10 Best Practices to Improve Recovery Objectives

V11A

Fabian Kessler,

Product Management Analyst,
Veeam Software

Contents

Introduction

When we talk about data backup and restore, we need to talk about two essential terms:

- Recovery Point Objective (RPO)
- Recovery Time Objective (RTO)

Knowing how much data loss is tolerable for business impact (RPO) and the maximum amount of time your business can tolerate to be offline (RTO) is important **(Picture 1)**. It defines how you build your backup strategy, how frequently the backup jobs will run and what type of backup is required.

This recovery objectives don't follow a general rule. Each company has other requirements for their business which defines their recovery objectives. If you want to read about RPO and RTO, our blog post "Why do recovery time and recovery point objectives matter?" explains more about RPO and RTO and how to gather the information required to define the values for your environment.

This white paper will show you best practices to help you to achieve and optimize your personal recovery objectives.

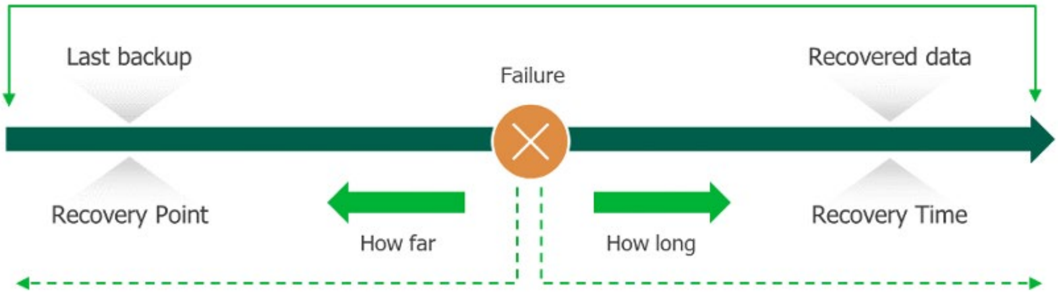

Picture 1 — RPO / RTO visualization

No. 1: Back up the right data the right way

Veeam® Backup and Replication™ can protect several types of source data from different hypervisors, physical machines, cloud environment, SaaS offerings and different types of applications. Before starting with the implementation, you need to think about the data you need to protect. The first step is to collect the characteristics of the data you want to protect. Do you need to protect applications? What sort of hypervisor is in use? How important is the data on this server for the business? What's your expected RPO/RTO for each system?

As soon as we are familiar with the production data, we can choose the right type of backup job to protect the data. We can choose from different types:

- VM-based backup
- Agent-based backup
- Enterprise application plug-ins
- NAS backups
- (Continuous) replication and Continuous Data Protection (CDP)

Consult the user guide to get familiar with the different type of jobs and their advantages and disadvantages before you make your decision.

Application-aware processing

Application-aware processing (AAIP) prepares applications on VM's or physical machines for a consistent backup.

With AAIP, depending on the application you will benefit from the following features with Veeam Backup & Replication:

- Log truncation
- Transaction log backup
- Pre-freeze/post-thaw scripts
- Item level restore directly in the console

To use AAIP, it must be enabled in the job settings and guest OS credentials with the necessary permission must be provided **(Picture 2)**.

With AAIP, you can protect many applications directly within the backup job settings.

application	operating system	backup job
Active Directory	Windows	vm and agent
Microsoft SQL	Windows	vm and agent
Microsoft Exchange	Windows	vm and agent
Microsoft Sharepoint	Windows	vm and agent
Oracle	Windows / Linux	vm and agent
MySQL	Linux	agent
PostgreSQL	Linux	agent

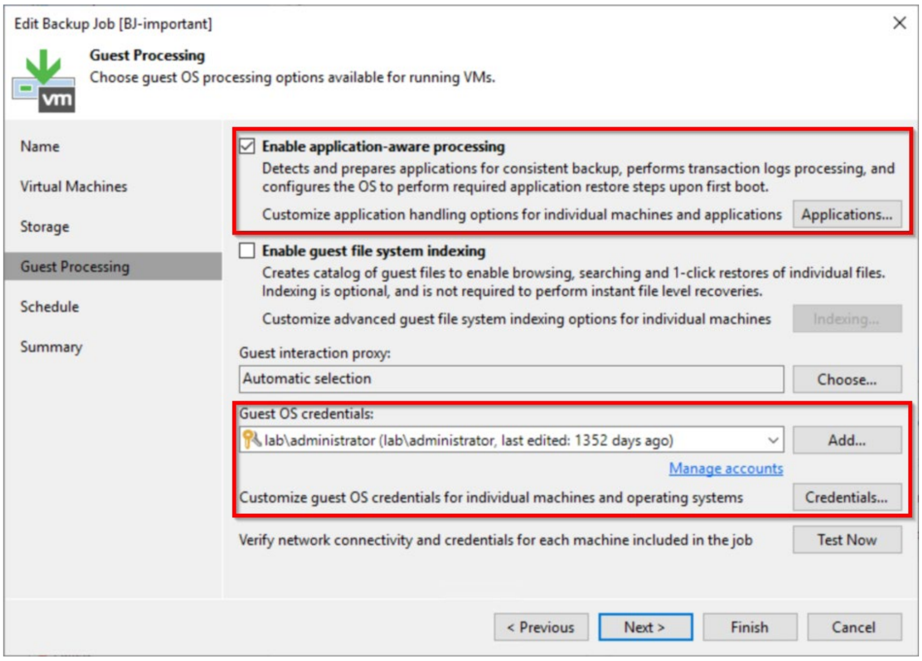

Edit Backup Job [BJ-important]

Guest Processing

Choose guest OS processing options available for running VMs.

☑ **Enable application-aware processing**

Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.

Customize application handling options for individual machines and applications

☐ **Enable guest file system indexing**

Creates catalog of guest files to enable browsing, searching and 1-click restores of individual files. Indexing is optional, and is not required to perform instant file level recoveries.

