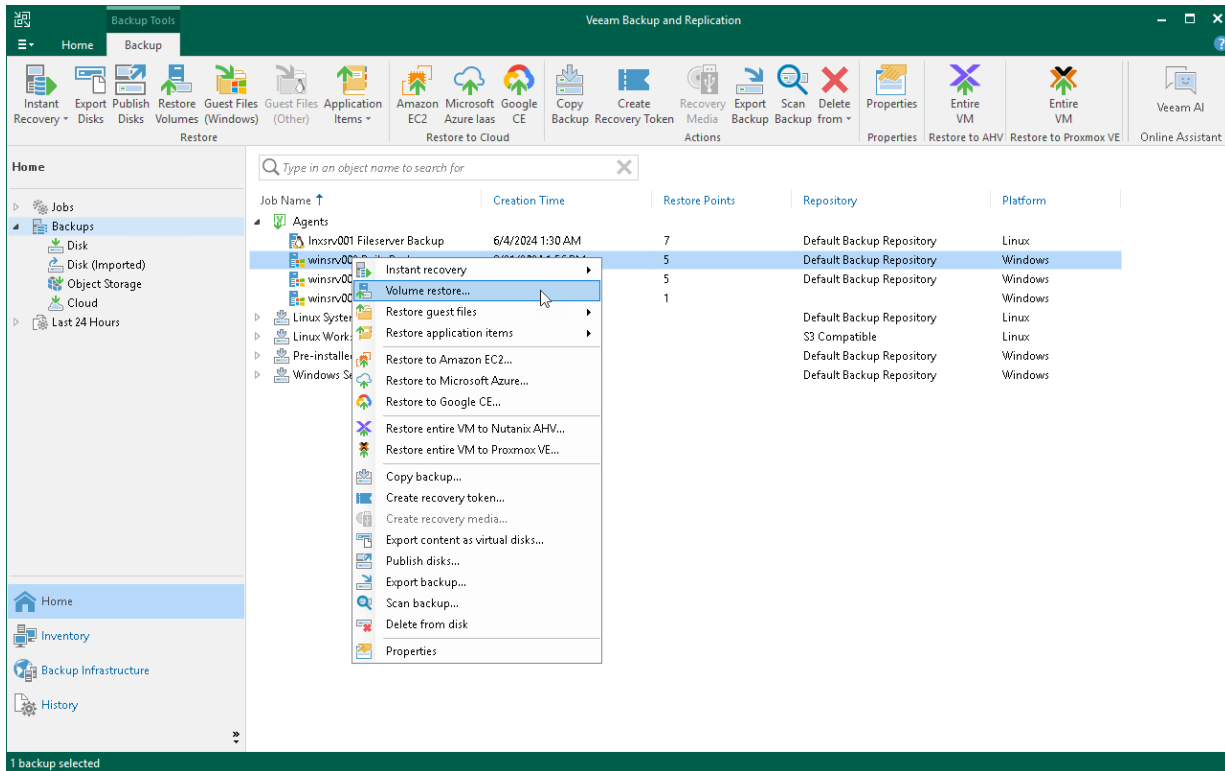
Volume-level restore has the following limitations:

- You cannot restore a system volume to a system volume of the original Veeam Agent computer or another computer with the running OS. To perform such restore, you need to boot the OS from the recovery image. To learn more, see [Restoring Data with Veeam Recovery Media](#). You can also restore a system volume to a non-system volume that has enough free space.
- You cannot restore a volume to a volume on which the Microsoft Windows swap file is hosted.

# Step 1. Launch Volume Level Restore Wizard

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

- Open the **Home** tab and click **Restore > Agent > Disk restore > Volume restore**. In this case, you will be able to select a backup of the necessary Veeam Agent computer at the **Backup** step of the wizard.
- Open the **Home** view. In the inventory pane, click the **Backups** node. In the working area, expand the necessary Veeam Agent backup, select the necessary computer in the backup and click **Restore Volumes** on the ribbon or right-click the computer and select **Volume restore**. In this case, you will proceed immediately to the **Restore Point** step of the wizard.





# Step 2. Select Backup

At the **Backup** step of the wizard, select a backup from which you want to recover data.

To quickly find the necessary backup, use the search field at the bottom of the window: enter a backup name or a part of it in the search field and click the **Start search** button on the right or press **[ENTER]**.

In the list of backups, Veeam Backup & Replication displays only volume-level backups created with Veeam Agent for Microsoft Windows. File-level backups are not displayed.

Volume Restore

Backup

Select a Veeam Agent backup to restore volumes from.

Backup

Restore Point

Disk Mapping

Secure Restore

Reason

Summary

Computer: winsrv002

Job name	Last restore point	Objects	Restore points
DB Backup	12/28/2022 4:48:25 PM	1	
Server Backup	12/30/2022 10:48:53 ...	1	
Windows Servers B...	1/11/2023 6:18:13 PM	1	
winsrv002 System...	7 days ago (6:18 PM ...	9	
filesrv03.tech.io...	1/12/2023 3:37:54 PM	1	8

Type in an object name to search for

< Previous

Next >

Finish

Cancel

## Step 3. Select Restore Point

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

By default, Veeam Backup & Replication uses the latest restore point. However, you can select any valid restore point to recover volumes to a specific point in time.

Veeam Backup & Replication displays restore points for volume-level backups only. For example, if you have run 3 job sessions to create a backup of all computer volumes and then changed the backup scope to file-level backup, Veeam Backup & Replication will display only 3 restore points in the list.

**Volume Restore**

**Restore Point**  
Select the desired restore point.

Backup

**Restore Point**

Disk Mapping

Secure Restore

Reason

Summary

Computer name: **winsrv002**

Data size: **98.6 GB**

Available restore points:

Created	Type
19 days ago (8:49 AM Friday 12/30/2022)	Increment
20 days ago (3:57 PM Thursday 12/29/2022)	Increment
21 days ago (2:48 PM Wednesday 12/28/2022)	Full

< Previous   **Next >**   Finish   Cancel

## Step 4. Map Restored Disks

At the **Disk Mapping** step of the wizard, select what volumes you want to restore and map volumes from the backup to volumes on the target computer.

### IMPORTANT

It is strongly recommended that you change disk mapping settings only if you have experience in working with Microsoft Windows disks and partitions. If you make a mistake, your computer data may get corrupted.

To select volumes for restore:

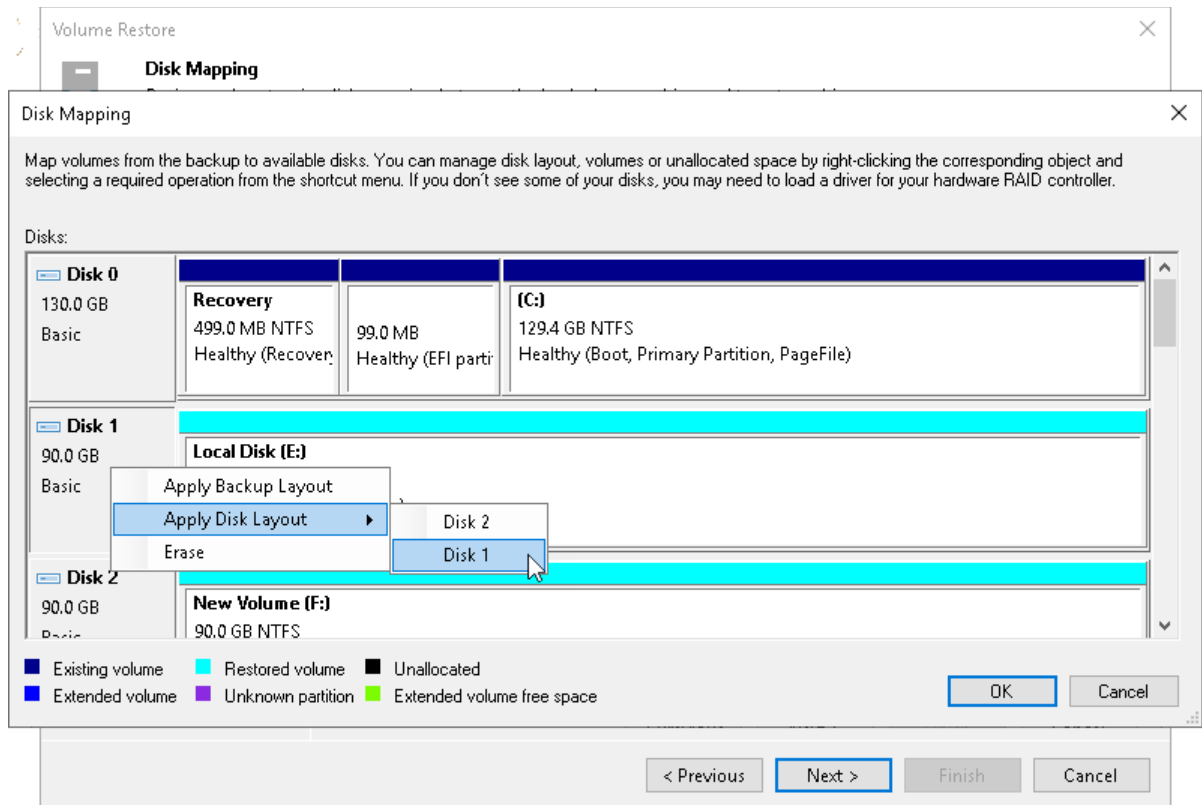
1. In the **Destination host** field, specify the target computer where you want to restore volumes. Click **Choose** and select the necessary computer. You can restore volumes only to computers that are added to the Veeam Backup & Replication inventory and run Veeam Agent for Microsoft Windows.
2. In the **Disk mapping** section, select check boxes next to volumes that you want to restore from the backup. By default, Veeam Backup & Replication restores volumes to their initial location and maps the restored volumes automatically. If the initial location is unavailable, a volume is restored to a disk of the same or larger size. If you want to map the restored volume to another computer disk, at the bottom of the wizard click **Customize disk mapping**.

### NOTE

If Veeam Backup & Replication cannot map a volume automatically, Veeam Backup & Replication will prompt you to perform disk mapping manually. To proceed to the **Disk Mapping** window, click **Yes**.

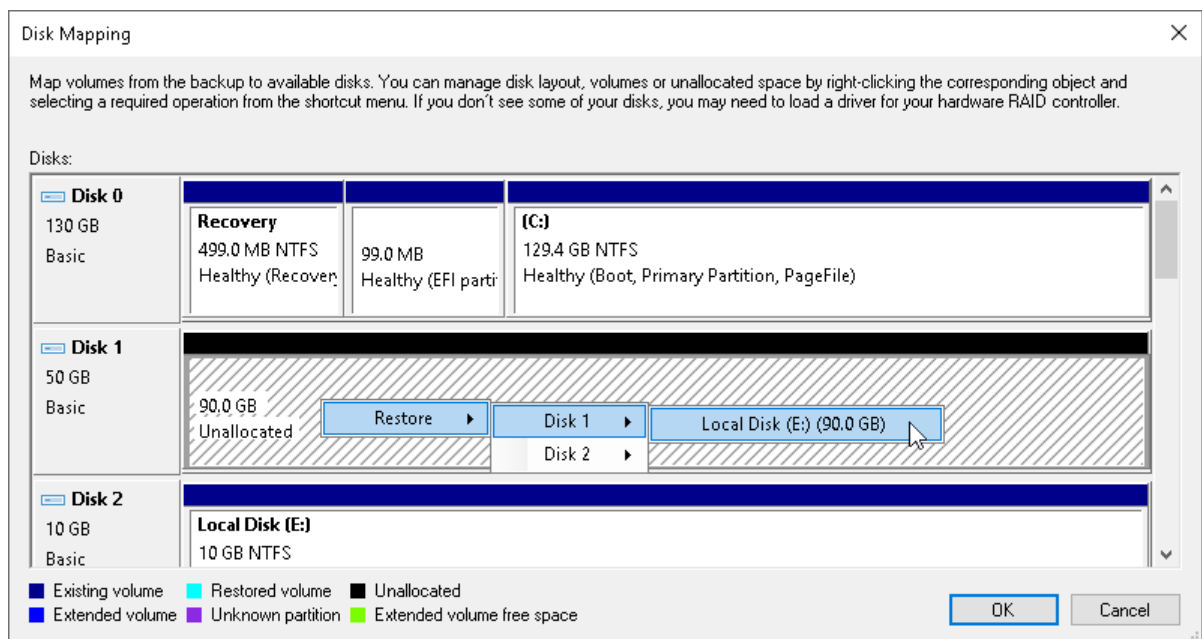
3. In the **Disk Mapping** window, specify how volumes must be restored:
  - Right-click the target disk on the left and select the necessary disk layout:
    - **Apply Backup Layout** – select this option if you want to apply to disk the settings that were used on your computer at the moment when you performed backup.
    - **Apply Disk Layout** – select this option if you want to apply to the current disk settings of another disk.

- **Erase** – select this option if you want to discard the current disk settings.



- Right-click unallocated disk space in the disk area on the right and select what volume from the backup you want to place on this computer disk.

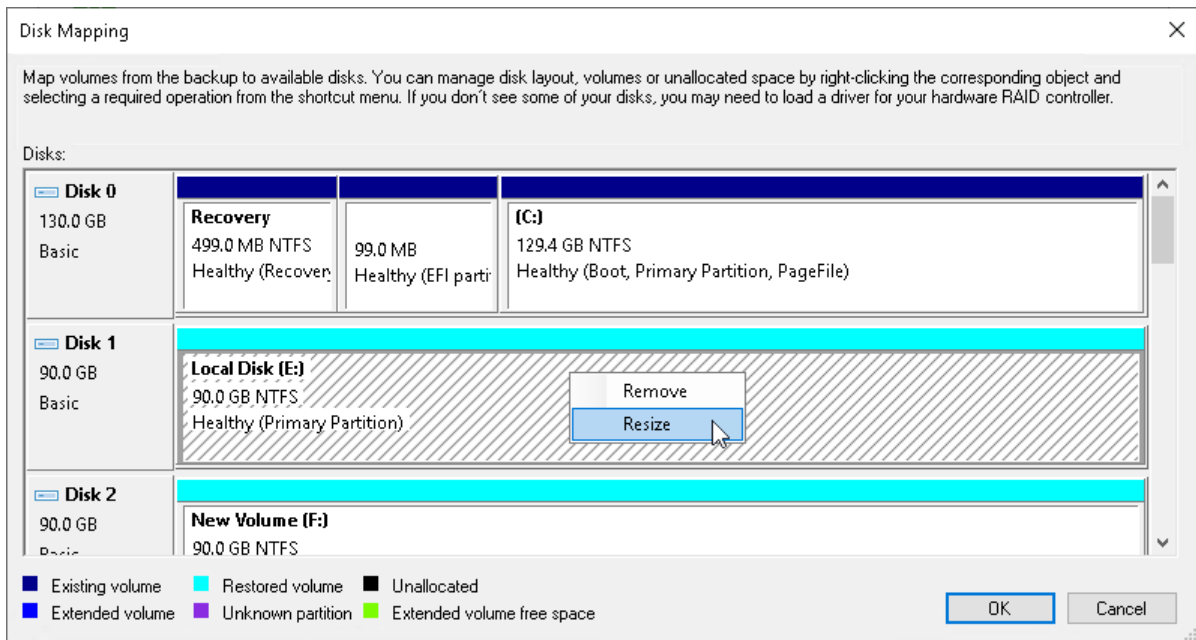
If you want to change disk layout configured by Veeam Backup & Replication, right-click an automatically mapped volume and select **Remove**. You will be able to use the released space for mapping volumes in your own order.



4. [For restore with volume resize] You can resize a volume mapped by Veeam Backup & Replication to a target computer disk. To resize a volume, right-click it in the **Disk Mapping** window and select **Resize**. With this option selected, you will pass to the [Volume Resize](#) window.

## NOTE

If you map a backup volume that is larger than the amount of available space on the target disk, Veeam Backup & Replication will prompt you to shrink the restored volume. After you agree and click **OK**, Veeam Backup & Replication will prepare to shrink the volume to the size of available disk space.



## Step 5. Resize Restored Volumes

At the **Disk Mapping** step of the wizard you can set the necessary size for the restored volumes. A volume will be shrunk or extended to the specified size during the process of data restore.


### NOTE

By default, Veeam Agent for Microsoft Windows displays volume size in megabytes (MB). This allows you to specify the desired size for the volume precisely. You can also choose to display volume size in gigabytes (GB). This may be helpful when you need to resize volumes on larger computer disks and want to simplify disk size calculations.

When you use GB as a volume size unit, you can specify volume size with integral numbers, for example, 1 GB, 60 GB or 200 GB, but not 0,8 GB, 60,5 GB or 200,7 GB. However, if the maximum volume size is in fact greater than the displayed value for less than 1 GB, Veeam Agent for Microsoft Windows will automatically add the exceeding amount of disk space to the extended volume. For example, if the maximum volume size is 60,2 GB, Veeam Agent for Microsoft Windows will display this size as 60 GB. When you specify 60 GB as a desired volume size, Veeam Agent for Microsoft Windows will extend the volume to 60,2 GB.

To resize a volume:

1. Specify a volume you want to resize:
  - a. Right-click a restored volume mapped to a target disk and select **Resize**.
  - b. [For volume shrink] Right-click unallocated disk space and select what volume from the backup you want to place on the computer disk. If the selected volume is larger than the amount of unallocated disk space, Veeam Backup & Replication will prompt you to shrink the restored volume.
2. In the **Volume Resize** window, select the volume size unit and specify the desired size for the restored volume.

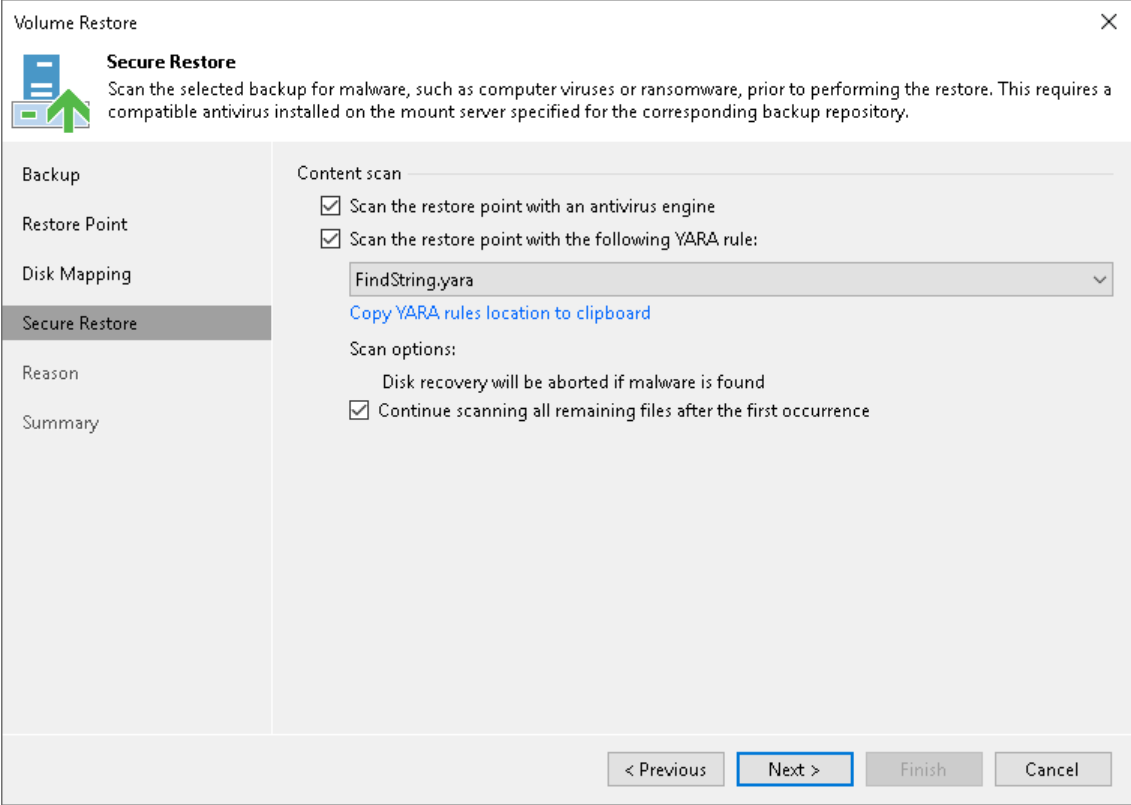


## Step 6. Specify Secure Restore Settings

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to perform secure restore. To learn more about secure restore, see the [Secure Restore](#) section in the Veeam Backup & Replication User Guide.

To specify secure restore settings:

1. In the **Content scan** section, specify the following:
  - a. If you want to scan the restored volume with antivirus software, select the **Scan the restore point with an antivirus engine** check box. For more information on antivirus scan, see the [Antivirus Scan \(Secure Restore\)](#) section in the Veeam Backup & Replication.
  - b. If you want to scan the restored volume with a YARA rule, select the **Scan backup content with the following YARA** rule check box and select a YARA rule from the drop-down list. By default, the YARA rules are located in the folder by the following path: C:\Program Files\Veeam\Backup and Replication\Backup\YaraRules.
2. In the **Scan options** section, select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue volume scan after the first malware threat is found. For information on how to view results of the antivirus scan, see the [Viewing Malware Scan Results](#) section in the Veeam Backup & Replication User Guide.



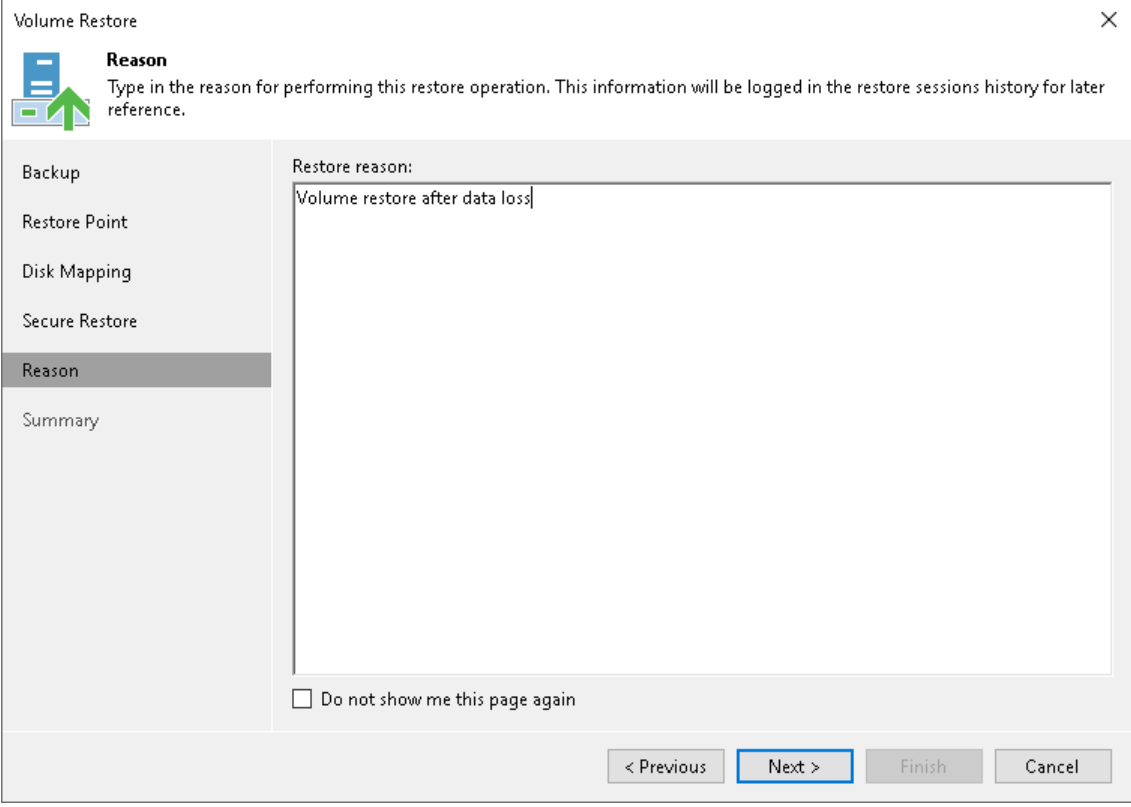
The screenshot shows the 'Volume Restore' wizard window, specifically the 'Secure Restore' step. The window has a title bar with 'Volume Restore' and a close button. On the left is a sidebar with icons and labels: 'Backup', 'Restore Point', 'Disk Mapping', 'Secure Restore' (highlighted), 'Reason', and 'Summary'. The main area is titled 'Secure Restore' and contains the following text: 'Scan the selected backup for malware, such as computer viruses or ransomware, prior to performing the restore. This requires a compatible antivirus installed on the mount server specified for the corresponding backup repository.' Below this, there are two sections: 'Content scan' and 'Scan options'. The 'Content scan' section has two checked checkboxes: 'Scan the restore point with an antivirus engine' and 'Scan the restore point with the following YARA rule:'. Below the second checkbox is a dropdown menu showing 'FindString.yara' and a link 'Copy YARA rules location to clipboard'. The 'Scan options' section has a label 'Disk recovery will be aborted if malware is found' and a checked checkbox 'Continue scanning all remaining files after the first occurrence'. At the bottom right are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 7. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the computer volume.

