

Data Protection for SAP HANA on Cisco HyperFlex Systems with Veeam Backup and Replication



Contents

Introduction	4
Purpose of this document	4
Test environment	4
Cisco Unified Computing System.....	5
Cisco HyperFlex system.....	6
Cisco HyperFlex HX-Series nodes.....	6
Cisco HyperFlex converged data platform software.....	7
Cisco HyperFlex HX Data Platform snapshot.....	7
Veeam Availability Suite	8
Backup	8
Restore.....	8
Instant virtual machine recovery	9
Virtual machine object recovery	9
Components.....	9
Backup server.....	10
Veeam repository sizing.....	11
Veeam Plug-in for SAP HANA	11
Solution design	12
Documents for installing the components	12
Configuring Cisco HyperFlex snapshot.....	13
Configure scheduled snapshots in VMware vSphere Client.....	13
Monitor snapshots in VMware vSphere.....	15
Revert the virtual machine	16
Configuring SAP HANA backup to file	19
Prepare the operating system for file backup.....	19
Configure file backup in SAP HANA Studio.....	20
Run a manual backup of SYSTEMDB to file	20
Run a manual backup of a tenant database to file	23
Recover a database from file	26
Configuring Veeam Backup and Replication managed snapshots with pre- and post-snapshot scripts.....	31
Download pre- and post-snapshot scripts	32
Create OS user for backup on the SAP HANA server.....	32
Create a secure HANA Database user store	34
Test the pre- and post-snapshot scripts	34
Configure a backup job in Veeam	35
Run a manual backup job.....	46
Restore a system from the Veeam backup repository	49
Configuring Cisco HyperFlex snapshot with virtual machine quiescing	56
Prepare the SAP HANA virtual machine	56
Configure scheduled snapshots in the VMware vSphere Client.....	56
Delete the backup catalog entry for deleted or replaced snapshots.....	59

Revert the virtual machine 60

Configuring SAP HANA Backint-based backup with the Veeam plug-in 60

 Configure the backup repository for backups from SAP HANA 61

 Install the Veeam Plug-in for SAP HANA 61

 Configure backup through Backint in SAP HANA Studio 64

 Perform manual backup from SAP HANA Studio..... 64

 Monitor backup progress and jobs in Veeam 67

 Back up the SAP HANA backup catalog 67

 Recover a database from Backint 69

Best practices 75

Conclusion 76

Introduction

With the general availability of SAP HANA running on hyperconverged systems such as the Cisco HyperFlex™ platform, organizations need to know the options available to them to protect their SAP HANA deployment and meet the service-level agreements (SLAs) required by the business. Cisco has partnered with Veeam to create this document to present the various options for protecting virtualized SAP HANA systems. Veeam Backup and Replication, a virtual machine-centric data protection suite, is certified with Backint for SAP HANA and can be integrated into the Cisco HyperFlex converged data platform to manage snapshot and replication tasks. This combination of products provides protection on the storage layer (Cisco HyperFlex system), infrastructure layer (Veeam suite), and application layer (Backint for SAP HANA).

Purpose of this document

This document describes data protection options for SAP HANA running on a Cisco HyperFlex system, including application-based, Cisco HyperFlex system-based, and Veeam Backup and Replication-managed approaches. This document does not replace the documentation available from SAP that discusses backup and recovery procedures, but instead serves as an extension to that documentation. This document focuses on the initial configuration of the various options and not on their daily operations. This document also does not provide detailed explanations of the various data protection methodologies. To use this document, you should have a basic understanding of data protection processes (backup, restore, and disaster recovery) and know the relevant terminology.

Note: The SAP HANA on Cisco HyperFlex installation and the Veeam Availability Suite on Cisco UCS installation are discussed in separate documents.

The following processes are covered:

- Automated Cisco HyperFlex native snapshots without SAP HANA integration
- SAP HANA backup to file
- Cisco HyperFlex native snapshots managed by Veeam with SAP HANA integration
- Manual Cisco HyperFlex native snapshots with SAP HANA integration
- SAP HANA Backint-based backup with Veeam Availability Suite for data and log files

Test environment

This section introduces the components and technologies used in the lab to test the processes described in this document.

Table 1 lists the hardware and software versions used in the test environment described in this document.

Test environment details

Layer	Component	Image or version
Computing	Cisco UCS 6332-16UP Fabric Interconnect pair	Release 4.0(1b)
	Cisco HyperFlex HX240 M5 Node	Release 4.0(1b)
	Cisco UCS S3260 M5 Storage Switch	Release 4.0(1b)
Network	Cisco Nexus® 9372PX Switch pair	Release 7.0(3)I2(2d)
Software	Cisco UCS® Manager	Release 4.0(1b)
	Cisco HyperFlex HX Data Platform Software	Release 3.5(1a)
	Veeam Availability Suite	Release 9.5 Update 4
	SAP HANA Platform Edition	Release 2.0 SP20

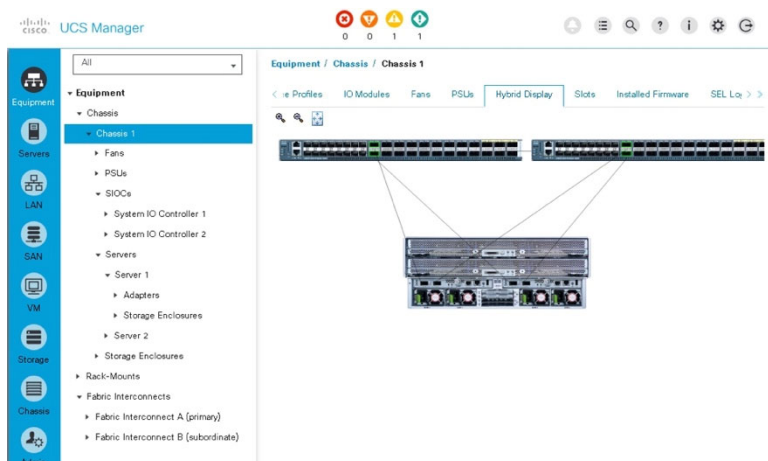
Cisco Unified Computing System

The Cisco Unified Computing System™ (Cisco UCS) is a state-of-the-art data center platform that unites computing, network, storage access, and virtualization resources into a single cohesive system.

Cisco UCS consists of these main resources:

- **Computing:** The system is based on an entirely new class of computing system that incorporates rack-mount and blade servers using Intel® Xeon® processor CPUs. The Cisco UCS servers offer patented Cisco® Extended Memory Technology to support applications with large data sets and allow more virtual machines per server.
- **Network:** The system is integrated onto a low-latency, lossless, 10- or 40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- **Virtualization:** The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access:** The system provides consolidated access to both SAN storage and network-attached storage (NAS) over the unified fabric. By unifying the storage access layer, Cisco UCS can access storage over Ethernet (with Network File System [NFS] or Small Computer System Interface over IP [iSCSI]), Fibre Channel, and Fibre Channel over Ethernet (FCoE). This approach provides customers with choice for storage access and investment protection. In addition, server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management for increased productivity.

Cisco UCS Manager



Cisco UCS consists of the following components:

- [Cisco UCS Manager](#) provides unified, embedded management of all software and hardware components in the Cisco Unified Computing System (Figure 1).
- [Cisco UCS 6000 Series Fabric Interconnects](#) are line-rate, low-latency, lossless, 10- or 40-Gbps Ethernet and FCoE interconnect switches that provide the management and communication backbone for Cisco UCS.
- [Cisco UCS 5100 Series Blade Server Chassis](#) supports up to eight blade servers and up to two fabric extenders in a 6-rack-unit (6RU) enclosure.

- [Cisco UCS B-Series Blade Servers](#) are Intel-based blade servers that increase performance, efficiency, versatility, and productivity.
- [Cisco UCS C-Series Rack Servers](#) deliver unified computing in an industry-standard form factor to reduce total cost of ownership (TCO) and increase agility.
- [Cisco UCS S-Series Storage Servers](#) deliver unified computing in an industry-standard form factor to address data-intensive workloads with reduced TCO and increased agility.
- [Cisco UCS adapters](#) with wire-once architecture offer a range of options to converge the fabric, optimize virtualization, and simplify management.

Cisco UCS is designed to deliver:

- Reduced TCO and increased business agility
- Increased IT staff productivity through just-in-time provisioning and mobility support
- A cohesive, integrated system that unifies the technology in the data center
- Industry standards supported by a partner ecosystem of industry leaders
- Unified, embedded management for easy-to-scale infrastructure

Cisco HyperFlex system

The Cisco HyperFlex system combines the industry-leading convergence of computing and networking resources provided by Cisco UCS with next-generation hyperconverged storage software to uniquely provide the computing resources, network connectivity, storage, and hypervisor platform needed to run an entire virtual environment. Everything is contained in a single uniform system.

Hyperconverged infrastructure offers some important advantages. It simplifies deployment and day-to-day management operations and increases agility, thereby reducing operating costs. Because hyperconverged storage can easily be managed by an IT generalist, it can also reduce technical costs on into the future that often otherwise arise when you implement complex systems that need dedicated management teams and skill sets.

The Cisco HyperFlex HX Data Platform is a purpose-built, distributed log-based file system that delivers high performance plus many data management and optimization features required in enterprise-class storage systems. This platform offers independent scaling of storage and computing resources; continuous data optimization through inline compression and deduplication; dynamic data distribution for increased data availability; and integrated native snapshots, rapid cloning, encryption, and virtual machine-level replication. This agile system is quick to deploy and easy to manage, it is scalable and flexible so that it can adapt to changing workloads, and it provides a high level of data security and availability.

Cisco HyperFlex HX-Series nodes

Cisco HyperFlex systems let you unlock the full potential of hyperconvergence and adapt IT to the needs of your workloads. The systems use an end-to-end software-defined infrastructure approach, combining software-defined computing in the form of Cisco HyperFlex HX-Series nodes, software-defined storage with the powerful Cisco HyperFlex HX Data Platform, and software-defined networking with the Cisco UCS fabric that integrates smoothly with the Cisco Application Centric Infrastructure (Cisco ACI™) platform. Together with a single point of connectivity and management, these technologies deliver a pre-integrated and adaptable cluster with a unified pool of resources that you can quickly deploy, scale, and manage to efficiently power your applications and your business.

A Cisco HyperFlex cluster requires a minimum of three HX-Series nodes. Data is replicated across at least two of these nodes, and a third node is required for continuous operation in the event of a single-node failure. The HX-Series nodes combine the CPU and RAM resources for hosting guest virtual machines, with the physical storage resources used by the Cisco HyperFlex

software. Each HX-Series node is equipped with one high-performance solid-state disk (SSD) for data caching and rapid acknowledgment of write requests. For maximum data capacity, each node also is equipped with spinning disks up to the platform’s physical capability.

Cisco HyperFlex converged data platform software

The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, distributed file system with a wide array of enterprise-class data management services. The data platform’s innovations redefine distributed storage technology, exceeding the boundaries of first-generation hyperconverged infrastructures. The data platform has all the features that you would expect of an enterprise shared storage system, eliminating the need to configure and maintain complex Fibre Channel storage networks and devices. The platform simplifies operations and helps ensure data availability. Enterprise-class storage features include the following:


- Replication of all written data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- Deduplication is always on, helping reduce storage requirements in the event that multiple operating system instances in client virtual machines result in large amounts of duplicate data.
- Compression further reduces storage requirements, reducing costs, and the log-structured file system is designed to store variable-sized blocks, reducing internal fragmentation.
- Thin provisioning allows large volumes to be created without requiring storage to support them until the need arises, simplifying data volume growth and making storage a “pay as you grow” proposition.
- Fast, space-efficient clones rapidly replicate virtual machines simply through metadata operations.
- Snapshots help facilitate backup and remote-replication operations: needed in enterprises that require always-on data availability.

Cisco HyperFlex HX Data Platform snapshot

The HX Data Platform uses metadata-based, zero-copy snapshots to facilitate backup operations. These space-efficient snapshots provide an excellent way to make frequent online backup copies of data without worrying about the consumption of physical storage capacity. Data can be moved offline or restored from snapshots instantaneously. Because Cisco HyperFlex systems are integrated with VMware vSphere and vCenter, HX Data Platform snapshots can be controlled with vSphere Snapshot Manager using the native vSphere interface.

If you create a regular virtual machine snapshot for backup, the workflow will create a VMware redo-log snapshot; the backup tool will read all data while VMware redirects all write operations to the redo-log file. After the backup and replica process is complete, the virtual machine snapshot deletion process is initiated. All data from the redo log must be written to the data files before the redo log can be deleted, as shown on the left side of Figure 2. The virtual machine snapshot deletion process can take a long time and consume a large amount of storage resources until the operation is finished.

Virtual machine backup: Standard operation and operation with Cisco HyperFlex snapshot integration



By integrating the Cisco HyperFlex platform, you completely avoid the use of VMware virtual machine snapshots. As shown on the right side of Figure 2, the workflow will create a Cisco HyperFlex snapshot. Veeam then reads all data out of the snapshots. After the backup and replica operations are complete, Cisco deletes the Cisco HyperFlex snapshot, with almost no impact on the virtual machine or the rest of the production environment.

Veeam Availability Suite

Veeam Availability Suite combines the backup, restore and replication capabilities of Veeam Backup and Replication with the advanced monitoring, reporting and capacity planning functions of Veeam ONE.

Backup

Veeam Backup and Replication operates at the virtualization layer and uses an image-based approach for virtual machine backup. To retrieve virtual machine data, no agent software needs to be installed in the guest OS. Instead, Veeam Backup and Replication uses vSphere snapshot capabilities and application-aware processing. When a new backup session starts, a snapshot is taken to create a cohesive point-in-time copy of a virtual machine, including its configuration, OS, applications, associated data, and system state. Veeam Backup and Replication uses this point-in-time copy to retrieve virtual machine data. Image-based backups can be used for different types of recovery, including full virtual machine recovery, virtual machine file recovery, instant virtual machine recovery, and file-level recovery.

Use of the image-based approach allows Veeam Backup and Replication to overcome the limitations of traditional backup processes. It also helps simplify recovery verification and the restore process: to recover a single virtual machine, you do not need to perform multiple restore operations. Veeam Backup and Replication uses a cohesive virtual machine image from the backup repository to restore a virtual machine to the required state without the need for any manual reconfiguration or adjustment.

With Veeam Backup and Replication, backup is a job-driven process in which one backup job can be used to process one or more virtual machines. The job is the configuration unit for the backup activity. Essentially, a job defines when, what, how, and where data is backed up. It indicates what virtual machines should be processed, what components should be used to retrieve and process virtual machine data, what backup options should be enabled, and where the resulting backup file should be saved. Jobs can be started manually by the user or scheduled to run automatically. The resulting backup file stores compressed and deduplicated virtual machine data. Compression and deduplication is performed by the Veeam proxy server.

Regardless of the backup method you use, the first run of a job creates a full backup of the virtual machine image. Subsequent job runs are incremental: Veeam Backup and Replication copies only those data blocks that have changed since the last backup job was run. To keep track of changed data blocks, Veeam Backup and Replication uses several approaches, including VMware's Changed Block Tracking (CBT) technology.

Restore

Veeam Backup and Replication offers a number of recovery options for various disaster recovery scenarios:

- Veeam Explorer enables you to restore single application items.
- Instant virtual machine recovery enables you to instantly start a virtual machine directly from a backup file.
- Full virtual machine recovery enables you to recover a virtual machine from a backup file to its original or another location.
- Virtual machine file recovery enables you to recover separate virtual machine files (virtual disks, configuration files, and so on).
- Virtual drive restore enables you to recover a specific hard drive of a virtual machine from the backup file and attach it to the original virtual machine or to a new virtual machine.
- Microsoft Windows file-level recovery enables you to recover individual Windows guest OS files (from File Allocation Table [FAT], New Technology File System [NTFS], and Resilient File System [ReFS] file systems).

- Multi-OS file-level recovery enables you to recover files from 15 different guest OS file systems.

Veeam Backup and Replication uses the same image-level backup process for all data recovery operations. You can restore virtual machines, virtual machine files and drives, application objects, and individual guest OS files to the most recent state or to any available restore point.

Instant virtual machine recovery

With instant virtual machine recovery, you can immediately restore a virtual machine to your production environment by running it directly from the backup file. Instant virtual machine recovery helps improve your recovery-time objective (RTO) and reduce disruption and downtime on production virtual machines. It is like having a temporary spare for a virtual machine. Users can remain productive while you troubleshoot the problem in the failed virtual machine.

When instant virtual machine recovery is performed, Veeam Backup and Replication uses the Veeam vPower technology to mount a virtual machine image to a VMware ESXi host directly from a compressed and deduplicated backup file. Because you do not need to extract the virtual machine from the backup file and copy it to production storage, you can restart a virtual machine from any restore point (incremental or full) in minutes.

After the virtual machine is back online, you can use VMware Storage vMotion to migrate the virtual machine back to production storage.

Virtual machine object recovery

Veeam Backup and Replication can help you to restore specific virtual machine files (.vmdk, .vmx, and others) if any of these files are deleted or the data store is corrupted. This option provides a great alternative to full virtual machine restoration, for example, when your virtual machine configuration file is missing and you need to restore it. Instead of restoring the whole virtual machine image to production storage, you can restore only the specific virtual machine file.

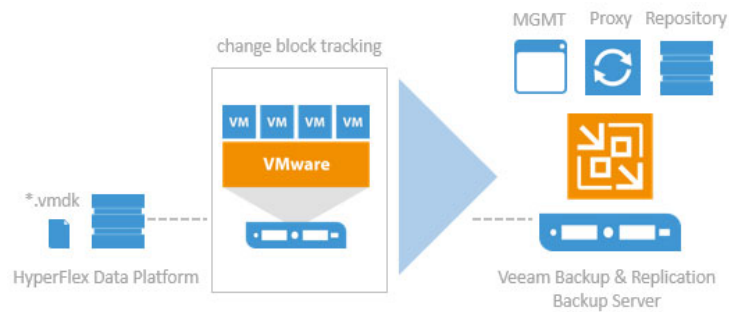
Another data recovery option provided by Veeam Backup and Replication is restoration of a specific hard drive of a virtual machine. If a virtual machine hard drive becomes corrupted for some reason (for example, by a virus), you can restore it from the image-based backup file to any known-good point in time.

Components

Veeam Availability Suite provides backup, restore, and replication capabilities plus advanced monitoring, reporting, and capacity planning functions. Veeam Availability Suite delivers everything you need to reliably ensure and manage your Cisco HyperFlex and VMware environment. Veeam Backup and Replication is a modular solution that lets you build a scalable backup infrastructure for environments of different sizes and configurations. The installation package of Veeam Backup and Replication includes a set of components that you can use to configure the backup infrastructure. Some components are mandatory and provide core functions, and some components are optional and can be installed to provide additional features to meet your particular business and deployment needs. You can co-install all Veeam Backup and Replication components on the same machine, physical or virtual, or you can set them up separately for a more scalable approach.

Figure 3 provides an overview on the main Veeam components.

Veeam Backup and Replication components



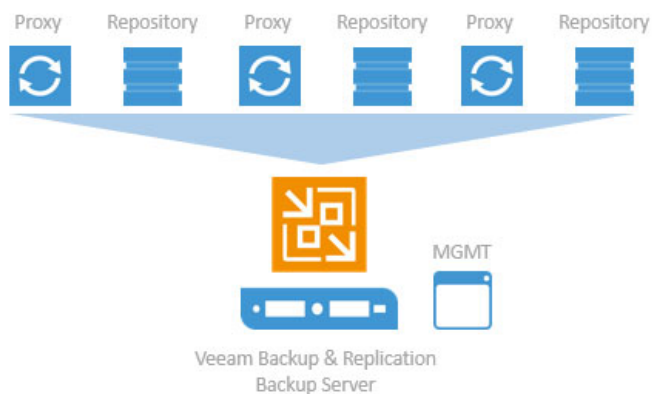
Backup server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup and Replication is installed. It is the core component in the backup infrastructure, filling the role of configuration and control center. The backup server performs all types of administrative activities:

- It coordinates backup, replication, recovery verification, and restore tasks.
- It controls job scheduling and resource allocation.
- It manages all proxy and repository servers.

It is used to set up and manage backup infrastructure components as well as specify global settings for the backup infrastructure (Figure 4).

Veeam backup server management



In addition to its primary functions, a newly deployed backup server also performs the roles of default backup proxy and backup repository.

The backup server uses the following services and components:

- Veeam Backup Service is a Windows service that coordinates all operations performed by Veeam Backup and Replication, such as backup, replication, recovery verification, and restore tasks. The Veeam Backup Service runs under the local system account or the account that has local administrator permissions on the backup server.
- Veeam Backup Shell provides the application user interface and allows user access to the application's functions.
- Veeam Guest Catalog Service is a Windows service that manages guest OS file system indexing for virtual machines and replicates system index data files to enable searches through guest OS files. Index data is stored in the Veeam backup catalog: a folder on the backup server. The Veeam Guest Catalog Service running on the backup server works in

conjunction with search components installed on Veeam Backup Enterprise Manager and (optionally) a dedicated Microsoft search server.

- Veeam Backup SQL Database is used by Veeam Backup Service, Backup Shell, and Guest Catalog Service to store data about the backup infrastructure, jobs, sessions, and so on. The database instance can be located on a Microsoft SQL Server installed either locally (on the same machine on which the backup server is running) or remotely.
- Veeam Backup PowerShell Snap-In is an extension to Microsoft Windows PowerShell 2.0. Veeam Backup PowerShell adds a set of cmdlets to allow users to perform backup, replication, and recovery tasks through the command-line interface (CLI) of PowerShell or run custom scripts to fully automate operation of Veeam Backup and Replication.
- Backup Proxy Services are a set of data movement services. The backup server runs these services in addition to dedicated services.

Veeam repository sizing

To estimate the amount of required disk space, you should know the following:

- Total size of the virtual machines being backed up
- Frequency of backup operations
- Retention period for backup files
- Whether jobs will use forward or reverse incremental backup processes

In addition, when testing is not possible beforehand, you should make assumptions about compression and deduplication ratios, change rate, and other factors. The following figures are typical for most deployments; however, you need to understand the specific environment to identify possible exceptions:

- Data reduction from compression and deduplication is usually 2:1 or greater. A ratio of 3:1 or greater is common, but you should always be conservative when estimating required space.
- The typical daily change rate is between 2 and 5 percent in a midsize or enterprise environment. However, this rate can vary greatly among servers, with some servers showing much higher values. If possible, run monitoring tools such as Veeam ONE to gain a better understanding of the actual change rates in your system.
- Include additional space for occasional full backups.
- Include additional space for backup chain transformation (forward forever incremental to reverse incremental).
- Include space equal to at least the size of a full backup multiplied by 1.25.

Using these numbers, you can estimate the required disk space for any job. In addition, you always should leave plenty of headroom for future growth, additional full backups, movement of virtual machines, and restoration of virtual machines from tape.

Note: A repository sizing tool that can be used for estimation is available at <http://vee.am/rps>. Note that this tool is not officially supported by Veeam. Nonetheless, it is heavily used by Veeam architects. The tool is regularly updated.

Veeam Plug-in for SAP HANA

The Veeam Plug-in for SAP HANA is an SAP certified Backint interface that allows you to perform backup and restore operations within the native SAP HANA tools such as SAP HANA Cockpit, SAP HANA Studio, and SQL commands. The plug-in allows you to run all SAP HANA Backint-operated backup modes such as full, differential, and incremental backups, in addition to predefined log backups. With the Veeam Plug-in for SAP HANA, you can restore your complete database to specific data points or to a timely defined data point or latest point in time.

The Veeam Plug-in for SAP HANA manages the SAP HANA Backint data stream from the SAP HANA database, adds deduplication and compression on the data stream, and transfers all the given data to one of the supported Veeam

repositories such as the SAP scale-out repository. For performance purposes, it also supports multistreaming for each SAP HANA service of more than 128 GB of data.

For restore operations, the plug-in delivers the required data streams from the Veeam repository back to the SAP HANA service and helps customers manage their SAP HANA environments much more easily.

Solution design


For you to follow the configuration steps in this document, your system must have the following components installed and configured:

- Cisco HyperFlex all-flash system
- One or more SAP HANA systems
- Veeam Backup and Recovery, installed on Cisco UCS C240 Rack Server, Cisco UCS S3260 Storage Server, and Cisco HyperFlex platform
- SAP HANA Studio


Figure 5 shows the components used in the system described in this document.

Components used in this document


VMware vSphere



SAP HANA Studio



Veeam Console



SAP HANA on Cisco HyperFlex



Veeam AS on Cisco UCS S3260



Documents for installing the components

Consult the following documents for detailed guidance in installing the components:

- Installing Cisco HyperFlex for SAP HANA: <https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/sap-hana-whitepaper.pdf>
- Installing Veeam on Cisco UCS: https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/data-protection-solutions/whitepaper_c11-739852.pdf
- Installing SAP HANA and SAP HANA Studio: https://help.sap.com/viewer/p/SAP_HANA_PLATFORM (select SAP HANA Master Guide)

Configuring Cisco HyperFlex snapshot

You can use Cisco HyperFlex snapshots to create point-in-time recovery points for the whole virtual machine directly on the storage platform. Because the snapshot is taken directly on the storage platform without any interaction with SAP HANA, there is no entry in the SAP HANA backup catalog or any internal snapshot that can be used as a recovery point. Under rare circumstances, such as when the snapshot is taken through a mass update in the SAP HANA database, SAP HANA may experience an inconsistent status and may not be able to start from the storage snapshot. Nevertheless, this kind of snapshot is a viable fast and simple option for starting the virtual machine in the event of a failure without overloading the SAP HANA backup catalog with snapshot entries.

This snapshot option should never be the only option used to protect SAP HANA and must be combined with other options discussed in this document.

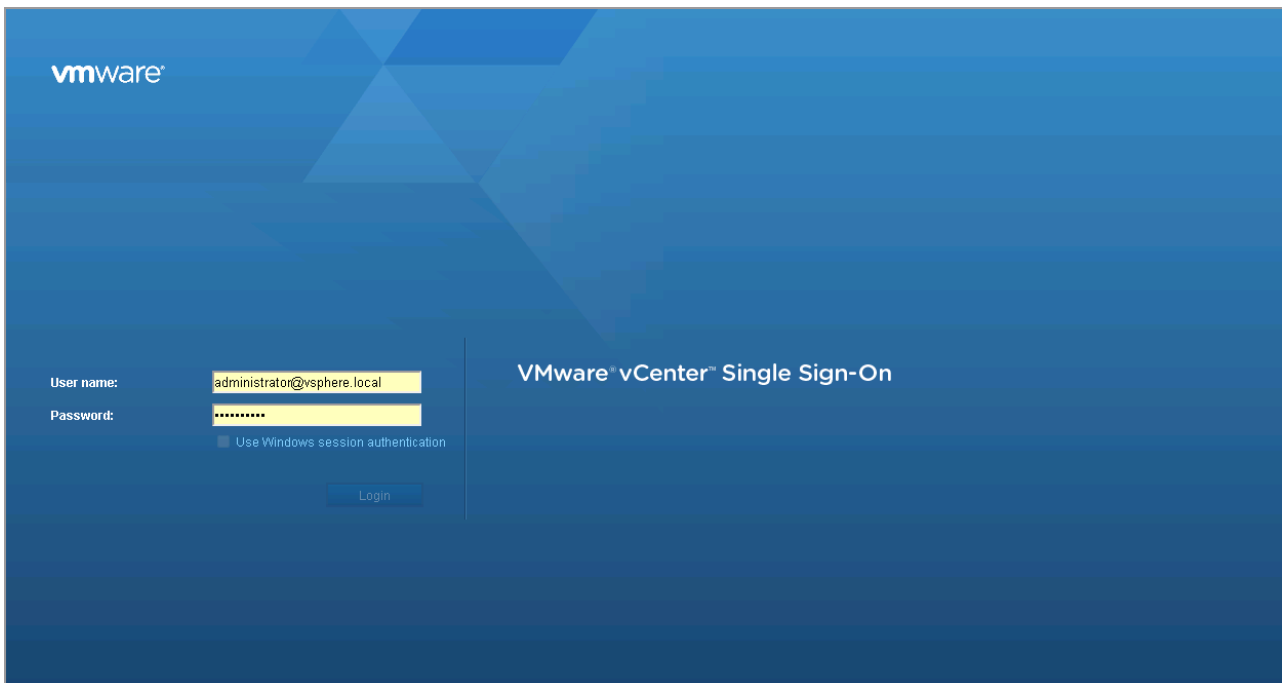
To configure this backup type, you perform the following high-level steps:

- Create a snapshot on the Cisco HyperFlex system for the virtual machine running SAP HANA.

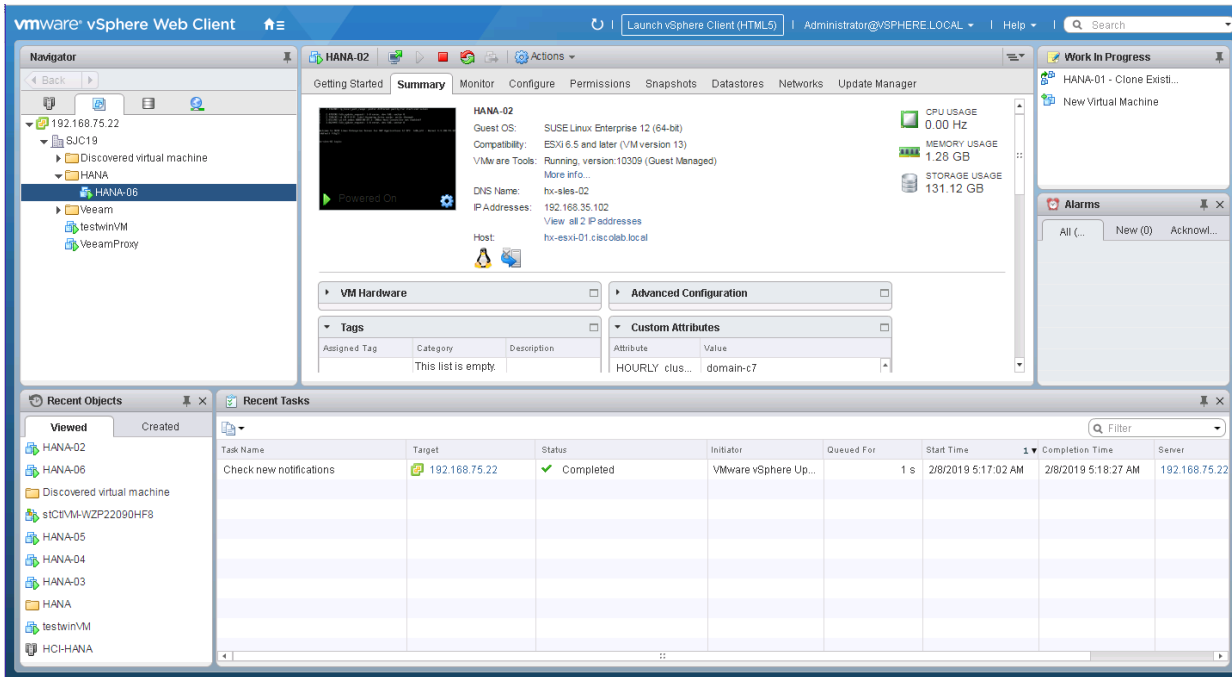
Configure scheduled snapshots in VMware vSphere Client

You first configure scheduled snapshots in the vSphere Client.

