



How to Backup & Restore SQL Server Databases with Veeam

Veeam Backup & Replication

Version 9.5

August, 2017



© 2017 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.

Important!

Please read the End User Software License Agreement before using the accompanying software program(s). Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

CONTENTS

CONTENTS	3
CONTACTING VEEAM SOFTWARE	5
OVERVIEW	6
WHAT'S NEW IN V9.5	7
PREPARING FOR SQL SERVER BACKUP AND RECOVERY	8
KNOW YOUR INFRASTRUCTURE RESOURCES	8
DETECT DATABASE LOGGING AND RECOVERY MODEL.....	9
DECIDE ON BACKUP AND RECOVERY TOOL.....	11
UNDERSTAND THE PROCESSES	14
How It Works: SQL Server and Transaction Log Backup	14
How It Works: SQL Server Database Restore	20
How It Works: Mount Operations	23
ADDRESS RECOVERY REQUIREMENTS	24
EXAMINE CONSIDERATIONS AND LIMITATIONS	25
PLAN FOR RECOVERY VERIFICATION.....	26
SETTING UP YOUR SQL BACKUP JOBS	27
OBJECTS TO PROCESS	27
Excluding Databases from Log Backup Process.....	28
APPLICATION-AWARE PROCESSING	29
TRANSACTION LOG HANDLING.....	31
LOG SHIPPING SERVER.....	34
SCHEDULING.....	35
SUPPORT FOR ALWAYSON AVAILABILITY GROUPS	36
CONFIGURING REQUIRED PORTS	38
MANAGING YOUR SQL BACKUP JOBS	39
STARTING, STOPPING, ENABLING AND DISABLING THE JOB	39
VIEWING JOB INFORMATION	40
RECONFIGURING JOBS WITH MICROSOFT SQL SERVER VMS.....	42
LOG BACKUP IMPORT.....	42
RESTORING SQL SERVER DATABASES	43
PREPARING FOR DATABASE RESTORE	44
Step 1: Check Required Ports	44
Step 2: Prepare Staging SQL Server	45
Step 3: Check User Account Permissions	47
Step 4: Examine Considerations and Limitations	49

Step 5: Extend Database Files from SQL Server Backup.....	49
FOLLOWING YOUR RESTORE SCENARIO	50
Restoring a Database To Current State	50
Restoring a Database to Specific Point in Time.....	54
Restoring a Database to State Before Selected Transaction.....	54
Restoring Database Schema and Data	56
APPENDIX. STATISTICS ON TRANSACTION LOG PROCESSING	60
REVISION HISTORY	63

CONTACTING VEEAM SOFTWARE

At Veeam Software we value the 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, please visit our Customer Center Portal at www.veeam.com/support.html 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 offices location, please visit www.veeam.com/contacts.html.

Online Support

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

- Full documentation set: www.veeam.com/documentation-guides-datasheets.html
- Community forum at forums.veeam.com

OVERVIEW

With database servers and other Tier 1 applications being virtualized, organizations need to ensure data **recoverability** to any point in time, minimizing impact on end users and supporting continuous business operation. Thus, backup and recovery strategies for highly transactional applications are of greater importance.

Veeam Backup & Replication offers you flexible options for backup and several approaches for restore of your virtualized SQL Servers and databases, including popular full restore of SQL Server virtual machine, file-level restore, and database restore using Veeam Explorer for Microsoft SQL Server. This document will describe the latter approach; to read about two others, refer to the Veeam Backup & Replication documentation at <https://www.veeam.com/documentation-guides-datasheets.html>.

This document is intended for backup administrators and other IT professionals involved in setting up SQL Server database backup and restore with Veeam solution.

Note:

SQL application items restore using U-AIR Microsoft SQL Recovery Wizard, described in the corresponding User Guide (https://www.veeam.com/veeam_backup_9_5_uair_wizard_user_guide_en_pg.pdf).

To properly configure a backup job that will create a transactionally-consistent backup and allow for future restore of your SQL Server databases, you should understand how the backup processes implemented by Veeam work.

Veeam Backup & Replication creates a backup of the whole SQL Server, including all SQL Server instances and DBs, which can be full, incremental or reverse incremental backup (depending on the selected backup method). As for native SQL means, they perform backup per database, creating full and differential backups. Remember that these native SQL backup types do not correspond to Veeam backup methods in any way. In this context, consider that Veeam Backup & Replication will do the following during application-aware image-level backup creation:

1. Create an image-level backup of your SQL Server VM, including all hosted instances and databases and their related settings, and store it to backup repository in Veeam's proprietary format: .VBK (full VM backup), .VIB (incremental), or .VRB (reversed incremental).
2. If transaction log handling option was selected, then Veeam Backup & Replication will obtain transaction log backup files (.BAK) created by native means and use Veeam log shipping server to store them to the location where VM backup is kept. These log backups will be stored in Veeam's proprietary format (.VLB).

These procedures, as well as Veeam log shipping server will be discussed later in this document.

Note:

To read more about backup methods and Veeam backup files, refer to the Veeam Backup & Replication User Guide at <https://www.veeam.com/documentation-guides-datasheets.html>. To read more about native backup means and Veeam approach since v8, you can refer to the <https://www.veeam.com/wp-sql-server-backup-with-veeam.html>.

WHAT'S NEW IN V9.5

The following new features and enhancements were implemented in this version Veeam Explorer for Microsoft SQL Server:

- Support for SQL Server 2016, including AlwaysOnGroups and system-versioned tables.
- Support for memory-optimized tables in SQL Server 2014 and *FileTable* tables in SQL Server 2012.
- UI enhancements, including ribbon menu, command-specific tabs, and support of display themes as in Veeam Backup & Replication console.

Starting with Veeam Backup & Replication 9.5 Update 2, you can also use Veeam Explorer for Microsoft SQL Server to restore from application-aware image backups created using Veeam Agent 2.0 for Microsoft Windows (*Server* Edition of the product is required). To learn more, see Edition Comparison (https://www.veeam.com/veeam_agent_windows_2_0_editions_comparison_ds.pdf) and Veeam Agent for Microsoft Windows User Guide (https://helpcenter.veeam.com/docs/agentforwindows/userguide/backup_job_vss_general.html).

PREPARING FOR SQL SERVER BACKUP AND RECOVERY

Before you start configuring SQL Server backup and recovery, you can use the following check list for preparations:

- Know Your Infrastructure Resources
- Detect Database Logging and Recovery Model
- Decide on Backup and Recovery Tool
- Understand the Processes
- Address Recovery Requirements
- Plan for Recovery Verification

Know Your Infrastructure Resources

When planning for your SQL Server backup and restore, examine your infrastructure and virtualized servers and answer the following questions:

1. ***What database servers are virtualized in your environment?***

Veeam Backup & Replication supports database restore for the following systems:

- Microsoft SQL Server 2016
- Microsoft SQL Server 2014
- Microsoft SQL Server 2012
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2008
- Microsoft SQL Server 2005 SP4

All editions of Microsoft SQL Server are supported; AlwaysOn Availability Groups are supported for Microsoft SQL Server 2012 and later.

2. ***What are your RTO & RPO for SQL Server databases?***

Frequency of the SQL Server backup depends on how much time you can afford to restore it. In general, the more often you back up, the shorter the restore time is.

3. ***How intensively your applications modify data in the database?***

This factor, together with required RTO may influence the backup window and job schedule, as well as space needed for storing transaction log backups.

4. ***What is your backup repository capacity?***

Repository should provide enough space to store the amount of data produced by database and transaction log backups. It is recommended that you estimate how much disk space a VM

backup and transaction log backup will occupy, and then plan for your repository accordingly. If you have Enterprise or Enterprise *Plus* Edition of Veeam Backup & Replication, you may use the scale-out backup repository.

Note:

Consider that in case of Enterprise license the scale-out repository can comprise up to 3 standard repositories, and Enterprise *Plus* – the unlimited number of standard repositories. To learn more about this feature, please refer to the Veeam Backup & Replication User Guide.

5. **Analyze how your backup infrastructure is organized and identify the locations of Veeam backup server, backup repository and production VMs.**

This will help you to decide on preferable machine for mount operation.

For mount operations, Veeam Explorer utilizes the corresponding service (Veeam Mount Service) that can run on Veeam backup server or on Veeam standalone console - both of them include Veeam Explorer in their setup. So, Veeam backup console can work as a mount server in the remote site, eliminating the need to deploy additional Veeam backup server in that site and minimizing traffic at restore. Consider the following recommendations:

- If repository and production (target) SQL Server VM are located in the same site with Veeam backup server, Veeam Explorer will utilize Mount Service running on Veeam backup server for mount operation.
- If repository and production (target) SQL Server VM are located in a remote site (separately from Veeam backup server), it can be reasonable to deploy Veeam Backup & Replication console in that remote site, and launch Veeam Explorer from that console, initiating mount operation locally in the site where console is running.

For more details on the mount operations, refer to the [How It Works: Mount Operations](#) section below.

Note:

If you plan to restore database items from an SQL Server VM running Microsoft Windows ReFS, consider that Veeam backup server or management console (that is, the machine where Veeam Explorer and the mount service are running) must be installed on the Microsoft Windows Server 2012 or later.

To restore from a server running Microsoft Windows ReFS 3.x, Veeam backup server or management console (that is, the machine where Veeam Explorer and the mount service are running) must be installed on the Microsoft Windows Server 2016.

Detect Database Logging and Recovery Model

SQL Server backup and restore processes go within the context of a database recovery model which defines the following:

- What types of backups and what restore scenarios will be supported for the database
- How the transaction log should be handled for that purpose

Microsoft SQL Server supports several logging and recovery models for its databases: *simple*, *full*, and *bulk-logged*.

- With **simple** recovery model specified for your database, you will be able to restore it only to the selected restore point; SQL Server will automatically truncate logs, and Veeam backup configuration will not affect any of them.
- With other models (**full** and **bulk-logged**), database transaction logs will not be truncated automatically by SQL Server, and Veeam offers different processing options for them.

To discover logging and recovery model for the database which you plan to backup and restore, use SQL Server Management Studio or contact your database administrator.

Note:

To read more about different logging and recovery models, refer to the [Recovery Models \(SQL Server\) MSDN article](#).

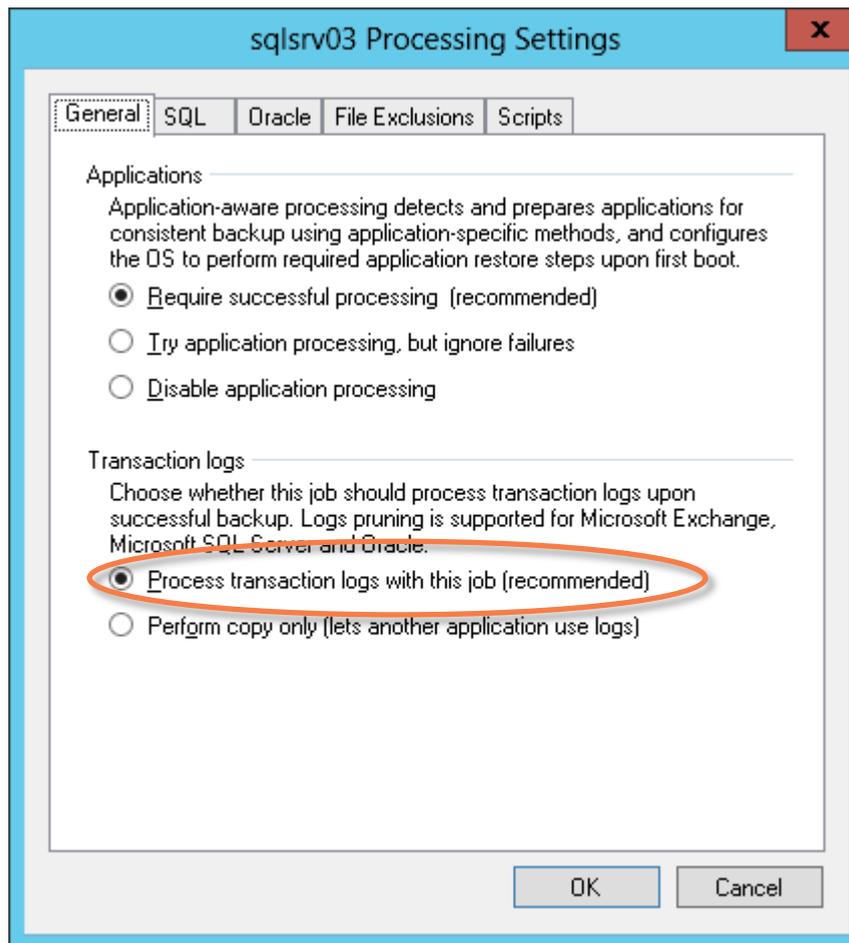
To learn how transaction log truncation works, refer to the [Transaction Log Truncation article](#).

Decide on Backup and Recovery Tool

You should also find out what tools are used to backup SQL Server VM hosting your database(s) and their transaction logs in your organization infrastructure.

Using Veeam

If you plan to use Veeam Backup & Replication for both image-level backup of SQL Server VM and database transaction log backup, then you will need to select the corresponding option (**Process transaction logs with this job**) when configuring VM processing settings in backup job. Job configuration will be discussed later in more details.



This option allows you to 'set and forget' an automated workflow for creating all the backups required for SQL Server database recovery using Veeam Backup & Replication. Thus, data can be quickly and easily recovered, following the scenario of your choice with Veeam tools.

With this option selected, you can configure SQL database log backup settings on the **SQL** tab, which becomes available to you.

Note: For other applications, this option triggers transaction log truncation.

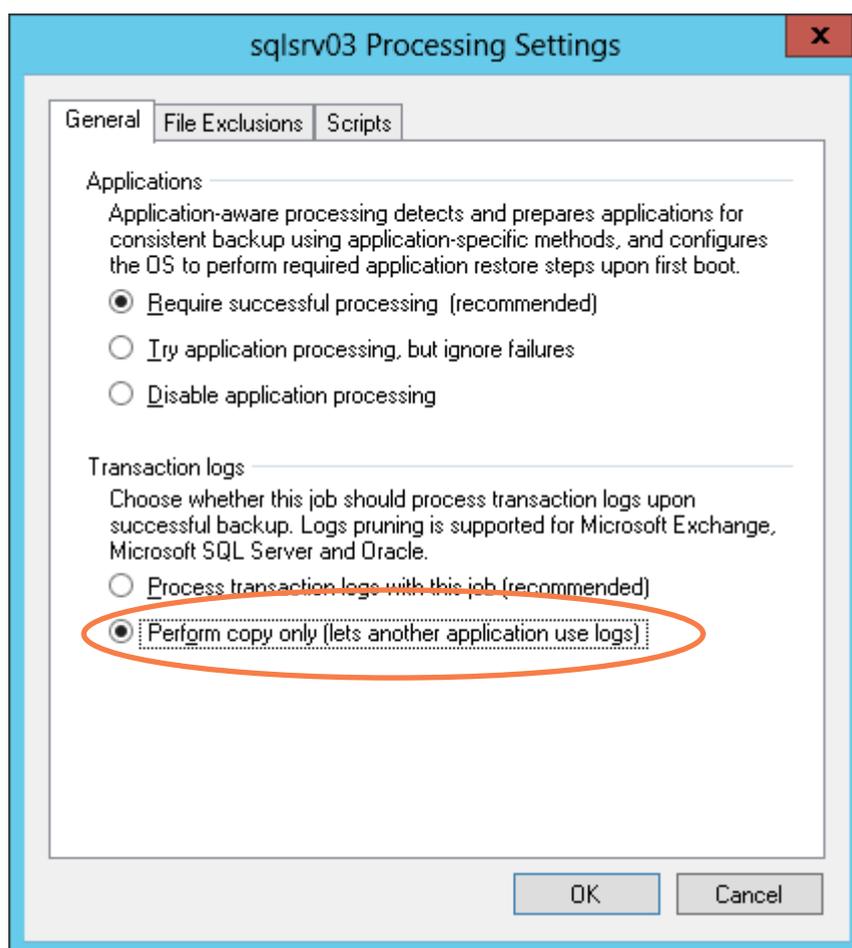
Using 3rd Party Tool

For large SQL Server deployments, frequent full database backups are not always acceptable since these operations would be very much time- and resource-taking. For this reason, organizations prefer to create full and differential database backups periodically, backing up transaction logs quite frequently. Therefore, to provide for proper restore, database administrators need to ensure they have the most recent full/differential backup and the associated sequence of transaction log backups at hand. When Veeam creates an image-level backup of SQL Server VM, using application-aware processing without special settings (that is, without COPY_ONLY flag), this is always treated by SQL as a full database backup.