Customize advanced guest file system indexing options for individual machines

Guest interaction proxy: Automatic selection

Guest OS credentials: lab\administrator (lab\administrator, last edited: 1352 days ago)

Customize guest OS credentials for individual machines and operating systems

Verify network connectivity and credentials for each machine included in the job

Picture 2 — general AAIP settings

Application-aware processing — Guest script processing

For applications which don't support Microsoft Volume Shadow Copy on Windows or don't have an integration in Linux machines, you can rely on pre-freeze/post-thaw scripts **(Picture 3)** to ensure that your applications are prepared for when the backup happens.

You may want to check with the vendor of the application if any scripts or documentation exists regarding how you can integrate Veeam with their application. For example, HCL Domino provides a step-by-step [implementation guide]() for their product.

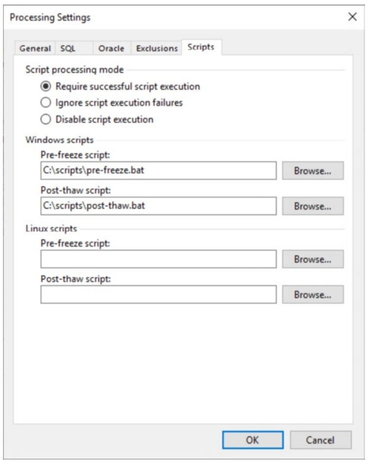

Processing Settings

Script processing mode

● Require successful script execution
○ Ignore script execution failures
○ Disable script execution

Windows scripts

Pre-freeze script: C:\scripts\pre-freeze.bat

Post-thaw script: C:\scripts\post-thaw.bat

Linux scripts

Pre-freeze script:

Post-thaw script:

Picture 3 — Guest scripting

Storage snapshots

Storage snapshots on your production storage give you low RPOs and RTOs, regardless of the backup software you use. You can orchestrate snapshots directly with the storage appliance console or with Veeam Backup & Replication. A storage snapshot does not just protect a single VM. When the snapshot is taken, all VMs on that storage volume are protected at the same time.

With Veeam Backup & Replication, you can benefit from both methods with supported storage appliances. Restoring a VM, files or application data from a storage snapshot **(Picture 4)** is as easy as it is from a Veeam backup.

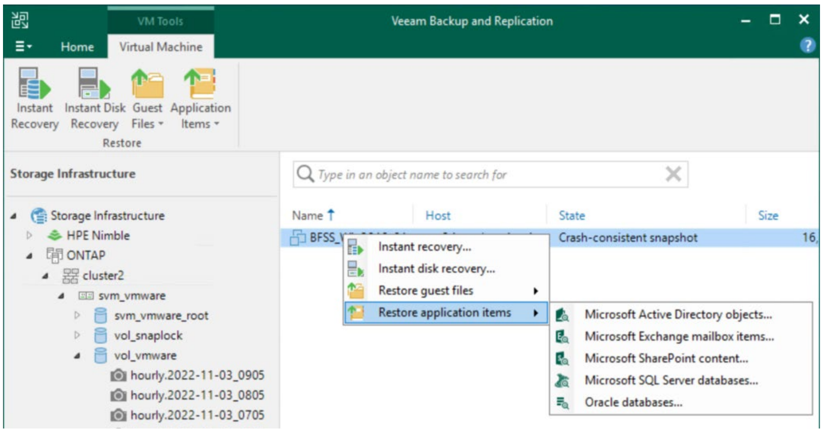

Veeam Backup and Replication

Picture 4 — Restoring from storage snapshot

Cloud-native snapshots

Like storage snapshots, you can use cloud-native snapshots for your workloads (VM, databases) in the cloud. With Veeam's cloud products, you can orchestrate this snapshot **(Picture 5)** and restore data from them.

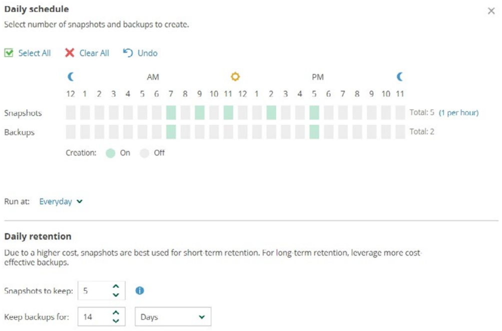

Daily schedule

Select number of snapshots and backups to create.

Snapshots — Total: 5 (1 per hour)

Backups — Total: 2

Run at: Everyday

Daily retention

Due to a higher cost, snapshots are best used for short-term retention. For long-term retention, leverage more cost-effective backups.

Snapshots to keep: 5

Keep backups for: 14 Days

Picture 5 — Cloud-native snapshots

No. 2: Use immutable backup storage

With the increasing number of ransomware attacks, it is important to use immutable or air-gapped backup storage. These days, it is no longer safe to rely solely on conventional storage. Store backups in a way that an attacker cannot delete them — that way you can protect your business from a complete data loss.

With Veeam Backup & Replication, you have several options to protect your backups from such attacks. You should choose at least one of them for your environment.

Hardened repository

A hardened repository is a repository on a Linux server that uses the immutability attribute in Linux file systems. Any supported backup file stored in this Linux repository is made immutable for the configured time. When immutability is enabled, all backups are protected for at least seven days **(Picture 6)**. GFS backups are automatically protected for the entire retention period. However, the the duration of immutability depends on the requirements for each individual organisation.

Configure the duration of immutability to be shorter than the retention period of your backup jobs or you get warnings in your backup job sessions.

There are a few requirements you need to consider when protecting the hardened repository against various attacks:

- Use a physical machine with internal hard drives. While it is possible to use a physical server with a SAN storage instead of internal hard drives, it creates additional potential security risks. The provided storage can be deleted on the admin interface of the storage appliance.
- Restrict access to SSH, remote management interfaces and physical consoles. With root permissions, an attacker can delete the whole volume / partition. This challenge exists with every software-based WORM solution. Also pre-configured appliances face this challenge: once an attacker is root on an appliance, he can delete everything.
- Consider blocking all unneeded network ports with the on-host firewall. If you want to use a hardware firewall, make sure it has the required throughput, otherwise it will be a bottleneck for your backup and recovery tasks.
- Do not install any other applications on the machine. Only the basic tools should be installed on the hardened repository. Each additional application could introduce new security risks to the server. For example, if you decide to install a monitoring agent, additional security hardening might be needed.

