veeam

Granular Recovery for Enterprise Apps: Deep Dive with Veeam

Dmitry Kniazev,
Technical Analyst,
Product Strategy

Karinne Bessette,
Technologist,
Product Strategy

Contents

Introduction: why comprehensive application recovery is critical for your data protection strategy? What problem are we solving?

At the heart of a solid backup strategy is recovery. Not having a solid recovery puts not just the backup implementation at risk, but organizations are going to have bigger problems facing data loss or incomplete recoveries. Should an organization not have the application recoverability it needs, everything is at risk. That means reputation, revenue and compliance.

The perception is that application recovery is very difficult. In fact, in some scenarios one of the *worst* things you can do is restore too much data. Let's analyze this a bit further.

The first scenario to consider is that how we choose to back up will determine (to an extent) what we can recover. This paper will go into the recovery with ease made possible by Veeam® Explorers™, but it must start with the versatility that comes with the Veeam backup.

There is a challenge in the market of moving beyond the two most primitive types of recovery: whole system and file level recovery. Whole system recovery makes good sense for scenarios like hardware failures or loss of a site. File recovery makes good sense for scenarios that happen daily, including many that may be handled by a helpdesk. But what about the data? SQL Server, Oracle, Active Directory and other enterprise applications drive businesses and simply having file-level or system-level recovery may not be the best way to recover. Take Active Directory for example, if a group policy object needs to be restored, restoring the whole domain controller isn't a good way to solve that problem.

Enter the versatility of the Veeam backup format. From a Veeam backup format, recoveries of many types can be done with ease for popular enterprise applications via integrated tools called Veeam Explorers. In this paper, DK and Karinne will provide some specific technical information on how a solid backup plan can lead to application recovery. Implementing a solid backup strategy that paves the way for enterprise application recovery, down to item-level, point-in-time or even specific transactions is the expectation today.

The benefit of Veeam Explorers

After reading this white paper, you'll walk away knowing what application explorers Veeam offers, which products use Veeam Explorers and how you can use Veeam Explorers to their full potential.

Before diving into the information section, it's essential to understand why Veeam Explorers set Veeam ahead in the industry and how impactful they are to backup administration. The heart of many modern businesses run on typical applications like Microsoft Exchange to coordinate critical business communications and they use Microsoft SharePoint to collaborate on data points. These applications not only enable businesses to orchestrate internal processes but allow them to collaborate with resources outside of the intranet infrastructure as well. The loss of/or inconsistency in data access can cause disorder and leave companies open to loss in revenue or productivity. This demand for data consistency is where Veeam Explorers come into play.

Veeam Explorers integrate into back-end applications in a way that allows data to restore not in a raw data format, but in a way that leaves application consistency intact. Veeam Explorers also allow you to export data from backups in the native formats for these applications. Veeam relives some of the pressure on backup administration through Veeam Explorers by ensuring that the interface is intuitive, which makes it easier to swiftly restore data, even if the backup administrator is not familiar with the application. The Veeam Explorers make Veeam the obvious backup choice, since Veeam Explorers give administrators a fast and effortless way to restore critical business data in a format that's integrated into the application.

Veeam products that leverage Veeam Explorers

Veeam Backup & Replication

Veeam Backup & Replication is one of Veeam's flagship products. Over time, Veeam has developed many technologies which power up over 100 features of Veeam Backup & Replication. Veeam first broke into the backup and recovery world with backup copy (to help meet the 3-2-1 Rule), vPower NFS for Instant VM Recovery® from backups, then followed up with automated VM testing with SureBackup , integrated native support for over 30 different storage systems and so much more, laying out foundation for what is leading backup and data management platform for all types of workloads today. This paper goes deeper particularly into granular recovery capabilities for such renowned applications as Active Directory, SQL, Exchange, SharePoint and Oracle but it is worth mentioning Veeam Backup & Replication supports granular restore for Windows file systems, Linux file systems and native storage snapshots for many vendors. Click here to learn more.

Managed Veeam Agents

Veeam Agents are persistent agents that Veeam deploys onto guest OSes in order to orchestrate a backup from within that guest OS for physical environments. Veeam Agent software differs from Veeam Backup & Replication virtual machine (VM) backup, where it's only necessary to deploy a runtime agent on the guest OS in order to get a consistent backup, making it agentless. Veeam offers two types of Veeam Agent deployments: managed and unmanaged. The difference is, if the Veeam Agent's backup is orchestrated by the Veeam Backup & Replication server (i.e., a managed server) or if the agent is installed on a machine as a standalone device (i.e., an unmanaged server).

Veeam Backup *for Microsoft Office 365*

This product was designed to back up Microsoft Office 365 Software as a Service (SaaS) deployments for both on-premises, cloud-only and hybrid deployments. Veeam Backup *for Microsoft Office 365* currently expands to cover Microsoft Exchange, Microsoft SharePoint, Microsoft OneDrive and Microsoft Teams. The two additional Veeam Explorers, Veeam Explorer *for Microsoft OneDrive* and Veeam Explorer *for Microsoft Teams*, are specific to this product because they are specific to the Microsoft SaaS environments they back up. Veeam Backup *for Microsoft Office 365* explorers are outside the scope of this white paper.