Note: Per Microsoft recommendation, starting with Veeam Backup & Replication 9.0, a Veeam backup job registers a SQL Server database backup by adding the corresponding record in the **system.db** system table of each SQL Server database. This feature is known as database labeling.

As a result, this full backup becomes a new base for the differential backups that follow, jeopardizing the sequence (see [this illustrated article](#)).

Thus, if you plan to preserve the chain of full and differential database backups (created by native means or by 3rd party backup tool) but you also want Veeam to back up your SQL Server VM, it is recommended that you trigger the COPY_ONLY flag for Veeam. For that, select the **Perform copy only** option.



This option indicates that a chain of database backups and transaction logs is created by native means or by 3rd party tool, and instructs Veeam to keep this sequence untouched. For that, Veeam will back up the specified SQL Server VM using the COPY_ONLY option for snapshot creation. Transaction logs

will not be affected in any way, so your database administrator can handle them with native or 3rd party tools.

Note: With the **Perform copy only** option selected, the **SQL** and **Oracle** tabs intended for log handling settings will be hidden, with their parameters unavailable to user.

The COPY_ONLY method also allows Veeam to support Always-On Availability Groups, creating backups of passive replicas.

For detailed description of SQL Server backup chain, you can read the article by Michael K. Campbell at <http://sqlmag.com/blog/breaking-backup-chain-redux-or-eating-crow> and the recommendations by Paul Randal at http://www.sqlskills.com/blogs/paul/backup-with-copy_only-how-to-avoid-breaking-the-backup-chain/.

Also, refer to the Veeam Backup & Replication User Guide and to the white paper by Tibor Karaszi for information on SQL Server backup: <https://www.veeam.com/wp-sql-server-backup-with-veeam.html>.

Understand the Processes

To create an application-aware backup of the SQL Server VM that enables you to restore the databases to the certain point in time or to specific transaction, you should set up a backup job with special application processing options, and enable SQL database transaction log handling. This section explains how such a job works, and provides recommendations on its configuration.

Note: These explanations refer only to the case when transaction log backup is turned ON for SQL Server VM backup job; otherwise, application-specific options will be unavailable.

In This Section

- [How It Works: SQL Server and Transaction Log Backup](#)
- [How It Works: SQL Server Database Restore](#)
- [How It Works: Mount Operations](#)

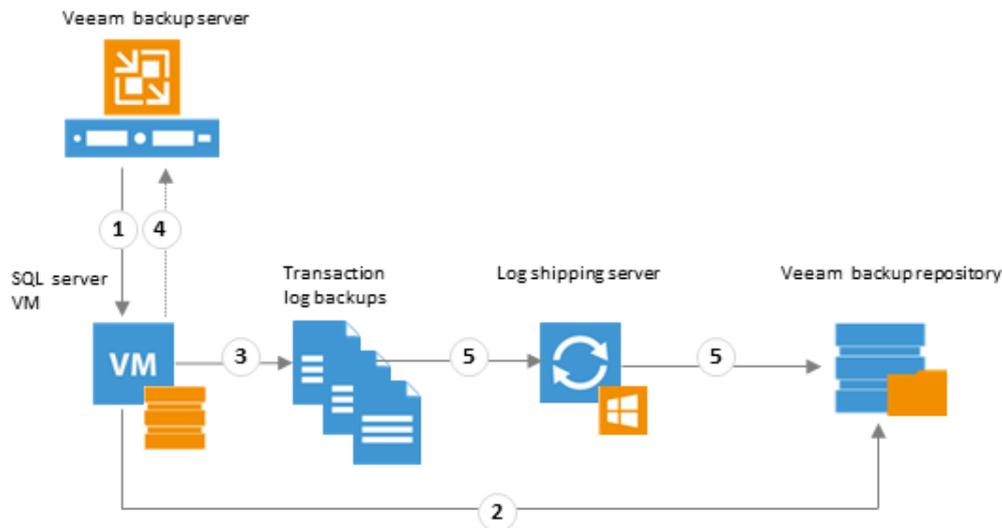
How It Works: SQL Server and Transaction Log Backup

If you plan to set up the SQL Server backup job to back up transaction logs, consider that it will comprise 2 jobs:

1. SQL Server image-level VM backup job ('parent') named *<job_name>*, for example, *Test Job*. This is the job you configure explicitly in the management console; its session starts on schedule or is started manually by user.
2. Transaction log backup job ('child') named with suffix: *<job_name> SQL Backup*, for example, *Test Job SQL Backup*. This job is created by Job Manager if it detects that 'parent' (VM backup) job is scheduled to back up at least one SQL Server with application-aware image processing switched on and transaction log backup enabled.

Transaction log backup job ('child') is triggered by VM backup ('parent') — this sequence ensures that VM (and database) restore point will be present when it comes to log replay. 'Child' job runs permanently in background, with the specified frequency of log shipping to repository. Its session data is kept in the Veeam Backup & Replication database and is displayed in the Veeam backup management console.

The whole process goes as follows:



Stages 1 and 2

When scheduled, Job Manager (Veeam component working on Veeam backup server) launches 'parent' job that creates an image-level backup of SQL Server VM and stores it to backup repository.

Stage 3

A new 'child' job session starts, and Veeam Backup & Replication installs a runtime component *Veeam Log Shipper Service* to the VM guest OS in order to support guest processing and log handing. This service runs during the 'child' job session and collects information about databases whose logs should be processed. The service also detects whether it is possible to store logs into the repository through direct access, or using a log shipping server. When 'child' job session ends, the service is stopped and removed from guest. Then a new session starts, and the service is installed again.

Transaction log backup is performed by SQL Server; this operation also includes log truncation so that the space can be reused. These log backups are stored as .BAK files in a temporary folder in the SQL Server VM guest file system.

Note:

Default location of the temporary folder is `%allusersprofile%\Veeam\Backup`. If it is necessary to use another location, you can specify it in the **SqlTempLogPath (STRING)** registry value that you can create under **HKLM\SOFTWARE\Veeam\Veeam Backup and Replication** registry key on your SQL Server VM (you may need to restart the Log shipper service on that VM for the change to take effect).

Stage 4

Every 15 minutes (default frequency) Job Manager component detects what databases currently exist on SQL Server, and maps this data with VM backup information kept in Veeam Backup & Replication database. This periodic mapping reveals the databases for which Veeam Backup & Replication must ship transaction logs to repository during this 15-min interval.

If there are any logs that for some reasons were not shipped to backup repository by Veeam during the previous interval(s), they will be also included in the processing list. To detect those remaining logs, Veeam enumerates .BAK files in the temporary folder.

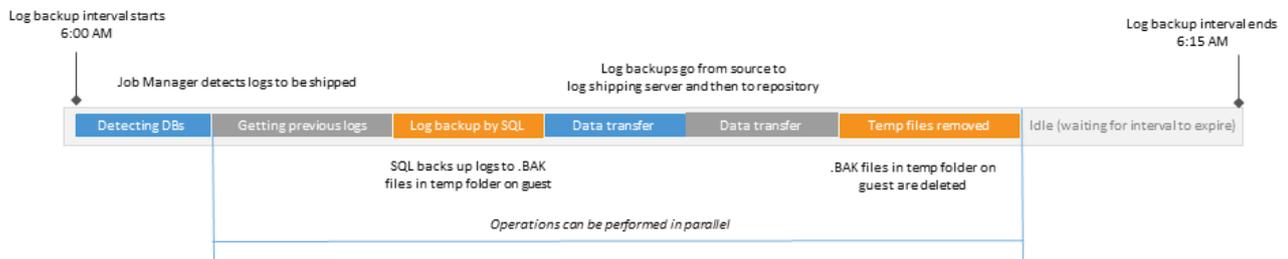
Stage 5

Transaction log backup files (.BAK) are transferred from temporary location on SQL Server to the target location in backup repository, either directly or via a log shipping server. Source-side Veeam transport service compresses log data to be transferred according to its built-in settings. On the repository side, data is compressed according to 'parent' job settings (see the "Data Compression and Deduplication" section of the User Guide for details).

As soon as data is copied to the target, transaction log backup files are deleted from the temporary folder on SQL Server.

Note: Transaction logs that for some reason were not processed during log backup interval remain in that temporary folder and are processed during the next log backup interval (see Stage 4).

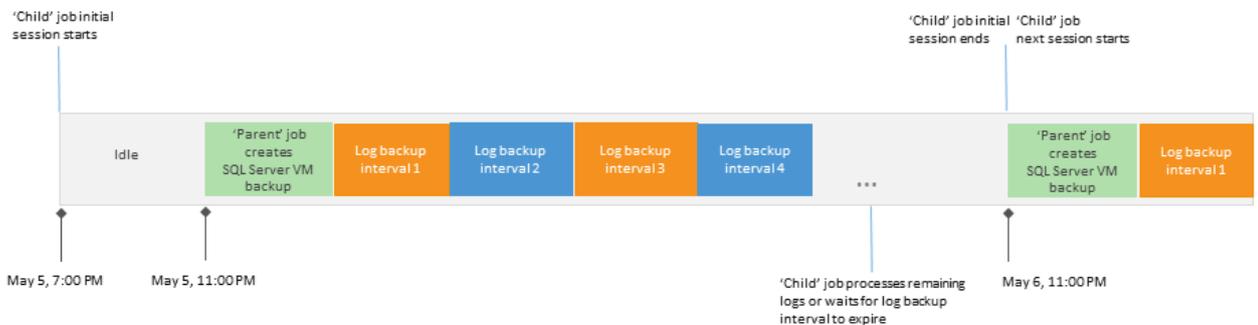
The figure below depicts in detail what goes on within a log backup interval for each SQL Server VM in the job (that has log backup option enabled).



Log Backup Sessions

A sequence of the 15-min (or custom) intervals between the runs of 'parent' job comprises a log backup ('child' job) session.

- Initial session starts at the moment when 'parent' job (image-level backup) schedule was enabled, then session starts with each start of the 'parent' job.
- Log backup ('child' job) continues working until the next start of the VM image-level backup ('parent'). At this point, 'child' job stops and then starts a new session, performing steps 1-5.
- Session ends just before the next run of the 'parent' job, and/or when this 'parent' job is disabled (using the menu command or by modifying job schedule).



Assume that VM image-level backup job was configured to run daily at 11:00 PM, starting on May, 5. The figure above illustrates the following course of action:

1. Initial 'child' job session starts as soon as 'parent' job schedule is enabled (no backups yet) – this happens at 7:00 PM on May, 5. The 'child' job remains in the *Idle* state, waiting for an image-level backup to arrive.
2. At 11:00 PM on May, 5 'parent' job runs on schedule, creating an image-level backup of SQL Server VM.
3. 'Child' job state is now *Working*, and a sequence of log backup intervals follows. If all log backups are successfully transferred to backup repository before the moment when a new session must start, the job will go into the *Idle* state and wait for remaining log backup interval(s) to expire.
4. At 11:00 PM on May, 6 next run of the 'parent' job takes place, and new 'child' job session starts.

Transaction log backup requires that at least one image-level backup of SQL Server VM is performed, so remember that:

- Transaction logs for newly appearing databases will be backed up only after the image-level backup of the corresponding SQL Server is performed.
- If a new log backup ('child') job session starts when no restore point has been yet created by image-level backup ('parent') job, and there are no previous logs to process, 'child' job will remain in the idle state, waiting for a new restore point to arrive.
- Transaction log backup will not take place after full SQL Server VM restore.

Backup Files and Retention

Backup Files

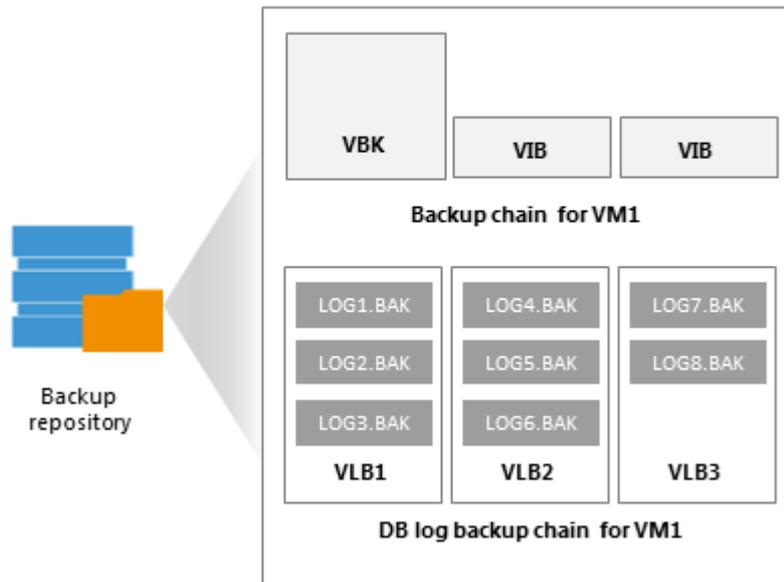
In the backup repository, logs are stored as .VLB files (Veeam proprietary format) co-located with corresponding SQL Server VM backups (.VBK/.VIB/.VRB files) in the repository folder (default is C:\backup\

At each start of the SQL Server backup job ('parent'), a new .VLB is created to store log backups in the repository:

- If the **Use per-VM backup files** option is selected for the repository, then Veeam will create a separate .VLB for each server processed by the job.
- If this option is cleared, then a single .VLB will be created for all servers processed by the job.

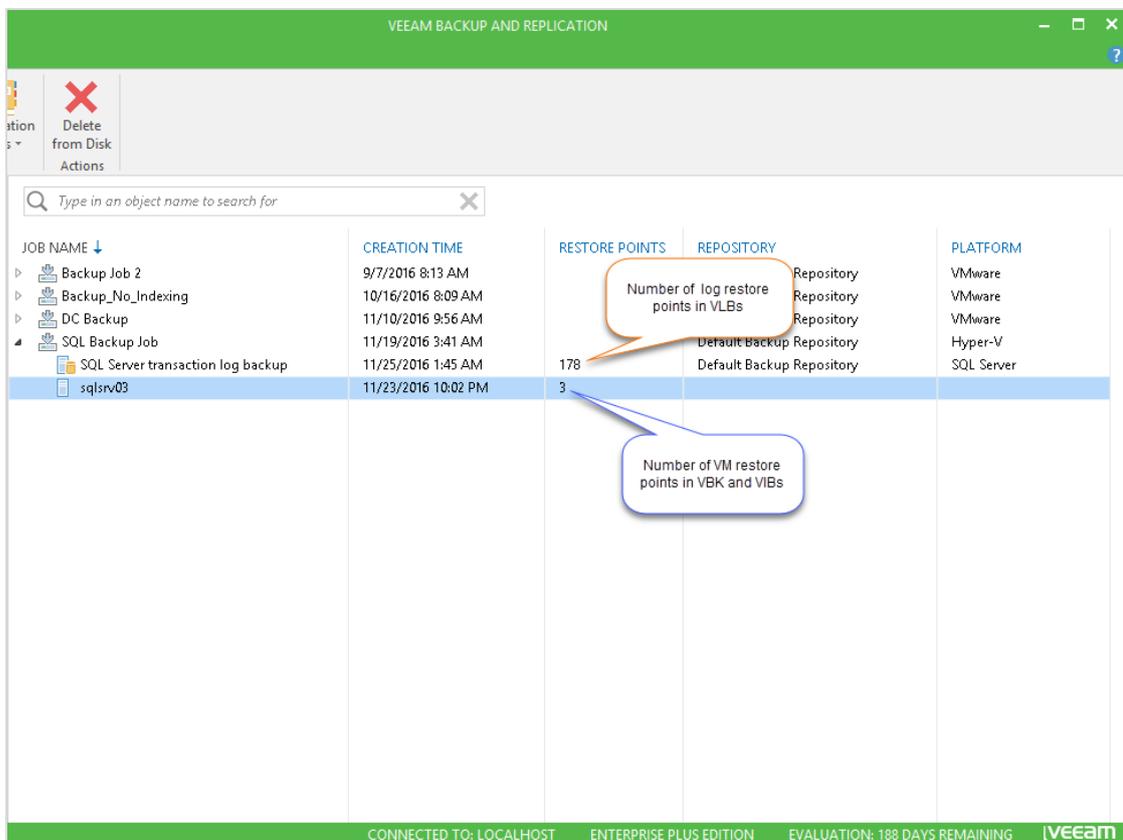
For example, if a job processes only one SQL Server, the repository will contain a number of .VLB files for it (a so-called chain).