### TIP

If you do not want to display the **Restore Reason** step of the wizard in future, select the **Do not show me this page again** check box.



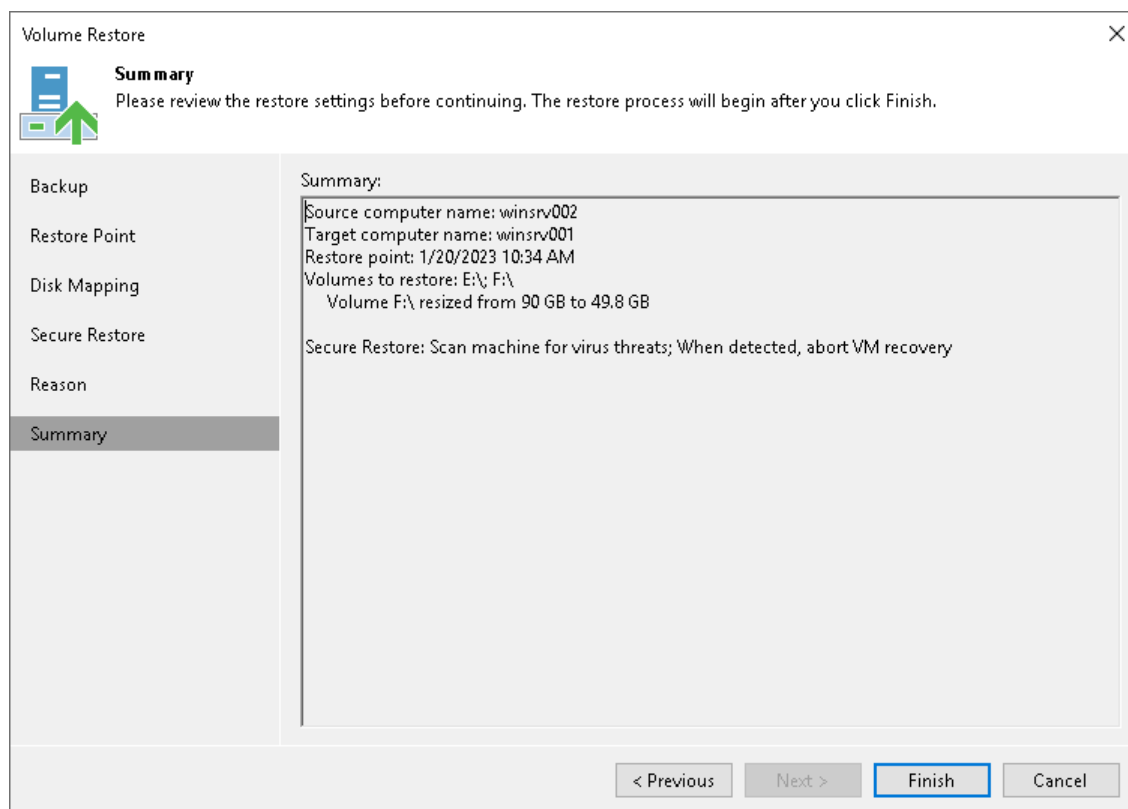
The screenshot shows the 'Volume Restore' wizard window. The title bar reads 'Volume Restore' with a close button. On the left is a sidebar with icons and labels for 'Backup', 'Restore Point', 'Disk Mapping', 'Secure Restore', 'Reason' (which is highlighted), and 'Summary'. The main area has a header 'Reason' with a sub-header 'Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.' Below this is a large text box labeled 'Restore reason:' containing the text 'Volume restore after data loss'. At the bottom of the main area is a checkbox labeled 'Do not show me this page again'. The bottom of the window contains four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.



## Step 8. Complete Restore Process

At the **Summary** step of the wizard, complete the procedure of volume-level restore.

1. Review settings of the restore process.
2. Click **Finish** to start the recovery process. Veeam Backup & Replication will perform partition re-allocation operations if necessary, restore the necessary volume data from the backup and overwrite volume data on the target computer with the restored data.



# Restoring Files and Folders

You can use the Veeam Backup & Replication console to restore individual files and folders from Veeam Agent backups.

The procedure of file-level restore from a Veeam Agent backup is similar to the same procedure for a VM backup. To learn more about file-level restore, see the [Restore from Linux, Unix and Other File Systems](#) section in the Veeam Backup & Replication User Guide.

## NOTE

Before you start file-level restore from a backup of a failover cluster, make sure that the cluster is added to a protection group in the Veeam Backup & Replication inventory. The failover cluster may be not present in the inventory, for example, in the following cases:

- The original protection group that contained the cluster was removed from Veeam Backup & Replication.
- You want to restore cluster data from a backup created on another backup server and imported in the Veeam backup console.

In this case, add the failover cluster whose data you want to restore to a protection group.

Before you perform file-level restore, [check prerequisites](#). Then use the **File Level Restore** wizard to restore the necessary files and folders.

1. [Launch the File Level Restore wizard](#).
2. [Select a machine](#).
3. [Select a restore point](#).
4. [Specify a restore reason](#).
5. [Verify restore settings](#).
6. [Finalize restore](#).

## Before You Begin

Before you begin the file-level restore process, check the following requirements and limitations:

## Requirements

Consider the following general requirements:

- You can restore files from the file systems listed in the [Guest OS File Restore](#) section in the Veeam Backup & Replication User Guide.
- The account that you use to start the Veeam Backup & Replication console and to connect to the backup server must have permissions and privileges described in the [Installing and Using Veeam Backup & Replication](#) section in the Veeam Backup & Replication User Guide.
- You can restore files from a backup that has at least one successfully created restore point.
- [For [restore to original location](#)] The mount server must have access to the Veeam Agent computer.

- [For [restore changes functionality](#)] This functionality is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise or Enterprise Plus editions of Veeam Backup & Replication are required.

## Requirements for ReFS

### NOTE

The information in this subsection is only applicable to restore from volume-level backups.

If you plan to restore files from a computer running Microsoft Windows ReFS, consider the following requirements for the Veeam Backup & Replication components involved in the restore process:

- The machine on which a mount point is created (for example, the mount server) must run Microsoft Windows Server 2012 OS or later.
- [For ReFS 3.x] If you plan to restore files from a computer running Microsoft Windows ReFS 3.x, the machine on which mount point is created must run Microsoft Windows Server 2016 OS or later and the ReFS version must be supported on it.
- The machine on which a mount point is created must run Microsoft Windows Server OS of the same version or newer than the Veeam Agent computer from which you plan to restore files.

To learn more about mount points creation, see the [Mount Points and Restore Scenarios](#) section in the Veeam Backup & Replication User Guide.

## Requirements for Data Deduplication

### NOTE

The information in this subsection is only applicable to restore from volume-level backups.

If you plan to restore files from a computer running Microsoft Windows Server 2012 OS or later and data deduplication is enabled for some volumes, consider the following for the Veeam Backup & Replication components involved in the restore process:

- The machine on which a mount point is created (for example, the mount server) must run Microsoft Windows Server 2012 OS or later.
- The machine on which a mount point is created must run Microsoft Windows Server OS of the same version or newer than the Veeam Agent computer from which you plan to restore files.
- Data deduplication must be enabled on the machine on which a mount point is created.

To learn in which scenarios on which machines mount points can be created, see the [Mount Points and Restore Scenarios](#) section in the Veeam Backup & Replication User Guide.

## Limitations

Consider the following limitations:

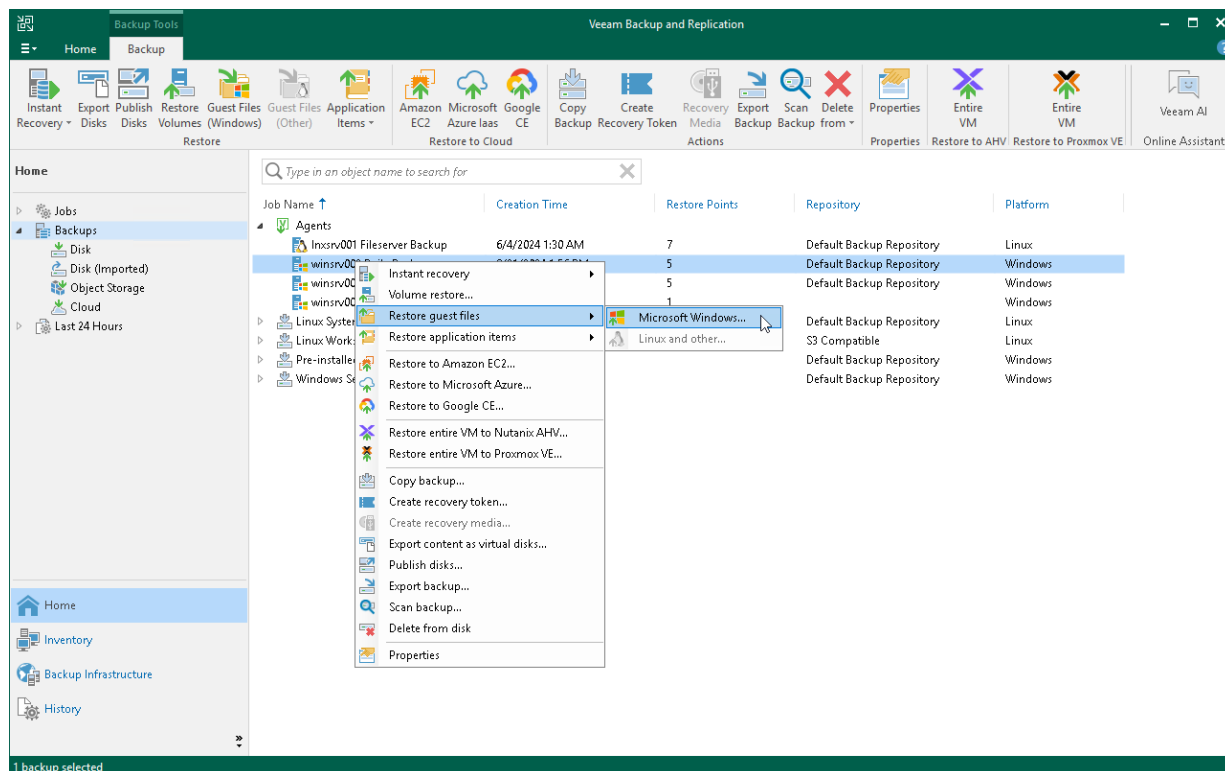
- You cannot restore pipes and other file system objects. File-level restore supports recovery of files and folders only.
- You can restore files from basic disks and dynamic disks (including simple, mirrored and striped volumes).

- [For restore from volume-level backups] Processing of reparse points is supported only for NTFS. Note that reparse points with reparse tag values other than `IO_REPARSE_TAG_MOUNT_POINT`, `IO_REPARSE_TAG_SYMLINK` and `IO_REPARSE_TAG_DEDUP` may be processed and restored incorrectly.
- The comparison functionality uses Veeam Deployer Service. This service is a 32-bit service. During the comparison, the service converts some 64-bit objects in 32-bit objects. That is why such objects are shown as deleted in the Veeam Backup browser, for example, some objects in the `Windows` folder.
- [For [permission restore](#)] Veeam Backup & Replication restores only permissions. Attributes such as Read-only, Encrypted and so on are not restored.
- [For [permission restore](#)] Permissions can be restored only for files and folders that are still present on the Veeam Agent computer. If files and folders are missing, restore fails.
- You cannot restore files encrypted with Windows EFS.

# Step 1. Launch File Level Restore Wizard

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

- Open the **Home** tab and click **Restore > Agent > Guest files restore > Microsoft Windows**. In this case, you will be able to select a Veeam Agent computer whose files and folders you want to restore at the [Machine](#) step of the wizard.
- Open the **Home** view. In the inventory pane, click the **Backups** node. In the working area, expand the necessary Veeam Agent backup, select the necessary computer in the backup and click **Guest Files (Windows)** on the ribbon. Alternatively, right-click the computer and select **Restore guest files > Microsoft Windows**. In this case, you will proceed immediately to the [Restore Point](#) step of the wizard.



# Step 2. Select Machine

At the **Machine** step of the wizard, select a backup from which you want to recover data.

To quickly find the necessary backup, use the search field at the bottom of the window: enter a Veeam Agent computer name or a part of it in the search field and click the **Start search** button on the right or press [Enter].

File Level Restore

Machine

Choose the machine you would like to restore.

Machine

Restore Point

Reason

Summary

Machine: winsrv002

Job name	Last restore point	Objects	Restore points
▶ Daily Backup	12/28/2022 4:48:25 PM	1	
▶ winsrv002 Daily Ba...	99 days ago (2:49 PM...		3
▶ System Backup	12/30/2022 10:48:53 ...	1	
▶ Weekly Backup	12/28/2022 4:49:27 PM	1	

▼

Type in an object name to search for

< Previous

Next >

Browse


Cancel

# Step 3. Select Restore Point

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

By default, Veeam Backup & Replication uses the latest restore point. However, you can select any valid restore point to recover data from.

File Level Restore

**Restore Point**  
Select the restore point to restore guest OS files from.

Machine

Restore Point

Reason

Summary

VM name: winsrv002

Original host: filesrv004.tech.local

VM size: 26.2 GB

Available restore points:

Created	Type	Backup
48 days ago (11:01 PM Thursday 2/16/...	Increment	Daily Backup
49 days ago (11:00 PM Wednesday 2/1...	Increment	Daily Backup
50 days ago (11:00 PM Tuesday 2/14/2...	Full	Daily Backup

< Previous

Next >

Browse

Cancel

## Step 4. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring data.

### TIP

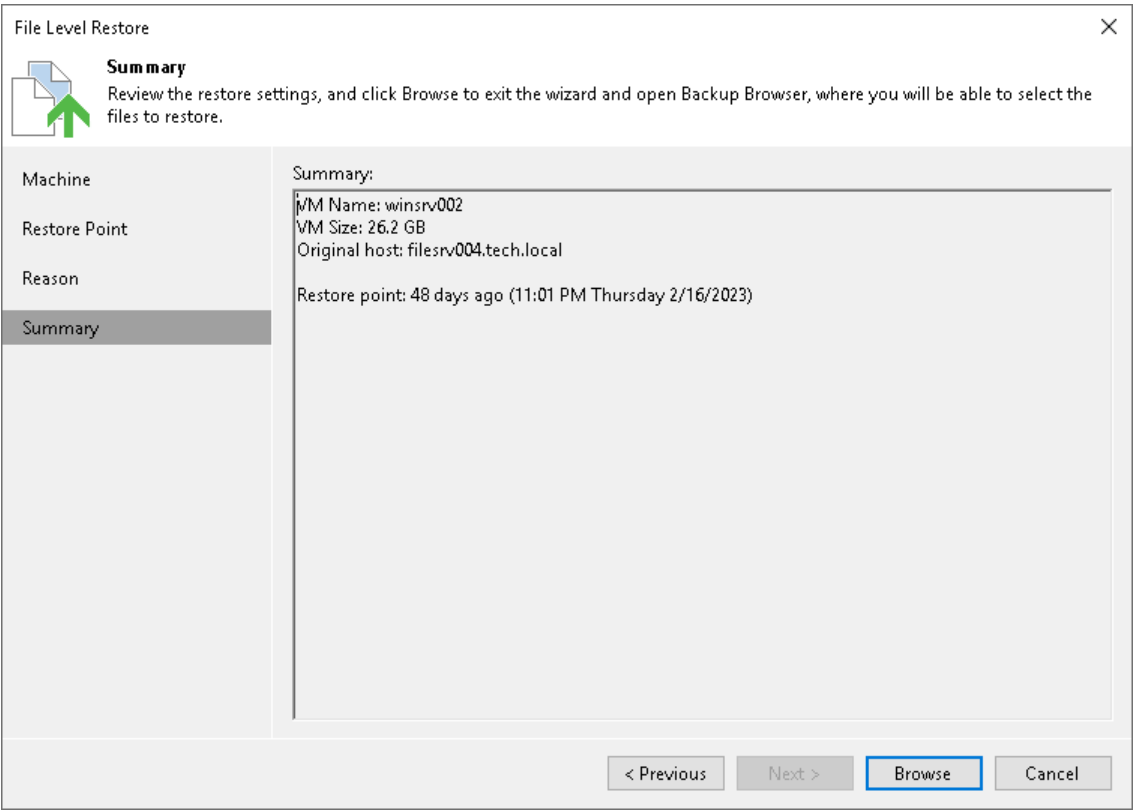
If you do not want to display the **Restore Reason** step of the wizard in future, select the **Do not show me this page again** check box.

The screenshot shows the 'File Level Restore' wizard window. The title bar reads 'File Level Restore'. On the left is a sidebar with four steps: 'Machine', 'Restore Point', 'Reason' (which is highlighted with a dark grey background), and 'Summary'. Above the sidebar, there is a 'Reason' section with a document icon and a green arrow pointing up, and the text: 'Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.' The main area of the wizard is titled 'Restore reason:' and contains a large text box with the text 'Restoring deleted files' entered. Below the text box is a checkbox labeled 'Do not show me this page again'. At the bottom of the window are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Browse', and 'Cancel'.



# Step 5. Verify Restore Settings

At the **Summary** step of the wizard, check details of the restore task and click **Browse** to close the wizard and open the Veeam Backup browser.



## Step 6. Finalize Restore

After you close the wizard, Veeam Backup & Replication opens the Veeam Backup browser with the file system tree of the restored Veeam Agent computer.


You can perform the following operations in the Veeam Backup browser:

- [Compare files and folders from a backup with the original files and folders.](#)
- [Restore only changed files and folders to the original location.](#)
- [Copy files and folders to the Veeam Backup & Replication console or network shared folder.](#)
- [Restore files and folders to the original location.](#)
- [Restore files and folders to another Veeam Agent computer.](#)
- [Restore permissions only.](#)
- [Launch application item restore.](#)
- [Open files in Microsoft Windows File Explorer.](#)

After you finish restoring files, [close the Veeam Backup browser.](#)

### NOTE

Consider the following:

- Folder symbolic links are displayed under the  icon.
- When restoring symbolic links, Veeam Backup & Replication restores only links, not the content they point to.
- Hard links are displayed and restored as files.

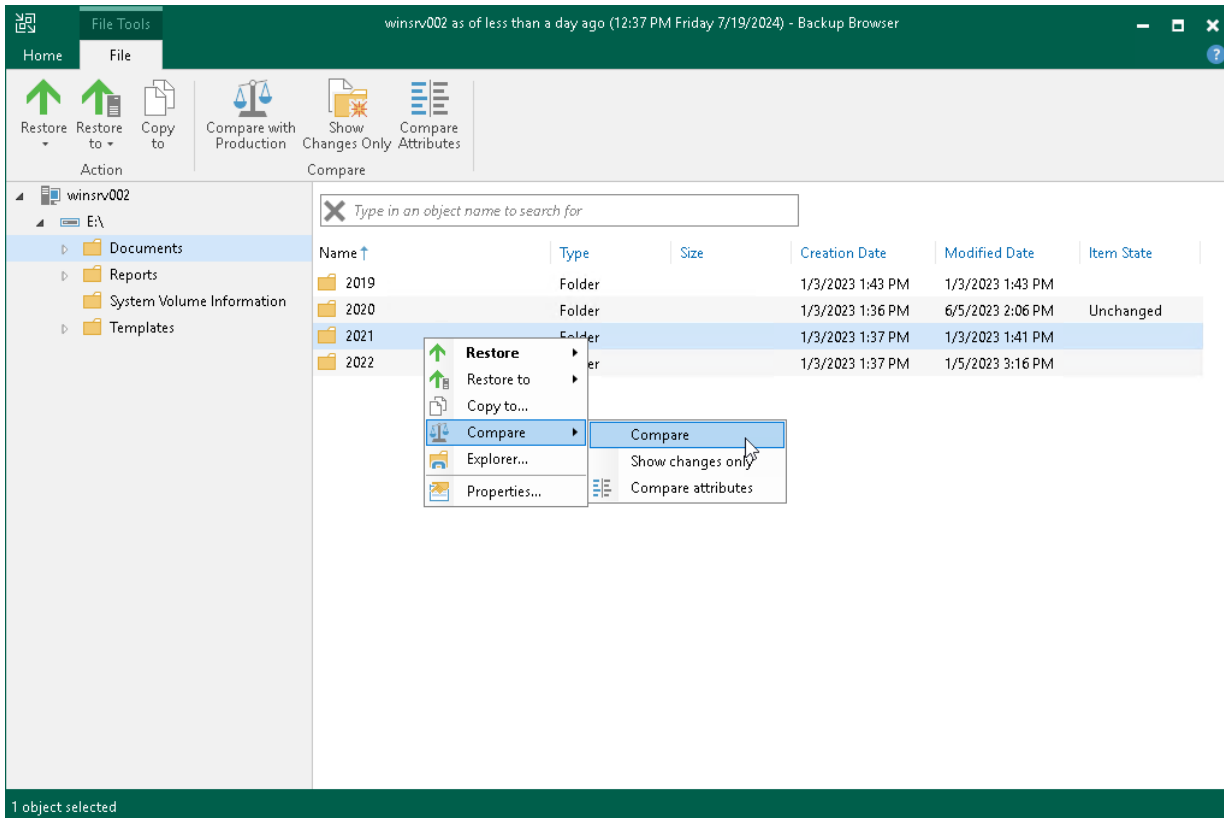
## Comparing Files and Folders

When Veeam Backup & Replication compares files, it compares their attributes. When Veeam Backup & Replication compares folders, it compares folder attributes, and also folder and file attributes inside the folder recursively. Once a file or folder with changed attributes is found, Veeam Backup & Replication stops comparing items and marks the folder as changed. If you further browse the folder in the compared state, Veeam Backup & Replication continues comparing non-compared files and folders. Veeam Backup & Replication compares the following attributes: Date Created, Date Modified, Size (only for files), Read-Only, Hidden, Archive, NTFS Encryption.

