Veeam SAP HANA Hyper-Availability Guide

Ali Salman

Solutions Architect, EMEA

Clemens Zerbe

Solution Architect SAP, Global Alliances

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Contents

About the Guide 3

Introduction 3

About SAP HANA Database 3
Data and Processing Layer 4
Persistence Layer 4
External Storage 4
Network Layer 4
About Veeam Backup & Replication 4
About Veeam Agent ***for Linux*** ***5***
SAP HANA on VMware 5

SAP HANA Backup & Recovery **6**

In-memory data and persistence disk 6
Why Backup 6
Backup Options & Tools 7
Backint API 7
HANA Storage Snapshot 8
Advantages and Disadvantages 9
Backint Backup 9
File Backup 9
Storage Snapshot 9
Backup Type 10
SAP HANA Recovery 10
Recovery Process 11

Veeam SAP HANA Backup & Restore **12**

Veeam Plugin for SAP HANA 12
Veeam Backup & Replication 13
Veeam Agent ***for Linux*** ***13***
Veeam SAP HANA Scripts 14
Installing Veeam Plug-in for SAP HANA 15
Configuring Veeam Plug-in for SAP HANA 15
Setting Permissions for Configuration File 16
Verifying Configuration of Veeam Plug-in for SAP HANA 16
Backup Database 17
Veeam Backup & Replication Console 21
Configuration of SAP HANA Virtual Machine Backup in Veeam 21

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam Backup Job 27
SAP HANA Studio. 27
Configuration of SAP HANA Physical Machine Backup 28
Veeam Agent ***for Linux*** Backup Job 32
SAP HANA Studio. 32
Veeam Backup & Replication Console 33
Veeam SAP HANA Recovery 40
SAP HANA Database Level Restore: 40
Virtual SAP HANA Restore — Veeam Virtual Disk-Level Restore 41
Virtual SAP HANA Restore — Instant Virtual Machine Restore 44
Performing Instant VM Recovery of SAP HANA VM 45
Physical SAP HANA Restore 52
Physical SAP HANA Recovery — Bare Metal & Volume Level 52
Physical SAP HANA Recovery — File-Level Restore 57
Restore operations from SAP HANA studio 59
Start SAP HANA System 61
Validate SAP HANA Operational Status 61

Appendix 62

About Veeam Software 62

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

About the Guide

This guide is intended for IT Administrators, Virtual Infrastructure Administrators, and any personnel responsible for the implementation of the product and operation.

This document describes the features included in Veeam® Backup & Replication™ and Veeam Agent for Linux for performing consistent backups for the SAP HANA application in accordance with Veeam best practices and SAP backup instructions.

This guide assumes that you have a good level of knowledge of Veeam Backup & Replication, Veeam Agent for Linux, SAP HANA and VMware vSphere.

Introduction

SAP HANA is an in-memory data platform that can be implemented both in the local data center (physical and virtual) and in the public cloud.

HANA is an acronym for the High-Performance Analytic Appliance and was designed so that analytical applications can be faster.

Enterprise-grade companies use SAP HANA worldwide, SAP HANA is the application directly related to customers core business functionality.

It's mission critical for the companies to have proper disaster recovery strategy for the SAP HANA. In case of any disaster, they can restore SAP HANA services and application with less efforts.

Veeam Hyper-Availability Platform™ can simplify the backup and recovery of virtual, physical and cloud workload and provides Availability to all mission-critical applications.

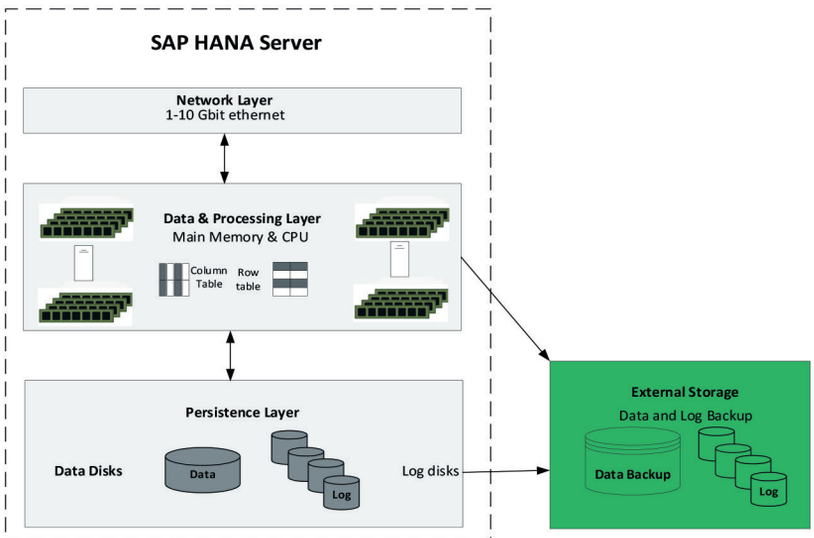
In this guide, we will explain how to backup and restore SAP HANA with Veeam Hyper-Availability Platform.

About SAP HANA Database

SAP HANA database server has three prominent internal blocks and one very important external building block:

- Network layer
- Data and processing layer
- Persistence layer
- External storage

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.



Data and Processing Layer

The data and processing layer is the most important building blocks of SAP HANA, represented in the hardware world by the main memory and the CPUs. As explained in the introduction, SAP HANA is an in-memory database, which means that the primary image of all the data is in RAM and all the internal algorithms are optimized to work on data that is present in memory.

Persistence Layer

As of any other database, SAP HANA could not be a database if it did not store data on a non-volatile medium. The SAP HANA systems come with enough disk storage to keep all the data and other required information. SAP HANA writes transactions logs synchronously and keeps a full data image in a synchronously updated data volume. The persistence layer is part of the Name Server (SYSTEMDB) and Index Server (tenant DBs) and it has data and transaction log volume for HANA system.

In other words, the persistence layer is the disk storage to store the SAP HANA database.

External Storage

An external storage medium can be used to store the logs and backup for protection along with other backup solutions.

Network Layer

SAP HANA infrastructures can be treated like other servers. Network requirements vary based on the several aspects of the system setup such as clustering. Generally, it is recommended to have a 10 Gigabit Ethernet.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

About Veeam Backup & Replication

Veeam Backup & Replication helps you achieve Hyper-Availability for ALL your virtual, physical and cloud-based workloads. With a single console, you can manage fast, flexible and reliable backup, recovery and replication of all your applications and data.

About Veeam Agent *for Linux*

As more IT organizations look to run increased amounts of Linux-based workloads in the public cloud, it is important to ensure they are backed up and can be recovered in order to avoid business disruption. However, backing up and recovering Linux instances is often cumbersome or expensive, requiring manual intervention and consuming an IT administrator's valuable time.

SAP HANA on VMware

The SAP HANA infrastructure can be virtualized on VMware Hypervisor vSphere ESXi, following is the list of VMware vSphere ESXi and SAP HANA version support matrix:

Single and Multiple SAP HANA VMs on a Single Physical Server

VMware vSphere Version	SAP HANA Release	SAP Note
vSphere 5.5(out of support since September 2018)	SAP HANA 1.0 SPS 07	1995460 and 2024433
vSphere 6.0	SAP HANA 1.0 SPS 11	2315348
vSphere 6.5 / 6.7	SAP HANA 1.0 SPS 12	2393917

SAP HANA Scale-Out Deployment Configurations

VMware vSphere Version	SAP HANA Release	SAP Note
vSphere 5.5	SAP HANA 1.0 SPS 09	2157587
vSphere 6.0	SAP HANA 1.0 SPS 09	2315348 (non-prod)
vSphere 6.5 / 6.7	SAP HANA 1.0 SPS 12	2393917

*Note: Nested virtualization is not supported by SAP. For more information, visit SAP HANA WIKI.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

SAP HANA Backup & Recovery

SAP HANA provides native tools to backup the SAP HANA database. These tools are available for SAP HANA users and backup vendors to use to take SAP HANA backup as per SAP standards.

In this section, we will discuss the available options to backup SAP HANA database and why backup is important to provide you with basic knowledge before moving on how to backup SAP HANA database by using Veeam.

In-memory data and persistence disk

SAP HANA is a memory database — all the business data resides in the main memory (e.g. master and transactional data) for maximum performance. SAP HANA also uses persistent storage to provide a fall-back in the event of fault or failure.

"SAP HANA persists in-memory data by using savepoints. During a savepoint operation, SAP HANA database flushes all changed data from memory to the data volumes. The data belonging to a savepoint represents a consistent state of the data on disk and remains so until the next savepoint operation has completed."

During normal database operation, changed data is automatically saved from memory to disk at regular savepoints. By default, savepoints are set to occur every five minutes, including during a backup.

Additionally, all data changes are recorded in the redo log buffer (the database recorded the changes into the redo log (in-memory) which will write it to log volumes consistently for all the changes to persistent data). When a database transaction is committed, the redo log buffer is saved to disk. Also, if the redo log buffer fills at any time, the redo log buffer is written to disk anyway, even if no commit has been sent.

Savepoints are also triggered automatically during following operations:

- Data backup
- Database shutdown
- Database restart

At any time, savepoints can also be trigger by executing the statement "ALTER SYSTEM SAVEPOINT."

Why Backup

There is a general misconception that backup is not required for an SAP HANA database. We will discuss why backup is required and it is mission-critical for SAP HANA disaster recovery.

As savepoints and redo log writing can protect SAP HANA databases in the event of failure, this raises the question of why would you backup if SAP HANA already has an automatic mechanism to protect from failure.

Let's focus on persistence storage which stores the savepoints and redo logs. In the case of persistent storage failure, all the savepoints and redo logs will be lost, and the database can no longer operate. This is one of the reasons why having a backup of the SAP HANA database is so important. Another example is to reset the database to an earlier point in time.

I hope you can understand the importance of the backup now. In the next section we will discuss options available to backup SAP HANA databases.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup Options & Tools

SAP HANA provides native tools to backup and restore SAP HANA databases:

- Backint API
- File System Backup
- HANA Snapshot (Storage Snapshot)

In the real world, most of the customers are using a mixed mode of backups to protect SAP HANA environments. We will cover the advantages and disadvantages of each backup mode.

Backint API

The Backint method is an API-based method — it's generally referred to SAP HANA backup using a third-party utility. It enables 3rd party tool vendors to directly connect their backup agents to the SAP HANA database. Backups are transferred via pipe from the SAP HANA database to the 3rd party backup agent, which runs on the SAP HANA database server and then sends the backups to the 3rd party backup server.

3rd party backup tools can be fully integrated with SAP HANA:

- Execution of backup and recovery from the SAP HANA cockpit, SAP HANA studio and via SAP HANA SQL commands
- Configuration of tool-specific parameters

The Backint backup can be configured through these options:

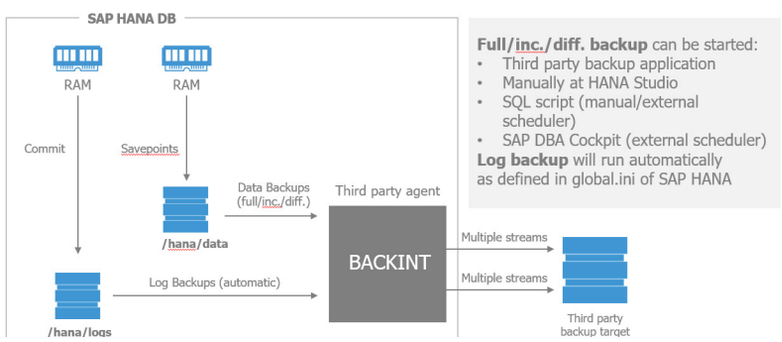
- Third-party backup application
- Manually at HANA Studio
- SQL Script (manual/external)
- SAP DB Cockpit (external scheduler)

The diagram below visualized the backup operations:

SAP HANA database stores the in-memory data into persistence storage

From the persistence storage, the third-party agent creates backup streams to the third-party backup target.

Store the backup on third-part backup storage.



File Backup

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

File System Backups protect complete data to a file. Also note that a File System Backup will put a load on the network and it is recommended to use shared backup storage so it can be available to all the nodes of the database.

You must remember to monitor file system fill level during the backup.

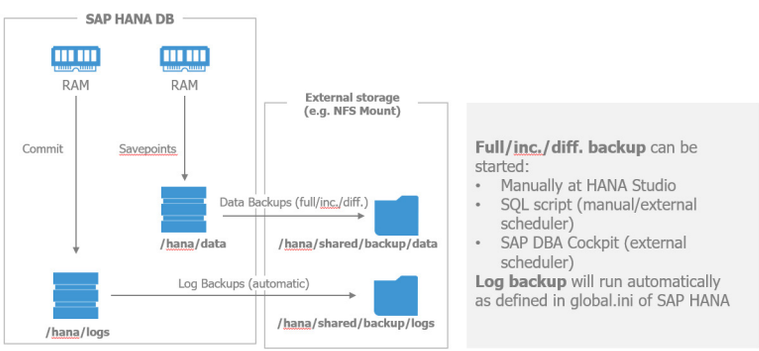
The File System Backups can be configured through these options:

- Manually at HANA Studio
- SQL Script (manual/external scheduler)
- SAP DBA Cockpit (external scheduler)
- Log backup can be set automatically

The diagram below visualized the backup operations:

1. SAP HANA database store the in-memory data into persistence storage

2. From persistence storage File System Backup store the backup file to external storage



HANA Storage Snapshot

Storage snapshots are the fastest option to safeguard the SAP HANA database without any performance overhead and impact The SAP HANA database must be online for storage snapshot backup.

The diagram below visualized the backup operations:

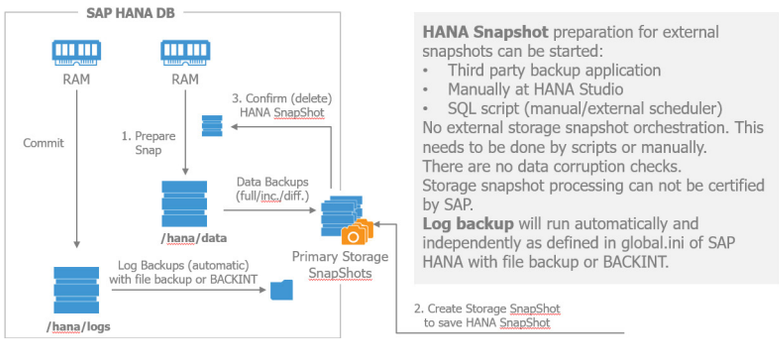
1. SAP HANA database stores the in-memory data into persistence storage

2. Prepare for a storage snapshot by initiating the SAP HANA Snapshot. During this process, the internal database snapshot is created, and reflects a consistent database state at the point in time it is created in the file system.

3. The storage snapshot is created based on the internal database snapshot

4. Confirm the SAP HANA Snapshot

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.



The method is not supported on multi-tenant environment.

Advantages and Disadvantages

We will have a close look on some advantages and disadvantages of the SAP HANA native backup tools.

Backint Backup

Advantages	Disadvantages	Backup Size	Backup Duration
Integrity checks at block level Integrated into existing data center infrastructure Backup tool offers additional features For example, encryption or de-duplication Backups are immediately available for recovery	Generates additional network load	Payload only	IO-bound (reading from data volume, writing to target) Network-bound (writing to backup server)

File Backup

Advantages	Disadvantages	Backup Size	Backup Duration
Integrity checks at block level	Requires additional storage Generates additional network load File system needs to be monitored (fill level) More time is needed to make backups available for recovery	Payload only	IO-bound (reading from data volume, writing to target) Network-bound (writing to target file system

Storage Snapshot

Advantages	Disadvantages	Backup Size	Backup Duration
Fast Generates negligible network load Can be encrypted (if data volume encryption is active)	No integrity checks at block level	Size of the data area (but is usually compressed or de-duplicated by storage tool)	Negligible

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Backup Type

The following are the type of backups available in the SAP HANA:

Option	Description
Complete	A complete data backup includes all the data that is required to recover the database to a consistent state
Incremental	Stores the data changed since the last full data backup (complete data backup or storage snapshot) or the last incremental or differential backup
Differential	Stores all the data changed since the last full data backup (complete data backup or storage snapshot)

SAP HANA Recovery

In the last section, we discussed about the options and type of backups for SAP HANA, but there is a saying, "Backup is nothing without restore." In this section we will discuss the recovery options available in SAP HANA systems.

For the recovery of SAP HANA services, it is possible to recover from a backup (full, incremental or differential), a point in time, or from the recovery of the transactions logs.

The following are the types of recovery in SAP HANA:

Option	Description
Most Recent Backup	Used for the recovery of the database to the closest possible time of the current one. For this recovery, it is necessary to have the data and logs available in the system's path to meet the objective
Point in Time	Used to recover databases at a specific point in time. For this recovery, it is necessary to have the data and logs available in the system's path to meet the objective
Specific Data Backup	Used to recover the database to a specific data backup. The specific backup file is necessary to perform this option
Specific Log Position	This type of recovery is an advanced option that can be used in exceptional cases where a previous recovery has failed

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Recovery Process

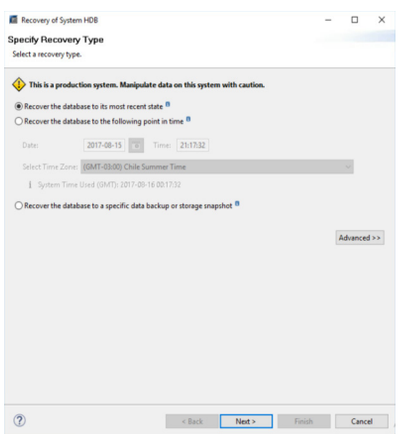
In this section we will discuss the recovery process and what options we have to restore the SAP HANA database.

