



Red Hat Virtualization 4.3

Product Guide

Introduction to Red Hat Virtualization 4.3

Red Hat Virtualization 4.3 Product Guide

Introduction to Red Hat Virtualization 4.3

Red Hat Virtualization Documentation Team
Red Hat Customer Content Services
rhev-docs@redhat.com

Legal Notice

Copyright © 2020 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document provides an introduction to Red Hat Virtualization.

Table of Contents

CHAPTER 1. INTRODUCTION TO RED HAT VIRTUALIZATION	3
1.1. RED HAT VIRTUALIZATION ARCHITECTURE	3
1.1.1. Self-Hosted Engine Architecture	3
1.1.2. Standalone Manager Architecture	4
1.2. RED HAT VIRTUALIZATION TERMINOLOGY	5
CHAPTER 2. RED HAT VIRTUALIZATION COMPONENTS	7
2.1. RED HAT VIRTUALIZATION MANAGER	7
2.2. HOSTS	7
2.3. STORAGE	7
2.4. DATA WAREHOUSE	8
2.5. METRICS STORE	8
2.6. NETWORKING IN RED HAT VIRTUALIZATION	9
CHAPTER 3. INSTALLING RED HAT VIRTUALIZATION	10
3.1. RED HAT VIRTUALIZATION INSTALLATION OVERVIEW	10
3.2. RED HAT VIRTUALIZATION INSTALLATION OPTIONS	10
3.2.1. Self-hosted engine (recommended)	10
3.2.2. Standalone Manager (manual)	10
CHAPTER 4. ACCESSING RED HAT VIRTUALIZATION	12
4.1. SUPPORTED INTERFACES FOR READ ACCESS	13
4.2. UNSUPPORTED INTERFACES	13

CHAPTER 1. INTRODUCTION TO RED HAT VIRTUALIZATION

Red Hat Virtualization is an enterprise-grade virtualization platform built on Red Hat Enterprise Linux. Virtualization allows users to easily provision new virtual servers and workstations, and provides more efficient use of physical server resources. With Red Hat Virtualization, you can manage your entire virtual infrastructure - including hosts, virtual machines, networks, storage, and users - from a centralized graphical user interface or REST API.

Table 1.1. Red Hat Virtualization Key Components

Component Name	Description
Red Hat Virtualization Manager	A service that provides a graphical user interface and a REST API to manage the resources in the environment. The Manager is installed on a physical or virtual machine running Red Hat Enterprise Linux.
Hosts	Red Hat Enterprise Linux hosts (RHEL hosts) and Red Hat Virtualization Hosts (image-based hypervisors) are the two supported types of host. Hosts use Kernel-based Virtual Machine (KVM) technology and provide resources used to run virtual machines.
Shared Storage	A storage service is used to store the data associated with virtual machines.
Data Warehouse	A service that collects configuration information and statistical data from the Manager.

For detailed technical information about Red Hat Virtualization, see the [Technical Reference](#).

1.1. RED HAT VIRTUALIZATION ARCHITECTURE

Red Hat Virtualization can be deployed as a self-hosted engine, or as a standalone Manager. A self-hosted engine is the recommended deployment option.

1.1.1. Self-Hosted Engine Architecture

The Red Hat Virtualization Manager runs as a virtual machine on self-hosted engine nodes (specialized hosts) in the same environment it manages. A self-hosted engine environment requires one less physical server, but requires more administrative overhead to deploy and manage. The Manager is highly available without external HA management.


The minimum setup of a self-hosted engine environment includes:

- One Red Hat Virtualization Manager virtual machine that is hosted on the self-hosted engine nodes. The RHV-M Appliance is used to automate the installation of a Red Hat Enterprise Linux 7 virtual machine, and the Manager on that virtual machine.
- A minimum of two self-hosted engine nodes for virtual machine high availability. You can use Red Hat Enterprise Linux hosts or Red Hat Virtualization Hosts (RHVH). VDSM (the host agent) runs on all hosts to facilitate communication with the Red Hat Virtualization Manager. The HA

services run on all self-hosted engine nodes to manage the high availability of the Manager virtual machine.

- One storage service, which can be hosted locally or on a remote server, depending on the storage type used. The storage service must be accessible to all hosts.