As described in the section above, during database log backup ('child') job session, transaction log backup is performed by native means of the SQL Server and stored as .BAK file to a temporary folder on the SQL Server VM guest file system. Then Veeam copies .BAK file to the current .VLB in the repository. When the new 'parent' job session starts, another .VLB is created, and the .BAK files that appear after that will be stored there during the 'child' job session. The resulting chain of .VLBs will look like shown below, depicted for a single SQL Server VM1:



Total number of all LOG<N> .BAK files stored at the moment in all VLBs is reported as a **number of restore points for the 'child' job** that backs up database logs. So, in the example above, the log backup job for SQL Server VM1 has created 8 restore points by the moment.

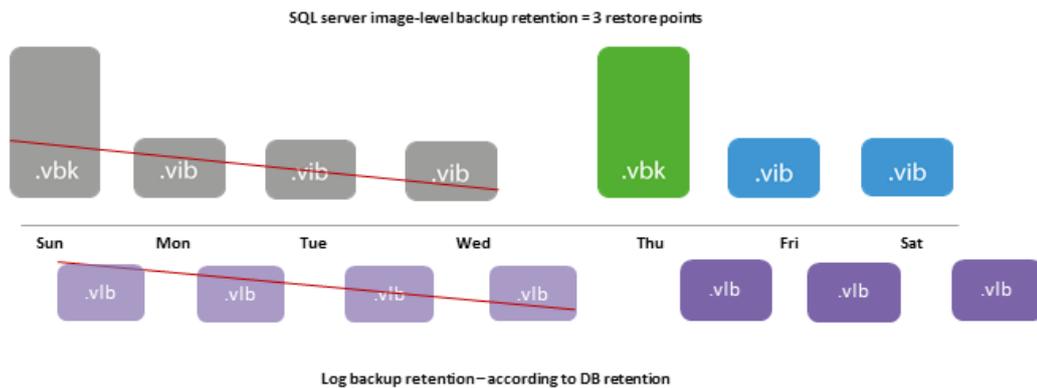
In the Veeam backup management console the total number of restore points is displayed in the **Restore Points** column of the preview pane - for both VM and log ('parent' and 'child') backup jobs:



Retention

By default, log backups retention policy is set to keep the number of restore points in accordance with the corresponding image-level backup of SQL Server VM. This allows you to have both image-level backup and all the necessary log backups at hand when you need to restore your database from VM backup to a desired state – database will be restored to the closest point before selected moment, and transaction log replay will then bring it to the desired state.

When a database restore point is removed due to VM backup retention settings, the corresponding chain of log backups is removed, too.



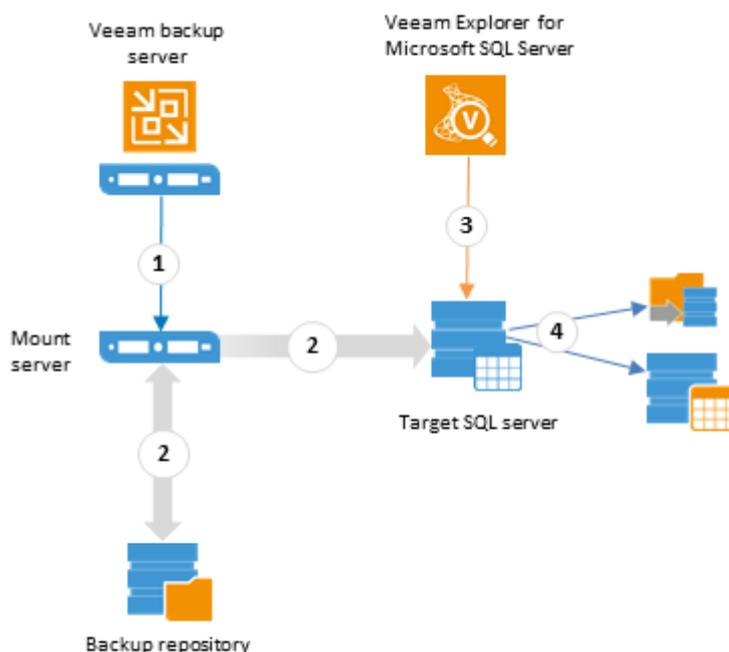
Another option allows you to keep only last <N> days. This retention setting can be used, for example, if you plan to save on storage space, saving log backups for the last few days – to be able to restore to some recent database state.

Important! If using this option, make sure that your SQL Server backup retention and log retention policies are consistent, so that corresponding restore point is always preserved - otherwise, the log replay will not be possible due to the database files missing.

How It Works: SQL Server Database Restore

Veeam Backup & Replication supports several methods of SQL Server restore from the backup. They are based on Veeam's patented vPower technology that allows you to run a VM directly from a compressed and deduplicated backup file, eliminating the need to extract the backup and copy it to production storage. Instead, Veeam exposes the VM files in a backup repository through a disk mount. Local disk space is used as a temporary write area while the VM is running from the backup. This approach is used to carry out the whole VM restore with Instant VM Recovery, Full VM Restore, as well as to perform VM disk recovery or file-level restore of the guest OS (see Veeam Backup & Replication User Guide for more information on restore options).

Veeam Backup & Replication also offers a convenient tool to extend the set of restore options – Veeam Explorer for Microsoft SQL Server, installed together with Veeam Backup & Replication server or standalone console as a part of Veeam setup. This tool helps you to restore databases from the SQL Server image-level backup to the original or different location, without a need to recover the whole server VM. The basic restore procedure includes the following steps:



1. The backup administrator uses Veeam Backup & Replication restore options to **mount** database files from the SQL Server's backup, making database hierarchy available for browsing and search.
2. Veeam Explorer for Microsoft SQL Server obtains SQL Server hierarchy information (instances and databases) and presents it to user, facilitating browsing and search. A user selects and follows the necessary database restore scenario. Veeam Explorer for Microsoft SQL Server supports several scenarios that are similar for database restore and export:
 - Restore/export to current restore point (the one currently mounted to Veeam backup server)
 - Restore/export to selected point in time

- Restore/export to the state before selected transaction.

The first scenario involves target SQL Server only; others involve target and staging SQL Servers, as follows:

	Restore scenario		Export scenario	
	To point in time	To selected transaction	To point in time	To selected transaction
Staging SQL Server	Not involved	Yes (involved; presents log to user)	Yes (involved)	Yes (involved; presents log to user, supports log replay)
Target SQL Server	Yes (involved; supports log replay)	Yes (involved; supports log replay)	Not involved	Not involved

3. Finally, the database is 're-created' on the target Microsoft SQL Server and becomes ready for use. If the export scenario was used, exported database can be then attached to the SQL Server you need.

Note: For more information, refer to https://helpcenter.veeam.com/docs/backup/explorers/vesql_hiw.html.

Veeam SQL Restore Service

In most scenarios, Veeam installs a runtime component named *Veeam SQL Restore Service* to the VM guest OS (target or staging server, depending on selected restore or export scenario) - to support restore activities on the SQL server VM guest.

Note: If you are restoring the database to the local SQL server instance, this runtime component is not installed on the SQL server guest OS.

This service runs during the restore session; it checks the rights assignment required for database restore, gets information about databases that should be restored, performs the necessary file operations (including database and transaction log copy) and so on. When restore session ends, the service is stopped and removed from guest. Then a new session starts, and the service is installed again.

The *Veeam SQL Restore Service* operates under the **Local System** account.

All service activities are logged to the `Veeam.SQL.Service_<timestamp>.log` file stored in the `Temp` folder of the system directory, next to the `Veeam.SQL.Service_<timestamp>.exe` file (runtime component installed per session). If you have enabled extended logging as described in [this Knowledge Base article](#), this log data will be stored in the Veeam Explorer for Microsoft SQL log.

Communication between Veeam Explorer and the service is performed using RPC; default TCP port range that should be open on the guest for inbound traffic includes ports **1025 - 1034**. If you need to change this port range, then do the following:

1. Open Veeam Explorer configuration file named `Config.xml` located at `C:\ProgramData\Veeam\Backup\SQLExplorer\` on the machine hosting Veeam Explorer (create the file if it is not present).

2. Locate the **<SQLExplorer>** tag (create the tag if it is not present):

```
<Veeam>
  <SQLExplorer>
    <Agent MinimumPort="1025" MaximumPort="1034"/>
  </SQLExplorer>
</Veeam>
```

3. Supply the necessary values for **Agent MinimumPort** and **MaximumPort**.

How It Works: Mount Operations

Starting with v9, Veeam Backup & Replication uses a mount server for mount operations during the restore process. This is a Windows server included in the list of managed servers and assigned a certain role in backup infrastructure: during file-level restore this server mounts VM guest files directly from the backup (kept in the repository) to itself, creating a “mount point” on its file system. As a result, VM guest files appear under `C:\veeamflr\ folder on the mount server and become available for browsing and granular restore.`

The key component of the mount server role is a special service responsible for mount functionality, named Veeam Mount Service. By design, this service is deployed on the following machines:

- On Veeam backup server
- On Veeam Backup & Replication standalone console
- On a managed server when it is assigned the role of a mount server associated with backup repository

Note:

If you plan to deploy both Veeam Backup & Replication server and standalone console, consider that Veeam Explorer for Microsoft SQL Server will be co-installed with both components.

Typical usage scenarios for mount operations are as follows:

- To provide for application item restore, VM backup is initially mounted to the machine where Veeam management console (standalone or co-installed with Veeam backup server) is deployed. This mount operation allows a user to utilize Veeam Explorer for browsing backup content and choosing the files/objects to be restored.
- After a ‘Restore’ command is initiated (typically, to the original VM running in production environment), selected data will be obtained from the repository using one more mount operation and transferred to target. This helps to decrease traffic between this repository and target VM and speed up recovery in distributed infrastructure.
- If a user performs application item restore via Enterprise Manager, and VM backup was created with guest indexing enabled, then just one mount operation will be performed, and the same mount point will be used for both browsing and restore. Data will be transferred between mount server and target VM.

For restore from backups of Microsoft SQL Server VMs, Veeam Backup & Replication creates an additional mount point on the original VM, and, if necessary, on a staging Microsoft SQL Server. Mount to the staging server is used when you perform restore to a specific transaction, or if Veeam Backup & Replication does not have information about databases (for example, if you initiate restore from storage snapshots).

To create a mount point on Microsoft Windows machines, Veeam Backup & Replication uses the iSCSI protocol. The remote machine (original VM) or staging server acts as an iSCSI initiator. The machine on which the Veeam Explorer runs acts as an iSCSI target. The iSCSI mount point is non-persistent — it is created only for duration of the restore process.

Address Recovery Requirements

Veeam supports several recovery scenarios and offers flexible options for them.

Note:

Database item-level recovery to production SQL Server is supported in Enterprise and Enterprise *Plus* editions only. Other editions support recovery via export. Please refer to the corresponding document for edition comparison:

https://www.veeam.com/veeam_availability_suite_9_5_editions_comparison_en_ds.pdf.

To choose a scenario, consider your organization's policies and requirements and decide on the following:

1. ***Will you need to recover your database to the moment when the certain Microsoft SQL Server VM restore point (backup or replica) was created, or to any point in time, maybe within the interval between two restore points?***

In the latter case, database will be restored to the closest VM restore point before the moment you specify, and then transaction log replay will bring the database to the necessary state. So, you will need to enable backup of Microsoft SQL Server transaction logs in the backup job settings.

2. ***Should you support more granular recovery and be able to roll back your databases to a state before undesired transaction?***

For the most granular restore, you will need to enable backup of the transaction logs and ensure the staging SQL Server availability and proper configuration.

3. ***Do you plan to restore your database to the original Microsoft SQL Server, or to a different server?***

Make sure that Veeam can communicate with target server over the network, and check that account you plan to use for restore has sufficient permissions. Also, check for proper target SQL Server version (later version database cannot be restored to earlier version server).

4. ***Will you mostly restore objects to their original location on a production SQL server?***

You can optimize data flow at restore - for that, examine the location of repository server and target SQL server to find out if they are located in the same site with Veeam backup server.

- If all these components are in the same site, Veeam Explorer will utilize Mount Service running on Veeam backup server for mount operation.
- If repository and production (target) SQL Server VM are located in a remote site (separately from Veeam backup server), then it can be reasonable to deploy Veeam Backup & Replication console in that remote site and launch Veeam Explorer from that console, initiating mount operation locally in the site where console is running.

5. ***Should you maintain database restores by yourself, or will you delegate restore rights to another user (group of users)?***

In the latter case, it is reasonable to use Enterprise Manager functionality for delegated restores of application items.

To implement the necessary scenario, you should, in particular, do the following:

1. Decide on the required level of restore granularity and configure the necessary SQL Server VM backup job settings.
2. Use the appropriate SQL Server VM file-level recovery option when choosing VM's restore point.
3. Provide user account with sufficient rights.
4. Examine considerations and limitations.

Examine Considerations and Limitations

1. Transaction log backup is not supported for Windows Server 2008 or earlier guests on Hyper-V 2012 R2.
2. If you plan to restore database items from an SQL Server VM running Microsoft Windows ReFS, consider that Veeam backup server or management console (that is, the machine where Veeam Explorer and the mount service are running) must be installed on the Microsoft Windows Server 2012 or later.
3. To restore from a server running Microsoft Windows ReFS 3.x, Veeam backup server or management console (that is, the machine where Veeam Explorer and the mount service are running) must be installed on the Microsoft Windows Server 2016.
4. SQL Server included in Microsoft SQL Server Failover Cluster cannot be used as a staging system.
5. By default, system databases (**master**, **model**, **msdb**) are skipped from transaction log processing and are not a part of Veeam Explorer restore workflow. These databases can be restored using file-level restore, as described in the Veeam Backup & Replication User Guide (https://helpcenter.veeam.com/backup/vsphere/performing_guest_restore.html). If you want to exclude other database(s) from transaction log processing workflow, please refer to this Veeam Knowledge Base article: <https://www.veeam.com/kb2104>. (Consider that exclusion configured this way will be treated as a global setting.)
6. By default, the **AUTO_CLOSE** option for SQL Server databases is set to **False**. However, check that this setting is **False** for the database you plan to restore (if **AUTO_CLOSE** is turned on, this may lead to database being skipped from processing).
7. Currently, point-in-time restore and restore to the state before selected transaction is not supported for replica VMs and for restore points created by backup copy job.
8. If you want to restore a database from AlwaysOn availability group node to the state as of prior to selected transaction, all nodes of the group should be located in the same time zone.
9. If you plan to restore an encrypted database using Veeam Explorer for Microsoft SQL Server, consider information provided in [this Knowledge Base article](#).
10. Table-level recovery is supported only for database tables with no external dependencies.
11. If you plan to restore database schema objects, consider that 'Replace' logic is not supported - that is, if an object with the same name exists in the production database and in the backup, then a backup object will not replace the existing one, but an object with a different name will be created instead, as described in the [Restoring Database Schema and Data](#). Also, consider that when objects are renamed, relationships between them are not renamed.
12. If both Microsoft SQL Server and Oracle Server are installed on one VM, and this VM is processed by a job with log backup enabled for both applications, Veeam Backup & Replication will back up only Oracle transaction logs. Microsoft SQL Server transaction logs will not be processed.

Plan for Recovery Verification

Veeam Backup & Replication automates and simplifies backup verification process, which is one of the most crucial parts of data management and protection. Using Veeam technologies, you can start the VMs directly from the backups in a fenced-off environment and perform reliability and availability testing as a routine part of the backup process.

Generally, to perform recovery verification testing of your SQL Server backup, you will need to create:

1. An application group required to verify full functionality of backed up VM
2. An isolated virtual lab where VM should be tested
3. A recovery verification job (SureBackup job)

For detailed information on SureBackup recovery verification, refer to the corresponding Veeam Backup & Replication User Guide:

- For VMware platform:
https://helpcenter.veeam.com/docs/backup/vsphere/surebackup_recovery_verification.html
- For Hyper-V platform:
https://helpcenter.veeam.com/docs/backup/hyperv/surebackup_recovery_verification.html

This process involves several types of resources, so to plan for them properly, in particular, in VMware environment, consider [these recommendations](#).

SETTING UP YOUR SQL BACKUP JOBS

This section describes how to set up your SQL Server backup job in Veeam Backup & Replication console; you can also use Veeam Backup Enterprise Manager (if deployed in your organization), as explained in its [User Guide](#).