We can perform the SAP HANA recovery process in the following ways by using the methods we have discussed in the last section on SAP HANA recovery options:

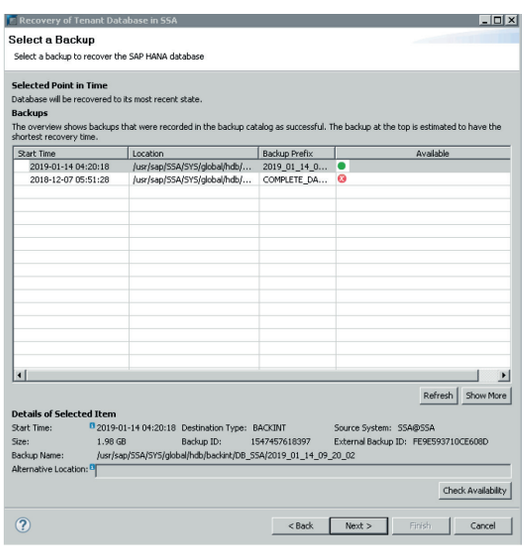
1. Restore the complete server from a backup image of the desired restore point and turn on the server.

2. Restore a point in time or most recent state of the database either by backint, a file backup or snapshot.

In this example, we will use SAP HANA studio recovery wizard to select the method and SAP HANA will automatically select the transaction log backups if it is required to recover at a point in time.



After selecting the desired restore method, SAP HANA Studio will display the backup in the catalogue and mark it with green if the restore points are available.



As we have provided the basic knowledge to understand the backup and restore process of SAP HANA, the next section will focus on how customers can leverage Veeam Hyper-Availability Platform for SAP HANA protection.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam SAP HANA Backup & Restore

Previously we described how SAP HANA backups and recovery works, in this section we will focus on the options Veeam provides to protect SAP HANA databases.

We will discuss the following Veeam products that protect SAP HANA environments:

Environment	Veeam
Virtual	Veeam Backup & Replication
Virtual & Physical	Veeam Plugin for SAP HANA
Physical	Veeam Agent *for Linux*

Veeam Plugin for SAP HANA

Veeam Plugin for SAP HANA is part of Veeam Backup & Replication Enterprise Plus Edition offers a full set of SAP HANA database protection:

- Certified SAP HANA Backint processing with the Veeam Plug-in for SAP HANA to protect the database directly: Database backup and restore.
- Protection of the SAP HANA server, HANA installation, and SAP HANA INI files with image-level or file-level functionality of Veeam Backup & Replication or Veeam Agent for Linux. Also, these backup methods offer options like SAP HANA snapshot scripting and storage snapshots.

If you want to use built-in Veeam Backup & Replication functionality, see the Veeam Backup & Replication documentation at https://www.veeam.com/documentation-guides-datasheets.html.

Use Veeam Plug-in for SAP HANA, if you want to guarantee the integrity of SAP HANA databases in your backups.

Veeam Plug-in integrates with SAP HANA's Backint. Backint is an API that allows third-party backup tools like Veeam to integrate their backup targets into the SAP HANA own backup process.

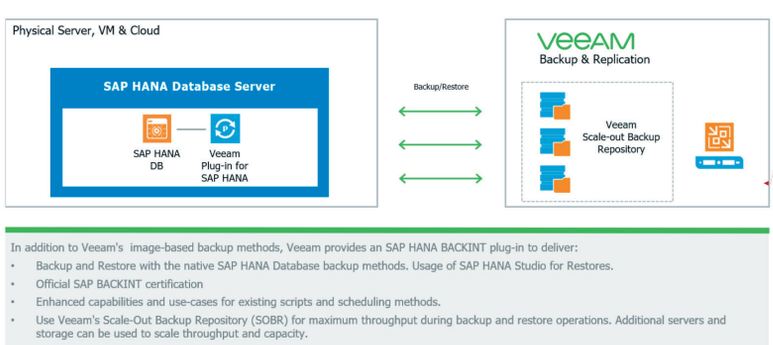
When you launch a backup with the Backint selection in a SAP HANA client, Veeam Plug-in for SAP HANA will create an application pipeline stream and transfer database files to Veeam backup repository in Veeam's proprietary format: .VAB (Veeam application backup).

The following diagram will show the backup operations:

1. HANA starts the backup.

2. Veeam Plugin for SAP HANA assigned multiple streams to backup target.

3. Data stores the Veeam backup repository.



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam Backup & Replication

Veeam Backup & Replication is the flagship product of Veeam Hyper-Availability Platform. The configuration of virtual environments is very simple as the SAP HANA server should be treated as virtual machine.

Veeam application-aware processing will use SAP HANA's **storage snapshot** method to prepare the database with pre/post freeze scripts along with the image-level backup of virtual machine and provide an option to purge the transactions logs.

The SAP HANA pre/post freeze scripts are available here:

https://github.com/VeeamHub/applications/tree/master/Freeze-Thaw%20Examples/sap-hana

The following diagram shows the backup operations:

1. Veeam Backup executes application-aware processing HANA pre-freeze script to prepare the internal snapshot of the SAP HANA database.

2. Preparation of the snapshot.

3. The data will be synchronized, the points of the transaction logs and internal SAP HANA backup catalogue will be updated to confirm the snapshot.

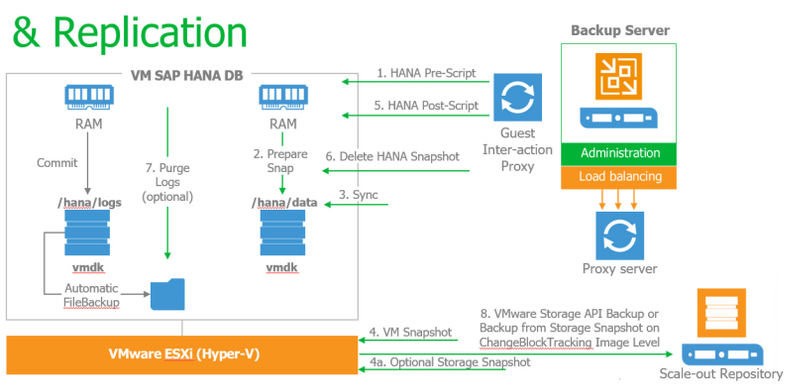
4. Veeam Backup Server will create a virtual machine snapshot.

5. Veeam Backup Server executes application-aware processing HANA post-freeze script.

6. Removal of HANA Snapshot to close the snapshot.

7. Optionally, purge the transactions logs.

8. Veeam data movers transfer the data to Veeam repository.



Veeam Agent *for Linux*

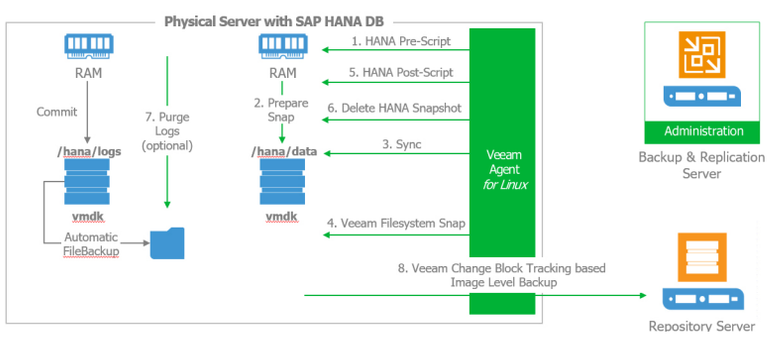
The configuration with Veeam Agent *for Linux* to configure a SAP HAN physical environment is very simple. In the same way as Veeam Backup & Replication virtual machine backup, it's possible to take image-level backup of physical servers with integration of SAP HANA via scripts.

Veeam Agent *for Linux* will use the backup method via HANA Snapshot to prepare the data on the disk made with scripts. As is known, Veeam Agent *for Linux* will perform the full backup (image) of the physical server.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

The following diagram shows the backup operations:

1. Starting the backup with Veeam Agent for Linux through the agent will execute the script "hana-pre-freeze.sh" to prepare the internal snapshot of the SAP HANA database
2. Snapshot the HANA to PREPARE the database for the backup.
3. The data will be synchronized, and the pointers of the transaction logs and the internal SAP HANA backup catalog will be updated to confirm the snapshot via CONFIRM.
4. Creation of snapshot of physical machine of the operative system.
5. Veeam Agent for Linux will execute the script "hana-post-thaw.sh"
6. Remove the HANA Snapshot.
7. If required and configured, you can purge the transaction logs as well.
8. Data transfer using "Change Block Tracking" to the Veeam repository.



Veeam SAP HANA Scripts

There are two script files to take a HANA storage snapshot backup of SAP HANA. Each script is responsible for the certain tasks as explained in the earlier section. There are additional configuration option available in the script to configure the parameters — if required, please read the instructions in the script here.

Script	Veeam
Pre-Freeze Script	Responsible for preparing the database for the HANA Snapshot.
Post-thaw Script	Responsible to delete the SAP HANA Snapshot.
HANA.Conf	Configuration file to provide parameters to SAP HANA script.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Installing Veeam Plug-in for SAP HANA

Veeam Plug-in for SAP HANA comes with the Veeam Backup & Replication installation package. If you deploy Veeam backup infrastructure for the first time, you can download Veeam Backup & Replication at: https://www.veeam.com/backup-replication-vcp-download.html

To install Veeam Plug-in, do the following:

1. Mount the Veeam Backup & Replication installation ISO file.

2. Go to the mounted ISO and copy the *.rpm file from /Plugins/SAP HANA/x64 to the SAP HANA server.

3. Run the following command to install the plug-in:

```
sudo rpm -i VeeamPluginforSAPHANA-9.5.4.2611.x86_64.rpm
```

If you have a SAP HANA scale-out cluster, repeat the installation on all cluster members.

For the installation and configuration steps, please visit: https://helpcenter.veeam.com/docs/backup/plugins/configure_sap_hana_plugin.html?ver=95u4

Configuring Veeam Plug-in for SAP HANA

The configuration of Veeam Plug-in for SAP HANA includes configuring integration settings between SAP HANA, Veeam backup server and a backup repository where the backup files will be stored. Veeam Plug-in uses SapBackintConfigTool to configure the integration settings. The wizard configures the SAP HANA Backint settings and creates the /opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml file.

To configure Veeam Plug-in for SAP HANA, do the following:

1. Launch the configuration tool with a user with database admin rights. Sudo rights are not needed if you have configured group access as described in the Required Permissions section.

```
sudo SapBackintConfigTool --wizard
```

2. Specify the Veeam Backup & Replication server (DNS name or IP address).

```
Enter backup server name or IP address: serv02.tech.local
```

3. Specify the port which will be used to communicate with the backup server. Default port: 10006.

```
Enter backup server port number: 10006
```

4. Specify credentials to authenticate with the Veeam Backup & Replication server.

```
Enter username: serv02\administrator

Enter password for serv02\administrator
```

5. Select the backup repository where you want to store the backup files. In the terminal dialog, you must enter the number of the repository from the list of available repositories.

```
Available backup repositories:

1. serv10_repo

2. serv07_repo

Enter repository number: 1

Configuration result:

SID SH2 has been configured
```

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

If you start the wizard for the first time on a SAP HANA scale-out cluster, the wizard will ask you for a cluster name that will be used within Veeam Backup & Replication user to identify the backup job for the cluster. Additional wizard runs within the SAP HANA scale-out cluster will not ask for this entry again.

Setting Permissions for Configuration File

The configuration of Veeam Plug-in for SAP HANA includes the configuration or creation of the SAP HANA Backint symlinks on all SAP HANA instances. To be able to do this for multiple SAP HANA instances at the same time, the configuration must be performed with sudo rights. Alternatively, you can use a member of the sapsys user group to configure the plug-in and set the symlink for SAP HANA instances where the user has access rights. You can repeat the wizard under another user to configure additional SAP HANA instances.

Alternatively, you can configure a Linux security group. You must add all Veeam Plug-in admins to this security group and set the following rights:

```
sudo chown root:<youradmingroup> /opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml
sudo chmod 664 /opt/veeam/VeeamPluginforSAPHANA/veeam_config.xml
```

Verifying Configuration of Veeam Plug-in for SAP HANA

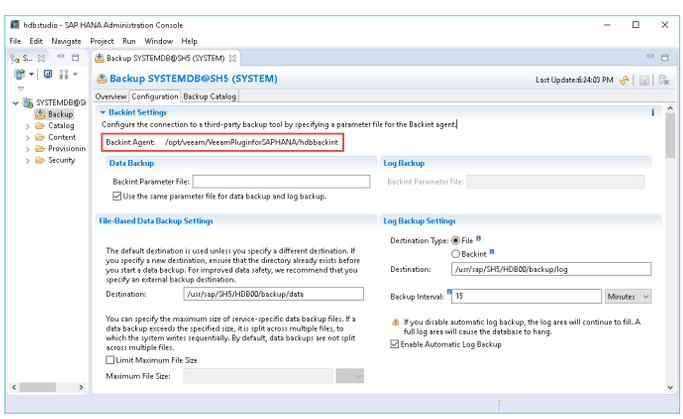
When you finish the plug-in configuration wizard, the plug-in creates a soft link in /hana/shared/<SID>/global/hdb/opt.

To verify the Backint configuration, do the following:

1. Connect to the database using SAP HANA Studio or Cockpit.

2. Go to Backup/Configuration.

3. In the Backint Agent field, make sure that the specified path leads to /opt/VeeamPluginforSAPHANA/hdbbackint.



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Changing Specific Parameter of Plug-in

If you want to edit a specific parameter of Veeam Plug-in for SAP HANA, use the *SapBackintConfigTool* command with the parameters below:

Parameter	Description
--help	Shows the list of tool parameters.
--show-config	Shows configuration parameters.
--wizard	Starts the wizard to configure the plug-in settings.
--set-credentials <"serv\username"> <password>	Specifies credentials to log in to Veeam Backup & Replication server.
--set-host <hostname>	Specifies the IP address or hostname of Veeam Backup & Replication server.
--set-port <port_number>	Specifies a port number that will be used to communicate with the Veeam Backup & Replication server.
--set-repository	Specifies the backup repository.
--set-restore-server	[for System Copy] Specifies the backup which will be copied.
--map-backup	Maps the imported backup.

Backup Database

To backup the database, there are following options:

SQL command including scripts

SAP HANA Studio

SAP HANA Cockpit

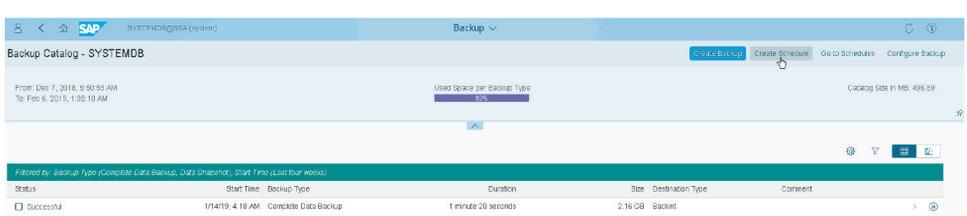
In the blog we will use SAP HANA cockpit to schedule the SAP HANA database backup, for other options and details please visit: https://helpcenter.veeam.com/docs/backup/plugins/sap_hana_backup.html?ver=95u4

Schedule Backup with SAP HANA Cockpit:

To schedule the Backint backup with SAP HANA Cockpit, do the following:

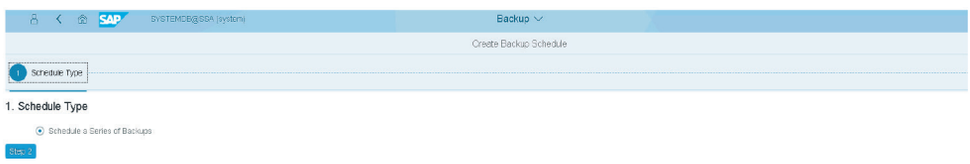
1. In the **System Overview** page, go to **Database Administration** and select **Manage Database Backups**.

2. At the **Backup Catalog** section, click **Create Schedule**

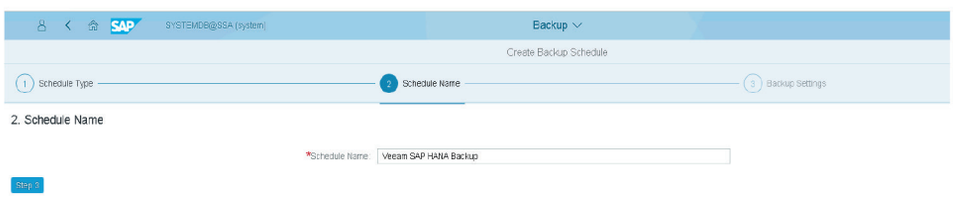


© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

3. Select **Schedule a Series of Backups**, Click **Step 2**



4. Provide the **Schedule Name**



5. Specify backup settings:

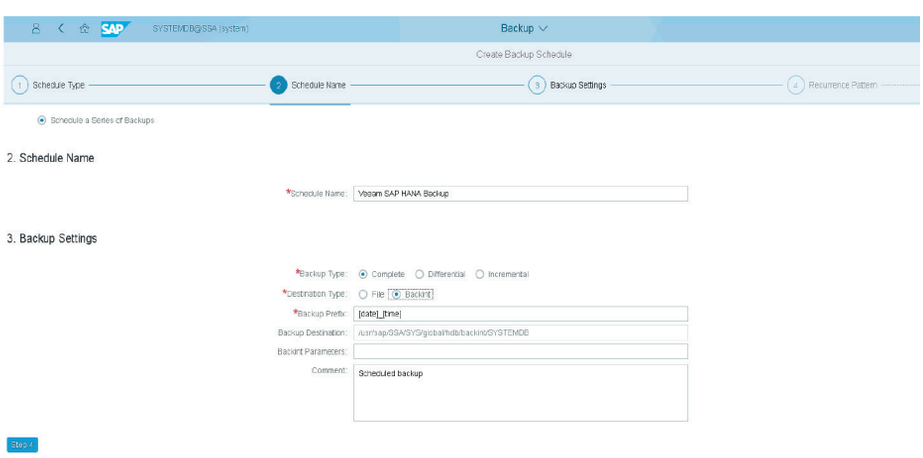
a. Select the required backup type:

- **Complete Data Backup:** A data backup includes all the data structures required to recover the database.
- **Incremental Data Backup**: An incremental backup stores the data changed since the last full data backup or the last delta backup (incremental or differential).
- **Differential Data Backup:** Differential backups store all the data changed since the last full data backup.

b. In the Destination Type setting, select Backint. With this option selected, the backup data will be written through Veeam Plug-in.

