



Red Hat
Virtualization (RHV)
Backups by SEP

Table of Contents

INTRODUCTION AND OVERVIEW

AGENT BASED BACKUP

IMAGE LEVEL BACKUP VIA RHV API

RHV BACKUP WITH SEP

BACKUP GRANULARITY

SETTING UP SEP ON RHV

BACKUP AND RECOVERY SCENARIOS

LAN BASED BACKUP

LAN BASED RECOVERY

LAN FREE BACKUP

LAN FREE RECOVERY

RHV CLUSTER

BACKUP DEDUPLICATION AND SITE-TO-SITE REPLICATION

PROTECTION OF NEWLY DEPLOYED VIRTUAL MACHINES

DATABASE AND APPLICATION VENDOR INSTRUCTIONS

SECURITY

CONCLUSION

Introduction and Overview

SEP backup and disaster recovery software is a robust, easy to manage, secure solution for businesses of any size. The SEP solution for Red Hat Virtualization (RHV) provides fast and efficient backup capabilities for RHV environments of any size. SEP leverages its patented Multi-Streaming Technology to simultaneously back up multiple streams of data, while simplifying and automating backups for all environments. Entire enterprise infrastructures can be managed by a single interface.

The SEP backup solution for RHV provides:

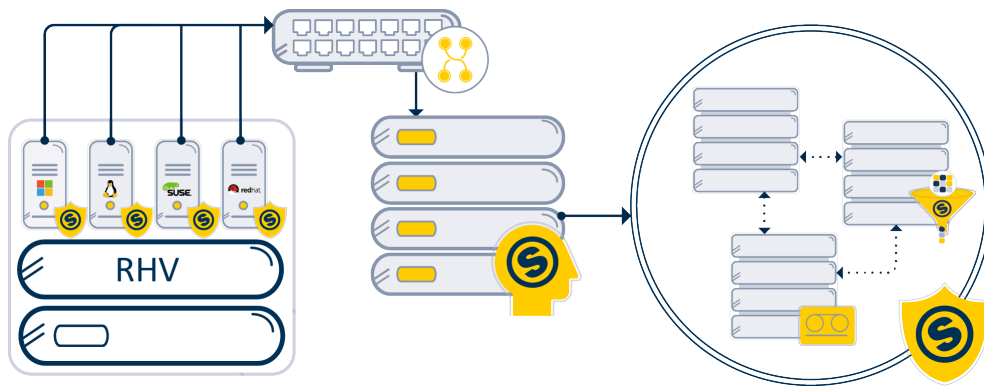
- Agentless backups through the RHV API
- Platform-independent functionality
- Built-in encryption capabilities for data while in flight and at rest
- A single enterprise management interface
- Built-in encryption capabilities

There are two different ways SEP can backup RHV virtual machines:

- Agent Based Backup
- Image Level Backup via RHV API

Agent Based Backup

Performing consistent Agent Based Backups of virtual machines running in a RHV environment is straightforward, where each VM is treated as if it were a physical machine and an agent is installed. Adding an agent to each VM will offer file level backups and restores.



[Figure 1] Process of Data Backup: The installation of a backup client software package must be performed on every virtual machine.

Pros:

- Adding an agent to each VM will offer file level backups and restores.
- Backup tasks can be scheduled.
- Backups can leverage deduplication and replication.

Cons:

- In the event of a Disaster Recovery, the VMs are treated like physical machines, which means the recovery process must be started with a bare metal disaster recovery solution (e.g. SEP BMR for Windows or Linux).
- Can be tedious and time consuming to install agents on every VM in large environments.

Image Level Backup via RHV API

The primary focus of this document is RHV Image Level Backups via the RHV API, where virtual machines are backed up without incurring downtime.

SEP's backup solution features integration with the RHV backup API, which was released with Red Hat Enterprise Virtualization 3.5.1. SEP is now capable of performing consistent, agentless backups of all RHV virtual machines, regardless of their workload or operating system. By leveraging the live snapshot capabilities built into RHV, reliable snapshot backups of virtual machines are assured without any downtime whether they are operating at the time of backup or not.

Pros:

- No downtime during backup
- Easy disaster recovery of virtual guests
- Optimized performance for virtual machines
- Fast LAN-free backup and restores are supported
- Requires minimum backup storage space due to SEP Si3 Deduplication technology
- Advanced backup functionality
- Backups and restores of any guest operating system can be performed

Cons:

- May need to add agents to support native, online backups of databases running on VMs

In the following sections, we discuss the implementation of this approach with SEP's backup solution.

This document does not cover any specific SEP installation instructions, but rather a general overview of a recommended configuration. Although installation is very simple, a good working knowledge of SEP is recommended.