To compare files and folders from a backup with the files and folders stored in the original location:

1. Select the necessary files and folders in the file system tree or in the details pane on the right. You can also select disks. In this case, Veeam Backup & Replication will compare files and folders stored on the disks.
2. Right-click one of the selected items and select **Compare > Compare** or click **Compare with Production** on the ribbon.

3. If prompted, in the **Credentials** window, specify user credentials to access the original location.



After the comparison, files and folders will have the following comparison states in the **Item State** column: *changed*, *unchanged*, *deleted*, *pending*, or *failed to compare*. The states are updated when you turn off and then turn on the comparison mode, and when you start restoring changes of files and folders. Note that when comparing symbolic links, Veeam Backup & Replication compares attributes of the links, not the attributes of files and folders which the symbolic link points to.

For files and folders in the comparison states, Veeam Backup & Replication provides other restore operations than for files and folders in the non-comparison state. For example, you can restore only changed files and folders. For more information, see [Restoring Changed Files and Folders](#).

#### TIP

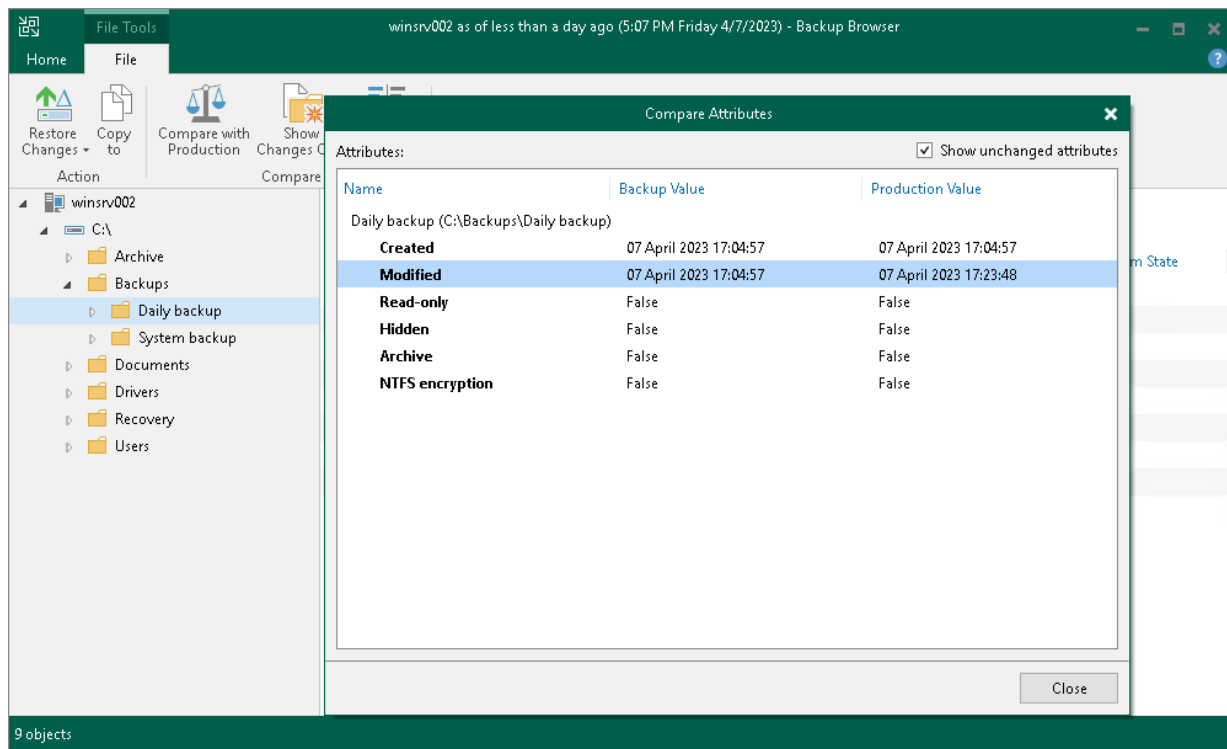
To show only changed files and folders (in the *changed* and *deleted* states), perform the compare operation, right-click any area in the Veeam Backup browser and select **Compare > Show changes only** or click **Show Changes Only** on the ribbon. Note that in this case the search field becomes unavailable. To show all files and folders, click the **Show changes only** option once again.

To switch off the comparison states, select an item in the comparison state and click **Compare > Compare** or click **Compare with Production** on the ribbon. Note that if you switch off comparison for child files and folders, comparison for parent folders will also be switched off.

You can view which attributes were changed for files and folders:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select **Compare > Compare attributes** or click **Compare Attributes** on the ribbon.

In the **Compare Attributes** window, Veeam Backup & Replication shows changed attributes. If you want to show all attributes, click the **Show unchanged attributes** check box at the top right corner. Note that Veeam Backup & Replication shows attributes maximum for 500 files and folders and shows attributes for the selected files and folders, not for the nested files.



## Restoring Changed Files and Folders

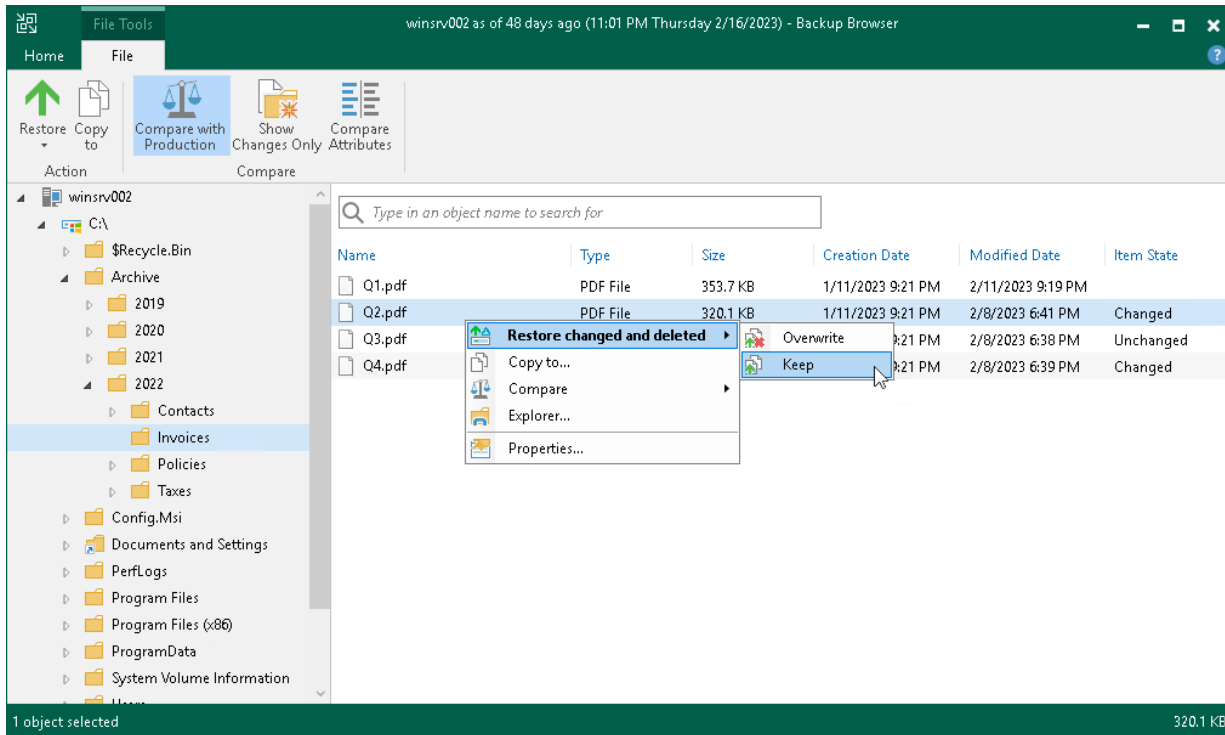
To restore only changed files and folders to the original location, do the following:

1. Select the necessary files and folders in the file system tree or in the details pane on the right. Note that at least one file or folder must be in a comparison state. Files and folders in the non-comparison state, Veeam Backup & Replication will compare automatically.
2. Right-click one of the selected items and select one of the following:
  - To overwrite the original files and folders with the ones restored from the backup, select **Restore changed and deleted > Overwrite**.
  - To save the files and folders restored from the backup next to the original ones, select **Restore changed and deleted > Keep**.

Veeam Backup & Replication will add the *RESTORED\_YYYYMMDD\_HHMMSS* postfix to the original names and store the restored items in the same folder where the original items reside.

Alternatively, you can select the same commands on the ribbon.

If you want to restore entire files and folders to the original location, see [Restoring to Original Location](#).



## Copying Files and Folders to Console or Shared Folder

To copy files and folders to the machine where the Veeam Backup & Replication console is installed or to a network shared folder:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and click **Copy to** or click **Copy to** on the ribbon.
3. In the **Choose Folder** window, select the necessary destination:
  - To recover files and folders to a folder on the machine where the Veeam Backup & Replication console is installed, click **Browse** to find the necessary folder.
  - To recover files and folders to a network shared folder, enter a path to the destination folder in the **Choose folder** field.
4. In the **Choose Folder** window, choose whether to preserve original NTFS permissions or not:
  - To keep the original ownership and security permissions for the restored items, select the **Preserve permissions and ownership** check box.

Veeam Backup & Replication will copy selected files and folders along with associated Access Control Lists, preserving granular access settings.

## IMPORTANT

To preserve permissions and ownership for the selected files and folders, you must run the Veeam Backup & Replication console under the Administrator account.

If you run the Veeam Backup & Replication console under any other account, you will see a warning message after you click **Copy to** at the step 2. In this case, close the **Choose Folder** window and click the **Click here to run as administrator** link in the warning window to elevate access rights.

- If you do not want to preserve the original ownership and access settings for the restored items, leave the **Preserve permissions and ownership** check box not selected.

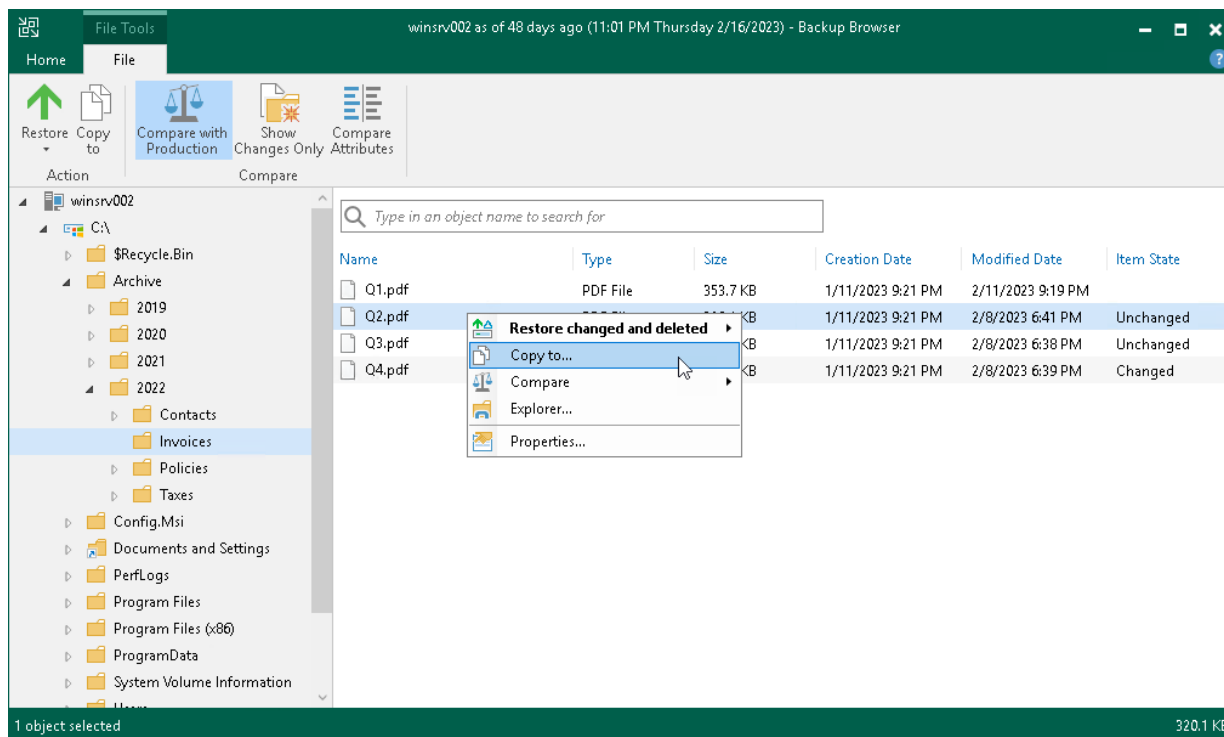
Veeam Backup & Replication will change security settings: the user who launched the Veeam Backup & Replication console will be set as the owner of the restored item, while access permissions will be inherited from the folder to which the restored item is copied.

5. If prompted, in the **Credentials** window, specify user credentials to access the destination location.

## NOTE

Consider the following:

- The copy to operation does not use the comparison states and copies all selected files and folders.
- When copying symbolic links, Veeam Backup & Replication copies the content which the links point to.



## Restoring Files and Folders to Original Location

To restore files and folders to the original location, do the following:

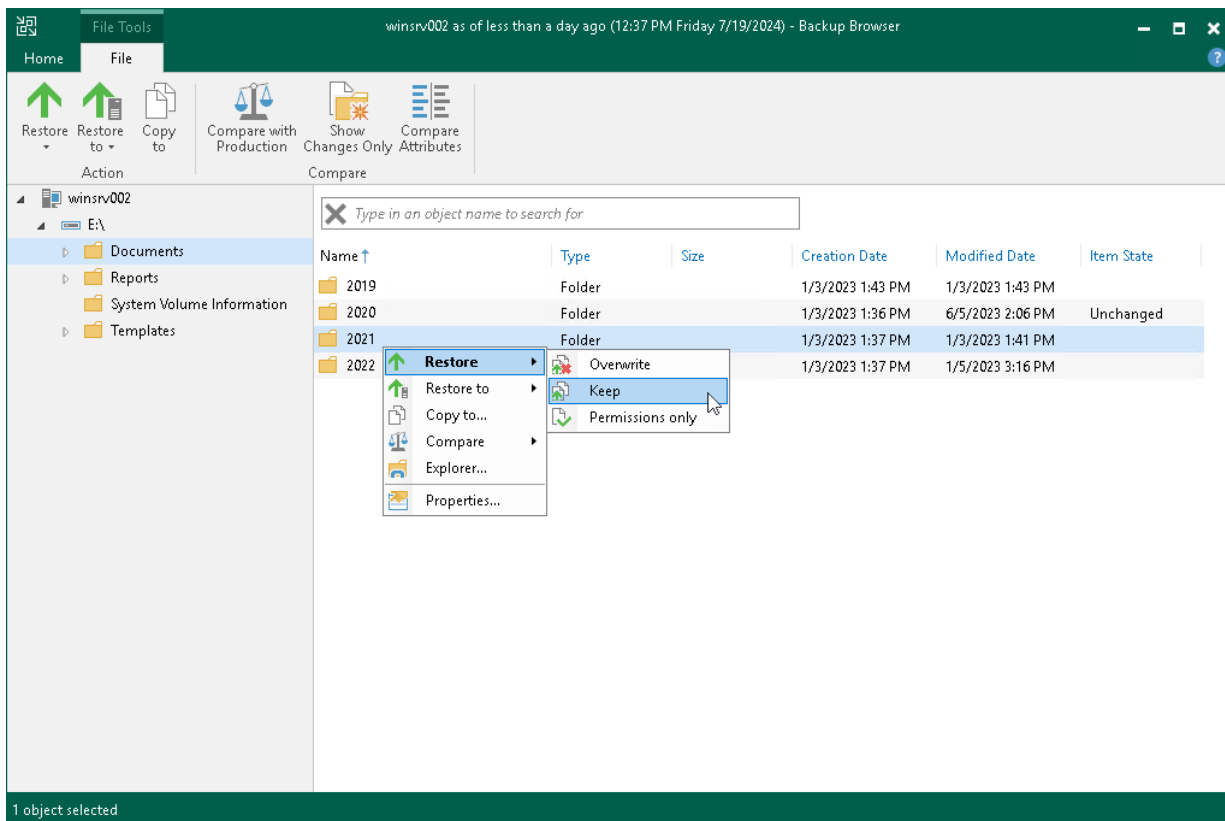
1. Select the necessary files and folders in the file system tree or in the details pane on the right.

2. Right-click one of the selected items and select one of the following:

- To overwrite the original files and folders with the ones restored from the backup, select **Restore > Overwrite**.
- To save the restored files and folders next to the original ones, select **Restore > Keep**.

Veeam Backup & Replication will add the *RESTORED\_YYYYMMDD\_HHMMSS* postfix to the original names and store the restored items in the same folder where the original items reside.

Alternatively, you can select the same commands on the ribbon.



## Restoring Files and Folders to Another Veeam Agent Computer

### NOTE

Consider the following:

- You can restore files and folders only to a computer that is managed by the same Veeam backup server that manages the Veeam backup repository where the backup resides.
- You cannot restore files and folders to cloud machines.
- You cannot restore files and folders that are in the comparison state.

To restore files and folders to another Veeam Agent computer, do the following:

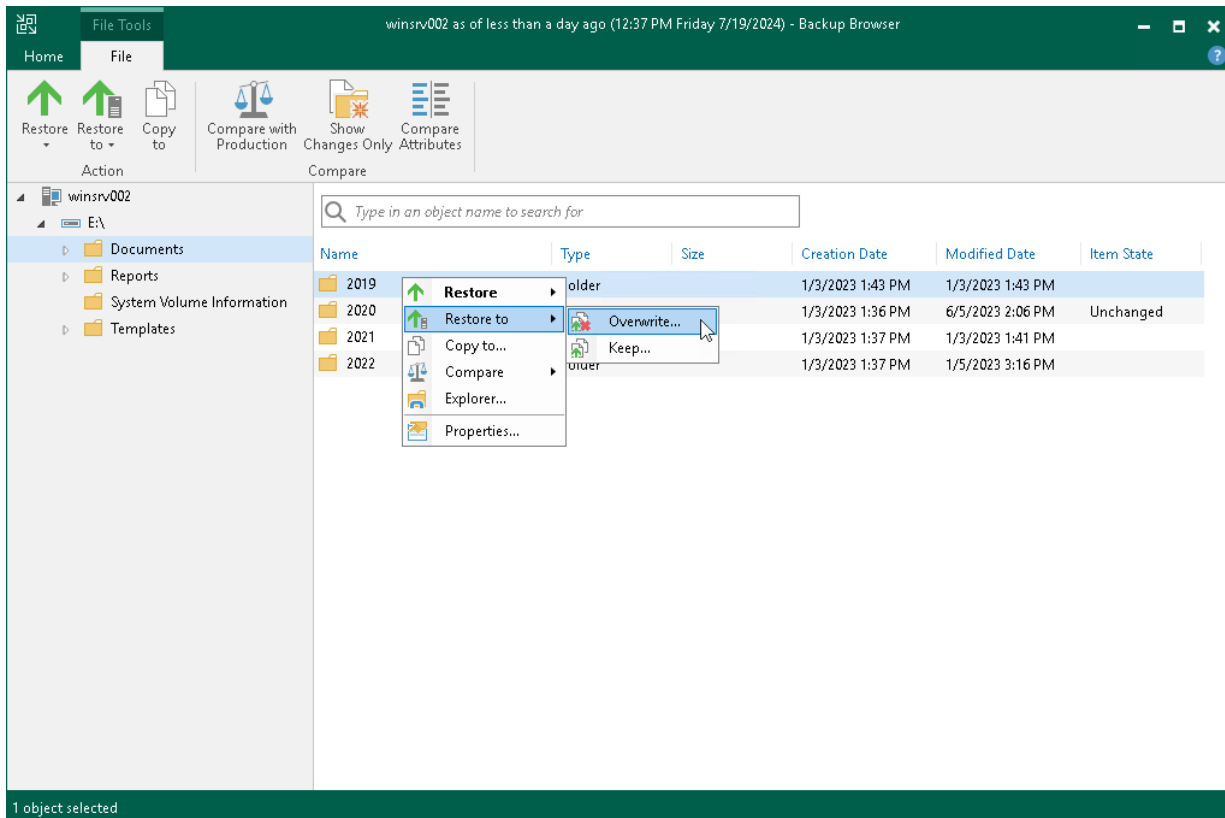
1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select one of the following:
  - If you want to overwrite files and folders with identical names on the target computer, select **Restore to > Overwrite**.

- If you want to keep files and folders with identical names on the target computer, select **Restore to > Keep**.

If there are items with identical names, Veeam Backup & Replication will add the *RESTORED\_YYYYMMDD\_HHMMSS* postfix to the original names and store the restored items on the target computer.

Alternatively, you can select the same commands on the ribbon.

3. In the **Select Host** window, select the target Veeam Agent computer.
4. If prompted, in the **Credentials** window, provide credentials to connect to the target computer.
5. In the **Choose Target Folder** window, specify a path to the folder where the restored objects will be saved.



## Restoring Permissions

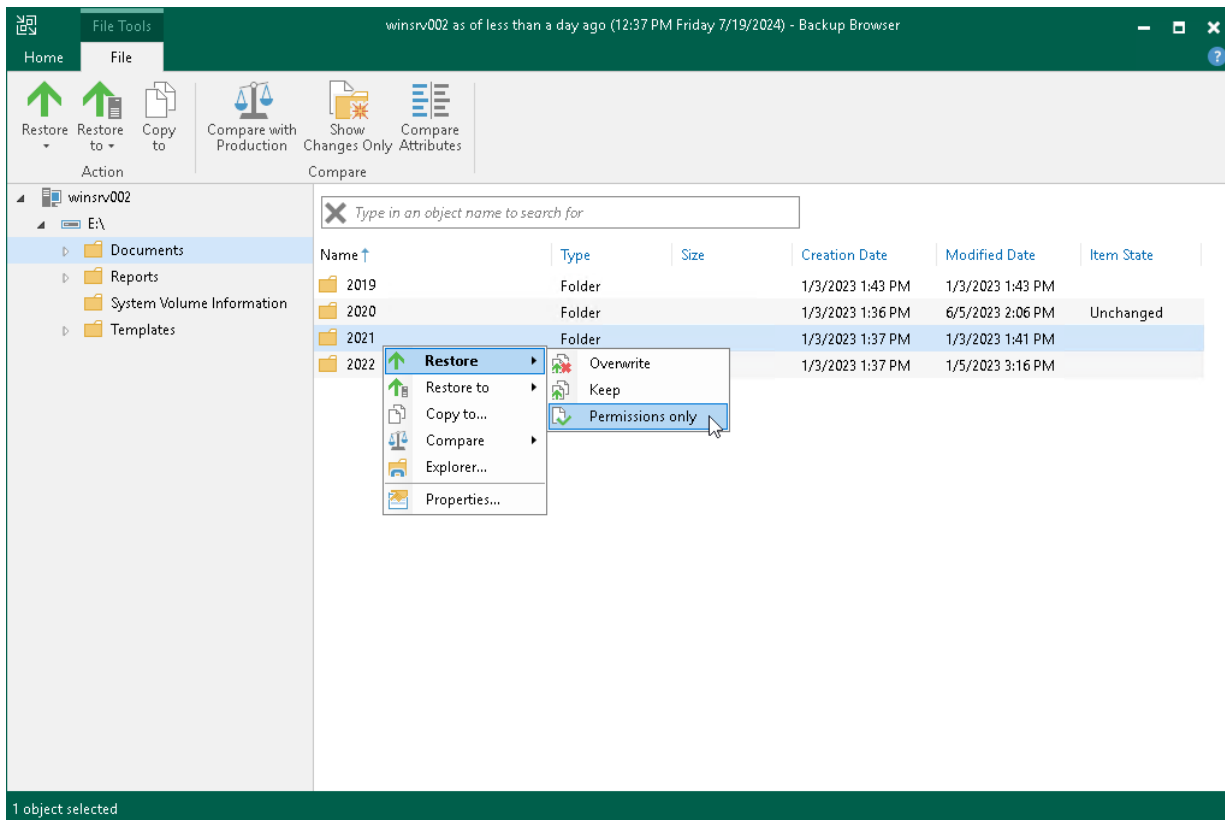
This functionality is available for files and folders in the non-comparison state.