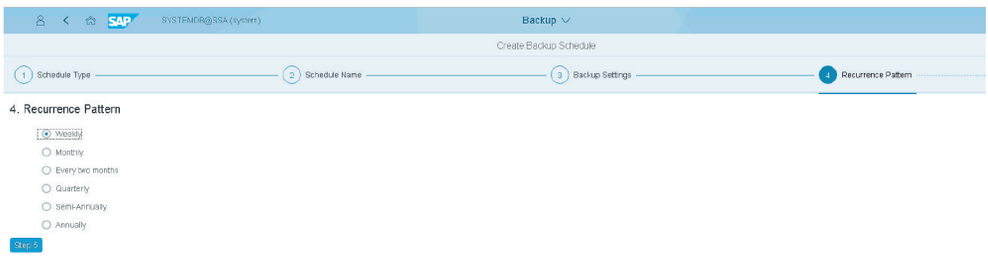
c. Change the default prefix for the backup file, if needed.

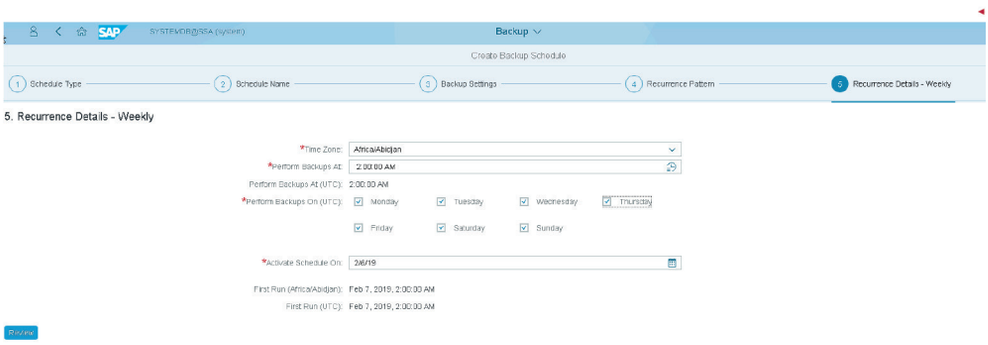
Select, **Step 4.**



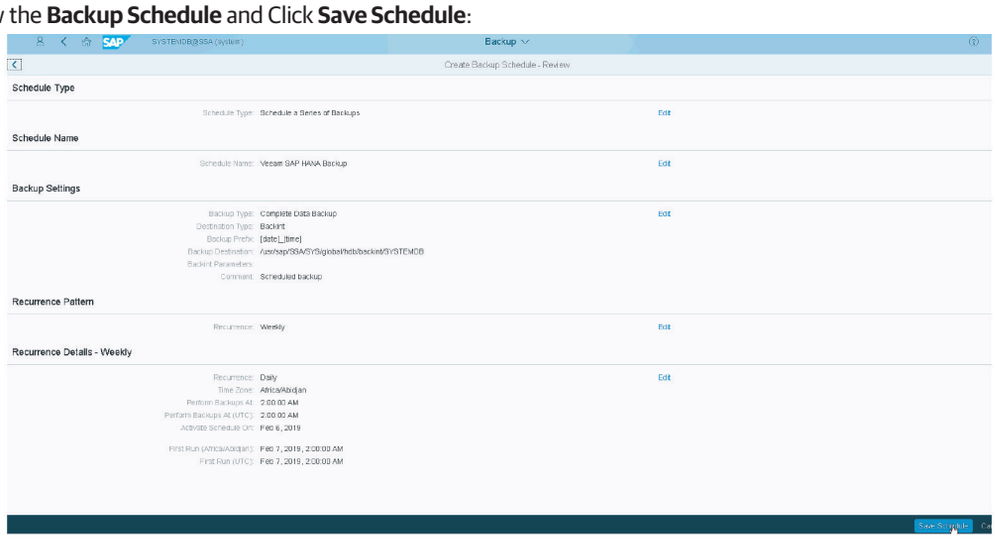
© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Select the desire **recurrence pattern** and **recurrence details**:





7. Review the **Backup Schedule** and Click **Save Schedule**:



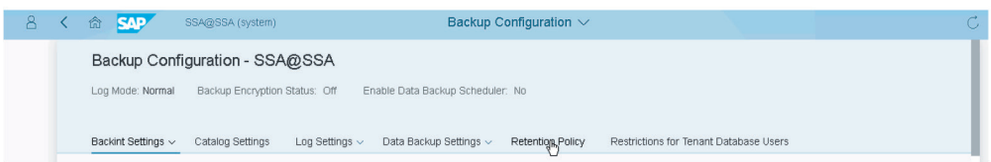
© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Retention backups with SAP HANA Cockpit:

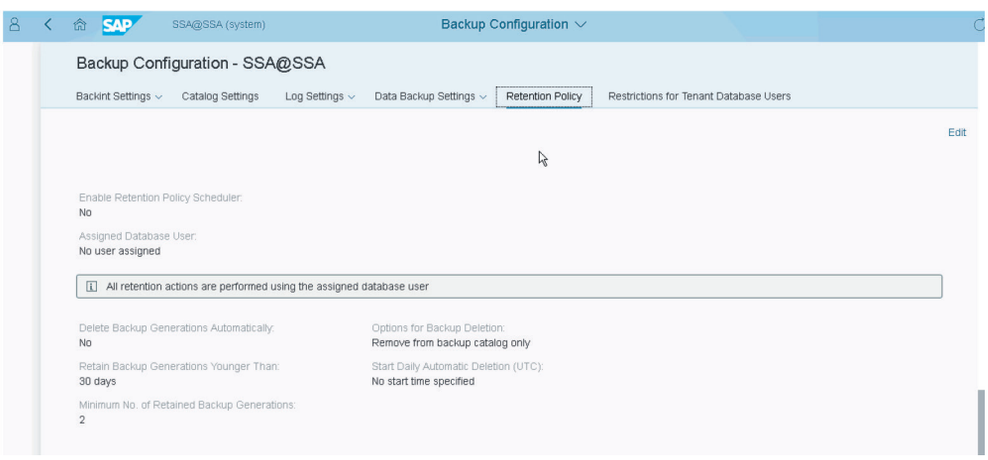
To retention backups with SAP HANA Cockpit, do the following:

1. In the **System Overview** page, go to **Database Administration** and select **Manage Database Backups**.

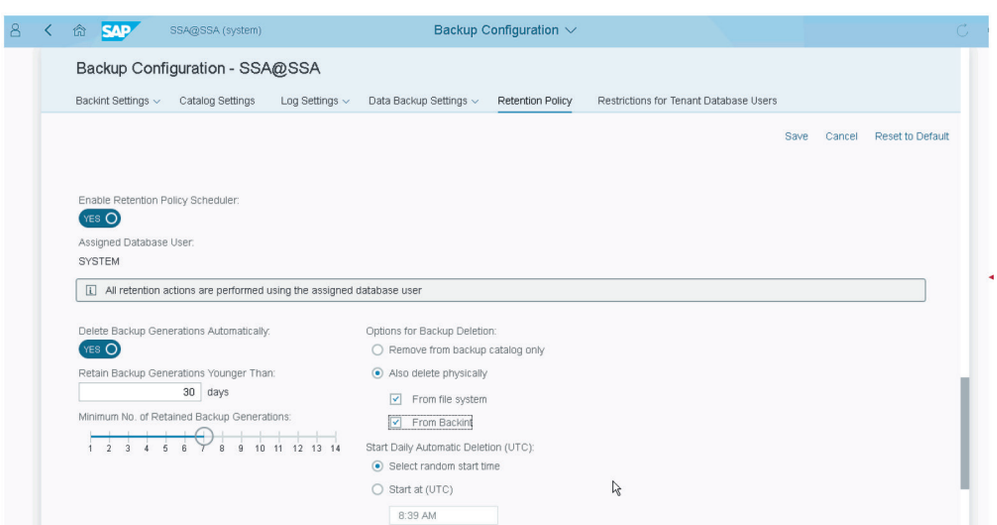
2. At the **Backup Catalog** section, click **Configure Backup.**



3. At the **Backup Configuration** section, click on **Retention Policy**.



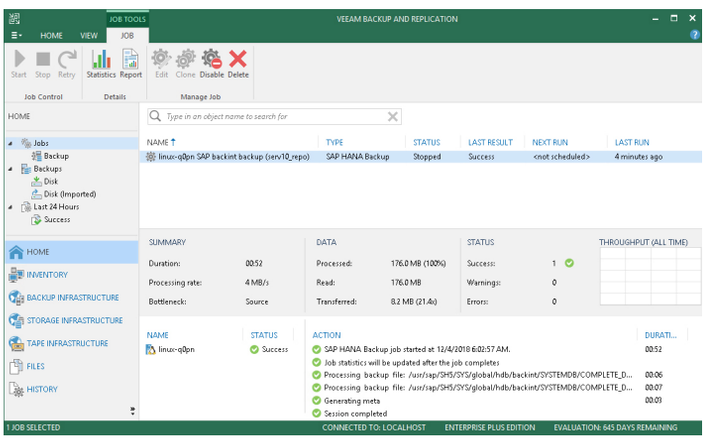
4. Change the values to your needed SLAs like e.g. 30 days retaining time for backups and at least 7 generation of backup like in the picture below. You also have the possibility to remove the backups from the catalog only or also remove the backups directly from our repository (option also delete physical from backint).



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam Backup & Replication Console

As we are using Veeam Backup & Replication repository as a target to store the backup file, we can view the progress of the backup job with the console as shown in the picture below:

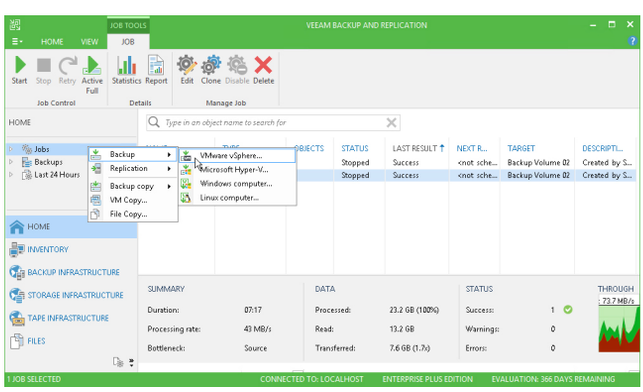


Configuration of SAP HANA Virtual Machine Backup in Veeam

In this section we will configure the SAP HANA virtual machine backup in Veeam Backup & Replication.

Step 1. Launch New Backup Job Wizard

On the **Home** tab, click **Backup Job** and select **VMware**.

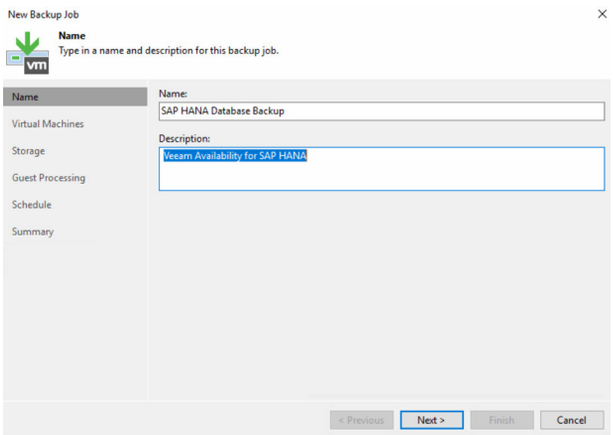


Step 2. Specify Job Name and Description

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

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

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

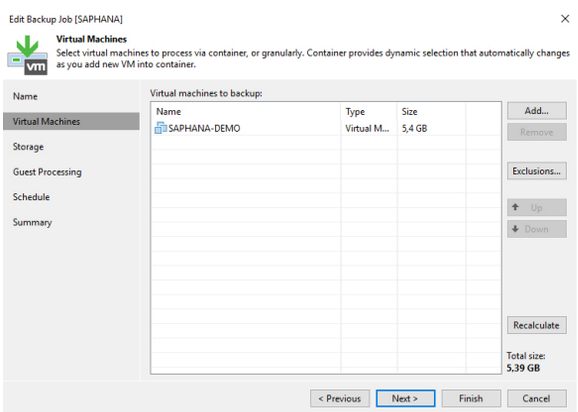


Step 3. Select SAP HANA Virtual Machine to Back Up

At the **Virtual Machines** step of the wizard, select the SAP HANA VM that you want to back up.

1. Click **Add**

2. Select the SAP HANA VM and click **Add**



Step 4. Specify Backup Storage Settings

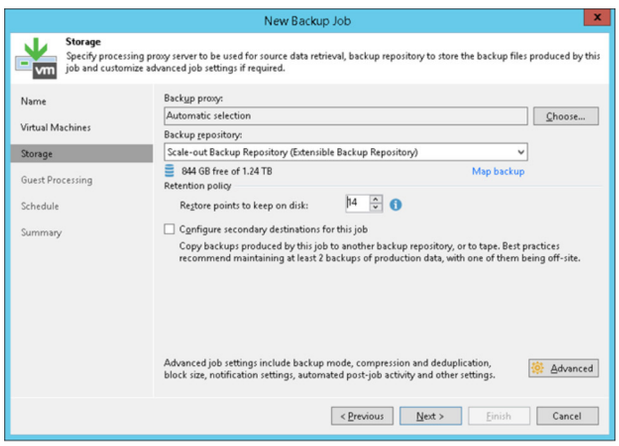
1. Click C**hoose** next to the Backup proxy field to select a backup proxy.

- If you choose **Automatic** selection, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and automatically assign an optimal backup proxy to process VMs in the job.
- If you choose **Use the selected backup proxy servers** specified below, you can explicitly select backup proxies that the job must use. It is recommended that you select at least two backup proxies to ensure that the backup job starts if one of the proxies fails or loses its connectivity to the source datastore.

2. From the **Backup repository** list, select a backup repository where the created backup files must be stored. When you select a backup repository, Veeam Backup & Replication automatically checks how much free space is available on the backup repository.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

3. In the **Retention policy** section, specify the number of restore points that you want to store on the backup repository. When this number is exceeded, the earliest restore point will be removed from the backup chain. The number of restore points doesn't correspond to the number of days to store restore points.



Note: Before you configure application aware processing on a SUSE or RedHat Linux system please make sure you have read the following Veeam knowledge base article.

https://www.veeam.com/kb1512

We also recommend using certificates instead of user/password combinations to logon to a remote linux system.

Note: We recommend creating a Veeam backup user for running the SAP HANA pre- and post-freeze scripts. This user does not need root/sudo rights but needs to be part of the sapsys group.

Step 5: Application-Aware Processing

1. To specify pre-freeze and post-thaw scripts for the job:

2. At the **Guest Processing,** Enable Application-Aware Processing, click **Applications** (Figure 1).

3. In the displayed list, select the VM and click **Edit**.

4. Click the **Scripts** tab.

5. In the **Script processing mode** section, specify the scenario for scripts execution:

- Select **Require successful script execution** if you want Veeam Backup & Replication to stop the backup process if the script fails. (Figure 2).

6. In the **Linux scripts** section, specify paths for SAP HANA pre- and post-freeze script (Figure 3).

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

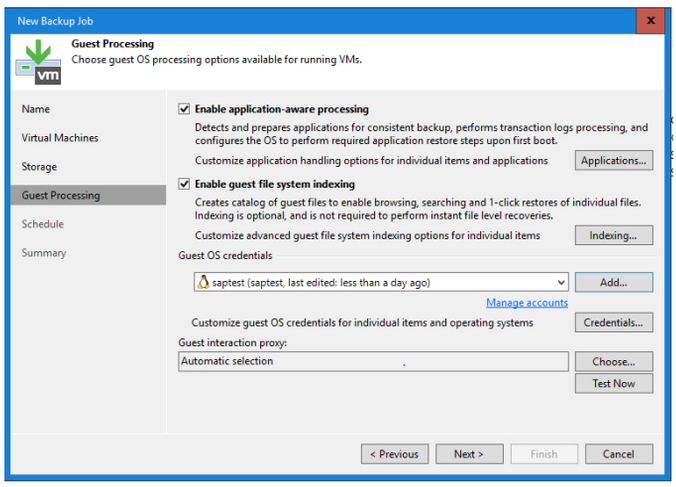


Figure 1. Enable Application-Aware Processing

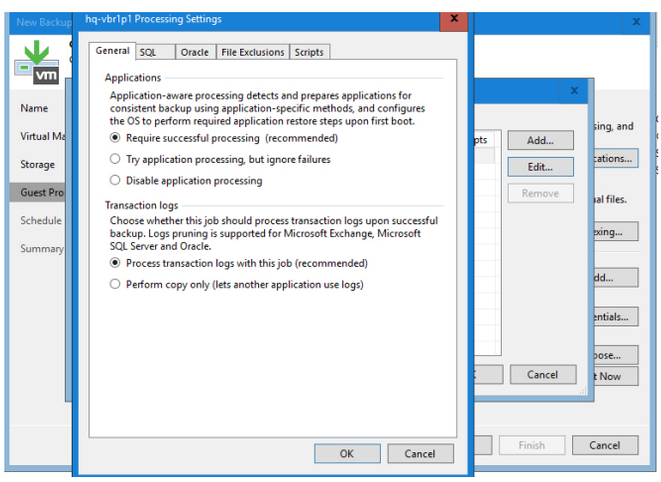


Figure 2. Select Script Option

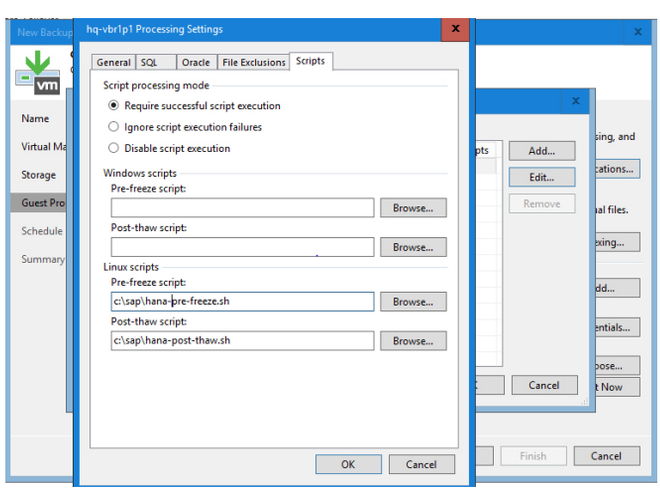


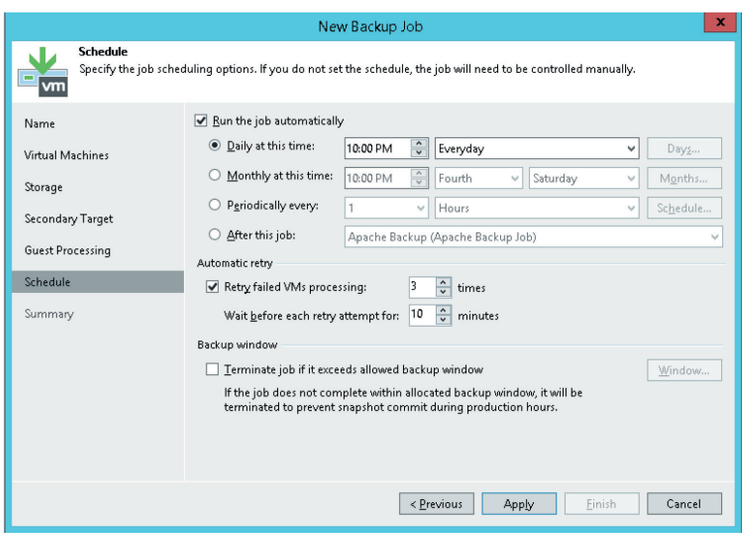
Figure 3. Select SAP HANA Pre- & Post-thaw script

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 6. Define Job Schedule

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

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to create the VM backup.
2. Define scheduling settings for the job.
3. In the **Automatic retry** section, define whether Veeam Backup & Replication must attempt to run the backup job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed VMs only. Enter the number of attempts to run the job and define time intervals between them. If you select continuous backup, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job runs.
4. In the **Backup window** section, define the time interval within which the backup job must complete. The backup window prevents the job from overlapping with production hours and ensures that the job does not provide unwanted overhead on the production environment. To set up a backup window for the job:
   a. Select the Terminate job if it exceeds allowed backup window check box and click Window.
   b. In the Time Periods section, define the allowed hours and prohibited hours for backup. If the job exceeds the allowed window, it will be automatically terminated.