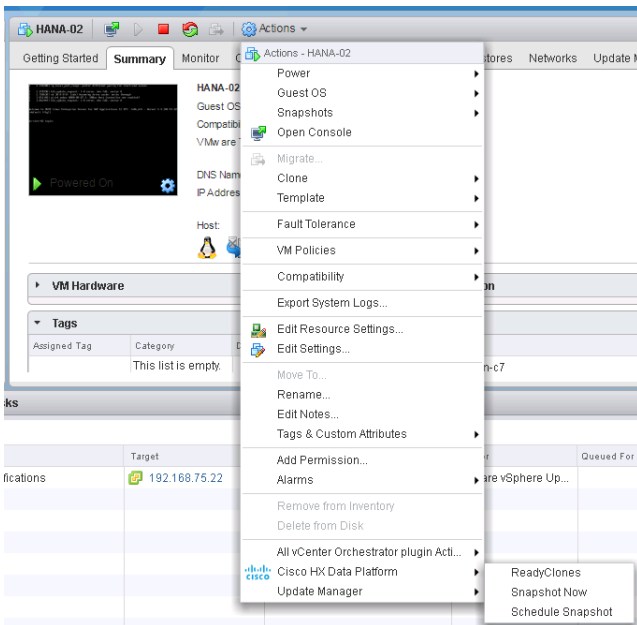
Connect to the vSphere Client from which the Cisco HyperFlex system is managed.



Select the virtual machine on which SAP HANA is running.



Choose Actions > Cisco HX Data Platform > Schedule Snapshot.



Select the checkbox for Enable Hourly Snapshot and specify the schedule and the number of snapshots that you want to keep that best suits your landscape. Then click OK.

The Cisco HyperFlex system will now automatically create snapshots based on the configured schedule.

Remember: These snapshots are not shown in the SAP HANA backup catalog.

Monitor snapshots in VMware vSphere

The snapshots are shown on the Snapshots tab of every individual virtual machine.

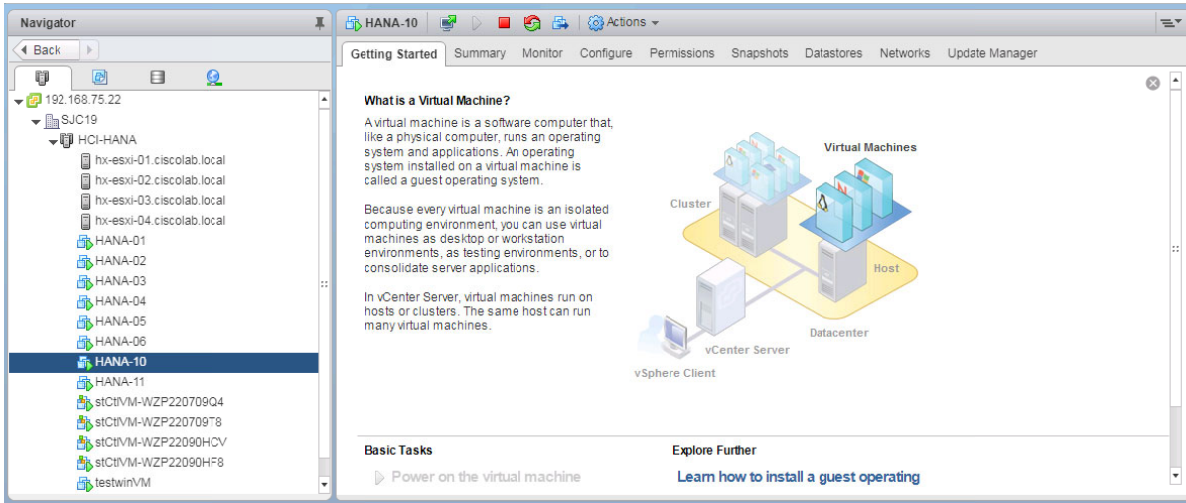
To monitor the snapshots, open the Snapshots tab on the virtual machine.

Snapshot details	
Name	VM-HOURLY-Fri-Feb-08-01:00:01-UTC-2019
Description	com.springpath.createSnapshotTask: task-3433
Created	Thursday, February 07, 2019 5:04:21 PM
Disk usage	7.21 GB
Snapshot the virtual machine's	No

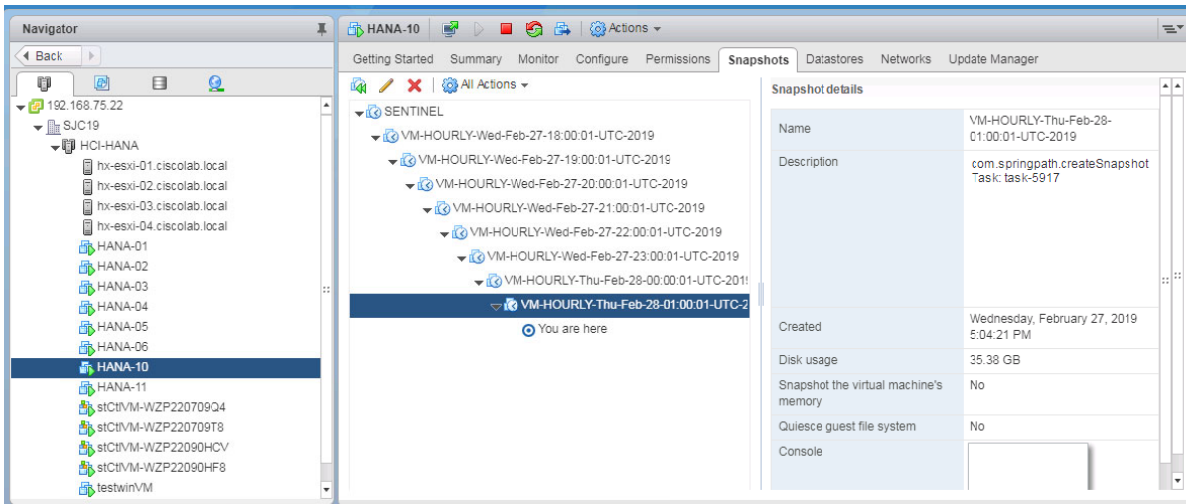
Revert the virtual machine

To restore the SAP HANA system from a snapshot, you can revert the virtual machine to that snapshot. If you want to preserve the actual state of the virtual machine, you should create a clone before reverting to a snapshot.

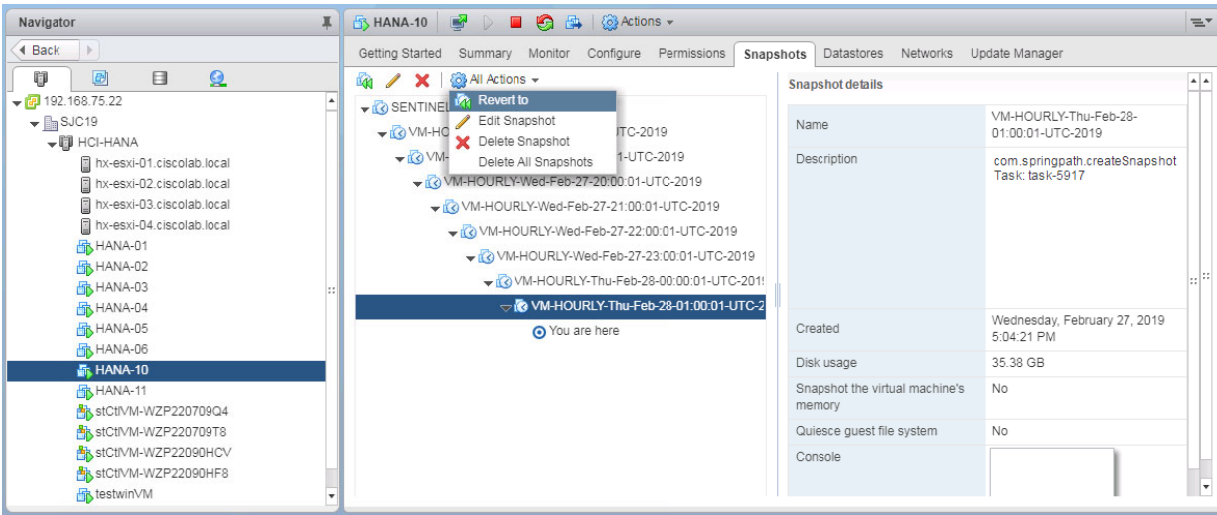
1. In the vSphere Web Client select the virtual machine that you want to revert.




2. Open the Snapshots tab and select the snapshot that you want to use. In most cases, you will want to use the latest snapshot.



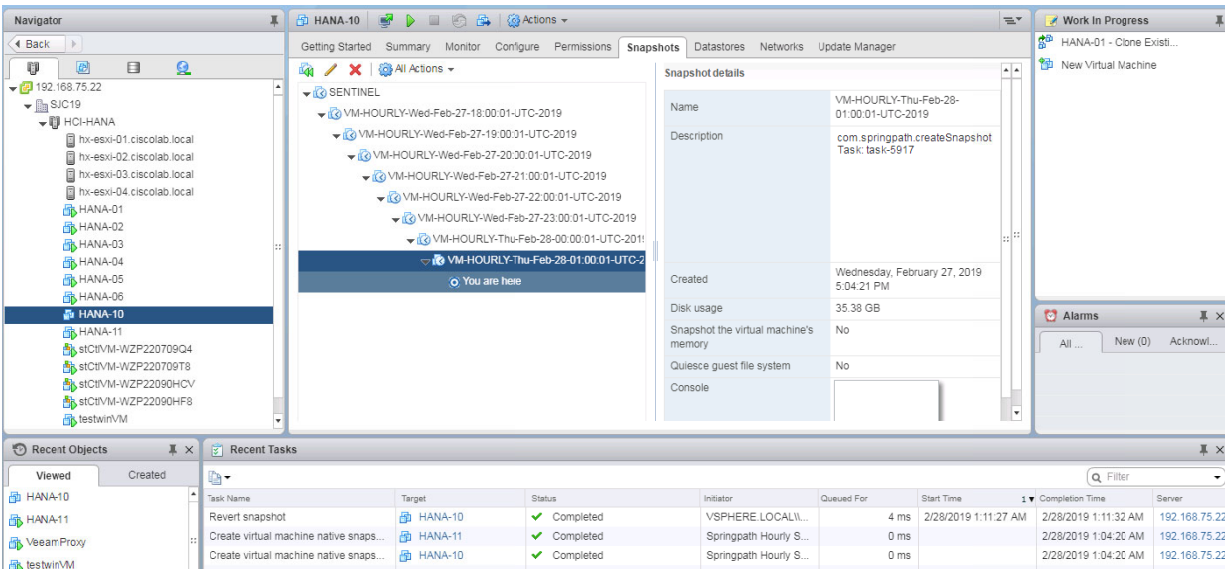
3. Click All Actions and choose “Revert to.”



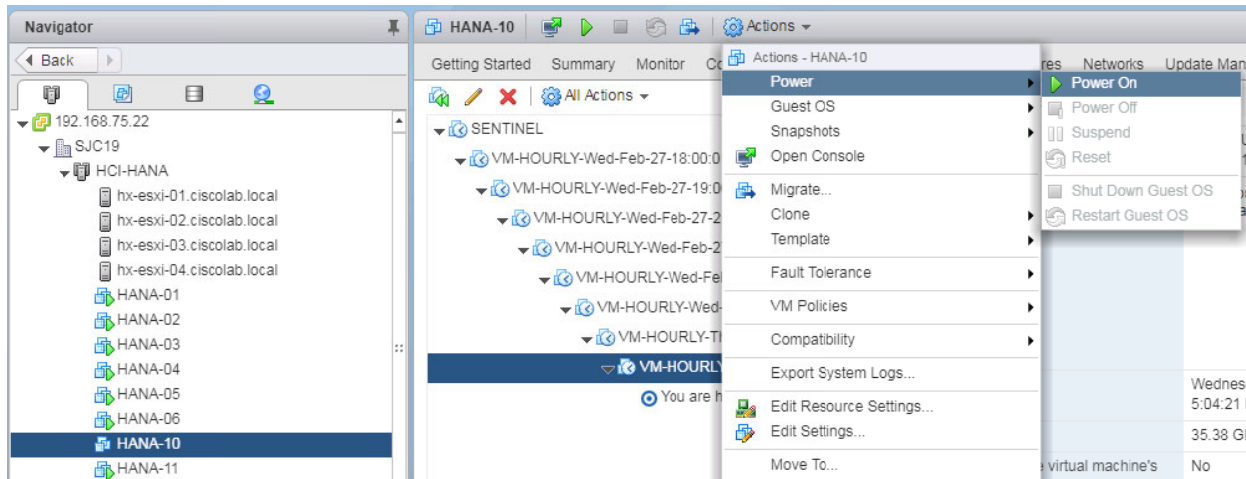
4. Confirm the process by clicking Yes.



On Cisco HyperFlex systems, the reversion process is very fast because a redo log does not need to be committed in the data files.



The power state of the virtual machine will automatically power off. You can now power on the virtual machine and start SAP HANA as usual.



Verify that SAP HANA is running before you try to access the database.

```
bwpadm@hx-sles-10:/usr/sap/BWP/HDB10> HDB info
```

USER	PID	PPID	%CPU	VSZ	RSS	COMMAND
bwpadm	24602	24601	0.0	15572	5468	-sh
bwpadm	24843	24602	0.0	13396	3416	_ /bin/sh /usr/sap/BWP/
bwpadm	24874	24843	0.0	36888	2824	_ ps fx -U bwpadm -
bwpadm	2814	1	0.0	21732	2988	sapstart pf=/hana/shared/
bwpadm	2823	2814	0.0	255776	49580	_ /usr/sap/BWP/HDB10/hx
bwpadm	2839	2823	2.4	15594892	13021944	_ hdbnameserver
bwpadm	3014	2823	0.2	3492084	425488	_ hdbcompileserver
bwpadm	3016	2823	0.3	8651028	7319292	_ hdbpreprocessor
bwpadm	3055	2823	2.4	14761732	11924760	_ hdbindexserver -p
bwpadm	3057	2823	0.4	5741536	1644464	_ hdbxsengine -port
bwpadm	3346	2823	0.2	3787280	445096	_ hdbwebdispatcher
bwpadm	2727	1	0.0	493908	28840	/usr/sap/BWP/HDB10/exe/sa

```
bwpadm@hx-sles-10:/usr/sap/BWP/HDB10>
```

You can connect to the database with the `hdbsql` command and check the status in detail. The established connection to the database indicates that the reversion process has likely succeeded.

```
bwpadm@hx-sles-10:/usr/sap/BWP/HDB10> hdbsql -i 10 -d BWP -u SYSTEM
Password:
```

```
Welcome to the SAP HANA Database interactive terminal.
```

```
Type: \h for help with commands
```

```
\q to quit
```

```
hdbsql BWP=>
```

Configuring SAP HANA backup to file

The simplest way to back up your data in SAP HANA is to use SAP HANA's backup-to-file option.

Cisco highly recommends that you do **not** store this backup file directly on the SAP HANA server. Otherwise, both the backup file and the original files may be lost if the SAP HANA server fails. You would also need to consider the number of backup versions that can be stored on the local file system before you run out of space. Cisco recommends using NFS mounted on a remote server or storage that is designed to store backup files and provides enough capacity to retain the number of backup versions that you need based on your retention-period requirements.

To configure this backup type, you perform the following high-level steps:

- Back up the data from SAP HANA to a remote file system.

Note: Do not use the local storage in the virtual machine as the backup destination.

Prepare the operating system for file backup

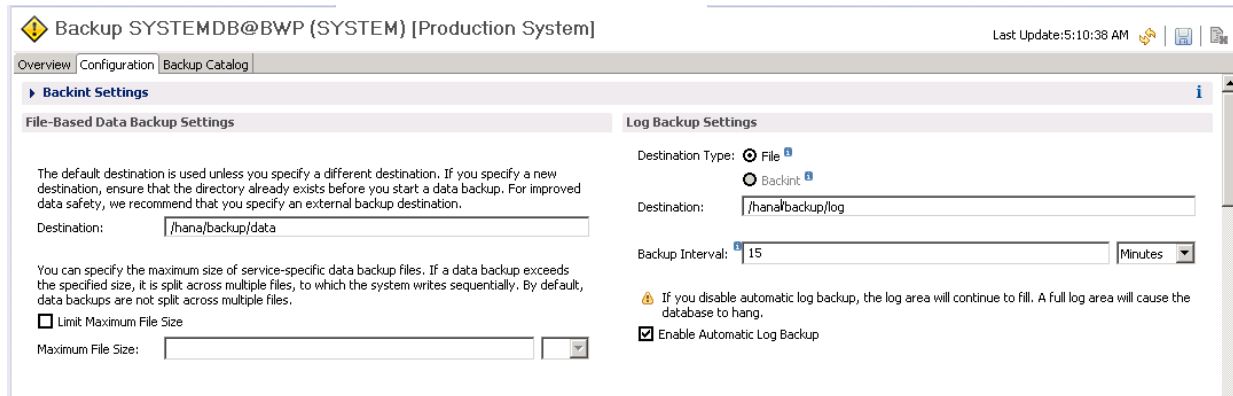
Enter the following commands to prepare the OS for file backup:

```
# mkdir /hana/backup
# mount <storage>:<share> /hana/backup
# mkdir /hana/backup/log
# mkdir /hana/backup/data
# chgrp -R sapsys /hana/backup
# chmod -R 774 /hana/backup
# ls -al
total 16
drwxrwxrwx  6 root root   4096 Feb 11 13:22 .
drwxr-xr-x 23 root root   4096 Jul 25  2018 ..
drwxrwxr--  4 root sapsys 4096 Feb 11 13:22 backup
drwxr-xr-x  4 root root     33 Feb  8 15:34 data
drwxr-xr-x  3 root root     17 Feb  8 15:34 log
drwxr-xr-x  3 root root   4096 Feb  8 15:34 shared
#
```

Configure file backup in SAP HANA Studio

You configure file backup in SAP HANA Studio.

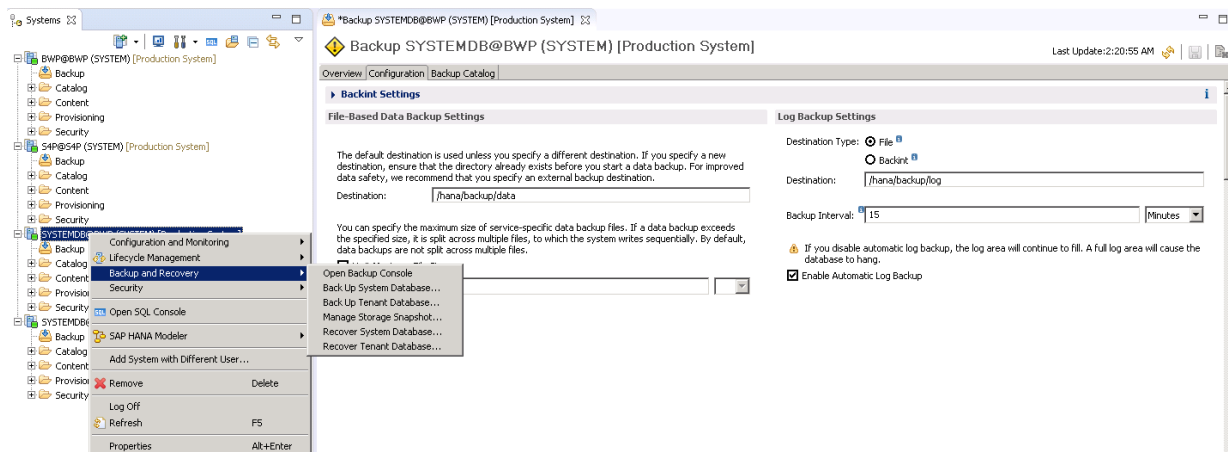
1. Log on to SAP HANA Studio and select the system database (SYSTEMDB) entry for the system that you want to back up. Open the Configuration tab for the backup settings.
2. Enter `/hana/backup/data` as the destination for File-Based Data Backup Settings and `/hana/backup/log` for Log Backup Settings and save the settings. Be sure that the checkbox for Enable Automatic Log Backups is selected.



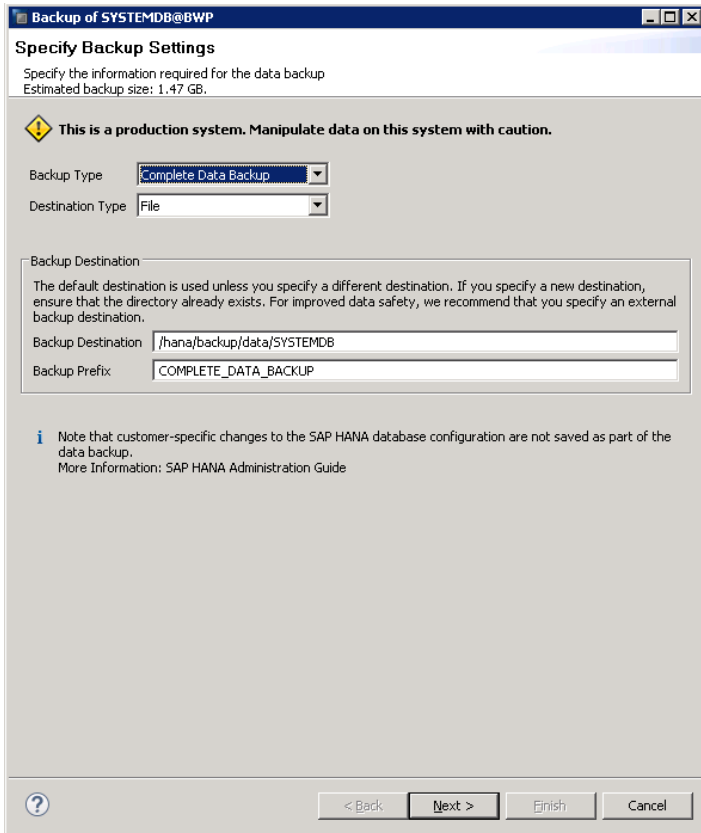
Run a manual backup of SYSTEMDB to file

Now manually test your configuration.

1. Right-click the SYSTEMDB entry and choose Backup and Recovery > Back Up System Database



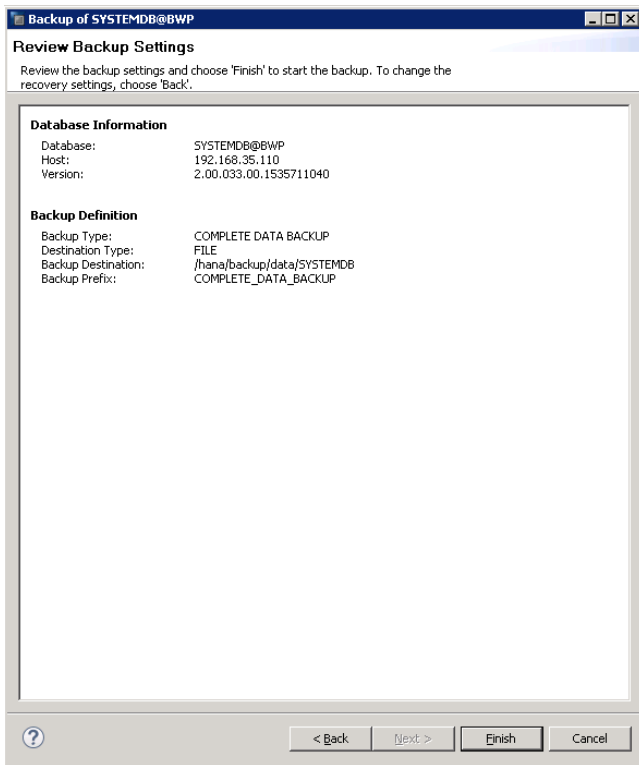
- In the Backup Type field, choose Complete Data Backup, and in the Destination Type field, choose File. Verify that the Backup Destination definition points to /hana/backup/data; if it does not, then the configuration was not saved properly. Click Next.



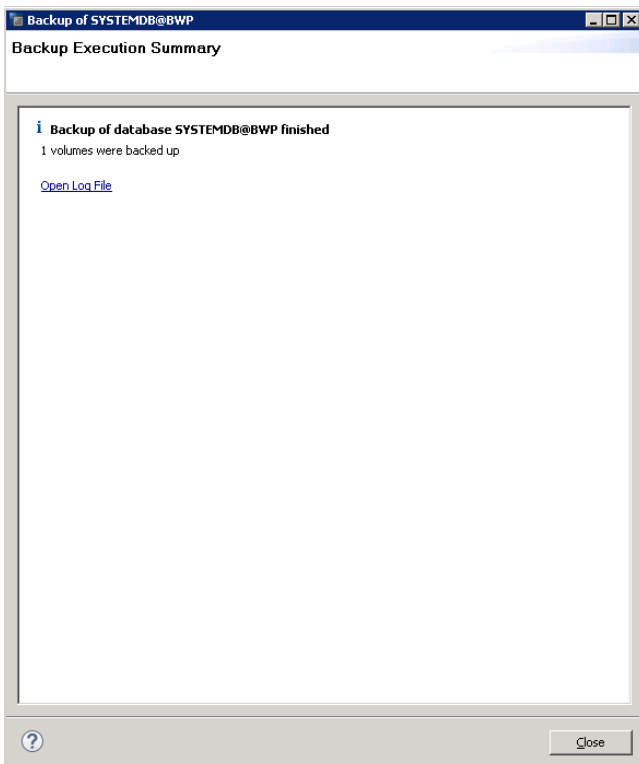
The screenshot shows a dialog box titled "Backup of SYSTEMDB@BWP" with the subtitle "Specify Backup Settings". It contains the following elements:

- A warning icon and text: "This is a production system. Manipulate data on this system with caution."
- Two dropdown menus: "Backup Type" set to "Complete Data Backup" and "Destination Type" set to "File".
- A "Backup Destination" section with a text input field containing "/hana/backup/data/SYSTEMDB" and a "Backup Prefix" field containing "COMPLETE_DATA_BACKUP".
- An information icon and text: "Note that customer-specific changes to the SAP HANA database configuration are not saved as part of the data backup. More Information: SAP HANA Administration Guide".
- Navigation buttons at the bottom: "< Back", "Next >", "Finish", and "Cancel".

- Review the backup job definition and click Finish.



4. After the backup job is finished, click Close.



5. Check the file system. The directory structure for the SYSTEMDB and tenant database is created and contains backup files for the SYSTEMDB data.

```
/hana/backup # ls -lR *
```

```
data:
```

```
total 8
```

```
drwxr-x--x 2 bwpadm sapsys 4096 Feb 11 14:56 DB_BWP
```

```
drwxr-x--x 2 bwpadm sapsys 4096 Feb 12 10:22 SYSTEMDB
```

```
data/DB_BWP:
```

```
total 0
```

```
data/SYSTEMDB:
```

```
total 1558168
```

```
-rw-r----- 1 bwpadm sapsys      155648 Feb 12 10:22 COMPLETE_DATA_BACKUP_databackup_0_1
```

```
-rw-r----- 1 bwpadm sapsys 1593843712 Feb 12 10:22 COMPLETE_DATA_BACKUP_databackup_1_1
```

```
log:
```

```
total 8
```

```
drwxr-x--x 2 bwpadm sapsys 4096 Feb 11 16:52 DB_BWP
```

```
drwxr-x--x 2 bwpadm sapsys 4096 Feb 11 16:48 SYSTEMDB
```

```
log/DB_BWP:
```

```
total 0
```

```
log/SYSTEMDB:
```


```
total 0
```

```
/hana/backup #
```

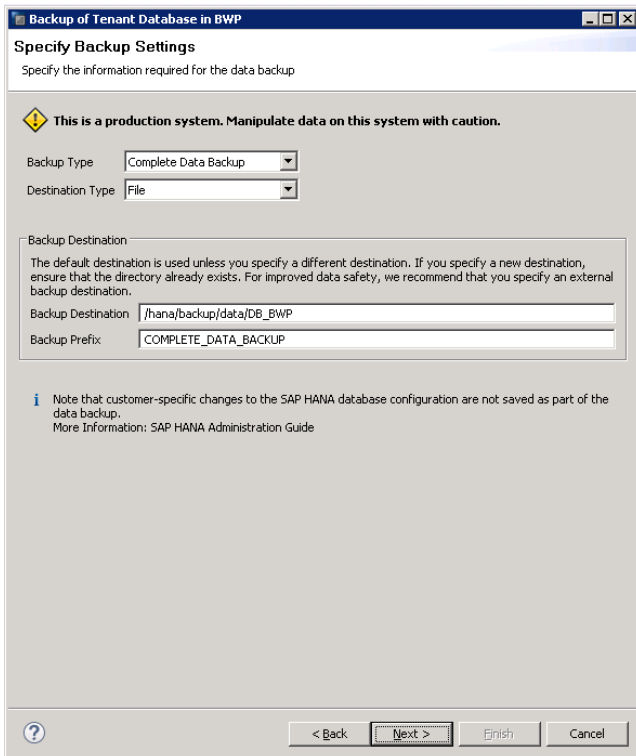
Run a manual backup of a tenant database to file

Now test the configuration for the tenant database.