In addition to security precautions, it is recommended to use XFS as a file system to benefit from our Fast Clone technology. With this technology you can create synthetic full backups without additional storage consumption.

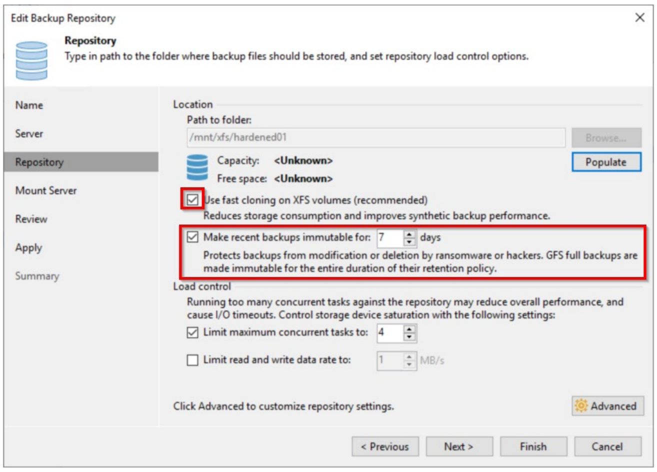

Repository

Type in path to the folder where backup files should be stored, and set repository load control options.

Location

Path to folder: /mnt/xfs/hardened01

Capacity: **<Unknown>**

Free space: **<Unknown>**

☑ Use fast cloning on XFS volumes (recommended)

Reduces storage consumption and improves synthetic backup performance.

☑ Make recent backups immutable for: 7 days

Protects backups from modification or deletion by ransomware or hackers. GFS full backups are made immutable for the entire duration of their retention policy.

Load control

Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:

☑ Limit maximum concurrent tasks to: 4

□ Limit read and write data rate to: 1 MB/s

Click Advanced to customize repository settings.

Picture 6 — Immutability for hardened repositories

Object Storage with Object Lock

Veeam Backup & Replication supports different types of object storage.

To protect the backups from accidental deletion, the object storage must support Object Lock. Object Lock enables Veeam Backup & Replication to store the backup data in such a way that it can only be deleted after the configured period. If you have decided to leverage object storage in your backup strategy, check in the Veeam Ready database which object storage appliances do support Object Lock with Veeam Backup & Replication.

[Veeam Ready — Object with immutability]

The implementation must follow our best practices:

- The data in the object storage bucket/container must be managed exclusively by Veeam, including retention and data management. Enabling lifecycle rules is unsupported and leads to backup and restore errors.
- Versioning and Object Lock must be enabled simultaneously for the bucket before you enable the immutability feature **(Picture 7)**. Any other approach will result in backup offloading errors and the inability to correctly interact with the backups in the bucket.
- Consult best practice guides from object storage manufacturers or vendors. Many object storage vendors recommend specific settings for how Veeam backup repositories and backup jobs should be configured for optimal performance.

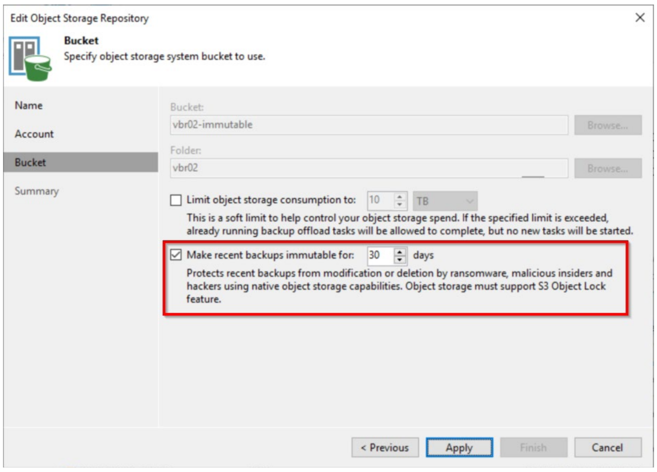

Edit Object Storage Repository

Bucket
Specify object storage system bucket to use.

Bucket: vbr02-immutable

Folder: vbr02

□ Limit object storage consumption to: 10 TB

This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.

☑ Make recent backups immutable for: 30 days

Protects recent backups from modification or deletion by ransomware, malicious insiders and hackers using native object storage capabilities. Object storage must support S3 Object Lock feature.

Picture 7 — Immutably for object storage

WORM tape

As an alternative to object storage, tapes are still a viable backup target. Tape backups give you air-gapped copies of your backups. For added protection, use WORM tapes to protect the tapes from being erased. WORM tapes can be added like standard tapes in the Veeam Backup & Replication console **(Picture 8)**.

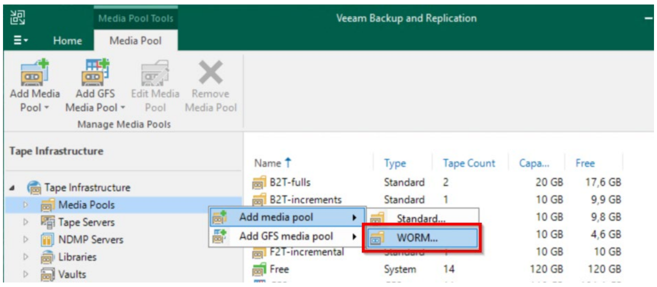

Name	Type	Tape Count	Capa...	Free
B2T-fulls	Standard	2	20 GB	17,6 GB
B2T-increments	Standard	1	10 GB	9,9 GB
			10 GB	9,8 GB
			10 GB	4,6 GB
F2T-incremental			10 GB	10 GB
Free	System	14	120 GB	120 GB

Picture 8 — WORM tapes

No. 3: Test your backups with SureBackup and Health Checks

Testing your backups is important.
Let's imagine the following scenario:

The operations team updated all the servers during a maintenance window. One of the servers only starts with a blue screen. The operations team calls the help desk and asks to restore the virtual machine from the last working state of the backup.

The help desk employee successfully restores the virtual machine. However, when booting up, he finds that the restored machine has the same problem as the original VM. After the restore, the machine still hangs with a blue screen.