Benefits of leveraging Veeam Explorers

Application-aware image processing

Application-aware image processing is a process that happens during backup that makes the application restore possible. As I'm sure you understand, a backup is a point-in-time preserved state of a machine and its documents made so that you can restore your documents or machine later on. There are four consistency levels to data backup.

The first is crash-consistent. Think of a crash-consistent state like if the power was pulled from the machine. If the server doesn't contain a highly transactional database, then you can expect that the device will power back on successfully, but there's always the chance you'll face a blue screen of death or corrupt applications and files due to something critical in memory being lost.

When application-aware image processing, a backup is taken in a VSS-consistent state. The VSS-consistent state backs the machine, then the backup will leverage an in-guest operation like VSS to do a memory dump and freeze applications into a shadow copy, which creates a clean image. Think of a VSS-consistent state as if you gracefully powered off a machine with the native operations of the operating system. In most cases, you can expect that, when you boot the device back up again, your files will be there and uncorrupted.

Veeam takes this a step further by using an in-guest runtime or with Veeam Agents' software-persistent agents to create an application-consistent state. For applications like Microsoft Exchange, SQL, SharePoint, Active Directory, Oracle Veeam gathers application specific data that will be used later for restore. Some examples happen when getting information — this includes SQL for database log mode and getting a consistent transaction state, domain controllers gathering the burflag key and other various state configurations, and Exchange DAG and Oracle database locations are recorded.

For applications backed up with application-aware image processing with a transaction process, Veeam records the transaction log state to create a transactionally consistent backup. When backups are done for SQL and Exchange, Veeam can be configured to truncate the transaction logs to prevent them from growing too large.

To learn more about the step-by-step process, check out this section of our user guide: https://helpcenter. veeam.com/docs/backup/vsphere/application_aware_processing.html?ver=100

Enabling this feature is very straight forward when you're creating a job. If you're already making a Veeam Backup & Replication job for a VM, you will find this setting on the guest processing tab. If you need to enable it on an existing job, right-click on that job and select edit to find guest processing.
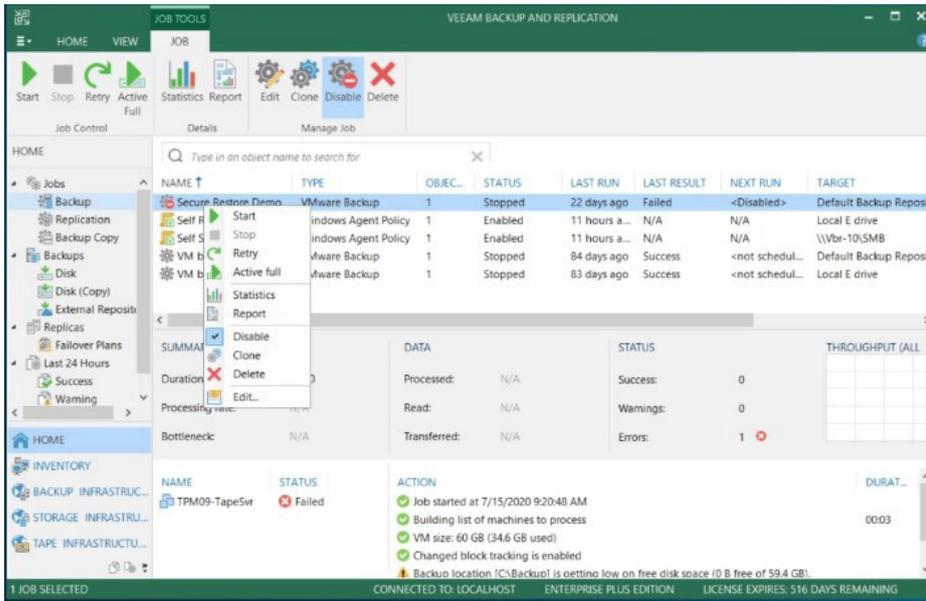
Once you've made it to the guest processing tab, you just need to check the box next to enable application-aware processing and select which guest OS credentials you need to log in and perform administrative operations on the machine.
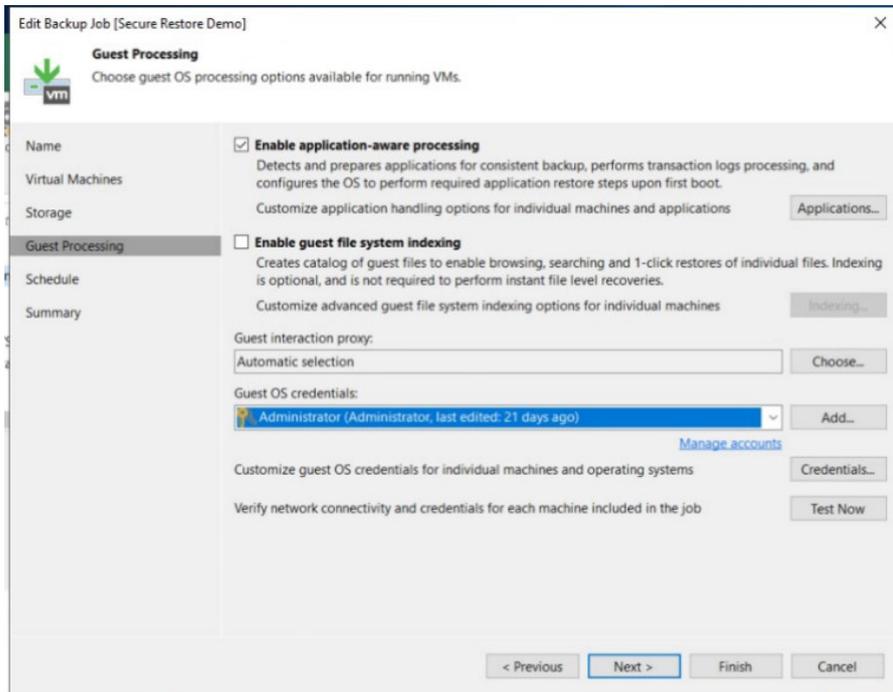
SQL and Oracle logs