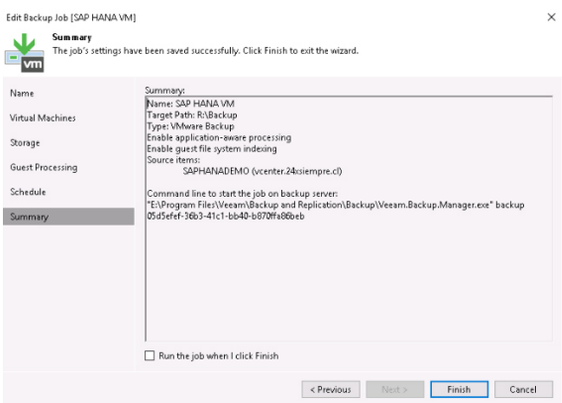


© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 7. Finish Working with the Wizard

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

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



Step 8. Run the Backup Job

To start a job:

1. Open the **Backup & Replication** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Start** on the ribbon or right-click the job and select **Start**.

Step 9. Validation

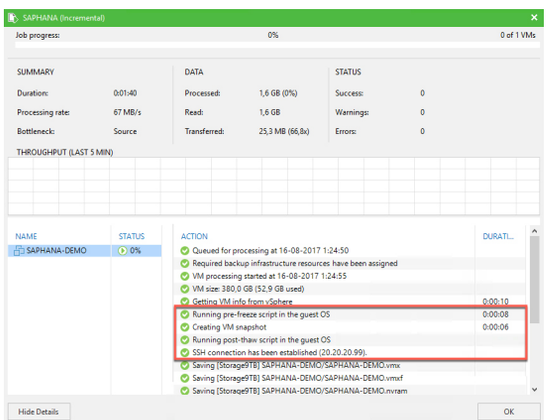
In this step, we will validate the successful HANA storage snapshot creation with Veeam pre-script and post-thaw script from the following:

1. Veeam Backup Job
2. SAP HANA Studio

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

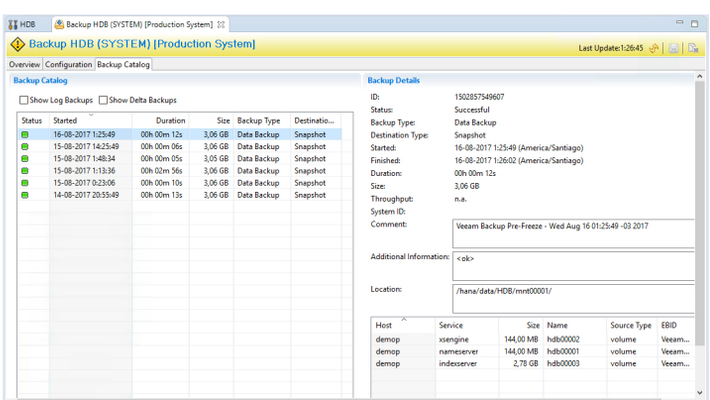
Veeam Backup Job

As marked in the picture below, you can validate the successful execution of pre- and post-thaw script from Veeam Backup & Replication Backup Job:



SAP HANA Studio

As we are using pre- and post-thaw script to create the HANA storage snapshot during the backup operations, we can validate the creation of HANA snapshots as shown in the picture below from SAP HANA Studio:



We have successfully completed the configuration of SAP HANA Virtual Machine backup with HANA Storage Snapshot options, now we will configure the backup of a physical SAP HANA server.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Configuration of SAP HANA Physical Machine Backup

In this section we will configure an SAP HANA physical machine backup. There are two ways of using the Veeam Agent for Linux:

- Managed by Veeam Agent for Linux
- Managed by VBR

Managed by Veeam Agent *for Linux*

To back up the SAP HANA physical server, Veeam Agent *for Linux* need to be installed on the SAP HANA physical server or

Step 1: Install the Veeam Agent for Linux

Download the Veeam software repository installation package (veeam-release) from the Veeam download page at https://www.veeam.com/downloads.html, and save the downloaded package on the computer where you plan to install the product.

1. Navigate to the directory where you have saved the veeam-release package and install the package with the following commands:

Operating Systems	Commands
SLES for SAP	zypper in ./veeam-release* && zypper refresh
RHEL for SAP HAN	rpm -ivh ./veeam-release* && yum check-update

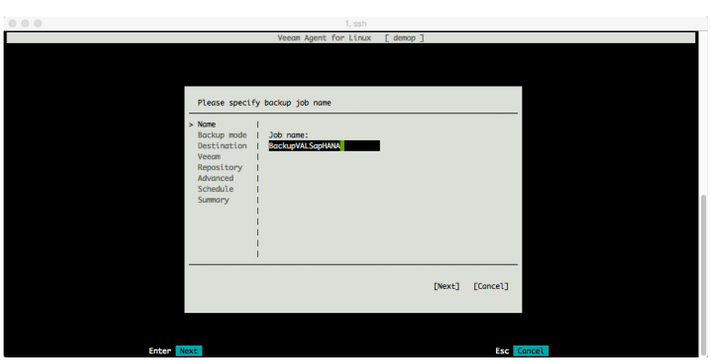
Step 2. Launch Backup Job Wizard

To launch the Backup Job wizard, do the following:

1. Launch the Veeam Agent *for Linux* control panel with the *veeam* or *veeamconfig* ui command.
2. Veeam Agent *for Linux* will display a welcome screen. Press the 'c' key to proceed to the Backup Job wizard and configure the backup job.

Step 3. Specify Backup Name

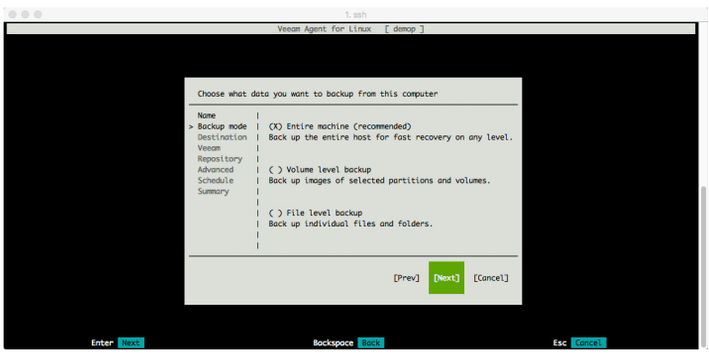
At the **Name** step of the wizard, in the **Job Name** field, type the name for the backup job and press **Enter.**



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 4. Select Backup Mode

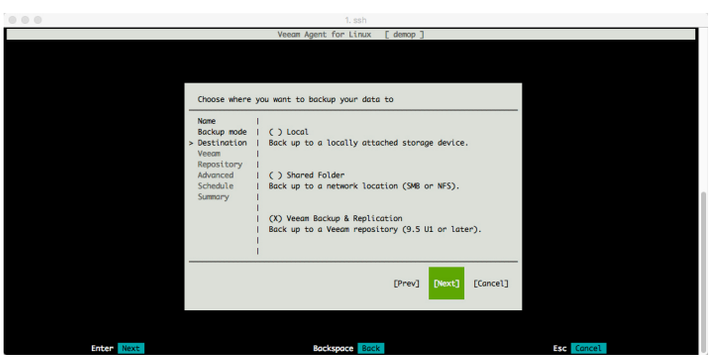
Select **Entire machine** option.



Step 5. Select Backup Destination

You can select one of the following options:

Type	Description
Local	Select this option if you want to save the backup on a removable storage device attached to the computer or on a local computer drive. With this option selected, you will pass to the Location step of the wizard.
Shared Folder	Select this option if you want to save the backup in a network shared folder. With this option selected, you will pass to the Network step of the wizard.
Veeam Backup & Replication	Select this option if you want to save the backup on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the Veeam step of the wizard.
Veeam Cloud Connect repository	Select this option if you want to create the backup on a cloud repository exposed to you by the Veeam Cloud Connect service provider. With this option selected, you will pass to the Service Provider step of the wizard.



In the guide, we are going to use **Veeam Backup & Replication server** as a **repository** to store the backup.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

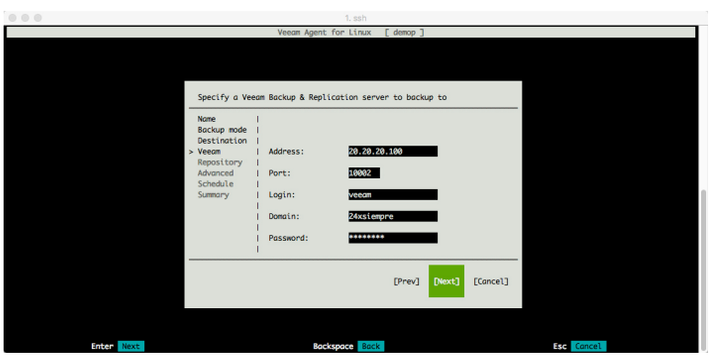
Step 6. Veeam Backup Repository Settings

In this step, we will supply the credentials and select the backup repository from Veeam Backup & Replication Server.

Specifying Backup Server Settings

Specify settings for the Veeam backup server that manages the target backup repository:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent *for Linux* must communicate with the backup repository. By default, Veeam Agent *for Linux* uses port 10006.
3. In the **Login** field, type in the username of the account that has access to the Veeam backup repository.
4. In the **Domain** field, type in the name of the domain in which the account that has access to the Veeam backup repository is registered, for example: DOMAIN.
5. In the **Password** field, type in the password of the account that has access to the Veeam backup repository.

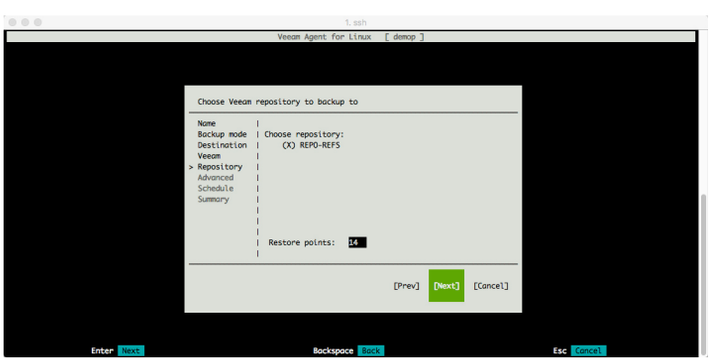


Selecting Backup Repository

Specify settings for the target backup repository

From the Choose repository list, select a backup repository where you want to store backups. The list displays only those backup repositories on which you have permissions to store data. To learn more, see Setting Up User Permissions on Backup Repositories.

In the Restore points field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent for Linux keeps the 14 latest restore points. After this number is exceeded, Veeam Agent for Linux will remove the earliest restore points from the backup chain.



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 7. Script Settings

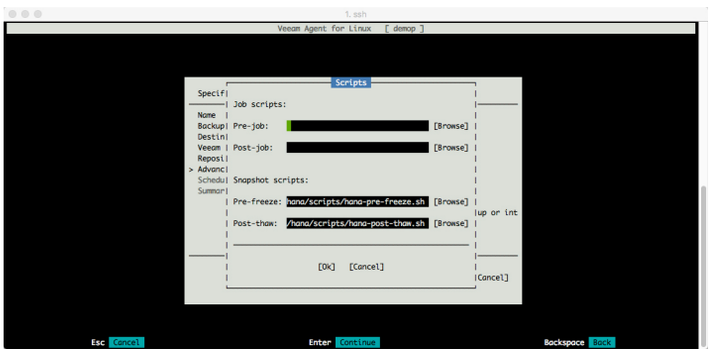
At the **Advanced** step of the wizard, specify script settings for the backup job:

1. In the **Scripts** window, in the **Snapshot scripts** section, specify custom scripts that you want to execute before Veeam Agent *for Linux* creates a snapshot of the backed-up volume and/or after the snapshot is created:

- In the **Pre-freeze** field, specify a path to HANA PRE-Freeze Script
- In the **Post-thaw** field, specify a path to HANA POST-Thaw Script

2. Switch to the **Ok** button and press **Enter**.

Note: Both scripts must be copied to an operating system path, for example, / hana / scripts /



Step 8. Specify Backup Schedule

At the **Schedule** step of the wizard, specify the schedule according to which you want to perform backup.

1. Make sure that the **Run the job automatically** check box is selected.

2. In the **Daily at** field, specify the time of day when the backup job must start.

3. In the list of days of the week, specify days on which the job must start. By default, Veeam Agent *for Linux* will start the backup job daily at the specified time. To exclude specific days from the schedule, clear the check boxes next to the necessary days

Step 9. Review Backup Job Settings

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

1. Review the settings of the configured backup job.

2. Select the start job now check box to start the job immediately.

3. Press **Enter** to exit the wizard.

Step 10. Validation

In this step, we will validate the successful HANA Storage Snapshot creation with Veeam pre- and post-thaw script from following:

1. Veeam Agent *for Linux* Backup Job.

2. SAP HANA Studio

3. Veeam Backup & Replication Console.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

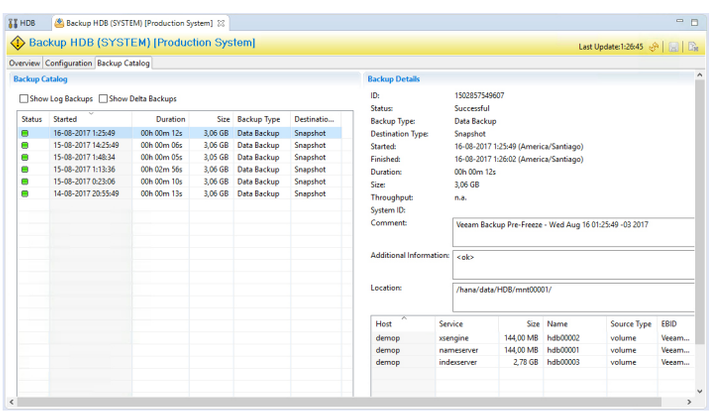
Veeam Agent *for Linux* Backup Job

You can check the progress of a backup job with Veeam Agent *for Linux* Console to validate if the backup job has been successfully completed as shown in the picture below:



SAP HANA Studio

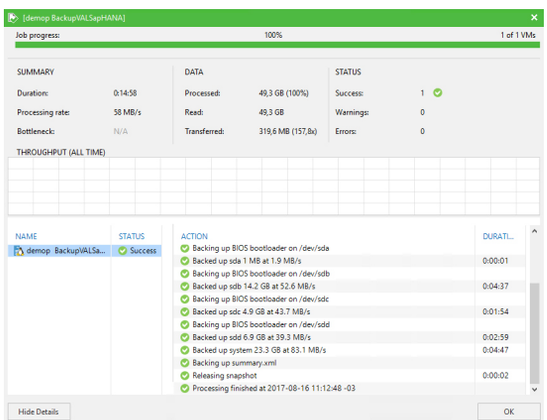
As we are using pre- and post-thaw script to create the HANA Storage Snapshot during the backup operations, we can validate the creation of a HANA Snapshot as shown in the picture below from SAP HANA Studio



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam Backup & Replication Console

As we are using Veeam Backup & Replication repository as a target to store the backup file, we can view the progress of the backup job with the console as shown in the picture below:



We have successfully created the backup jobs for SAP HANA virtual and physical environment and also validated the creation of a HANA Snapshot. In the next section we will cover the recovery procedure.

Managed by Veeam Backup & Replication

The steps are logically the same. Install the agent via VBR console, create a backup job and integrate the pre and post script.