Figure 1.1. Self-Hosted Engine Red Hat Virtualization Architecture




1.1.2. Standalone Manager Architecture

The Red Hat Virtualization Manager runs on a physical server, or a virtual machine hosted in a separate virtualization environment. A standalone Manager is easier to deploy and manage, but requires an additional physical server. The Manager is only highly available when managed externally with a product such as Red Hat's High Availability Add-On.

The minimum setup for a standalone Manager environment includes:

- One Red Hat Virtualization Manager machine. The Manager is typically deployed on a physical server. However, it can also be deployed on a virtual machine, as long as that virtual machine is hosted in a separate environment. The Manager must run on Red Hat Enterprise Linux 7.
- A minimum of two hosts for virtual machine high availability. You can use Red Hat Enterprise Linux hosts or Red Hat Virtualization Hosts (RHVH). VDSM (the host agent) runs on all hosts to facilitate communication with the Red Hat Virtualization Manager.
- One storage service, which can be hosted locally or on a remote server, depending on the storage type used. The storage service must be accessible to all hosts.

Figure 1.2. Standalone Manager Red Hat Virtualization Architecture



1.2. RED HAT VIRTUALIZATION TERMINOLOGY

- **Cluster** - A cluster is a set of physical hosts that are treated as a resource pool for virtual machines. Hosts in a cluster share the same network infrastructure and storage. They form a migration domain within which virtual machines can be moved from host to host.
- **Data Center** - A data center is the highest level container for all physical and logical resources within a managed virtual environment. It is a collection of clusters, virtual machines, storage domains, and networks.
- **Events** - Alerts, warnings, and other notices about activities help the administrator to monitor the performance and status of resources.
- **HA Services** - The HA services include the **ovirt-ha-agent** service and the **ovirt-ha-broker** service. The HA services run on self-hosted engine nodes and manage the high availability of the Manager virtual machine.
- **High Availability** - High availability means that a virtual machine is automatically restarted if its process is interrupted, either on its original host or another host in the cluster. Highly available environments involve a small amount of downtime, but have a much lower cost than fault tolerance, which maintains two copies of each resource so that one can replace the other immediately in the event of a failure.
- **Host** - A host, or hypervisor, is a physical server that runs one or more virtual machines. Hosts are grouped into clusters. Virtual machines can be migrated from one host to another within a cluster.
- **Host Storage Manager (HSM)** - Any non-SPM host in the data center that can be used for data operations, such as moving a disk between storage domains. This prevents a bottleneck at the SPM host, which should be used for shorter metadata operations.

- **Logical Network** - A logical network is a logical representation of a physical network. Logical networks group network traffic and communication between the Manager, hosts, storage, and virtual machines.
- **Remote Viewer** - A graphical interface to connect to virtual machines over a network connection.
- **Self-Hosted Engine Node** - A self-hosted engine node is a host that has self-hosted engine packages installed so that it can host the Manager virtual machine. Regular hosts can also be attached to a self-hosted engine environment, but cannot host the Manager virtual machine.
- **Snapshot** - A snapshot is a view of a virtual machine's operating system and all its applications at a point in time. It can be used to save the settings of a virtual machine before an upgrade or before installing new applications. In case of problems, a snapshot can be used to restore the virtual machine to its original state.
- **Storage Domain** - A storage domain is a logical entity that contains a standalone image repository. Each storage domain is used to store virtual disks or ISO images, and for the import and export of virtual machine images.
- **Storage Pool Manager (SPM)**- The Storage Pool Manager (SPM) is a role assigned to one host in a data center. The SPM host has sole authority to make all metadata changes for the data center, such as the creation and removal of virtual disks.
- **Template** - A template is a model virtual machine with predefined settings. A virtual machine that is based on a particular template acquires the settings of the template. Using templates is the quickest way of creating a large number of virtual machines in a single step.
- **VDSM** - The host agent service running on the hosts, which communicates with the Red Hat Virtualization Manager. The service listens on TCP port 54321.
- **Virtual Machine** - A virtual machine is a virtual workstation or virtual server containing an operating system and a set of applications. Multiple identical virtual machines can be created in a **Pool**. Virtual machines are created, managed, or deleted by power users and accessed by users.
- **Virtual Machine Pool**- A virtual machine pool is a group of identical virtual machines that are available on demand by each group member. Virtual machine pools can be set up for different purposes. For example, one pool can be for the Marketing department, another for Research and Development, and so on.