Logs are records of transactions that happened to the database before its current state. These logs are native functions of the application that Veeam leverages to provide more granular restore and backup options. When backing up a server that has transaction logs configured to databases, Veeam takes extra steps to gather references to this data so that it can be used later for restores. In addition to the backup on the machine, there's also the option to configure a job that runs between machine backup job runs in order to do a log backup.

Log backups will have their own schedule outside of machine backups and will only back up transaction logs that are generated between cycles. The log backup produces a file that's labeled with the extension VLB and sits alongside the machine backup files but have their own retention cycle. Log backups are carried out by a machine called a log shipping server. Oracle and SQL both produce transaction records that serve this same function, but they do it in different ways.

However, not all SQL databases produce transaction logs, and not all databases are recommended to have transaction logs, so not all databases have transaction log backups. There are three recovery models that a SQL database can be: full, bulk-logged or simple. The only recovery models that produce transaction logs are full and bulk-logged recovery logs. Configuring your SQL database recovery models and the differences between these two transaction log collection methods are outside the scope of this white paper, SQL administrators should be familiar with this information.

The Veeam Explorers we offer

Each Veeam Explorer starts the same way, by selecting your desired restore point and going through the restore launch wizard to mount the backup point to the mount server. After the backup point is mounted, these explorers separate in their operations and go beyond just a file restore to an application restore. Depending on what application you are restoring, some additional steps or services in the background will be kicked off to emulate application-specific files. This section will expand more on what each Veeam Explorer does for each enterprise application.

Veeam Explorer *for Microsoft Active Directory*

Veeam Explorer *for Microsoft Active Directory* emulates the backup file in a way that displays the Active Directory database in a hierarchy of objects contained within. These objects include not only the groups, users and computers used but also group policy object organizational units, Active-Directory-integrated DNS records and more. This Veeam Explorer can also restore a user's Active Directory object relationship with the corresponding mailbox object.

Before launching Veeam Explorer *for Microsoft Active Directory*, you'll need to check a few things to ensure a successful launch. The operating system that the restore explorer launched from must be the same operating system version or higher than the one that's on the domain controller you plan on restoring. To emulate the Active Directory database, we use the service dynamic link library that's installed when you install Active Directory services. The esent.dll installed for these services needs to be the same version of the Active Directory service that's used on the domain controller you're installing.

Here's how it works:

1. The user launches the Active Directory application restore wizard for the desired domain controller. In the wizard, the user picks the point in time they want from the available backups.

2. Veeam mounts the backup file to the mount server that's specified in the repository settings.

3. The server in which Veeam Explorer *for Microsoft Active Directory* is launched leverages the dynamic link library service (esent.DLL) to emulate the Active Directory database from the mounted backup file.

4. Once the database file is obtained, the Veeam Explorer *for Active Directory* displays all the objects that are available for restore and export from the backup file.
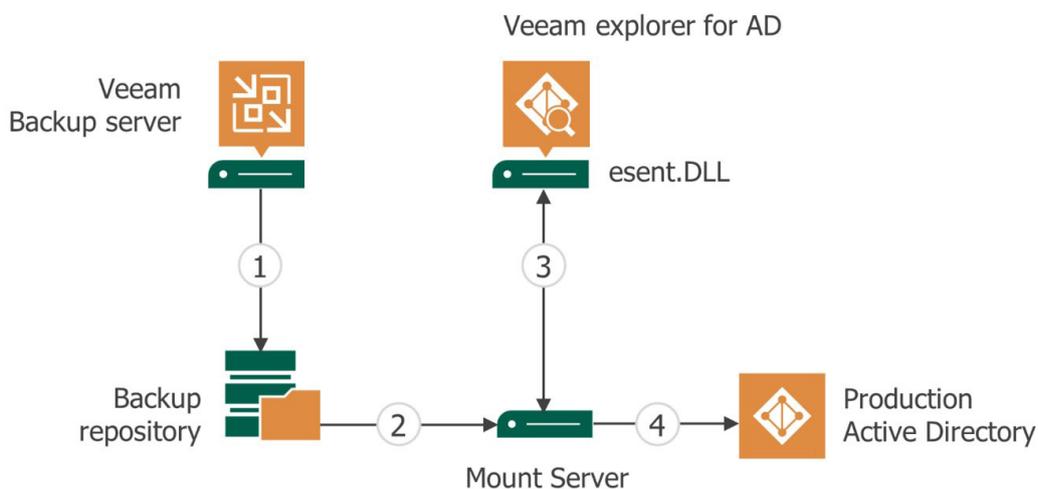
Veeam Explorer *for Microsoft SQL*