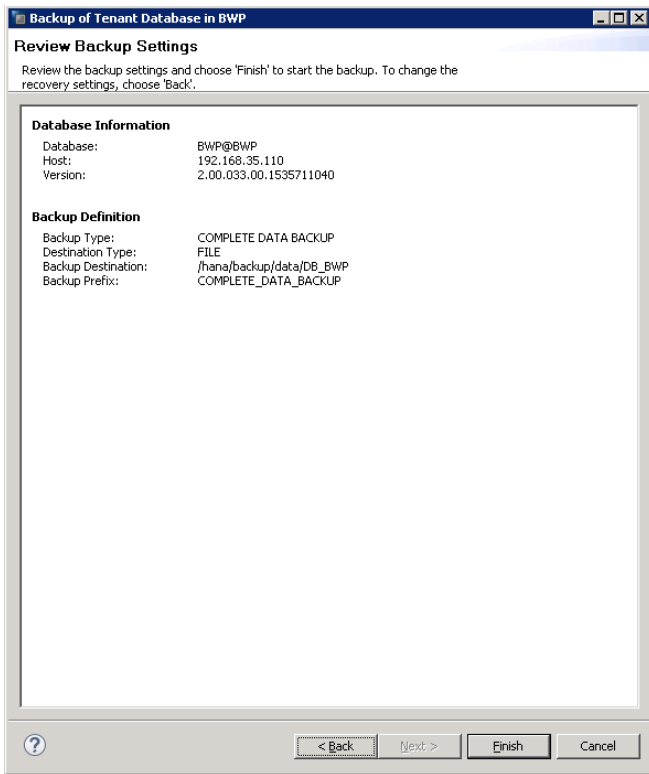
1. Select the SYSTEMDB entry and right-click it. Choose Backup and Recovery > Back Up Tenant Database.
2. In the new window that appears, select the database that you want to back up and click Next.



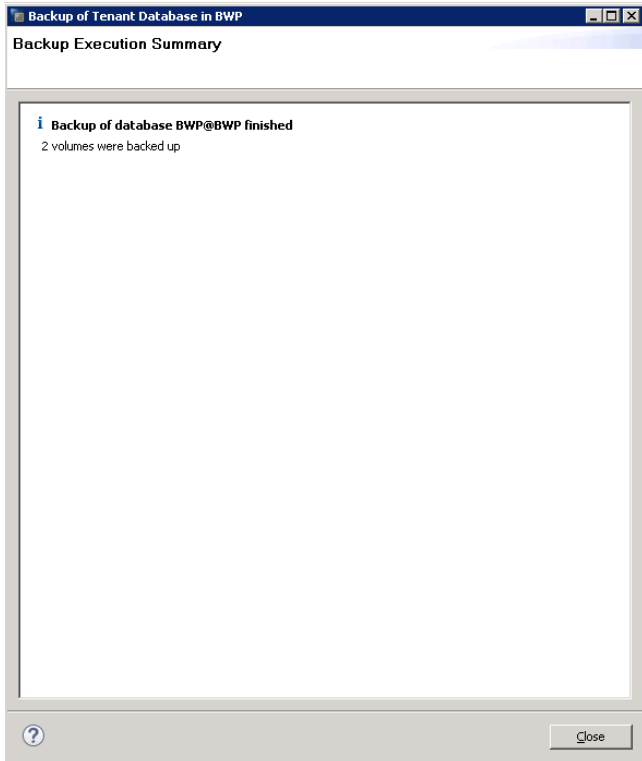
3. In the Backup Type field, select Complete Data Backup, and in the Destination Type field, select File.
4. The Backup Destination field should already point to /hana/backup/data. Click Next.



5. Review the backup job definition and click Finish.



6. After the backup job is finished, click Close.



7. Check the file system.

```
/hana/backup # ls -lR *  
data:
```

```
total 8
```

```
drwxr-x--x 2 bwpadm sapsys 4096 Feb 12 10:40 DB_BWP
drwxr-x--x 2 bwpadm sapsys 4096 Feb 12 10:22 SYSTEMDB
```

```
data/DB_BWP:
```

```
total 1771380
```

```
-rw-r----- 1 bwpadm sapsys      155648 Feb 12 10:40 COMPLETE_DATA_BACKUP_databackup_0_1
-rw-r----- 1 bwpadm sapsys    83894272 Feb 12 10:40 COMPLETE_DATA_BACKUP_databackup_2_1
-rw-r----- 1 bwpadm sapsys 1728061440 Feb 12 10:40 COMPLETE_DATA_BACKUP_databackup_3_1
```

```
data/SYSTEMDB:
```

```
total 1558168
```

```
-rw-r----- 1 bwpadm sapsys      155648 Feb 12 10:22 COMPLETE_DATA_BACKUP_databackup_0_1
-rw-r----- 1 bwpadm sapsys 1593843712 Feb 12 10:22 COMPLETE_DATA_BACKUP_databackup_1_1
```

```
log:
```

```
total 8
```

```
drwxr-x--x 2 bwpadm sapsys 4096 Feb 11 16:52 DB_BWP
drwxr-x--x 2 bwpadm sapsys 4096 Feb 11 16:48 SYSTEMDB
```

```
log/DB_BWP:
```

```
total 0
```

```
log/SYSTEMDB:
```

```
total 0
```

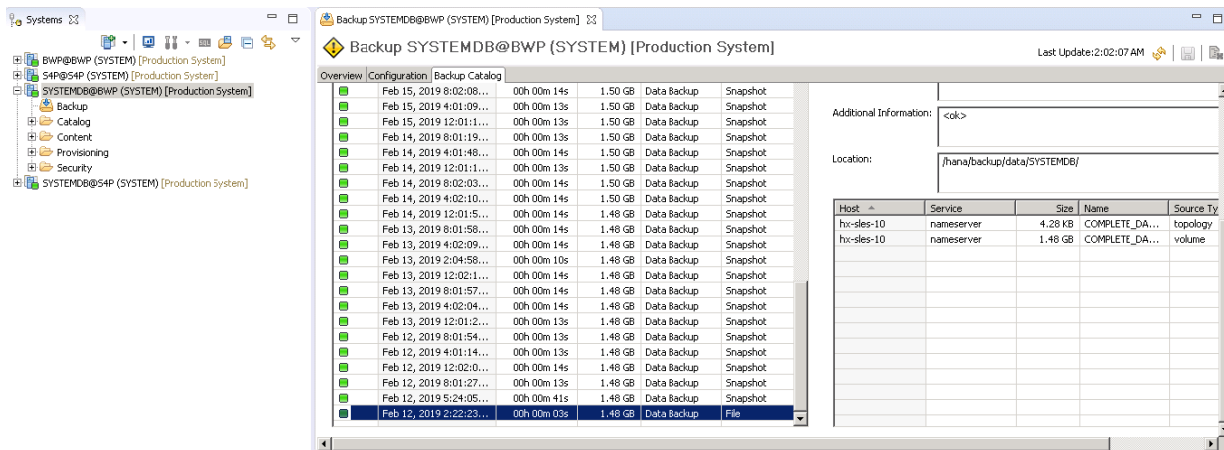
```
/hana/backup #
```

To schedule a backup-to-file operation, use the SAP documentation related to the applications you are using.

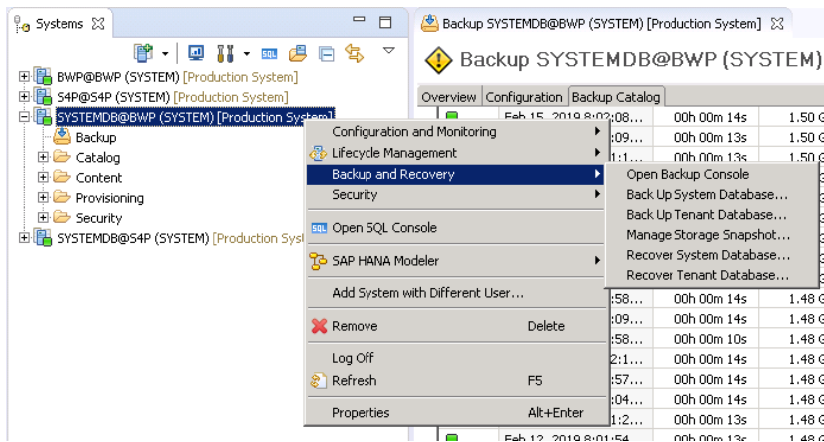
Recover a database from file

You can recover the existing database from a file or install a new database and recover the data from a file. The process for both operations is basically the same. Refer to the recovery-related documentation from SAP at <http://help.sap.com/hana>.

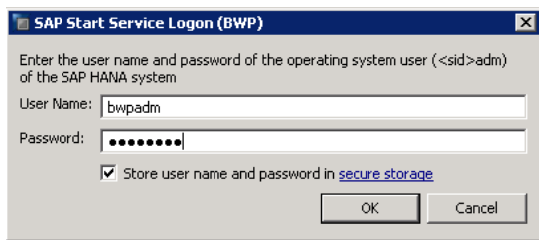
1. Open SAP HANA Studio. Select the database that you want to recover and open the backup catalog.




2. Right-click the database that you want to recover and choose Backup and Recovery. Then choose either Recover System Database or Recover Tenant Database.



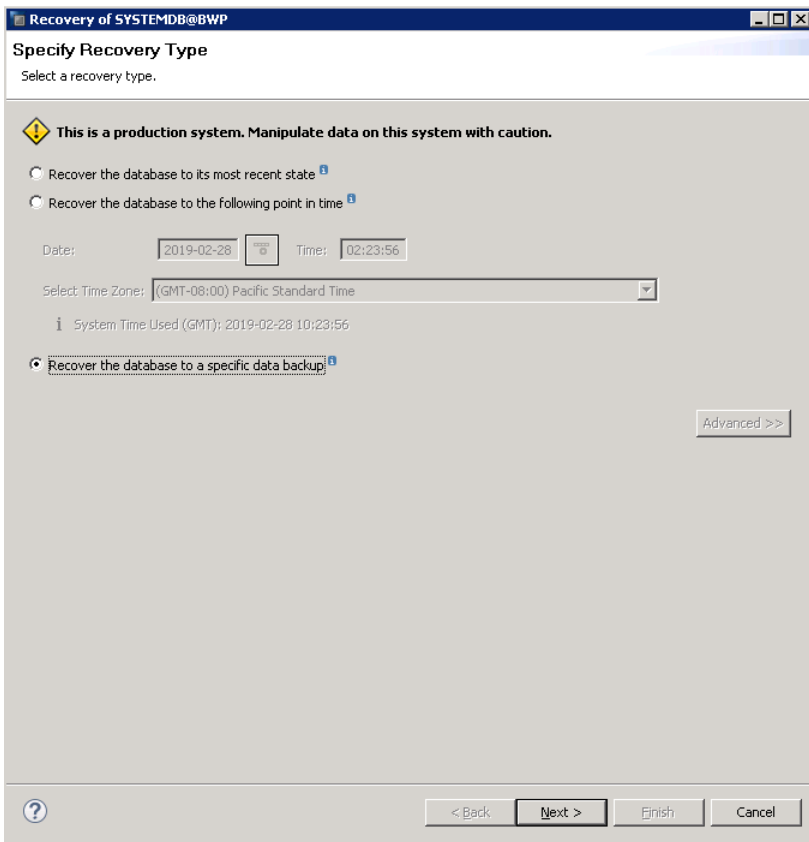
3. Enter an operating system user and password with the required permissions: for instance, enter <sid>adm. Then click OK.




4. If the system is running, click OK to confirm that the system can be shut down.



5. Select the recovery type that best meets your needs and click Next.

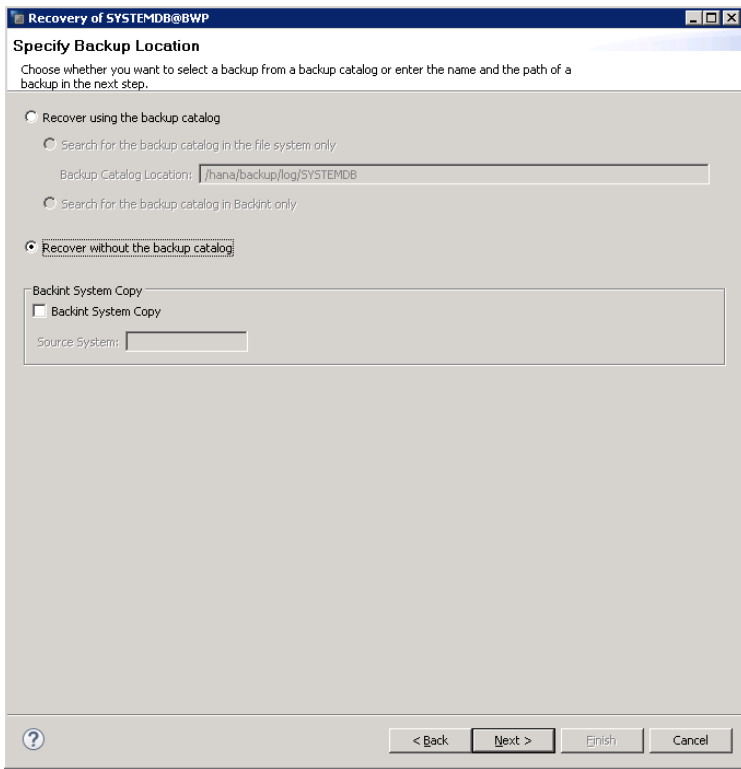


6. Click OK to confirm that you want to recover the data.



7. Specify whether you want to select the backup location from a backup catalog, on a file system or through Backint, or whether you want to perform recovery without a backup catalog. For details about the options, see the documentation from SAP.

Select the option that best meets your needs and click Next.



Recovery of SYSTEMDB@BWP

Specify Backup Location

Choose whether you want to select a backup from a backup catalog or enter the name and the path of a backup in the next step.

Recover using the backup catalog

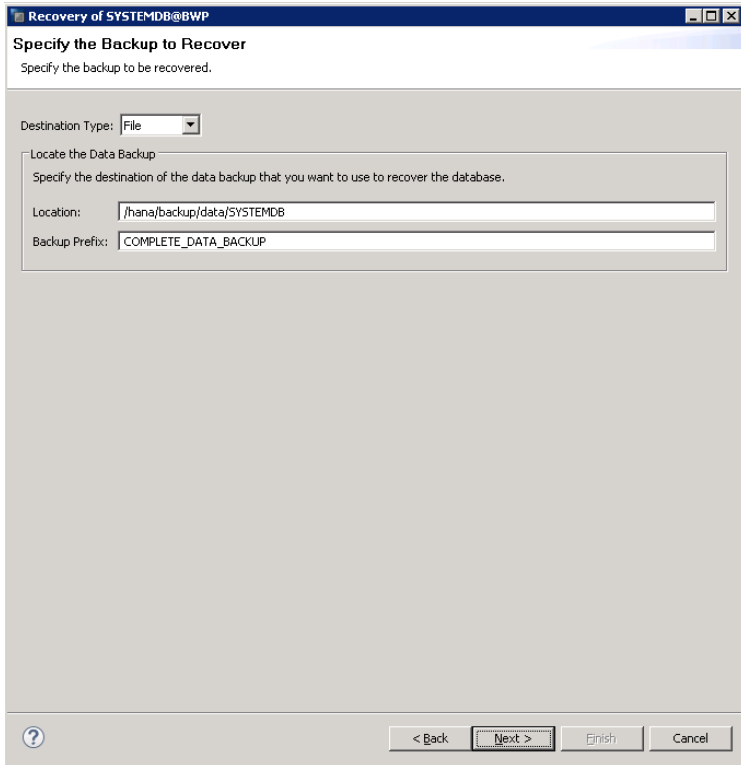
- Search for the backup catalog in the file system only
Backup Catalog Location: /hana/backup/log/SYSTEMDB
- Search for the backup catalog in Backint only

Recover without the backup catalog

- Backint System Copy
Source System: _____

Buttons: < Back, Next >, Finish, Cancel

8. In the Destination Type field, choose File, and specify the location and prefix of the backup files. Click Next



Recovery of SYSTEMDB@BWP

Specify the Backup to Recover

Specify the backup to be recovered.

Destination Type: File

Locate the Data Backup

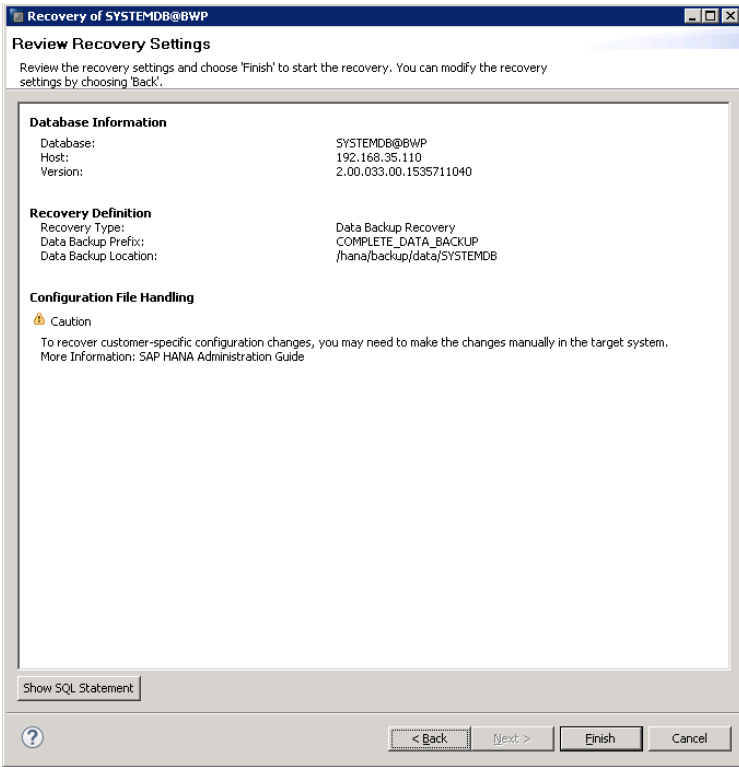
Specify the destination of the data backup that you want to use to recover the database.

Location: /hana/backup/data/SYSTEMDB


Backup Prefix: COMPLETE_DATA_BACKUP

Buttons: < Back, Next >, Finish, Cancel

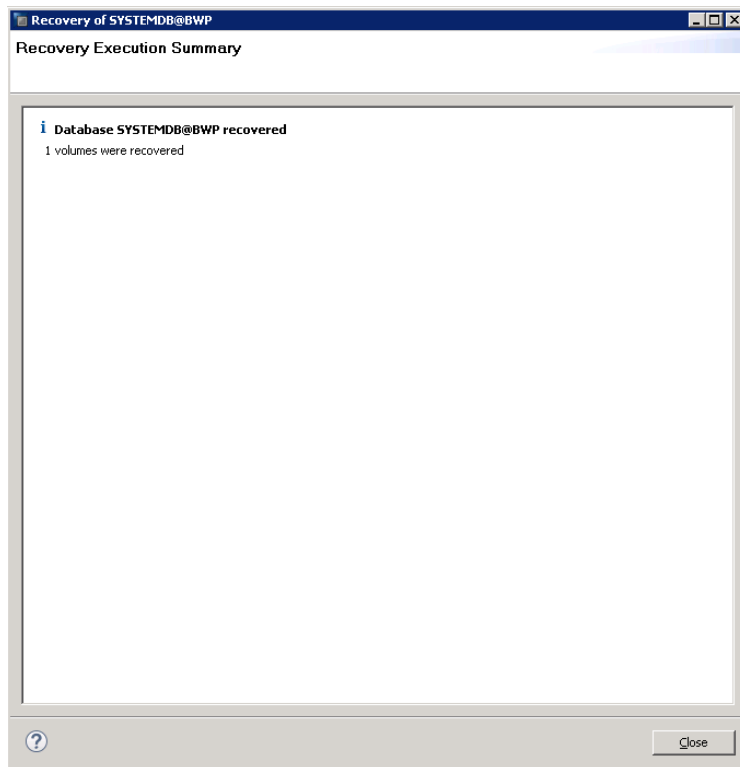
9. Review the information and click Finish.



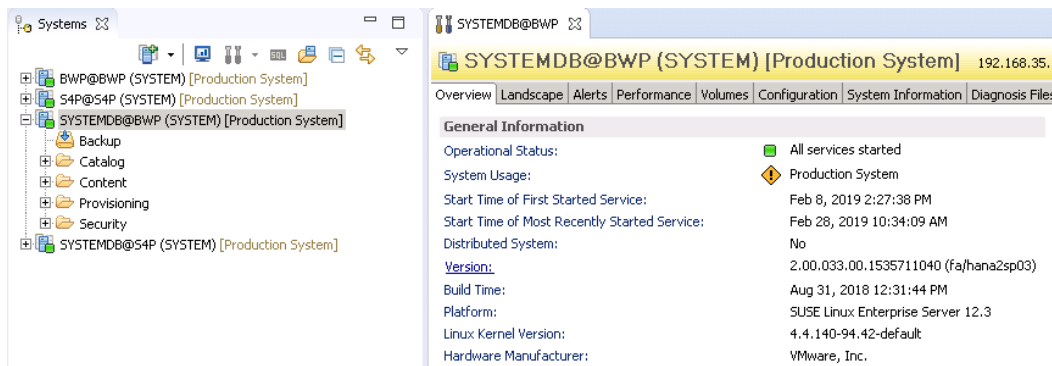
10. To confirm that you want to recover the database, click Yes.



11. After the recovery process is finished, click Close.



The database will start automatically after recovery.



Configuring Veeam Backup and Replication managed snapshots with pre- and post-snapshot scripts

Veeam Backup and Replication integrates into the Cisco HyperFlex HX Data Platform and can run commands before taking a snapshot and after the snapshot is finished. Veeam provides pre-freeze and post-thaw scripts for SAP HANA. These are located on GitHub under MIT license.

To configure this backup type, you perform the following high-level steps:

1. Run the pre-freeze script on the SAP HANA server.
2. Create a snapshot on the Cisco HyperFlex system for the virtual machine.
3. Run the post-thaw script on the SAP HANA server.
4. Back up the data from the Cisco HyperFlex system to the Veeam backup repository.

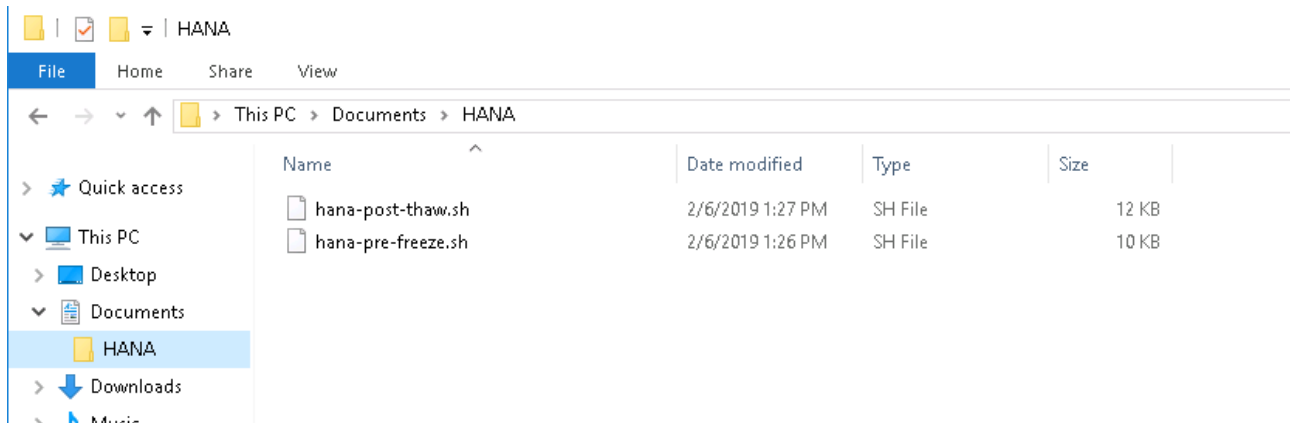
5. Delete the snapshot on the Cisco HyperFlex system.

To restore the system from the snapshot, you need to restore the data from the Veeam backup repository first or use the instant virtual machine recovery function in Veeam.

Download pre- and post-snapshot scripts

Download the pre-freeze and post-thaw scripts from the Veeam area of GitHub:

<https://github.com/VeeamHub/applications?files=1>. For example, for SAP HANA, download the hana-post-thaw and hana-pre-freeze commands.



Create OS user for backup on the SAP HANA server

You should create an OS user for backup operations on all SAP HANA systems. This user should have read access to the installed SAP HANA configuration files. In the following example, the user is named veeam and is part of the sapsys group.

Note: The connection from the Veeam Backup and Recovery (VBR) server to the SAP HANA server is established with this backup user, and the commands run on the SAP HANA server are elevated with a **sudo** command to root.

```
# useradd -d /home/veeam -G sapsys -m -u 2000 -p ***** veeam
```

Communication between the Veeam server on Windows and the SAP HANA system on Linux is performed through Secure Shell (SSH). As a best practice, you should use shared keys for authentication instead of a user and password combination. The easiest way to create a private and public key pair is with the **ssh-keygen** command in Linux.

1. In Linux, enter the following command:

```
# ssh-keygen
```

```
Generating public/private rsa key pair.
```

```
Enter file in which to save the key (/root/.ssh/id_rsa): /tmp/id_rsa
```

```
Enter passphrase (empty for no passphrase): *****
```

```
Enter same passphrase again: *****
```

```
Your identification has been saved in /tmp/id_rsa.
```

```
Your public key has been saved in /tmp/id_rsa.pub.
```

```
The key fingerprint is:
```

```
SHA256:6ABAonu72Pcn1l/t0CZTQTi2ZDHISVoOVadnrCpyS5k root@hx-sles-10
```

```
The key's randomart image is:
```

```
+---[RSA 2048]-----+
```

```
|+.      .++++OO |
|o.      =+ *= . |
|. .      .+.O+ .|
|. .      .o o |
|. . . . S   o |
|. . o     . B |
|. . . . E + |
| o .. o o = B |
|. o. o.o . . . |
```

```
+----- [SHA256] -----+
```

```
#
```

```
#
```

```
# mkdir /home/veeam/.ssh
```

```
# chown veeam /home/veeam/.ssh
```

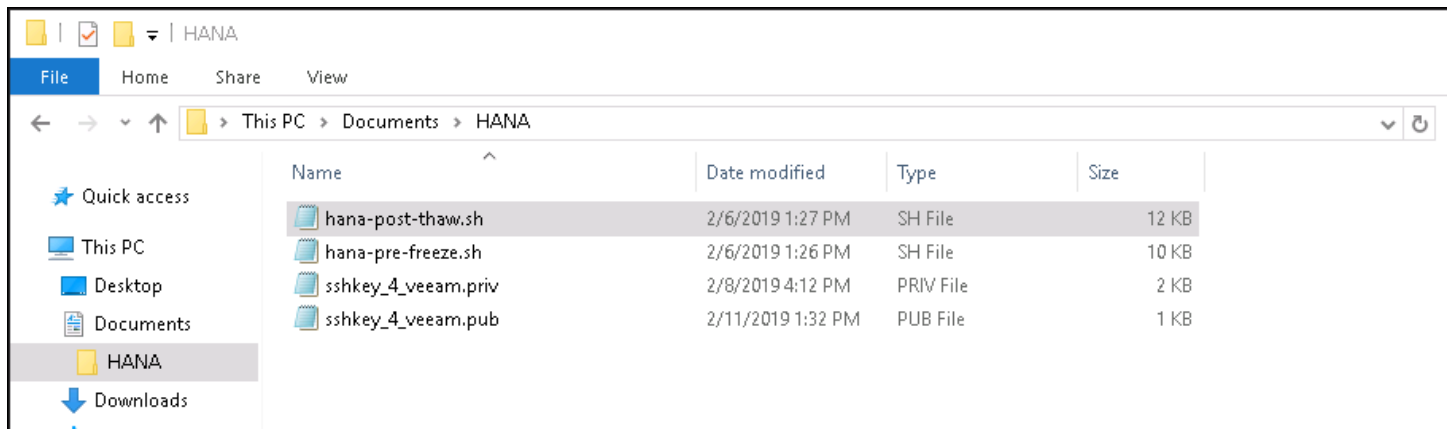
```
# chmod 700 /home/veeam/.ssh
```

```
# cat /tmp/id_rsa.pub > /home/veeam/.ssh/authorized_keys
```

```
#
```

2. Copy the private key file to the VBR server as **sshkey_4_veeam.priv**.

3. Copy the public key file to the VBR server as **sshkey_4_veeam.pub**.



Name	Date modified	Type	Size
hana-post-thaw.sh	2/6/2019 1:27 PM	SH File	12 KB
hana-pre-freeze.sh	2/6/2019 1:26 PM	SH File	10 KB
sshkey_4_veeam.priv	2/8/2019 4:12 PM	PRIV File	2 KB
sshkey_4_veeam.pub	2/11/2019 1:32 PM	PUB File	1 KB

4. Add the Veeam user to the sudoers file on all SAP HANA servers.

```
# visudo
```

```
...
```

```
veeam ALL=(ALL) NOPASSWD: ALL
```

```
....
```

```
#
```

Depending on your local security policies, you may want to lock this user from manual logon. Refer to the documentation for your Linux distribution for information about the process.

Create a secure HANA Database user store

The best way to grant user root access to the SAP HANA system to run the pre-freeze and post-thaw scripts is to maintain the SAP HANA secure user store.

Run this command as user root on the SAP HANA system:

```
# /hana/shared/<SID>/hdbclient/hdbuserstore set HDB<NR><hostname>:3<NR>13@SYSTEMDB
<USER><PASSWORD>
```

<SID> = SAP HANA SID.

<NR> = SAP HANA system number.

<HOSTNAME> = hostname of the server SAP HANA is running on.

<USER> = Database user with privileges to create SAP HANA internal snapshots.

<PASSWORD> = Password for the database user.