CHAPTER 2. RED HAT VIRTUALIZATION COMPONENTS

2.1. RED HAT VIRTUALIZATION MANAGER

The Red Hat Virtualization Manager provides a graphical user interface and a RESTful API to manage the resources of the Red Hat Virtualization environment. In a standalone Manager environment, the Manager is installed on a Red Hat Enterprise Linux 7 physical machine, or a virtual machine hosted in a separate environment. In a self-hosted engine environment, the Manager is installed as a virtual machine hosted on self-hosted engine nodes in the same environment it manages.

Native high availability for the Manager is only available in a self-hosted engine environment. At least two self-hosted engine nodes are required for high availability.

See: [Administration Guide](#)

2.2. HOSTS

Hosts, also known as hypervisors, are the physical servers on which virtual machines run. Full virtualization is provided by using a loadable Linux kernel module called Kernel-based Virtual Machine (KVM). KVM can concurrently host multiple virtual machines running either Windows or Linux operating systems. Virtual machines run on the host machine and are managed remotely by the Red Hat Virtualization Manager.

Red Hat Virtualization supports two types of host: Red Hat Virtualization Host and Red Hat Enterprise Linux. You can use either or both in your Red Hat Virtualization environment, depending on your requirements.

Red Hat recommends that you install at least two hosts and attach them to the Red Hat Virtualization environment. If you attach only one host, you cannot access features such as migration and high availability.

Red Hat Virtualization Host (RHVH)

Red Hat Virtualization Host is a minimal operating system based on Red Hat Enterprise Linux, built for easy management, easy maintenance, and simple deployment. It is distributed as an ISO file from the Customer Portal and contains only the packages required for the machine to act as a host.

Red Hat Enterprise Linux

Red Hat Enterprise Linux servers with the appropriate subscriptions attached can be used as hosts. These hosts are more customizable than RHVH.

See: [Hosts](#) in the *Administration Guide*.

2.3. STORAGE

Setting up storage and attaching it to your Red Hat Virtualization environment is a prerequisite before you can start creating end-user virtual machines. Red Hat Virtualization has three types of storage domains; however, only the data domain is now fully supported:

- The data domain contains all the data associated with virtual machines. The data domain supports all storage types that are supported for use with Red Hat Virtualization. See [Storage Types](#) in the *Planning and Prerequisites Guide*. For information on uploading to a data domain, see [Uploading Images to a Data Storage Domain](#) in the *Administration Guide*.
- The ISO domain is a deprecated storage domain type that was used to store ISO files for

installing a virtual machine operating system or additional applications, such as the Windows guest agents and drivers. Virtual machine images can now be uploaded to data domains instead. For information on uploading to an ISO domain, see [The ISO Uploader Tool](#) in the *Administration Guide*.

- The export domain is a deprecated storage domain type that was used as a temporary storage repository for moving images between data centers and Red Hat Virtualization environments. This is now done by importing data storage domains.

The ISO and export domains only support file-based storage types (NFS, POSIX, or GlusterFS). The ISO domain supports local storage when used in a local storage data center.

See: [Storage](#) in the *Administration Guide*.

2.4. DATA WAREHOUSE

The Red Hat Virtualization Manager includes a data warehouse that collects monitoring data about hosts, virtual machines, and storage. Data Warehouse, which includes a database and a service, must be installed and configured along with the Manager setup, either on the same machine or on a separate server.

The Red Hat Virtualization installation creates two databases:

- The Manager database (**engine**) is the primary data store used by the Red Hat Virtualization Manager. Information about the virtualization environment like its state, configuration, and performance are stored in this database.
- The Data Warehouse database (**ovirt_engine_history**) contains configuration information and statistical data which is collated over time from the Manager database. The configuration data in the Manager database is examined every minute, and changes are replicated to the Data Warehouse database. Tracking the changes to the database provides information on the objects in the database. This enables you to analyze and enhance the performance of your Red Hat Virtualization environment and resolve difficulties.