To restore permissions for files and folders, do the following:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.



2. Right-click one of the selected items and select **Restore > Permissions only** or select **Restore > Permissions only** on the ribbon.



## Launching Application Item Restore

If you are restoring files from a Veeam Agent computer where the supported applications are installed, you can also launch application item restore directly from the Veeam Backup browser. Veeam Backup & Replication lets you restore items and objects from the following applications:

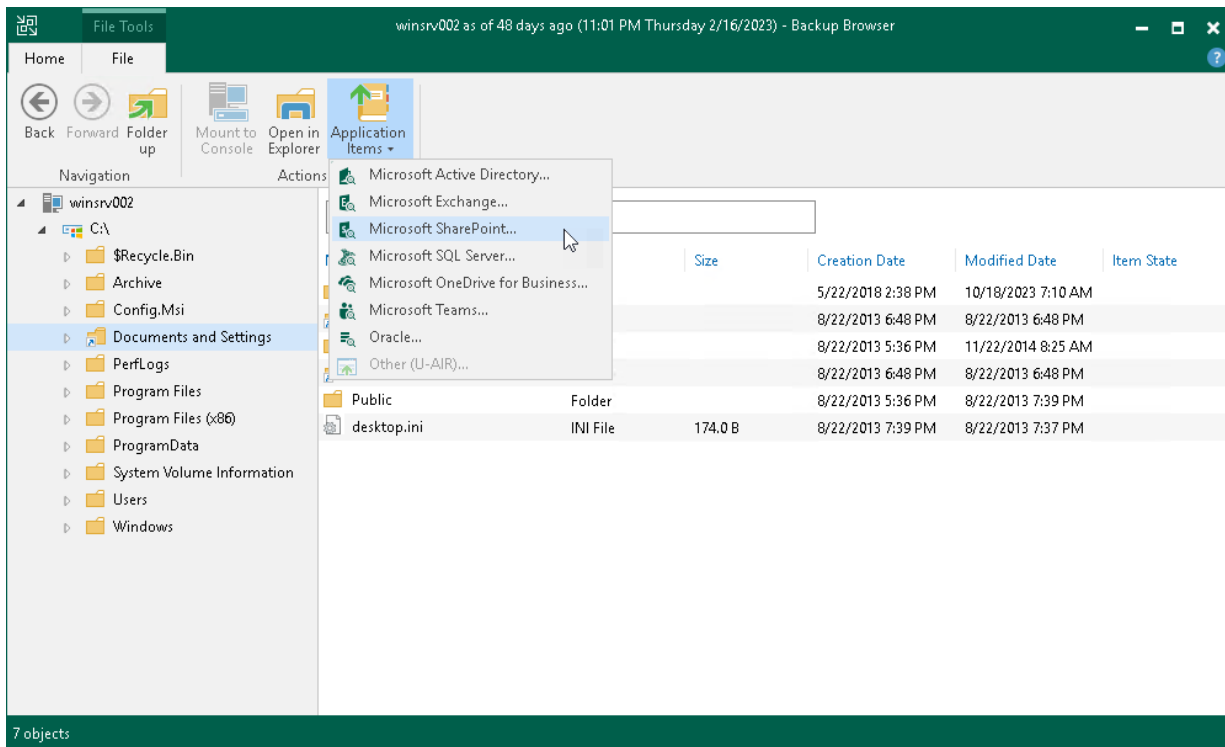
- Microsoft Active Directory
- Microsoft Exchange
- Microsoft SharePoint
- Microsoft SQL Server
- Oracle

To restore application items, Veeam Backup & Replication uses special tools called Veeam Explorers.

To launch application item restore, do the following:

1. On the ribbon, switch to the **Home** tab.
2. Click **Application Items** and select the required application.

3. In the opened Veeam Explorer, perform the necessary operations. For more information on Veeam Explorers, see the [Veeam Explorers User Guide](#).



## Working with Microsoft Windows File Explorer

You can use Microsoft Windows File Explorer to work with restored files and folders:

1. On the ribbon of the Veeam Backup browser, switch to the **Home** tab and click **Mount to Console** to mount the Veeam Agent computer disks to the Veeam Backup & Replication console.
2. To open Microsoft Windows File Explorer, do the following:
  - Click **Open in Explorer** on the Veeam Backup browser ribbon or right-click the necessary folder and select **Explorer**.
  - Click **File Explorer** in the **Start** menu of the machine where Veeam Backup & Replication console is installed. Browse to the `C:\VeeamFLR\<machinename>\<volume n>` folder where the disks of the machine are mounted and find the necessary files.

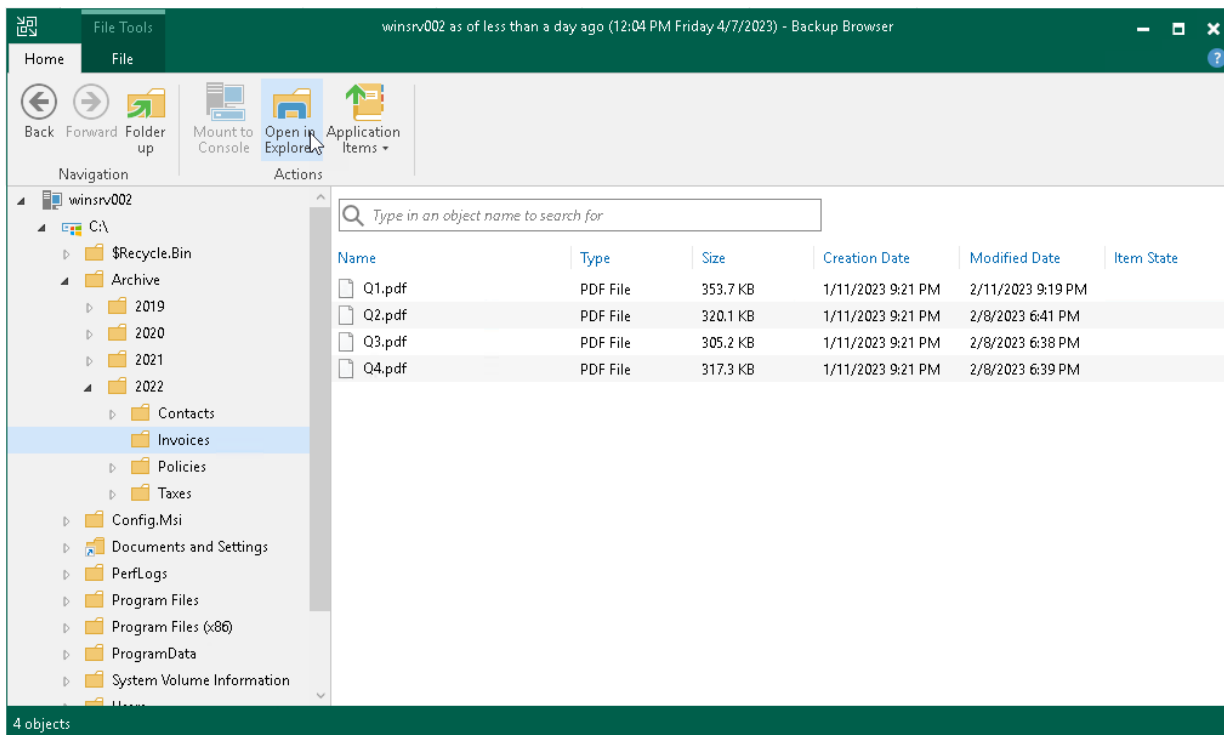
### NOTE

The **Mount to Console** button is not available if the mount point is already created on the Veeam Backup & Replication console.

It is recommended that you use Microsoft Windows File Explorer only to view file content, not to restore files. For file restore, use Veeam Backup browser. This browser has the following advantages:

1. You can browse the Veeam Agent computer file system ignoring the file system ACL settings.
2. You can preserve permissions and ownership during file-level restore.

If you open the Veeam Agent computer file system in Microsoft Windows Explorer, these capabilities are not available. For more information, see [Microsoft Docs](#).



## Closing Veeam Backup Browser

You can browse Veeam Agent computer files only while the Veeam Backup browser is open. After the Veeam Backup browser is closed, Veeam Backup & Replication unmounts Veeam Agent computer disks from the machine where the Veeam Backup & Replication console is installed and from the mount server (if you have restored Veeam Agent computer files to the original location).

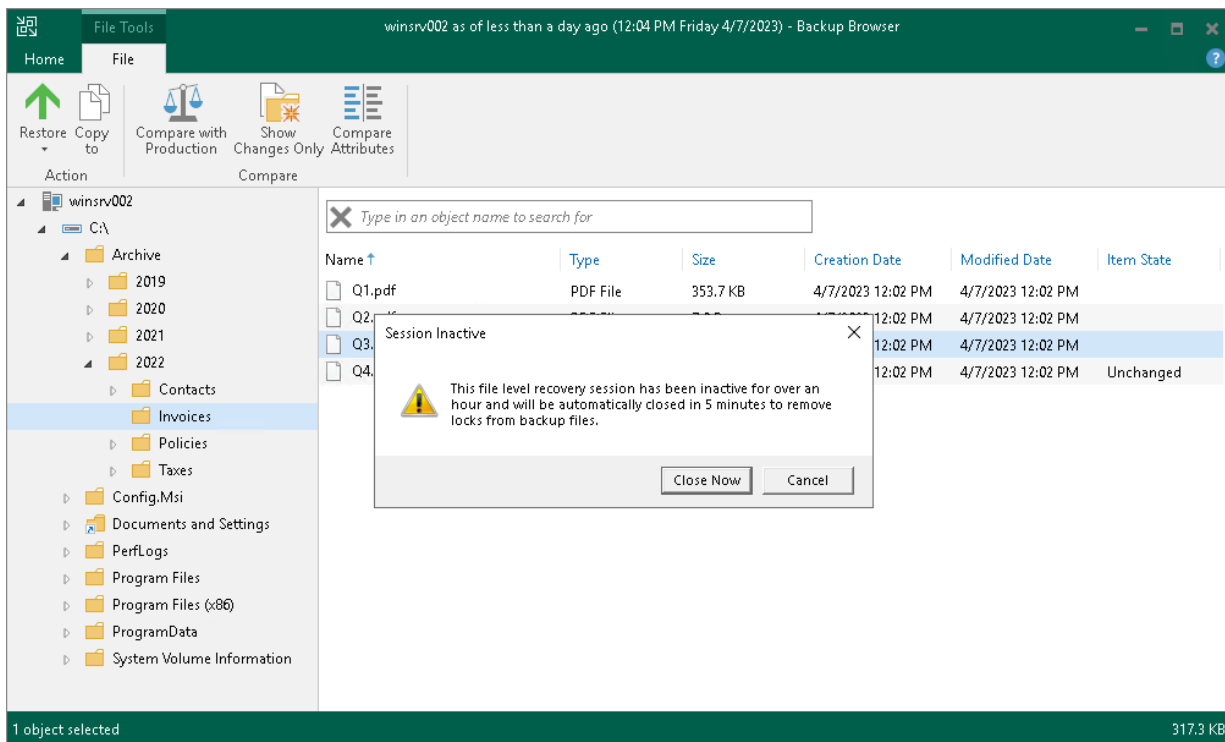
It is recommended that you close the Veeam Backup browser after you have finished restoring Veeam Agent computer files. When the Veeam Backup browser is open, the backup file whose computer file system is displayed in the browser is locked in the backup repository. As a result, some scheduled operations that use this backup file may fail.

Veeam Backup & Replication checks if there is any activity in the Veeam Backup browser with an interval of 5 minutes. If the user or Veeam Backup & Replication components and services do not perform any actions for 30 minutes, Veeam Backup & Replication displays a warning that the Veeam Backup browser is to be closed in 5 minutes.

After the warning is displayed, you can perform one of the following actions:

- You can close the Veeam Backup browser manually.
- You can click **Cancel** to postpone the close operation. In this case, the Veeam Backup browser will remain open for 5 minutes. After this period expires, Veeam Backup & Replication will display the warning again.

- You can perform no action at all. In this case, the Veeam Backup browser will close automatically in 5 minutes.

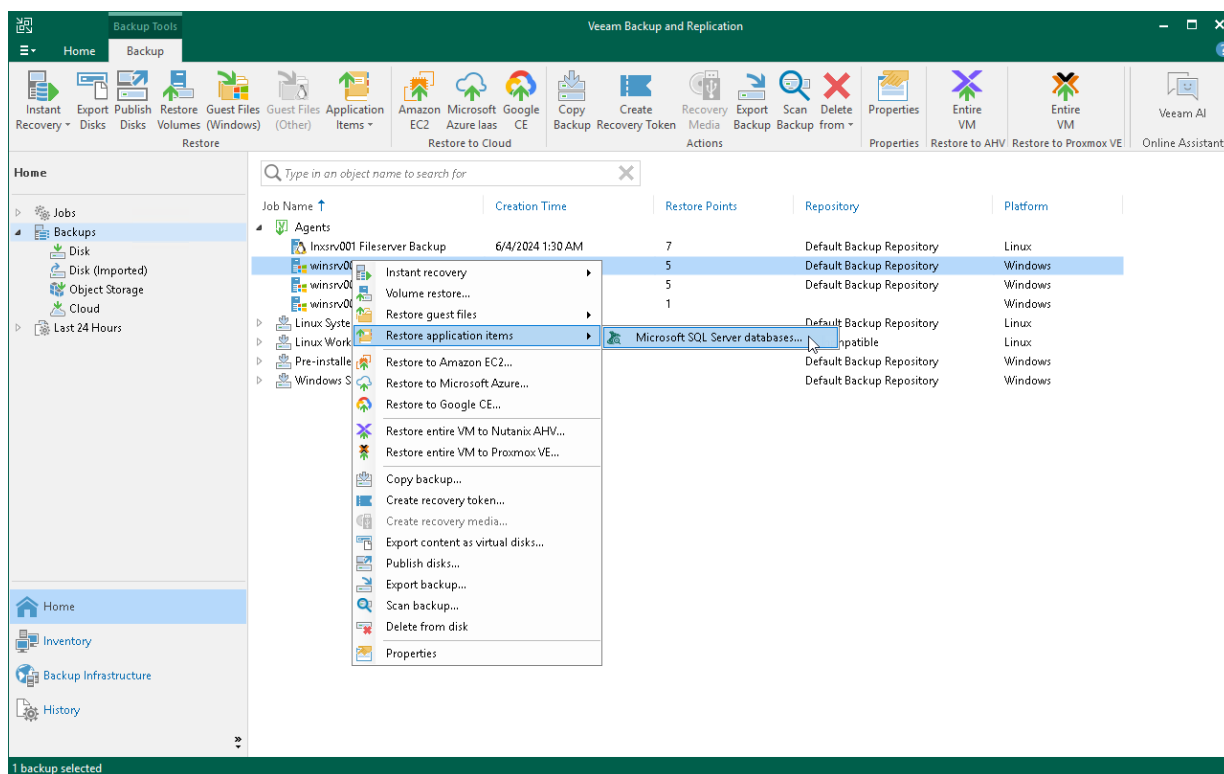


# Restoring Application Items

You can use Veeam Explorers to restore application items from backups created using Veeam Agent for Microsoft Windows. Veeam Backup & Replication lets you restore items and objects from the following applications:

- Microsoft Active Directory
- Microsoft Exchange
- Microsoft SharePoint
- Microsoft SQL Server
- Oracle

The procedure of application item restore from a Veeam Agent backup does not differ from the same procedure for a VM backup. To learn more, see the [Application Item Restore](#) section in the Veeam Backup & Replication User Guide.



# Exporting Disks

You can restore computer disks from Veeam Agent backups created using Veeam Agent for Microsoft Windows and convert them to disks of the VMDK, VHD or VHDX format.

During disks restore, Veeam Backup & Replication creates standard virtual disks that can be used by VMware vSphere and Microsoft Hyper-V VMs.

- When you restore a disk in the VMDK format, Veeam Backup & Replication creates a pair of files that make up the VM virtual disk: a descriptor file and file with the virtual disk content.
- When you restore a disk in the VHD/VHDX format, Veeam Backup & Replication creates a file of the VHD or VHDX format.

You can save converted disks locally on any server or SMB share added to the backup infrastructure or place disks on a datastore connected to an ESXi host (for VMDK disk format only). VMDK disks can be restored as thin provision and thick disks:

- Disks restored to a datastore are saved in the thin provisioned format.
- Disks restored to a server are saved in the thick provisioned format.

Veeam Backup & Replication supports batch disk restore. For example, if you choose to restore 2 computer disks, Veeam Backup & Replication will convert them to 2 virtual disks and store these disks in the specified location.

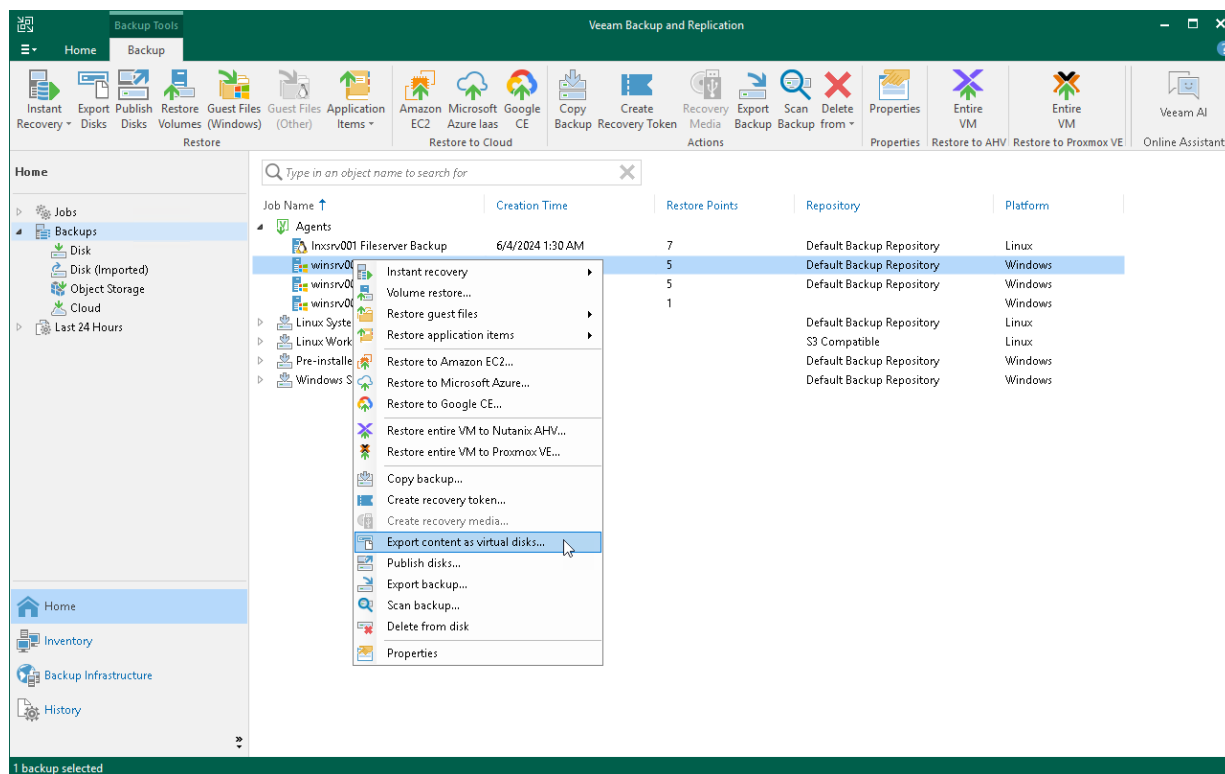
To restore disks and convert them to the VMDK, VHD or VHDX format, perform the following steps in the **Export Disk** wizard:

# Step 1. Launch Export Disk Wizard

To launch the **Export Disk** wizard, do either of the following:

- Open the **Home** tab and click **Restore > Agent > Disk restore > Export disk**. In this case, you will be able to select a backup of the necessary Veeam Agent computer at the **Backup** step of the wizard.
- Open the **Home** view. In the inventory pane, click the **Backups** node. In the working area, expand the necessary Veeam Agent backup, select the necessary computer in the backup and click **Export Disks** on the ribbon or right-click a computer in the backup and select **Export content as virtual disks**.

In this case, you will pass immediately to the **Restore Point** step of the wizard.



## Step 2. Select Backup

At the **Backup** step of the wizard, select a backup from which you want to restore disks. In the list of backups, Veeam Backup & Replication displays all backups that are currently hosted on the Veeam backup repository and Veeam Cloud Connect repository.

Export Disk

Backup

Select a backup to export disk content from.

Backup

Restore Point

Disks

Target

Secure Restore

Reason

Summary

Machine: filesrv03.tech.local

Job name	Last restore point	Objects	Restore points
DB Backup	12/28/2022 4:48:25 PM	2	
MySQL DB Backup	12/30/2022 10:48:53 ...	1	
PosgreSQL DB Bac...	1/12/2023 3:37:54 PM	1	
filesrv02.tech.lo...	less than a day ago (1...	3	
filesrv03.tech.lo...	less than a day ago (1...	2	
Workstations Back...	12/28/2022 4:49:27 PM	1	

Type in an object name to search for

< Previous

Next >

Finish

Cancel



# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point from which you want to restore disks. In the list of points, Veeam Backup & Replication displays all restore points that have been created. Make sure that you select a restore point that relates to the selected backup.

Export Disk

Restore Point

Select the restore point to export disks from.

Backup

Restore Point

Disks

Target

Secure Restore

Reason

Summary

Computer name: **filesrv03.tech.local**

Data size: **43.3 GB**

Available restore points:

Created	Type
less than a day ago (12:37 PM Thursday 1/19/2023)	Increment
less than a day ago (11:15 AM Thursday 1/19/2023)	Full

< Previous

Next >

Finish

Cancel

## Step 4. Select Disks

At the **Disks** step of the wizard, select check boxes next to those disks that you want to export.

Export Disk

Disk

Select one or more disks to export.

Backup

Restore Point

Disks

Target

Secure Restore

Reason

Summary

Disks:

Disk name	Size	Volumes	
<input checked="" type="checkbox"/> Disk 2	20 GB	Local Disk (E:)	

Select All

Clear All

< Previous

Next >

Finish

Cancel

## Step 5. Select Destination and Disk Format

At the **Target** step of the wizard, select the destination for disk export and format in which you want to save the resulting virtual disk.