For example, enter the following command to grant access for user SYSTEM on the SAP HANA system with SID S4P and system number 10 running on host hx-sles-10:

```
# /hana/shared/S4P/hdbclient/hdbuserstore set HDB10 hx-sles-10:31013@SYSTEMDB SYSTEM
*****
```

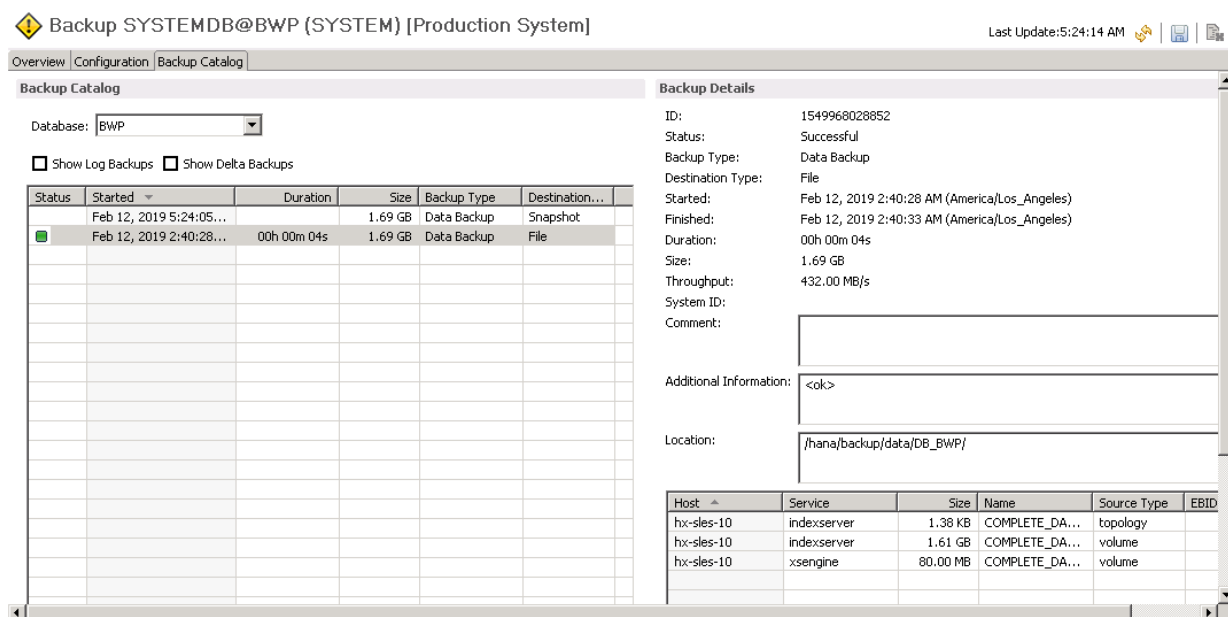
Test the pre- and post-snapshot scripts

Copy hana-pre-freeze.sh and hana-post-thaw.sh to the SAP HANA server and run the files manually to verify that they work. In some cases, you may need to run the **dos2unix** command to remove hidden characters in the scripts. Copy the tested new version back to the VBR server.

1. Run the pre-freeze script manually on the SAP HANA server.

```
veeam@hx-suse-10 :~ > sudo ./hana-pre-freeze.sh
```

In SAP HANA Studio, the snapshot is listed but without status and duration information.



Backup SYSTEMDB@BWP (SYSTEM) [Production System] Last Update: 5:24:14 AM

Overview Configuration Backup Catalog

Backup Catalog

Database: BWP

Show Log Backups Show Delta Backups

Status	Started	Duration	Size	Backup Type	Destination...
	Feb 12, 2019 5:24:05...		1.69 GB	Data Backup	Snapshot
■	Feb 12, 2019 2:40:28...	00h 00m 04s	1.69 GB	Data Backup	File

Backup Details

ID: 1549968028852
 Status: Successful
 Backup Type: Data Backup
 Destination Type: File
 Started: Feb 12, 2019 2:40:28 AM (America/Los_Angeles)
 Finished: Feb 12, 2019 2:40:33 AM (America/Los_Angeles)
 Duration: 00h 00m 04s
 Size: 1.69 GB
 Throughput: 432.00 MB/s
 System ID:
 Comment:

Additional Information: <ok>

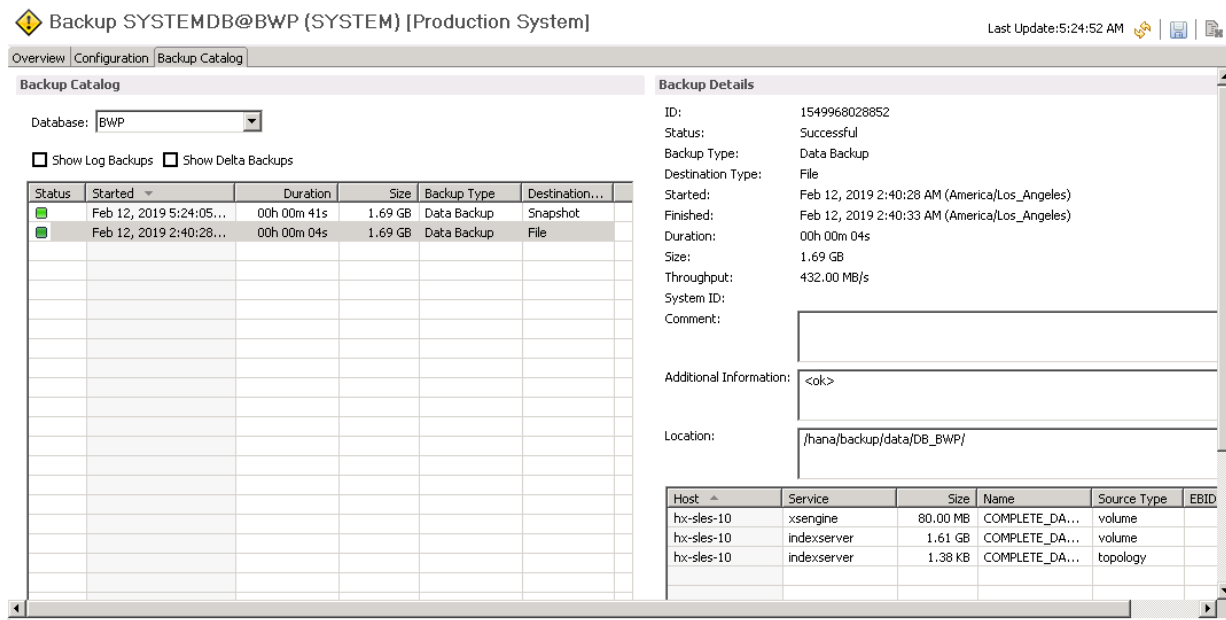
Location: /hana/backup/data/DB_BWP/

Host	Service	Size	Name	Source Type	EBID
hx-sles-10	indexserver	1.38 KB	COMPLETE_DA...	topology	
hx-sles-10	indexserver	1.61 GB	COMPLETE_DA...	volume	
hx-sles-10	xsengine	80.00 MB	COMPLETE_DA...	volume	

2. Run the post-thaw script manually on the SAP HANA server.

```
veeam@hx-suse-10 :~ > sudo ./hana-post-thaw.sh
```

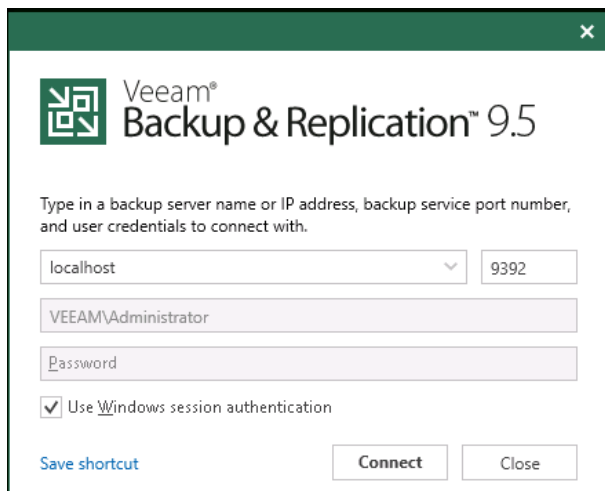
In SAP HANA Studio, the snapshot is now listed with status and duration information.



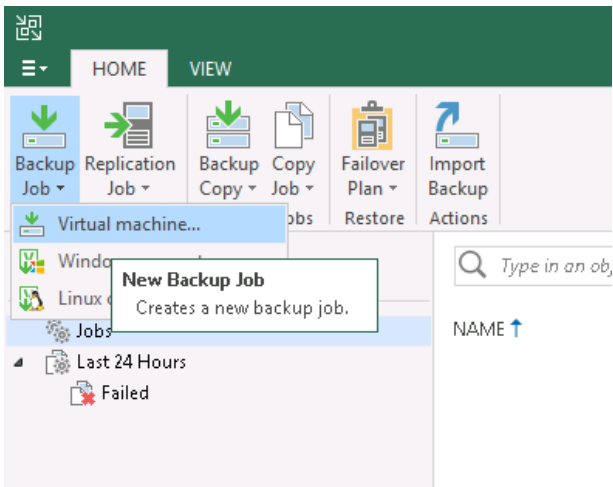
Configure a backup job in Veeam

Now that the SAP HANA server is prepared and the scripts have been tested, you can configure the backup job in Veeam.

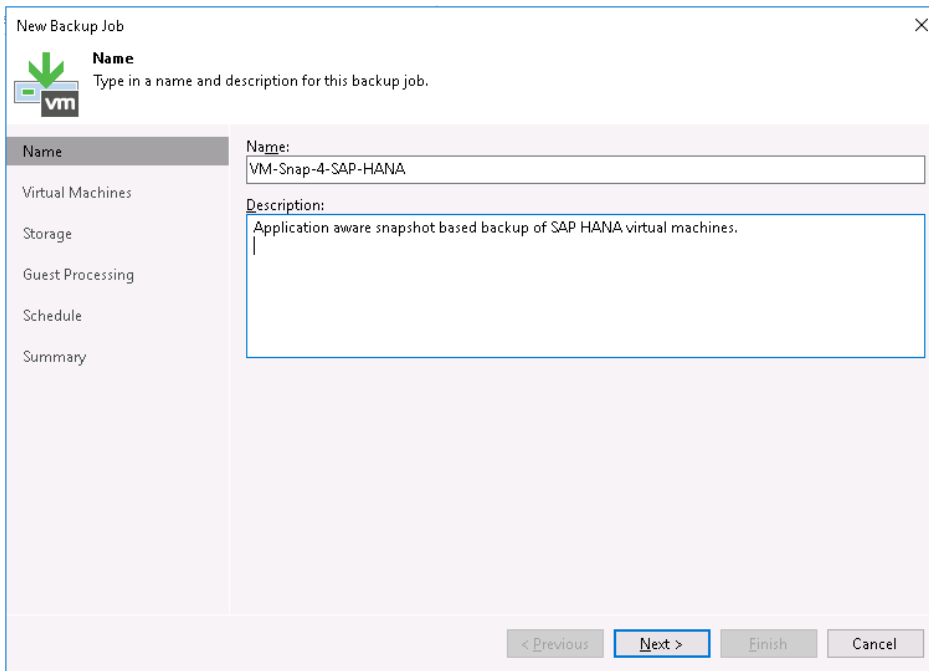
1. Open the Veeam Backup and Replication console.



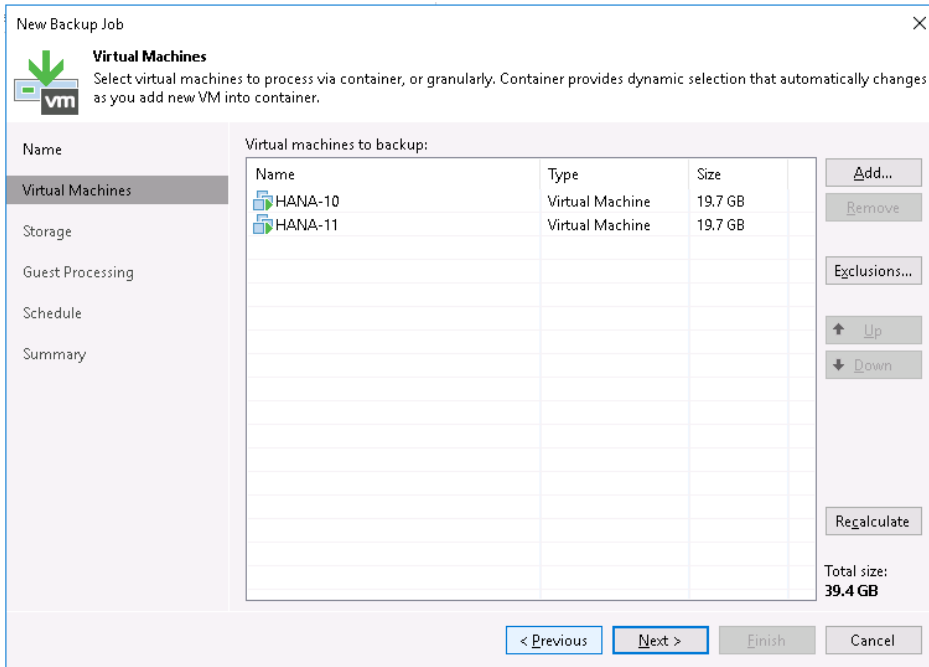
2. Click Backup Job and choose “Virtual machine.”



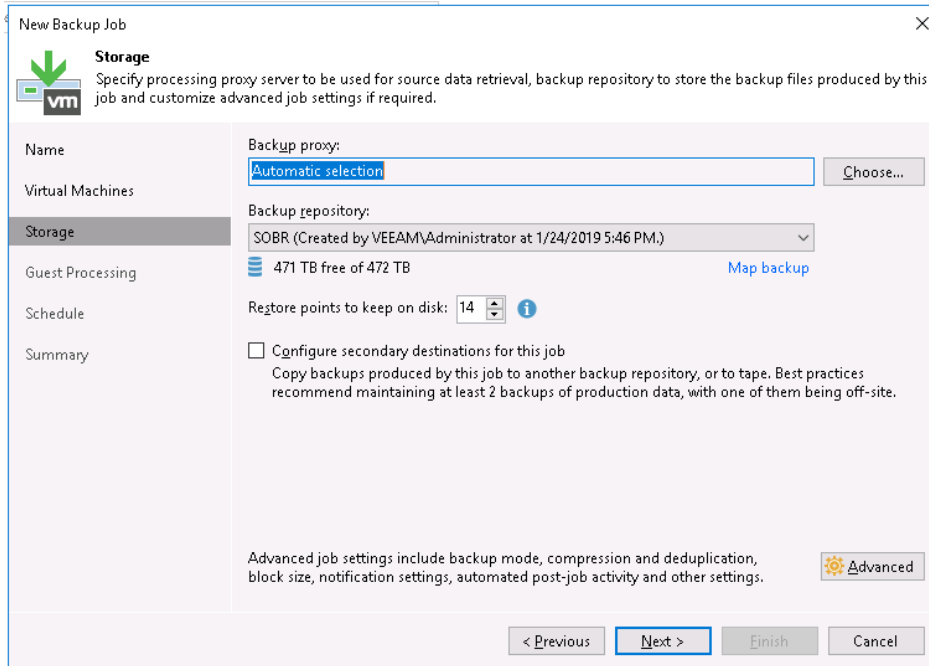
3. Enter a name for the backup job and click Next.



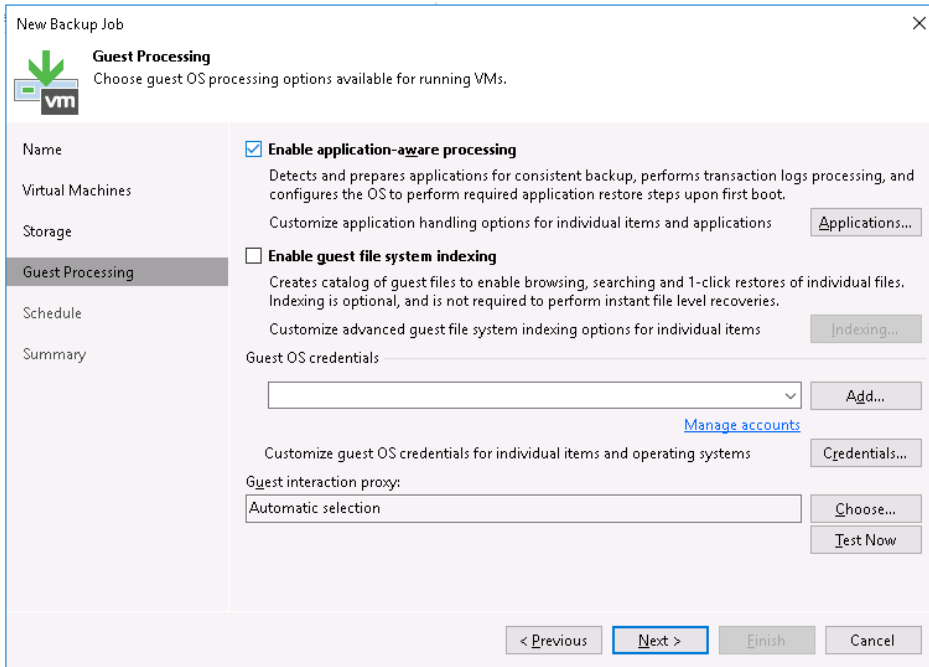
4. Add the virtual machines to be backed up by this job and click Next.



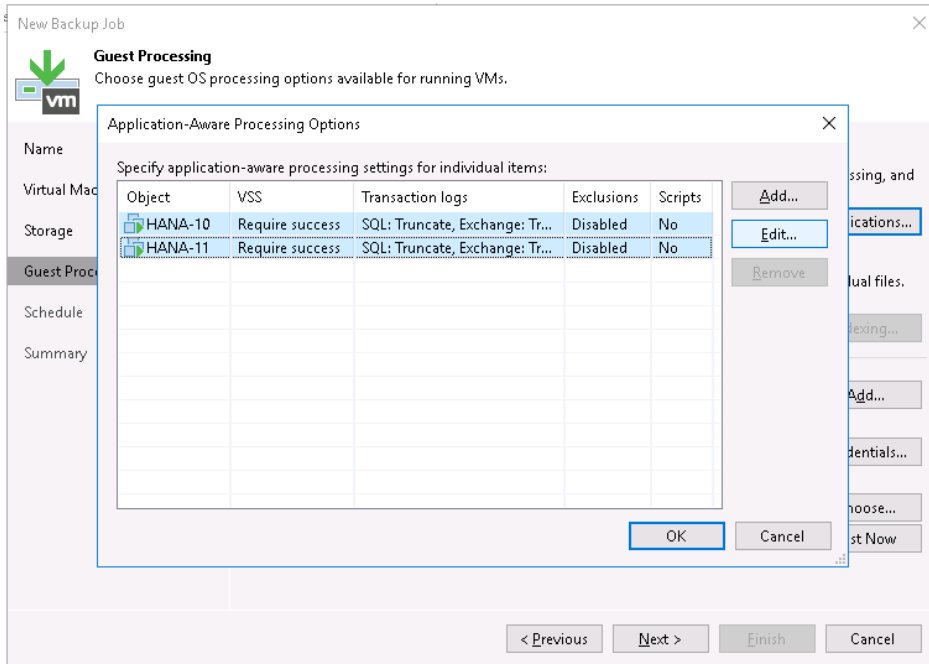
5. Select the backup proxy and backup repository for this backup job and click Next.



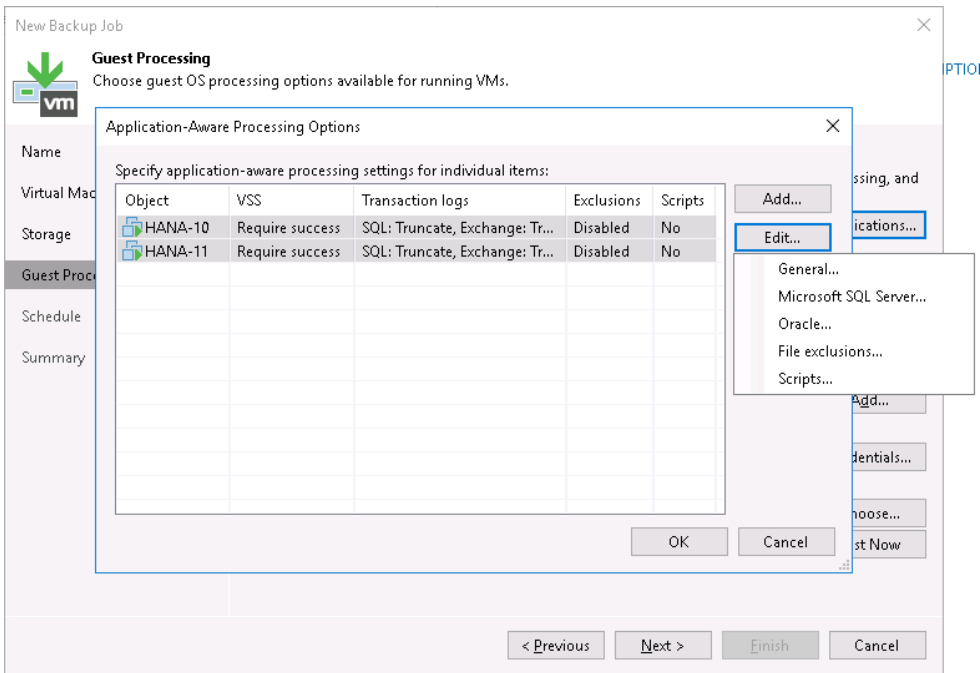
6. Select the checkbox for “Enable application-aware processing” and click Applications.



7. Select all objects and click Edit.

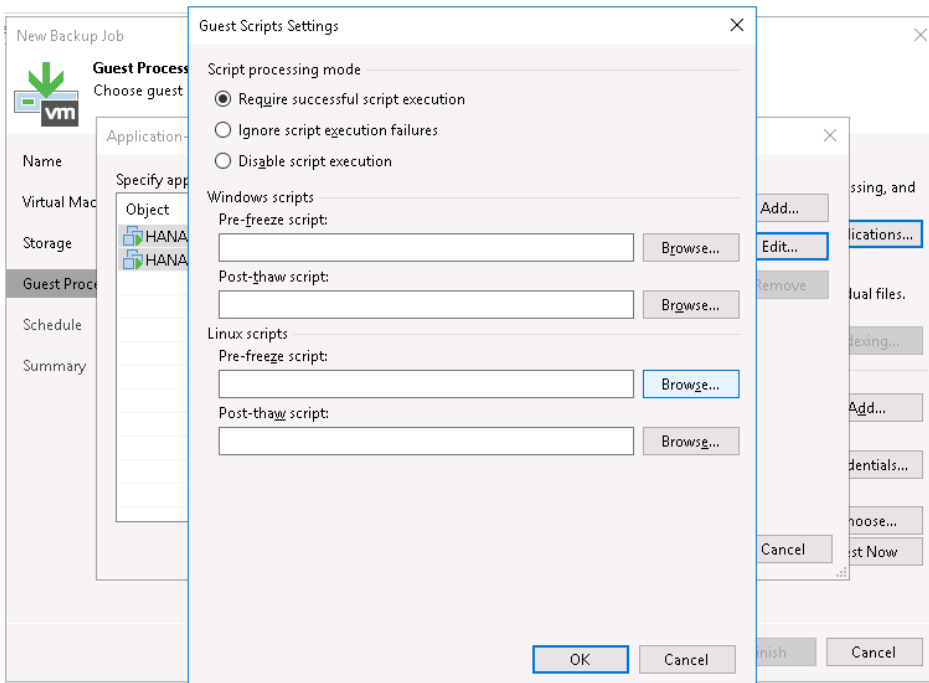


8. Choose Scripts.

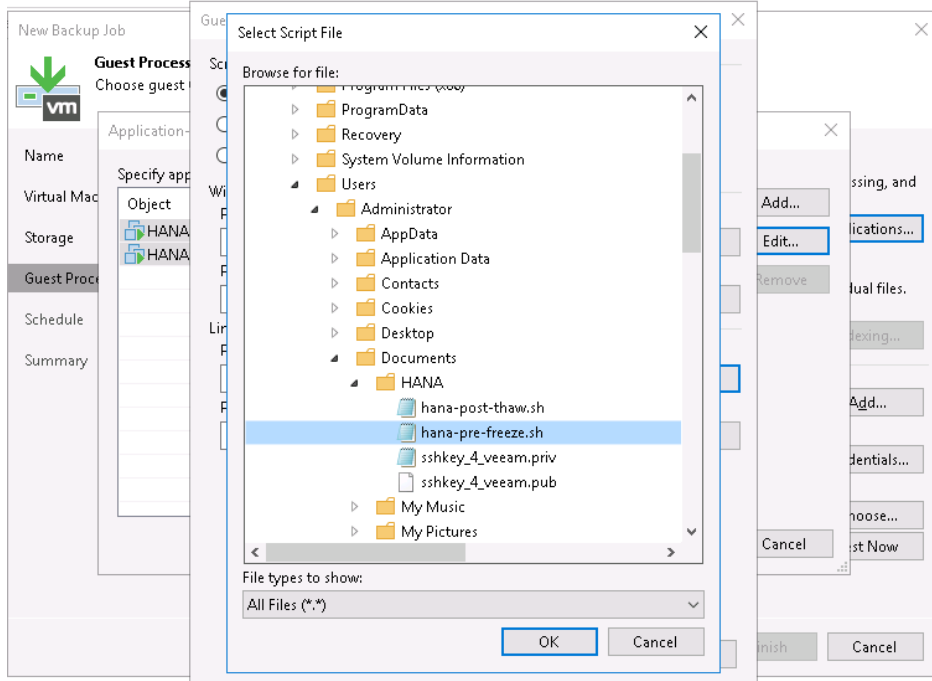


9. Select “Require successful script execution.”

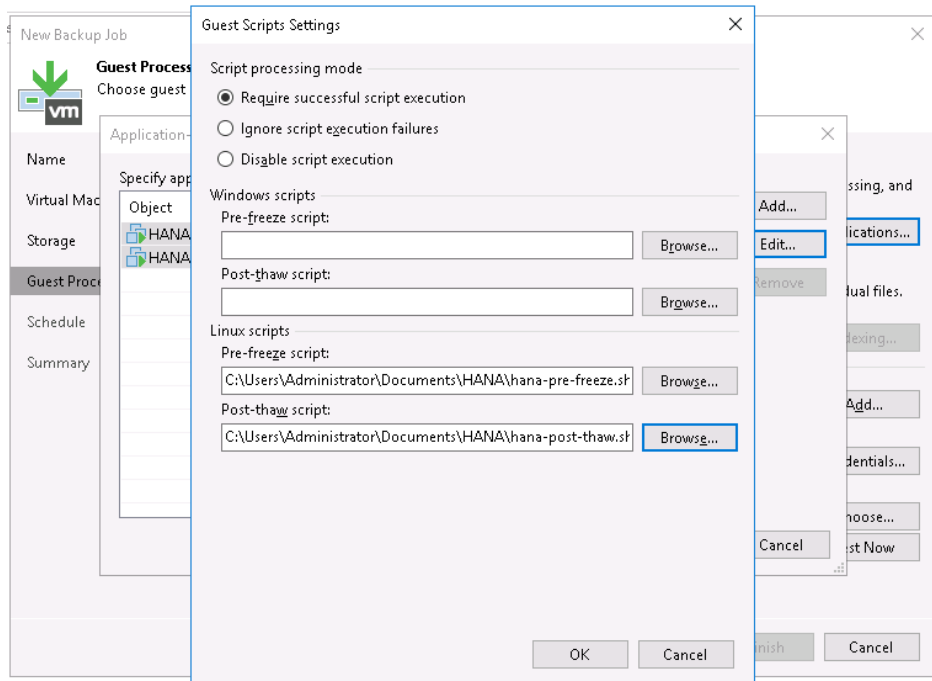
10. In the “Linux scripts” section, for “Pre-freeze script,” click Browse.



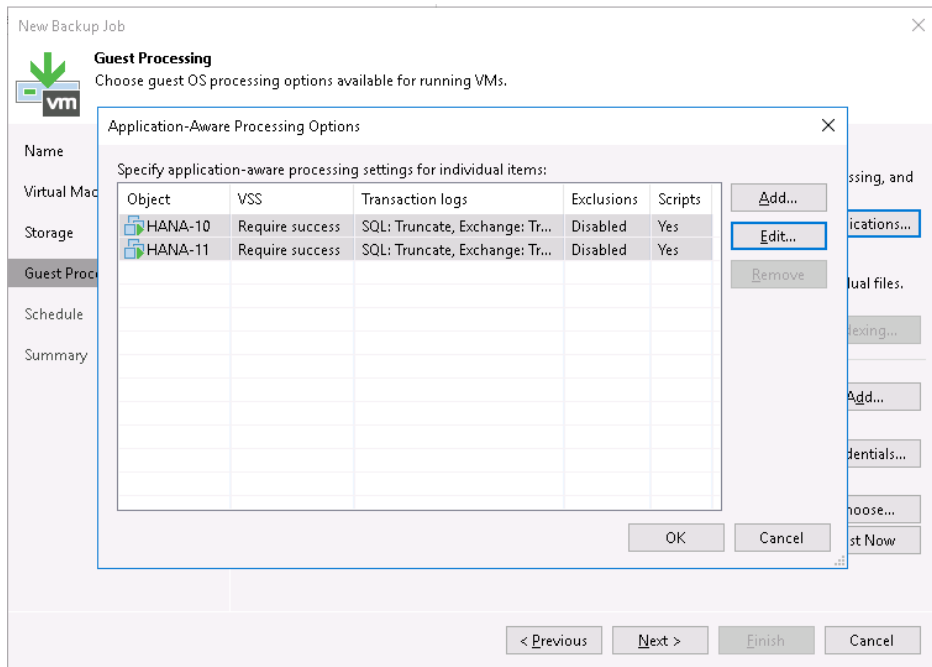
11. Select the hana-pre-freeze.sh script local to the VBR server and click OK.