Note: The following document on how Veeam backups SQL Server VMs can be also recommended (currently available for v8): https://www.veeam.com/veeam_backup_8_0_sql_backup_guide_pg.pdf.

In This Section

- [Objects to Process](#)
- [Application-Aware Processing](#)
- [Transaction Log Handling](#)
- [Log Shipping Server](#)
- [Scheduling](#)
- [Support for AlwaysOn Availability Groups](#)

Objects to Process

When deciding on the virtualized SQL Server(s) to be included in the image-level backup ('parent') job, consider the following:

- All databases of a single SQL Server instance will be processed sequentially, one by one - to minimize load on the SQL Server guest OS.
- If you want the log backup ('child') job to skip certain database (that is, exclude it from log processing workflow), you can set this database logging and recovery model to **simple**.

For SQL Server 2012 and later, Veeam Backup & Replication supports backup and recovery for AlwaysOn Availability Groups (as described later in this document).

Note: AlwaysOn Availability Group requires that the SQL Server instances reside on Windows Server Failover Clustering (WSFC) nodes.

Excluding Databases from Log Backup Process

If you need to exclude a database from transaction log processing (without excluding the database itself from application-aware VSS processing), you can use the following methods:

- Set database logging model to *simple* - then no log file will be created for this database, and log processing job ('child') will not be run for it.
- Create a list of databases to be excluded from transaction log processing. For that, you will need to create the following registry entries under **HKLM\SOFTWARE\Veeam\Veeam Backup and Replication**:

Registry entry	Type	Default value	Description	Examples
SqlBackupInstanceDatabaseDelimiter	REG_SZ	:	Specified character string will be treated as a delimiter between <i>InstanceName</i> and <i>DatabaseName</i> in a list entry.	sql_instance_name: - skip all databases from such SQL server instance (for example, sql_2014:) :db_name - skip all databases with such name on all instances (for example, :master).
SqlBackupInstanceDatabasePairsDelimiter	REG_SZ	;	Specified character string will be treated as a delimiter between list entries.	sql_2014;;master
SqlBackupDatabasesToSkip	REG_SZ	None	Instructs Veeam to exclude SQL databases from transaction logs backup process. Databases and instances should be listed using the delimiters described by the <i>SqlBackupInstanceDatabaseDelimiter</i> and <i>SqlBackupInstanceDatabasePairsDelimiter</i>	sql_instance_name:db_name - skip everything which contains such instance and database name sqlserver:msdb;sqlserver:model;:master;sqlserver:

If a database name contains default delimiter character(s), you can use any other character or character string as a delimiter when creating a list of exclusions. For example, if

SqlBackupInstanceDatabaseDelimiter value is set to **###123###**, then

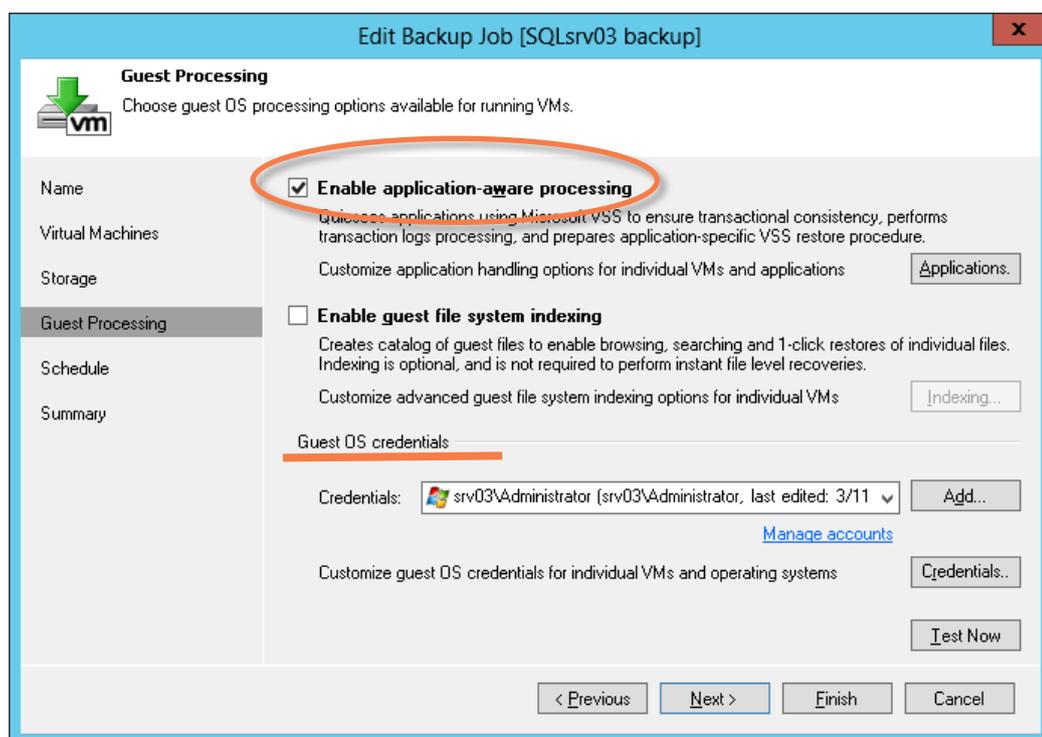
SqlBackupDatabasesToSkip will contain a list of entries like this:

sql_instance_name###123###db_name_1;sql_instance_name###123###db_name_2

Application-Aware Processing

To configure transaction log handling settings, do the following on the **Guest Processing** step of the backup job wizard.

1. Select the **Enable application-aware processing** check box.
2. In the **Guest OS credentials** section, specify the account that will be used to perform the log handling operations on guest operating system of the SQL Server VM, including truncation and backup (i.e. shipment to repository).



For log backup, this account should have the sysadmin server role assigned on that SQL Server VM. See also [this Veeam KB article](#). This is the recommended setting; however, if you need to provide minimal permissions to the account performing backup operation, you can assign the following:

- o SQL Server instance-level roles: **dbcreator** and **public**
- o Database-level roles: **db_backupoperator**, **db_denydatareader**, **public**; for system databases (*master*, *model*, *msdb*) - **db_backupoperator**, **db_datareader**, **public**
- o Securables: **view any definition**, **view server state**

For truncation of SQL Server 2012, SQL Server 2014 and SQL Server 2016 database transaction logs, this account should have the **db_backupoperator** database role (minimal required) or the **sysadmin** server role.

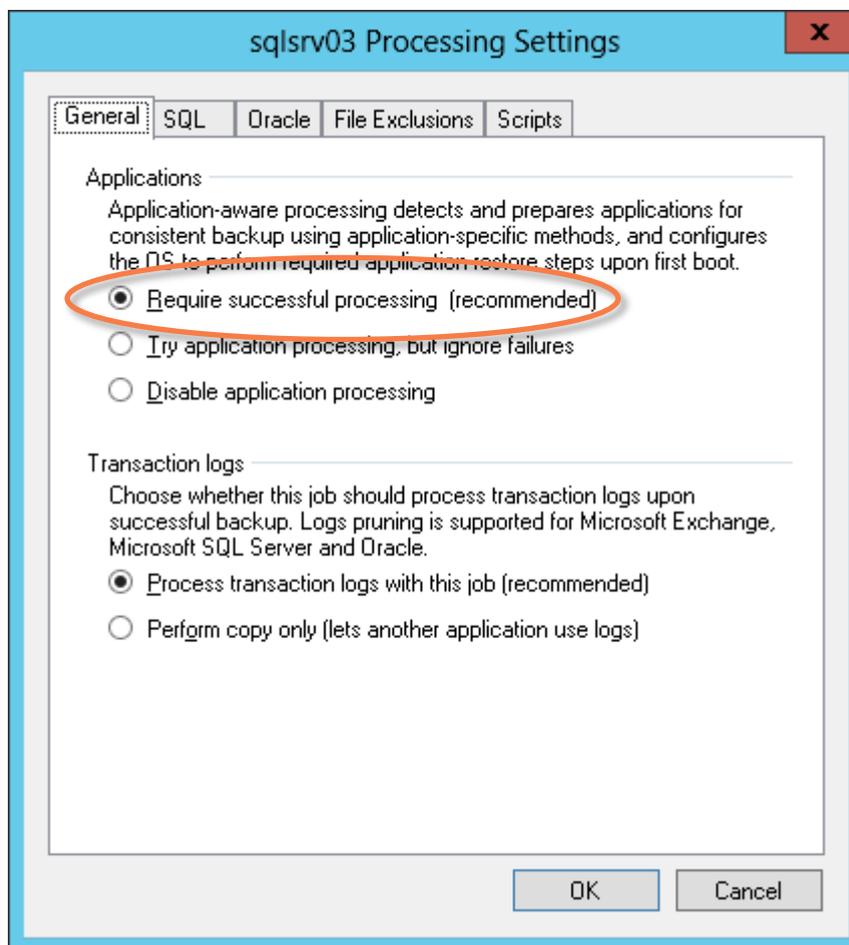
Note:

If you want transaction logs to be truncated, note that in case log truncation with the specified account is not a success, Veeam will try to perform it using NT AUTHORITY\SYSTEM account, so for SQL Server 2016, 2014 or 2012 make sure it has sufficient rights (see [this Veeam Knowledge Base article](#) for more information).

As for SQL Server 2008 and 2008 R2, default settings in these versions allow for database log truncation by local SYSTEM account (however, if they were modified, make sure this account is permitted to truncate logs).

Also, make sure that **Deny log on locally** and **Deny log on through Terminal Services** are turned OFF for the corresponding account (these can be turned on, in particular, due to group policy settings).

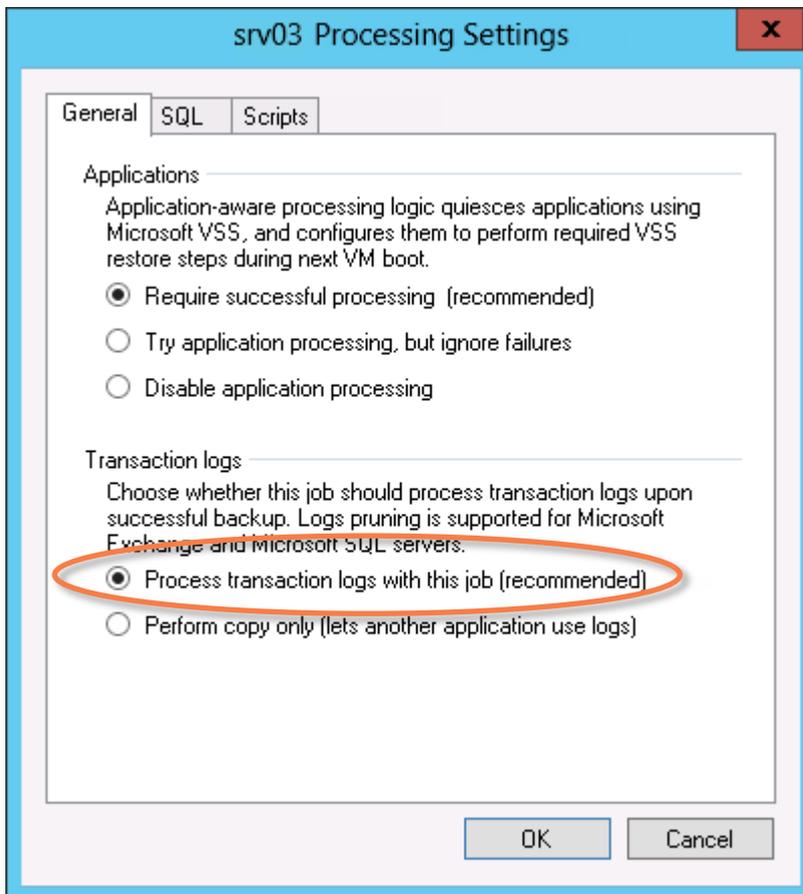
3. Then click the **Application** button, select the SQL Server VM in the list of processed machines, and click **Edit** to open its processing settings.
4. On the **General** tab, in the **Application** section, select **Require successful processing (recommended)** option – this will instruct Veeam to stop the backup process if any VSS errors occur during application quiescence.

**Note:**

If you want to continue the backup process even if VSS errors occur, you can select **Try application processing but ignore failures** – this option is recommended to guarantee completion of the job; however, the created backup image will be only crash-consistent.

Transaction Log Handling

To configure transaction log handling settings, in the **Transaction** log section on the **General** tab make sure you have selected the **Process transaction logs with this job** options. Then the **SQL** tab will become available to you.



Important! If both Microsoft SQL Server and Oracle Server are installed on one VM, and this VM is processed by a job with log backup enabled for both applications, Veeam Backup & Replication will back up only Oracle database logs. Microsoft SQL Server transaction logs will not be processed.

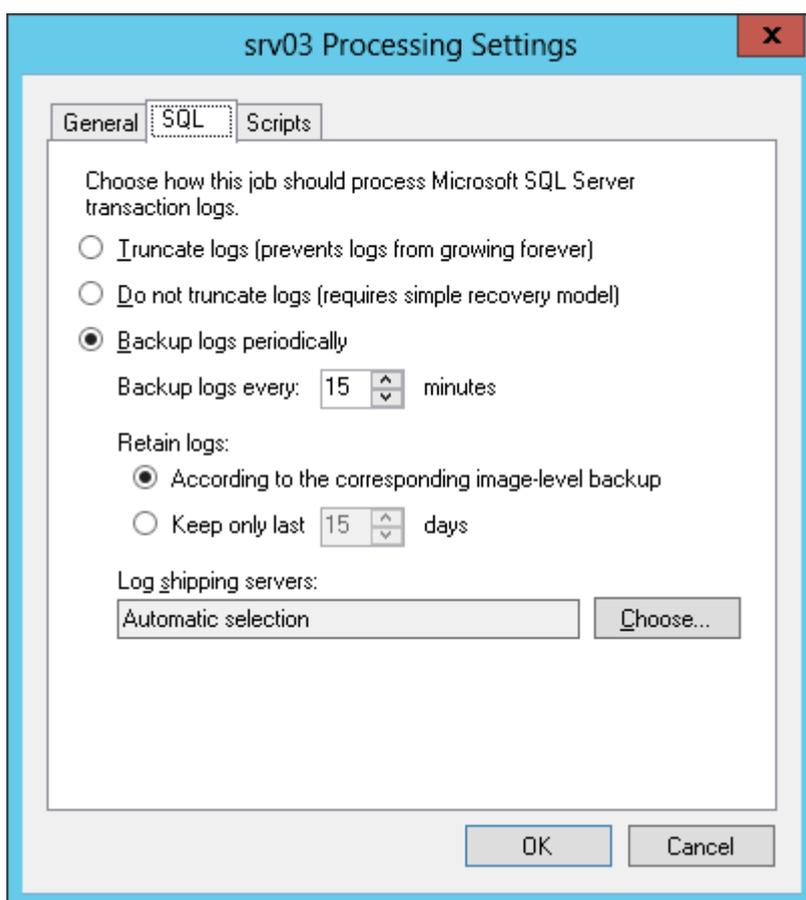
Veeam Backup & Replication supports the following options:

- **Truncate logs (prevent logs from growing forever)** – with this option selected, Veeam Backup & Replication runtime process responsible for coordination of application-aware activities will wait for the backup to complete, and then it will trigger truncation of transaction logs (using VSS). If truncation of transaction logs is not possible for some reason, the logs will remain untouched in the VM guest OS until the next start of the Veeam runtime process.

Important! If selected, this option will allow you to restore your database to the specified restore point only; restore to the certain point in time or to the state prior to specific transaction will not be possible. Also, remember that **Truncate logs** option requires that user account under which Veeam will access guest OS (and, correspondingly, transaction logs) should have sufficient permissions — see the [Application-Aware Processing](#) section above for details.

- **Do not truncate logs (requires simple recovery model)** – with this option selected, Veeam will preserve logs (if any) on the original SQL Server. Nor they will be truncated, neither backed up by Veeam Backup & Replication. Thus, if you use this option, your database administrator will have to take care of database logs. Applicable restore scenario – database restore to the state as of currently selected VM restore point.
- **Backup logs periodically** – with this option selected, Veeam Backup & Replication will periodically ship transaction logs to the backup repository and store them with SQL Server VM backup.

This backup option will allow you to use any restore scenario: to the state as of VM restore point, to the certain point in time, or to the state prior to specific transaction. The process of transaction log shipment is described in detail in the next section. You can also specify how often Veeam will process these transaction logs, uploading them to repository, set the retention and specify preferred log shipping server for data transfer.