Step 1. Agent installation

Ensure that the communication (ssh login) into the SAP HANA operating system is working well. Please also refer to

https://www.veeam.com/kb1512

https://www.veeam.com/kb2061

https://helpcenter.veeam.com/docs/backup/hyperv/credentials_manager_linux_pubkey.html?ver=95u4

We also recommend using certificates instead of user/password combinations to logon to a remote linux system.

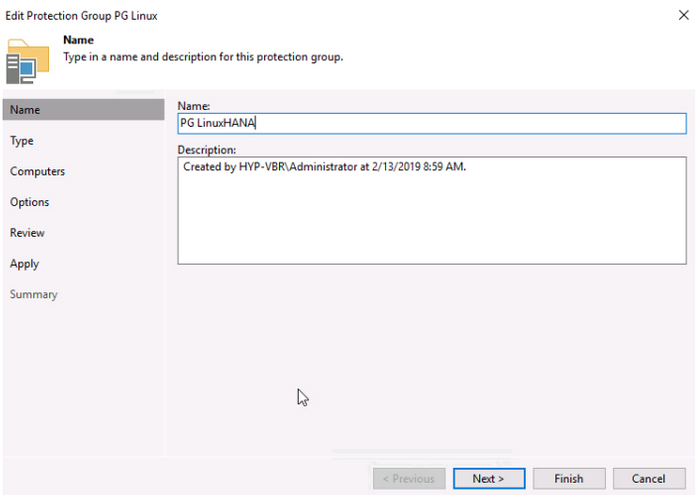
Note: We recommend creating a Veeam backup user for running the SAP HANA pre- and post-freeze scripts. This user does not need root/sudo rights but needs to be part of the sapsys group.

If the agent was installed properly it should have an entry for your new system in the infrastructure section with the agent installed.

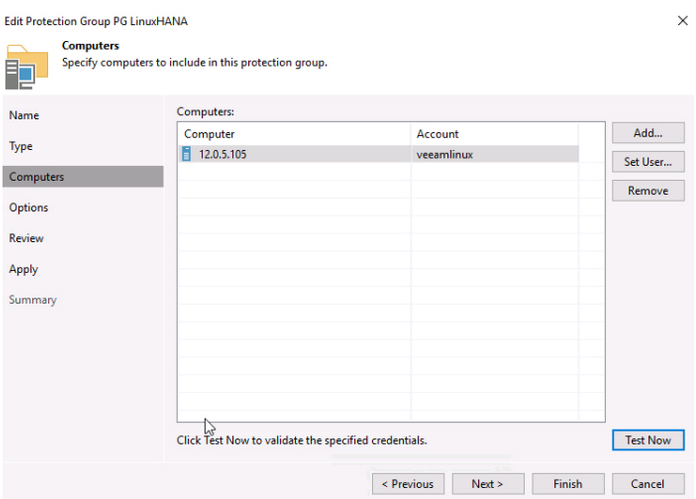


© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

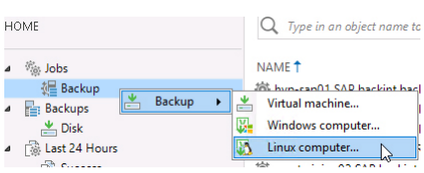
Step 2. Create a Protection group with all your SAP HANA systems



Step 3. Add your SAP HANA systems

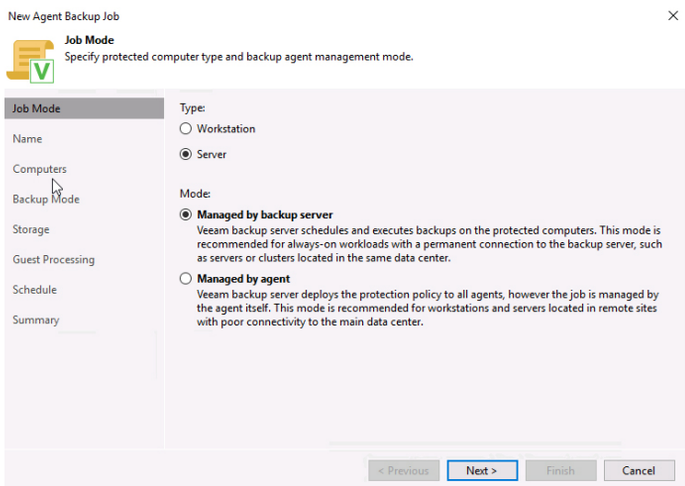


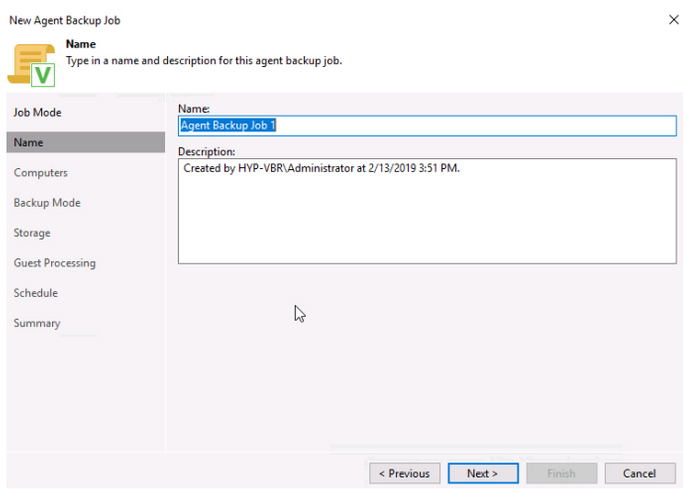
Step 4. Create a Linux computer Agent based backup job



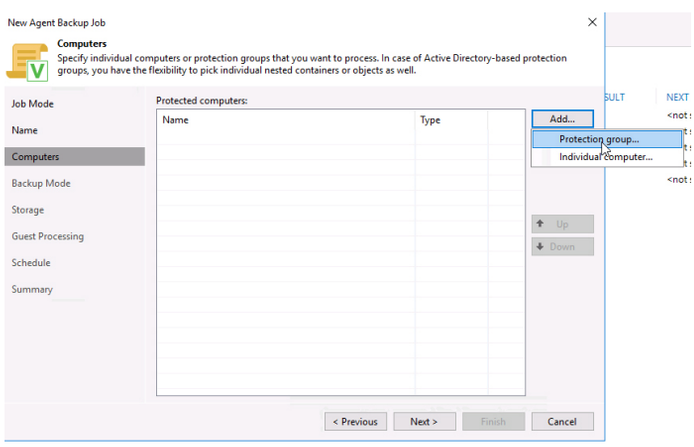
Select „Managed by backup server"

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

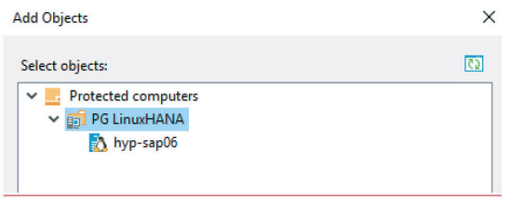


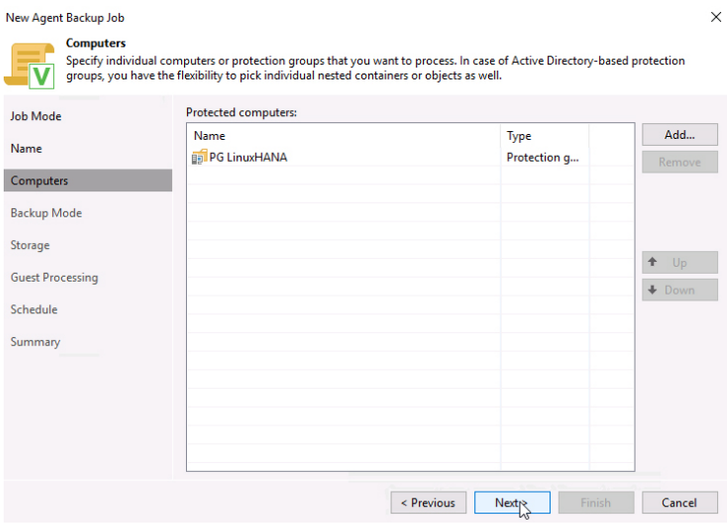


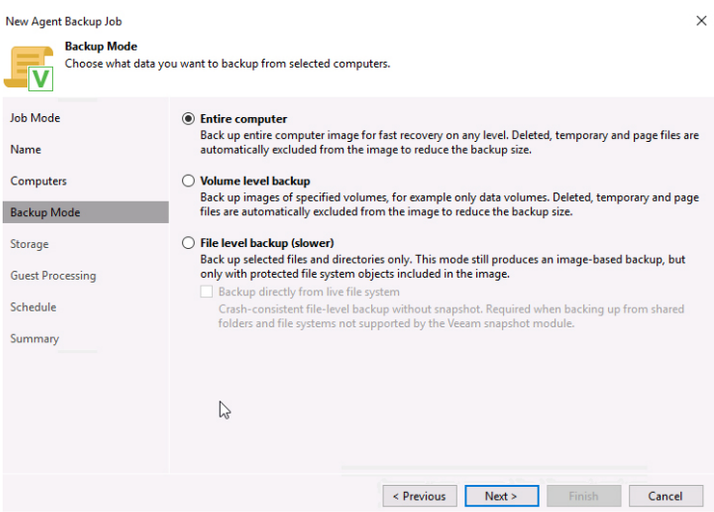
Give your job a name and add your previously created protection group



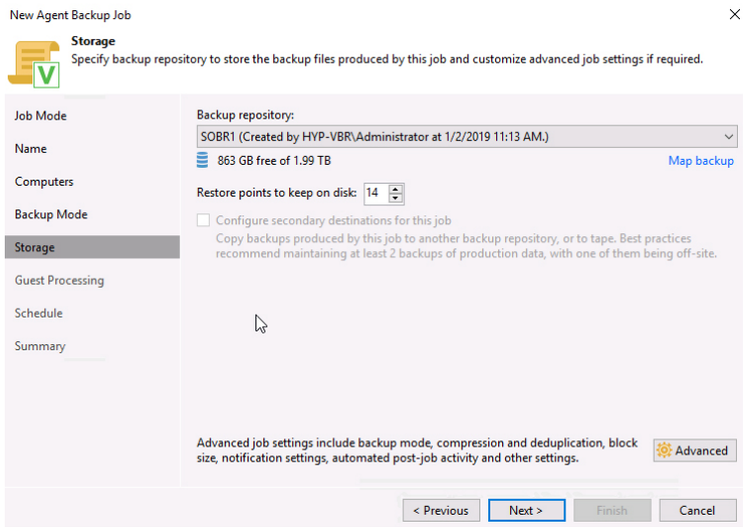
© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

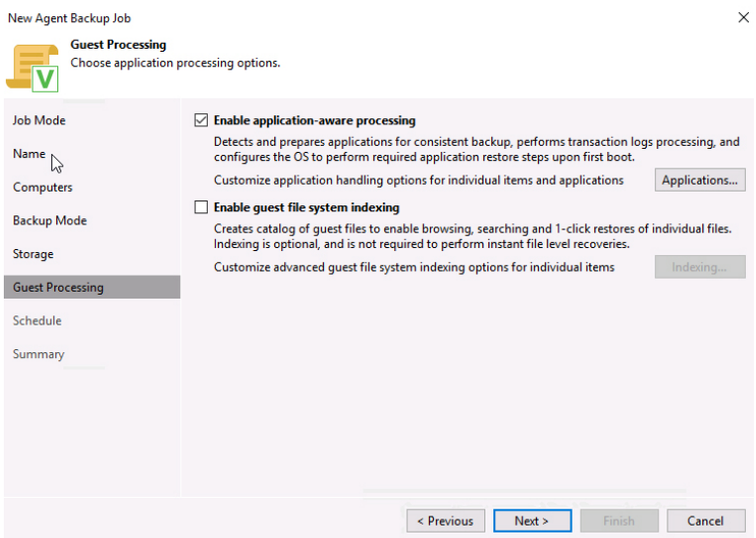




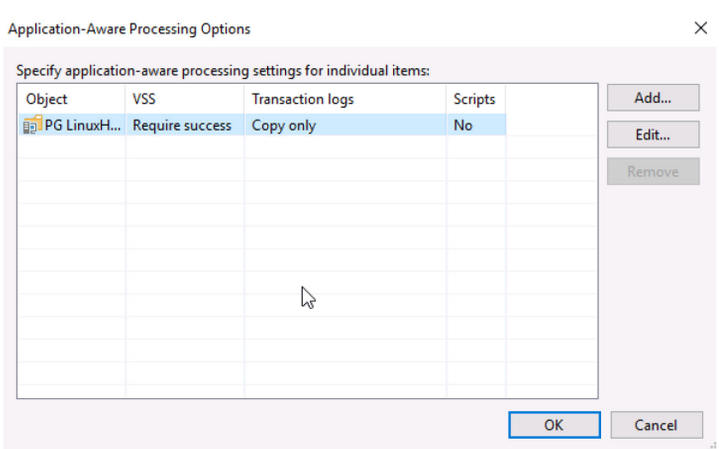


© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

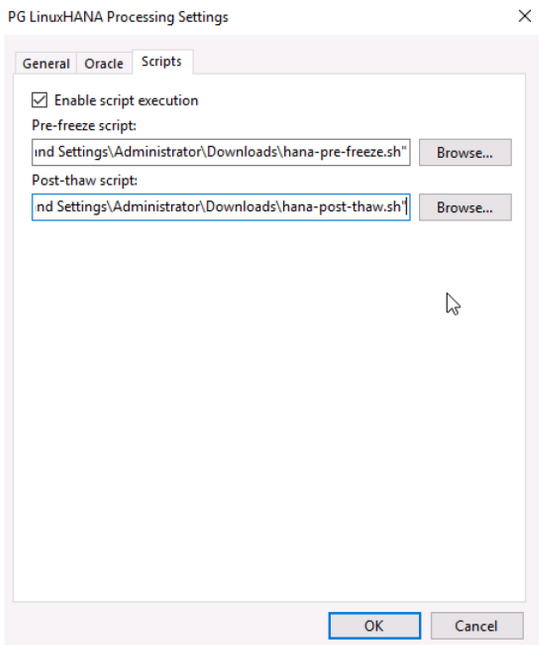


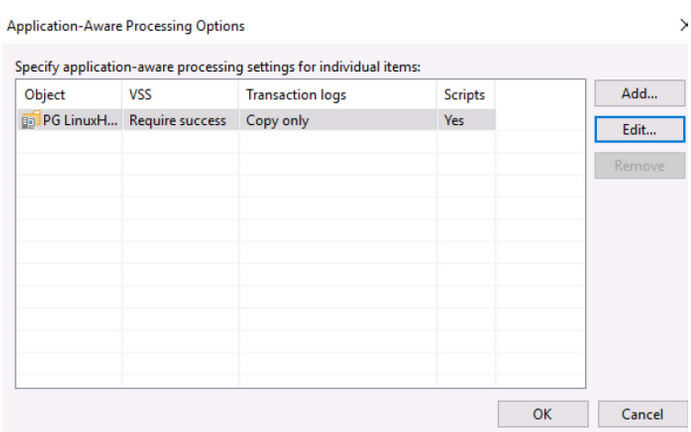


Enable application aware processing and edit/add the pre & post scripts.



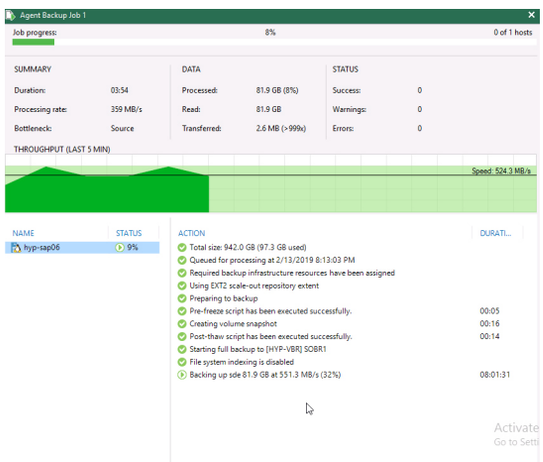
© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

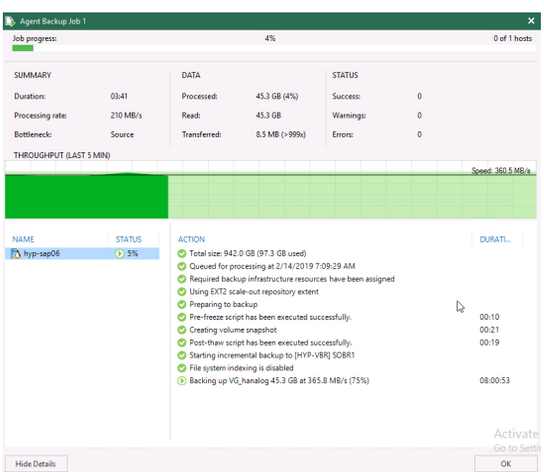




© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

If everything will run smoothly it should look like this:





© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Veeam SAP HANA Recovery

Veeam provides various options to restore data. In this guide we will perform the recoveries of SAP HANA using the following options:

Environment	Recovery Type
Virtual or Physical	SAP HANA Database Restore.
Virtual Machine	Veeam Virtual Disk Level Restore.
	Complete Virtual Machine Restore — Instant VM Recovery.
Physical Server	Volume-Level Restore / Bare Metal Restore.
	File-Level Restore.

For the SAP HANA database restore, we will use SAP HANA Studio.

SAP HANA Database Level Restore:

With the configured Veeam Plug-in for SAP HANA you can restore your databases from the backups that reside in the Veeam backup repository. All restore operations are performed in the SAP HANA side. To restore databases, you can use SAP HANA Cockpit, SAP HANA Studio, or SQL commands including scripts.

Keep in mind that examples provided below are for demonstrating purposes only. To see the full restore functionality of SAP HANA tools, see the SAP HANA Recovery section of SAP HANA Administration Guide.