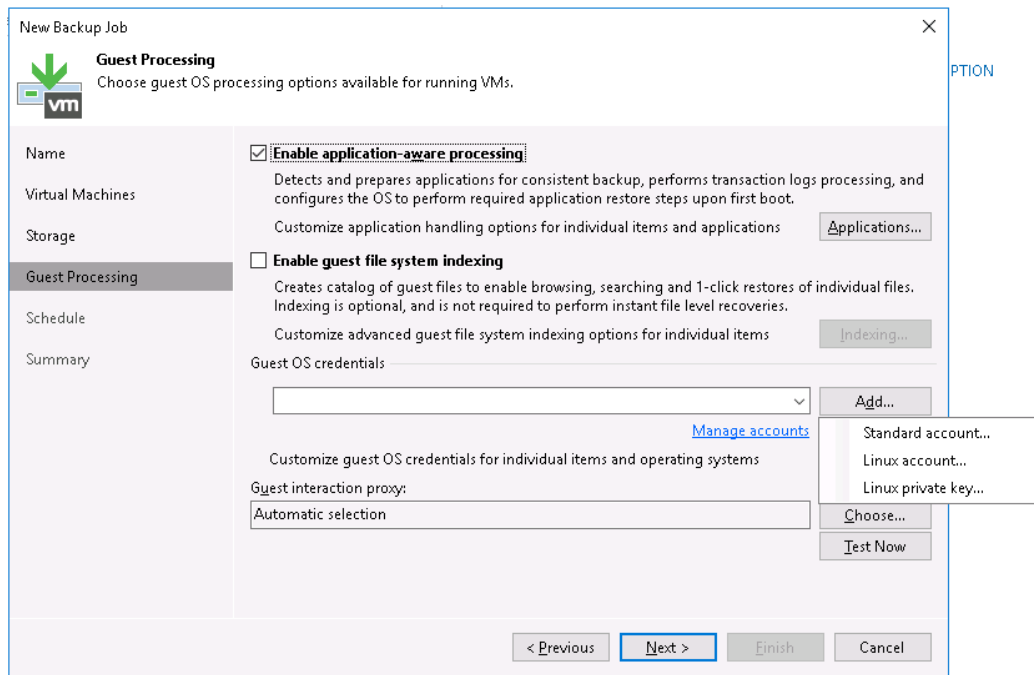
12. Repeat this step for the “Post-thaw script” field and click OK.



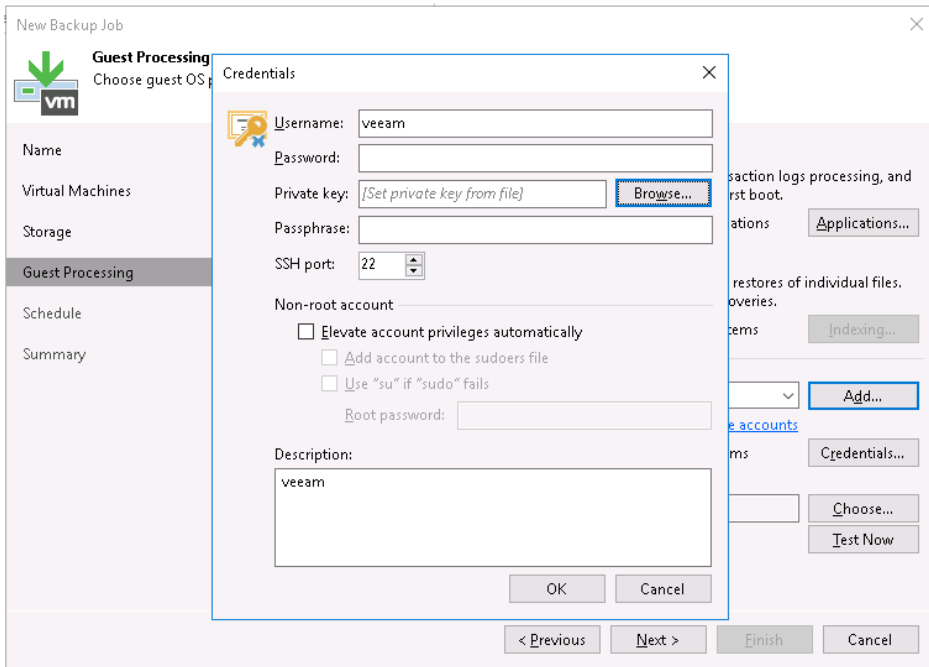
13. Click OK.



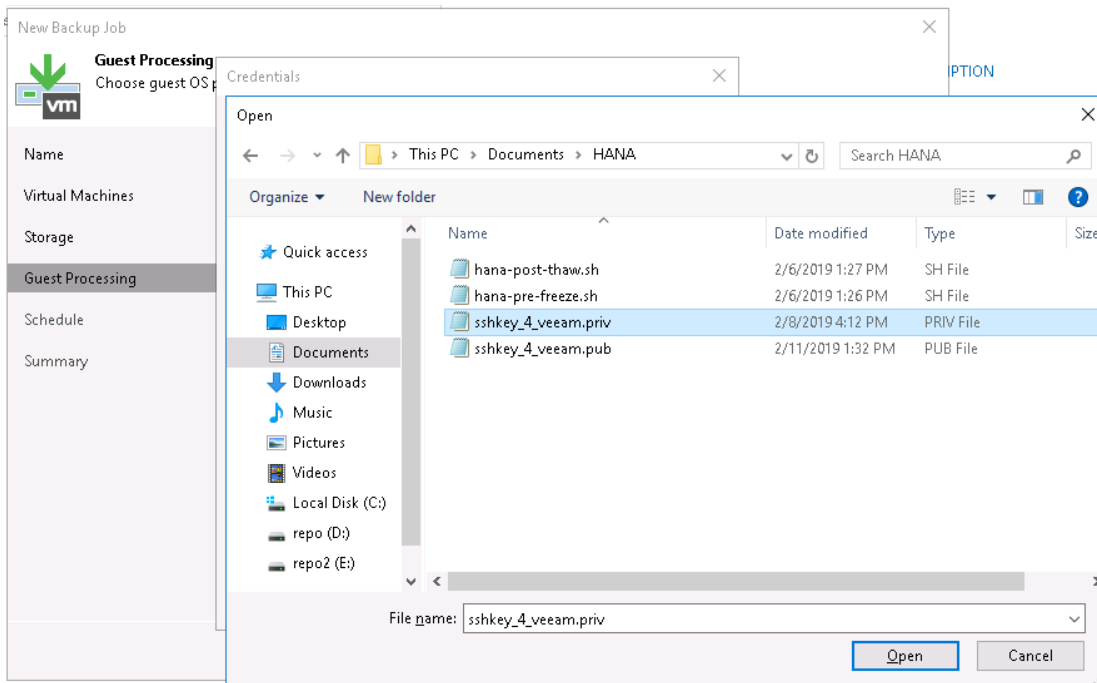
14. In the Guest OS credentials section, click Add and choose “Linux private key.”



15. In the Username field, enter the name of the backup user on the SAP HANA server (for example, **veeam**) and click Browse to find the private key file.

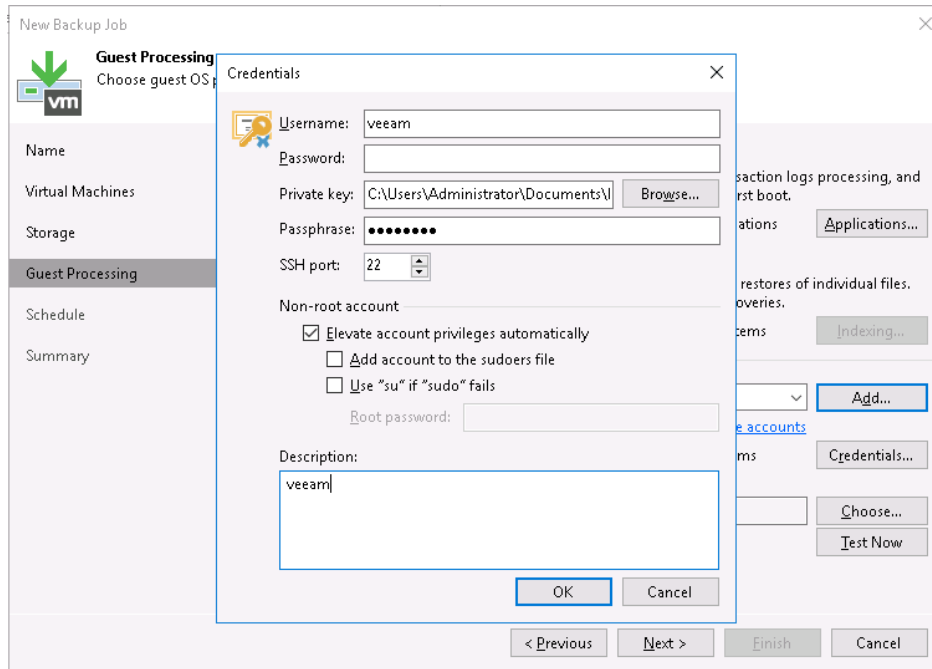


16. Select the private key file stored on the VBR server and click Open.

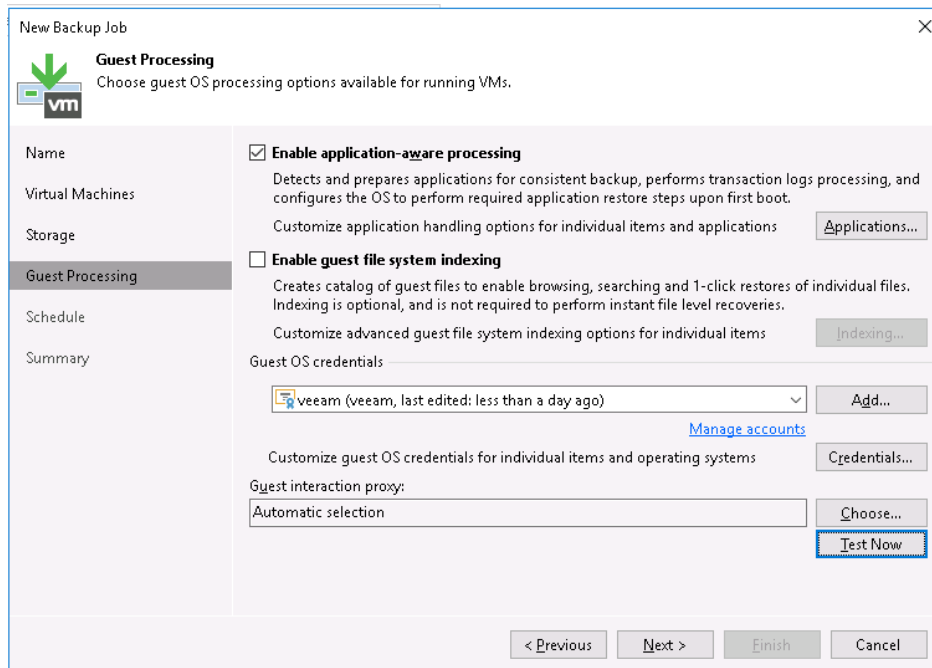


17. Enter the phrase defined at the time of the key creation.

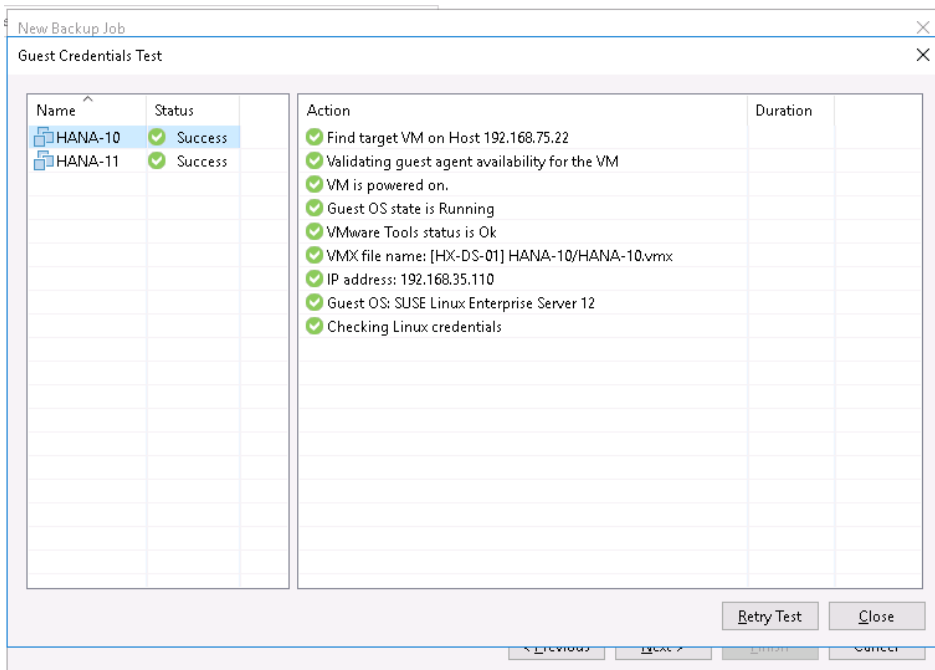
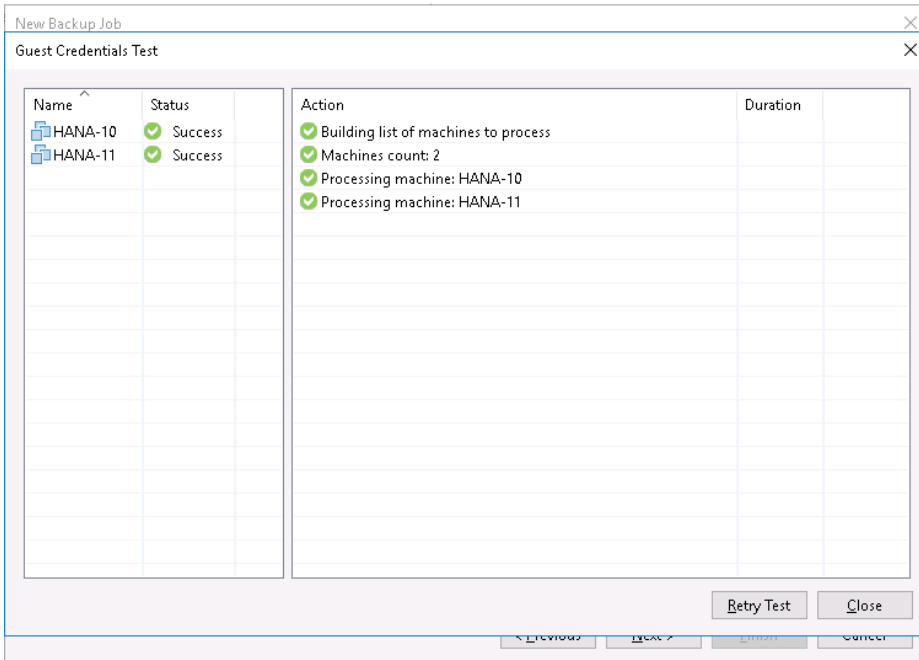
18. Select the check box for “Elevate account privileges automatically.” If you do not select this option, you will see an error message later in the creation process. Click OK.



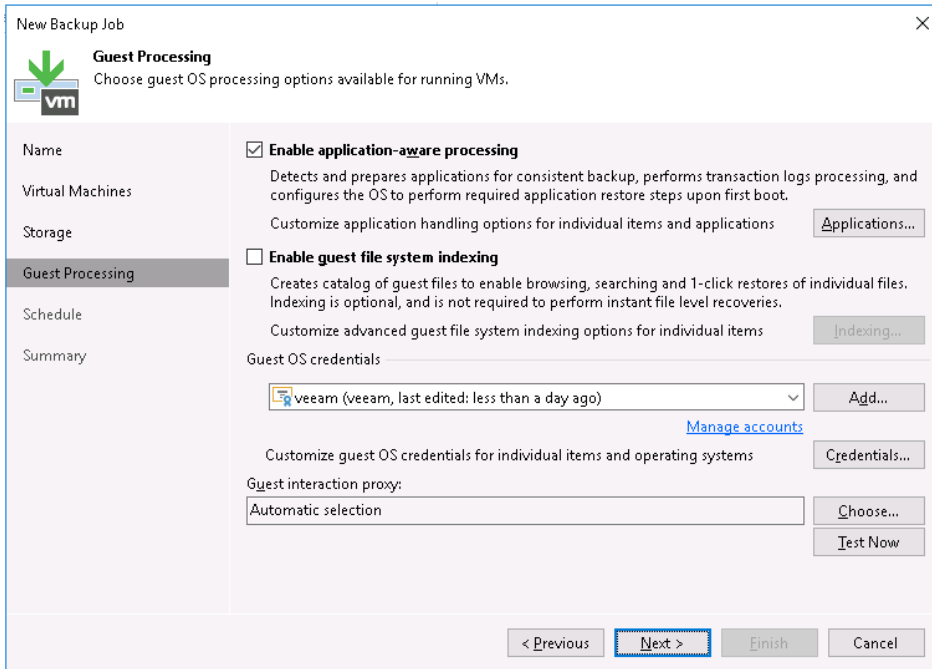
19. Click Test Now.



20. Veeam now tests the communication with the SAP HANA servers, but it does not run the scripts. Click Close after the test is finished.



21. Click Next.



New Backup Job

Guest Processing
Choose guest OS processing options available for running VMs.

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Enable application-aware processing
Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.
Customize application handling options for individual items and applications: [Applications...](#)

Enable guest file system indexing
Creates catalog of guest files to enable browsing, searching and 1-click restores of individual files. Indexing is optional, and is not required to perform instant file level recoveries.
Customize advanced guest file system indexing options for individual items: [Indexing...](#)

Guest OS credentials:

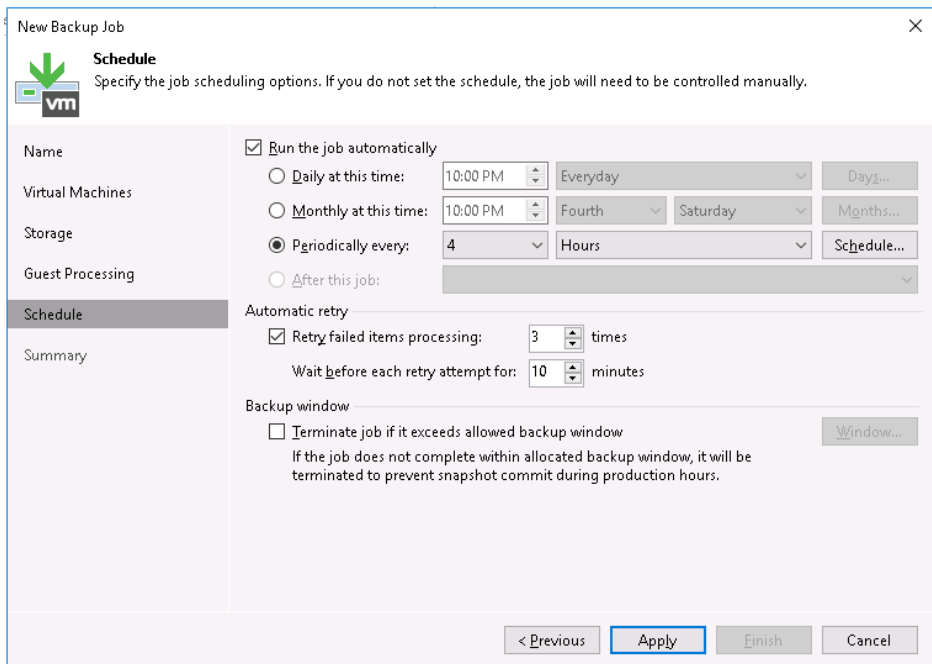
[veeam](#) (veeam, last edited: less than a day ago) [Add...](#)
[Manage accounts](#)

Customize guest OS credentials for individual items and operating systems: [Credentials...](#)

Guest interaction proxy:
Automatic selection [Choose...](#) [Test Now](#)

[< Previous](#) [Next >](#) [Finish](#) [Cancel](#)

22. Select the checkbox for “Run the job automatically.” Select “Periodically every” and choose the time that best meets your needs: for example, 4 hours. Click Apply.



New Backup Job

Schedule
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Run the job automatically

Daily at this time: 10:00 PM [Everyday](#) [Days...](#)

Monthly at this time: 10:00 PM [Fourth](#) [Saturday](#) [Months...](#)

Periodically every: 4 [Hours](#) [Schedule...](#)

After this job: [Schedule...](#)

Automatic retry

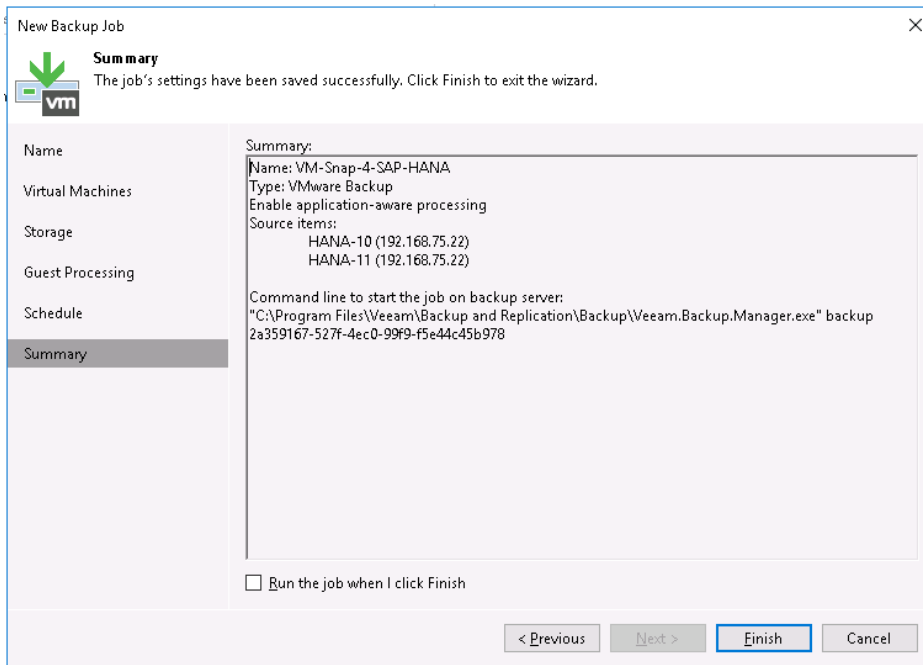
Retry failed items processing: 3 times
Wait before each retry attempt for: 10 minutes

Backup window

Terminate job if it exceeds allowed backup window
If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours. [Window...](#)

[< Previous](#) [Apply](#) [Finish](#) [Cancel](#)

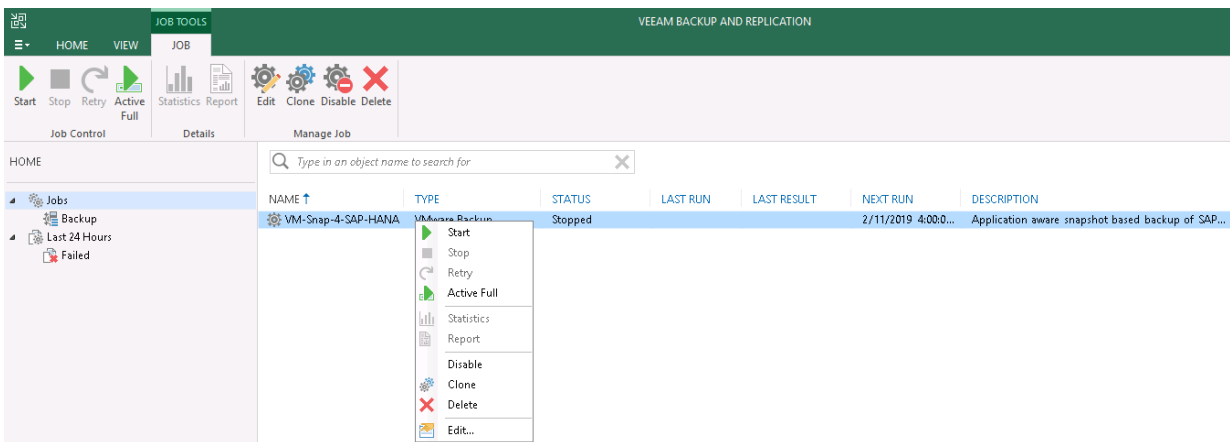
23. Click Finish.



Run a manual backup job

Now manually run a backup job.

1. Right-click the backup job you created and choose Start.



The job details are shown in the lower section of the Veeam console window.

2.

SUMMARY		DATA		STATUS		THROUGHPUT (LAST 5 MIN)					
Duration:	01:15	Processed:	0.0 B (0%)	Success:	0						
Processing rate:	0 KB/s	Read:	0.0 B	Warnings:	0						
Bottleneck:	Detecting	Transferred:	0.0 B	Errors:	0						

NAME	STATUS	ACTION	DURATI...
HANA-10	0%	Job started at 2/11/2019 2:55:01 PM	
HANA-11	0%	Building list of machines to process	00:04
		VM size: 1.3 TB (39.4 GB used)	
		Changed block tracking is enabled	
		Processing HANA-10	00:52
		Processing HANA-11	00:07
		All VMs have been queued for processing	00:00

3. Click a virtual machine name to see the job details for this virtual machine.

SUMMARY		DATA		STATUS		THROUGHPUT (LAST 5 MIN)					
Duration:	01:56	Processed:	0.0 B (0%)	Success:	0						
Processing rate:	0 KB/s	Read:	0.0 B	Warnings:	0						
Bottleneck:	Detecting	Transferred:	0.0 B	Errors:	0						

NAME	STATUS	ACTION	DURATI...
HANA-10	0%	Queued for processing at 2/11/2019 3:58:27 PM	
HANA-11	0%	Required backup infrastructure resources have been assigned	
		Using Backup Repository 1 scale-out repository extent	
		VM processing started at 2/11/2019 3:58:32 PM	
		VM size: 690.0 GB (29.7 GB used)	
		Storage initialized	00:44
		Getting VM info from vSphere	00:05
		Running pre-freeze script in the guest OS	00:06
		Creating HyperFlex snapshot	00:14
		Running post-thaw script in the guest OS	00:00

SUMMARY		DATA		STATUS		THROUGHPUT (LAST 5 MIN)					
Duration:	02:46	Processed:	64.4 GB (4%)	Success:	0						Speed: 1.2 GB/s
Processing rate:	2 GB/s	Read:	64.4 GB	Warnings:	0						
Bottleneck:	Source	Transferred:	13.5 GB (4.8x)	Errors:	0						

NAME	STATUS	ACTION	DURATI...
HANA-10	5%	Creating HyperFlex snapshot	00:14
HANA-11	3%	Running post-thaw script in the guest OS	00:00
		Saving [HX-DS-01] HANA-10/HANA-10.vmcx	00:00
		Saving [HX-DS-01] HANA-10/HANA-10.nvram	00:00
		Using backup proxy VMware Backup Proxy for disk Hard disk 1 [nfs]	00:00
		Using backup proxy VMware Backup Proxy for disk Hard disk 2 [nfs]	00:00
		Using backup proxy VMware Backup Proxy for disk Hard disk 3 [nfs]	00:00
		Hard disk 2 (130.0 GB) 16.2 GB read at 416 MB/s [CBT]	00:43
		Hard disk 1 (60.0 GB) 15.9 GB read at 408 MB/s [CBT]	00:42
		Hard disk 3 (500.0 GB) 15.6 GB read at 399 MB/s [CBT]	00:42

The backup file is already listed in the SAP HANA backup catalog as a snapshot-based backup. It is listed after the post-thaw script is finished. The data is transferred from the Cisco HyperFlex HX Data Platform to the Veeam backup repository after the entry in the SAP HANA backup catalog is complete.

The scripts provided by Veeam work for all databases on the virtual machine and work with multiple system IDs (SIDs) as well as multitenant databases.

Backup SYSTEMDB@BWP (SYSTEM) [Production System] Last Update: 7:53:07 AM

Overview | Configuration | Backup Catalog

Backup Catalog

Database: SYSTEMDB

Show Log Backups Show Delta Backups

Status	Started	Duration	Size	Backup Type	Destination...
█	Feb 11, 2019 6:56:37...	00h 56m 10s	1.47 GB	Data Backup	Snapshot

Backup Details

ID: 1549896997279
 Status: Successful
 Backup Type: Data Backup
 Destination Type: Snapshot
 Started: Feb 11, 2019 6:56:37 AM (America/Los_Angeles)
 Finished: Feb 11, 2019 7:52:47 AM (America/Los_Angeles)
 Duration: 00h 56m 10s
 Size: 1.47 GB
 Throughput: n.a.
 System ID:
 Comment: Veeam Backup Pre-Freeze - Mon Feb 11 14:56:37 UTC 2019

Additional Information: <ok>

Location: /hana/data/mnt00001/

Host	Service	Size	Name	Source Type	EBID
hx-sles-10	nameserver	1.47 GB	hdb00001	volume	Veeam

Backup SYSTEMDB@BWP (SYSTEM) [Production System] Last Update: 8:04:53 AM

Overview | Configuration | Backup Catalog

Backup Catalog

Database: BWP

Show Log Backups Show Delta Backups

Status	Started	Duration	Size	Backup Type	Destination...
█	Feb 11, 2019 6:56:37...	00h 56m 10s	1.67 GB	Data Backup	Snapshot

Backup Details

ID: 1549896997279
 Status: Successful
 Backup Type: Data Backup
 Destination Type: Snapshot
 Started: Feb 11, 2019 6:56:37 AM (America/Los_Angeles)
 Finished: Feb 11, 2019 7:52:47 AM (America/Los_Angeles)
 Duration: 00h 56m 10s
 Size: 1.67 GB
 Throughput: n.a.
 System ID:
 Comment: Veeam Backup Pre-Freeze - Mon Feb 11 14:56:37 UTC 2019

Additional Information: <ok>

Location: /hana/data/mnt00001/

Host	Service	Size	Name	Source Type	EBID
hx-sles-10	xsengine	80.00 MB	hdb00002.00003	volume	Veeam
hx-sles-10	indexserver	1.59 GB	hdb00003.00003	volume	Veeam

NAME ↑	TYPE	STATUS	LAST RUN	LAST RESULT	NEXT RUN	DESCRIPTION
VM-Snap-4-SAP-HANA	VMware Backup	Stopped	19 minutes ago	Success	2/11/2019 8:00:0...	Application aware snapshot based backup of SAP...