The following table lists possible database logging models and Veeam backup options, describing all combinations.

SQL DB logging model	Veeam option		
	Truncate logs	Do not truncate logs	Backup logs periodically
Simple	Databases in this mode are skipped from this type of processing.	Applicable option for this mode.	Databases in this mode are skipped from this type of processing. Log files do not grow (and do not need to be backed up), so

SQL DB logging model	Veeam option		
	Though it is a default option, there is no sense in applying it in this case.		this option does not make sense in this case.
Full	Applicable option. Veeam performs "backup to NUL" for log files on guest.	Applicable but not recommended to use without native or 3rd party means of log truncation or backup – otherwise, logs will increase in size.	Applicable option. Log backup files (.BAK) are copied from the temporary folder on SQL Server to Veeam repository. As soon as data is copied to target, .BAK files are deleted from source.
Bulk-logged	Applicable option. Veeam performs "backup to NUL" for log files on guest.	Applicable but not recommended to use without native or 3rd party means of log truncation or backup – otherwise, logs will increase in size.	Applicable option. Log backup files (.BAK) are copied from the temporary folder on SQL Server to Veeam repository. As soon as data is copied to target, .BAK files are deleted from source.

Please be aware that in some situations transaction log handling options – **Truncate logs** or **Backup logs** will not be applied to SQL Server VM, and corresponding operations will not be performed. This can happen in the following cases:

- Transaction logs created for this SQL Server databases for the last 7 days were backed up by another job that exists on Veeam backup server
- OR-
- At the moment, this SQL Server VM is being processed by another job that is performing log backup but has not placed these logs into backup repository yet

Note: Under such circumstances, transaction logs will not be truncated for this SQL Server databases (in both cases).

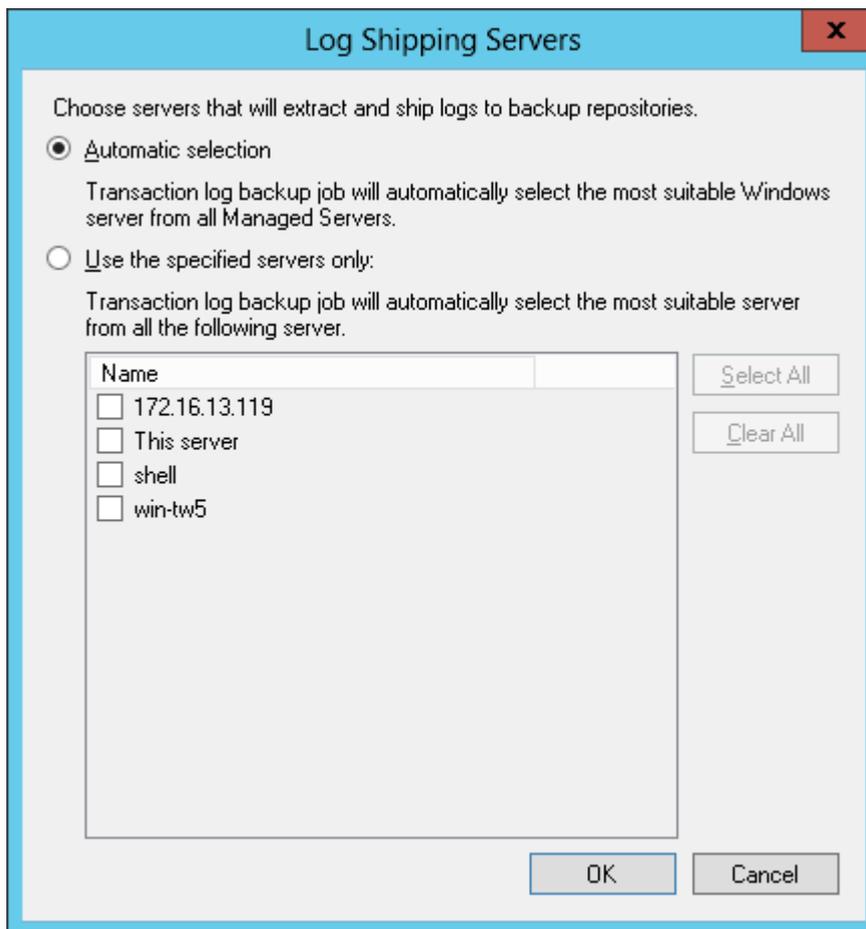
A situation may occur when you need to remove SQL Server VM from the backup job – please refer to the [corresponding section](#) of this document for the steps to be taken in this case.

Options available on other tabs of the **Processing Settings** dialog are described in detail in the corresponding sections of the Veeam Backup & Replication User Guide.

Log Shipping Server

For each SQL Server VM guest OS, Veeam Backup & Replication defines the way it will transport log backup files to the backup repository. Two transport types are supported:

- Direct data transfer between the VM guest OS and repository. This is a recommended method as it does not involve additional resources and puts less load on the VM guest OS.
- Alternatively, if it is not possible to establish direct connection, you can configure Veeam Backup & Replication to use a Microsoft Windows host added to the backup Infrastructure as a log shipping server. You can select server(s) from the list, or instruct Veeam Backup & Replication to automatically choose an optimal log shipping server from the Managed Servers (that is, added to backup infrastructure in Veeam backup console). The choice will be based on two criteria: possible data transfer method(s) and network availability.



Possible data transfer methods for log shipping server are:

- *Network* – this is a default method which can be used for both platforms. It uses the VM guest OS to obtain files and then transfer them over the network. To offload the VM guest OS, logs are created one by one (not simultaneously); 1 log creation request is created for every DB.
- *VIX* - this method can be used for VMware VMs only. It allows you to obtain guest files bypassing the network. In this case, for each SQL instance, 1 log creation request is created for all DBs (grouped by instance).

Network availability criteria are based on the location of the Microsoft Windows host and VM guest OS. Otherwise, Veeam Backup & Replication detects if any of the following is applicable to the Microsoft Windows host:

- A. Microsoft Windows host is located on the same ESX(i) host as the VM guest OS.
- B. Microsoft Windows host and the VM guest OS are in the same sub-network
- C. Microsoft Windows host and the VM guest OS are in different sub-networks (the production infrastructure is isolated from the backup infrastructure).

So, when choosing as log shipping server for log transfer, top priority goes to the Windows host "A" that supports the Network data transfer mode.

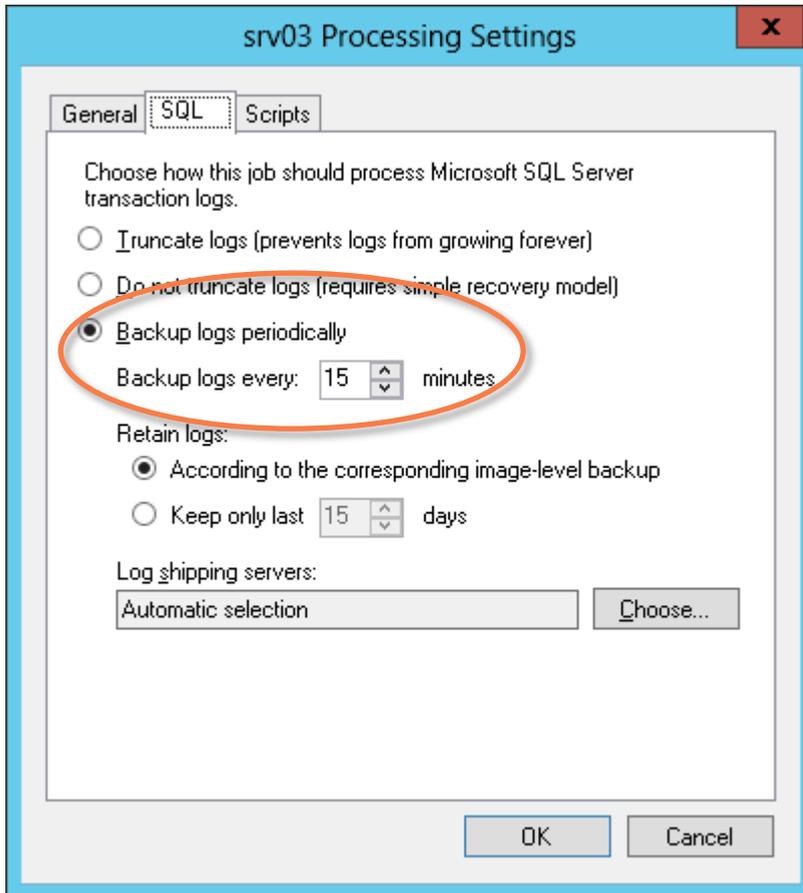
Important! You may not want some of your busy managed hosts to participate in the process – if so, use manual server selection when assigning the role of log shipping servers. It is also recommended that you have not only one dedicated Windows host for log shipment, but also a redundant host for availability purposes.

Veeam Backup & Replication automatically re-detect available hosts, so if a log shipping server fails for some reason, processing will be performed by another server.

Scheduling

To prevent a backup job from overlapping with production hours and ensure it does not provide unwanted overhead on your virtual environment, you can schedule SQL Server VM backup to run during off-peak hours, or configure a specific backup window -a period of time on week days when backup jobs are permitted to run. If the job exceeds the allowed window, it will be automatically terminated.

With a scheduled SQL Server VM backup job, transaction log backup job runs permanently in background. By default, logs are backed up to repository every 15 minutes. You can modify this setting on the **SQL** tab:



Consider the following when you set up log backup interval:

- Log processing for a database will not start until the last log for the previously processed database has been moved away from the VM guest OS.
- It is recommended to set up log backup frequency with respect to actual time required to transfer log data to repository. Otherwise, you will be getting an error message in the job session data, and SLA misses will be also reported in job statistics. For example, if you set up the job to backup logs every 5 minutes, but actually the process takes 15 minutes, the SLA will be missed 2 times.
- You can also use network and repository traffic throttling capabilities.

Support for AlwaysOn Availability Groups

AlwaysOn Availability Groups allow you to increase fault tolerance between active and hot-standby databases without involving shared physical disks, which is quite important for virtualization of SQL Servers. Veeam Backup & Replication supports AlwaysOn Availability Groups for virtualized SQL Server 2012 and later.

Image-level Backup of SQL Server VM

During image-level backup of an SQL Server VM, Veeam requests and analyzes information about databases that are included in the AlwaysOn Availability Groups. Depending on the retrieved information, a VSS snapshot will be created with or without COPY_ONLY flag. The COPY_ONLY flag for VSS snapshot is triggered if the VM represents a secondary node for at least one AlwaysOn Availability Group (and this group is not in the *Resolving* state).

Veeam Backup & Replication also detects what cluster the database belongs to. If the backup job does not include all VMs from the cluster, a warning message will be issued.

Retrieved information is saved for further log identification.

Transaction Log Backup

Transaction log backup can be performed only for those databases that were successfully backed up with the corresponding settings, either on the primary or on the secondary node of AlwaysOn Availability Group.

Transaction logs processing interval may be the same, or may differ through the VMs included in AlwaysOn Availability Group – in the latter case, Veeam Backup & Replication will use minimal value (default is 15 min). At each log processing interval, Veeam Backup & Replication chooses the AlwaysOn Availability Group node for which transaction logs will be backed up.

Important! Logs are backed up for the only one AlwaysOn Availability Group node selected.

To become a subject for log backup, the node must meet the following criteria:

- Required Veeam components can be installed on that node (in particular, the VM must be running)
- If there are any logs remaining in the temporary folder on that node of AlwaysOn Availability Group, this means they were not backed up to the destination repository during the previous log backup session, so this node must be processed first.
- Databases in the AlwaysOn Availability Group(s) for that node were successfully backed up (for the last two processing intervals).
- Node should be available over the network (preferred data transfer method) or via VIX (for vSphere platform).
- The VM should be included in the list of preferred nodes (for backup) retrieved from SQL Server. If there are no preferred nodes, then any node can be chosen.

Considerations and Limitations

If you want to restore a database from AlwaysOn availability group node to the state as of prior to selected transaction, all nodes of the group should be located in the same time zone.

If you are using Basic Availability Groups on Microsoft SQL Server 2016 Standard Edition, consider that the secondary replica node cannot be backed up because the secondary node is not readable. For more information, see <https://blogs.msdn.microsoft.com/alwaysonpro/2016/05/02/sql-server-2016-alwayson-availability-group-enhancements-basic-availability-group-on-standard-edition/>.

CONFIGURING REQUIRED PORTS

The following table describes network ports that must be opened to ensure proper communication with the runtime coordination process deployed inside the Microsoft SQL Server guest OS for application-aware processing and transaction logs shipping of SQL Server VM.

From	To	Protocol	Port	Notes
Veeam backup server or Guest interaction proxy (in Enterprise and Enterprise <i>Plus</i> editions)	Microsoft SQL Server VM guest OS	TCP, UDP	135, 137-139, 445	Ports required to deploy the runtime coordination process on the VM guest OS.
		TCP	49152-65535 (for Microsoft Windows 2008 and newer)	Dynamic RPC port range used by the runtime coordination process deployed inside the VM guest OS for application-aware processing (when working over the network).* For more information, see http://support.microsoft.com/kb/929851/en-us .
		TCP	6167	[For Microsoft SQL Server transaction logs shipping.] Port used by the runtime process on the VM guest OS from which transaction logs are collected.
Microsoft SQL Server VM guest OS	Veeam backup server or Guest interaction proxy (in Enterprise and Enterprise <i>Plus</i> editions)	TCP	49152-65535 (for Microsoft Windows 2008 and newer)	Dynamic RPC port range used by the runtime coordination process deployed inside the VM guest OS for application-aware processing (when working over the network).* For more information, see http://support.microsoft.com/kb/929851/en-us .
	Log shipping server	TCP	2500 - 5000	[For Microsoft SQL Server transaction logs shipping.] Default range of ports used by Veeam data mover service for data transmission over the network.

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

MANAGING YOUR SQL BACKUP JOBS

As explained above, with log backup switched on, SQL Server backup job comprises 2 jobs:

1. 'Parent' job, which performs image-level VM backup and is explicitly created and named by user as **<job_name>**, for example, *MyTestJob*.
2. 'Child' job, which performs transaction log processing and is named automatically using SQL Backup suffix: **<job_name> SQL Backup**. So, for our example it will be *MyTestJob SQL Backup*.

Child job is created by Job Manager if it detects that there is a parent job, scheduled to backup at least one SQL Server VM, with application-aware image processing enabled.

Parent job session starts on schedule or is started manually by user; child job (transaction log processing) follows it - to ensure that full backup is periodically created, and log backups then will be of use to restore the VM to the required state. Child job (transaction log processing) runs permanently in background with the specified interval. It can be stopped or disabled under certain conditions.

Starting, Stopping, Enabling and Disabling the Job

Starting

- Image-level backup job ('parent') can be started automatically upon schedule or manually when you click **Start** on the Veeam Backup & Replication toolbar.
- Log backup job ('child') is initially started when you enable schedule for 'parent' image-level backup job, and then it works continuously in the background, with new session beginning each time the 'parent' job is started.

Important! For the 'child' job to run, the 'parent' job must be enabled and have its schedule configured for automatic job runs. If schedule is disabled, and you force image-level SQL Server backup job start manually by clicking **Start**, no log processing will be performed for SQL Server.

Stopping/Disabling

- If you need to stop transaction log processing for all VMs included in the parent job (SQL Server VM image-level backup), select this parent job and use the **Disable** command from the ribbon or shortcut menu. You can also clear **Run the job automatically** check box in the job properties.
The Job Manager instructs the child job to complete the remaining steps of log processing for all VMs included in the parent job, and changes the status of both jobs to *Disabled* (no re-start will be performed after that).
- To re-activate transaction log processing for all VMs in the parent job, clear the **Disable** property for that parent job in its shortcut menu, or release the **Disable** tool button on the toolbar.
- If you need to disable log processing for specific VM, go to the job settings and switch log backup off for that VM on the SQL tab of the VM **Processing Settings** dialog.
- If you need to manually stop both image-level and log backup jobs, then use **Stop** command from the toolbar or the parent job's shortcut menu, and then **Disable** parent job.

- Instant **Stop** is also performed for both jobs when the Veeam Backup service is stopped, for example, during the upgrade.

Note: Transaction log processing job state is also changed to **Disabled** when all VMs' processing intervals are over, or when the backup of SQL Server VM is started.

To delete both jobs from configuration, select SQL Server VM backup (parent job) and use the **Delete** command from the ribbon or shortcut menu. Child job(s) will be deleted automatically.