For instructions on the SAP HANA recovery, see:

Recovering Databases Using SQL Commands

Recovering Databases with SAP HANA Studio

Recovering SYSTEMDB with SAP HANA Cockpit

Recovering Tenant Databases with SAP HANA Cockpit

Recovering Databases to Other Servers

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Virtual SAP HANA Restore — Veeam Virtual Disk-Level Restore

In the Veeam Virtual-Level Restore, we will use the Quick Rollback option to restore the change blocks only in the SAP HANA VM.

The diagram below shows the restore operations:

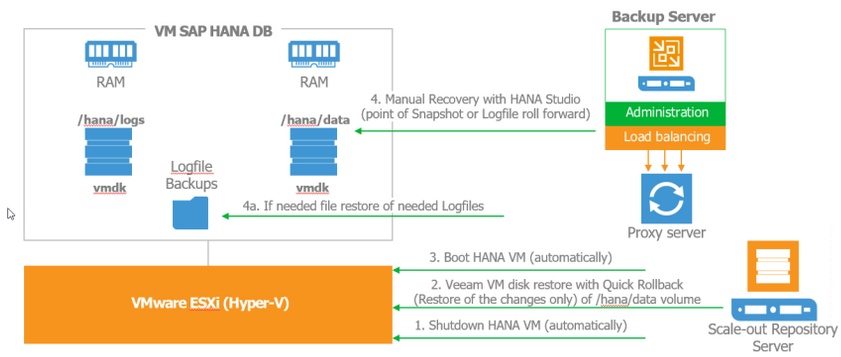
1. Shutdown HANA VM (automatically)

2. Veeam VM disk restore with Quick Rollback (Restore of the changes only) of /hana/data_volume.

3. Boot HANA VM (automatically)

4. Perform the restore from the SAP HANA studio. In step 4a, if required, we can perform file-level restore for the needed logs.

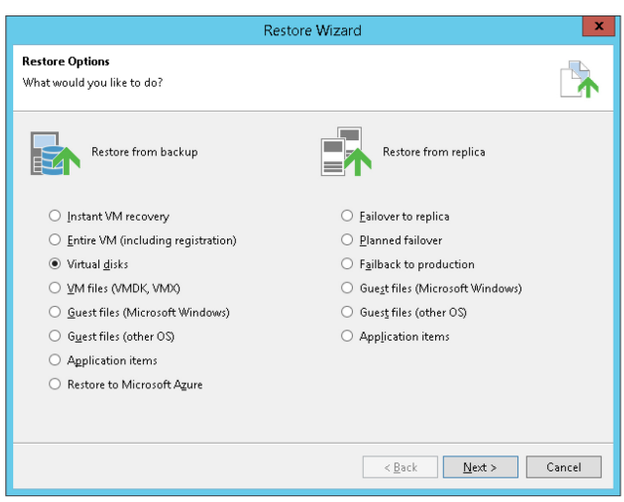
Point of snapshot/backup restore + optional log file roll forward with Veeam VM disk restore



Step 1. Launch Virtual Disk Restore Wizard

To launch the **Virtual Disk Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware vSphere backup**. In the **Restore from backup** section, select **Virtual disks**.
- Open the **Home** view, in the inventory pane select **Backups**. In the working area, expand the necessary backup, select the VM whose disks you want to restore and click **VM Files** > **Virtual Disk** on the ribbon.
- Open the **Home** view, in the inventory pane select **Backups**. In the working area, expand the necessary backup job, right-click the VM whose disks you want to restore and select **Restore virtual disk**



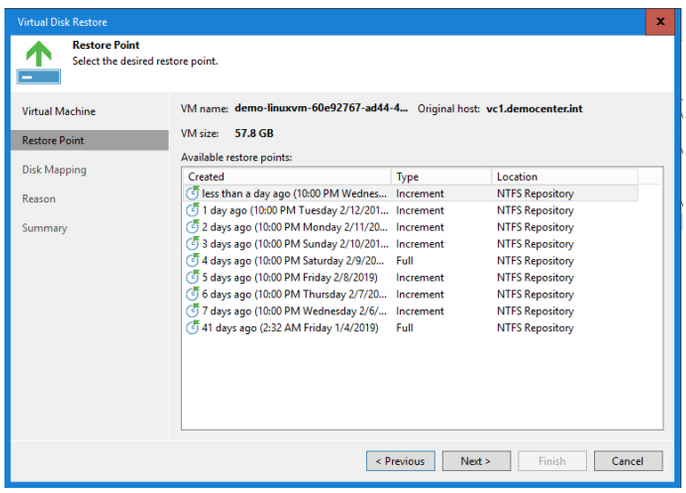
© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 2. Select VM

At the **Virtual Machine** step of the wizard, select the **SAP HANA VM**

Step 3. Select Restore Point

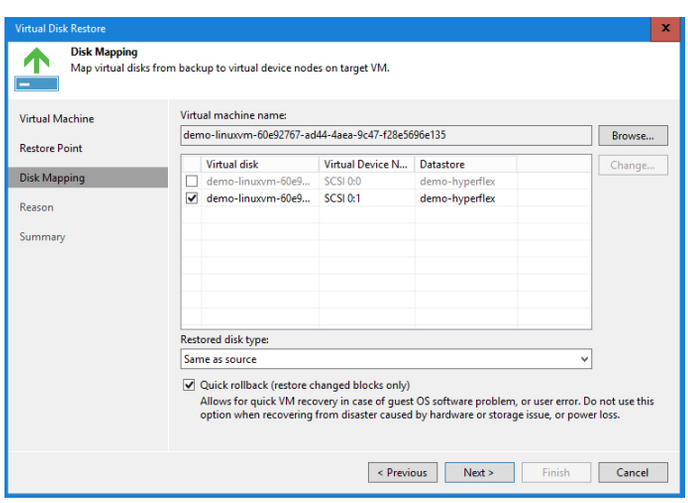
At the **Restore Point** step of the wizard, select the restore point from which you want to restore the SAP HANA VM Disk.



Step 4. Select Virtual Hard Disks to Restore

At the **Disk Mapping** step, select which virtual hard disks to restore, choose a VM to which the disks must be attached and define additional restore settings.

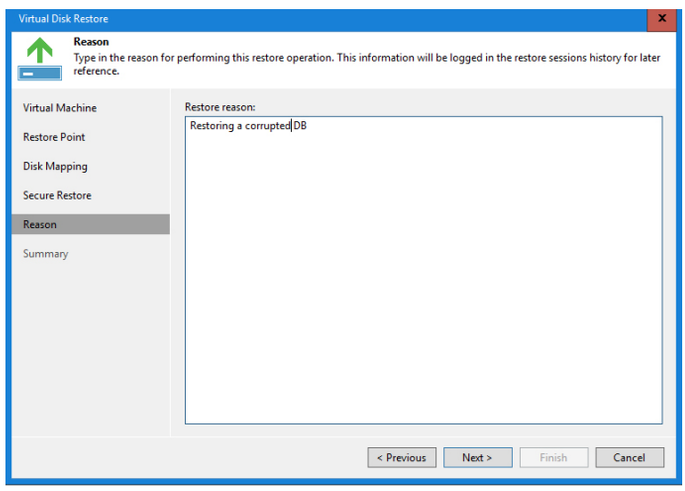
1. In the Restored Disk Type: **Select same as source.**
2. Select **Quick rollback** (restore changed blocks only) checkbox. -*Veeam Backup & Replication will query CBT to get data blocks that are necessary to revert the VM disk to an earlier point in time, and will restore only these data blocks from the backup. Quick rollback significantly reduces the restore time and has little impact on the production environment*.
3. Select **Next.**



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 5. Specify Restore Reason

At the Reason step of the wizard, enter a reason for restoring VM disk(s). The information you provide will be saved in the session history so you can reference it later.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, complete the VM disk(s) restore.

1. Review details for the restore task.
2. To start a VM immediately after the restore process, select the **Power on VM after restoring** check box.
3. Click the **Pick proxy to use** link to select backup proxies over which VM data must be transported to the target datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.

- If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and will automatically assign optimal proxy resources for processing VM data.
  During the restore process, VM hard disks are processed simultaneously. Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VM hard disk processing.
- If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that must be used for restoring. It is recommended that you select at least two backup proxies to ensure that VM hard disks are recovered if one of backup proxies fails or loses its connectivity to the target datastore during restore.

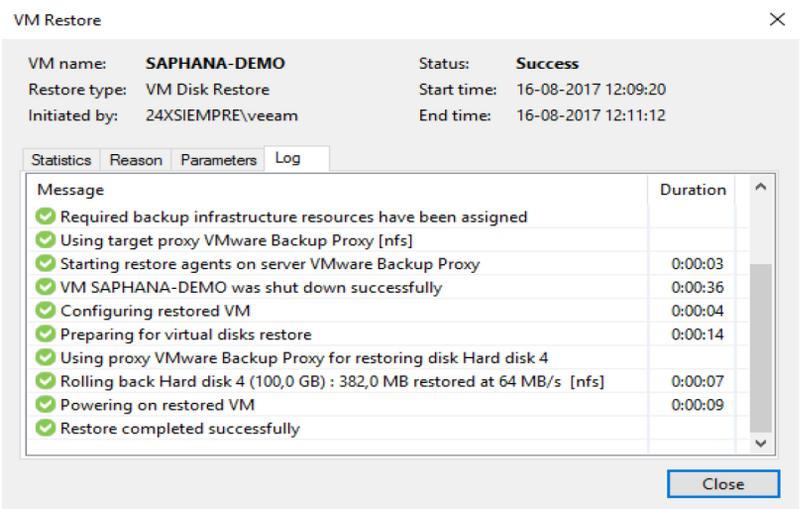
4. Click **Finish** to start VM disk(s) restore.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 7. Restore Progress

You can view the restore progress from the Veeam Backup & Replication console.

1. Open the **History** view. In the inventory pane, select **Restore**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.



Virtual SAP HANA Restore — Instant Virtual Machine Restore

With Instant VM Recovery®, you can immediately restore a VM into your production environment by running it directly from the backup file. Instant VM Recovery helps improve recovery time objectives (RTO), and minimize disruption and downtime of production VMs. It is like having a "temporary spare" for a VM: Users remain productive while you can troubleshoot an issue with the failed VM.

When Instant VM Recovery is performed, Veeam Backup & Replication uses the Veeam vPower® technology to mount a VM image to an ESX(i) host directly from a compressed and deduplicated backup file. Since there is no need to extract the VM from the backup file and copy it to production storage, you can restart a VM from any restore point (incremental or full) in a matter of minutes.

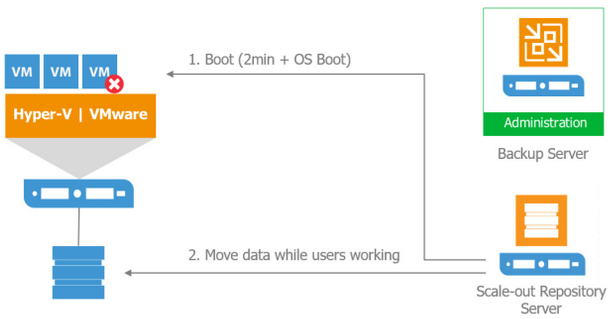
The archived image of the VM remains in read-only state to avoid unexpected modifications. By default, all changes to virtual disks that take place while the VM is running, are logged to auxiliary redo logs residing on the NFS server (backup server or backup repository). These changes are discarded as soon as a restored VM is removed, or merged with the original VM data when VM recovery is finalized.

In case of complete SAP HANA VM failure, Instant VM Recovery help reduces the RTO for SAP HANA systems restore.

The following diagram shows the Instant VM Recovery operations:

1. Veeam Backup & Replication uses the Veeam vPower technology to mount a VM image to an ESX(i) host directly from a compressed and deduplicated backup file. Since there is no need to extract the VM from the backup file and copy it to production storage, you can restart a VM from any restore point (incremental or full) in a matter of minutes.
2. Online data migration to production storage.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.



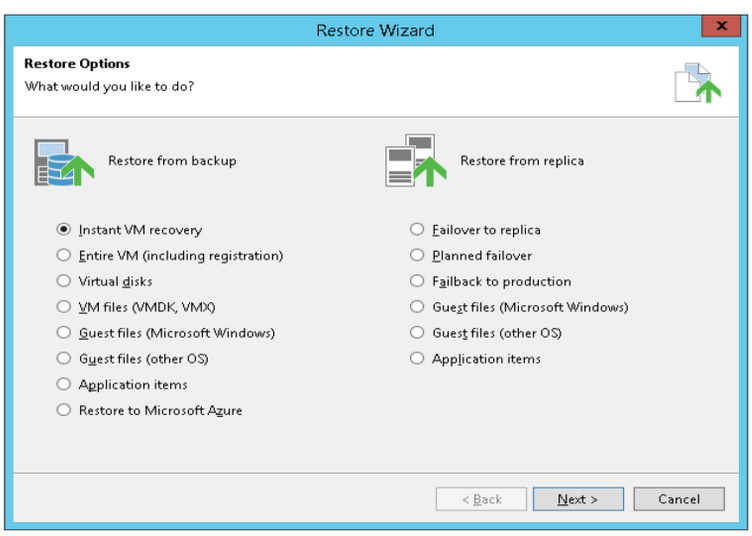
Performing Instant VM Recovery of SAP HANA VM

Step 1. Launch Instant VM Recovery Wizard

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

- On the **Home** tab, click **Restore** and select **VMware vSphere backup**. In the **Restore from backup** section, select **Instant VM Recovery**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, select the VM you want to restore and click **Instant VM Recovery** on the ribbon.

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



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 2. Select SAP HANA VMs

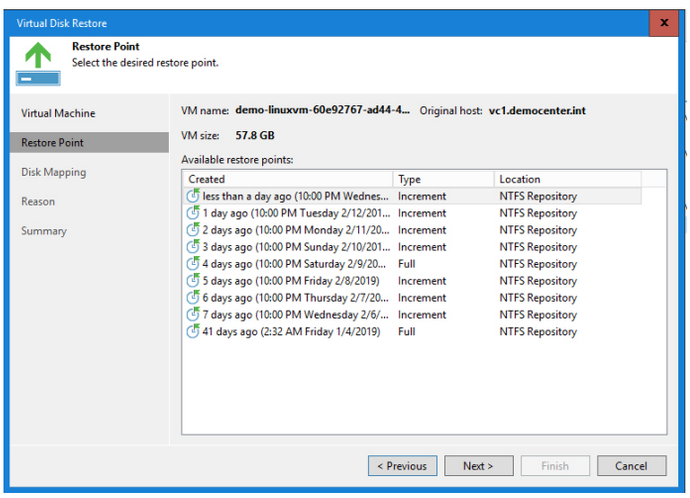
At the **Virtual Machines** step of the wizard, select the SAP HANA VM that you want to back up.

1. Click **Add**

2. Select the SAP HANA VM and click **Add**.

Step 3. Select Restore Point

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



Step 4. Select Recovery Mode

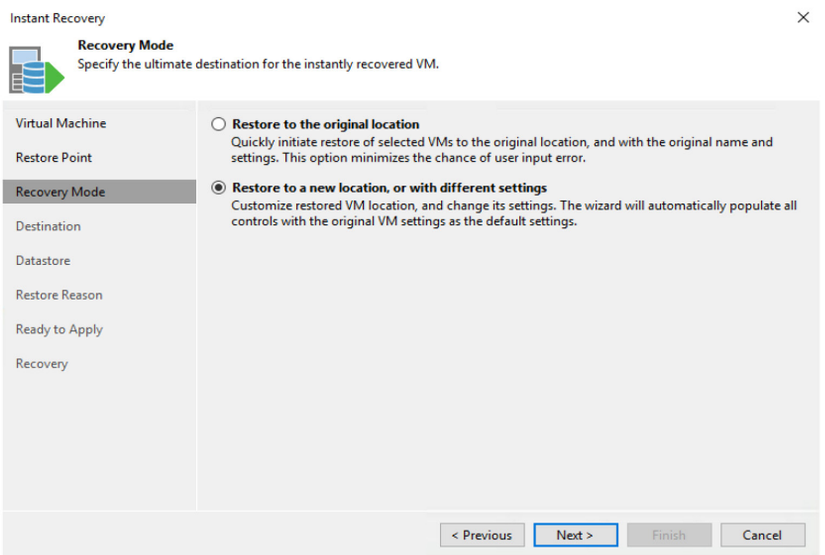
At the **Recovery Mode** of the wizard, choose the necessary restore mode:

Select **Restore to the original location** if you want to restore the VM with its initial settings and to its original location. If this option is selected, you will pass directly to the Reason step of the wizard.

Select **Restore to a new location, or with different settings** if you want to restore the VM to a different location and/or with different settings (such as VM location, network settings, format of restored virtual disks and so on). If this option is selected, the Instant Recovery wizard will include additional steps for customizing VM settings.

In the guide, we are selecting the option to restore to a new location.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

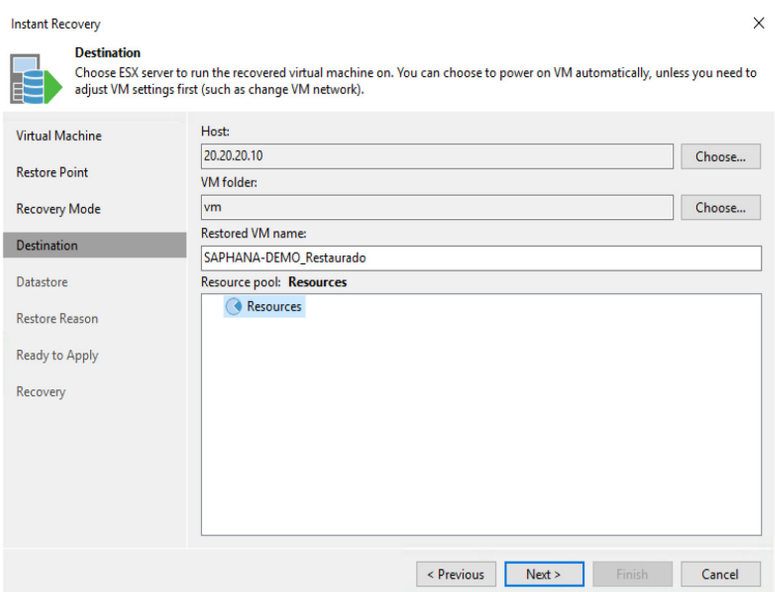


Step 5. Select Destination for Restored VM

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

Select a destination for the restored VM:

1. In the **Host** field, specify a host on which the VM must run.
2. In the **VM folder** field, specify a folder to which the restored VM must be placed.
3. In the **Restored VM name** field, enter a name under which the VM must be restored and registered. By default, the restored VM has the name of the original VM. If you are restoring the VM to the same ESX(i) host or the same data center where the original VM is registered and the original VM still resides there, it is recommended that you change the VM name to avoid conflicts.
4. In the **Resource pool** list, select a resource pool to which the VM must be placed.