SUMMARY	DATA	STATUS	THROUGHPUT (ALL TIME)
Duration: 17:17	Processed: 1.3 TB (100%)	Success: 2 ✓	
Processing rate: 2 GB/s	Read: 1.3 TB	Warnings: 0	
Bottleneck: Source	Transferred: 59.0 GB (23.4%)	Errors: 0	

NAME	STATUS	ACTION	DURATI...
HANA-10	Success	Using backup proxy VMware Backup Proxy for disk Hard disk 1 [nfs]	00:00
HANA-11	Success	Using backup proxy VMware Backup Proxy for disk Hard disk 2 [nfs]	00:00
		Using backup proxy VMware Backup Proxy for disk Hard disk 3 [nfs]	00:00
		Hard disk 2 (130.0 GB) 130.0 GB read at 489 MB/s [CBT]	04:33
		Hard disk 1 (60.0 GB) 60.0 GB read at 432 MB/s [CBT]	02:23
		Hard disk 3 (500.0 GB) 500.0 GB read at 607 MB/s [CBT]	14:04
		Removing HyperFlex snapshot	00:11
		Finalizing	00:00
		Busy: Source 96% > Proxy 26% > Network 17% > Target 1%	
		Primary bottleneck: Source	
		Network traffic verification detected no corrupted blocks	
		Processing finished at 2/11/2019 4:15:16 PM	

If a backup job, such as the transport of the data from the Cisco HyperFlex HX Data Platform to the Veeam backup repository, fails, the entry in the SAP HANA backup catalog will still exist. Remove snapshot entries from the catalog if the snapshot or the backup on the Veeam repository server does not exist to prevent information from being incorrectly read.

Restore a system from the Veeam backup repository

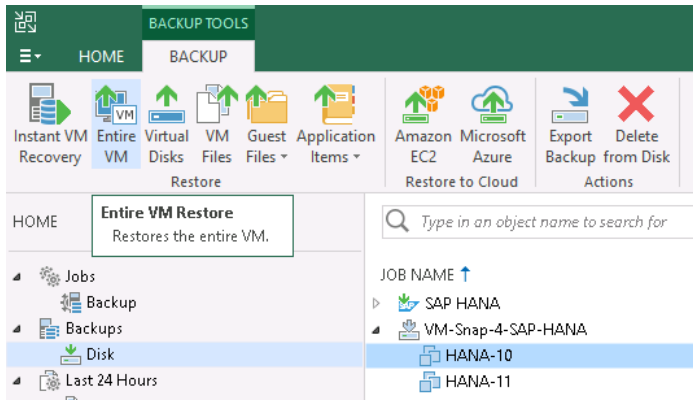
To restore the data from the Veeam backup repository, you initiate the restore process from the Veeam console.

1. Open the Veeam console and on the Backup tab navigate to Backups > Disk. In the right pane, select the virtual machine that you want to restore.

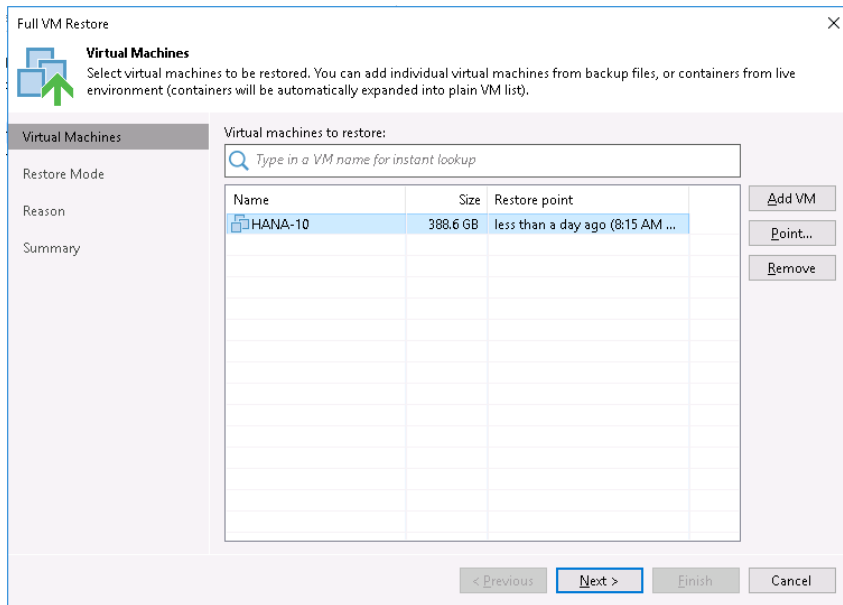
The screenshot shows the Veeam Backup and Replication console interface. The 'BACKUP TOOLS' menu is open, showing options like 'Instant VM Recovery', 'Entire VM', 'Virtual Disks', 'VM Files', 'Guest Files', 'Application Items', 'Amazon EC2', 'Microsoft Azure', 'Export Backup from Disk', and 'Delete'. The main pane shows a tree view on the left with 'Disk' selected under 'Backups'. The right pane displays a table of backup jobs:

JOB NAME ↑	CREATION TIME	RESTORE POINTS	REPOSITORY	PLATFORM
SAP HANA				
VM-Snap-4-SAP-HANA	2/23/2019 12:00 AM		SOBR	VMware
HANA-10	2/28/2019 8:15 AM	33		
HANA-11	2/28/2019 8:15 AM	33		

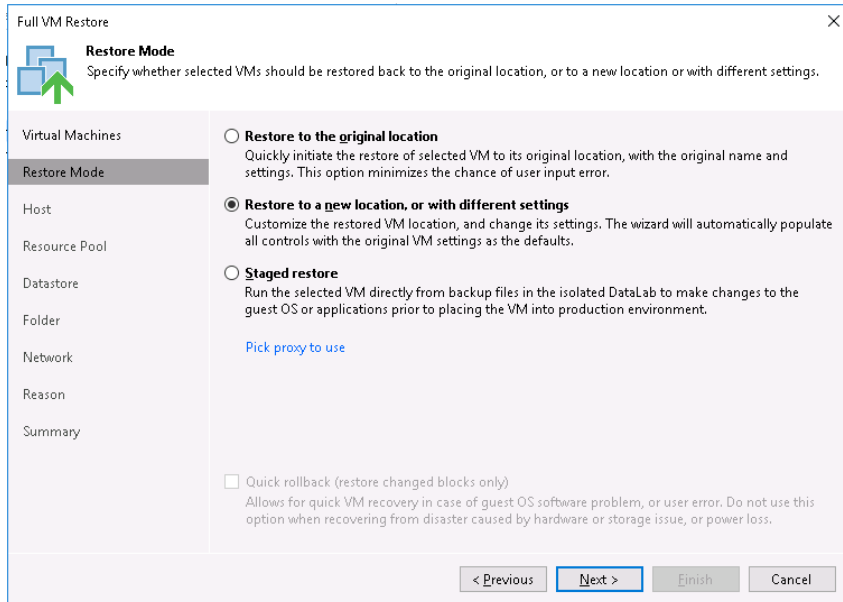
2. Select Entire VM. This option will restore all files before the virtual machine can be started.



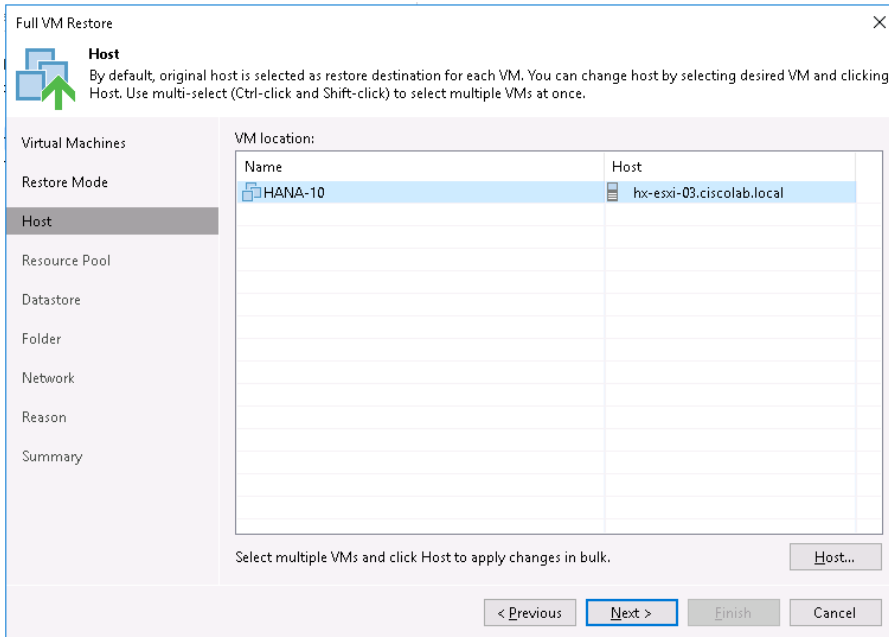
3. Click Next.



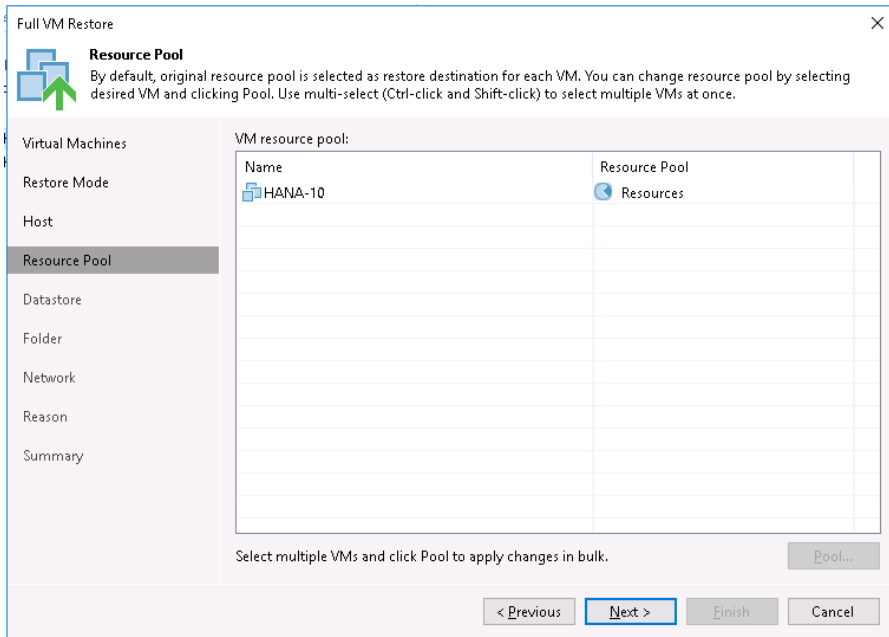
4. For Restore Mode, select “Restore to a new location, or with different settings” and click Next.



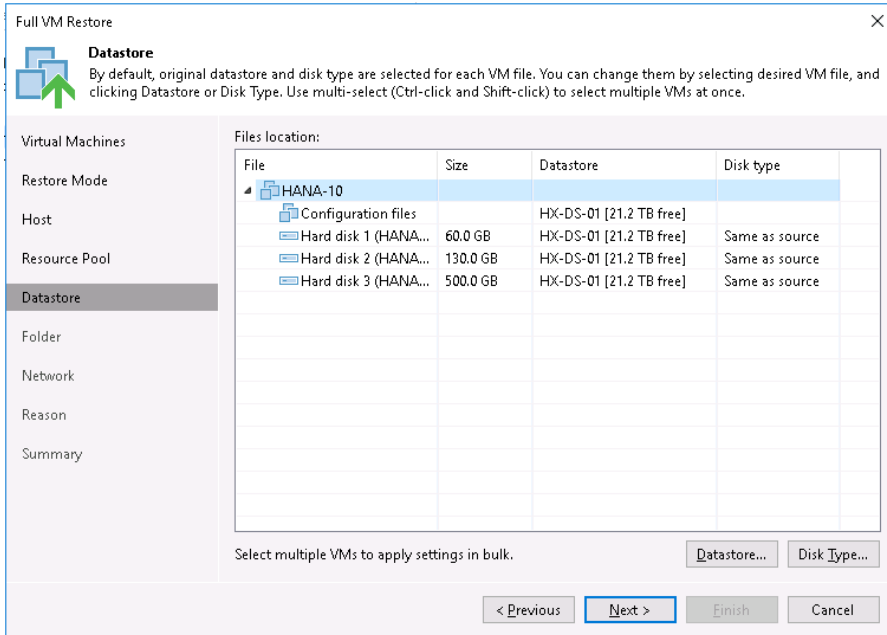
5. If required, change the host. Click Next.



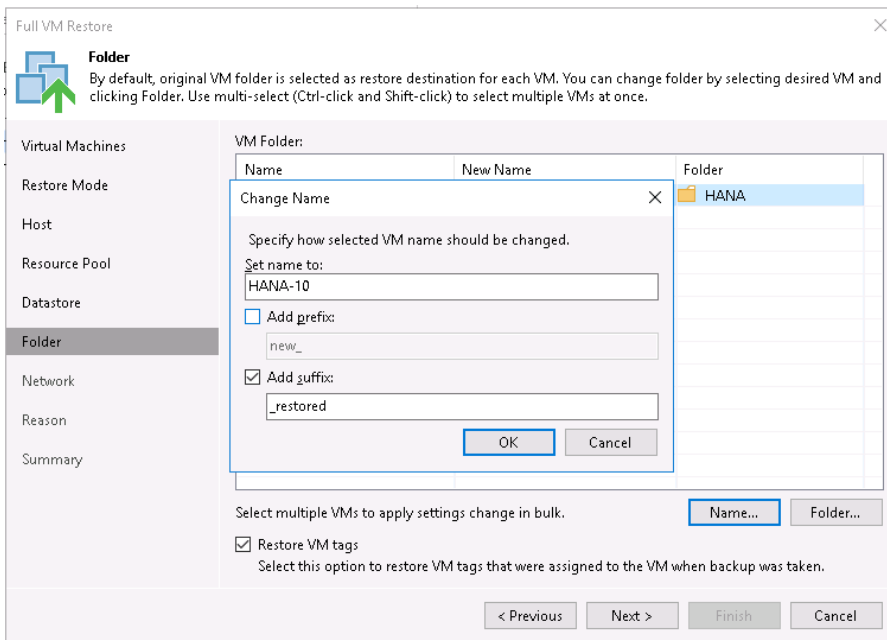
6. If required, change the resource pool. Click Next.



7. If required, change the data store. Click Next.



8. Click Name and specify the name of the virtual machine after the restore operation and click OK.



9. Click Next.

Full VM Restore
✕

Folder

By default, original VM folder is selected as restore destination for each VM. You can change folder by selecting desired VM and clicking Folder. Use multi-select (Ctrl-click and Shift-click) to select multiple VMs at once.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Reason

Summary

VM Folder:

Name	New Name	Folder
HANA-10	HANA-10_restored	HANA

Select multiple VMs to apply settings change in bulk. Name... Folder...

Restore VM tags
Select this option to restore VM tags that were assigned to the VM when backup was taken.

< Previous
Next >
Finish
Cancel

10. If required, change the network connectivity settings. Click Next.

Full VM Restore
✕

Network

By default, we will connect the restored VM to the same virtual networks as the original VM. If you are restoring to a different location, specify how networks map between original and new locations.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Reason

Summary

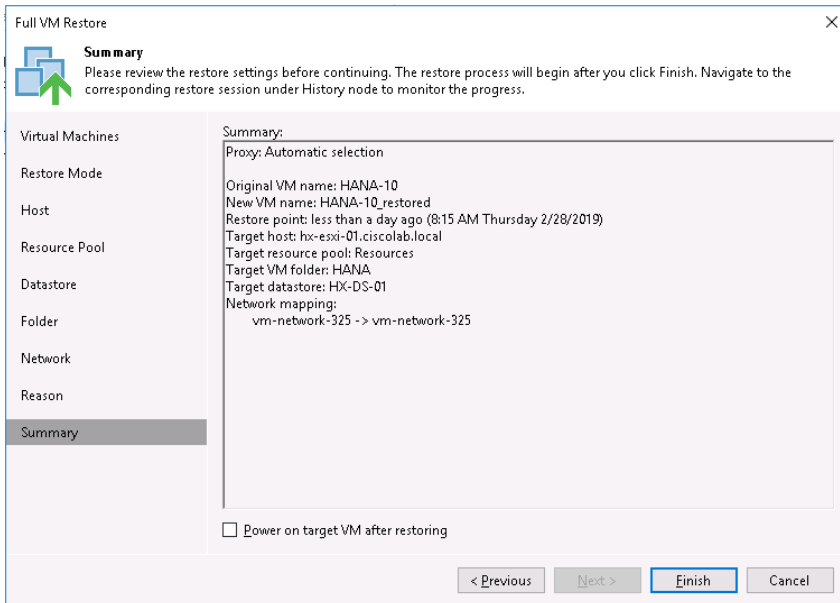
Network connections:

Source	Target
HANA-10	
vm-network-325	vm-network-325

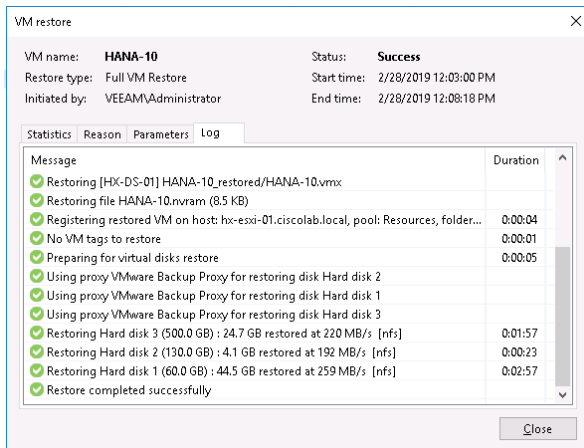
Select multiple VMs to apply settings change in bulk. Network... Disconnect

< Previous
Next >
Finish
Cancel

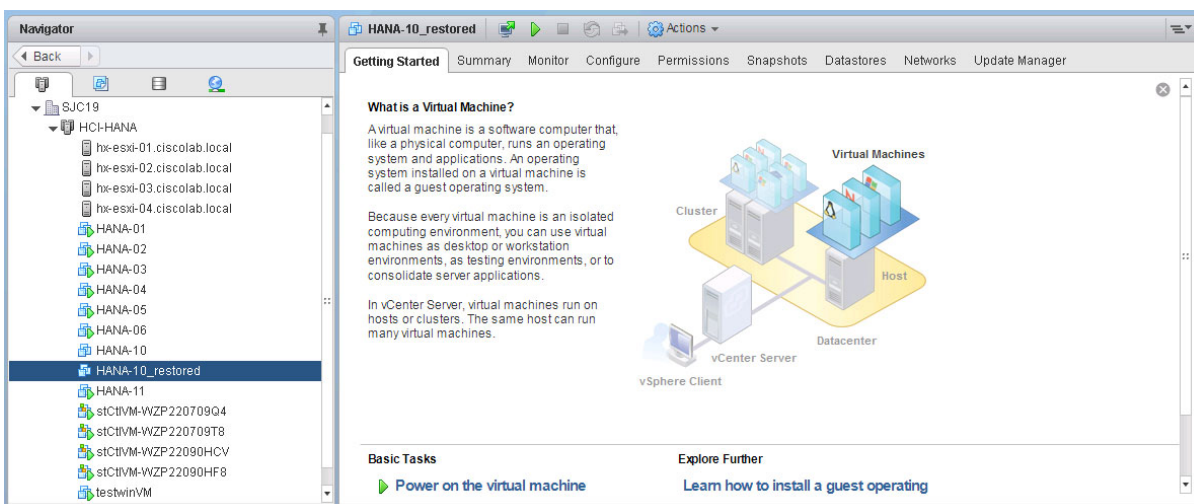
11. Review the settings and click Finish.



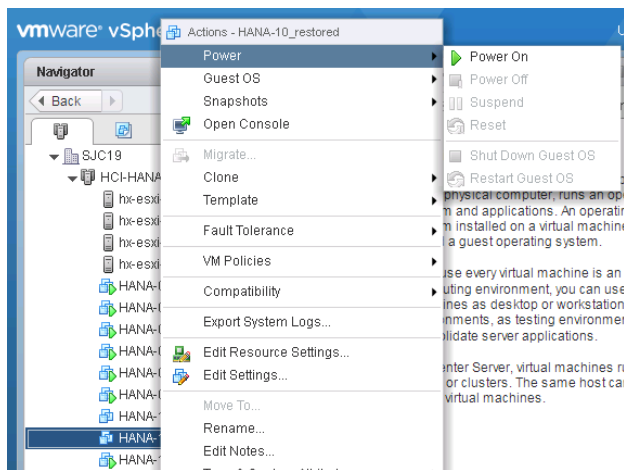
12. The virtual machine will be restored with the specified settings. After the process is finished, click Close.



The restored virtual machine is powered off on the target system.



13. Right-click the restored virtual machine and choose Power > Power On.



14. Log on to the virtual machine and check the status of the SAP HANA system.

```
bwpadm@hx-sles-10:/usr/sap/BWP/HDB10> HDB info
```

USER	PID	PPID	%CPU	VSZ	RSS	COMMAND
bwpadm	2650	2022	0.0	106964	3756	sshd: bwpadm@pts/0
bwpadm	2651	2650	0.4	15452	5340	_ -sh
bwpadm	2756	2651	0.0	13396	3348	_ /bin/sh /usr/sap/
bwpadm	2787	2756	0.0	36888	3080	_ ps fx -U bwpa
bwpadm	2021	1	6.0	428384	28432	/usr/sap/BWP/HDB10/exe/sa
bwpadm	2104	2021	6.1	21724	4440	_ sapstart pf=/usr/sap/
bwpadm	2111	2104	0.0	21732	3052	_ sapstart pf=/usr/
bwpadm	2119	2111	1.6	256036	49284	_ /usr/sap/BWP/
bwpadm	2135	2119	199	6418120	3710492	_ hdbnameuse
bwpadm	2284	2119	0.0	96216	5876	_ /usr/sap/
bwpadm	2288	2284	0.3	1499424	16916	_ SapBa
bwpadm	2319	2119	8.2	3342772	363932	_ hdbcompil
bwpadm	2321	2119	7.8	3673344	402736	_ hdbprepro
bwpadm	2365	2119	24.0	4520732	999740	_ hdbxsengi
bwpadm	2623	2119	0.0	96216	5836	_ /usr/sap/
bwpadm	2640	2119	0.0	96216	5880	_ /usr/sap/
bwpadm	1845	1	0.0	36752	4536	/usr/lib/systemd/systemd
bwpadm	1867	1845	0.0	88140	1660	_ (sd-pam)

```
bwpadm@hx-sles-10:/usr/sap/BWP/HDB10>
```

Configuring Cisco HyperFlex snapshot with virtual machine quiescing

For manual snapshots or snapshots taken using the VMware API, VMware Tools integration and OS quiescing can be used to create SAP HANA-aware snapshots by calling the same `hana-pre-freeze.sh` and `hana-post-thaw.sh` scripts as used by Veeam Backup and Replication. The main difference between the Veeam-managed approach and the manual or API approach is that the data will not be backed up in the Veeam repository. Instead, the data from the snapshot will remain in the Cisco HyperFlex storage.

One use for this option is maintenance work, such as a SAP HANA upgrade or data housekeeping.

This snapshot option should not be the only option used to protect SAP HANA. It must be combined with other options discussed in this document.

Unlike for the Veeam-managed option, for the VMware Tools option you need to copy the pre-freeze and post-thaw scripts directly to the SAP HANA servers. Refer to the VMware documentation for virtual machine quiescing for details.

To configure this backup type, you perform the following high-level steps:

1. Run the pre-freeze script on the SAP HANA server.
2. Quiesce the disk I/O.
3. Create a snapshot on the Cisco HyperFlex system for the virtual machine.
4. Run the post-thaw script on the SAP HANA server.

Prepare the SAP HANA virtual machine

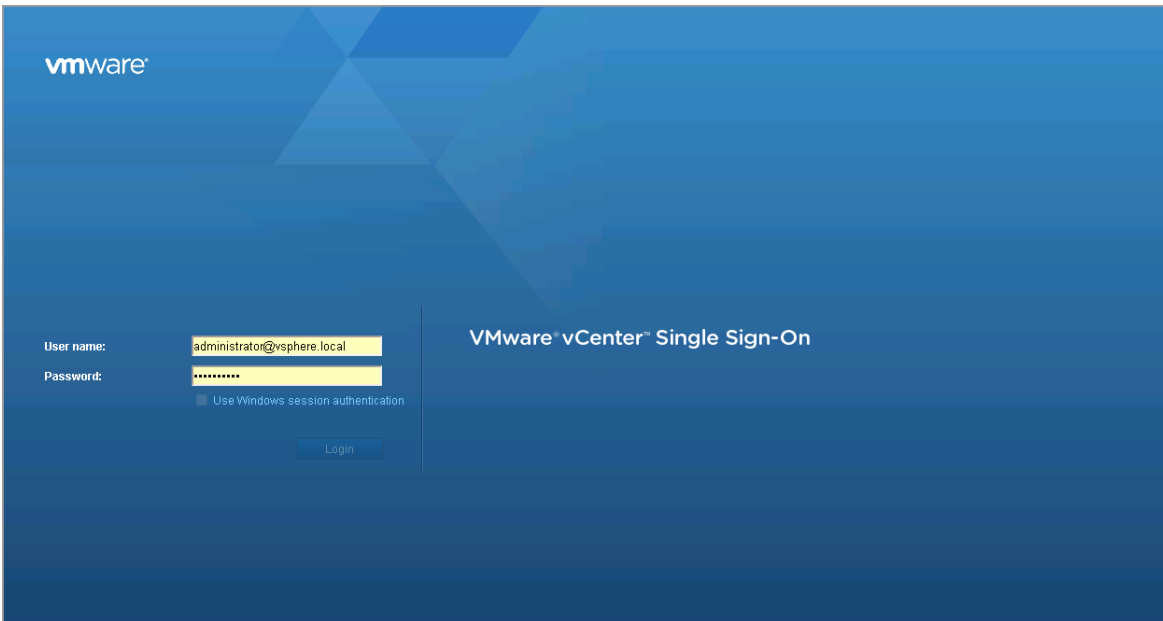
You first need to copy the scripts to the SAP HANA virtual machine.

1. Copy `hana-pre-freeze.sh` to `/usr/sbin/pre-freeze-script` and copy `hana-post-thaw.sh` to `/usr/sbin/post-thaw-script` on the SAP HANA system.
2. Repeat the preceding step for all virtual machines on which this option will be used.

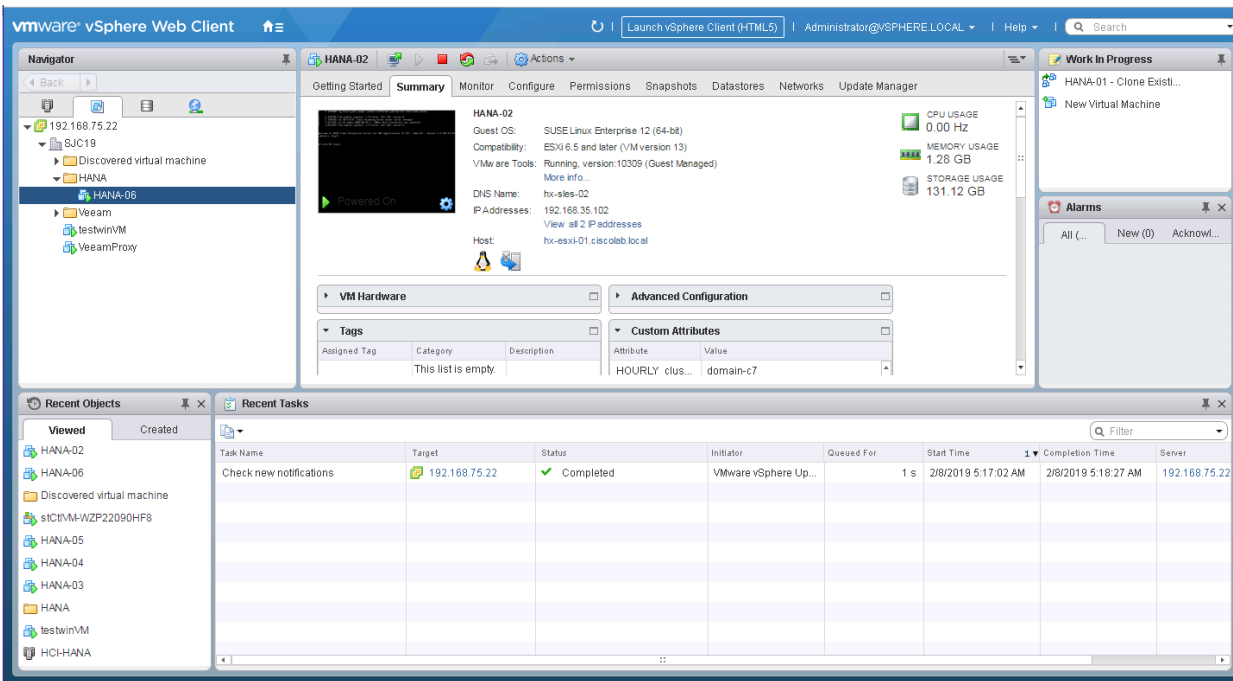
Configure scheduled snapshots in the VMware vSphere Client

Configure the snapshots for the vSphere Client.