Veeam Explorer *for Microsoft SQL* is leveraged not only to restore entire SQL databases but also to specific objects and data, including stored procedures, tables, functions and more. If the SQL database logging type is full or bulk, you can also do point-in-time data restores through log replay. If you're restoring to a previous state from the time the backup was taken, then the restore would require the SQL staging server that was mentioned in the above staging server section of this paper. If a staging server is not required, the target SQL server is leveraged for the restore processes. The same logic applies if you are exporting the backup in place of restoring to a production machine.

Here's how it works without staging server:

1. The user launches the SQL application restore wizard for the SQL machine. In the wizard, there's an option to select which restore point is needed.

2. Veeam mounts the backup file to the mount server that's specified in the repository settings.

3. From the mount server, the data will be mounted to the target SQL server through a temporary iSCSI connection.

4. Veeam Explorer *for Microsoft SQL* retrieves and displays the SQL server hierarchy that's in the application explorer from the target SQL server.

5. Once the information has populated, the user can navigate to any object they need for restore or export.
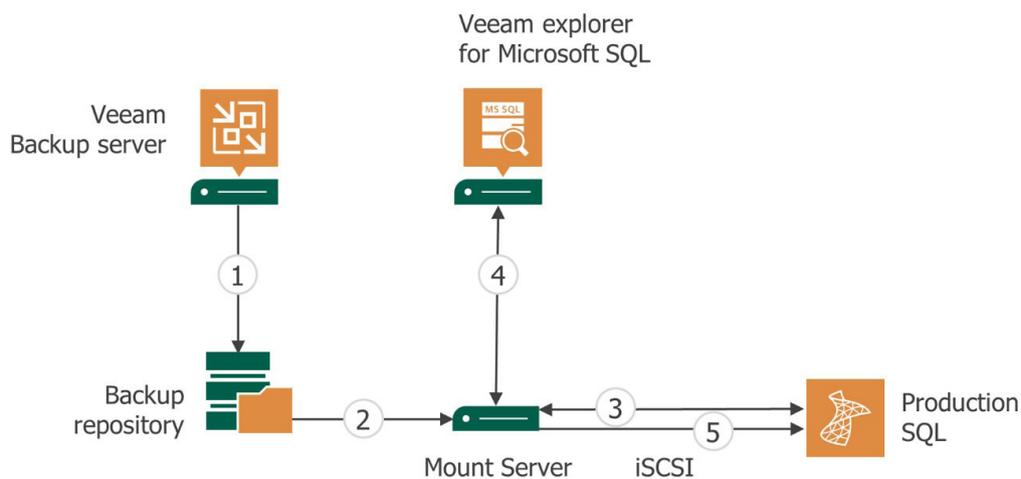
Here's how it works with a staging server:

1. The user launches the SQL application restore wizard for the SQL machine. In the wizard, there's an option to select which restore point is needed.

2. Veeam then mounts the backup file to the mount server that was specified in the repository settings.

3. From the mount server, server-mounted files are mounted to the configured SQL staging server through a temporary iSCSI connection.

4. Veeam Explorer *for Microsoft SQL* retrieves and displays the SQL server hierarchy in the application explorer from the SQL staging server.

5. Once the desired database object is selected for restore and it's required point-in-time, the data is sent from the mount server to the target SQL server with any log replays if needed.
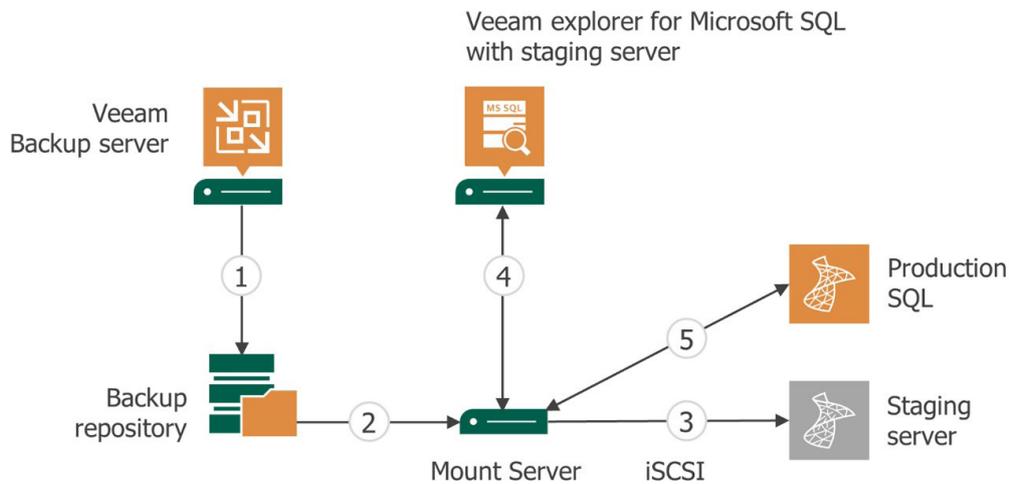
Veeam Explorer *for Microsoft Exchange*

Veeam Explorer *for Microsoft Exchange* is leveraged for Veeam Backup & Replication and Veeam Backup *for Microsoft Office 365*. The options for restore are all the same, which allows you to restore from multiple mailboxes down to individual emails. Beyond mail, Veeam can restore contacts, calendar items, notes, tasks and most items associated with Exchange. Export options for data include both .msg and .pst file formats.