RHV Backup with SEP

Since SEP communicates directly with the RHV hypervisor, it does not matter what operating system is running on the VM. Windows, Linux or any other operating system supported by the RHV hypervisor is treated as a virtual guest and backed up via direct communication with the hypervisor.

Backup Granularity

Backups can be as granular as desired. SEP can back up everything from the VM configuration, the individual disk of a single VM, up to all virtual machines on a RHV cluster.

Setting up SEP on RHV

Required components for this solution are:

- SEP Backup Server
- SEP Remote Device Server (RDS)
- DataMover (SEP Backup Server/RDS can be used for this task)
- Backup Storage (backup target)

Setting up SEP for RHV is easy and users can be up and running and ready to back up their first RHV VM in just 10 minutes after installing the SEP backup solution. No agents are required on the RHV manager, hypervisors or virtual machines – all that is required is a VM within the RHV environment to run SEP and suitable backup storage. API integration means SEP can look up and reference all VM names on the RHV server, making configuration simple and straightforward.

Backup and Recovery Scenarios

SEP RHV backups can operate in two modes: (1) LAN based backup, where data is transferred over the Local Area Network and (2) LAN free backup, where backups are sent directly from the host to the backup target. Users will need to decide which is best suited for their objectives and specific use case.

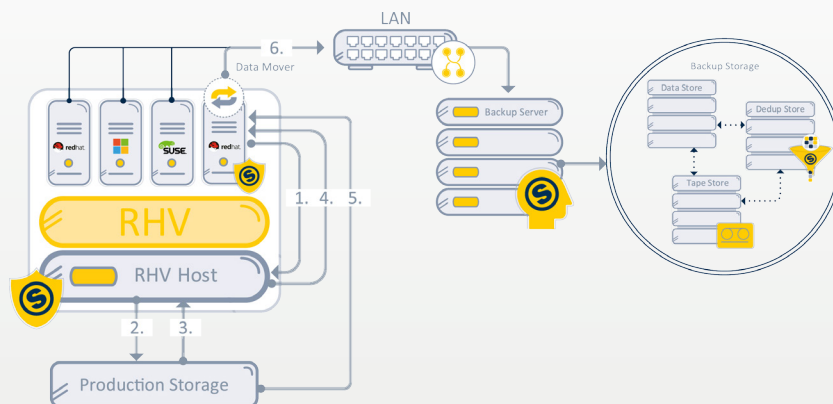
LAN Based Backup

The first scenario assumes that the SEP server is external to the RHV environment and that the backup storage is attached directly to the backup server. SEP is hardware agnostic, so any type of storage that is recognized by an operating system is supported. Backups can be sent to any kind of disk or tape, most commonly, SATA, SCSI, FC or other SAN/NAS based storage.

When the SEP server is external to the RHV environment, a single VM must be dedicated to the SEP backup environment within the RHV datacenter. This VM is referred to as the SEP DataMover. The SEP backup environment consists of an installation of a SEP Remote Device Server running on Red Hat Enterprise Linux (version 7 or later).

Note, that it is possible for the SEP server itself to reside within the RHV environment and in that instance, the SEP backup server can also be the DataMover.


Figure 2 below displays the LAN Based Backup architecture and the basic data flow. As you can see, all backup and restore data flows over the LAN, which must be capable of handling the amount of data transfers required to back up the VMs on a regular basis.



[Figure 2] The steps to back up a VM are as follows:

- [1.] A snapshot of the VM to be backed up is created. This enables consistent backups of the VMs running without interruption.
- [2.][3.][4.] Once the snapshot is successfully created, the VM configuration is exported from the snapshot and backed up as the first stream in the SaveSet.
- [5.][6.] The VM disks contained in the snapshot are then attached to the DataMover VM using the backup API. Once attached, the data on the disks is backed up. After the backup has completed successfully, the snapshot is removed.

LAN Based Recovery



[Figure 3] The restoration process includes the following steps:

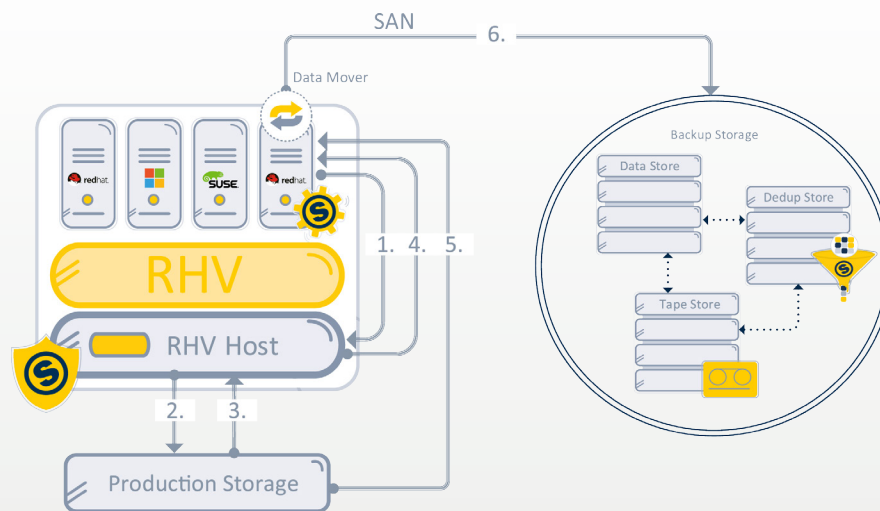
- [1.] The first step is restoring the VM configuration to the SEP DataMover VM, then the target VM is created.
- [2.] In the second step the disks on the target VM are created and attached to the SEP DataMover VM, allowing the data restore to occur.
- [3.] Once the restore is successful, the disks are detached from the SEP DataMover VM and attached to the target VM. The recreated VM can then be put into production.

LAN Free Backup

The second possible architecture for implementing native Red Hat Virtualization backups with SEP involves using the SAN storage connected directly to the RHV cluster. Although LAN Free Backups work in a similar way to LAN Based Backups, LAN Free Backups can be advantageous because large volumes of backup data are not traversing the LAN but rather are transferred across the dedicated storage interface.

The storage on the SAN uses the same architecture as a LAN Based Backup – there is an SEP DataMover VM within the RHV cluster and the SEP Backup Server may reside externally or be combined with the DataMover VM. The only difference with a LAN Free Backup, is that the storage is attached directly to the SEP DataMover VM using the Storage Area Network.

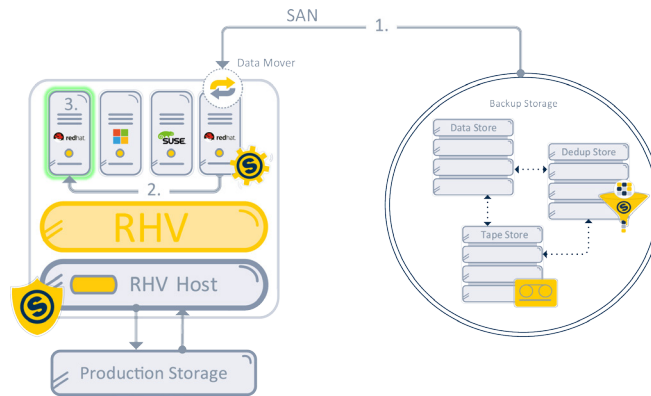
Figure 4 below illustrates the LAN Free Backup architecture and the basic data flow. It shows all backup and restore data flow over the SAN.



[Figure 4] The steps to back up a VM are as follows:

- [1.] A snapshot of the VM to be backed up is created. This enables consistent backups of the VMs running without interruption.
- [2.] [3.] [4.] Once the snapshot is successfully created, the VM configuration is exported from the snapshot and backed up as the first stream in the SaveSet.
- [5.] [6.] The VM disks contained in the snapshot are then attached to the DataMover VM using the backup API. Once attached, the data on the disks is backed up. After the backup has completed successfully, the snapshot is removed.

LAN Free Recovery

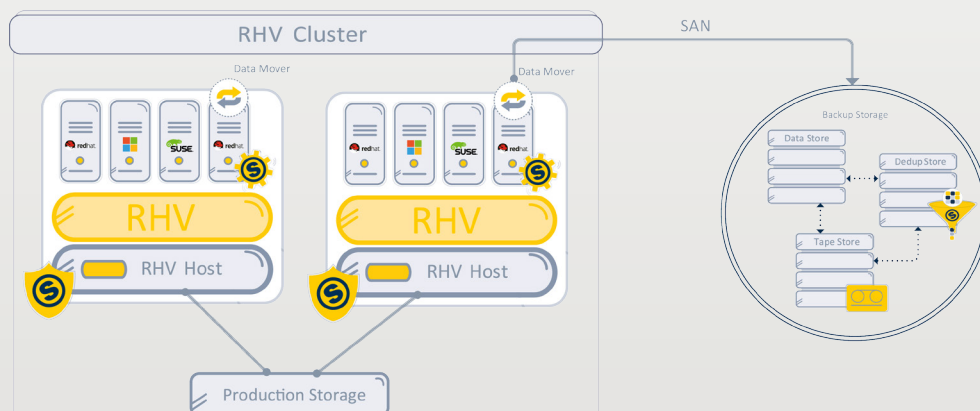


[Figure 5] The restoration process includes the following steps:

- [1.] The first step is restoring the VM configuration to the SEP DataMover VM, then the target VM is created.
- [2.] In the second step the disks on the target VM are created and attached to the SEP DataMover VM, allowing the data restore to occur.
- [3.] Once the restore is successful, the disks are detached from the SEP DataMover VM and attached to the target VM. The recreated VM can then be put into production.