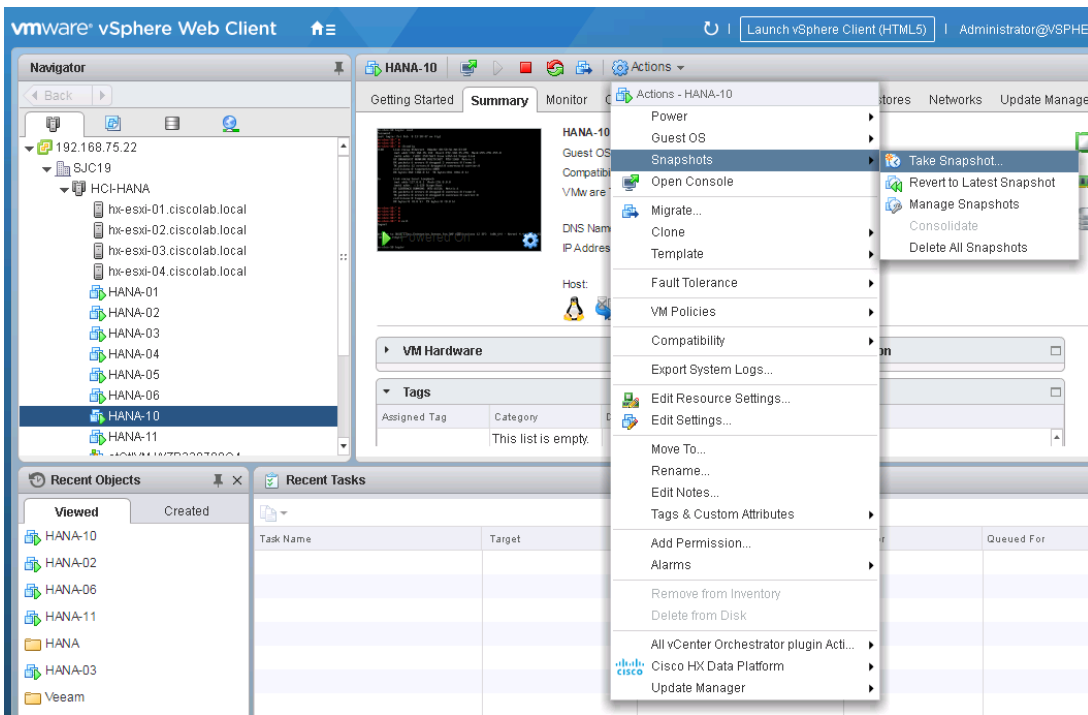
1. Connect to the vSphere Client from which the Cisco HyperFlex system is managed.



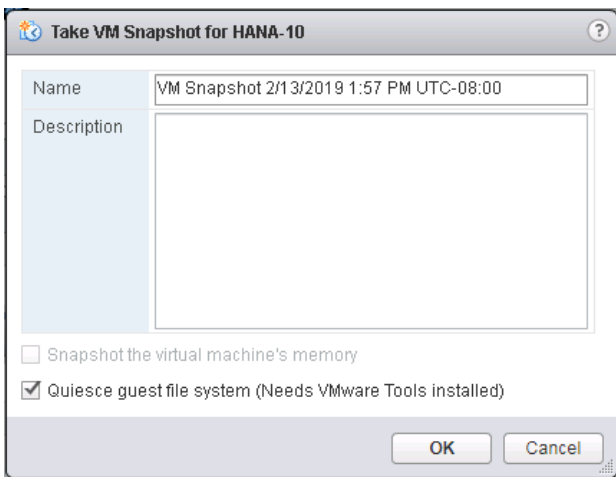
2. Select the virtual machine on which SAP HANA is running.




3. Choose Actions > Snapshots > Take Snapshot.



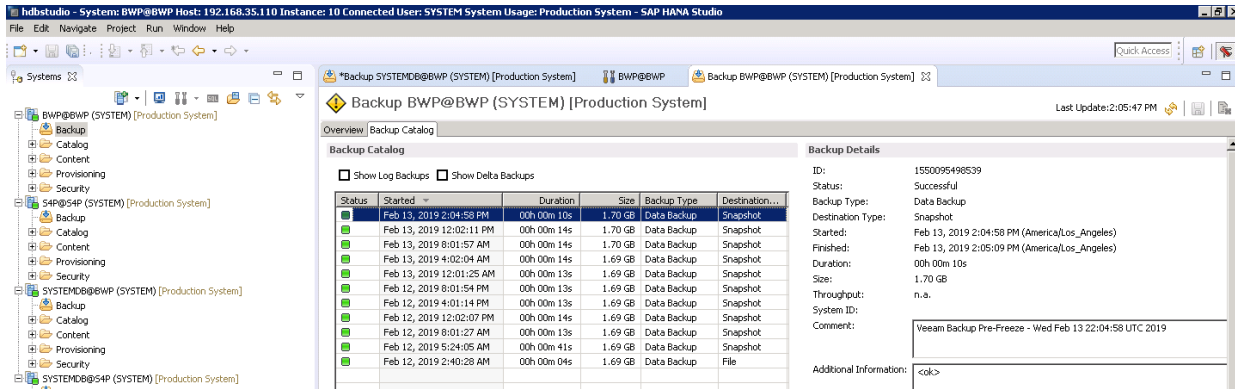
4. Deselect the check box for “Snapshot the virtual machine’s memory” and select the check box for “Quiesce guest file system.” Click OK.



The snapshot is taken and is displayed in the vSphere client.



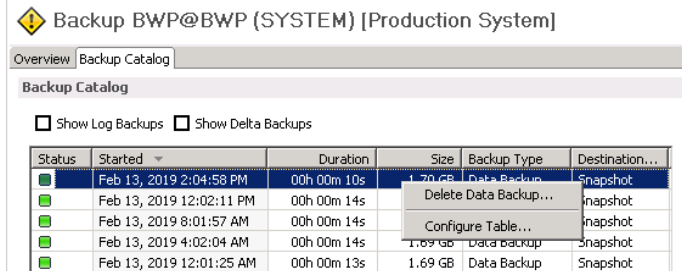
The snapshot also is listed in the SAP HANA backup catalog.



Delete the backup catalog entry for deleted or replaced snapshots

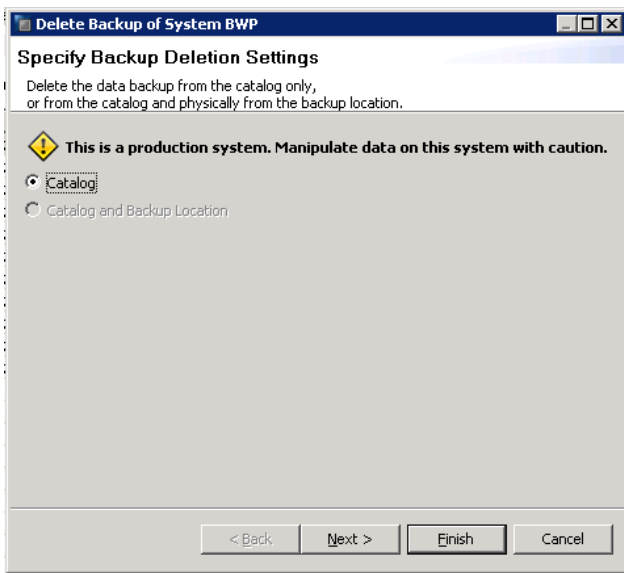
If a snapshot on the Cisco HyperFlex system is deleted or replaced, be sure to delete the entry from the SAP HANA backup catalog.

1. Select the entry and right-click it.

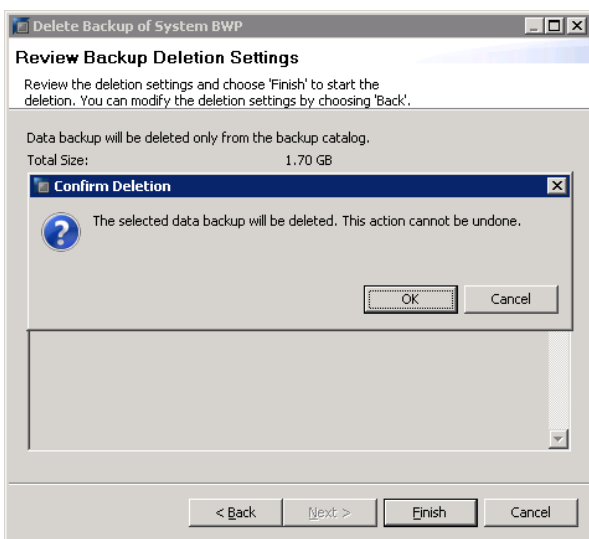


For snapshot-based backups, only the catalog entry can be deleted.

2. Click Finish.



3. Confirm the deletion, and after the entries are deleted click Finish.



Revert the virtual machine

To revert the virtual machine, you use the same process presented in the section “Configuring Cisco HyperFlex snapshot” earlier in this document.

Configuring SAP HANA Backint-based backup with the Veeam plug-in

With the release of Veeam Availability Suite 9.5 Update 4, the Veeam Plug-in for SAP HANA is available and can be used to create application-based backups as recommended by SAP. The plug-in creates a link between SAP HANA and the Veeam repository. The backup and restore jobs are managed by the SAP administrator within the SAP system. Veeam serves only as a storage repository.

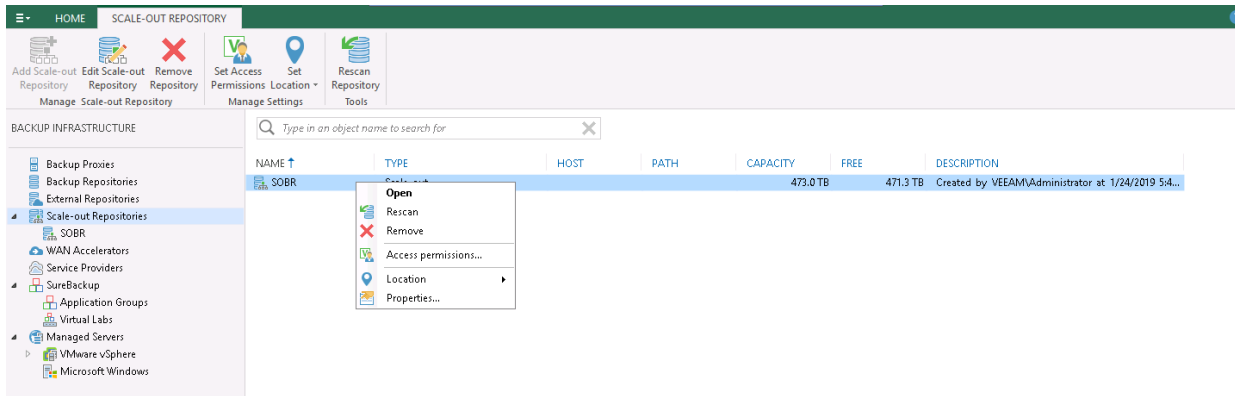
To configure this backup type, you perform the following high-level steps:

- Back up the SAP HANA database to the Veeam repository server using Backint.

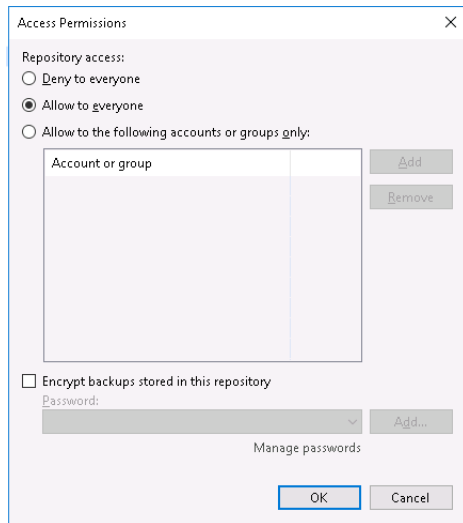
Configure the backup repository for backups from SAP HANA

To allow the SAP HANA system to write to a Veeam backup repository, you must grant access permissions in the Veeam system.

1. Log on to the Veeam Backup and Recovery console and under Backup Infrastructure, select either Backup Repositories or Scale-out Repositories. Then select the repository that will be used for SAP HANA backups.
2. Right-click the repository name and choose “Access permissions.”



3. Select “Allow to everyone” or maintain the account or group configuration based on the Veeam documentation. Click OK.

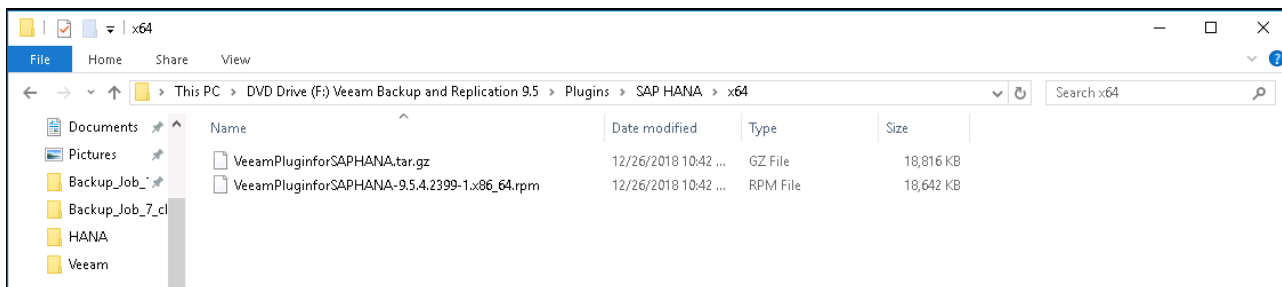


Your work at the Veeam console is complete.

Install the Veeam Plug-in for SAP HANA

The Veeam Plug-in for SAP HANA is located on the Veeam Backup and Recovery DVD.

1. Copy the SAP HANA plug-in from the Veeam DVD to the Linux system that has SAP HANA installed.



2. Install the plug-in on the SAP HANA system as user root.

```

hx-sles-11:~/Veeam # ls -l
total 37508
-rw-r--r-- 1 root root 19088526 Feb 11 16:24 VeeamPluginforSAPHANA-9.5.4.2399-
1.x86_64.rpm
-rw-r--r-- 1 root root 19267035 Feb 11 16:24 VeeamPluginforSAPHANA.tar.gz
hx-sles-11:~/Veeam #

```

```

hx-sles-11:~/Veeam # zypper in VeeamPluginforSAPHANA-9.5.4.2399-1.x86_64.rpm
Refreshing service 'SUSE_Linux_Enterprise_Server_for_SAP_Applications_12_SP3_x86_64'.
Retrieving repository 'SLE-12-SP3-SAP-Updates' metadata .....[done]
Building repository 'SLE-12-SP3-SAP-Updates' cache .....[done]
Retrieving repository 'SLE-HA12-SP3-Updates' metadata .....[done]
Building repository 'SLE-HA12-SP3-Updates' cache .....[done]
Retrieving repository 'SLES12-SP3-Updates' metadata .....[done]
Building repository 'SLES12-SP3-Updates' cache .....[done]
Loading repository data...
Reading installed packages...
Resolving package dependencies...

```

The following NEW package is going to be installed:

```
VeeamPluginforSAPHANA
```

The following package has no support information from its vendor:

```
VeeamPluginforSAPHANA
```

1 new package to install.

Overall download size: 18.2 MiB. Already cached: 0 B. After the operation, additional 48.3 MiB will be used.

```

Retrieving package VeeamPluginforSAPHANA-9.5.4.2399-1.x86_64
(1/1), 18.2 MiB ( 48.3 MiB unpacked)

```

```
VeeamPluginforSAPHANA-9.5.4.2399-1.x86_64.rpm:
```

```
Package is not signed!
```

```
VeeamPluginforSAPHANA-9.5.4.2399-1.x86_64 (Plain RPM files cache): Signature verification failed [6-File is unsigned]
```

```
Abort, retry, ignore? [a/r/i] (a): i
```

```
Checking for file conflicts: .....[done]
(1/1) Installing: VeeamPluginforSAPHANA-9.5.4.2399-1.x86_64 .....[done]
Additional rpm output:
Run "SapBackintConfigTool --wizard" to configure the Veeam Plug-in for SAP HANA
```

```
hx-sles-11:~/Veeam #
hx-sles-11:~/Veeam # SapBackintConfigTool --wizard
Enter backup server name or IP address: veeam
Enter backup server port number 10006:
Enter username: Administrator
Enter password for Administrator:*****
Available backup repositories:
1. Default Backup Repository
2. SOBR
Enter repository number: 2
Configuration result:

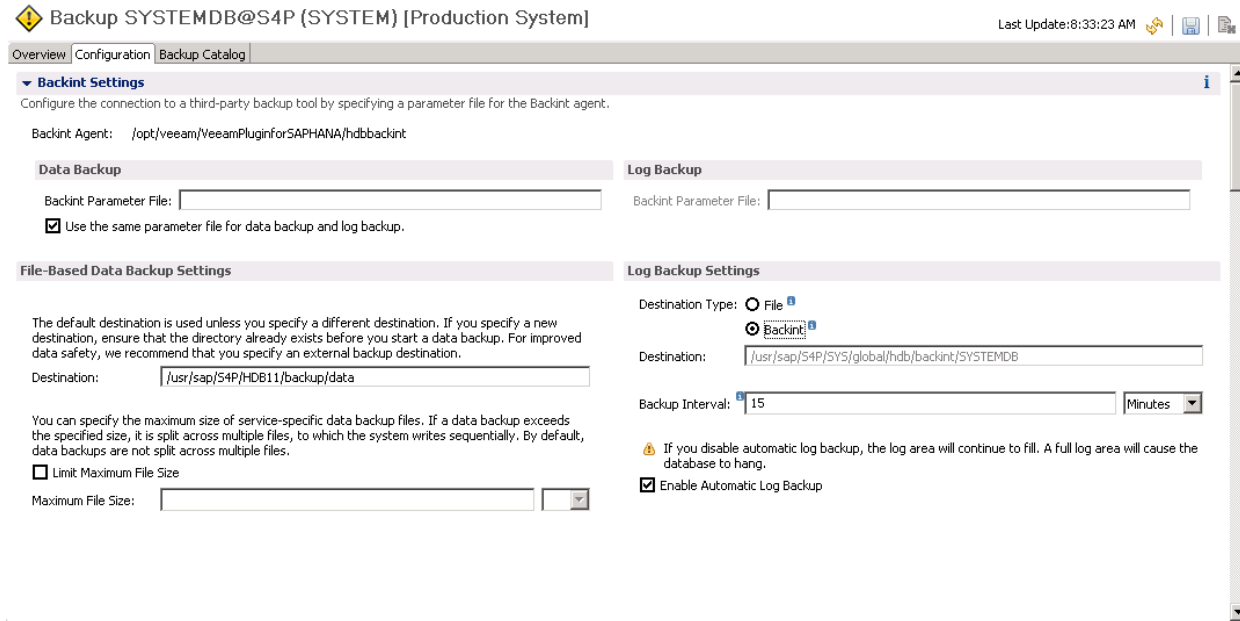
SID S4P has been configured

hx-sles-11:~/Veeam #
```

Configure backup through Backint in SAP HANA Studio

In SAP HANA Studio, under Configuration > Backint Settings, Backint Agent is now pointing to /opt/veeam/....

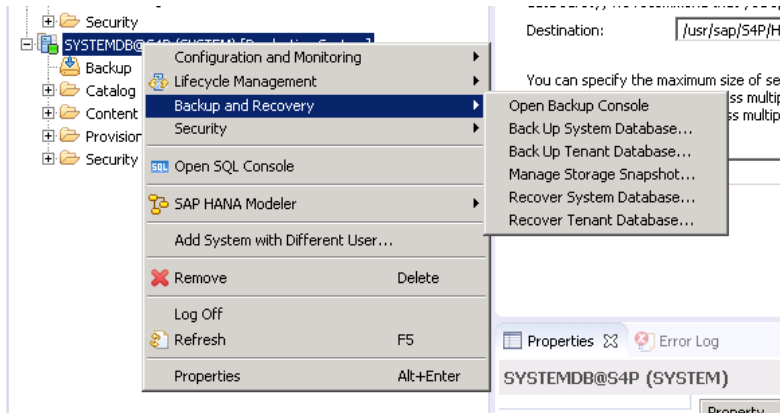
1. To enable the SAP HANA system to also back up the log files to Veeam, change Destination Type from File to Backint.
2. Save the new configuration.



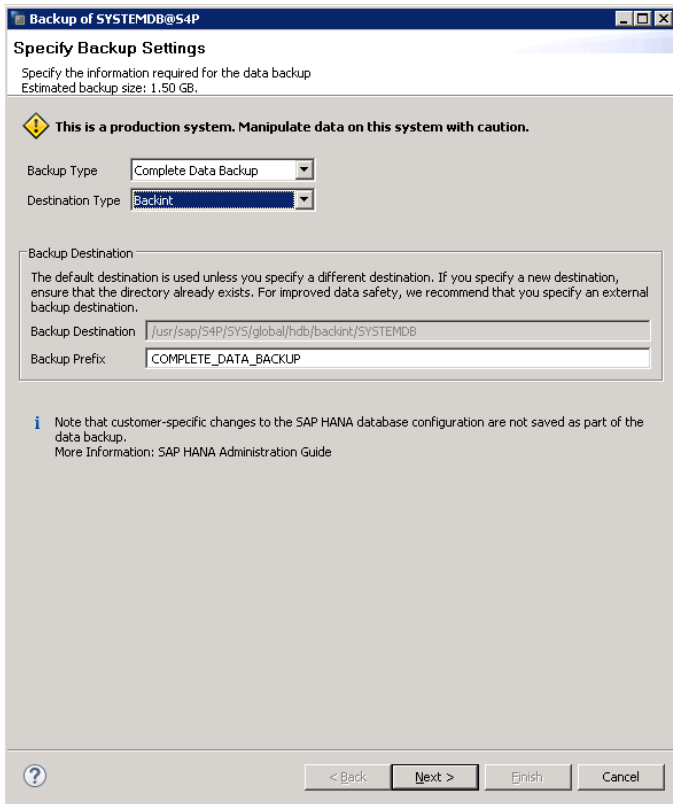
Perform manual backup from SAP HANA Studio

A simple way to test the new configuration and also establish the backup job configuration within Veeam is to initiate a manual backup of the system database.

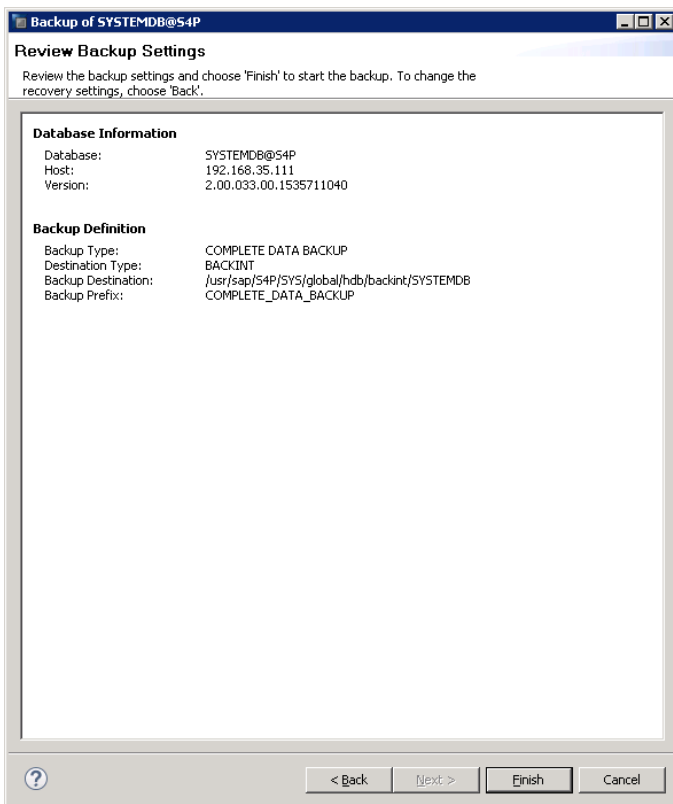
1. In SAP HANA Studio, right-click the SYSTEMDB file of the changed system and choose Back Up System Database.

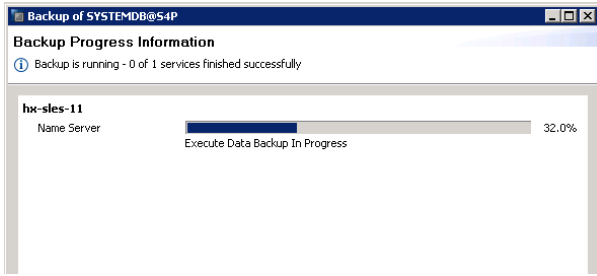


2. In the Destination Type field, choose Backint and click Next.

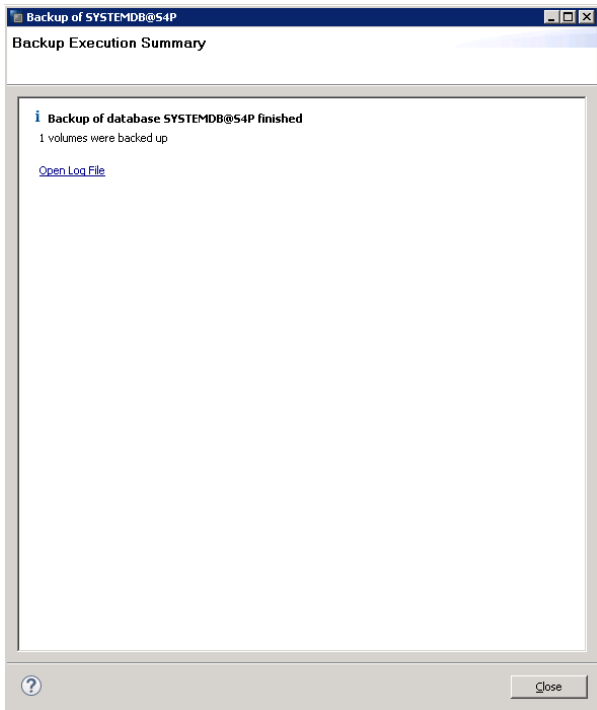


3. Review the job summary and click Finish.

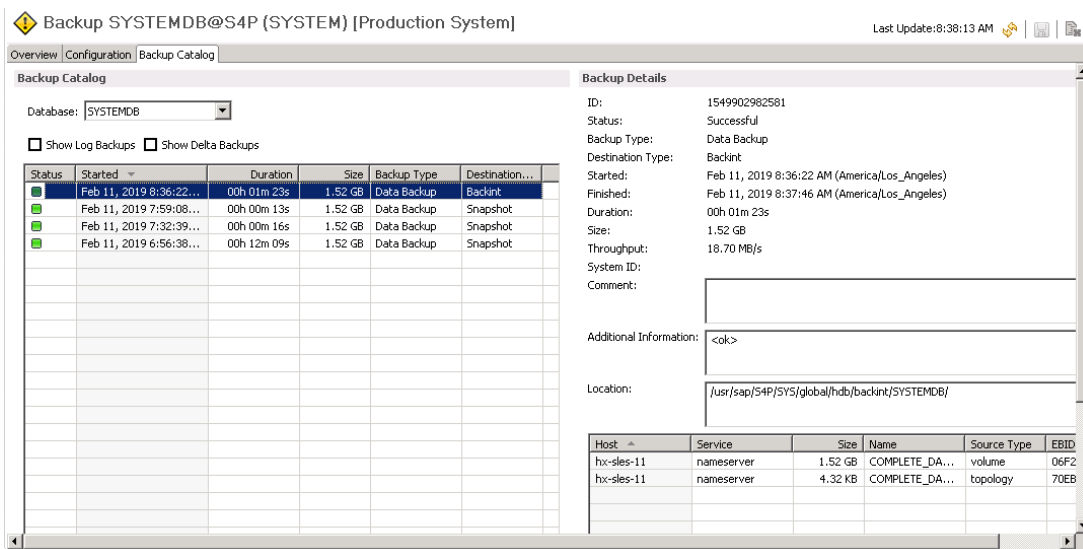




4. After the backup job is finished, click Close to close the window.



In the backup catalog, a new entry with Backint as the destination is listed.



Monitor backup progress and jobs in Veeam

In the Veeam console, a new job is listed as “<hostname> SAP backing backup (<Repository>).” This location is used for all backup jobs from this host. The information about the database or backup type is shown in the job details.

NAME ↑	TYPE	STATUS	LAST RUN	LAST RESULT	NEXT RUN	DESCRIPTION
hcx-sles-11 SAP backint backup (SOBR)	SAP HANA Backup	Stopped	2 minutes ago	Success	<not scheduled>	SAP HANA backint job
VM-Snap-4-SAP-HANA	VMware Backup	Stopped	41 minutes ago	Success	2/11/2019 8:00:0...	Application aware snapshot based backup of SAP...

SUMMARY	DATA	STATUS	THROUGHPUT (ALL TIME)
Duration: 01:24	Processed: 1.5 GB (100%)	Success: 1 ✔	
Processing rate: 18 MB/s	Read: 1.5 GB	Warnings: 0	
Bottleneck: Source	Transferred: 594.1 MB (2.6x)	Errors: 0	

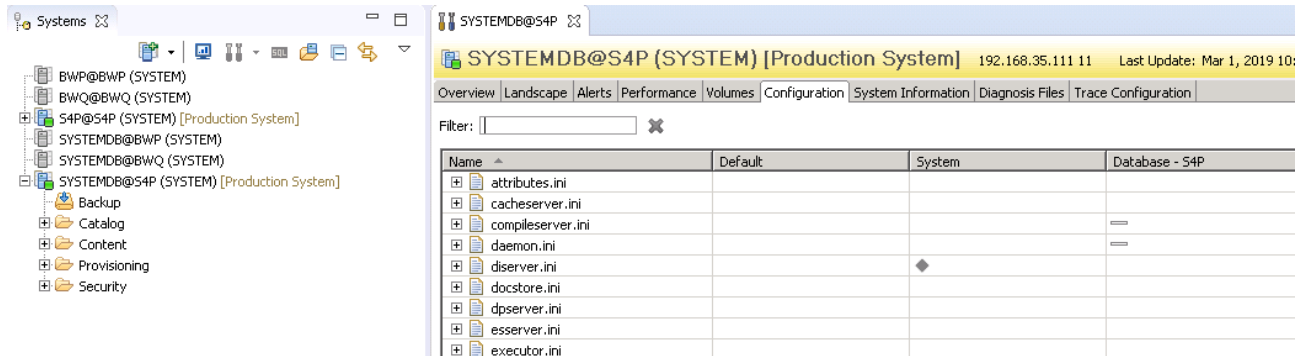
NAME	STATUS	ACTION	DURATI...
hcx-sles-11	✔ Success	<ul style="list-style-type: none"> ✔ SAP HANA Backup job started at 2/11/2019 4:36:59 PM. ✔ Backup is currently in progress, statistics will be updated upon job completion ✔ Processing backup file: /usr/sap/SAP/SYS/global/hdb/backint/SYSTEMDB/COMPLETE_DATA_BACKUP_databackup_0_1 ✔ Processing backup file: /usr/sap/SAP/SYS/global/hdb/backint/SYSTEMDB/COMPLETE_DATA_BACKUP_databackup_1_1 ✔ Creating metadata ✔ Session completed 	01:24

Back up the SAP HANA backup catalog

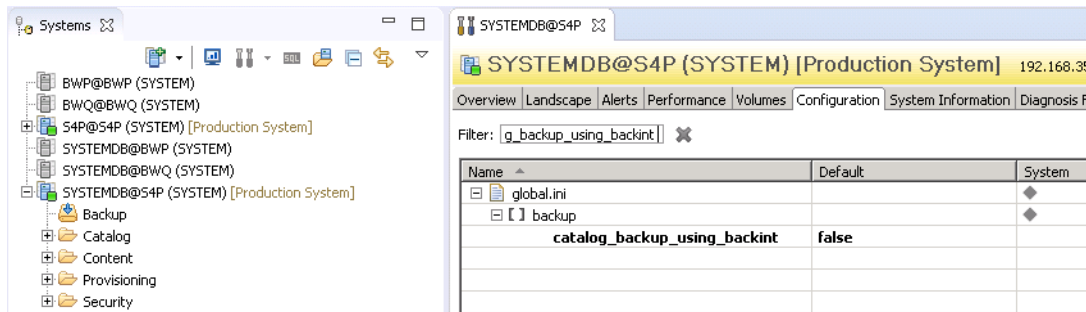
With the Veeam Plug-in for SAP HANA Backint installed, the default location for the backup catalog backup is still the former destination: located at the file location for the log backup. You should change the location to Backint.