After the backup is mounted to the mount server, Veeam uses an extensible storage engine. An extensible storage engine is an indexed and sequential access method technology that allows us to recover the Exchange database structure from the mounted backup file.

How it works with Veeam Backup & Replication:

1. The user launches the Exchange application restore wizard for the Exchange machine. In the wizard, there's an option to select the restore point you need.

2. Veeam mounts the backup file to the mount server that's specified in the repository settings.

3. From the mounted backup on the mount server, the Exchange database (.edb or .adb) is acquired and added to the Veeam Explorer *for Microsoft Exchange* inventory.

4. Under the stores list in the application restore console, you can see the added database(s) and content emulated by using the extensible storage engine.

5. The desired data can be selected and restored either directly back to the production Exchange server or exported as need.
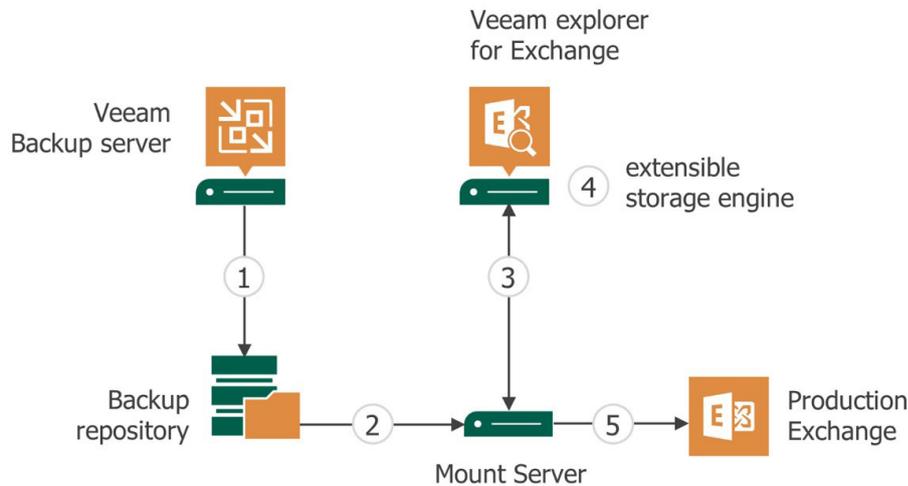
Keep in mind that, when using Veeam Explorer *for Microsoft Exchange* to restore to an on-premises Exchange server, standard rules and permissions still apply. For example, if a user's mailbox and account were deleted, to restore the mailbox to the original location, the users' Active Directory account would need to be restored first. Also, with native Exchange restore, the user account that was used for the restore process needs to have Exchange impersonation permission, even if they are a domain administrator.

Veeam Explorer *for Microsoft SharePoint*

Veeam Explorer *for Microsoft SharePoint* uses a SQL staging server to emulate the SQL backend database for SharePoint to pull the content database, which allows for the restore and export of SharePoint documents, libraries and sites. This Veeam Explorer is also used in both Veeam Backup & Replication for on-premises SharePoint and Veeam Backup *for Microsoft Office 365* for Office 365 SharePoint online. If you only need to pull data from the back-end SQL instance, you can leverage Veeam Explorer *for Microsoft SQL* to complete the restore.

How it works for Veeam Backup & Replication:

1. After selecting the machine, you want to restore from and completing the SharePoint wizard, the backup file that contains the data for the SharePoint site will be mounted to the repository's specified mount server.

2. Once mounted to the mount server, an iSCSI connection is created from the SQL staging server to the mount server. After the connection is made, the SharePoint content database is mounted to the staging server.

3. Once the content database (.mdf) has been mounted to the SQL staging server, Veeam Explorer *for Microsoft SharePoint* displays the SharePoint content databases in the Veeam Explorer.

4. After the data is selected for restore, an additional mount point is made on the target SharePoint machine where the data needs to be restored. If the data is exported, there is no other mount process that takes place.