Viewing Job Information

You can view summary information on the latest SQL Server VM and transaction log backup sessions under the **Last 24 hours** node in the Backup & Replication view.

Information on all sessions of both jobs is available in the **History** view, with the **Jobs > Backup** node selected in the tree on the left.

Information is presented in the following way:

- SQL Server VM image-level backup (parent job) sessions are displayed with the **<job_name>** as **Job Name** and **Backup** as **Session Type**.
- Database transaction log backup (child job) sessions are displayed with the **<job_name> SQL Backup** as **Job Name** and **SQL Log Backup** as **Session Type**.

In the **Backup & Replication** view, you can view summary information on the image-level and log backups created by these jobs: for that, select the **Backups** node in the navigation tree, and expand the parent job in the working area on the right:

JOB NAME	CREATION TIME	RESTORE POINTS	REPOSITORY	PLATFORM
Backup Job 2	9/7/2015 8:13 AM		Default Backup Repository	VMware
Backup_No_Indexing	10/16/2015 8:09 AM		Default Backup Repository	VMware
DC Backup	11/10/2015 9:56 AM		Default Backup Repository	VMware
SQL Backup Job	11/19/2015 3:41 AM		Default Backup Repository	Hyper-V
SQL Server transaction log backup	11/25/2015 1:45 AM	178	Default Backup Repository	SQL Server
sqlsrv03	11/23/2015 10:02 PM	3		

Under the job node, you can find information on the **Transaction log backup**, including:

- The last backup creation time

- The number of restore points for transaction log of all SQL Servers for which the logs were backed up by this job
- Destination repository where log backups were stored
- SQL Server (as the **Platform** of the job)

You can also view statistics and reports separately for image-level and log backup jobs. To see the latest session data, use the **Backup & Replication** view, and under the **Jobs** node select the SQL Server backup job you need.

- To view SQL Server VM backup session data, select **Backup Statistics** from the job's shortcut menu, or **Statistics** from the ribbon menu.
- To view SQL transaction log backup session data, select **SQL Statistics** from the job's shortcut menu, or from the ribbon menu.

To examine previous sessions, use the **History** view, select the necessary job and click **Statistics** on the ribbon. Detailed description of statistics information is provided in the [Appendix](#).

Tip:

If you have Veeam ONE deployed in your environment, you can also generate the "SQL Backup Job Historical Information" report described at https://helpcenter.veeam.com/docs/one/reporter/sql_backup_job_information.html.

Reconfiguring Jobs with Microsoft SQL Server VMs

In some situations, you may need to reconfigure a backup job that processes a Microsoft SQL Server VMs and ships transaction logs. For example, you may want to create a separate backup job to process the virtualized database, and delete the VM running the database from the previously created job.

When you configure a new job, mind the restriction on transaction logs shipping.

By default, the new backup job that processes the VM will not ship transaction logs if transaction logs for this VM have been shipped for the last 7 days by another backup job on the same backup server. To overcome this restriction, you can reconfigure the backup jobs.

The procedure below is required only if you do not want to delete the initial backup job.

1. Remove the Microsoft SQL Server VM from the initial backup job.

As an alternative, you can disable transaction logs shipping in the job properties: in the guest processing settings, select the **Do not truncate** or **Perform copy-only** option.

2. In the registry on the backup server, locate the **SqlBackupLogsAgeDaysToSkipLogBackup** registry key and set its value to **0**
-or-
delete transaction logs for the Microsoft SQL Server VM from the backup repository.
3. Create a new backup job to process the Microsoft SQL Server VM and enable transaction log shipping in the job properties.
4. It is strongly recommended that you delete the **SqlBackupLogsAgeDaysToSkipLogBackup** registry key in a week. In this case, if the VM is added to several jobs, transaction logs will not be shipped to different backup repositories.

Log Backup Import

Import of log backups without VM backups is not supported (as there will be no restore point). You can either import .VBM file (database and log backup), or import .VBK file (VM backup) – Veeam will browse for corresponding log backups and import them, too.

Also, consider that .VLB files cannot be copied to DR site by your backup copy job.

RESTORING SQL SERVER DATABASES

To recover an SQL Server database with Veeam, you do not need to fully restore and start the virtual machine hosting the SQL Server. Instead, Veeam Backup & Replication data recovery functionality is used to mount the selected restore point (image-level backup or replica) of Microsoft SQL Server VM to the mount server, and then you select how you want your database to be restored.

1. You can use Veeam Explorer for Microsoft SQL Server that supports the following scenarios:
 - Restore database to the original or different location, to the state as in the selected restore point of SQL Server VM.
 - Restore database to the selected point in time.
 - Restore database to the state prior to undesired transaction.
 - Restore database schema and data to the states listed above.
 - Restore all databases hosted on the SQL Server or instance ('bulk restore') to the state as in the selected restore point (backup), or to the selected moment in time.

Note: This tool is described in detail in the sections below and in the [Veeam Explorers Series](#) document.

2. Authorized restore operators can recover the necessary databases using Veeam Backup Enterprise Manager. Two scenarios are supported:
 - Restore database to the original or different location, to the state as in the selected restore point of SQL Server VM.
 - Restore database to selected point in time.

Note: For details on configuring and using Enterprise Manager, refer to its documentation at <https://helpcenter.veeam.com/docs/backup/em/introduction.html>.

3. Finally, users can utilize Veeam Application-Item Restore wizard (U-AIR) to configure a request for database recovery to selected restore point (backup or replica), as well as for recovery of a table or a query result. Three scenarios are supported:
 - Restore database schema objects from backup into production database by running automatically created scripts.
 - Restore database tables to the production database, or save them to the specified location.
 - Execute a custom query against the database from the backup to restore specific data back to the production Microsoft SQL Server database or save it as data files to a specific folder.

Note: For details on using AIR wizard, refer to U-AIR documentation at https://www.veeam.com/veeam_backup_9_5_uair_wizard_user_guide_en_pg.pdf.

The next sections will give you several recommendations on how to carry out the restore procedure using Veeam Explorer for Microsoft SQL Server; Enterprise Manager capabilities will be also covered in brief.

Preparing for Database Restore

Before starting SQL Server database restore process, you should take preparatory steps and check the corresponding prerequisites.

- Step 1: Check Required Ports
- Step 2: Prepare Staging SQL Server
- Step 3: Check User Account Permissions
- Step 4: Examine Considerations and Limitations
- Step 5: Recover Database Files from SQL Server Backup

Step 1: Check Required Ports

The following network ports must be opened to ensure proper communication of the Veeam backup server with the VM guest OS system during application-item restore.

From	To	Protocol	Port	Notes
Machine running Veeam Explorer (this can be Veeam backup server or standalone console)	Microsoft SQL VM guest OS	TCP	1433,1434 and other	Port used for communication with the Microsoft SQL Server installed inside the VM during application-item restore. Exact port numbers depends on configuration of your Microsoft SQL server. To learn more, see http://msdn.microsoft.com/en-us/library/cc646023.aspx#BKMK_ssde .
		UDP	1434	Port used by Microsoft SQL Server Browser Service. To learn more, see https://technet.microsoft.com/en-us/library/ms181087(v=sql.130).aspx .
		TCP	1025 - 1034	Default RPC port range for the runtime component installed on target or staging server guest OS to support restore operations in most scenarios. Port is open only during application item restore (Veeam Explorer session). To learn how to modify default port range, see How It Works section of the Veeam Explorers document.

The following table describes range of ports opened by Veeam Backup & Replication for iSCSI traffic during restore to the original VM. See [How It Works: Mount Operations](#) for details.

From	To	Protocol	Port	Notes
Target remote machine to which application items are restored, or staging Microsoft SQL Server	Mount server associated with the backup repository	TCP	3260 - 3270	<p>Range of ports opened by Veeam Backup & Replication for iSCSI traffic during restore to the original VM. See How It Works: Mount Operations for details.</p> <p>Ports are open only during application item restore (Veeam Explorer session).</p>

Step 2: Prepare Staging SQL Server

To restore or export database to the state as of the current restore point, no extra configuration is needed. However, as described in the [How It Works](#) section, to support other scenarios, Veeam Explorer for Microsoft SQL Server needs a staging Microsoft SQL Server. In particular, it is required for the following operations:

- Restore to the state before selected transaction
- Export to selected point in time
- Export to the state before selected transaction

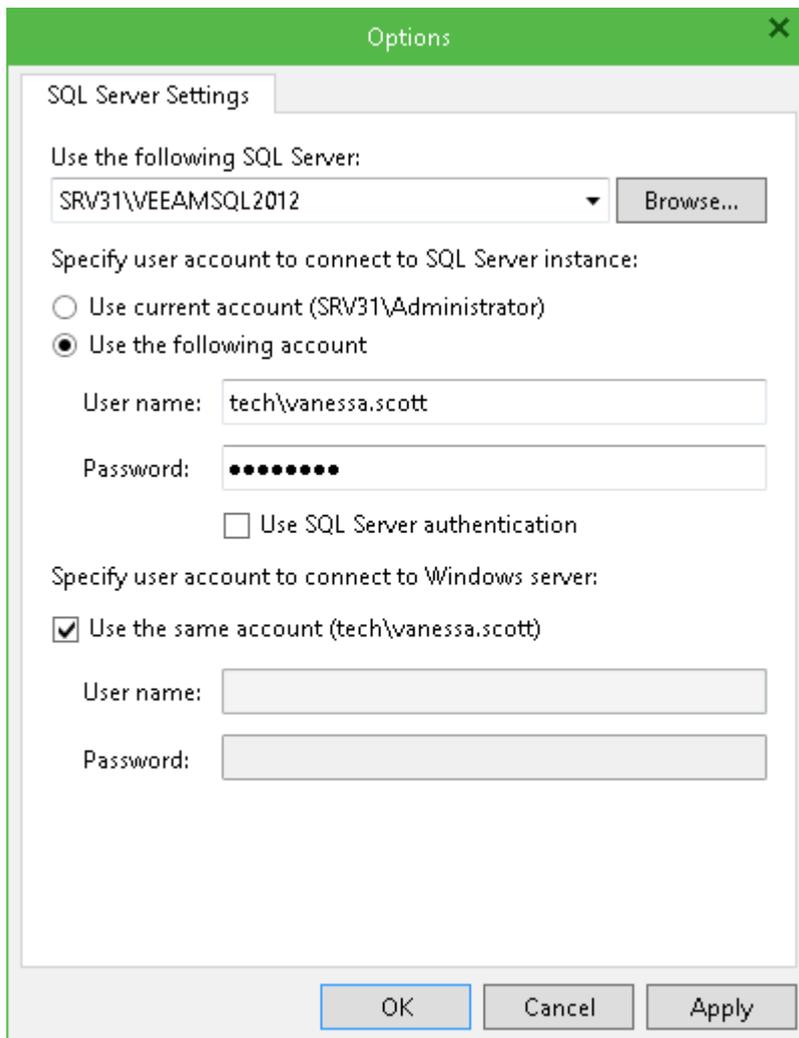
By default, local Microsoft SQL Server deployed with Veeam backup server is used for that purpose. You can use another SQL Server as a staging system.

Important!

- Remember that your staging system must have the same or later version as the original SQL Server.
- SQL Server included in Microsoft SQL Server Failover Cluster cannot be used as a staging system.
- Make sure the accounts you specify have sufficient [permissions](#) to connect to Windows machine and to access SQL server.
- Consider that specific ports on staging server will be used during restore session, as described in the [Check Required Ports](#) section.

To specify another staging server:

1. Launch Veeam Explorer for Microsoft SQL Server, open main menu and click **Options**.
2. Select an SQL Server instance from the list, or click **Browse**.
3. Specify user account that will be used to access this SQL server - this can be current account (the one under which Veeam Explorer is running) or another account.
4. Specify account that will be used to access Windows machine where that SQL server runs. You can also use current account or another account.
5. Click **Apply**, then close the dialog.



If you plan to add database(s) to Veeam Explorer scope manually, consider the following when setting up the staging SQL Server:

- If this SQL Server belongs to an untrusted domain, connection will not be possible.
- If this SQL Server belongs to a trusted domain, then only SQL Server authentication is possible.
- If this SQL Server belongs to the same domain as the machine where Veeam Explorer runs, then both Windows and SQL Server authentication methods are possible. In this case, if staging SQL Server is running under local account and you plan to use Windows authentication for connection, you will need to configure delegation settings as follows:
 - a) In **Active Directory Users and Computers**, select the necessary staging SQL Server.
 - b) Open its properties and go to the **Delegation** tab. Select **Trust this computer for delegation to specified services only** and **Use any authentication protocol** options for the **cifs** service on the computer where Veeam Explorer runs.
 - c) After that, restart the staging SQL Server.
 - d) Select the domain user account that you want to use for connection to the staging SQL Server, and in its properties on the **Account** tab make sure the **Account is sensitive and cannot be delegated** check box is cleared.

Note: Contact your Active Directory administrator, if necessary.

Step 3: Check User Account Permissions

Requirements for Veeam Explorer Operations

Check the following permissions and contact your database administrator if necessary:

1. To create an application-consistent backup, the user account that you specify for guest processing of the Microsoft SQL Server VM in the backup job should have the **sysadmin** fixed role assigned on that SQL Server. Otherwise, no database content will be visible in Veeam Explorer. See also [this Veeam KB article](#).
This is the recommended setting; however, if you need to provide minimal permissions to the account performing backup operation, you can assign them as recommended in the [Application-Aware Processing](#) section of this guide.
2. The account used to run Veeam Explorer for Microsoft SQL Server should have sufficient permissions for the folder where you plan to export the database files: **Read** and **Write** are minimal recommended.
3. The account you will use to access the target Microsoft SQL Server where database will be restored needs **sysadmin** fixed role on that server. This also refers to the account that will be used to access the staging server (if necessary).
4. The account you plan to use for connection to the Windows machine (where database log backup files will be copied for further log replay) will need sufficient permissions to access the administrative share on that machine: **Read** and **Write** are minimal required. For restore scenarios, that machine is your target SQL Server.

Delegating Restore Rights via Enterprise Manager

If you plan to delegate rights to perform database restore operations to personnel in charge, you should assign an Enterprise Manager role to them and configure role settings accordingly:

1. Log in to Enterprise Manager with administrative rights, go to **Configuration** tab.
2. To create a role for restore operator, follow the steps described in the [Configuring Security Settings](#) section of the Veeam Backup Enterprise Manager User Guide.
3. To allow for database restore, make sure the **Microsoft SQL Server databases** checkbox is selected for restore operator's account in the **Account** dialog.
4. If you need to prevent restore operators from overriding production databases at restore, select the **Deny in-place database restores (safer)** check box.

ACCOUNT

Account type: User

Account: tech\wscott

Role: Restore Operator

Restore scope:

- All virtual machines
- Selected virtual machines only Choose

Allow restore of:

- Entire virtual machines
- Guest files
- Allow in-place file restores only
- Allow restore of files with these extensions only:
- Microsoft Exchange items
- Databases
 - Microsoft SQL Server databases
 - Oracle databases
 - Deny in-place database restores (safer)

OK Cancel

Authorized users will be able to log in to Enterprise Manager and access the **Items** tab with database restore commands and options.

Step 4: Examine Considerations and Limitations

1. If you are using SQL Server Express edition (for example, SQL Server 2012 Express included in Veeam setup) as a staging system, consider that databases that exceed 10 GB cannot be attached to this SQL Server due to Express edition limitations ([http://msdn.microsoft.com/en-us/library/cc645993\(v=SQL.110\).aspx](http://msdn.microsoft.com/en-us/library/cc645993(v=SQL.110).aspx)). If you plan to recover large databases following a scenario that requires a staging server, then it is necessary to configure a fully functional commercial version of Microsoft SQL Server as a staging system.
2. By default, system databases (**master, model, msdb**) are skipped from transaction log processing and are not a part of Veeam Explorer restore workflow. These databases can be restored using file-level restore, as described in the Veeam Backup & Replication User Guide (https://helpcenter.veeam.com/docs/backup/vsphere/guest_file_recovery.html). If you want to exclude other database(s) from transaction log processing workflow, please refer to this Veeam Knowledge Base article: <https://www.veeam.com/kb2104>. (Consider that exclusion configured this way will be treated as a global setting.)
3. Currently, point-in-time restore and restore to the state before selected transaction is not supported for replica VMs and for restore points created by backup copy job.
4. If you want to restore a database from AlwaysOn availability group node to the state as of prior to selected transaction, all nodes of the group should be located in the same time zone.
5. If you plan to restore an encrypted database using Veeam Explorer for Microsoft SQL Server, consider information provided in [this Knowledge Base article](#).
6. Table-level recovery is supported only for database tables with no external dependencies.
7. If you plan to restore database schema objects, consider that 'Replace' logic is not supported - that is, if an object with the same name exists in the production database and in the backup, then a backup object will not replace the existing one, but an object with a different name will be created instead, as described in the Restoring Database Schema and Data. Also, consider that when objects are renamed, relationships between them are not renamed.