To calculate an estimate of the space and resources the **ovirt_engine_history** database will use, use the [RHV Manager History Database Size Calculator](#) tool. The estimate is based on the number of entities and the length of time you have chosen to retain the history records.

See: [Data Warehouse Guide](#)

2.5. METRICS STORE

Metrics Store collects logs and metrics from Red Hat Virtualization. The data is transferred from Red Hat Virtualization to OpenShift where it is stored and aggregated in Elasticsearch and saved in indexes. The data can then be analyzed and visualized in Kibana.

- [Elasticsearch](#) is a distributed, RESTful search and analytics engine that lets you perform and combine many types of searches.
- [Kibana](#) is an open source analytics and visualization platform designed to work with Elasticsearch. You can easily perform advanced data analysis and visualize your data in a variety of charts and tables.

See: [Metrics Store User Guide](#) and [Metrics Store Installation Guide](#)

2.6. NETWORKING IN RED HAT VIRTUALIZATION

Operations such as storage, host management, user connections, and virtual machine connectivity all rely on a well-planned and well-configured network to deliver optimal performance. Setting up networking is a vital prerequisite for a Red Hat Virtualization environment. Planning for your projected networking requirements and implementing your network accordingly is much simpler than discovering your networking requirements through use and altering your network configuration retroactively.

Red Hat Virtualization separates network traffic by defining logical networks. Logical networks define the path that a selected network traffic type must take through the network. They are created to isolate network traffic by functionality or to virtualize a physical topology.

The **ovirtmgmt** logical network is created by default and labeled as the **Management** network. The **ovirtmgmt** logical network is intended for management traffic between the Red Hat Virtualization Manager and hosts. You can define additional logical networks to segregate:

- General virtual machine traffic
- Storage-related traffic (such as NFS or iSCSI)
- Virtual machine migration traffic
- Virtual machine display traffic
- Gluster storage traffic

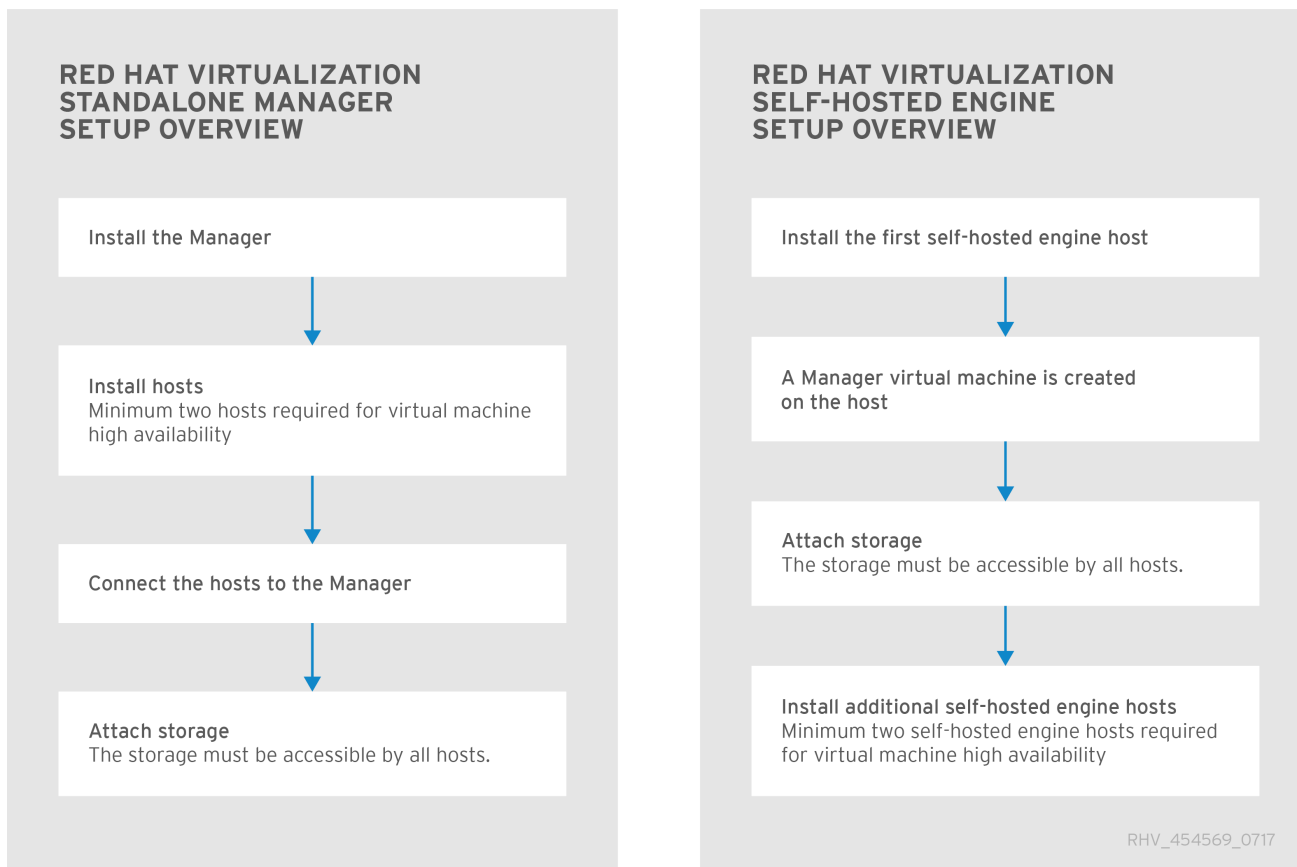
See: [Logical Networks](#) in the *Administration Guide*.