Unfortunately, the help desk has to inform the operations team that a restore was not possible.

Could this have been prevented? Yes, if the recovery had been tested beforehand.

It is recommended that you regularly test your backups for recoverability. There are several reasons why a recovery should be tested. One of them is what happened in the scenario above. The virtual machine was already in a damaged state and was restored in this way. Manual testing is possible, but neither effective nor economical. With Veeam Backup & Replication, you can automate this testing. SureBackup® jobs can be used to check your backups on a regular basis.

If you have decided to use SureBackup jobs, you should consider the following points:

- SureBackup's performance depends mainly on the start-up time of the tested machines. This requires good (random) read performance from the backup repository used.
- SureBackup jobs are based on Instant VM Recovery sessions. Test your backup repository to see if it can provide the performance to run multiple VMs simultaneously with Instant VM Recovery®.
- SureBackup can be used on any Hyper-V or vSphere Host. It does not have to be the original environment. If possible, use a non-productive infrastructure. This way you do not use the CPU, memory and network capacity of the production host for your tests.

A second reason backups need to be tested is if the data on the backup storage becomes corrupted (bit rot). This situation can occur especially with unstable backup storage. While a SureBackup job with CRC testing enabled can detect such problems, you can also enable health checks within a job. If corrupted data was found, health check marks the restore point as corrupted in the database and tries to fix it with a new restore point. According to our customers, health checks run about three-four times faster in V11a than in previous versions. Running the health check on a daily or weekly basis will help you to discover corrupted backup data.

In addition to SureBackup jobs and health checks, you can use the Data Integration API to automate the tests with your own application. [V10: Reuse your data with the new Data Integration API]

Extra Tip: *With Veeam ONE™, you can create a weekly report that gives you an overview of all SureBackup jobs.*

[Recovery Verification Overview — Veeam ONE Reporting Guide]

No. 4: Ensure that the hardware is fast enough for backup and restore

To achieve the defined RPO and RTO goals, sufficient hardware is required. This applies not only to the backup hardware used, but also to components in your production environment that could become a potential bottleneck.

Potential bottlenecks

Since every backup application processes a large amount of data, it is important that the data flow is efficient and that all resources involved in the backup process are used optimally.

To meet the desired RPO and RTO, consider the following:

- Does my productive environment provide me with enough I/O and read/write speed to meet my RPO/RTO requirements?
- How is the productive environment connected to the backup environment in terms of networking?
- What transport method does my proxy server use?
- Have I allocated sufficient resources to my backup proxy?
- Do I have suitable storage as the primary backup target?

Veeam Backup & Replication provides advanced statistics on the efficiency of the data flow and lets you identify bottlenecks **(Picture 9)** in data transfer.

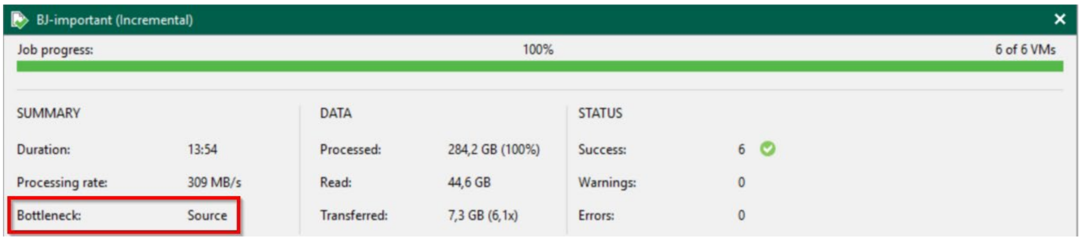

BJ-important (Incremental)

Job progress: 100% 6 of 6 VMs

SUMMARY		DATA		STATUS	
Duration:	13:54	Processed:	284,2 GB (100%)	Success:	6 ✓
Processing rate:	309 MB/s	Read:	44,6 GB	Warnings:	0
Bottleneck:	Source	Transferred:	7,3 GB (6,1x)	Errors:	0

Picture 9 — Identify bottleneck

Network

Ensure that the required bandwidth is available. The Veeam proxy server and the primary Veeam backup repository should be in the same location as the production environment to optimize backup and restore tasks.

It is not recommended to perform these tasks over a slow network connection, as this could affect the performance of the backup job or the recovery process.

Proxy server

Decide on the optimal transport method. With Veeam Backup & Replication, you can choose between three transport methods for backups and restores:

- Direct storage access (optional: Backup from storage snapshots)
- Virtual application (HotAdd)
- Network (NBD)

Perform backup and restore tests with different transport modes to find out which mode offers the best performance for your environment.

Network mode is usually the slowest method for backups and restores, especially if your ESXi servers only have 1Gbit management interfaces. If you still need to use network mode, make sure you use a 10Gig NIC for the ESXi server's management port.

SAN storage usually has poor single-stream performance but good parallel-stream performance. So, for restoring large single discs, hot-add is usually the fastest method. We recommend deploying at least one hot-add proxy for restores if you use other transport modes in your backup jobs.

For more information on our transport methods, see our white paper 10 Best Practices for VMware vSphere Backups.

Backup repository

To optimize backup and restore times, use a physical server with internal hard drives and an enterprise raid controller as the primary backup storage. This will give you the best performance for backups and restores.

If you still choose to run your backup repository on a virtual machine, consider a dedicated virtualisation environment to ensure that your backup repositories are separated from the production environment.

To optimize recovery times with a virtual repository, do not use VMFS-based discs for your backup repository on a vSphere host. You will need to set up a new ESXi host to open the VMFS datastore in the event of a complete server failure. This additional step takes time which is better invested in recovering your data.

Deduplication appliances are not recommended as primary backup storage. However, deduplication appliances can be suitable as secondary backup storage or for long-term storage.

Testing

Test backups and restores in your environment to see if your hardware is fast enough for your defined recovery goals. Detect potential bottlenecks early.

Test the best transport mode for your environment. Perform test recoveries with different restore modes before putting your backup system into production.

No. 5: Choose the right restore mode

When recovery of data or entire machines becomes necessary, you can choose between different recovery modes. Veeam offers 91 different recovery scenarios with Veeam Backup & Replication v11. These are summarised in our recovery scenario poster. Note, however, that not every recovery method is suitable for every situation. Especially when it comes to restoring entire virtual machines.