1. From the **Server** list, select a server on which the resulting virtual disks must be saved. If you plan to save the disks in the VMDK format on a datastore, select an ESXi host to which this datastore is connected.
2. In the **Path to folder** field, specify a folder on the server or datastore where the virtual disks must be placed.
3. Select the export format for disks:
  - **VMDK** – select this option if you want to save the resulting virtual disk in the VMware VMDK format.
  - **VHD** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHD format.
  - **VHDX** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHDX format (supported by Microsoft Windows Server 2012 and later).
4. Click **Disk type** to specify how the resulting disk must be saved:
  - [For VMDK disk format] in the thin provisioned, lazy zeroed thick provisioned, or eagerly zeroed thick provisioned format
  - [For VHD and VHDX disk formats] in the dynamic or fixed format
5. [For export of a VMDK disk to an ESXi host] Click the **Pick proxy to use** link to select backup proxies over which backup data must be transported to the target datastore.

## NOTE

Consider the following:

- If you have selected to store the resulting virtual disk in a datastore, you will be able to save the virtual disk in the VMDK format only. Other options will be disabled.
- If you have selected to store the resulting virtual disk on the server running Microsoft Windows Server OS and in the VMDK format, you will be able to save the virtual disk in the lazy zeroed thick provisioned format only.

The image shows a screenshot of the 'Export Disk' dialog box in Veeam Backup & Replication. The 'Target' tab is selected in the left-hand navigation pane. The main area contains the following fields and options:

- Server:** A dropdown menu showing 'filesrv004'.
- Path to folder:** A text box containing 'C:\File\_Share\Veeam' and a 'Browse...' button.
- Export format:** Three radio button options:
  - ☐ **VMDK**: This virtual disk type is used by VMware products such as VMware Workstation, or VMware vSphere. Maximum VMDK disk size is 62TB. Pick proxy to use.
  - ☒ **VHD**: This virtual disk type is used by Microsoft products such as Microsoft Hyper-V or Microsoft Azure. Maximum VHD disk size is 2TB.
  - ☐ **VHDX**: This virtual disk type is used by more recent versions of Microsoft products such as Microsoft Hyper-V. Maximum VHDX disk size is 64TB.
- Disk type:** A dropdown menu showing 'Dynamic'.

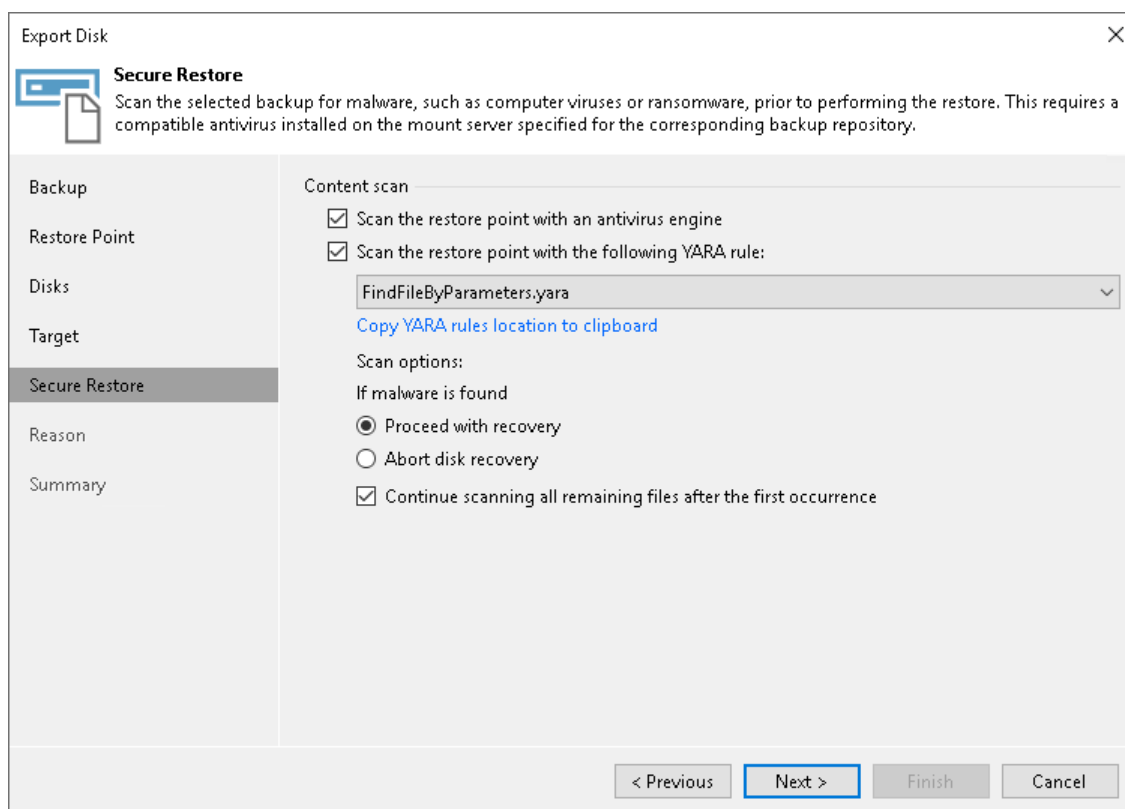
At the bottom of the dialog, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 6. Specify Secure Restore Settings

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to perform secure restore — scan restored disk data with antivirus software before restoring the disk. To learn more about secure restore, see the [Secure Restore](#) section in the Veeam Backup & Replication User Guide.

To specify secure restore settings:

1. In the **Content scan** section, specify the following:
  - a. To use antivirus software as a scan engine, select the **Scan the restore point with an antivirus engine** check box. For more information on antivirus scan, see the [Antivirus Scan \(Secure Restore\)](#) section in the Veeam Backup & Replication User Guide.
  - b. To use a YARA rule as a scan engine, select the **Scan the restore point with the following YARA rule** check box and choose a YARA rule from the drop-down list. By default, the YARA rules are located in the folder by the following path: C:\Program Files\Veeam\Backup and Replication\Backup\YaraRules.
2. Instruct Veeam Backup & Replication what to perform in case malware is found:
  - Select **Proceed with recovery** if you want to continue the recover process, despite the found malware threat.
  - Select **Abort disk recovery** if you want to stop the recovery process after the first malware threat is found.
3. In the **Scan options** section, select the **Continue scanning all remaining files after the first occurrence** check box if you want the antivirus software to continue volume scan after the first malware threat is found. For information on how to view results of the antivirus scan, see the [Viewing Malware Scan Results](#) section in the Veeam Backup & Replication User Guide.

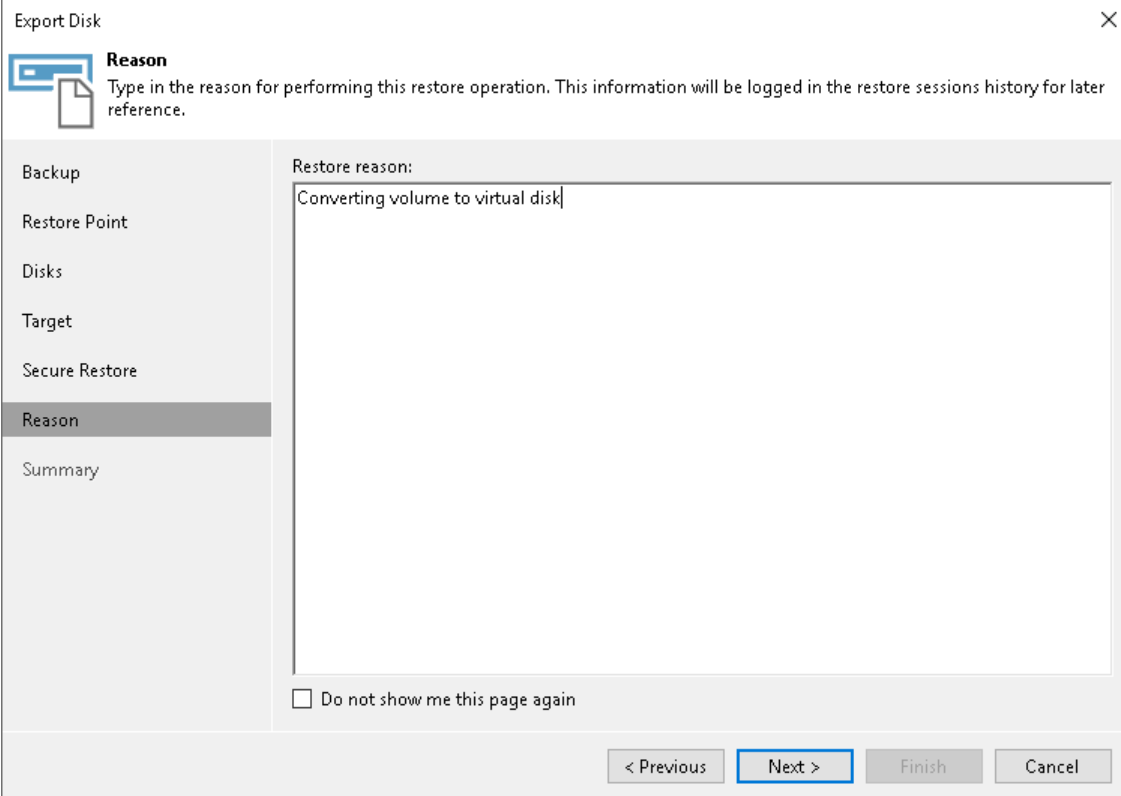


## Step 7. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the computer volume.

### TIP

If you do not want to display the **Restore Reason** step of the wizard in future, select the **Do not show me this page again** check box.

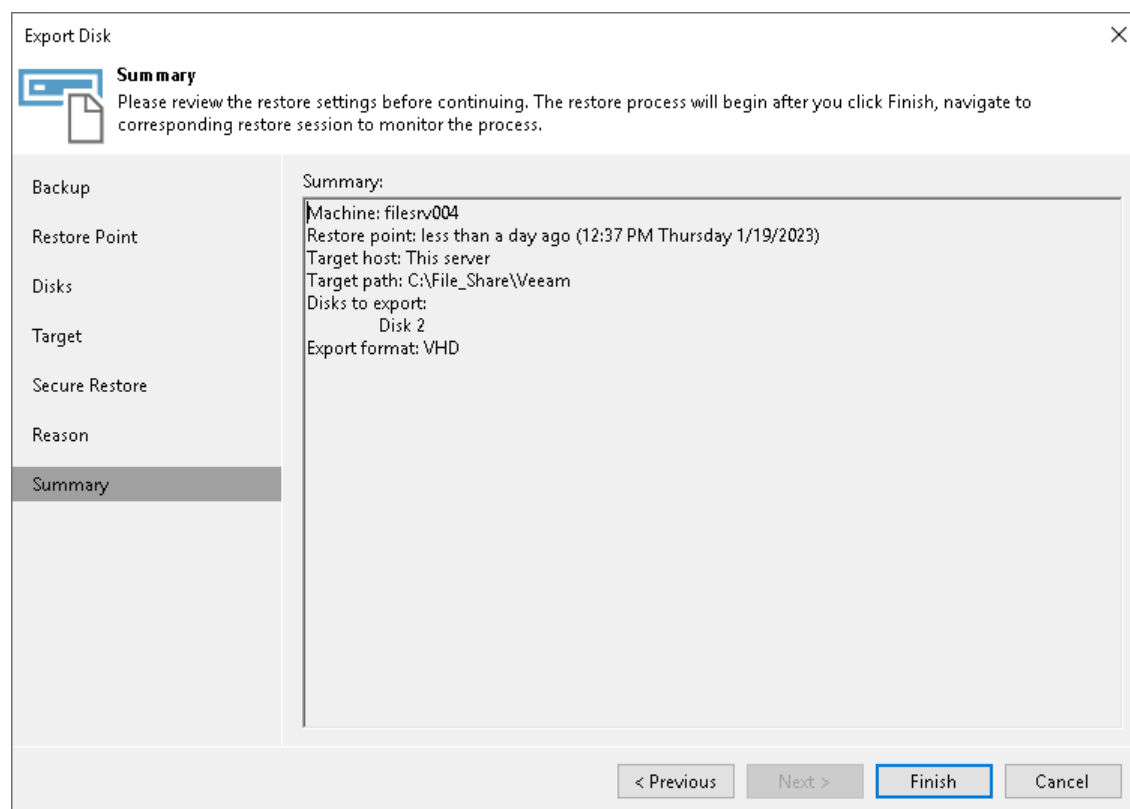


The screenshot shows the 'Export Disk' wizard window. The title bar says 'Export Disk' with a close button. The window has a sidebar on the left with the following steps: Backup, Restore Point, Disks, Target, Secure Restore, Reason (selected), and Summary. The main area is titled 'Reason' and contains the instruction: 'Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.' Below this is a large text box labeled 'Restore reason:' containing the text 'Converting volume to virtual disk'. At the bottom of the main area is a checkbox labeled 'Do not show me this page again'. The bottom of the window has four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 8. Complete Restore Process

At the **Summary** step of the wizard, complete the disk restore procedure.

1. Review details for the disk to be restored.
2. Click **Finish** to start the restore procedure and exit the wizard.



# Publishing Disks

You can use the Veeam backup console to publish disks from backups created by Veeam Agent backup jobs and backup copy jobs.

## TIP

You can publish disks using the PowerShell console. To learn more, see the [Disk Publishing \(Data Integration API\)](#) section in the Veeam PowerShell Reference.

Disk publishing allows you to save time by getting backup content of one or multiple disks instead of all disks from a backup. This technology gives read-only access to data and helps if you want to analyze data of your backup. For example, look for specific documents or usage patterns, or perform antivirus scan of backed-up data.

For Microsoft Windows-based Veeam Agent computers, disk publishing uses the iSCSI protocol. After the publishing, the target server can access the backup content using the iSCSI initiator and read the necessary data from the disk.

To learn more, see the [Disk Publishing](#) section in the Veeam Backup & Replication User Guide.

## Performing Disk Publish

Before you publish disks, [check prerequisites](#). Then use the **Publish Disks** wizard.

1. [Launch the wizard](#).
2. [Select a Veeam Agent computer whose disks you want to publish](#).
3. [Select a restore point](#).
4. [Select disks](#).
5. [Specify the target server](#).
6. [Specify a reason for disk publishing](#).
7. [Finish working with the wizard](#).

## Before You Begin

Before you publish disks, check the following requirements and limitations:

- The necessary ports must be opened on the target server. For more information, see [Ports](#).
- The target server must support the file system of the disk that you plan to publish.
- If data deduplication is enabled for some disks in a backup, data deduplication must be enabled on the target server.
- The target Microsoft Windows server must support the same ReFS version or later than the version used on the Veeam Agent computer from which you plan to publish disks. For more information on which OSes support which ReFS, see the [ReFS versions and compatibility matrix](#).

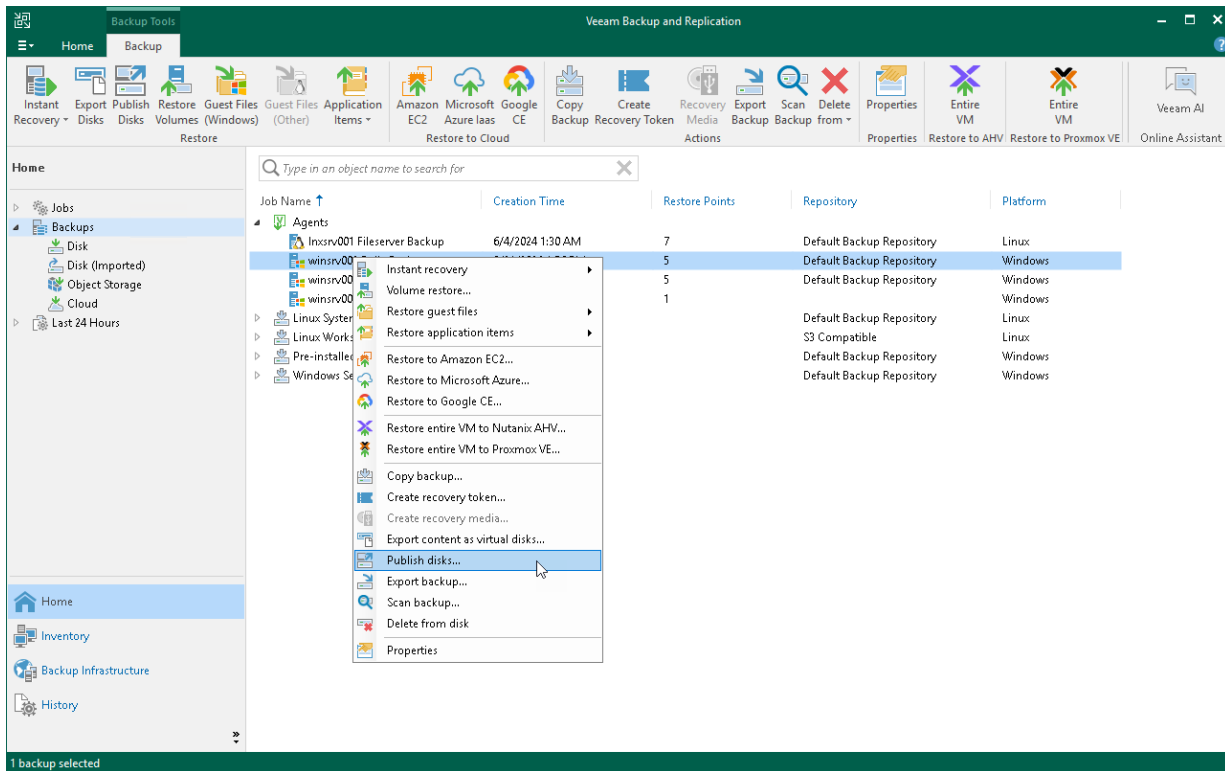
For the full list of limitations, see the [Considerations and Limitations](#) section in the Veeam Backup & Replication User Guide.



## Step 1. Launch Publish Disks Wizard

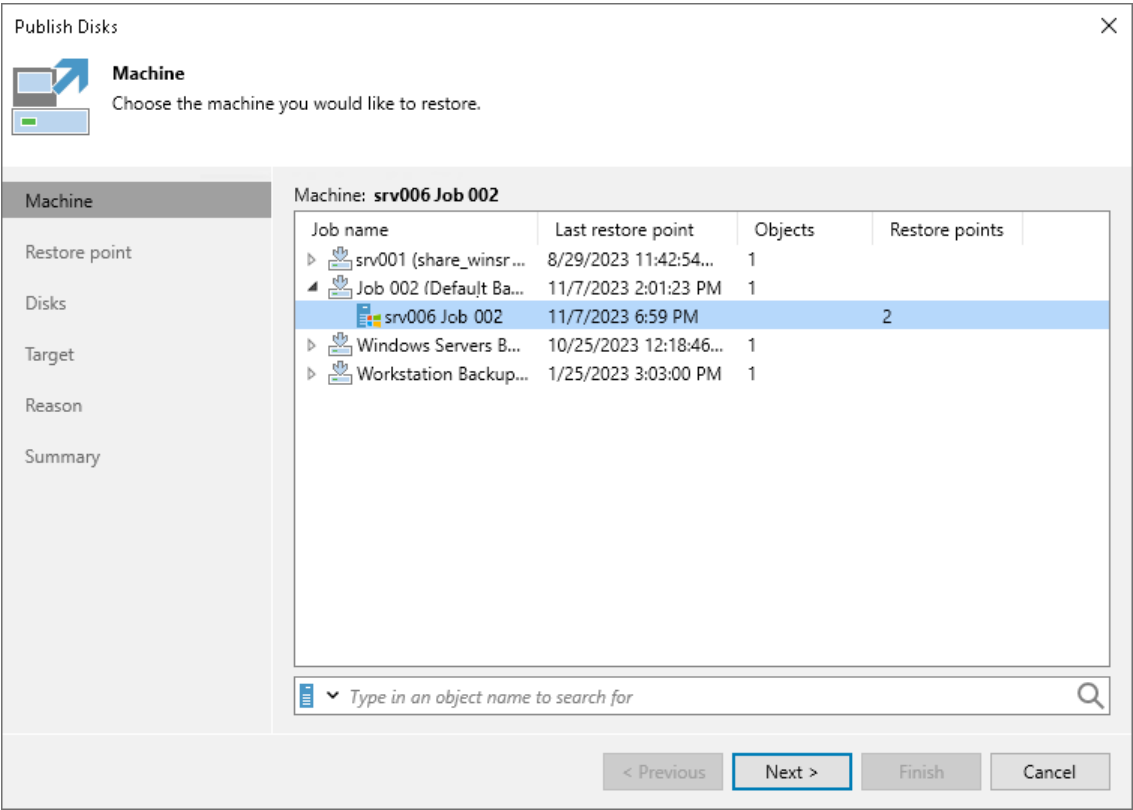
To launch the **Publish Disks** wizard, do either of the following:

- On the **Home** tab, click **Restore > Agent > Disk Restore > Publish disk**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary Veeam Agent backup, select a computer whose disks you want to publish and click **Publish Disks** on the ribbon. Alternatively, you can right-click the computer and select **Publish disks**. In this case, you will proceed to the [Restore point](#) step of the wizard.



## Step 2. Select Computer

At the **Machine** step of the wizard, expand a backup and select a Veeam Agent computer whose disks you want to publish.



### Step 3. Select Restore Point

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

Machine

Restore point

Disks

Target

Reason

Summary

Computer name: **srv006 Job 002**

Data size: **949 MB**

Available restore points:

Created	Type	Backup
less than a day ago (6:59 PM...	Increment	Job 002
less than a day ago (4:01 PM...	Full (M)	Job 002

< Previous

Next >

Finish

Cancel

# Step 4. Select Disks

At the **Disks** step of the wizard, select a check box next to the disks that you want to publish. Click **Select All** if you want to select all disks from the backup.

Machine

Restore point

**Disks**

Target

Reason

Summary

Disks:

Disk name	Size	Volumes
<input checked="" type="checkbox"/> Disk 2	9.98 GB	Local Disk (E:)
<input type="checkbox"/> Disk 1	106 GB	Local Disk (F:)

Select All

Clear All

< Previous

Next >

Finish

Cancel

## Step 5. Select Target Server

At the **Target** step of the wizard, select a Microsoft Windows server that will have access to the disk content.

You can select one of the following types of servers:

- A server added to the backup infrastructure.

If you want to add a new backup server to the backup infrastructure at this step, click **Add**. In this case, you will be able to add a new Microsoft Windows server. To learn more, see the [Adding Microsoft Windows Servers](#) section in the Veeam Backup & Replication User Guide.

- A temporary server. In this case, select *Specify a different host* from the drop-down list. In the **Target Server** window, specify the following settings:
  - a. In the **Host name** field, specify a server name or IP address of the server.
  - b. Select the account from the **Credentials** list. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add a new account in the Credentials Manager. To learn more, see the [Credentials Manager](#) section in the Veeam Backup & Replication User Guide.
- The original server. In this case, select *Original server* from the drop-down list.

If prompted, specify credentials for the target server.

The screenshot shows the 'Publish Disks' wizard window with the 'Target' step selected. The window title is 'Publish Disks' with a close button (X) in the top right corner. On the left, there is a sidebar with icons and labels: 'Machine', 'Restore point', 'Disks', 'Target' (highlighted), 'Reason', and 'Summary'. The main area is titled 'Target' with the instruction 'Select a target server to mount disks to.' Below this, there is a 'Target server:' label followed by a dropdown menu showing 'VBR12-01.tech.local (temporary added host)' and a small downward arrow. To the right of the dropdown is an 'Add...' button. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

# Step 6. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for publishing disks.

TIP

If you do not want to show this page, select the **Do not show me this page again** check box. If you further will want to return this page, follow the instructions described in [this Veeam KB article](#).