To change the location for the catalog backup from File to Backint, you need to change the parameter **catalog_backup_using_backint** in the global.ini file from **false** to **true**. An easy way to change parameters is to use SAP HANA Studio.

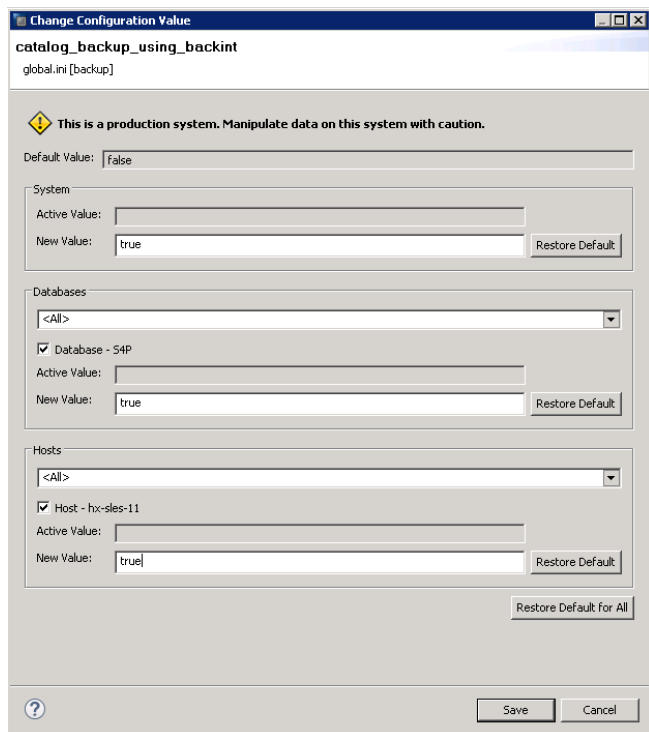
1. Open SAP HANA Studio and double-click the SYSTEMDB entry of the system that you want to change. Select the Configuration tab at the right.




2. In the Filter field, enter `catalog_backup_using_backint` and press Enter.



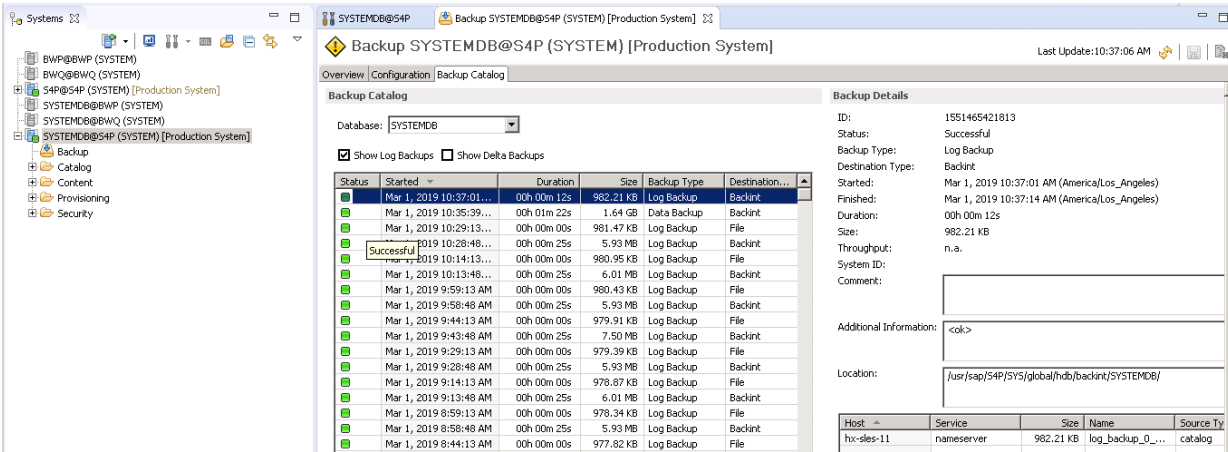
3. Double-click the field with the value “false.”
4. In the new window that appears, enter `true` as the new value in all sections. Click Save.



The parameter becomes active without a system restart.



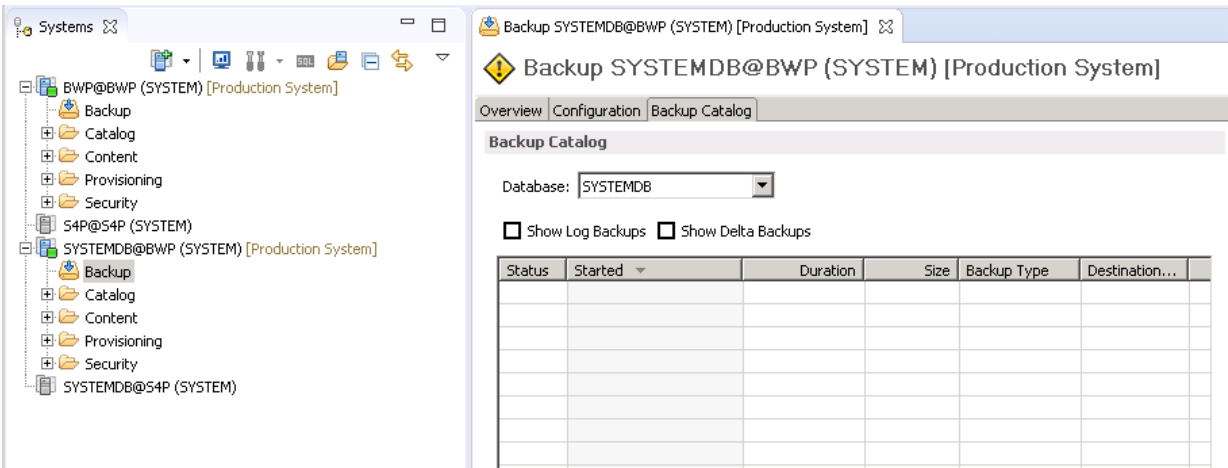
The next catalog backup will automatically use the Backint integration.



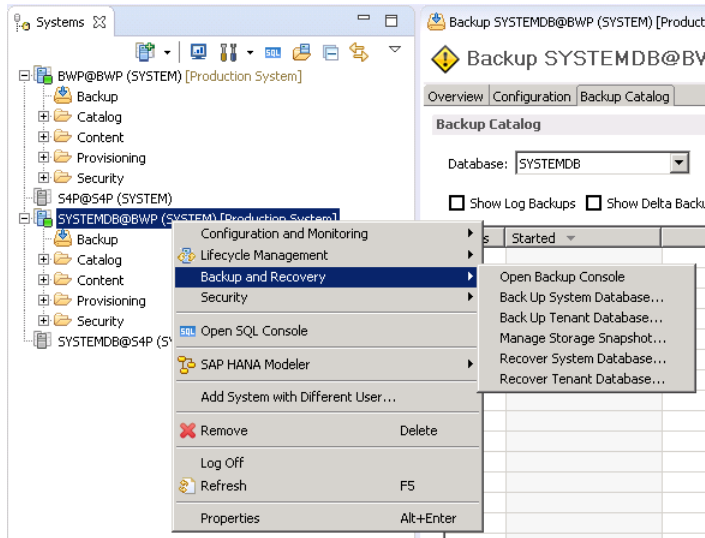
Recover a database from Backint

To demonstrate recovery from Backint, a new SAP HANA system with the SID BWP was installed. To enable the Veeam Plug-in for SAP HANA Backint, the plug-in was installed on the system and linked to the SAP HANA system as described earlier in this document.

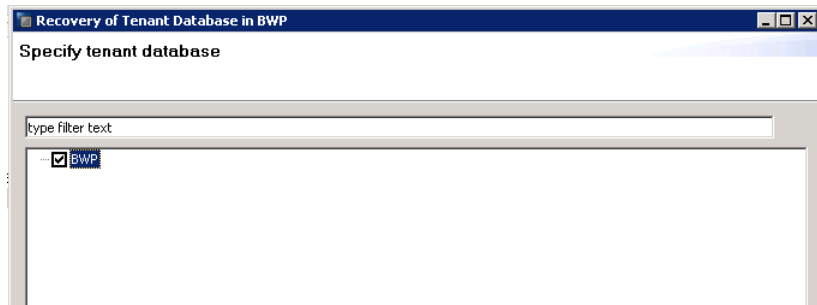
In the new system, the backup catalog is empty.



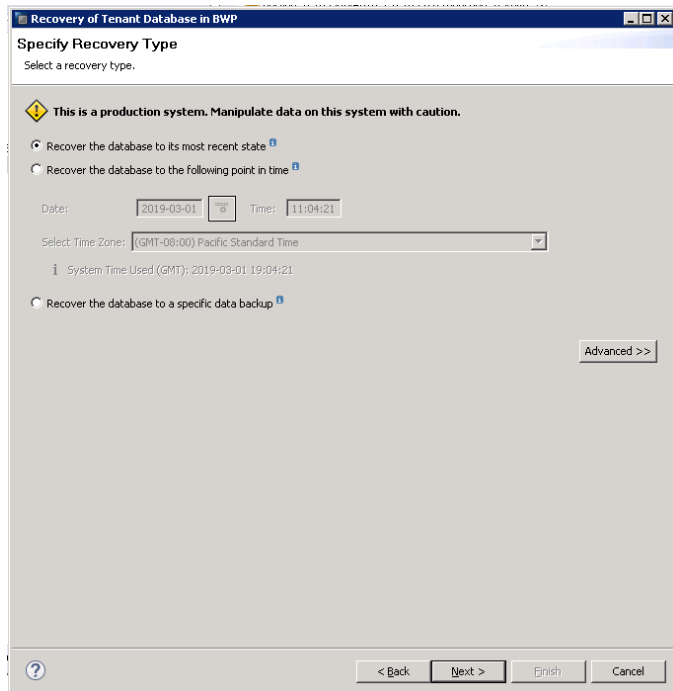
1. Right-click SYSTEMDB@BWP and select Recover Tenant Database.



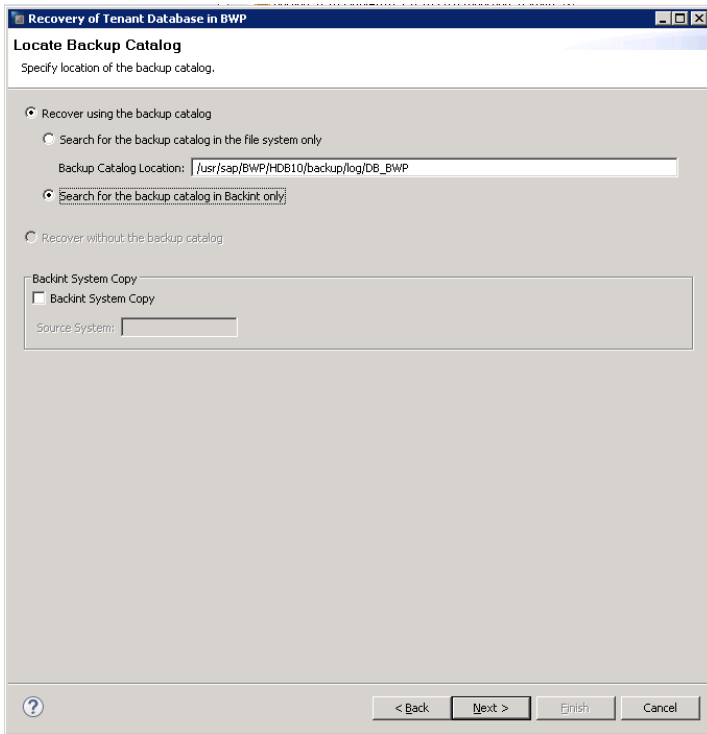
2. Select BWP and click Next.




3. Select the recovery mode that best meets your needs. In most cases, recovery to the most recent state is the best choice. Click Next.



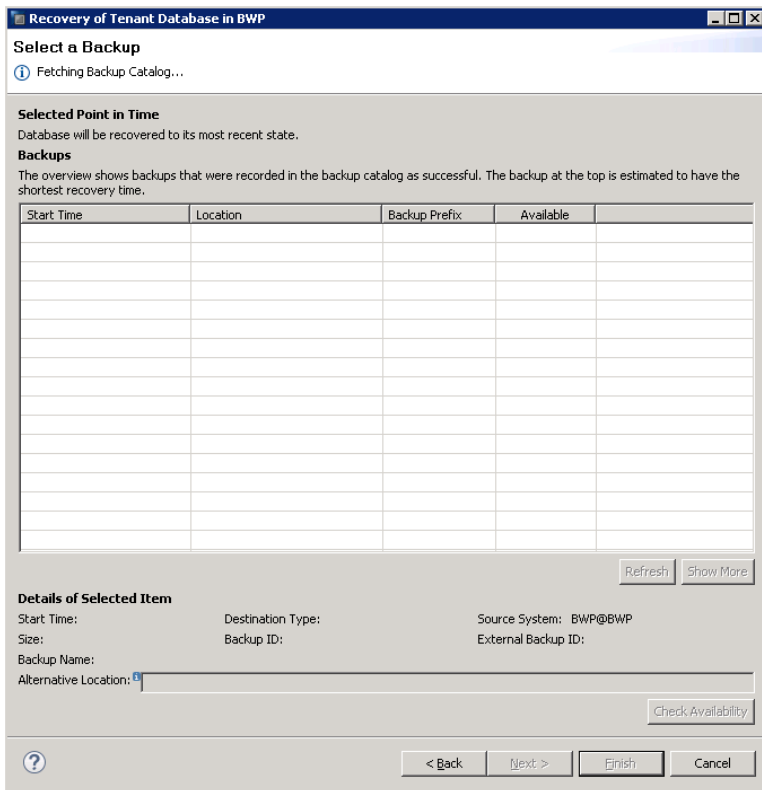
4. Select “Recover using the backup catalog” and “Search for the backup catalog in Backint only.” Click Next.



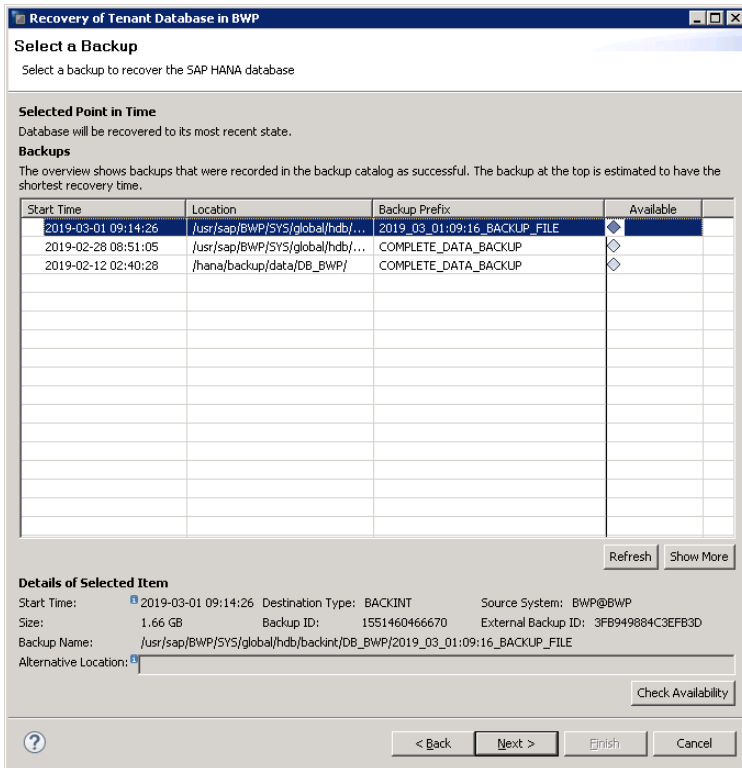
5. Accept the warning telling you that the system will be stopped if it is a production system.



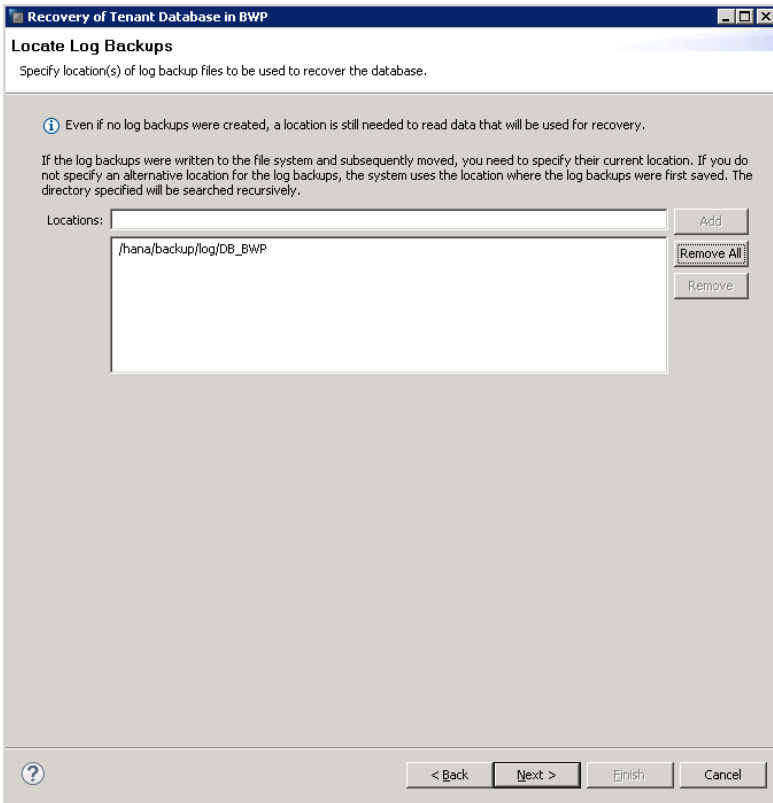
The system fetches data from the VBR server using Backint.



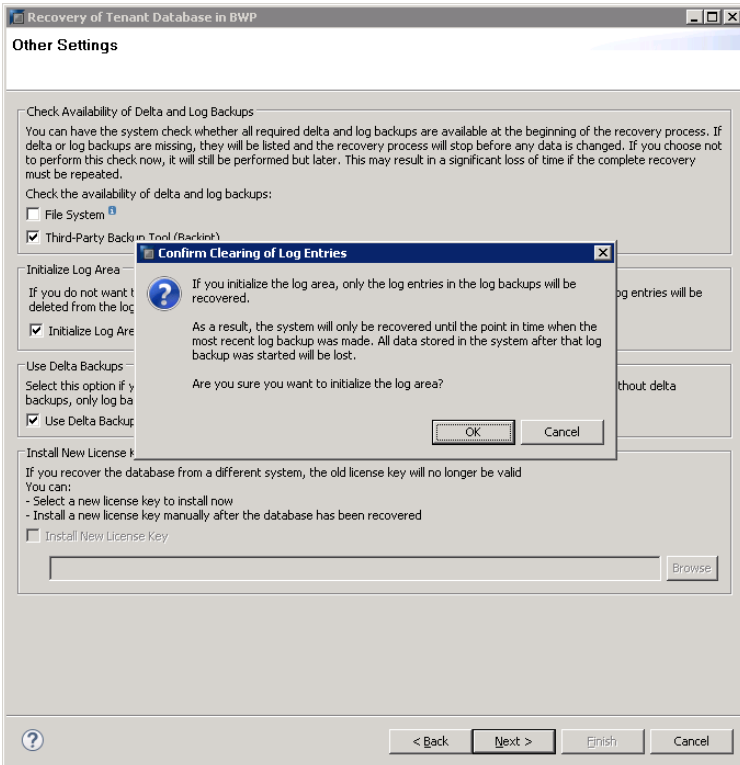
6. Select the backup version that you want to restore and click Next.



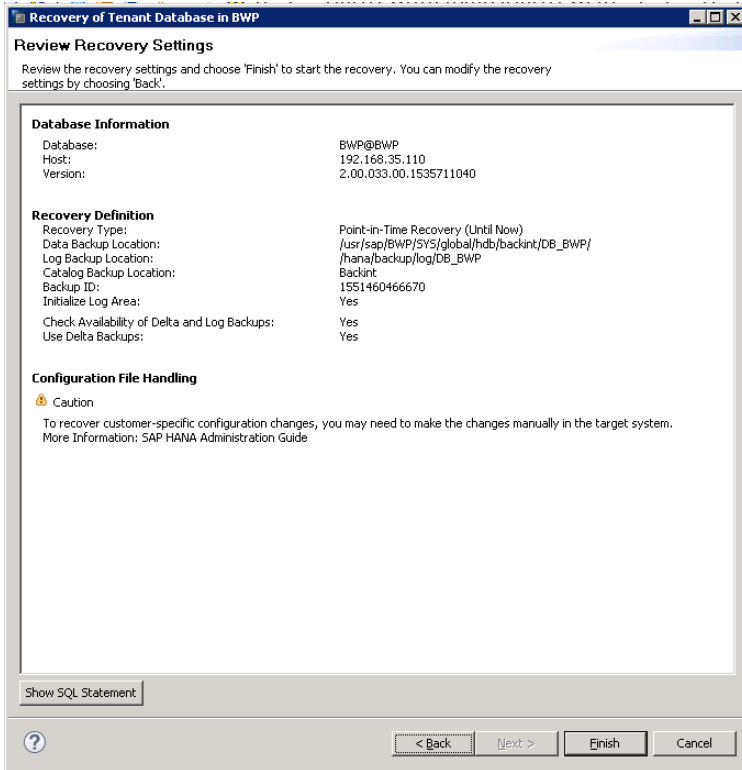
7. Enter the location for the database log file backups on the file system: for example, /hana/backup/log/DB_BWP. Click Next.



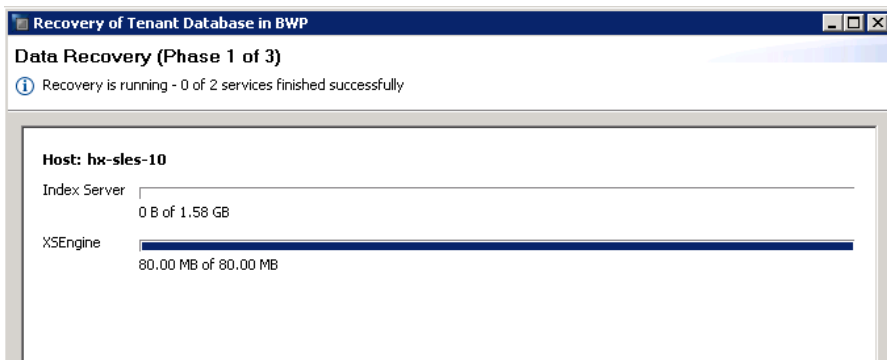
8. Select Third-Party Backup Tool (Backint) and Initialize Log Area. Confirm the selection and click Next.



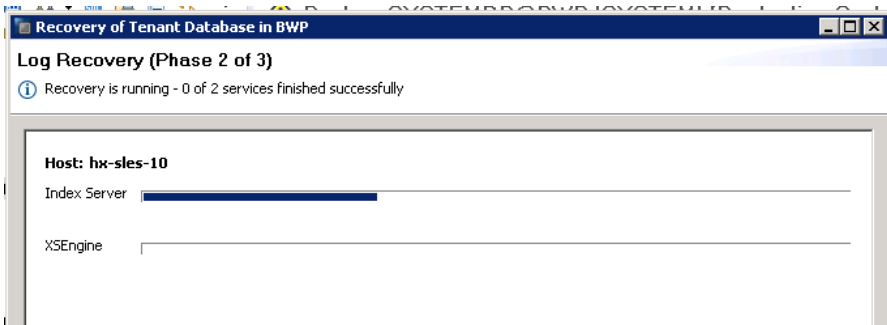
9. Review the settings and click Finish.



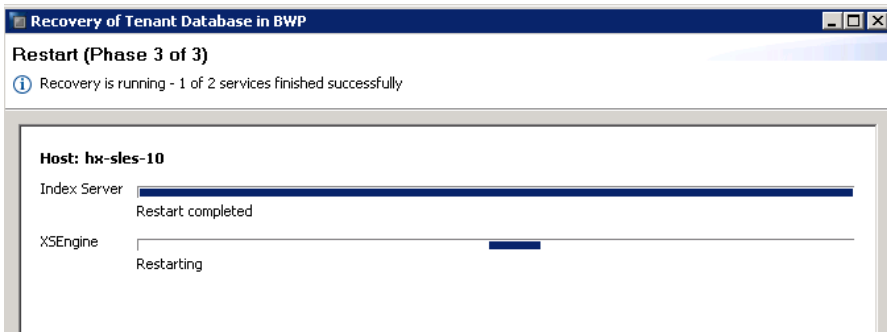
First the selected data backup file is used to recover the database.



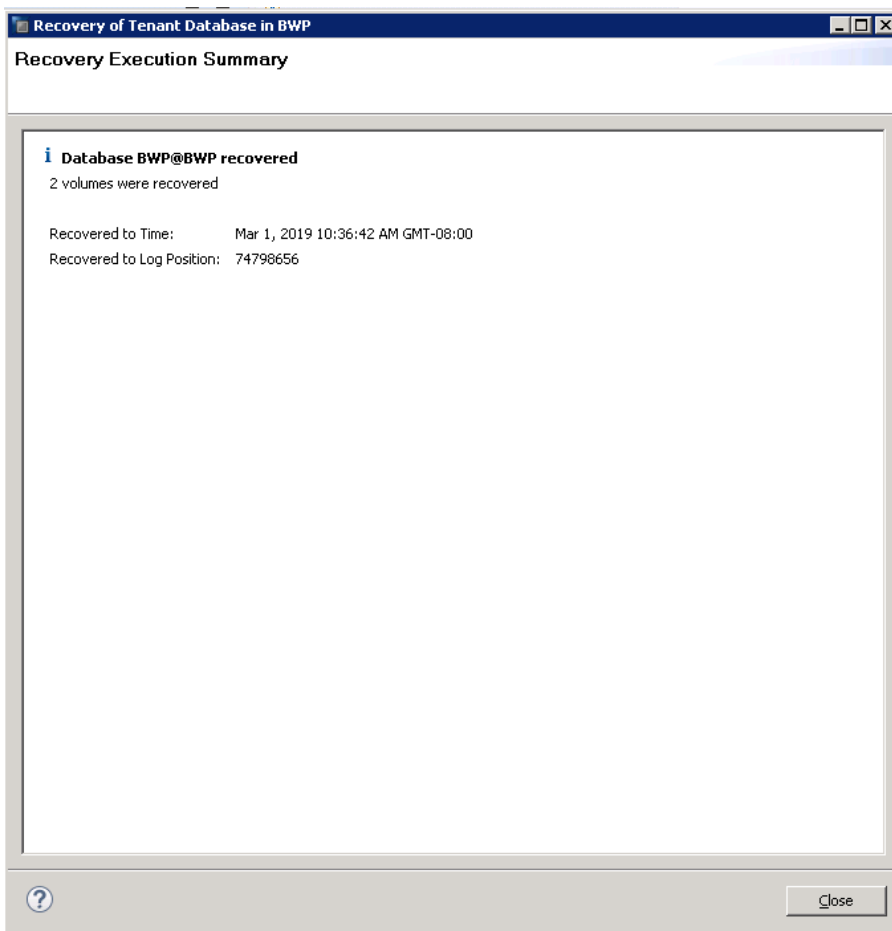
Next, the available log file backups are applied to the database.



Then the database is started.



10. Click Close to finish the recovery process.



Best practices

One best practice is recommended for all customers: Combine two or more of the methodologies described in this document to reduce risk and improve recovery-point objective (RPO) and RTO.

Here is one approach for combining the various methodologies:

- Create Cisco HyperFlex snapshots every hour.
 - Provides a simple restart point for the whole virtual machine
 - Does not place an entry in the backup catalog

- Create Veeam-managed snapshot-based backups every 4 hours.
 - Provides an application-aware restart point for the whole virtual machine
 - Places an entry in the backup catalog
 - Requires full virtual machine recovery
- Create a Backint-based backup once a day or week.
 - Provides application-based backup and recovery
 - Used in combination with the other options, allows a roll forward of the database

Conclusion

The seamless integration of Veeam technology with the Cisco HyperFlex HX Data Platform provides a best-in-class solution for SAP HANA systems that can meet business SLA RPO/RTO requirements and help ensure application consistency. The Cisco HyperFlex hyperconverged platform is an excellent choice for running SAP HANA.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)