All mount processes are temporary and are removed at the end of the data transfer.
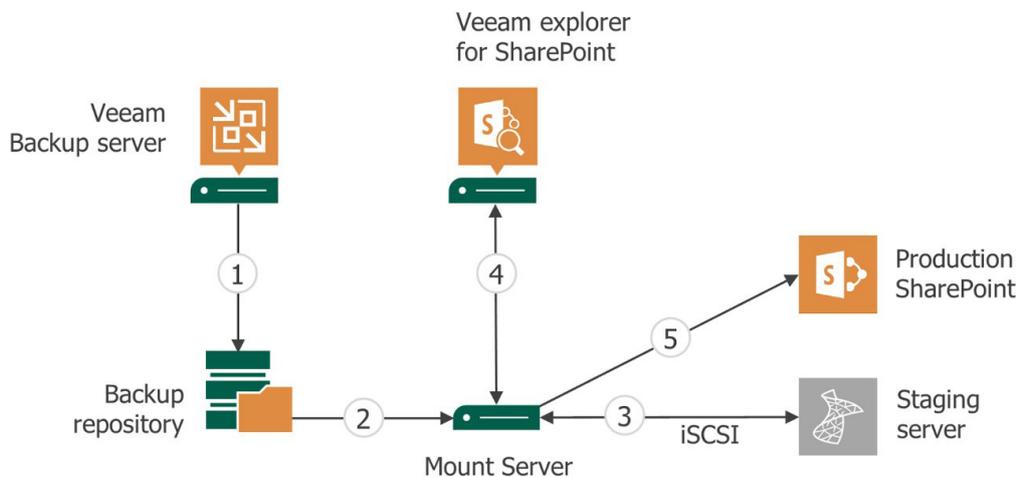
Veeam Explorer *for Oracle*

Veeam Explorer *for Oracle* can be leveraged for both the Windows and Linux versions of Oracle. This Veeam Explorer allows you to be able to restore or export Oracle databases. When restoring a database to an Oracle server, a bash shell needs to be installed on the target Oracle server to perform the restore. There are two paths you can take when you restore an Oracle database; one where an Oracle staging server is required and one where it is not. If you are not using application-aware image processing to back up the Oracle server, then a staging server is needed to pull Oracle database paths. If you're trying to restore to a point-in-time with log replays, the Oracle staging server can be used to replay the logs. The process for both Windows and Linux Oracle is the same, but in place of an iSCSI protocol, fuse is used. For restoring to the Oracle server, we leverage native RMAN (i.e., the Oracle recovery manager) to recover and recreate the database.

When Veeam Explorer *for Oracle* is launched:

1. After completing the restore wizard for Veeam Explorer *for Oracle*, a request is sent to the backup server to supply information on the Oracle databases in the backup file.

   1.1. If the job was processed with application-aware image processing, then the backup server will send back the requested data from its configuration database.

   1.2 If the job wasn't processed with application-aware image processing, the backup needs to be mounted to an Oracle staging server with either iSCSI for Windows or fuse for Linux Oracle.

2. This database information is passed to the server where Veeam Explorer *for Oracle* is running, and the databases are displayed within the console.

3. After selecting the desired database to restore, you have the option to restore this database to a production Oracle machine or export the database for manual upload.
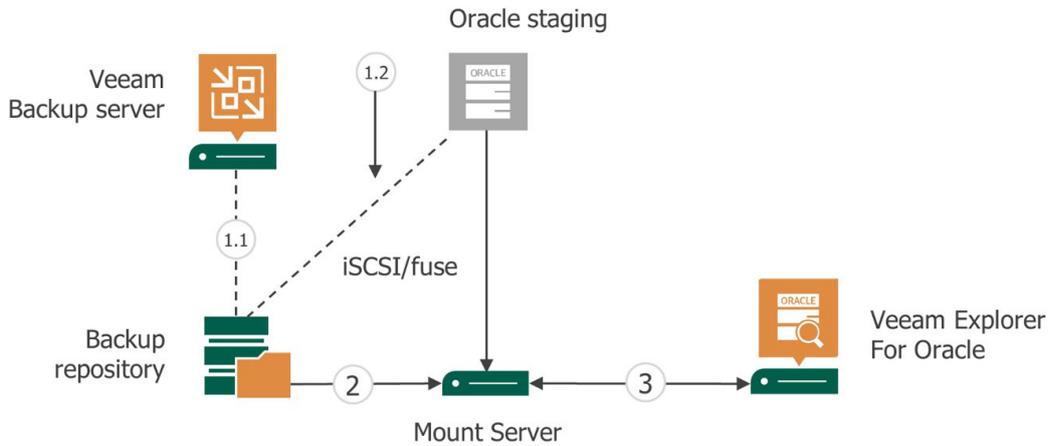
How restores are performed for a current state restore:

1. Once the database you want to restore and the target Oracle server is selected, an iSCSI (or fuse) connection is made from the mount server to the target Oracle server.

2. After the connection is complete, commands are sent from the server that's running the Veeam Explorer to the native recovery manager that's on the target Oracle server.