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 6. Select Destination for Virtual Disk Updates

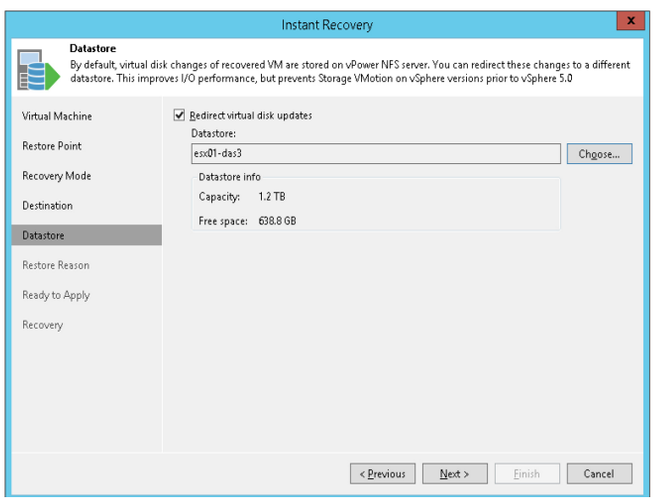
It is highly recommended to use this step for SAP HANA Virtual machines to improve performance and redirect the redo logs to faster storage.

At the **Datastore** step of the wizard, you can select where redo logs must be written when the VM is running from the backup. By default, redo logs are stored directly on the backup server. However, you can store redo logs on any datastore in the virtual environment. Redirecting redo logs improves recovery performance but makes Storage vMotion not possible for ESX(i) 5.x and earlier.

To redirect redo logs:

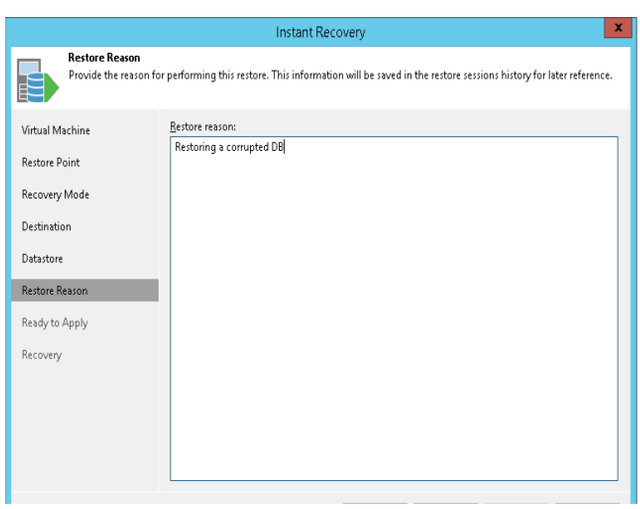
1. Select the **Redirect virtual disk** updates check box.

2. Choose the datastore.



Step 7. Specify Restore Reason

At the **Restore Reason** step of the wizard, enter a reason for performing Instant VM Recovery for the VM. The information you provide will be saved in the session history and you can reference it later.

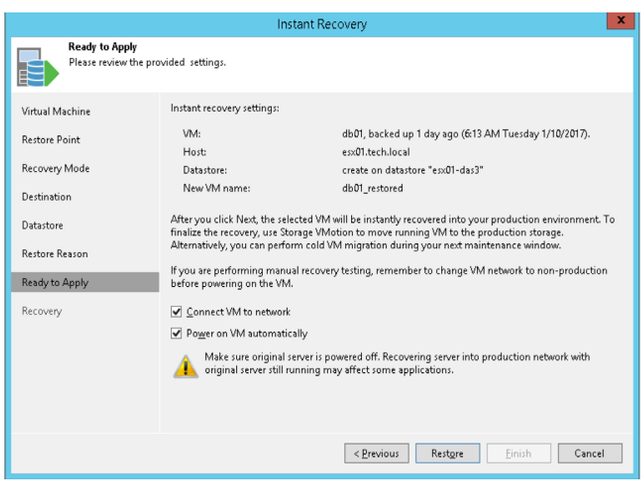


© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 8. Verify Instant VM Recovery Settings

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

1. If you are recovering a production VM that has failed and you want to restore it with initial network settings, select the **Connect VM to network** check box. If you are recovering a VM for testing disaster recovery while the initial VM is still running, leave this check box not selected. Before you power on such a VM, you will have to manually change VM network settings: Disconnect the VM from the production network and connect it to an isolated non-production network to avoid conflicts.
2. To start a VM immediately after recovery, select the **Power on VM automatically** check box. If you are recovering the VM to the production network, make sure that the initial VM is powered off to avoid conflicts.
3. Check the settings you have specified for Instant VM Recovery and click **Next**. Veeam Backup & Replication will recover the VM on the selected ESX(i) host.

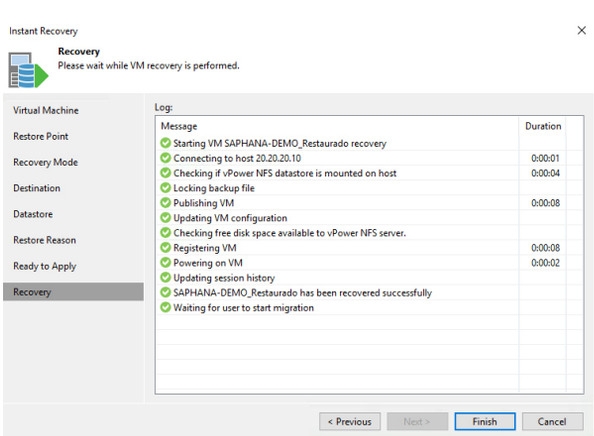


Step 9. Finalize Instant VM Recovery

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

To check the progress of Instant VM Recovery and view session details as shown in the picture below:

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



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 10. Migrating Recovered VM

To migrate a recovered VM to the production environment:

1. Open the H**ome** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM and select **Migrate to production** (Figure 1).
4. Veeam Backup & Replication will launch the Quick Migration wizard (Figure 2). During migration, Veeam Backup & Replication will restore the VM from the backup file and additionally move all changes that were made while the VM was running from the backup in the Instant Recovery mode.

At the Destination step of wizard, select the following (Figure 3).

- Click **Choose** to the Host or Cluster field and select an ESX(i) host or cluster if you like to relocate the restored VM.
- Click **Choose** next to the **Resource pool** field and select the target resource pool
- Click **Choose** and select the folder.
- Click **Choose** and select the datastore. Veeam Backup & Replication displays only those datastores that are accessible by the selected ESX(i) host. If you have chosen to restore the VM to a cluster, Veeam Backup & Replication will display only shared datastores.

5. Select Infrastructure Components for Data Transfer (Figure 4).

- Click **Choose** next to the **Source proxy** and **Target proxy** fields to select backup proxies for migration. In the **Backup Proxy** window, you can choose automatic proxy selection or assign proxies explicitly.

6. Select which migration mechanism to use (Figure 4).

- VMware vMotion — If you want to use VMware vMotion to relocate the VMs, leave the **Force Veeam quick migration** check box not selected. Veeam Backup & Replication will attempt to use the VMware vMotion mechanism to migrate the selected VMs. If VMware vMotion cannot be used for some reason (for example, if using it can cause data loss or if you do not have a VMware vSphere license for this functionality), Veeam Backup & Replication will fail over to its native migration mechanism.
- Veeam Quick Migration — If you do not want to use VMware vMotion, select the **Force Veeam quick migration** check box. Veeam Backup & Replication will use its native migration mechanism.

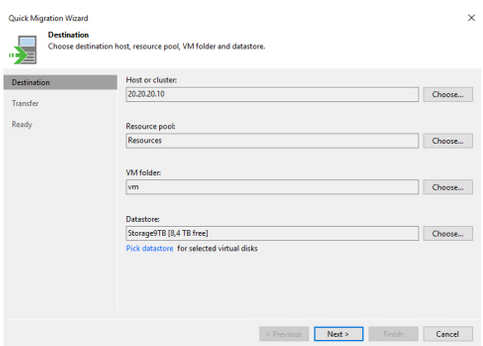


Figure 3. Destination

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

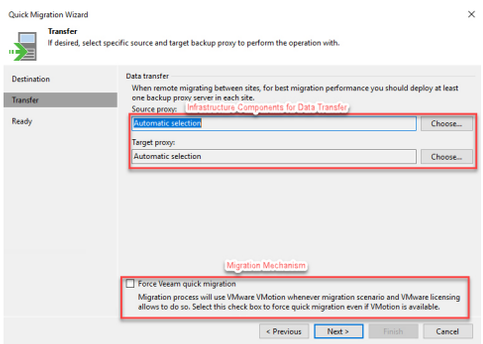


Figure 4. Infrastructure Components & Data Migration

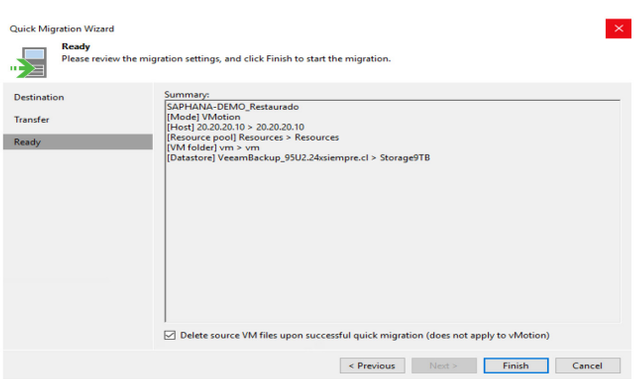
Step 11. Finish Working with Wizard

At the **Ready** step of the wizard.

1. Review details of the quick migration task.
2. By default, when VM migration completes successfully, Veeam Backup & Replication waits for a heartbeat signal from the VM on the target host. If the heartbeat is received, the original VM on the source host is deleted. Note that you cannot use this option if you selected to migrate with VMware vMotion.

- If you disable the **Delete source VM files upon successful migration** option, the source VM will not be deleted. All jobs to which the VM is added will continue to process the source VM.
- If you enable the **Delete source VM files upon successful migration** option, the source VM will be deleted. All jobs to which the VM is added will switch to the migrated VM. The backup chain will be continued, thus, the next job session for the VM will be incremental.

3. Click **Finish** to close the wizard and start the migration process.



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Physical SAP HANA Restore

In this section, we will perform the recovery of an SAP HANA system running on a physical server.

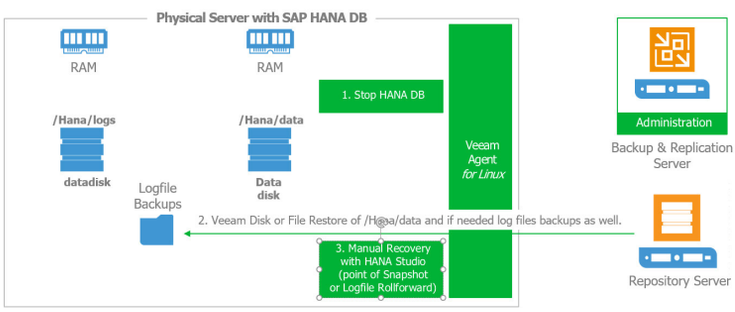
Veeam Agent for Linux provides following restore scenarios:

Failure	Restore Type
Complete Server Failure	Bare Metal Restore
Database Level Failure — Point in Time Transaction Restore	Volume-Level Restore

The below diagram shows the restore operations for SAP HANA physical server:

1. Stop HANA Database (In case of point in time restore) Skip the step for bare metal restore.
2. Recover disk volume, files or logs if necessary or complete server.
3. Point of Snapshot restore from SAP HANA studio.

HANA snapshot backup with Veeam Agent for Linux backups



Physical SAP HANA Recovery — Bare Metal & Volume Level

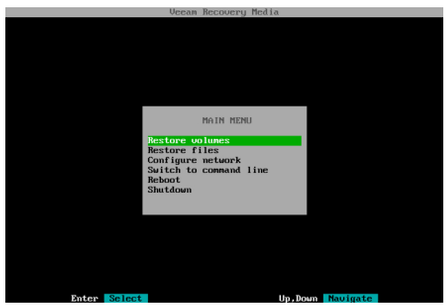
During the recovery of an SAP HAN system, we will use Veeam Recovery media *for Linux* to perform the volume level or bare metal.

Step 1. Boot from Veeam Recovery Media

To boot from the **Veeam Recovery Media**:

1. **[For CD/DVD/BD]** Power on your computer. Insert the media with the recovery image to the drive and power off the computer.
   [For removable storage device] Attach the removable storage device with the recovery image to your computer.
2. Start your computer.
3. Wait for Veeam Agent *for Linux* to load files from the **Veeam Recovery Media.**
4. After the recovery image OS has loaded, make sure that **network settings** are specified correctly and configure the network adapter if necessary.
5. Choose the **Restore Volume** recovery option.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

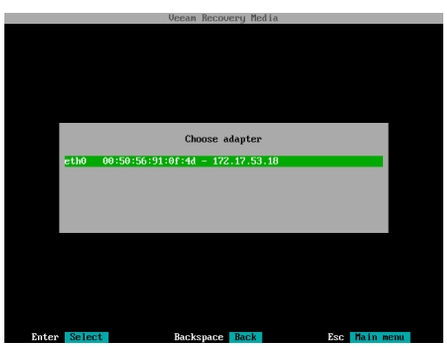


Step 2. Configure Network Settings

To open the **Network settings** dialog, in the Veeam Recovery Media main menu, select the **Configure network** option and press **Enter**.

Veeam Agent *for Linux* will display network adapters that are available on the system. If there is a DHCP server in your network, Veeam Agent *for Linux* will configure the network settings automatically and display the IP address assigned to the network adapter. You can then press **Esc** to return to the Veeam Recovery Media main menu and launch the Volume Restore wizard.

You can manually configure TPC/IP v4 settings for network adapters if necessary.



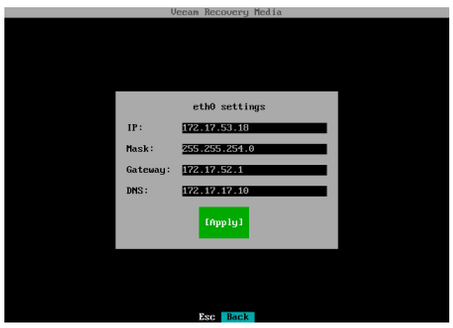
In this guide we will manually configure the TCP/IP v4 settings. To configure network settings:

1. In the **Choose adapter** list, select the network adapter that you want to use to connect to the network shared folder or Veeam backup repository where the backup resides and press **Enter**.
2. In the **Configure adapter** dialog, select the **Manual** option and press **Enter**.
3. In the **Adapter settings** dialog, specify the following network settings:

- IP address
- Subnet mask
- Default gateway
- DNS server

4. Select the **Apply** button and press **Enter**.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.



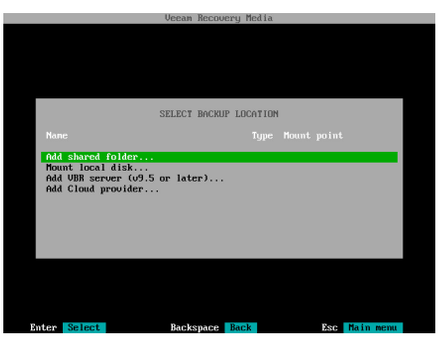
Step 3. Select Backup Location

At the **Select backup location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

To recover data from backup, you need to mount the backup storage on which the backup file resides to the recovery image OS file system. Veeam Agent *for Linux* automatically mounts external USB drives that are connected to the computer and displays them in the list of available backup locations. You can select the necessary device and press **Enter** to pass to the Browse for backup files step of the wizard.

If the backup file is located in a network shared folder, on a local drive or on a Veeam backup repository, select one of the following options:

- **Add shared folder** — select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the Mount shared folder step of the wizard.
- **Mount local disk** — select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. With this option selected, you will pass to the Select local disk step of the wizard.
- **Add VBR server** — select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the Specify backup server parameters step of the wizard.
- **Add Cloud provider** — select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the Specify Cloud provider parameters step of the wizard.

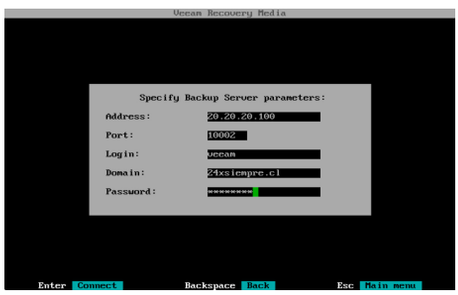


In this guide we will use Veeam Backup & Replication server for the restore operations.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Specify settings for the Veeam backup server that manages the backup repository where the backup file resides:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent *for Linux* must communicate with the backup repository. By default, Veeam Agent *for Linux* uses port 10002.
3. In the **Login** field, type in the username of the account that has access to the Veeam backup repository.
4. In the **Domain** field, type in the name of the domain in which the account that has access to the Veeam backup repository is registered, for example: DOMAIN.
5. In the **Password** field, type in the password of the account that has access to the Veeam backup repository.
6. Press **Enter**. Veeam Agent *for Linux* will connect to the Veeam backup server, and you will pass immediately to the Backup step of the wizard.