Publish Disks

Reason

Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

Machine

Restore point

Disks

Target

**Reason**

Summary

Restore reason:

Compare disk content

☐ Do not show me this page again

< Previous

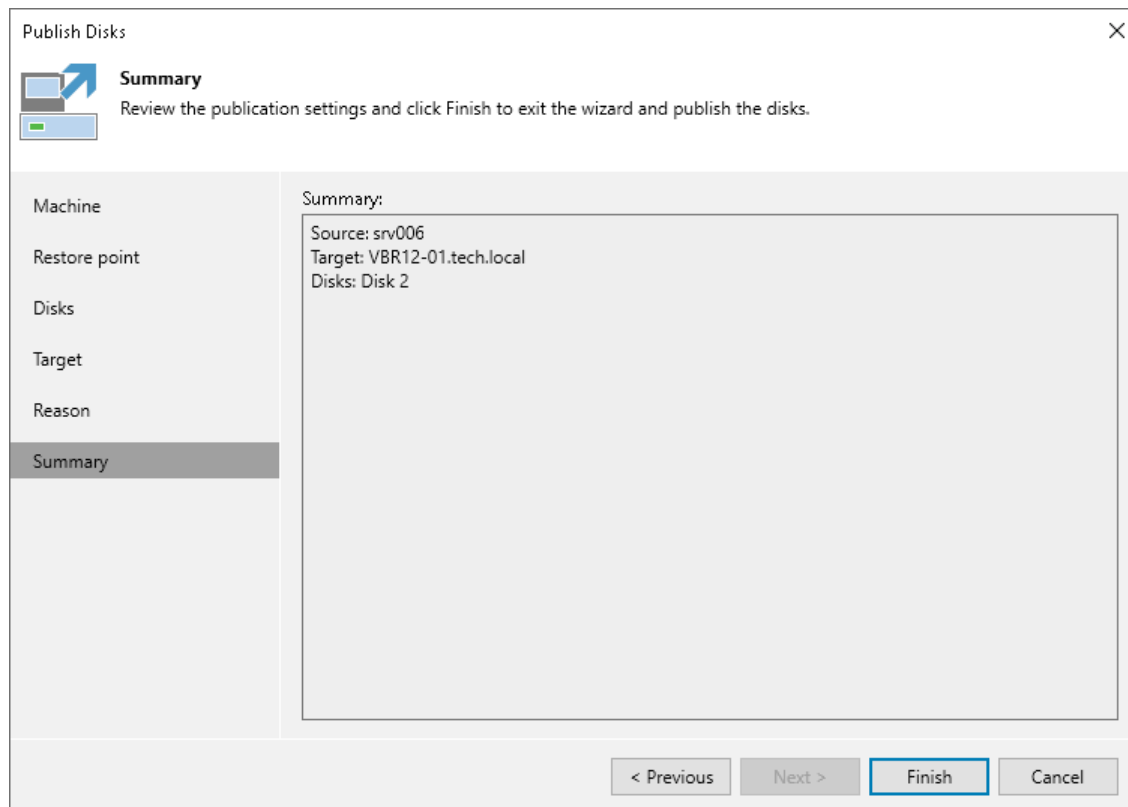
Next >

Finish

Cancel

## Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.



## What You Do Next

After the disks are published, go to the `C:\VeeamFLR\` folder on the target server to browse disks content.

After you started a disks publishing session, you can view the session statistics or stop the session from the Veeam backup console. To learn more, see [Managing Publishing Disks Session](#).

## Managing Publishing Disks Session

After you started a publishing session, you can check details about the session or stop it.

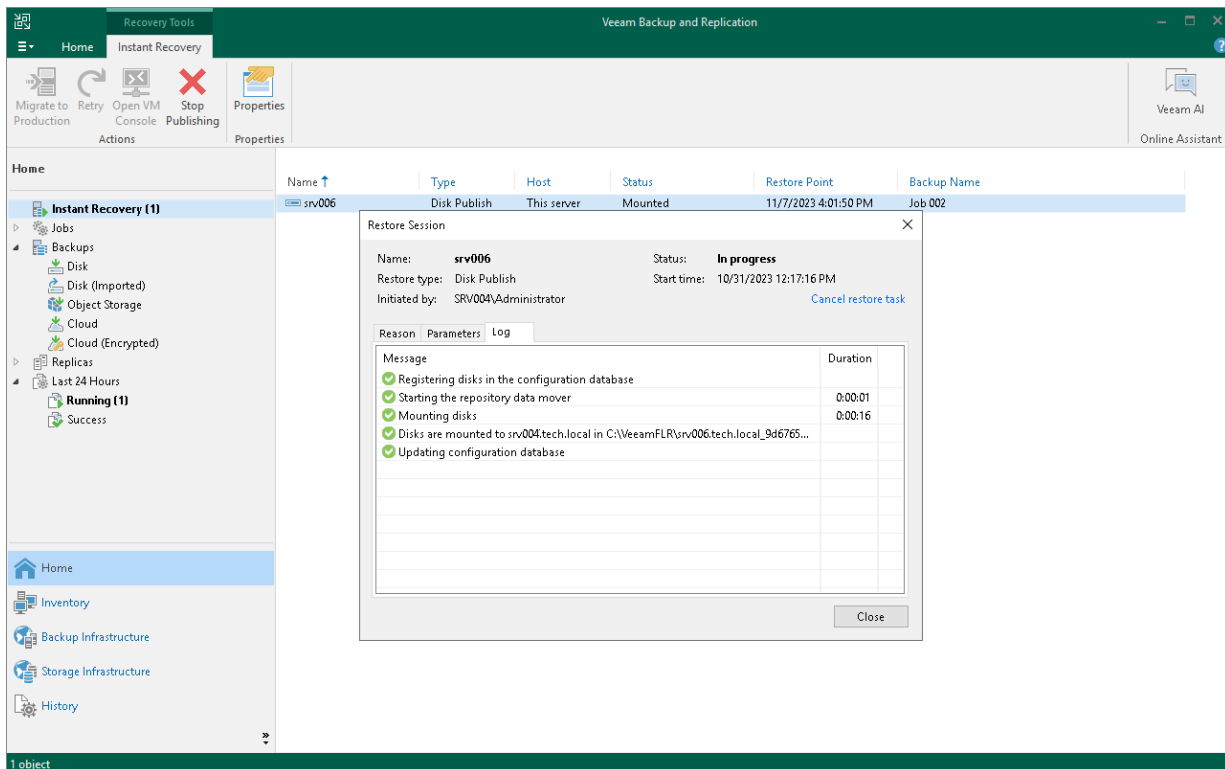
### Viewing Statistics on Publishing Session

To view publishing session statistics, do one of the following:

- Open the **Home** view. In the inventory pane, select **Instant Recovery**. In the working area, select the necessary publishing session and click **Properties** on the ribbon. Alternatively, right-click the session and **Properties**.
- Open the **Home** view. In the inventory pane select **Last 24 hours**. In the working area, double-click the necessary publishing session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.
- Open the **History** view. In the inventory pane select **Restore**. In the working area, double-click the necessary publishing session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

The publishing statistics provides the following data:

- At the top of the **Restore Session** window, Veeam Backup & Replication shows general session statistics. It includes a name of the Veeam Agent computer whose disk you want to publish, a name of the backup server which initiated the publishing session, a user name of the account under which the session was started, session status and duration details.
- The **Reason** tab shows the reason for the publishing session.
- The **Parameters** tab shows information about the target server, the Veeam Agent computer whose disks you publish and the restore point selected for publishing.
- The **Log** tab shows the list of operations performed during the session.



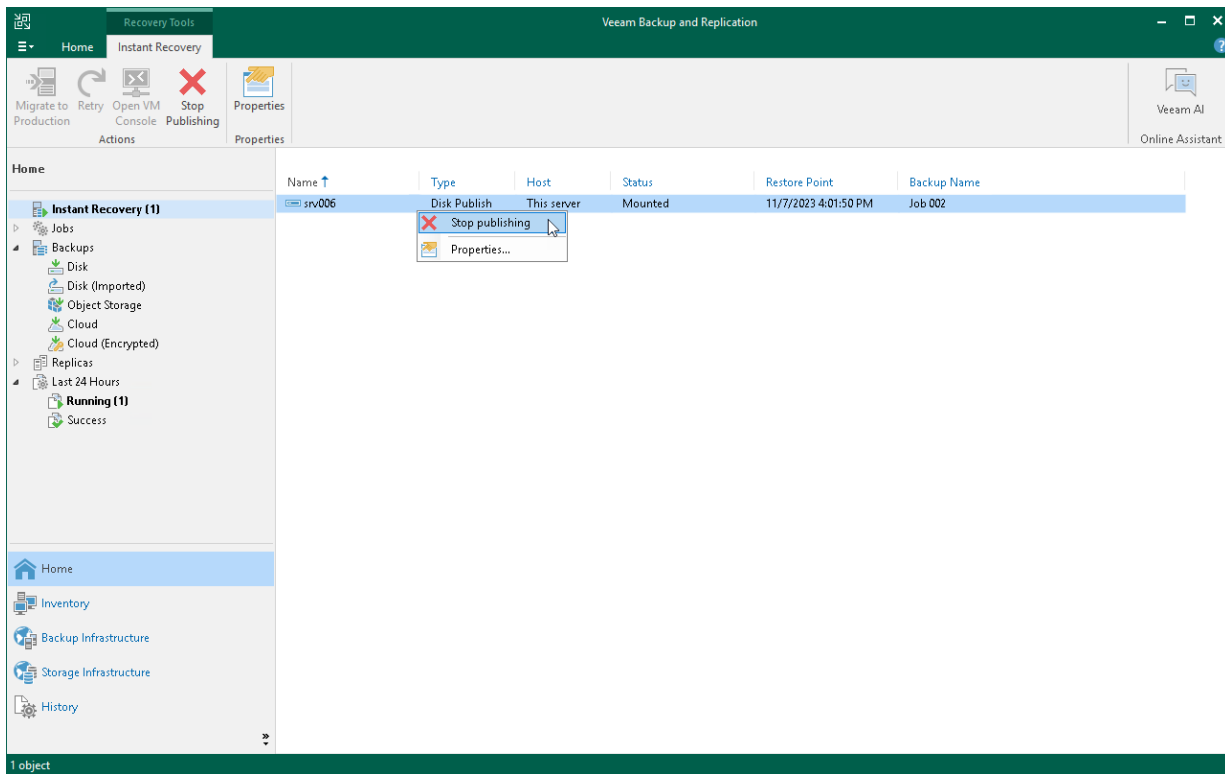
## Stopping Publishing Session

To stop a publishing session, do one of the following:

- Open the **Home** view. In the inventory pane select **Instant Recovery**. In the working area, double-click the necessary publishing session and click **Cancel restore task** in the **Restore Session** window. Alternatively, you can select the necessary publishing session and click **Stop Publishing** on the ribbon or right-click the session and click **Stop Publishing**.
- Open the **Home** view. In the inventory pane select **Last 24 hours**. In the working area, double-click the necessary publishing session and click **Cancel restore task** in the **Restore Session** window. Alternatively, you can select the necessary publishing session and click **Stop** on the ribbon or right-click the session and click **Stop session**.

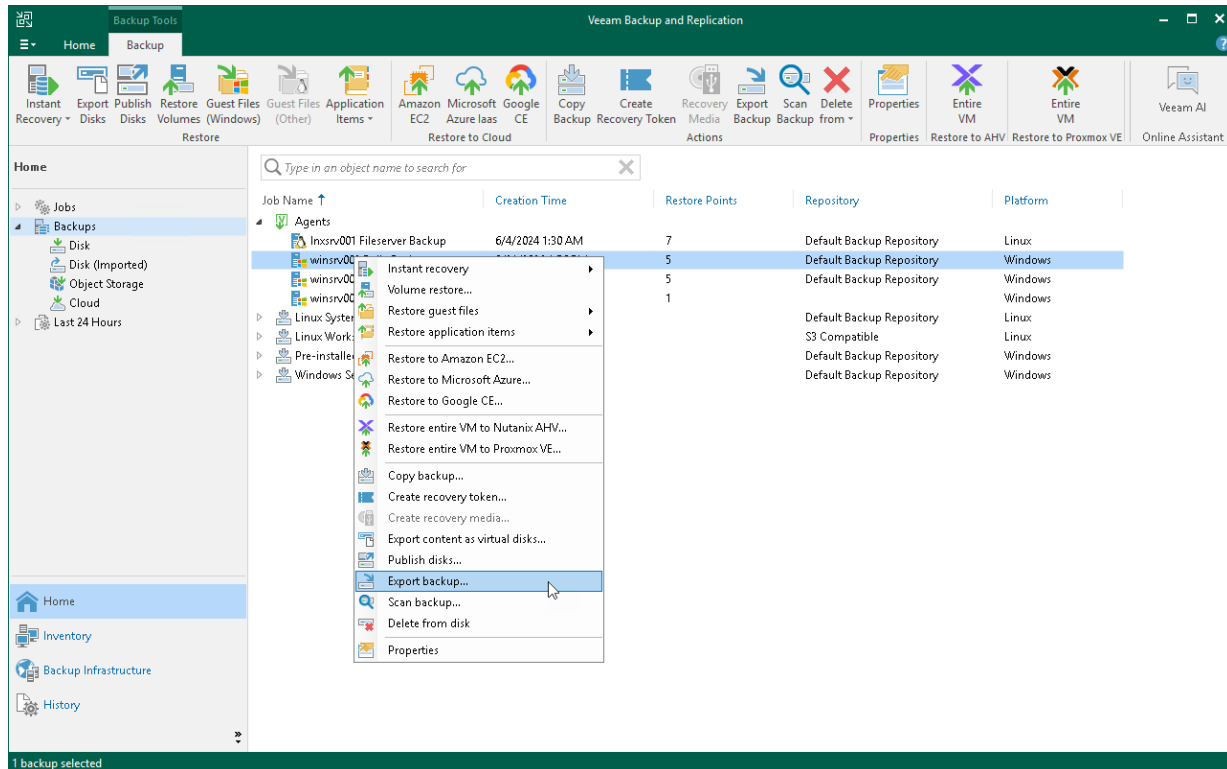


- Open the **History** view. In the inventory pane select **Restore**. In the working area, select the necessary publishing session and double-click it. In the **Restore Session** window, click **Cancel restore task**. Alternatively, you can right-click the publishing session and click **Stop session**.



# Exporting Restore Point to Full Backup File

You can restore data from a specific restore point in a Veeam Agent backup and export this data to a standalone full backup file. The procedure of Veeam Agent backup export does not differ from the same procedure for a VM. To learn more, see the [Exporting Backups](#) section in the Veeam Backup & Replication User Guide.



# Performing Administration Tasks

You can manage Veeam Agent backup jobs and backups created with these jobs. Veeam Backup & Replication allows you to perform the following administration tasks:

- [Import Veeam Agent backups.](#)
- [Enable and disable Veeam Agent backup jobs.](#)
- [View Veeam Agent backup job statistics.](#)
- [Delete Veeam Agent backup jobs.](#)
- [View Veeam Agent backup properties.](#)
- [Create recovery token.](#)
- [Copy Veeam Agent backups.](#)
- [Remove Veeam Agent backups.](#)
- [Delete Veeam Agent backups.](#)
- [Configure global settings.](#)
- [Assign roles to users.](#)

# Importing Veeam Agent Backups

You may need to import a Veeam Agent backup in the Veeam Backup & Replication console in the following situations:

- The Veeam Agent backup is stored on a drive managed by another computer (not the Veeam backup server).
- The Veeam Agent backup is stored in a backup repository managed by another Veeam backup server.
- The Veeam Agent backup has been removed in the Veeam Backup & Replication console.

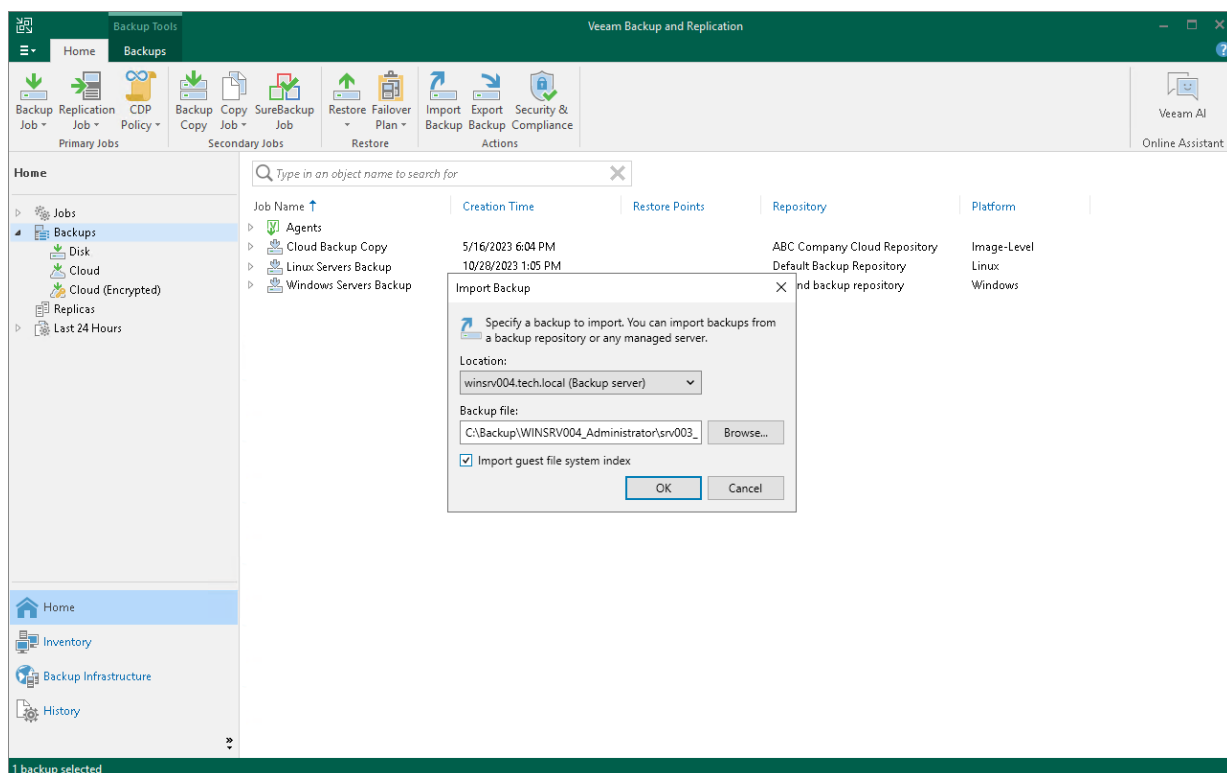
After importing, the Veeam Agent backup becomes available in the Veeam Backup & Replication console. You can restore data from such backup in a regular manner.

Before importing a backup, check the following prerequisites:

- The computer or server from which you plan to import the backup must be added to Veeam Backup & Replication. Otherwise you will not be able to access backup files.
- To be able to restore data from previous backup restore points, make sure that you have all incremental restore points in the same folder where the full backup file resides.

To import a Veeam Agent backup:

1. In Veeam Backup & Replication, click **Import Backup** on the **Home** tab.
2. From the **Computer** list, select the computer or server on which the backup you want to import is stored.
3. Click **Browse** and select the necessary VBM or VBK file. If you select the VBM file, the import process will be notably faster. We recommend that you use the VBK files for import only if a corresponding VBM file is not available.
4. Click **OK**. The imported backup will become available in the **Home** view, under the **Backups > Disk (imported)** node in the inventory pane.




# Importing Encrypted Backups

You can import Veeam Agent backups that were encrypted by Veeam Backup & Replication or Veeam Agent for Microsoft Windows.

To import an encrypted backup file:

1. On the **Home** tab, click **Import Backup**.
2. From the **Computer** list, select the host on which the backup you want to import is stored.
3. Click **Browse** and select the VBM or VBK file.
4. Click **OK**. The encrypted backup will appear under the **Backups > Disk (encrypted)** node in the inventory pane.
5. In the working area, select the imported backup and click **Specify Password** on the ribbon, or right-click the backup and select **Specify password**.
6. In the **Password** field, enter the password for the backup file. If you changed the password one or several times while the backup chain was created, you need to specify the latest password. For Veeam Agent backups, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

If you enter correct password, Veeam Backup & Replication will decrypt the backup file. The backup will be moved under the **Backups > Disk (imported)** node in the inventory pane.



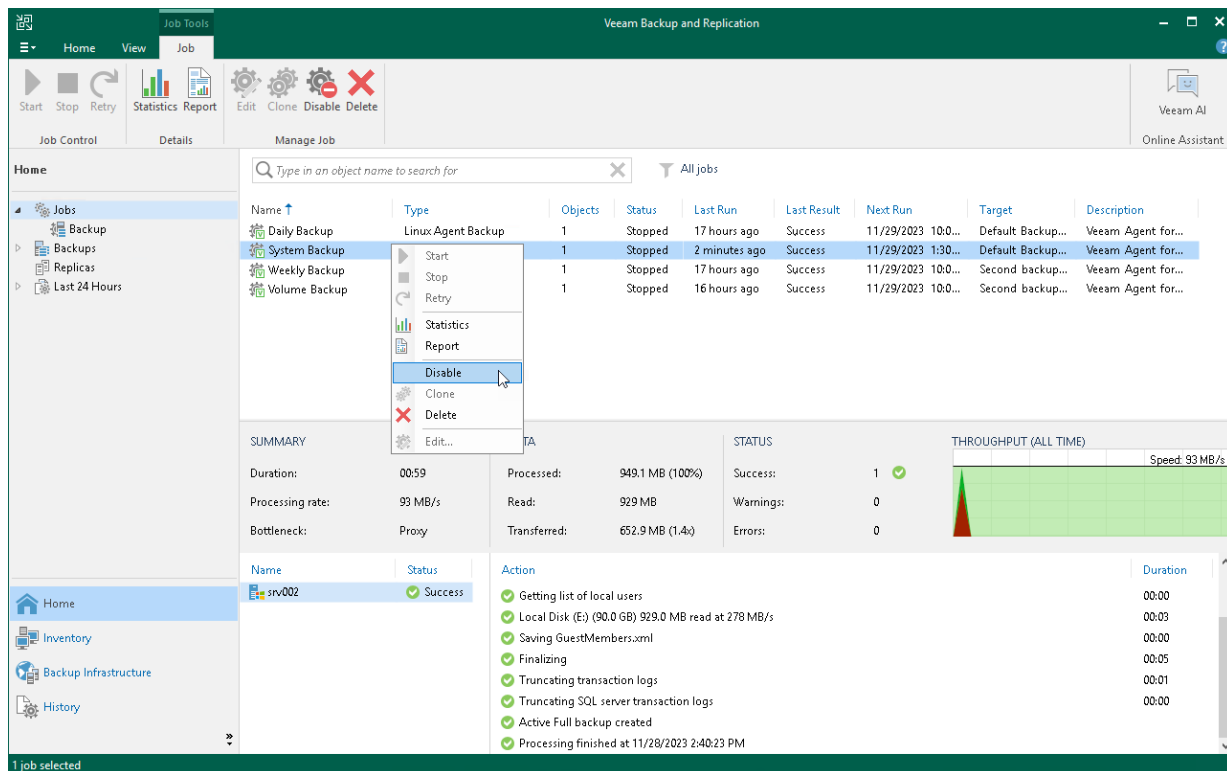
# Enabling and Disabling Veeam Agent Backup Jobs

You can disable and enable Veeam Agent jobs in Veeam Backup & Replication.