3. Once the database files successfully transfer to the target Oracle server, the database is initialized using the re-create function that's within the native recovery manager on the Oracle server. This process makes the database is available for use.
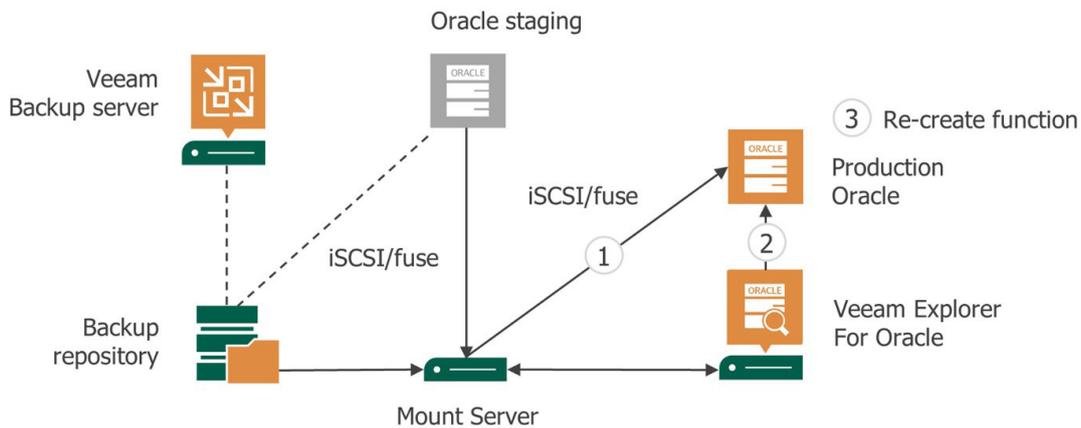
The following is how point-in-time restores are performed.

A point-in-time restore uses the redo logs to bring a database to a specified point-in-time. This means that ARCHIVELOG needs to be turned on for the database.

1. Once the database is located for restore and the desired point-in-time is selected, the wizard will locate the closest log backup point from before the time that was specified.

2. After your point-in-time is selected, a connection is established from the mount server to the target Oracle server. With this connection, the database files from the last backup of the database that happened within the time frame specified are transferred to the target Oracle server.

3. Next, the redo logs from the log backups are replayed to the restored Oracle database to bring the database back to the time specified.
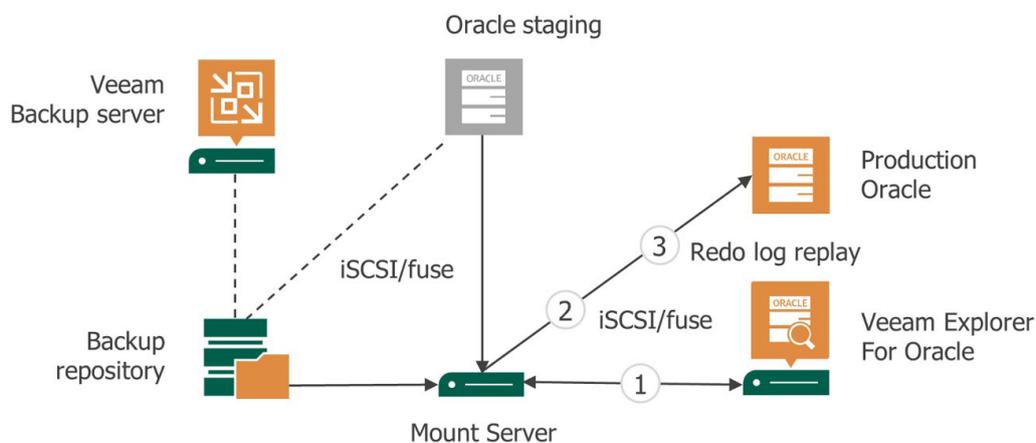
How restores are performed for a previous transaction state:

This option allows the Oracle database to be restored to a specific transaction that occurred on that database. This offers the most granular restore type but needs to use an Oracle staging server to display the transactions for selection. Here, ARCHIVELOG needs to be turned on for the database.

1. Once the database is selected for a transactional state restore, a connection is established from the mount server to the staging Oracle server. Then, the Oracle database that's within the backup is mounted to the staging server.

2. Next, the redo logs from the selected point-in-time are transferred to the Oracle staging server and replayed.

3. The transaction history is then displayed within Veeam Explorer *for Oracle* so that specific unwanted transactions can be removed from the restore scope.

4. Once the transaction restore scope is specified, a connection is made from the mount server to the target production Oracle server. Through this connection, the backed-up database files are mounted and copied over to the target Oracle server.

5. Last, the transaction logs are replayed to bring the database to the desired state in time.

Supported versions

We support almost every application version that the application vendor still has in their support cycle. Still, it's always best to check the latest supported versions list from our site to make sure Veeam Explorers support the application version that you're running. Below is the most up-to-date list of applications that Veeam supports at the time that this white paper was written, but check the user guide system requirements section ([Supported Applications]) https://helpcenter.veeam.com/docs/backup/vsphere/system_requirements.html?ver=100#guest) section before getting started to confirm what is supported at the time you're installing the software.

Check out the supported application chart here: (https://helpcenter.veeam.com/docs/backup/vsphere/system_requirements.html?ver=100#guest)

Alternatives to using Veeam Explorers

There are some scenarios in which the application that's being restored is not supported by one of our integrated restore explorers, has no applicable staging server or even cases in which application-aware processing was not enabled in the backup. When this happens, Veeam has some alternative options.

The first and easiest option would be to leverage our Instant VM Recovery or disk recovery options. Instant VM Recovery uses the same vPower technology that's found in many other features like recovery verification, staged restore, U-AIR and multi-OS file-level restore. By using Instant VM Recovery, Veeam emulates a connection from the backup on the repository to the host in which the machine in the backup is published. Once the machine is published as a VM, it can be powered on and accessed to pull out any information that uses the native application recovery. If you already have a server with the application present, you just need to access the disk where the database exists, and then instant disk recovery can be leveraged. Instant disk recovery publishes a disk to the host and can be attached to any VM.