Step 5: Extend Database Files from SQL Server Backup

You can use any Veeam Backup & Replication restore option to extend database files from the SQL Server's backup and make them available for Veeam Explorer. This can be done automatically through a wizard, or manually – in the latter case, you should then add the database to Veeam Explorer scope. Refer to the [corresponding section](#) of Veeam Explorers guide for details.

Following Your Restore Scenario

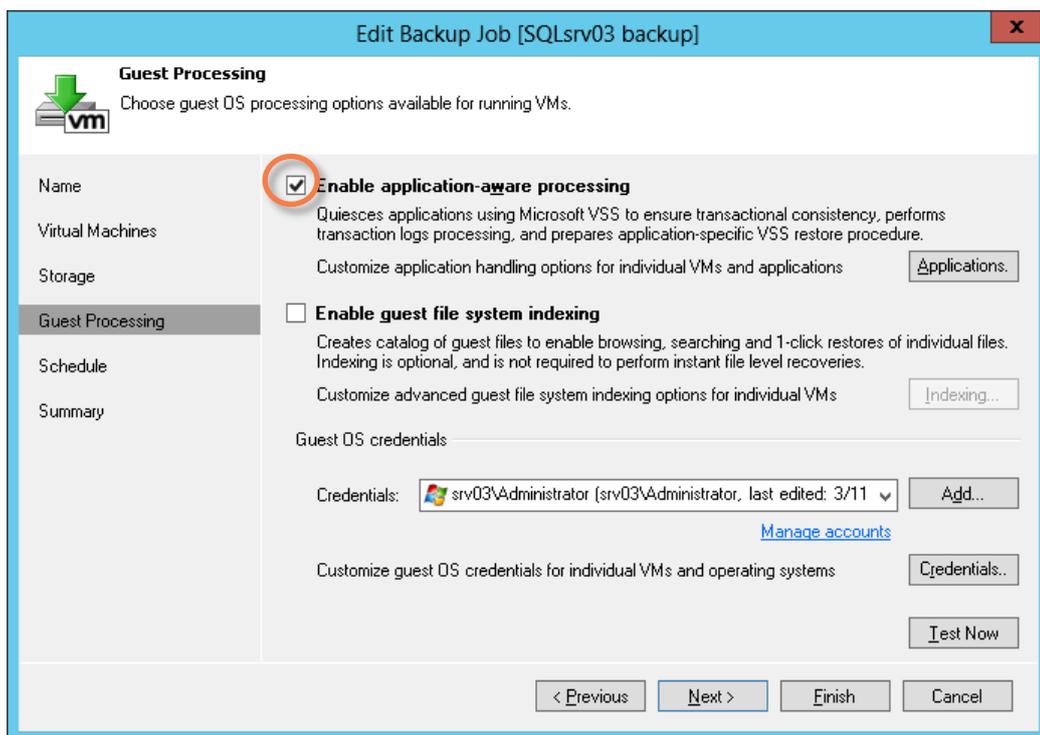
With Veeam Explorer for Microsoft SQL Server, you can restore SQL server database(s) to the necessary state, using the original or different server instance as a destination. For that, the Restore Wizard is used, where you can choose target server and restore options. Also, you can use 1-Click Restore feature to quickly recover your database(s) to the original location.

- Restoring a Database to Current State
- Restoring a Database to Specific Point in Time
- Restoring a Database to State Before Selected Transaction
- Restoring Database Schema and Data

Restoring a Database To Current State

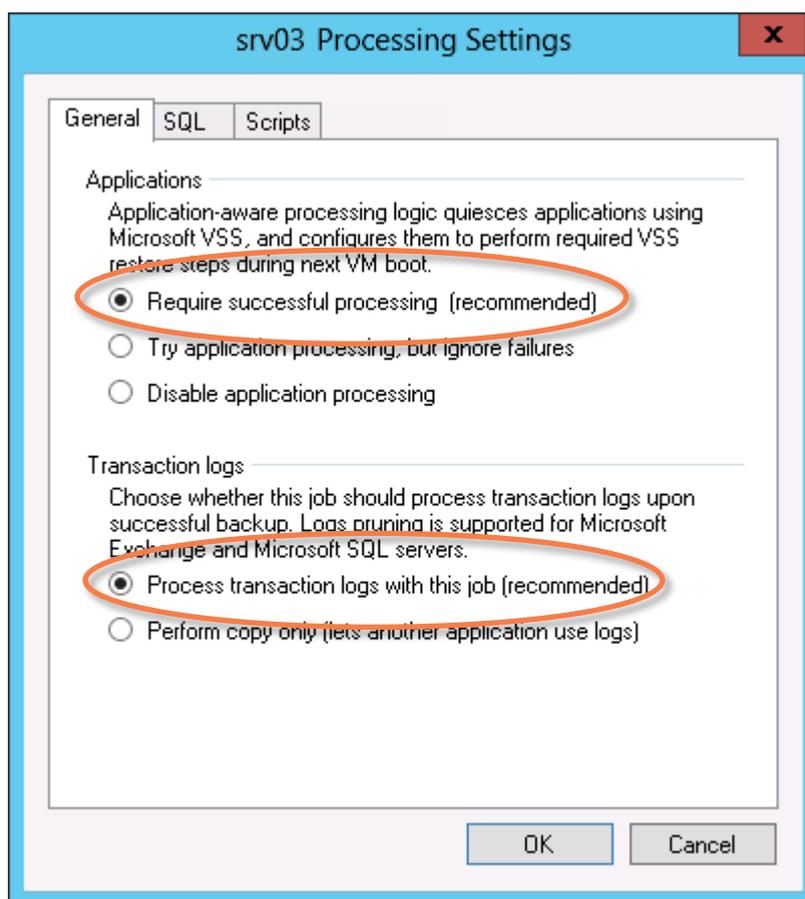
If you plan to restore a database using this scenario, consider that you will need an application-consistent image-level backup (or replica) of your SQL Server VM, successfully created by Veeam backup (or replication) job. Remember the following:

1. Any logging and recovery model can be set for this database.
2. In the Veeam job settings, the **Enable application-aware image processing** check box should be selected in **Guest Processing** options:



3. On the **General** tab of the advanced guest processing settings dialog, the following options should be selected:
 - **Require successful processing (recommended)** – in **Applications** section

- **Process transaction logs with this job (recommended)** – in **Transaction logs** section



4. Any transaction log handling option can be selected on the **SQL** tab.

Refer to the https://helpcenter.veeam.com/docs/backup/explorers/vesql_bu_job_settings.html to learn more.

Also, consider that such a scenario has 2 options:

- Restore to original location
- Restore to custom location

Restore to Original Location

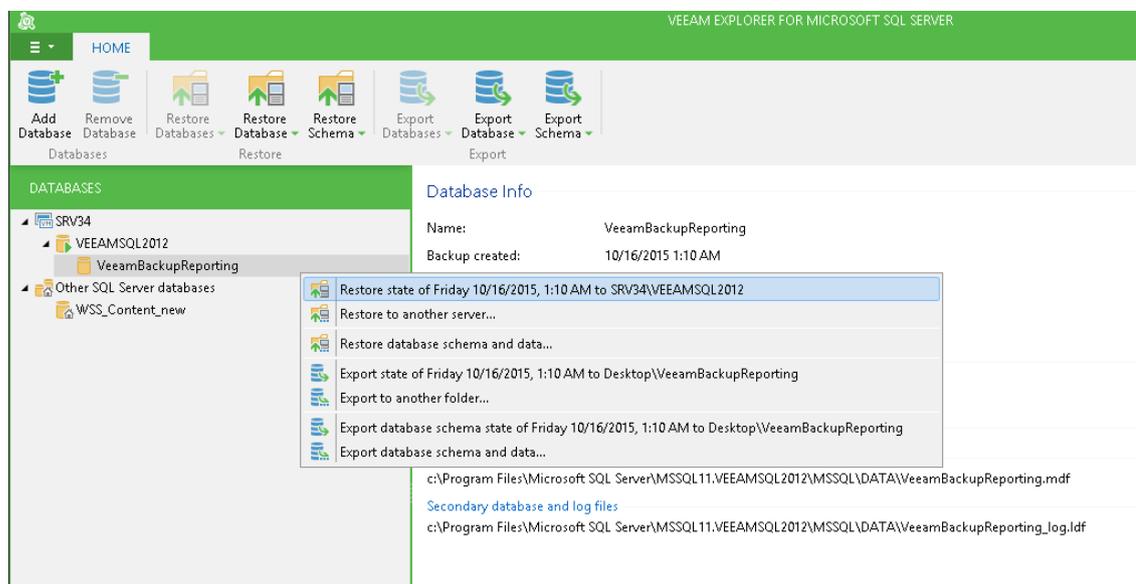
Use 1-Click Restore if you need to quickly restore a database to the original location (that is, the SQL Server instance that was originally hosting the database). Restore process will be performed with the following settings:

- Database will be restored to the current restore point (that is, to the moment when currently mounted backup or replica of SQL Server VM was created).
- Database will be restored to the original VM (files will be copied to the original location and then attached to the original SQL Server instance).
- By default, Veeam will first use current user credentials and Windows authentication method to connect to target SQL Server. If these credentials are insufficient, then Veeam will try using

the account specified for guest OS access in the SQL Server backup job settings. If it does not suit either, you will be prompted to enter the credentials manually.

Important! The account used for connection to SQL Server should have administrative rights on that server.

To launch 1-Click Restore from Veeam Explorer, select the required database in the navigation tree on the left, then select **Restore Database > Restore <current_state_date> to <server>\<instance>** from the toolbar or from the node's shortcut menu:

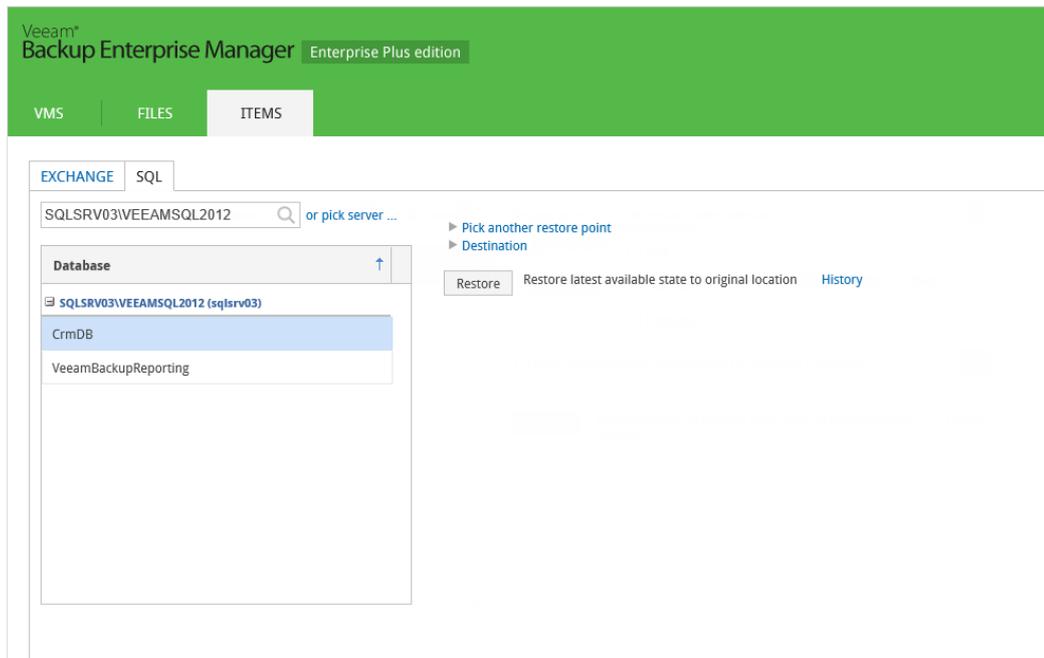


You can also restore multiple databases at once – for that, select the required server or instance in the navigation tree on the left, then:

- If restoring multiple databases of the selected SQL Server, select **Restore <current_state> to <server_name>** - in this case, all databases will be restored to corresponding instances on that (original) server.
- If restoring multiple databases of the selected instance, select **Restore <current_state> to <server_name>\<instance_name>** - in this case, all databases will be restored to the specified instance on the original server.

If you are to launch 1-Click Restore in Enterprise Manager:

1. Log in to Enterprise Manager with sufficient rights, then on the Items tab click **SQL**.



2. Enter the name of SQL Server hosting the database you need to restore; use the *server_name\instance_name* format. Alternatively, click the link to **pick server** from the list of available SQL Server VM backups.
3. Select the database you need and click the **Restore** button.

Restore to Custom Location

To restore database to the state as of in currently selected restore point, but to a server different from the original, follow the steps described in [this section of Veeam Explorers document](#). If using Enterprise Manager, follow the steps described in [this section of Veeam Enterprise Manager User Guide](#).

Restoring a Database to Specific Point in Time

This scenario allows you to recover your database to specific point in time (which may differ from current SQL Server restore point). Database will be recovered to the closest VM restore point before the moment you specify, and then transaction log replay will be performed on the target server to bring the database to the necessary state. So, to implement this scenario, ensure the following:

1. SQL Server transaction logging is enabled - logging and recovery model must be set to full or bulk-logged.
2. Transaction log handling in the SQL Server VM backup job settings had been configured to keep log backups: the **Backup logs** option should be selected.
3. You have SQL Server application-aware backup created successfully.

In such a scenario you can also choose to restore to the original or different location.

- To restore via Veeam Explorer, follow the instructions [provided in Veeam Explorers documentation](#).
- To restore via Enterprise Manager, follow the instructions [provided in Enterprise Manager documentation](#).

Note:

- Point- in-time restores from replicas and backups created with Backup Copy jobs are not supported).
- Point-in-time-restores are not supported for Hyper-V VMs which were backed up with VSS **copy_only** option enabled.

Restoring a Database to State Before Selected Transaction

This scenario offers the greatest granularity of the database recovery: you can select the undesired operation from the list of database operations (available for the time period you have specified) – for example, user deletion. The database will be restored to the state prior to the transaction containing the specified operation, using log replay on target server.

A list of SQL database operations and their display names as presented to user for selection when fine-tuning database restore is provided in the [here](#).

Note:

This restore scenario can be implemented only using Veeam Explorer for Microsoft SQL Server – Enterprise Manager does not support restore to the state before transaction.

To implement such a scenario, check the following:

1. SQL Server transaction logging is enabled - logging and recovery model must be set to **Full** or **Bulk-logged**.
2. Transaction log handling in the SQL Server VM backup job settings had been configured to keep log backups: the **Backup logs** option was selected.
3. You have SQL Server application-aware backup created successfully.
4. All nodes of the same AlwaysOn Availability Group (if applicable) are located in the same time zone.

Then follow the instructions provided in the [Restoring a Database to Specific Transaction](#) section of Veeam Explorers document.

Note: Under certain circumstances, restore to the state before selected transaction may fail with the *"The specified STOPAT time is too early"* error when the restore is performed from secondary node of AlwaysOn availability group. As a workaround, perform restore from the primary node.