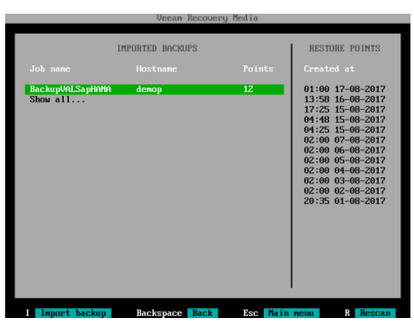
Step 4. Select Backup and Restore Point

The **Backup** step window comprises of two panes:

- The **Imported backups** pane on the left displays information about backup: Host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported backups** pane, ensure that the backup from which you want to recover data is selected and press **Enter**.
2. In the **Restore points** pane, select with the **Up** and **Down** keys the restore point from which you want to recover data and press **Enter**.



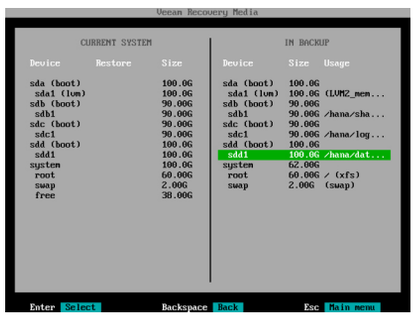
© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Step 5. Restore the Volume

The **In backup** pane of the **Veeam Recovery Media** wizard contains a list of disks and volumes in the backup. You can select volumes in the backup that you want to restore to your computer and specify mapping rules for these volumes.

To map a source volume to a target volume:

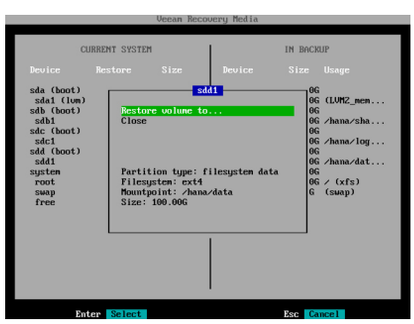
1. In the **In backup** pane, select a volume in the backup whose data you want to recover and press **Enter**. (For BMR restore you can select the option to restore whole disk, Veeam will perform the restore of all volumes from a disk in the backup.)



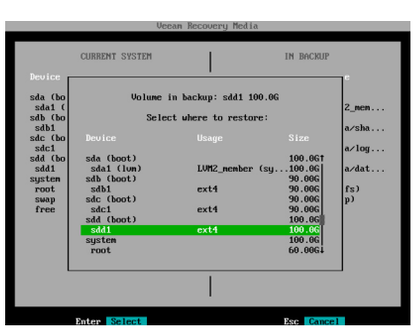
2. Veeam Agent *for Linux* will display a window with information on the selected volume (partition type, file system type, mount point and volume size) and a list of available operations:

- **Restore volume to** — select this option if you want to restore the selected volume to your computer.
- **Close** — select this option if you want to close the window and select another volume.

Select **Restore volume** to option.

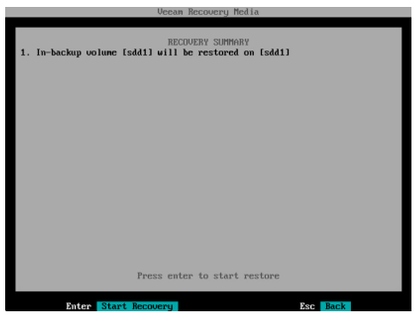


3. Veeam Agent *for Linux* will display a list of volumes on your computer. Select the volume that you want to restore and press **Enter**.

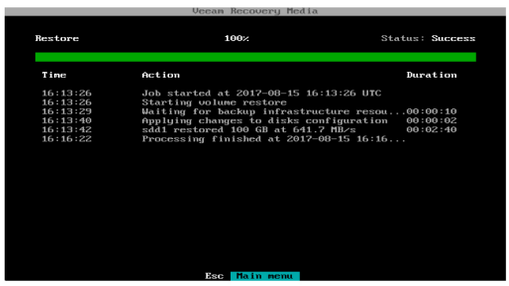


© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

4. Review the specified recovery settings.



5. Press **Enter** to start the volume-level restore process.



Physical SAP HANA Recovery — File-Level Restore

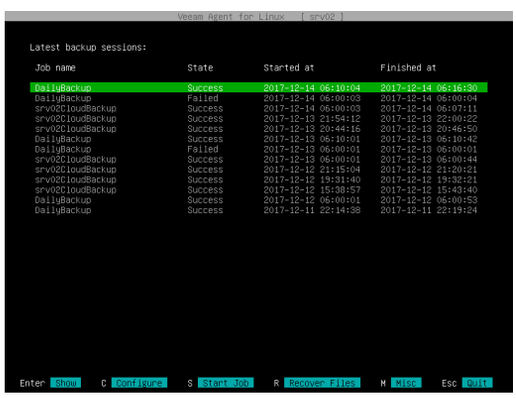
With Veeam file-level restores, users can restore individual files (data or logs) from SAP HANA backup.

Step 1. Launch File-Level Restore Wizard

To launch the **File-Level Restore** wizard, do the following:

1. Launch the Veeam Agent *for Linux* control panel with the veeam or veeamconfig ui command.

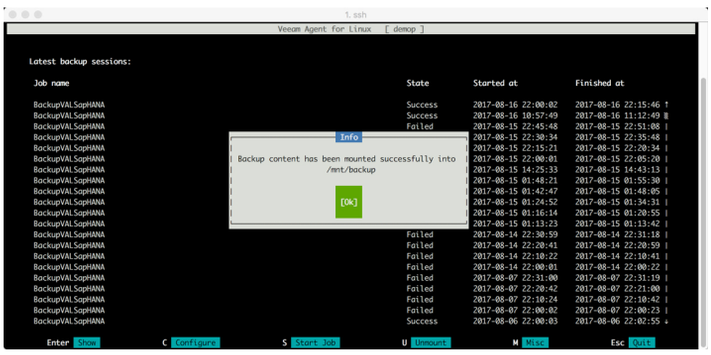
2. In the Veeam Agent *for Linux* control panel, press the "**r**" key to proceed to the File-Level Restore wizard.



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

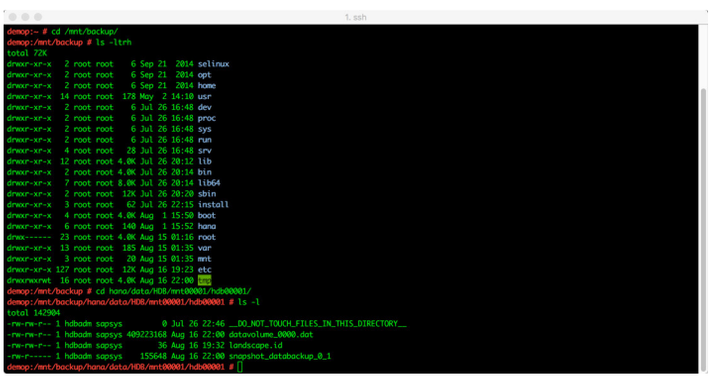
Step 2. Select Backup and Restore Point

Select the backup and restore point:



Step 3. Save Restored Files

When the backup file content is mounted to the /mnt/backup directory in the computer's file system, you can use Linux command line utilities or your preferred file browser to work with restored files and directories. You can browse for files and directories in the mounted backup and copy SAP HANA files and directories that you want to restore to their initial location or to a new location.



© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Restore operations from SAP HANA studio

After the SAP HANA Server is successfully restored, we will use SAP HANA studio to perform the HANA Storage Snapshot restore. Keep in mind for HANA 2.0 with MDC you have to recover the SYSTEM DB first and afterwards your single tenant. So, you need to run this process twice.

- Open **SAP HANA Studio**.
- Right-click on the SYSTEM@SID , Point to **Backup And Recovery**, Select **Recover System Database** (Figure 2).
- Specify **Recovery Type** – **Select Recover the database** to a **specific data backup** option (Figure 3).
- Select based on your environment **Recover using the backup catalogue** – file or backint(Figure 4).
- Select a backup to restore, select the available snapshot or backup for restore operations (Figure 5).
- Review the **recovery settings** and click **finish** (Figure 6).

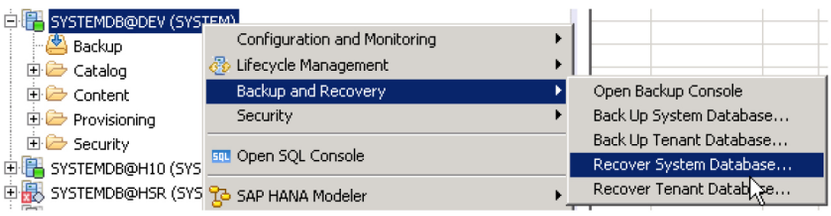


Figure 1. Verify SAP HANA services have stopped

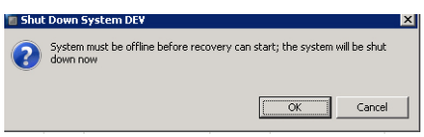


Figure 2. SYSTEM DB will be taken offline! Attention all you tenants will be resetted to the same stage and taken offline too

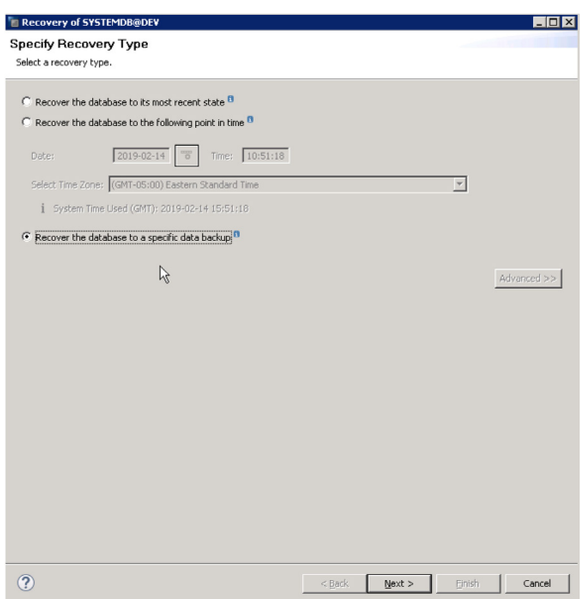


Figure 3. Recovery Type

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

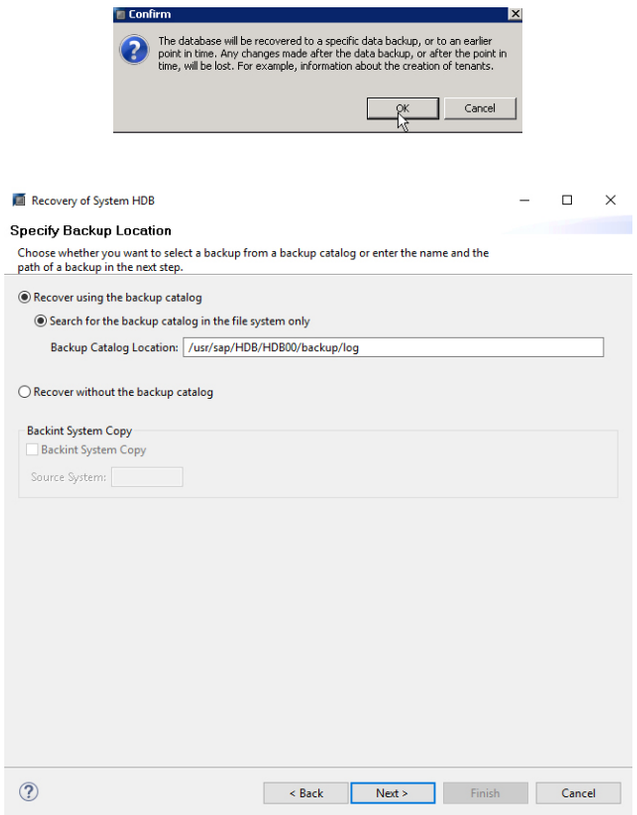
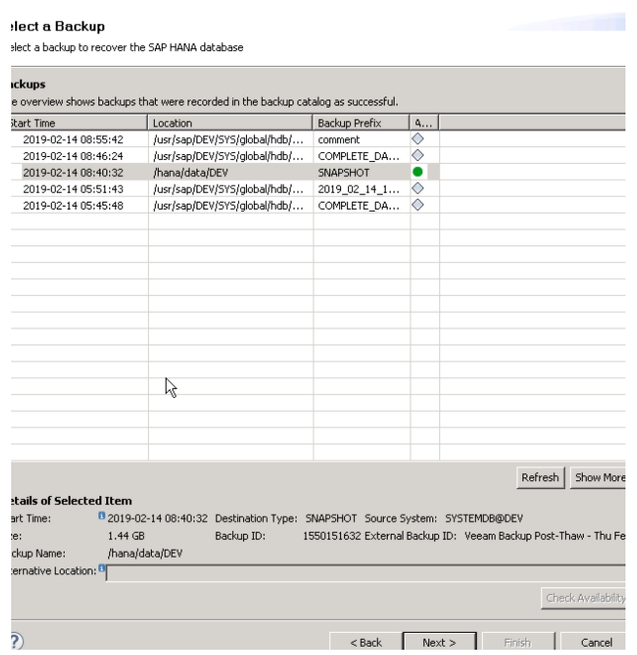


Figure 4. Specify Backup Location



Important (if your snapshot is marked red) : For File & Device restores — you need /hana/data/SID and /hana/log/SID restored.

Do not forget to do the same for your tenant!

Figure 6. Review Recovery Settings

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Start SAP HANA System

1. Open **SAP HANA Studio.**

2. Right-click on the **HDB(System),** point to **Configuration and Monitoring,** Select **Start System** (Figure 1).

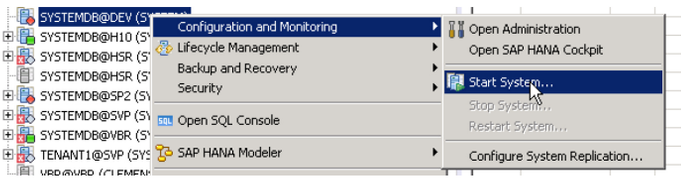


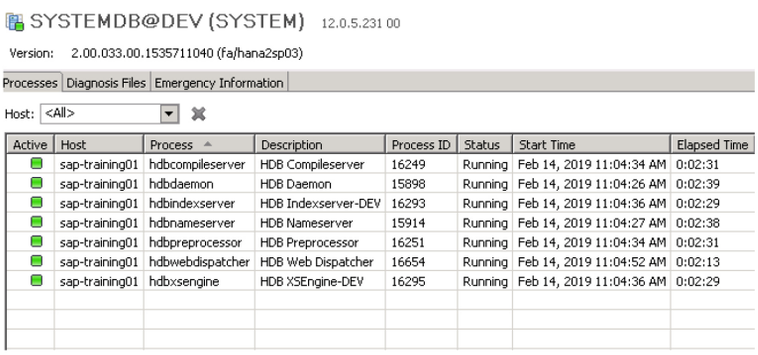
Figure 1. Start System

Validate SAP HANA Operational Status

Note: Keep in mind that you might have to start y

1. Open **SAP HANA Studio,** Click on **Processes** to see if all the processes are running as expected (Figure 1).

2. Click on **SYSTEM@SID** to see the **operational status**, operational status should be **all services started** (Figure 2).



Active	Host	Process	Description	Process ID	Status	Start Time	Elapsed Time
	sap-training01	hdbcompileserver	HDB Compileserver	16249	Running	Feb 14, 2019 11:04:34 AM	0:02:31
	sap-training01	hdbdaemon	HDB Daemon	15898	Running	Feb 14, 2019 11:04:26 AM	0:02:39
	sap-training01	hdbindexserver	HDB Indexserver-DEV	16293	Running	Feb 14, 2019 11:04:36 AM	0:02:29
	sap-training01	hdbnameserver	HDB Nameserver	15914	Running	Feb 14, 2019 11:04:27 AM	0:02:38
	sap-training01	hdbpreprocessor	HDB Preprocessor	16251	Running	Feb 14, 2019 11:04:34 AM	0:02:31
	sap-training01	hdbwebdispatcher	HDB Web Dispatcher	16654	Running	Feb 14, 2019 11:04:52 AM	0:02:13
	sap-training01	hdbxsengine	HDB XSEngine-DEV	16295	Running	Feb 14, 2019 11:04:36 AM	0:02:29

Figure 1. Validate SAP HANA services execution

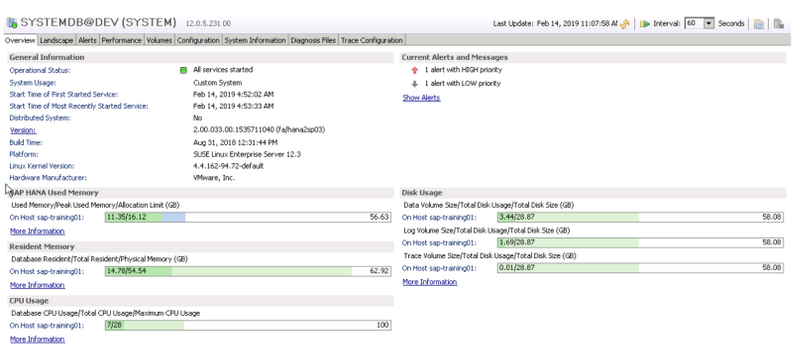


Figure 2. Operational Status

The SAP HANA Database has been successfully restored with Veeam Hyper-Availability Platform.

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Appendix

You can find the latest versions of the scripts under:

https://github.com/VeeamHub/applications



About Veeam Software

Veeam recognizes the new challenges companies across the globe face in enabling the Always-On Enterprise™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of Availability for the Always-On Enterprise™ by helping organizations meet recovery time and point objectives (RTPO™) of less than 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified recoverability, leveraged data and complete visibility.

Veeam Availability Suite™, which includes Veeam Backup & Replication™, leverages virtualization, storage and cloud technologies that enable the modern data center to help organizations save time, mitigate risks and dramatically reduce capital and operational costs, while always supporting the current and future business goals of Veeam customers.

Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit www.veeam.com

© 2019 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Intelligent Data Management

fuels the ***Fortune 500***



VEEAM

Intelligent Data Management
for the Hyper-Available Enterprise