When you disable the job, you prohibit the user to store the resulting backup in the backup repository. If the user starts a disabled job manually or the job starts by schedule, the job session will fail and report the " *The job has been disabled by the Veeam Backup & Replication administrator*" error. To let Veeam Agent store backups in the backup repository again, you must enable the disabled job.

To disable or enable the scheduled backup job in Veeam Backup & Replication:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Disable** on the ribbon, or right-click the necessary job in the working area and select **Disable**. To enable the disabled job, click **Disable** on the toolbar, or right-click the job and select **Disable** once again.

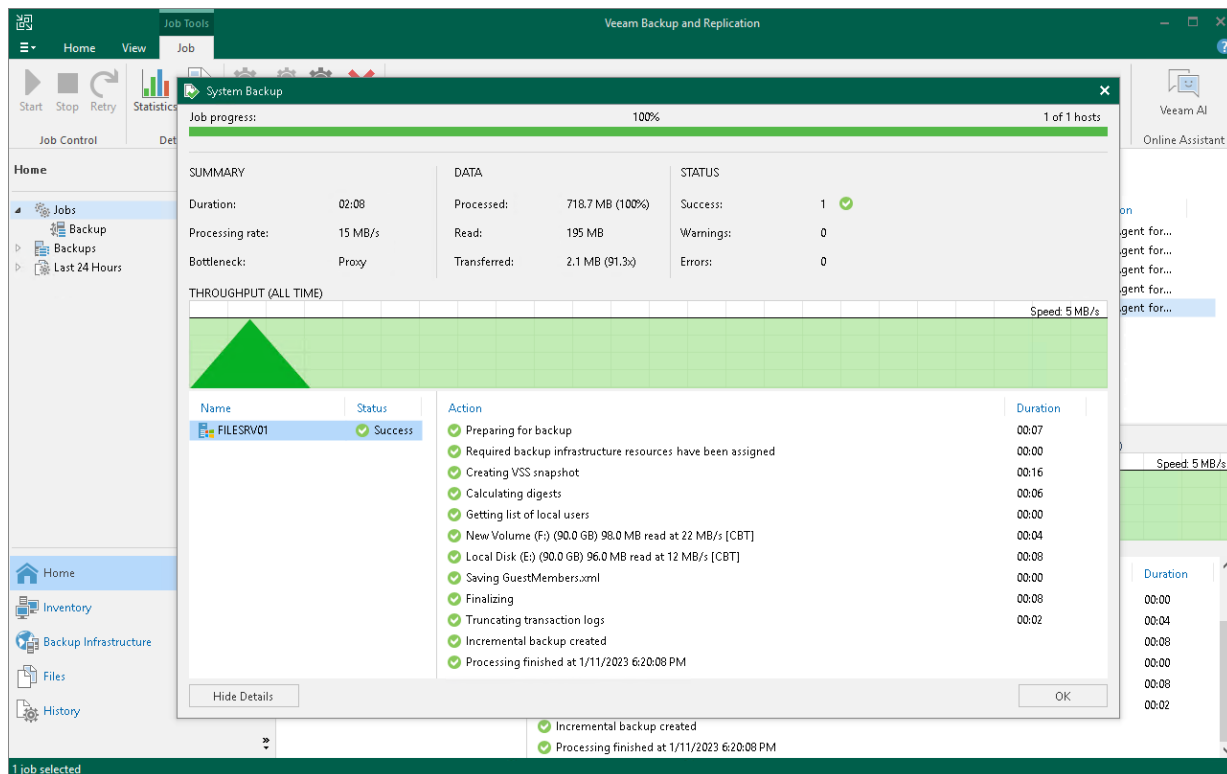


# Viewing Veeam Agent Backup Job Statistics

You can view statistics about Veeam Agent backup jobs in the Veeam Backup & Replication console. Veeam Backup & Replication displays statistics for Veeam Agent backup jobs in the similar way as for regular backup jobs. The difference is that the list of objects included in the job contains a Veeam Agent machine instead of one or several VMs.

To view Veeam Agent backup job statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. In the working area, select the necessary Veeam Agent backup job and click **Statistics** on the ribbon, or right-click the job and select **Statistics**.



# Deleting Veeam Agent Backup Jobs

You can delete Veeam Agent backup jobs.

When you delete a Veeam Agent backup job, Veeam Backup & Replication removes all records about the job from its database and console. When the user starts a new Veeam Agent backup job session manually or the job starts automatically by schedule, the job will appear in the Veeam Backup & Replication console again, and records about a new job session will be stored in the Veeam Backup & Replication database.

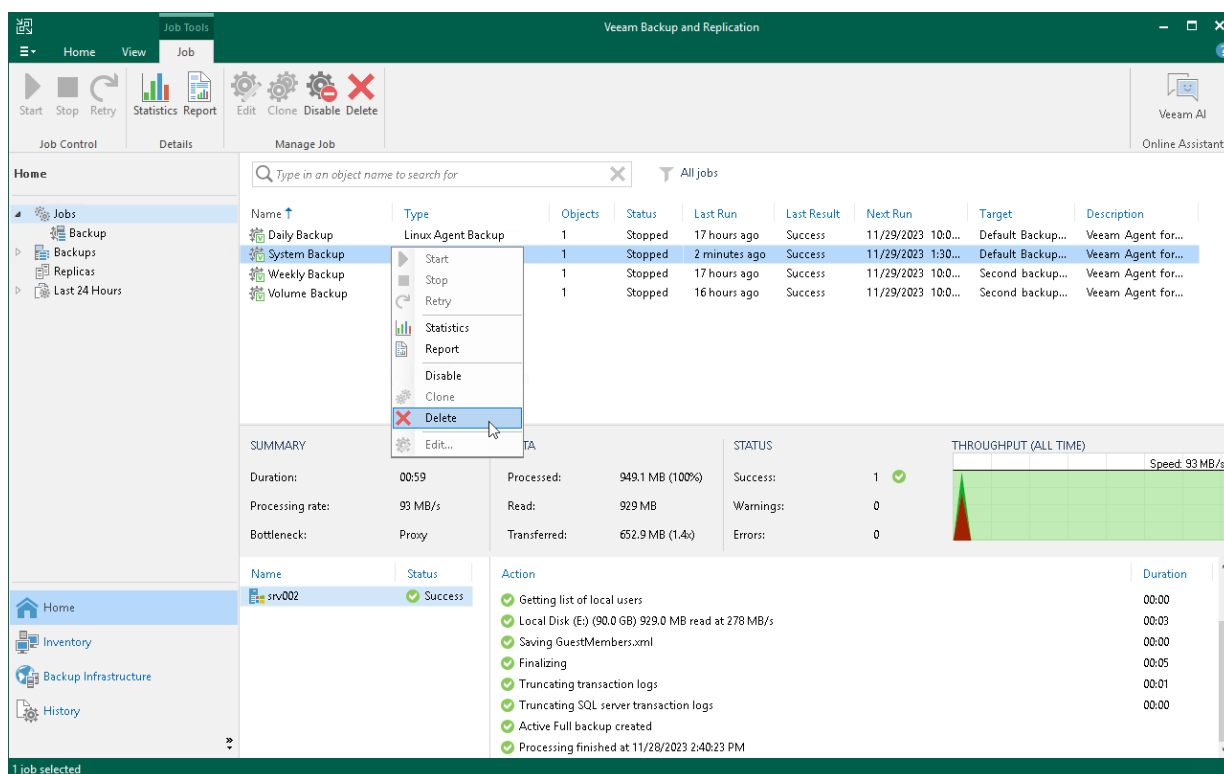
## NOTE

When you delete a Veeam Agent backup job, the backup files become orphaned and can be deleted by the background retention. For more information about the background retention, see the [Background Retention](#) section in the Veeam Backup & Replication User Guide.

To prevent the job from starting permanently, you must delete the job and unassign access rights permissions for this user from the backup repository. To completely delete the job, you must perform this operation in Veeam Agent on the Veeam Agent computer.

To remove a job:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Delete** on the ribbon, or right-click the necessary job in the working area and select **Delete**.



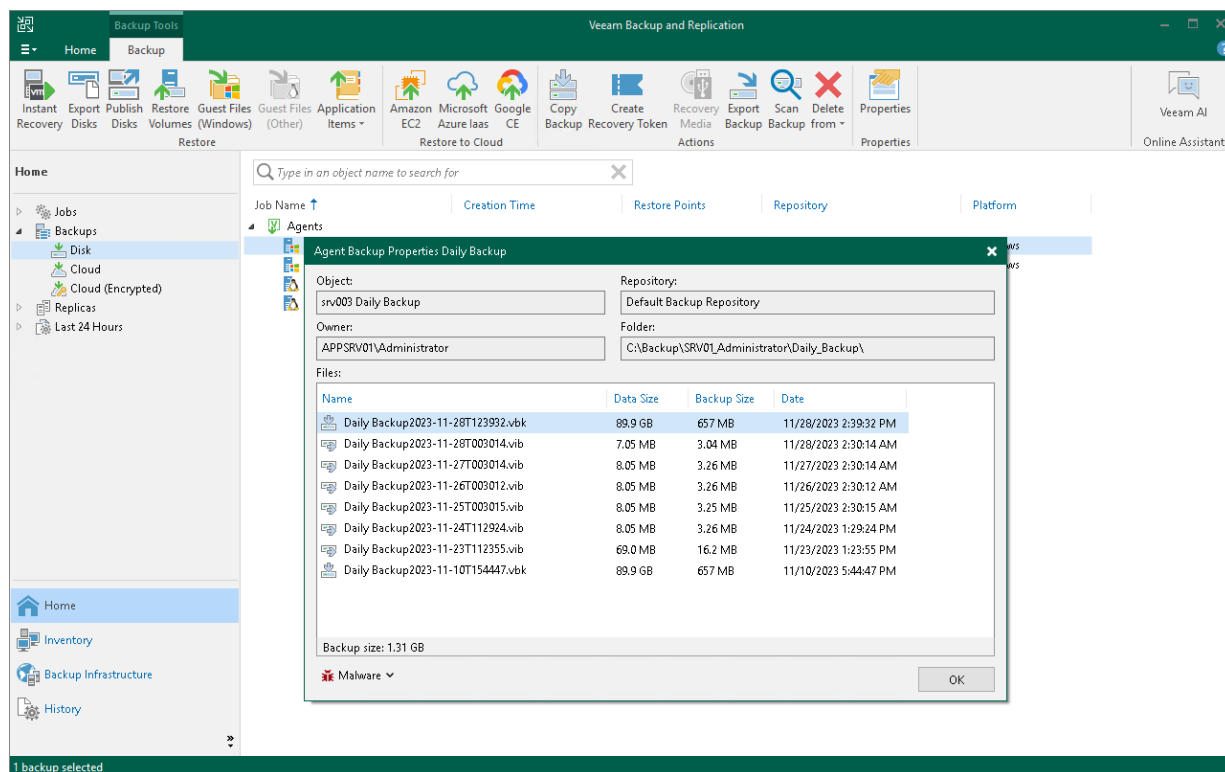


# Viewing Veeam Agent Backup Properties

You can view statistics about Veeam Agent backups.

To view Veeam Agent backup statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click **Disk** under the **Backups** node.
3. In the working area, expand the **Agents** node, select the necessary backup and click **Properties** on the ribbon, or right-click the backup and select **Properties**.



# Creating Recovery Token

If you want to recover volumes or an entire computer protected with Veeam Agent, you can use the **Create recovery token** operation.

You can generate the recovery token on the Veeam Backup & Replication side. Then, on the computer side, with this recovery token get access to the backup and recover data that are stored in the backup.

To learn more, see [Backup Server Settings](#).

## Considerations and Limitations

Before creating a recovery token, consider the following prerequisites and limitations:

- Recovery tokens stay valid for 24 hours.
- You can recover files and folders from the selected backups only.
- During recovery, Veeam Backup & Replication does not stop backup operations.
- You cannot create a recovery token for backups stored in Veeam Cloud Connect repository.

## Generating Recovery Token

To create a recovery token on the Veeam Backup & Replication side:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. In the working area, right-click the backup and select **Create recovery token**.

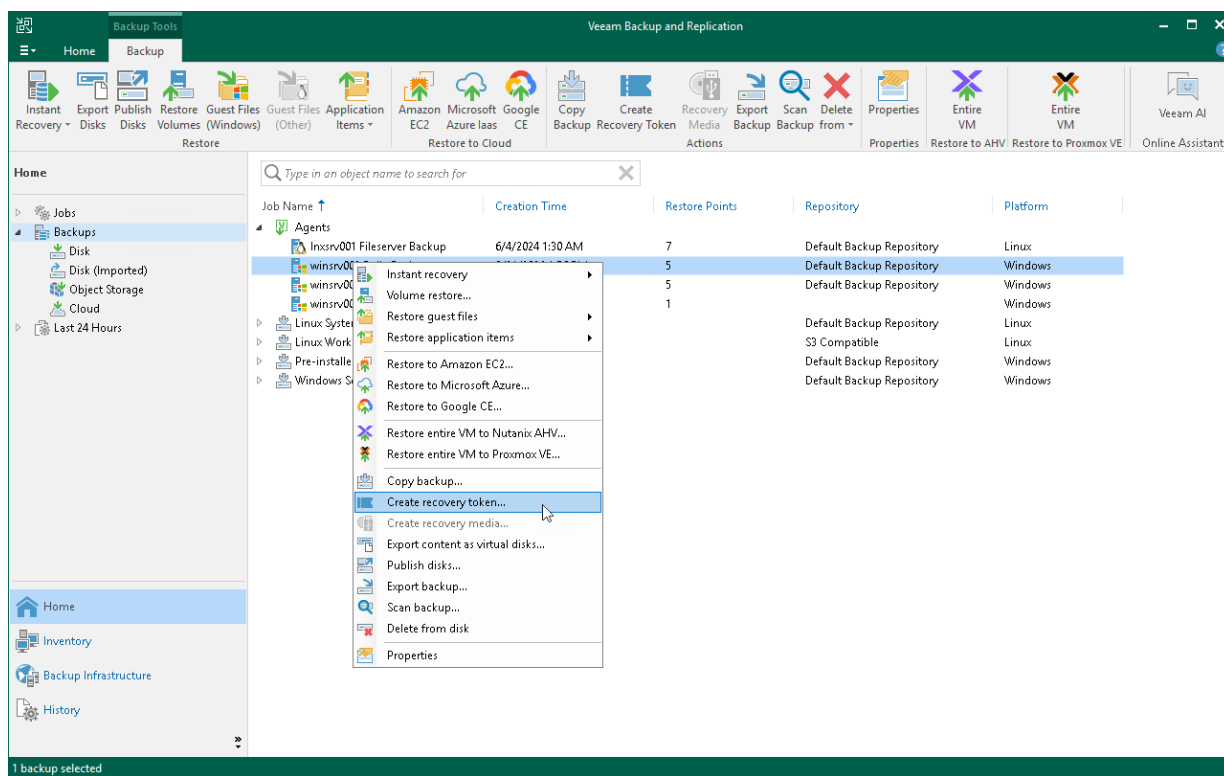
You can create a recovery token for several backups. To do this, press and hold [Ctrl], select multiple backups, right-click one of the selected backups and select **Create recovery token**.

4. In the **Create Recovery Token** window, click **Create**.

You can modify the existing recovery token using the PowerShell console. To learn more, see the [Working with Tokens](#) section in the Veeam PowerShell Reference.

## TIP

Alternatively, you can get access to the backup using user credentials. To learn more, see [Backup Server Settings](#).



# Copying Veeam Agent Backups

Veeam Backup & Replication offers the copying backup functionality that can be helpful if you want to copy backups of a backup job to another repository, local or shared folder.

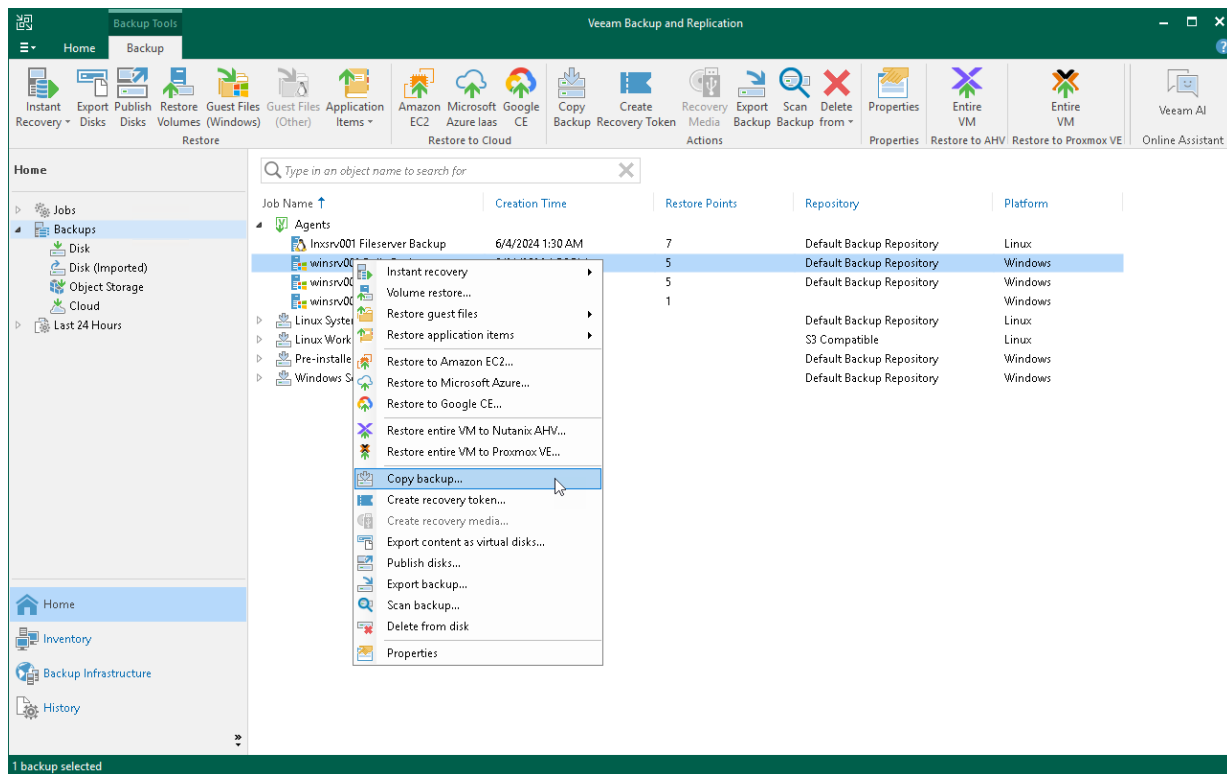
When Veeam Backup & Replication performs the copy operation, it disables the job, copies files to the target location and then enables the job. After the copy operation finishes, the copied backups are shown in a node with the **(Exported)** postfix in the inventory pane.

## NOTE

You cannot copy backups to/from object storage repositories.

To copy a backup, do the following:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. In the working area, right-click the backup and select **Copy backup**.



# Removing Veeam Agent Backups

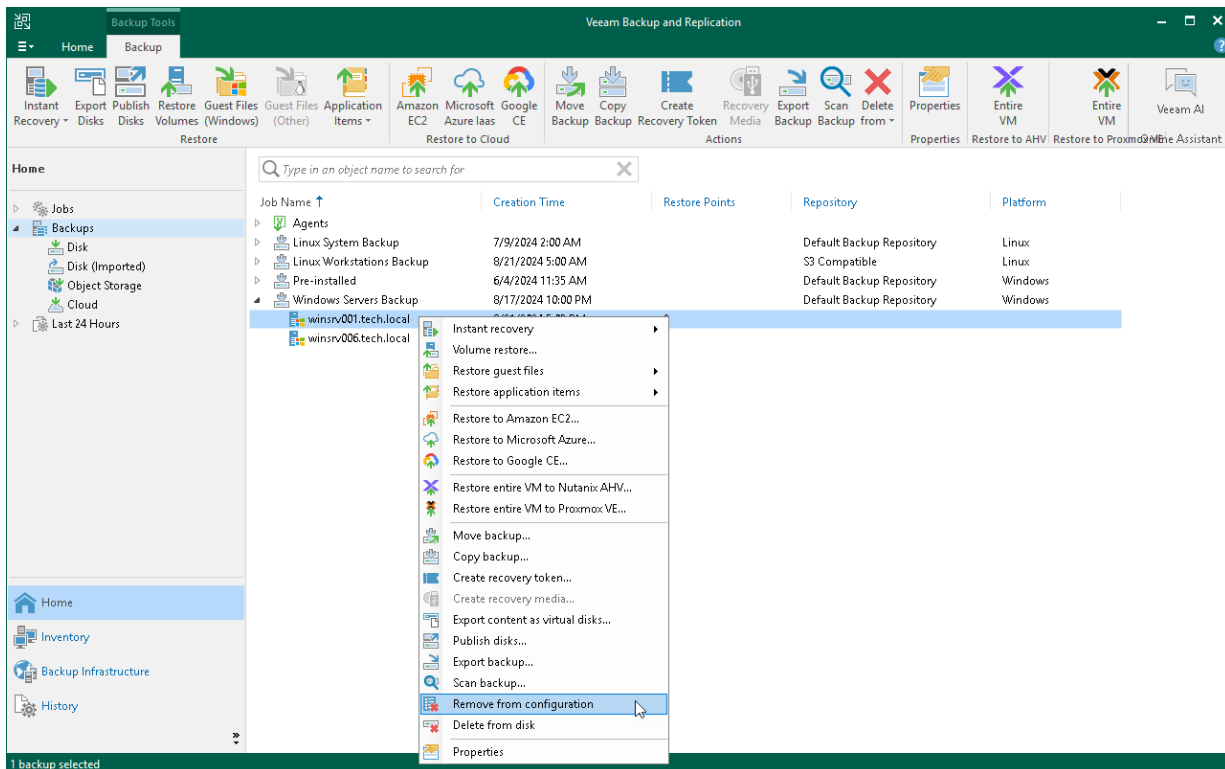
If you want to remove records about Veeam Agent backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation. When you remove a Veeam Agent backup from configuration, the actual backup files remain in the backup repository. You can import the backup to the Veeam Backup & Replication at any time later and restore data from it.

## IMPORTANT

Consider the following:

- Removing backups from configuration is designed for experienced users only. Consider using the [Delete from disk](#) operation instead.
- Do not remove a Veeam Agent backup from configuration if Veeam Agent is set up to use the backup cache and the backup cache contains one or several restore points that are not uploaded to the target location yet. If you remove such a backup and then import it in the Veeam Backup & Replication console, the backup will receive the new ID in the configuration database. As a result, Veeam Agent will become unable to upload restore points from the backup cache to the target location and to create new restore points in the backup cache. To continue creating backups in the Veeam backup repository, you will need to delete restore points from the backup cache and run the backup job to create a new restore point in the backup repository.

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. Press and hold the [Ctrl] key, select the backup, right-click the backup and select **Remove from configuration**.