Tips

If you plan to restore a database which is not so highly-transactional, consider that displaying transaction records on the **Fine-tuning the Restore Point** step of the wizard may take several minutes. This happens because, by default, Veeam Veeam Explorer loads the number of transaction logs required to get 200 transactions back and 200 forward from the specified point in time. So, in case of too few transactions within one log file, more transaction logs will be loaded to reach that number. This value can be changed, as follows:

1. Go to the `%programdata%\Veeam\Backup and Replication\SQLExplorer` folder on the machine where Veeam Explorer runs, and locate or create a configuration file named **Config.xml** in that folder.
2. Locate or add the **FineTuningMaxLoadedItems** configuration parameter to the file content and specify the value you need. For example, for Veeam Explorer to load 50 transactions back and 50 forward from the selected point in time, specify the following:

```
<Veeam>
    <SQLExplorer>
        <Show FineTuneMaxLoadedItems="50" />
    </SQLExplorer>
</Veeam>
```

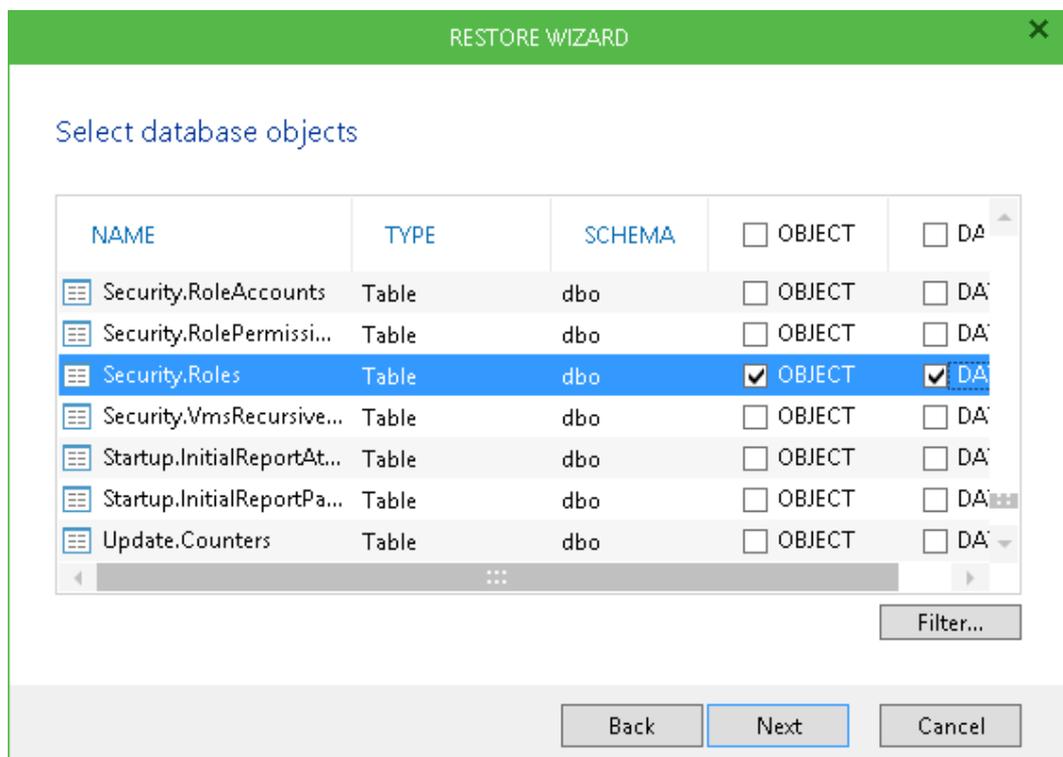
3. Save and close the configuration file.

Restoring Database Schema and Data

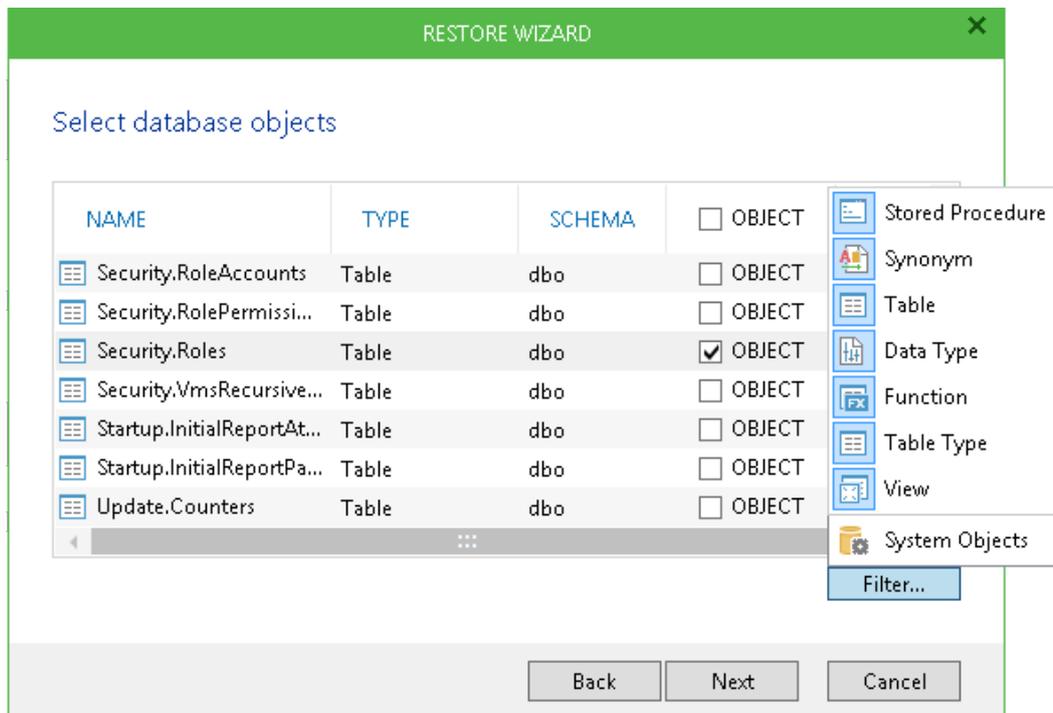
Veeam Explorer allows you to restore certain database schema objects and data, using similar scenarios (restore to the current state, to specific point in time, or to the state before selected transaction).

Do the following:

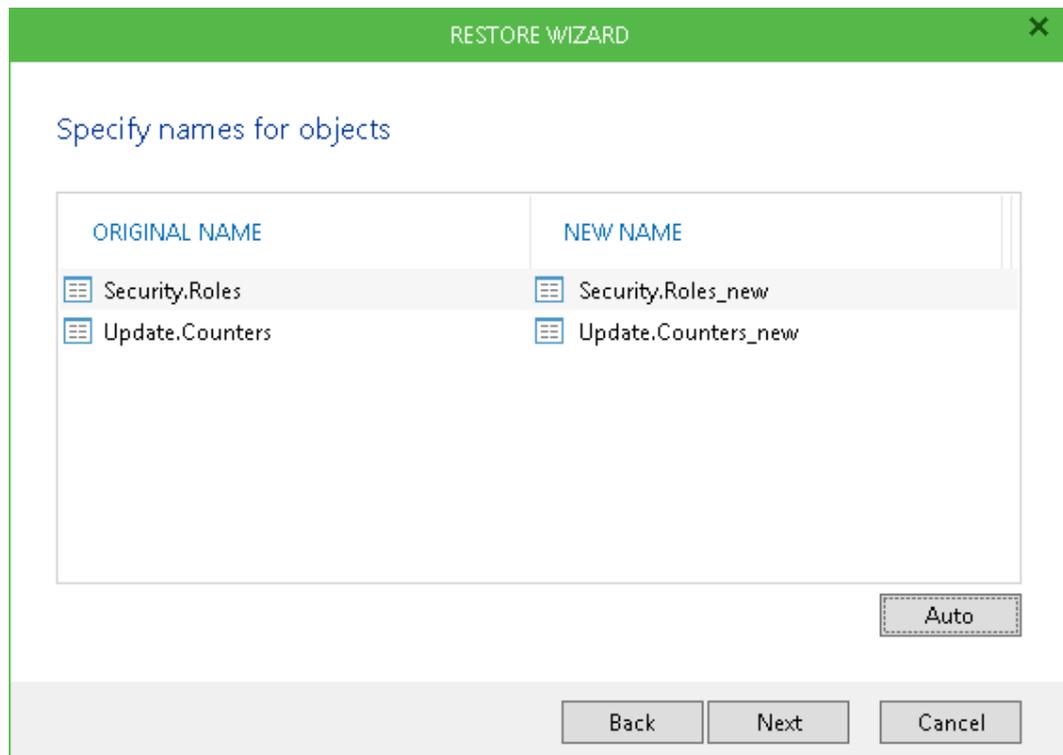
1. Click **Restore Schema > Restore database schema and data** on the toolbar, or use the database's shortcut menu command.
2. Then go through the steps of the restore wizard, specifying the necessary options for the scenario you need.
3. At the **Select database objects** step, review the list of database objects. Use the **Object** and **Data** check boxes to specify whether a database object and/or data should be restored.



To display only specific objects, click **Filters** and select the object type:



4. Specify object names for the objects after restore. When proceeding to this step, Veeam Explorer checks for object name in backup to exist in the target database:
 - If no duplicates exist, new name will be the same as the original name.
 - If any duplicates are found, you will get a message asking to provide different new name. You can click **Auto** to create new names automatically - then Veeam Explorer will use the original names followed by suffix (for example, **<name>_newsuffix**). You can edit the name if necessary and click **Auto** to check for duplicates again. If the check is a success, you can proceed to the next step.



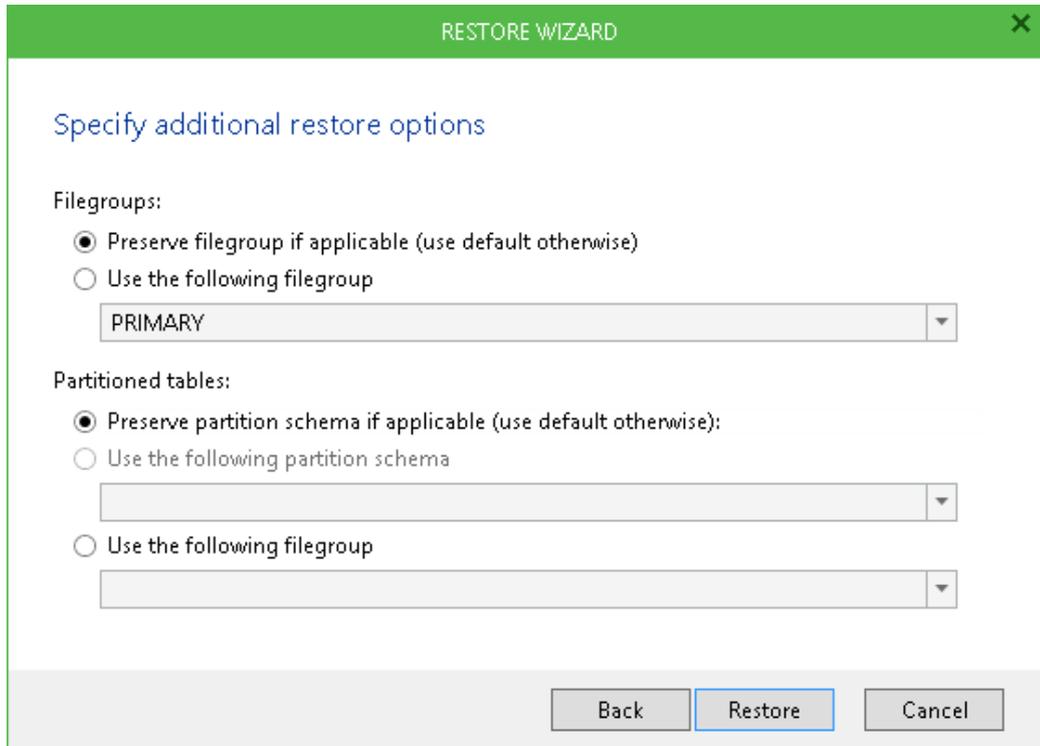
5. Next, specify additional restore options.

Select how the filegroups should be restored for selected schema object(s):

- **Preserve filegroup** - with this option selected, original filegroup will be restored. If this scenario is not applicable (for example, no such filegroup exists on target), then default filegroup configured on target will be used.
- **Use the following filegroup** - select the necessary filegroup on target.

Select how partitioned tables should be restored:

- **Preserve partition schema** - with this options selected, the table will be restored to the original partition schema (if applicable; if not - default partition schema on target will be used)
- **Use the following partition schema** - select the necessary partition schema on target.
- **Use the following filegroup** - alternatively, select the filegroup on target to be used at restore.



Then click **Restore** and wait for the restore process to complete.

APPENDIX. STATISTICS ON TRANSACTION LOG PROCESSING

In the detailed statistics window of the SQL transaction log backup, you can examine the overall statistics of logs backup (child job), as well as per-VM information.

SQL Backup Job SQL Backup

Last period (all VMs)

Databases	RPO	Status
Protected: 1	SLA: 100%	Success: 1
Unprotected: 0	Misses: 0	Warning: 0
Excluded: 3	Max delay: 00:00:00	Errors: 0

Throughput (all time)

Throughput (all time)									

NAME	STATUS
sqlsrv03	Success

Latest session

Duration:	0:00:22	Read:	188.0 KB
Bottleneck:	Target	Transferred:	7.2 KB

Last period

RPO		Sessions	
SLA:	100%	Success:	96
Misses:	0	Warning:	0
Max delay:	0:00:00	Errors:	0

Duration		Log size	
Average:	0:00:22	Average:	190.1 KB
Maximum:	0:02:10	Maximum:	256.0 KB
Sync interval:	0:15:00	Total:	17.8 MB

Errors Warnings Success

ACTION ↓	DURATION
✓ New transaction log backup interval started at 9/10/2015 9:46:09 PM (15 minutes)	
✓ Enumerating SQL Server databases	
✓ Performing transaction log backup for VEEAMSQL2012\VeeamBackupReporting	
✓ Saving transaction logs backup file to repository (VEEAMSQL2012\VeeamBackupReporting)	0:00:05
✓ Bottleneck: Target (Network); Data read: 188.0 KB; Processing rate: 8.5 KB/s	
✓ Transaction log backup interval expired	0:13:46
✓ Removing guest components	0:00:01
✓ Network traffic verification detected no corrupted blocks	
✓ Processing finished at 9/10/2015 10:00:21 PM	

Hide Details OK

Overall statistics on log backups for all VMs included in SQL Server image-level backup job ('parent') is displayed in the upper part of the statistics window.

The **Last period (all VMs)** section contains statistical data for selected session of this image-level backup job.

In the **Databases** column, you can view the following information:

- **Protected** – number of databases that were backed up at least once during the last session
- **Unprotected** – number of databases that failed to back up during the last session
- **Excluded** – databases excluded from processing. This may be due to any of the following reasons: registry key value (list of exclusions), database status is *Offline*, database recovery model is set to Simple, database is read-only, database was deleted after the latest full backup.

Note:

Unprotected databases do not comprise **Excluded**, as they have different reasons for being non-processed.

In the **RPO** column, you can view the following information:

- **SLA value** – how many log backup intervals completed in time with successful log backup (calculated as % of total number of intervals)
- **Misses** – how many intervals were missed (number of intervals)
- **Max delay** – actual difference between configured log backup interval and time actually required for log backup. If exceeded, a warning is issued.

In the **Status** column, the following information is displayed (per job): number of VMs processed successfully, with warnings or with errors.

The **Latest session** section includes the following information for the selected VM for the latest log processing interval:

- **Duration** – duration of log shipment from guest to repository by now since the start of current log processing interval
- **Bottleneck** – the operation with greatest duration in the last completed interval; this can happen due to the following:

Reason (displayed name)	Slowing-down operation
Log backup	Saving .BAK files to a temporary location on VM guest file system
Network	Uploading log files to log shipping server
Target	Saving files to target repository

- **Read** — amount of data read from temporary folder on VM guest
- **Transferred** – amount of data transferred to target repository

The **Last period** section includes the following statistics of log backups per VM for the latest sequence of log processing intervals (that is, log backup job session).

- The **RPO** column includes statistics on log processing interval (calculated as described above)
- The **Sessions** column includes statistics of log backups per VM, calculated as follows:
 - **Success** – number of intervals when all database logs were backed up successfully

- **Warning** – number of sequential intervals with failed log processing (if not more than 4 intervals in a sequence)
- **Errors** – number of sequential intervals with failed log processing (more than 4 intervals in a sequence)
- The **Duration** column includes the following information:
 - **Average** – average duration of log data transfer (through all intervals in the session)
 - **Max** – maximal duration of log data transfer (through all intervals in the session)
 - **Sync interval** – duration of periodic intervals specified for log backup in the parent job settings (default is 15 min)
- The **Log size** column includes the following information:
 - **Average** – average amount of data (through all intervals) read from VM guest
 - **Max** – maximal amount of data (over all 15-min intervals) read from VM guest
 - **Total** – total amount of data written to repository

Tip: To filter out the corresponding actions, use the **Errors**, **Warnings** and **Success** filter buttons.

Note: Statistics on transaction log processing is updated periodically, simultaneously for VM and for the log backup job

REVISION HISTORY

Revision #	Date	Change Summary
Revision 1	8/6/2017	Initial version of the document for Veeam Backup & Replication 9.5.