CHAPTER 3. INSTALLING RED HAT VIRTUALIZATION

3.1. RED HAT VIRTUALIZATION INSTALLATION OVERVIEW

The following diagram outlines the setup procedure for a standalone Manager deployment and a self-hosted engine deployment.

Figure 3.1. Setup Overview of Red Hat Virtualization



3.2. RED HAT VIRTUALIZATION INSTALLATION OPTIONS

3.2.1. Self-hosted engine (recommended)

You can install Red Hat Virtualization as a self-hosted engine using either the [Cockpit web interface](#) or the command line. Cockpit is the recommended method.

Because the Manager installation is automated, the Manager database (**engine**) and Data Warehouse service and database (**ovirt-engine-history**) can only be installed locally. If you want to host any databases or services on remote machines, you must migrate them after the installation is complete. Instructions for this are included in both of the following guides.

- To install a self-hosted engine using Cockpit (recommended), see [Installing Red Hat Virtualization as a self-hosted engine using the Cockpit web interface](#).
- To install a self-hosted engine using the command line, see [Installing Red Hat Virtualization as a self-hosted engine using the command line](#).

3.2.2. Standalone Manager (manual)

You can only install a standalone Manager using the command line.

During installation, you can specify whether to install the Manager database (**engine**) and Data Warehouse service and database (**ovirt-engine-history**) locally or remotely. You do not have to choose the same option for both; one database can be hosted on a remote machine, while the other is hosted on the Manager machine.

Local databases can be manually configured, or automatically configured by the installation script. A remote Manager database must be manually configured. A remote Data Warehouse database can be configured automatically only if the Data Warehouse service is installed on the same machine (if not, the database must be configured manually).

- To install a standalone Manager with local databases, see [Installing Red Hat Virtualization as a standalone Manager with local databases](#). This guide assumes that all databases and services will be installed on the same machine as the Manager.
- To install a standalone Manager with remote databases, see [Installing Red Hat Virtualization as a standalone Manager with remote databases](#). This guide assumes that all databases and services will be installed on one additional machine separate from the Manager.

Mixing local and remote databases

Installing one database on the Manager machine and the other database remotely, and the various combinations thereof, is not specifically documented. For example:

- The Manager and Manager database are installed on the *Manager machine*, and the Data Warehouse service and database are installed on *Server 1*. In this case, you must configure the Manager locally, then install Data Warehouse separately.
- The Manager and the Data Warehouse service are installed on the *Manager machine*, and the Manager database and Data Warehouse database are installed on *Server 1*. In this case, you must manually configure the Manager and Data Warehouse databases, then configure the Manager and the Data Warehouse service.
- The Manager, Manager database, and Data Warehouse database are installed on the *Manager machine*, and the Data Warehouse service is installed on *Server 1*. In this case, you must manually configure the Data Warehouse database, then configure the Manager, then install the Data Warehouse service.
- The Manager and the Data Warehouse service are installed on the *Manager machine*, the Manager database is installed on *Server 1*, and the Data Warehouse database is installed on *Server 2*. In this case, you must manually configure the Manager and Data Warehouse databases, then configure the Manager and the Data Warehouse service.
- The Manager is installed on the *Manager machine*, the Manager database is installed on *Server 1*, and the Data Warehouse service and database are installed on *Server 2*. In this case, you must manually configure the Manager database, then configure the Manager, then install Data Warehouse.