To decide on the right method of restoration, one can ask three questions:

1. What is to be restored?
2. Why am I performing this recovery?
3. How much time do I have for the recovery (RTO)?

The following examples explains a few restore method besides the normal Guest OS File level restores.

Operating system or application failure after installing updates

Most likely only the operating system drive will have issues after an update of the operating system or an application. Restoring all drives would overwrite data newer than the backup. So, only the affected drive must be restored.

The recommended option is to restore the virtual hard drive with Quick Rollback enabled **(Picture 10)**. Quick Rollback allows you to write back changed blocks only (and not the entire hard drive). It must be activated manually in the restore wizard.

NOTE: *Don't use Quick Rollback in case of a problem caused by hardware or storage issue, or due to a power loss. Quick Rollback only works for restores to the original location.*

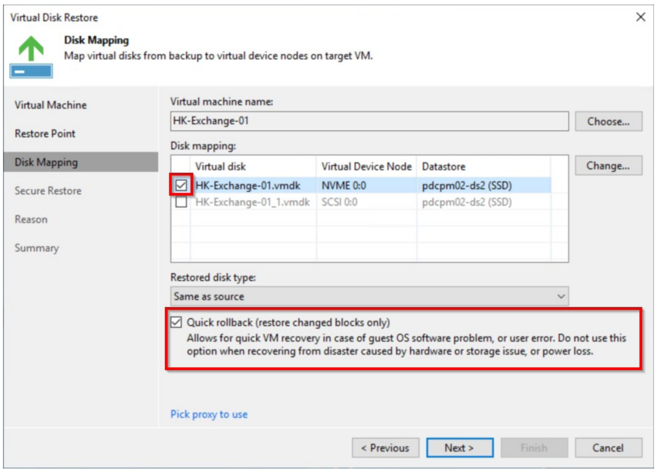

Virtual Disk Restore

Disk Mapping

Map virtual disks from backup to virtual device nodes on target VM.

Virtual machine name: HK-Exchange-01

Disk mapping:

	Virtual disk	Virtual Device Node	Datastore
☑	HK-Exchange-01.vmdk	NVME 0:0	pdcpm02-ds2 (SSD)
□	HK-Exchange-01_1.vmdk	SCSI 0:0	pdcpm02-ds2 (SSD)

Restored disk type: Same as source

☑ Quick rollback (restore changed blocks only)

Allows for quick VM recovery in case of guest OS software problem, or user error. Do not use this option when recovering from disaster caused by hardware or storage issue, or power loss.

Picture 10 — Restore with quick rollback

Malware/ransomware

If you have been attacked by ransomware or malware, it is important that you don't start immediately with the recovery. Once you have recovered the machines that you believe have been affected, there is a good chance that these machines will be reinfected.

After an attack, assume that all systems have already been affected. It is important that you shut down your backup server and production environment as soon as possible. Consult a security expert before you turn a system back on and restore the data. If you do not have your own security department, contact a security company or your security insurance company if they have one.

Also, check the contents of your backups for traces of malware. With Veeam Backup & Replication, you can leverage Instant VM Recovery Sessions or Secure Restore **(Picture 11)** to check the restore point for malware. If you use Secure Restore, the recovery time is significantly higher and chances are you will not achieve your desired recovery time goals.

If you need to restore hundreds or thousands of VMs, consider restoring everything to an isolated environment. The forensic team can check the machines individually and then release them to the production environment.

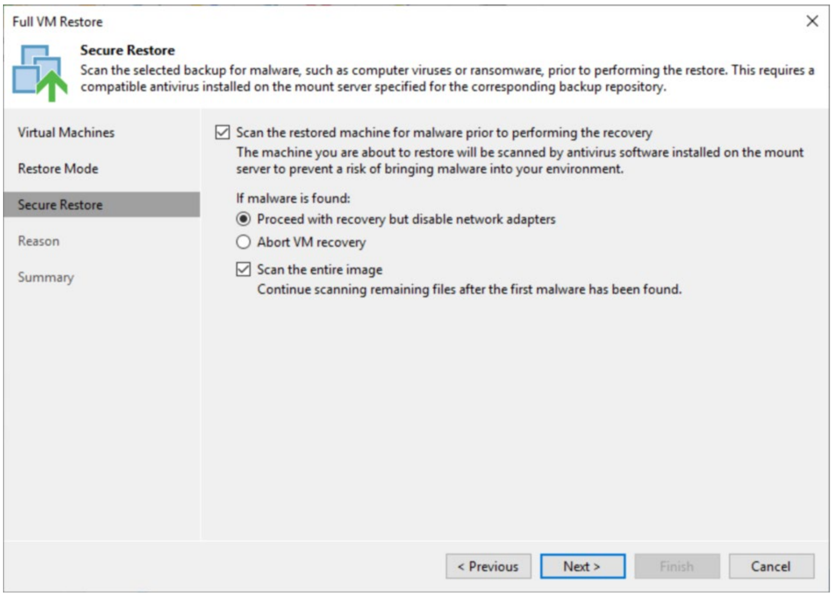

Full VM Restore

Secure Restore

Scan the selected backup for malware, such as computer viruses or ransomware, prior to performing the restore. This requires a compatible antivirus installed on the mount server specified for the corresponding backup repository.

☑ Scan the restored machine for malware prior to performing the recovery

The machine you are about to restore will be scanned by antivirus software installed on the mount server to prevent a risk of bringing malware into your environment.

If malware is found:

● Proceed with recovery but disable network adapters

○ Abort VM recovery

☑ Scan the entire image

Continue scanning remaining files after the first malware has been found.

Picture 11 — Secure restore

No. 6: Have spare hardware ready (or service)

When an emergency has occurred, there is not much time to think about where you can get the hardware to recover your data. In the event of a fire in the server room or theft, you will have lost some or all of your devices. Therefore, it is important to be prepared for this eventuality. Being prepared means thinking about which platform you want to restore your servers to and how you will access the restored data.