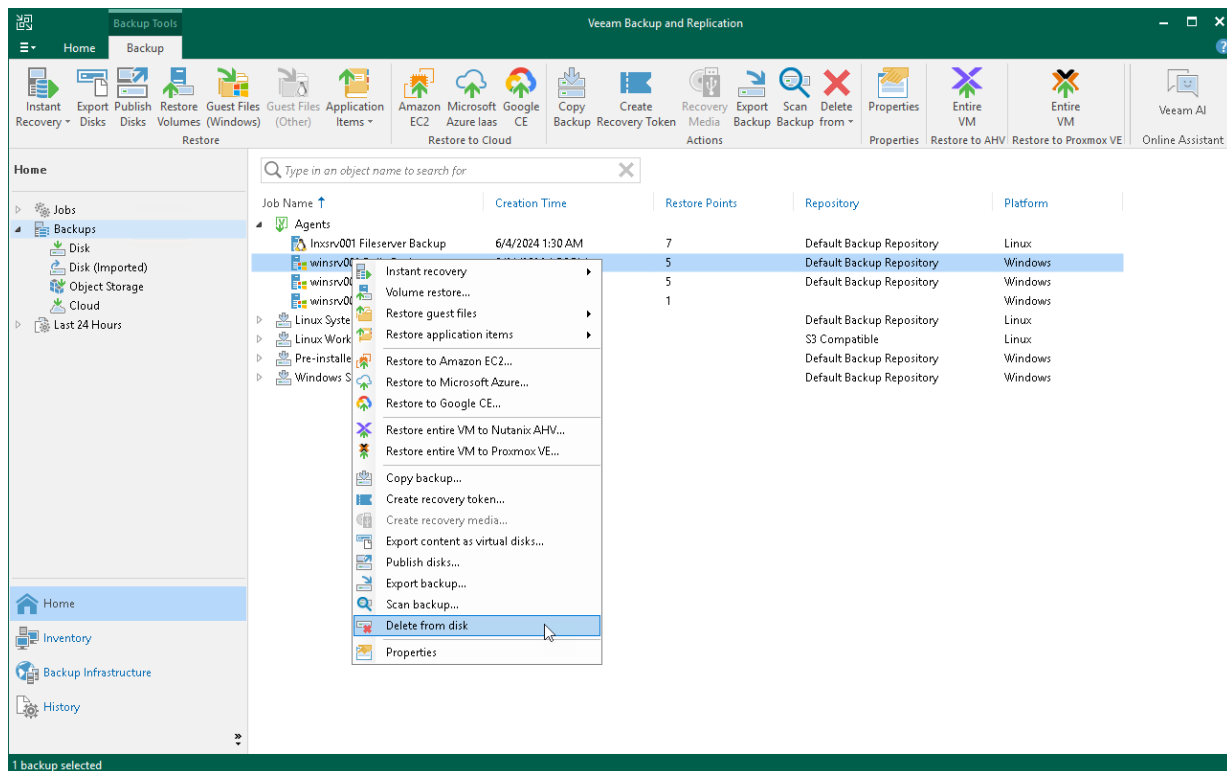


# Deleting Veeam Agent Backups from Disk

If you want to delete records about backups from the Veeam Backup & Replication console and configuration database and, additionally, delete backup files from the backup repository, you can use the **Delete from disk** operation.

To remove a Veeam Agent backup from the backup repository:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. Select the necessary computer backup and click **Delete from > Disk** on the ribbon or right-click the computer and select **Delete from disk**.



# Configuring Global Settings

Global settings configured on the Veeam backup server apply to Veeam Agent backup jobs as well. You can:

- Configure network throttling settings so that Veeam Agent backup job does not consume all network resources. To learn more, see the [Specifying I/O Settings](#) topic in the Veeam Backup & Replication User Guide.
- Configure the following global notification settings to get alerted about the Veeam Agent backup job results:
  - Email notifications. Veeam Agent for Microsoft Windows sends email notifications on every type of backup tasks, such as backup job sessions started automatically by schedule, backup job sessions started from the command line and ad-hoc backup tasks. To learn more, see the [Specifying Email Notification Settings](#) section in the Veeam Backup & Replication User Guide.
  - SNMP notifications. To learn more, see the [Specifying SNMP Settings](#) section in the Veeam Backup & Replication User Guide.

# Assigning Roles to Users

User roles configured on the Veeam backup server apply to Veeam Agent backup jobs as well.

To learn more, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.



# Automating Veeam Agent for Windows Operations

You can automate Veeam Agent for Microsoft Windows operations with Veeam Agent Configurator.

Veeam Agent Configurator is a tool that provides a command-line interface for Veeam Agent for Microsoft Windows. With Veeam Agent Configurator, you can perform data protection and administrative operations for Veeam Agent for Microsoft Windows from the command line, create custom scripts or integrate Veeam Agent for Microsoft Windows with third-party applications.

Veeam Agent Configurator is available in the Workstation and Server editions of Veeam Agent for Microsoft Windows.

Veeam Agent Configurator comes with Veeam Agent for Microsoft Windows. The `Veeam.Agent.Configurator.exe` file is placed in the product folder on the computer protected with Veeam Agent for Microsoft Windows, by default, `C:\Program Files\Veeam\Endpoint Backup`.

For more information, see the [Veeam Agent Configurator Reference](#).

# Appendix A. Veeam Agent Events

Veeam Agent for Microsoft Windows logs its events to event logs on the computer where the product is installed. Events can be used for monitoring the backup job activity and alerting about the backup status.

The table below lists all events logged by Veeam Agent for Microsoft Windows.

Event ID	Name	Description	Event Log	Source	Severity
110	Backup Job Started	Veeam Agent <jobname> has been started [by user <username>].	Veeam Agent	Veeam Agent	Information
115	Backup Job Resumed	Backup Job <jobname> has been resumed.	Veeam Agent	Veeam Agent	Information
190	Backup Job Finished	Veeam Agent <jobname> finished with <job status>. Job details: <additional information about the job results> <sup>1</sup> .	Veeam Agent	Veeam Agent	Information Warning Error
191	Backup Job Retry	Veeam Agent <jobname> finished with Error and will be retried. Job details: <additional information about the job results> <sup>1</sup> .	Veeam Agent	Veeam Agent	Warning
195	Synchronization Finished	Synchronization for cached restore points finished with <job status>. <Additional information about synchronized restore points>.	Veeam Agent	Veeam Agent	Information Warning Error
196	Destination Changed	Backup job destination has been switched from <target> to Backup Cache.	Veeam Agent	Veeam Agent	Information
197	Backup Cache Deleted	Backup cache has been deleted by <username>.	Veeam Agent	Veeam Agent	Information

Event ID	Name	Description	Event Log	Source	Severity
1074	Computer shut down <sup>2</sup>	The process C:\Windows\system32\Shutdown.exe (<jobname>) has initiated the shutdown of computer <jobname> on behalf of user NT AUTHORITY\SYSTEM for the following reason: No title for this reason could be found Reason Code: 0x800000ff Shutdown Type: shutdown Comment: Computer was shut down after successful backup by Veeam Agent for Microsoft Windows.	System	User32	Information
4010	License Installed	License key for Veeam Agent for Windows has been installed.	Veeam Agent	Veeam Agent	Information
4020	License Expiring	License key for Veeam Agent is about to expire in <N> days.	Veeam Agent	Veeam Agent	Warning
4030	License Expired	License key for Veeam Agent has expired.	Veeam Agent	Veeam Agent	Error
4025	License in Grace Period	Your license has expired and needs to be renewed. Veeam Agent will continue to operate for the duration of the grace period.	Veeam Agent	Veeam Agent	Error
4040	License Support Expiring	Support contract for Veeam Agent is about to expire in <N> days.	Veeam Agent	Veeam Agent	Warning
4050	License Support Expired	Support contract for Veeam Agent has expired. Contact Veeam sales representative to renew your support contract.	Veeam Agent	Veeam Agent	Error
4060	Product Edition Changed	Veeam Agent for Windows edition has been changed from <previousedition> to <currentedition>.	Veeam Agent	Veeam Agent	Information
10010	Restore Point Created	'<computername>' restore point has been created.	Veeam Agent	Veeam Agent	Information

Event ID	Name	Description	Event Log	Source	Severity
10050	Restore Point Removed	Restore point for '<computername>' has been removed according to the configured retention policy.	Veeam Agent	Veeam Agent	Information
23010	Backup Job Created	The <jobname> backup job has been created.	Veeam Agent	Veeam Agent	Information
23050	Backup Job Modified	The <jobname> backup job has been modified.	Veeam Agent	Veeam Agent	Information
23090	Backup Job Deleted	EndpointBackup <jobname> has been deleted.	Veeam Agent	Veeam Agent	Information
23051	Agent Modified	Veeam Agent option <optionname> has been changed. <sup>3</sup>	Veeam Agent	Veeam Agent	Information
23501	Agent Updated	Veeam Agent has been successfully updated.	Veeam Agent	Veeam Agent	Information
		Failed to update Veeam Agent.			Error
23110	Backup Mode Changed	The <jobname> backup mode has been changed from <previousmode> to <currentmode>.	Veeam Agent	Veeam Agent	Information
23120	Backup Source Updated	The <jobname> backup job source objects have been updated.	Veeam Agent	Veeam Agent	Information
26010	USB Device Ejected	Target USB device has been successfully ejected.	Veeam Agent	Veeam Agent	Information
178	Managed mode Enabled	Veeam Agent has been switched to managed mode. <sup>4</sup>	Veeam Agent	Veeam Agent	Information
179	Managed mode Disabled	Veeam Agent has been switched to free mode. <sup>4</sup>	Veeam Agent	Veeam Agent	Information
201	Read-only mode Enabled	Read only UI access has been enabled. <sup>4</sup>	Veeam Agent	Veeam Agent	Information
202	Read-only mode Disabled	Read only UI access has been disabled. <sup>4</sup>	Veeam Agent	Veeam Agent	Information

<sup>1</sup> Job details contain information about the reason for completing the job with the Warning or Error status.

<sup>2</sup> The event is triggered if the user has instructed Veeam Agent for Microsoft Windows to shut down the computer on successful backup.

<sup>3</sup> The event is triggered if the user has changed any Veeam Agent for Microsoft Windows setting other than backup job settings.

<sup>4</sup> The event can be triggered if Veeam Agent for Microsoft Windows is managed by a Remote Monitoring and Management platform, for example, LabTech.

# Appendix B. Moving Veeam Agent Backups

A Veeam Agent backup job starts a new backup chain when you change a target location in the backup job settings. If you do not want the backup job to start a new backup chain, you can map a backup job to the previously made backups and continue the existing backup chain. For more information about backup chains, see [Backup Chain](#).

This appendix describes the following cases of moving Veeam Agent backups:

- [Move a Veeam Agent backup to the Veeam backup repository](#).  
In this section, scenarios describe how to move a locally stored backup to a Veeam backup repository and how to move a backup from one Veeam backup repository to another and continue the existing backup chain.
- [Move a Veeam Agent backup to the Veeam Cloud Connect repository](#).  
In this section, scenarios describe how to move a locally stored backup to a backup repository used as a cloud repository and continue the existing backup chain.

## TIP

This appendix describes how to move backups created by Veeam Agent operating in the standalone mode. To learn how to move backups created by Veeam Agent operating in the managed mode, see the following topics:

- If Veeam Agent is managed by Veeam Backup & Replication, see the [Moving Backup](#) section in the Veeam Agent Management Guide.
- If Veeam Agent is managed by Veeam Service Provider Console, see the [Moving Veeam Agent Backups](#) section in the Guide for Service Providers.

For more information about Veeam Agent operation modes, see [Standalone and Managed Operation Modes](#).

# Moving Veeam Agent Backups to Veeam Backup Repository

This topic describes how to place existing Veeam Agent backups in a Veeam backup repository so that the Veeam Agent backup job continues the existing backup chain.

This topic covers the following scenarios:

- [Move a backup stored locally to the Veeam backup repository.](#)
- [Replace the initial Veeam backup repository with another Veeam backup repository.](#)

## NOTE

Keep in mind that the following users participate in implementation of both scenarios:

- Veeam Agent user performs all actions in the Veeam Agent control panel.
- Veeam Backup Administrator performs all actions in the Veeam Backup & Replication console.

## Moving Locally Stored Backup to Veeam Backup Repository

In this scenario, a backup job is targeted at a local folder. The backup job ran at least once.

This scenario can be helpful when you need to back up a large amount of data to the Veeam backup repository but the network connection is slow. In this case, you can create a backup locally and move it to the Veeam backup repository.

To move locally stored backup files to the Veeam backup repository, the Veeam Backup Administrator must perform the following steps:

1. Browse to the folder where Veeam backup repository stores backup files. For example, `C:\Backup`.  
For more information, see the [Configure Backup Repository Settings](#) section in the Veeam Backup & Replication User Guide.
2. In the `C:\Backup` folder, create a new folder in the following format:  
`C:\Backup\<domain_name>_<user_name>\<folder_name>`, where:
  - `<domain_name>` – domain name that Veeam Agent uses to connect to the Veeam backup repository.
  - `<user_name>` – user name that Veeam Agent uses to connect to the Veeam backup repository.
  - `<folder_name>` – name of the target folder to store backups. You can use the name of the original backup job or specify a new name for the folder.For example: `C:\Backup\TECH_Administrator\System_Backup`.
3. Using external tools, move all backup files to the folder created at the step 2.
4. Rescan the Veeam backup repository. To learn how to rescan the Veeam backup repository, see the [Rescanning Backup Repositories](#) section in the Veeam Backup & Replication User Guide.  
As a result, the moved backup will appear under the **Backups > Disk (Imported)** node in the inventory pane.

## NOTE

If the backup file is encrypted, you must decrypt the backup before proceeding to the next step. To learn more, see the [Decrypting Data with Password](#) section in the Veeam Backup & Replication User Guide.

After the Veeam Backup Administrator moved the backup files to the target folder, the Veeam Agent user must perform the following steps:

1. Edit the Veeam Agent backup job:
  - a. At the **Name** step of the wizard, make sure that the name of the backup job is the same as the name of the folder you created at the step 2.
  - b. At the **Destination** step of the wizard, target the backup job at the Veeam backup repository.
  - c. At the **Backup Server** step of the wizard, select a backup repository where you want to store backups and specify settings for the Veeam backup server that manages the target backup repository. For more information, see [Veeam Backup Repository Settings](#).
  - d. At the **Backup Repository** step of the wizard, click the **Map backup** link and select the moved backup file.
  - e. At the **Summary** step of the wizard, make sure that the *Run the job when I click Finish* check box is cleared and click **Finish** to save changes. As a result, the backup will be moved from the **Backups > Disk (Imported)** node to the **Backups > Disk** node in the inventory pane.
2. Run the Veeam Agent backup job. The backup job will continue the backup chain for the full backup that was moved to the Veeam backup repository.

## Replacing Initial Veeam Backup Repository with Another Veeam Backup Repository

In this scenario, the backup job is targeted at a Veeam backup repository and you want to replace the initial repository with another Veeam backup repository.

This scenario can be helpful if you want to move backed up data to a repository that has more storage capacity.

To replace the Veeam Backup & Replication repository, the Veeam Backup Administrator must perform the following steps:

1. Delete the existing Veeam Agent backup job from the Veeam Backup & Replication console. For more information, see [Deleting Veeam Agent Backup Jobs](#).

## NOTE

Deleting the job from the Veeam Backup & Replication console does not remove the backup job from the Veeam Agent control panel.

2. Delete configuration database records about backups created by the Veeam Agent backup job. For more information, see [Removing Veeam Agent Backups](#).
3. Browse to the folder where the target Veeam backup repository stores backup files. For example, `C:\Backup`.

For more information, see the [Configure Backup Repository Settings](#) section in the Veeam Backup & Replication User Guide.



4. In the `C:\Backup` folder, create a new folder in the following format:

`C:\Backup\<domain_name>_<user_name>\<folder_name>`, where:

- o `<domain_name>` – domain name that Veeam Agent uses to connect to the Veeam backup repository.
- o `<user_name>` – user name that Veeam Agent uses to connect to the Veeam backup repository.
- o `<folder_name>` – name of the target folder to store backups. You can use the name of the original backup job or specify a new name for the folder.

For example: `C:\Backup\TECH_Administrator\System Backup`

5. Using external tools, move all backup files to the folder created at the step 4.
6. Rescan the Veeam backup repository. To learn more, see the [Rescanning Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

As a result, the moved backup will appear under the **Backups > Disk (Imported)** node in the inventory pane.

#### NOTE

If the backup file is encrypted, you must decrypt the backup before proceeding to the next step. To learn more, see the [Decrypting Data with Password](#) section in the Veeam Backup & Replication User Guide.

After the Veeam Backup Administrator moved the backup files to the target folder, the Veeam Agent user must perform the following steps:

1. Edit the Veeam Agent backup job:
  - a. At the **Name** step of the wizard, make sure that the name of the backup job is the same as the name of the folder you created at the step 4.
  - b. At the **Backup Server** step of the wizard, select a new backup repository and specify settings for the Veeam backup server that manages the target backup repository. For more information, see [Veeam Backup Repository Settings](#).
  - c. At the **Backup Repository** step of the wizard, click the **Map backup** link and select the moved backup file.
  - d. At the **Summary** step of the wizard, make sure that the *Run the job when I click Finish* check box is cleared and click **Finish** to save changes. As a result, the backup will be moved from the **Backups > Disk (Imported)** node to the **Backups > Disk** node in the inventory pane.
2. Run the Veeam Agent backup job.

# Moving Veeam Agent Backups to Veeam Cloud Connect Repository

This topic describes how to place existing Veeam Agent backups in a cloud repository in the Veeam Cloud Connect infrastructure so that the Veeam Agent backup job continues the existing backup chain. For more information about Veeam Cloud Connect, see the [Veeam Cloud Connect Guide](#).

The scenarios described in this section can be helpful when you need to backup a large amount of data to the Veeam Cloud Connect repository but the network connection is slow. In this case, you can create a backup locally and move it to the Veeam Cloud Connect repository.

This topic covers the following scenarios:

- [Move a backup to a backup repository used as a cloud repository](#)
- [Move a backup to a scale-out backup repository used as a cloud repository](#)

## NOTE

Keep in mind that the following users participate in implementation of both scenarios:

- Service provider moves backup files to the target repository.
- Veeam Backup Administrator rescans the target repository in the Veeam Backup & Replication console.
- Veeam Agent user performs actions in the Veeam Agent control panel.

## Moving Backup to Backup Repository

Use this scenario if you want to move a locally stored backup to a backup repository used as a cloud repository. For information about backup repositories, see the [Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

To move locally stored backup files to the cloud repository, the service provider must perform the following steps:

1. Browse to the folder where the cloud repository stores backup files. For example, `C:\Backup`.

For more information, see the [Configure Backup Repository Settings](#) section in the Veeam Backup & Replication User Guide.

2. In the `C:\Backup` folder, create a new folder in the following format:

- [If you work with tenant accounts] `C:\Backup\<tenant_name>\<folder_name>`
- [If you work with subtenant accounts]  
`C:\Backup\<tenant_name>\Users\<subtenant_name>\<folder_name>`,

where:

- `<tenant_name>` — name of the tenant account that you use to connect to the service provider. For more information, see the [Tenant Account Credentials](#) section in the Veeam Cloud Connect Guide.
- `<subtenant_name>` — name of the subtenant account that you use to connect to the service provider. For more information, see the [Managing Subtenant Accounts on SP Side](#) section in the Veeam Cloud Connect Guide.

- o `<folder_name>` – name of the target folder to store backups. You can use the name of the original backup job or specify a new name for the folder.

For example, `C:\Backup\ABC_Company\System_Backup` or  
`C:\Backup\ABC_Company\Users\John_Smith\System_Backup`.

3. Using external tools, move all backup files to the folder created at the step 2.

After the service provider moved the backup files to the backup folder, the Veeam Backup Administrator must rescan the cloud repository. To learn how to rescan a repository, see the [Rescanning Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

The information about the added backup will be displayed in the rescan job session window.

After the Veeam Backup Administrator rescanned the backup repository, the Veeam Agent user must perform the following steps:

1. Edit the Veeam Agent backup job:
  - a. At the **Destination** step of the wizard, target the backup job at the Veeam Cloud Connect Repository.
  - b. At the **Service Provider** step of the wizard, specify settings for the cloud gateway to connect to the service provider. For more information, see [Specifying Service Provider Settings](#).
  - c. At the **Credentials** step of the wizard, specify settings for the tenant or subtenant account that you want to use to connect to the service provider. For more information, see [Specifying User Account Settings](#).
  - d. At the **Backup Resources** step of the wizard, click the **Map backup** link and select the moved backup.
  - e. At the **Summary** step of the wizard, click **Finish** to save changes.
2. Run the Veeam Agent backup job. The backup job will continue the backup chain for the full backup that was moved to the cloud repository.

## Moving Backup to Scale-Out Backup Repository

Use this scenario if you want to move a locally stored backup to a scale-out backup repository used as a cloud repository. For information about scale-out backup repositories, see the [Scale-Out Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

### IMPORTANT

Before you perform the move, make sure that the backup job name, the target backup folder name, the names of VBM files and paths to VBM files meet the following requirements:

- The names contain only allowed alphanumeric characters: a-z, A-Z, 0-9.
- The names contain only allowed special characters: \_ - . + = @ ^.
- The names of VBM files and paths to VBM files do not contain spaces. If the names or the paths contain spaces, replace them with underscores.

To move locally stored backup files to the cloud repository, the service provider must perform the following steps:

1. Browse to the folder where the cloud repository stores backup files. For example, `C:\Backup`.

For more information, see the [Configure Backup Repository Settings](#) section in the Veeam Backup & Replication User Guide.

2. In the `C:\Backup` folder, create new folders for each extent of the scale-out backup repository. The folders must be in the following format:
  - [If you work with tenant accounts]  
`C:\Backup\<extent_name>\<tenant_name>\<folder_name>`
  - [If you work with subtenant accounts]  
`C:\Backup\<extent_name>\<tenant_name>\Users\<subtenant_name>\<folder_name>`,

where:

- `<extent_name>` – name of the backup repository you use as an extent. For more information, see the [Add Performance Extents](#) section in the Veeam Backup & Replication User Guide.
- `<tenant_name>` – name of the tenant account that you use to connect to the service provider. For more information, see the [Tenant Account Credentials](#) section in the Veeam Cloud Connect Guide.
- `<subtenant_name>` – name of the subtenant account that you use to connect to the service provider. For more information, see the [Managing Subtenant Accounts on SP Side](#) section in the Veeam Cloud Connect Guide.
- `<folder_name>` – name of the target folder to store backups. You can use the name of the original backup job or specify a new name for the folder.

For example, `C:\Backup\Extent_1\ABC_Company\System_Backup` or  
`C:\Backup\Extent_1\ABC_Company\Users\John_Smith\System_Backup`.

3. Using external tools, copy the VBM file to both folders created at the step 2.
4. Using external tools, copy the VBK and VIB files to the folders created at the step 2 according to the placement policy specified for the backup repository. For more information, see the [Backup File Placement for Performance Tier](#) section in the Veeam Backup & Replication User Guide.

After the service provider moved the backup files to the backup folder, the Veeam Backup Administrator must rescan the cloud repository. To learn how to rescan the repository, see the [Rescanning Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

The information about the added backup will be displayed in the rescan job session window.

After the Veeam Backup Administrator rescanned the backup repository, the Veeam Agent user must perform the following steps:

1. Edit the Veeam Agent backup job:
  - a. At the **Destination** step of the wizard, target the backup job at the Veeam Cloud Connect Repository.
  - b. At the **Service Provider** step of the wizard, specify settings for the cloud gateway to connect to the service provider. For more information, see [Specifying Service Provider Settings](#).
  - c. At the **Credentials** step of the wizard, specify settings for the tenant or subtenant account that you want to use to connect to the service provider. For more information, see [Specifying User Account Settings](#).
  - d. At the **Backup Resources** step of the wizard, click the **Map backup** link and select the moved backup.
  - e. At the **Summary** step of the wizard, click **Finish** to save changes.
2. Run the Veeam Agent backup job. The backup job will continue the backup chain for the full backup that was moved to the cloud repository.