Another option is to use Veeam DataLabs™ technology. Many of Veeam›s products and features use DataLabs, such as SureBackup, SureReplica, On-Demand Sandbox™ and Veeam Availability Orchestrator. Leveraging one of the DataLab features is the best choice when dealing with machines that have necessary dependencies. When working with Veeam DataLabs, an isolated environment is spun up where the desired machine and dependent machines are able to communicate without affecting the production environment. This is where the machines can be accessed, and the necessary data funneled out.

Lastly, if the only thing you need from a machine is a file that can be accessed and extracted from a traditional file explorer, Veeam offers file-level restore. This file restore works in two ways, one that's for Windows machine files and another that's for OS file systems, which covers most Linux and non-Windows file systems.

Summary

Thanks for taking the time reading this paper, where we presented an overview of multiple ways Veeam can provide for enterprise application recovery, down to item-level, point-in-time or even specific transactions. For any IT professional out there it's important to realize that a backup strategy is only as solid as its recoveries are. You don't want to spend time recovering the entire VM when you just need a single email, as well as you don't want to spend time thinking IF you can recover when the site failure has occurred. In today's rapidly changing world, it's imperative for your backup solution to provide a robust data protection no matter where the workloads reside or how granular you need the recovery to be performed.

Veeam Backup & Replication is a reliable 4-in-1 solution for backups, replicas, storage snapshots and CDP replicas that helps businesses worldwide to modernize data protection and eliminate downtime.

Don't take our word for it and take Veeam Backup & Replication for a spin with FREE 30-day trial with no feature limitations. Achieve faster, more flexible recovery and retention options, secure ransomware protection and unprecedented resiliency for any size business.

Resource articles

Veeam Explorers user guide

https://helpcenter.veeam.com/docs/backup/explorers/explorers_introduction.html?ver=100

Best practices for restores:

https://bp.veeam.com/vbr/VBP/4_Operations/O_Restores/

Supported applications — Veeam Agents:

https://helpcenter.veeam.com/docs/backup/agents/supported_platforms.html?ver=100

Supported applications — Veeam Backup & Replication:

https://helpcenter.veeam.com/docs/backup/vsphere/system_requirements.html?ver=100#guest

Application-aware processing

https://helpcenter.veeam.com/docs/backup/vsphere/application_aware_processing.html?ver=100

Configuring Veeam DataLabs:

https://www.veeam.com/wp-activate-data-datalabs-configuration.html

Instant VM Recovery:

https://helpcenter.veeam.com/docs/backup/vsphere/instant_recovery.html?ver=100

Oracle RMAN Plug-in:

https://helpcenter.veeam.com/docs/backup/plugins/rman_plugin.html?ver=100

SQL backup processes:

https://helpcenter.veeam.com/docs/backup/howtosql/how2_sql_understand_process.html?ver=95u4

About the Authors

Dmitry Kniazev (or just DK) is a Technical Analyst on the Product Strategy team. He has been at Veeam since 2015 helping establish Veeam Vanguard program, which is a top tier influencer community and evangelizing Veeam solutions. You can find him hosting technical sessions, webinars and podcasts. DK is a Veeam Certified Engineer and a vExpert currently focusing on Veeam Agent for Linux and Google Cloud. Follow Dmitry on Twitter @vmdk_co.

Karinne Bessette is a technologist at Veeam on the product strategy team and is based in the U.S. This individual has a strong technical background in network and security. She developed project management skills to self-manage tasks in a goal-oriented fashion and is well-versed in Office 365 and Azure platforms. At Veeam, Karinne has worked as a customer support engineer that supports Veeam products. In her time at Veeam, she has written and updated documentation, lead a team as a technical senior, and has pushed to obtain technical knowledge that's outside just the scope of her role. Her credentials include Cisco Security (CCNA) and Cisco Routing & Switching (CCNA), CompTIANetwork+, Security +, CompTIA A+, Project +, and Linux Professional Institutes Essentials, Microsoft AZ-900 and AZ-103, Veeam VMCE2019 and VMCE2020, and VMware vExpert2020. Follow Karinne on @RinBytes or @Veeam.

About Veeam

Veeam® is the leader in Backup solutions that deliver Modern Data Protection™. Veeam provides a single platform for modernizing backup, accelerating hybrid cloud, and securing data. With 400,000+ customers worldwide, including 82% of the Fortune 500 and 69% of Global 2,000, Veeam customer-satisfaction scores are the highest in the industry at 3.5x the average. Veeam's 100-percent channel ecosystem includes global partners, as well as HPE, NetApp, Cisco and Lenovo as exclusive resellers. Veeam has offices in more than 30 countries. To learn more, visit https://www.veeam.com or follow Veeam on Twitter @veeam.

veeAM

NEW

V11

Eliminate Data Loss
Eliminate Ransomware

#1 Backup and Recovery

Try now:
https://www.veeam.com/backup-replication-virtual-physical-cloud.html