Depending on your business requirements and financial possibilities, you will have probably a second data center in a disaster recovery site. But, what can you do if you can't provide one?

Where do you want to restore?

Purchasing new hardware after a server failure can take time. It can take weeks or months to get new hardware. So, you need to think about alternatives to overcome this obstacle.

If you do not have replacement hardware or a second location to run your virtual machines, you can use the public cloud or find a local service provider to run your machines.

How can you access the recovered data?

Planning for recovery also means thinking about how your staff can access the data after a disaster.

We should think about the client devices and how they can access the server. For the end-user devices, there are two options.

Bring your own device/home office.

In the event of a disaster, consider having your staff work from their home office or bring their own device into the office. For these options, usage policies and security measures should have been established prior to the disaster.

Workplace as a service

In the case of a strict “no home office” policy, you can turn to a “workplace as a service” provider. With this option, you can rent a complete workplace infrastructure for your staff to be up and running in the first few days after the disaster.

No. 7: Avoid chicken-egg issues

The last thing you want to deal with when disaster strikes is chicken-egg problems. There are some important points to consider if you want to keep your RTO window small.

Encrypted backups

Your company has followed best practices to encrypt all backups. Now the following scenario occurs:

A fire has burnt down the entire building. The backup server has been lost. Fortunately, the company has replacement hardware at its remote site and the backup data still exists in its cloud-based object storage. The backup data is encrypted with a strong key so that
it cannot be accessed if it is stolen.

The password to access the encrypted backup files was stored in a password application on one of the failed servers.

The backup files are available, but the content cannot be accessed without knowing the password. Without the encryption key, their backup data is worthless. What could be done differently to avoid such a situation?

- Keep the encryption password off-site and in a safe place. The password must be available in case you have lost your backup server and you need to restore data from encrypted backup files. One option is to print out the encryption password and keep it in a safe place, like a safe deposit box or fireproof safe.
- Use password loss protection. With Veeam Enterprise Manager, authorized users can recover data even if the encryption password is lost. If the password is lost, the backup server provides a challenge key for Enterprise Manager. Using asymmetric encryption with a public/private key pair, Enterprise Manager generates a response that the backup server can use to unlock the backup file without the password being available. If you use Enterprise Manager for password recovery scenarios, you should place it in a secure location or in a cloud VM.

If password loss protection was successfully enabled, you see as an indicator the note “Loss protection enabled” **(Picture 12)** in your backup job and configuration backup settings.

Configuration backup

After the fire event, you decide to install a new Veeam Backup & Replication Server to be able to restore your environment from the cloud-based object storage. The configuration backup **(Picture 13)** will allow you to restore the configuration from the lost backup server to a newly installed backup server. To be able to do that, the configuration backup must be available in case you lose your backup server. A few ideas on how to achieve that:

- Use a repository in your secondary location for the configuration backup job
- Use File Copy Jobs to copy the configuration backup to a second location
- Use File to Tape Jobs to have a copy of the configuration backup on Tape

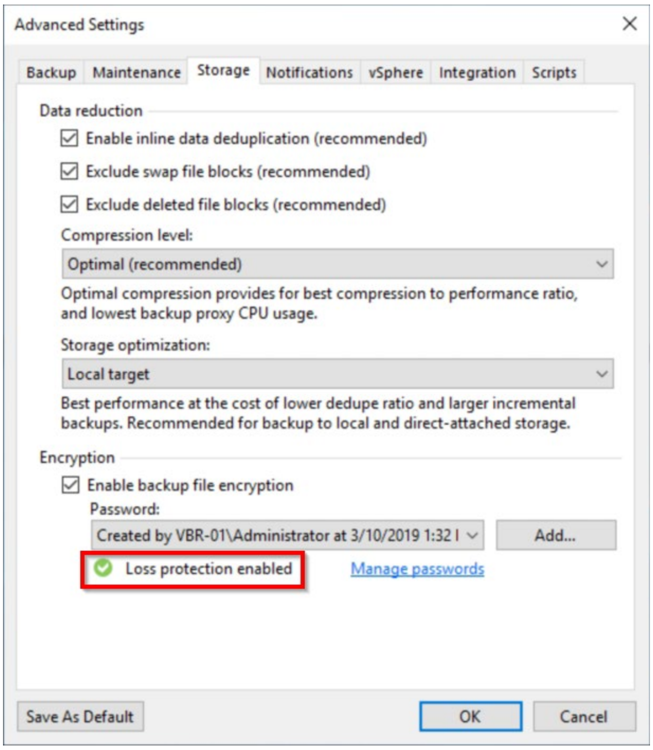

Advanced Settings

Backup | Maintenance | Storage | Notifications | vSphere | Integration | Scripts

Data reduction

☑ Enable inline data deduplication (recommended)

☑ Exclude swap file blocks (recommended)

☑ Exclude deleted file blocks (recommended)

Compression level: Optimal (recommended)

Optimal compression provides for best compression to performance ratio, and lowest backup proxy CPU usage.

Storage optimization: Local target

Best performance at the cost of lower dedupe ratio and larger incremental backups. Recommended for backup to local and direct-attached storage.

Encryption

☑ Enable backup file encryption

Password: Created by VBR-01\Administrator at 3/10/2019 1:32 I

Loss protection enabled

Picture 12 — Password loss protection enabled

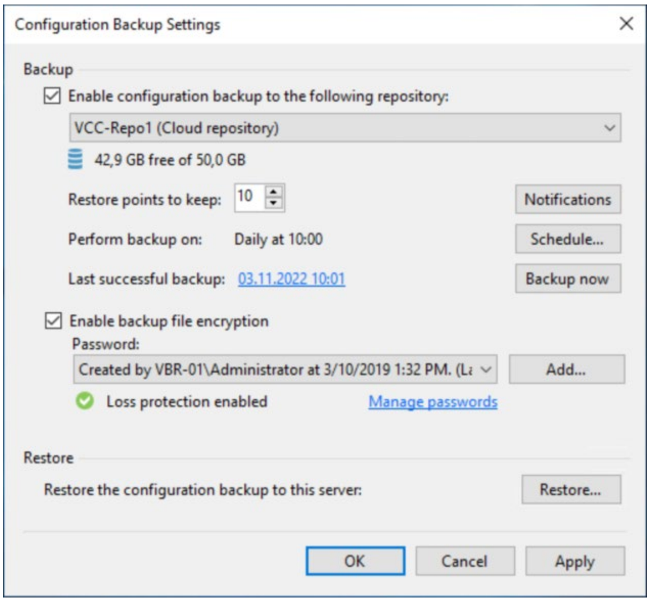

Configuration Backup Settings

Backup

☑ Enable configuration backup to the following repository:

VCC-Repo1 (Cloud repository)

42,9 GB free of 50,0 GB

Restore points to keep: 10

Perform backup on: Daily at 10:00

Last successful backup: 03.11.2022 10:01

☑ Enable backup file encryption

Password: Created by VBR-01\Administrator at 3/10/2019 1:32 PM. (Lε

Loss protection enabled

Restore

Restore the configuration backup to this server:

Picture 13 — Configuration backup

No. 8: Test a disaster recovery plan

When an emergency occurs, a quick response is expected. The right specialists, both internal and possibly external, must be ready. Replacement hardware may need to be brought into service. Servers or applications must be restored in the correct order. Access to data must be restored.

Since the recovery of applications, services or the entire data center can be stressful, it is advisable to create a disaster recovery plan. The plan should contain all the information you need to recover from a disaster that may one day occur, including:

- Servers and components that are important for the application or service
- Dependencies on other applications or services
- The correct recovery order for the components
- How to check if the recovery was successful
- How to return from emergency operation to normal operation

This plan should not only be documented, but tested, as you want to avoid faulty processes during the actual recovery. Doing that manually can be really time intensive, so it should be automatized.

Veeam offers a solution for orchestrating and testing disaster recovery plans. The product is called Veeam Disaster Recovery Orchestrator.

This product allows you to configure your documented disaster recovery plan as orchestration plan, whether you protect your environment with backup jobs **(Picture 14)**, replication jobs or storage replication. This allows you to automate the recovery and verification of your enterprise VMs, applications and services.

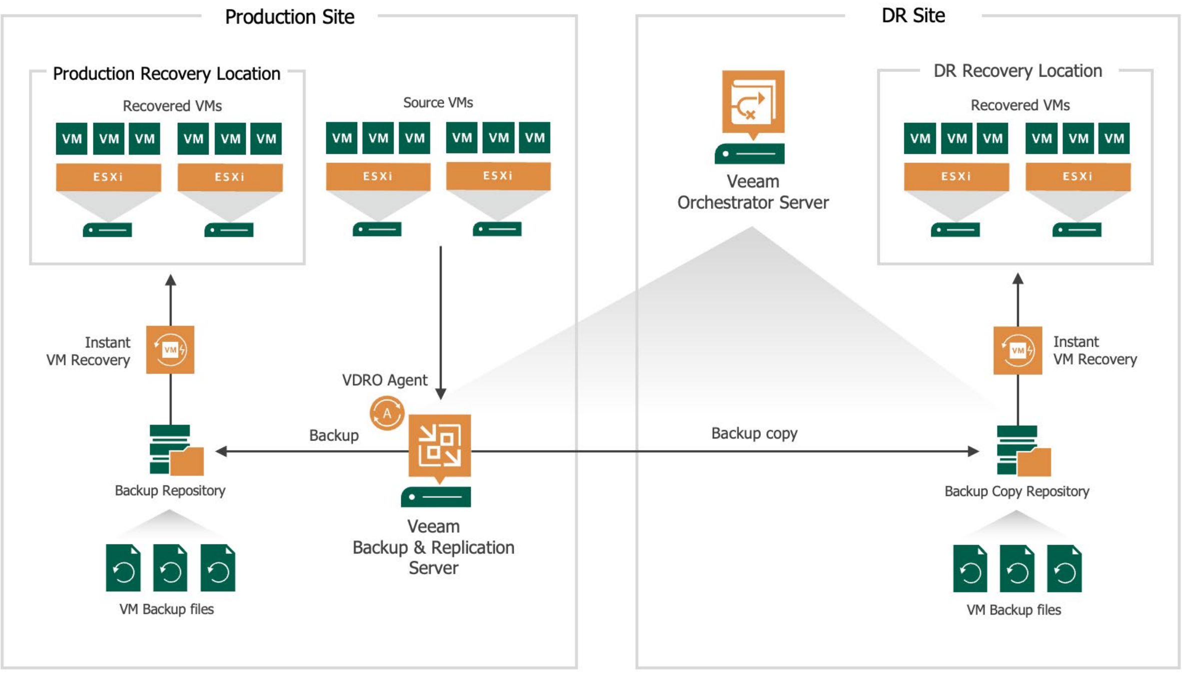

Picture 14 — Veeam Disaster Recovery Orchestrator

No. 9: Empower application owners with self-service capabilities

The restore does not necessarily have to be done by the backup administrator. It can be delegated to an application or server owner. If the backup operator performs the restore, at some point he or she must involve the application or server owner to perform a functional check and/or additional tasks before returning the server to production status. Coordinating the proper recovery of data, functional testing and additional tasks can lead to longer recovery times.

To reduce this timeframe, provide application and server owners with the right tools to perform the recovery themselves.

Veeam Enterprise Manager

With Enterprise Manager, Veeam offers a solution to delegate restores to application and VM owners. You can configure restore scopes for the users and user roles **(Picture 15)**.

Besides VM restores, Guest OS File Restore and some applications are also supported.

We recommend the following best practices for Enterprise Manager:

- If possible, Enterprise Manager should be installed at a recovery site. This makes it accessible in the event of a disaster recovery scenario. Since it is also needed for the password loss protection feature, installing it at a recovery site has advantages in case you lose the primary backup server and the password to decrypt the backup files.
- Enterprise Manager can be made virtually highly available with hypervisor features such as vSphere HA or similar technology.
- To further protect your backup server, the Enterprise Manager portal should be installed in a DMZ zone.