CHAPTER 4. ACCESSING RED HAT VIRTUALIZATION

Red Hat Virtualization exposes a number of interfaces for interacting with the components of the virtualization environment. Many of these interfaces are fully supported. Some, however, are supported only for read access or only when your use of them has been explicitly requested by Red Hat Support.

=== Supported Interfaces for Read and Write Access

Direct interaction with these interfaces is supported and encouraged for both read and write access:

Administration Portal

The Administration Portal is a graphical user interface provided by the Red Hat Virtualization Manager. It can be used to manage all the administrative resources in the environment and can be accessed by any supported web browsers.

See: [Administration Guide](#)

VM Portal

The VM Portal is a graphical user interface provided by the Red Hat Virtualization Manager. It has limited permissions for managing virtual machine resources and is targeted at end users.

See: [Introduction to the VM Portal](#)

Cockpit

In Red Hat Virtualization, the Cockpit web interface can be used to deploy a self-hosted engine environment, as well as perform other administrative tasks on a host. It is available by default on Red Hat Virtualization Hosts, and can be installed on Red Hat Enterprise Linux hosts.

See:

- [Installing Cockpit on Red Hat Enterprise Linux hosts](#) .
- [Installing Red Hat Virtualization as a self-hosted engine using the Cockpit web interface](#) .

REST API

The Red Hat Virtualization REST API provides a software interface for querying and modifying the Red Hat Virtualization environment. The REST API can be used by any programming language that supports HTTP actions.

See: [REST API Guide](#)

Software Development Kit (SDK)

The Python, Java, and Ruby SDKs are fully supported interfaces for interacting with the Red Hat Virtualization Manager.

See:

- [Python SDK Guide](#)
- [Java SDK Guide](#)
- [Ruby SDK Guide](#)

Ansible

Ansible provides modules to automate post-installation tasks on Red Hat Virtualization.

See: [Automating Configuration Tasks using Ansible](#) in the *Administration Guide*.

Self-Hosted Engine Command Line Utility

The **hosted-engine** command is used to perform administrative tasks on the Manager virtual machine in self-hosted engine environments.

See: [Administering the Manager Virtual Machine](#) in the *Administration Guide*.

VDSM Hooks

VDSM hooks trigger modifications to virtual machines, based on custom properties specified in the Administration Portal.

See: [VDSM and Hooks](#) in the *Administration Guide*.

4.1. SUPPORTED INTERFACES FOR READ ACCESS

Direct interaction with these interfaces is supported and encouraged only for read access. Use of these interfaces for write access is not supported unless explicitly requested by Red Hat Support. **Red Hat Virtualization Manager History Database::** Read access to the Red Hat Virtualization Manager history (**ovirt_engine_history**) database using the database views specified in the *Data Warehouse Guide* is supported. Write access is **not** supported.

Libvirt on Hosts

Read access to **libvirt** using the **virsh -r** command is a supported method of interacting with virtualization hosts. Write access is **not** supported.

4.2. UNSUPPORTED INTERFACES

Direct interaction with these interfaces is **not** supported unless your use of them is explicitly requested by Red Hat Support :

The **vds**m-client Command

Use of the **vds**m-client command to interact with virtualization hosts is **not** supported unless explicitly requested by Red Hat Support . // review above

Red Hat Virtualization Manager Database

Direct access to, and manipulation of, the Red Hat Virtualization Manager (**engine**) database is **not** supported unless explicitly requested by Red Hat Support . // review above



IMPORTANT

Red Hat Support will not debug user-created scripts or hooks except where it can be demonstrated that there is an issue with the interface being used rather than the user-created script itself. For more general information about Red Hat's support policies see [Production Support Scope of Coverage](#).