RHV Cluster

SEP fully supports backup of all virtual machines in a high availability RHV Cluster environment. The SEP components communicate directly with the RHV host resource pool so in the event any backup tasks are not successful, they are able to fail-over to VMs on a different RHV cluster host. SEP is able to find any virtual guest anywhere on the cluster and initiate a backup.



Backup Deduplication and Site-to-Site Replication


The SEP backup server stores the backup data in SEP (disk based) DataStores, tapes, VTLs and tape loaders. As an efficient and cost effective way to significantly reduce the amount of required disk backup storage, SEP offers deduplication and replication technology.

With SEP, administrators have the ultimate flexibility to use storage hardware (including tape loaders) from any desired storage vendor.

Si3 Deduplication and Replication is a hardware independent technology, fully integrated with SEP, that enables highly-efficient, variable, inline, block-level, source and target side data deduplication along with fast, bandwidth-saving, site-to-site replication. All virtual machines backed up to a Si3 Deduplication DataStore are deduplicated and compressed to minimize the data footprint. Similar virtual guests are only stored once in the Si3 backup storage, making it an excellent choice for storing multiple VM backups that would otherwise take up vast amounts of storage space. Equally useful is the ability to perform highly-efficient, block-level replication of data to another Si3 Deduplication DataStore, thereby providing an efficient way of replicating data to an unlimited number of DR sites.

SEP replication is designed to consistently replicate data between SEP Remote Device Servers (RDS), including the Backup Server. The replication searches for new blocks of data on the source media pool and replicates those changes to a target media pool according to the defined schedule. Only the changed data blocks are sent over the network and backed up to the target server, which significantly reduces the amount of data transferred over the network. This feature enables administrators to set up scheduled restores of RHV virtual guests at a disaster recovery site as cloud standby systems. Note, that SEP replication can be performed from “n” SEP RDS Servers (or SEP Backup Servers) to “n” SEP RDS Servers (or SEP Backup Servers).

To configure a SEP Si3 Deduplication DataStore, simply select it from a dropdown menu and input the desired size.



Protection of Newly Deployed VMs

The protection of dynamically virtualized RHV environments requires special attention. Newly deployed VMs must be set up as new backup tasks. SEP offers a specially designed script to easily detect all new virtual guests of the RHV hypervisor in order to set up new tasks.

Database and Application Vendor Instructions

One advantage of SEP's integration with RHV is that any VM workload can be snapshotted, while it is online and in production and the backup will be a verifiable copy of the VM.

In addition, SEP users can add database modules to backup specific databases running on the RHV VMs to get native, online, granular backups of any SEP supported database.

Regardless of workloads, the deployment process is simple for all VMs whether they are a Windows based Exchange server or other groupware solution, or a Linux based VM running a database like SAP or Oracle.

Security

SEP is capable of providing many options to enhance security and assist with meeting compliance requirements:

- The backup data streams can be encrypted to protect access to the guest virtual machine backup from unauthorized users.
- Logs can be sent to a syslog server to maintain a single centralized location for tracking information.
- Notifications can be sent to email addresses to provide immediate alerts for completed and/or failed jobs.

These options are all part of the SEP notification system that can be used to restore backup information into a centralized audit system.

Conclusion

SEP's native Red Hat Virtualization backup provides a simple and efficient solution for RHV environments of all sizes. Data can either be transferred using an existing LAN or directly to a SAN. Configuration is quick and easy and when combined with other SEP features, such as deduplication and replication, users can maximize their storage space and be assured their data is safe, geographically dispersed, secure and easily accessible.

SEP features an intuitive GUI that can run locally on the backup server or remotely from a client, making management simple and efficient. The built-in reporting also makes monitoring backup consistency and job status equally simple and straightforward.

Restoring VMs is quick and easy and can also be automated to restore to other VM pools for warm standbys. Everything from the VM configuration to the virtual disks that hold data is restored so the VMs are back up and running as fast as the data can be transferred from the backup target to the RHV cluster. Backups can be replicated to other sites for DR purposes using several replication technologies, making SEP a core part of any multi-site disaster recovery strategy. SEP leverages its patented Multi-Streaming, deduplication and replication technologies to offer the fastest and most efficient backup and disaster recovery solution on the market for Red Hat Virtual environments of any size.