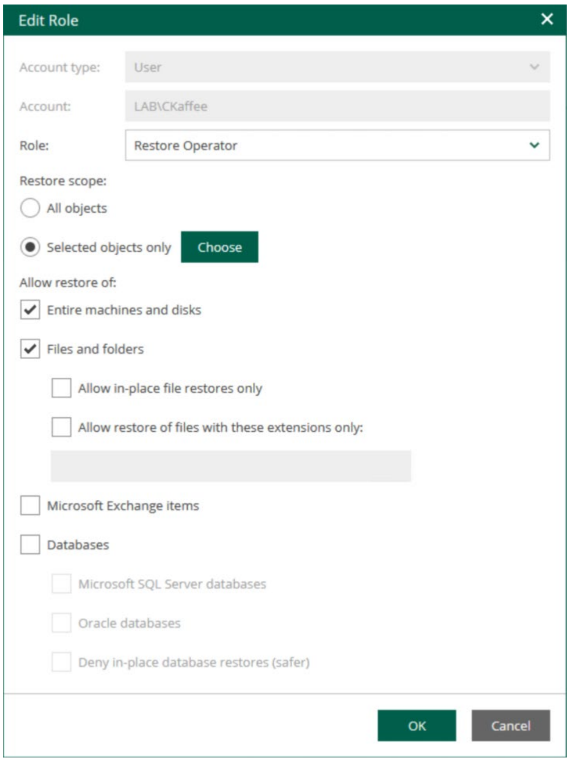

Account type: User

Account: LAB\CKaffee

Role: Restore Operator

Restore scope:

○ All objects

● Selected objects only

Allow restore of:

☑ Entire machines and disks

☑ Files and folders

□ 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)

Picture 15 — Enterprise Manager restore scope configuration

Enterprise Application Plug-ins

Veeam plug-ins for enterprise applications allow the configuration of backup jobs as well as recovery to be placed in the hands of the application owner. These plug-ins use the native backup tools of the respective application to back up and restore data from a Veeam Backup Repository.

Veeam offers plugins for various applications:

- Veeam Plug-in *for SAP HANA*
- Veeam Plug-in *for Oracle RMAN*
- Veeam Plug-in *for SAP on Oracle*

For extra protection and to get a second copy of the backups, you can use a “backup copy job” to a second repository or send the backup files to tape with a “File to Tape job”. If you have a scale-out repository, offloading to object storage with immutability is supported.

No. 10: Use current versions of Veeam

New product versions of Veeam Backup & Replication offer performance and security enhancements, as well as additional support for new versions of the platforms and applications being backed up. With each update, recovery options are optimized and new features are added.

Veeam Backup & Replication V10

With Veeam Backup & Replication V10 recovery features such as next-generation Instant Recovery engine and Instant Database Recovery have been introduced.

The next-generation Instant Recovery engine enables you to initiate Instant Recovery for multiple VMs at once with much better performance comparted to the previous version. V10 also brought Instant Recovery for individual disks. If the problem occurs only on one disk, there is no need to recover the entire VM.

Veeam Backup & Replication V11

Along with Veeam Backup & Replication V11, Continuous Data Protection (CDP) has been released. CDP allows you to replicate virtual machines from one vSphere Cluster to another vSphere Cluster with a minimum RPO of two seconds **(Picture 16)**. With CDP, you can do immediate recoveries to a latest state or desired point in time. This new opportunity minimizes data loss and eliminates downtime in your environment.

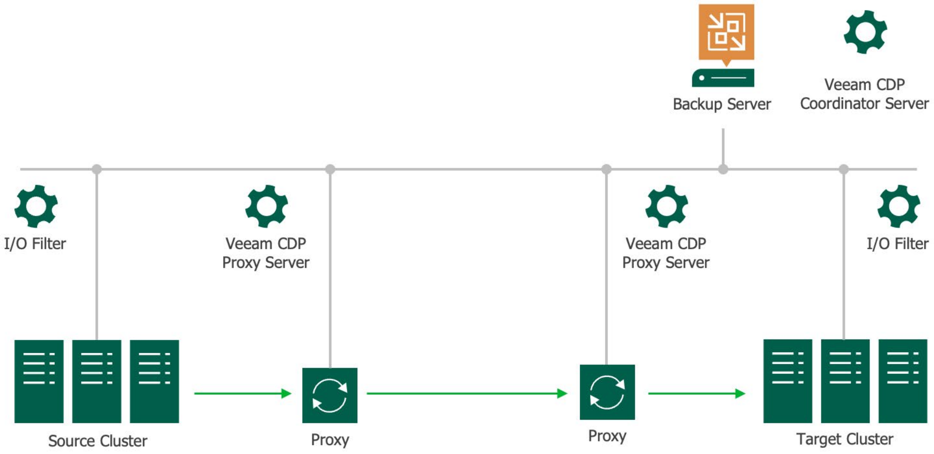

Picture 16 — Continuous data protection

Conclusion

Following these best practices will ensure that you can meet your recovery point objectives and recovery time objectives. You will be prepared to recover your data with no more data loss than expected.

Veeam's backup and recovery products can help you follow the best practices. To learn more, download a 30-day FREE trial today!

To conclude the white paper, here are the important key points summarized again:

- Use the right hardware to achieve the goals of your backup strategy
- Use immutable and/or air gapped backup storage
- Have and test your disaster recovery plan

About the Author

Fabian Kessler is a member of the Veeam product management team. Before that, he worked as a System Engineer at a Veeam Service Provider in Switzerland. There, he was responsible for designing and managing the Veeam environment (Veeam Backup & Replication, Veeam Cloud Connect and Veeam Backup *for Microsoft 365*).

About Veeam Software

Veeam® is the leader in backup, recovery and data management solutions that deliver Modern Data Protection. We provide a single platform for cloud, virtual, SaaS, Kubernetes and physical environments. Our customers are confident their apps and data are protected and always available with the most simple, flexible, reliable and powerful platform in the industry. Veeam protects over 400,000 customers worldwide, including more than 82% of the Fortune 500 and over 60% of the Global 2,000. Veeam's global ecosystem includes 35,000+ technology partners, resellers and service providers, and alliance partners and has offices in more than 30 countries. To learn more, visit www.veeam.com or follow Veeam on LinkedIn @veeamsoftware and Twitter @veeam.