



Red Hat Virtualization 4.0

Administration Guide

Administration Tasks in Red Hat Virtualization

Red Hat Virtualization 4.0 Administration Guide

Administration Tasks in Red Hat Virtualization

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

Legal Notice

Copyright © 2018 Red Hat.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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, 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 Software Collections 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 book contains information and procedures relevant to Red Hat Virtualization administrators.

Table of Contents

PART I. ADMINISTERING AND MAINTAINING THE RED HAT VIRTUALIZATION ENVIRONMENT	6
CHAPTER 1. GLOBAL CONFIGURATION	7
1.1. ROLES	7
1.2. SYSTEM PERMISSIONS	11
1.3. SCHEDULING POLICIES	16
1.4. INSTANCE TYPES	23
1.5. MAC ADDRESS POOLS	26
CHAPTER 2. DASHBOARD	30
2.1. PREREQUISITES	30
2.2. GLOBAL INVENTORY	30
2.3. GLOBAL UTILIZATION	32
2.4. CLUSTER UTILIZATION	33
2.5. STORAGE UTILIZATION	34
PART II. ADMINISTERING THE RESOURCES	35
CHAPTER 3. QUALITY OF SERVICE	36
3.1. STORAGE QUALITY OF SERVICE	36
3.2. VIRTUAL MACHINE NETWORK QUALITY OF SERVICE	37
3.3. HOST NETWORK QUALITY OF SERVICE	39
3.4. CPU QUALITY OF SERVICE	41
CHAPTER 4. DATA CENTERS	42
4.1. INTRODUCTION TO DATA CENTERS	42
4.2. THE STORAGE POOL MANAGER	43
4.3. SPM PRIORITY	43
4.4. USING THE EVENTS TAB TO IDENTIFY PROBLEM OBJECTS IN DATA CENTERS	44
4.5. DATA CENTER TASKS	44
4.6. DATA CENTERS AND STORAGE DOMAINS	48
4.7. DATA CENTERS AND PERMISSIONS	50
CHAPTER 5. CLUSTERS	53
5.1. INTRODUCTION TO CLUSTERS	53
5.2. CLUSTER TASKS	53
5.3. CLUSTERS AND PERMISSIONS	77
CHAPTER 6. LOGICAL NETWORKS	80
6.1. LOGICAL NETWORK TASKS	80
6.2. VIRTUAL NETWORK INTERFACE CARDS	88
6.3. EXTERNAL PROVIDER NETWORKS	94
6.4. LOGICAL NETWORKS AND PERMISSIONS	98
6.5. HOSTS AND NETWORKING	100
CHAPTER 7. HOSTS	112
7.1. INTRODUCTION TO HOSTS	112
7.2. RED HAT VIRTUALIZATION HOST	113
7.3. RED HAT ENTERPRISE LINUX HOSTS	113
7.4. SATELLITE HOST PROVIDER HOSTS	114
7.5. HOST TASKS	114
7.6. HOST RESILIENCE	134
7.7. HOSTS AND PERMISSIONS	145

CHAPTER 8. STORAGE	148
8.1. UNDERSTANDING STORAGE DOMAINS	149
8.2. PREPARING AND ADDING NFS STORAGE	149
8.3. PREPARING AND ADDING LOCAL STORAGE	152
8.4. ADDING POSIX COMPLIANT FILE SYSTEM STORAGE	154
8.5. ADDING BLOCK STORAGE	156
8.6. IMPORTING EXISTING STORAGE DOMAINS	164
8.7. STORAGE TASKS	171
8.8. STORAGE AND PERMISSIONS	177
CHAPTER 9. WORKING WITH RED HAT GLUSTER STORAGE	180
9.1. RED HAT GLUSTER STORAGE NODES	180
9.2. USING RED HAT GLUSTER STORAGE AS A STORAGE DOMAIN	181
9.3. CLUSTERS AND GLUSTER HOOKS	190
CHAPTER 10. POOLS	196
10.1. INTRODUCTION TO VIRTUAL MACHINE POOLS	196
10.2. VIRTUAL MACHINE POOL TASKS	196
10.3. POOLS AND PERMISSIONS	208
10.4. TRUSTED COMPUTE POOLS	210
CHAPTER 11. VIRTUAL DISKS	213
11.1. UNDERSTANDING VIRTUAL MACHINE STORAGE	213
11.2. UNDERSTANDING VIRTUAL DISKS	213
11.3. SETTINGS TO WIPE VIRTUAL DISKS AFTER DELETION	215
11.4. SHAREABLE DISKS IN RED HAT VIRTUALIZATION	216
11.5. READ ONLY DISKS IN RED HAT VIRTUALIZATION	216
11.6. VIRTUAL DISK TASKS	217
11.7. VIRTUAL DISKS AND PERMISSIONS	229
CHAPTER 12. EXTERNAL PROVIDERS	232
12.1. INTRODUCTION TO EXTERNAL PROVIDERS IN RED HAT VIRTUALIZATION	232
12.2. ADDING EXTERNAL PROVIDERS	233
12.3. EDITING EXTERNAL PROVIDERS	249
12.4. REMOVING EXTERNAL PROVIDERS	249
PART III. ADMINISTERING THE ENVIRONMENT	250
CHAPTER 13. BACKUPS AND MIGRATION	251
13.1. BACKING UP AND RESTORING THE RED HAT VIRTUALIZATION MANAGER	251
13.2. BACKING UP AND RESTORING VIRTUAL MACHINES USING THE BACKUP AND RESTORE API	258
CHAPTER 14. ERRATA MANAGEMENT WITH RED HAT SATELLITE	262
CHAPTER 15. USERS AND ROLES	264
15.1. INTRODUCTION TO USERS	264
15.2. INTRODUCTION TO DIRECTORY SERVERS	264
15.3. CONFIGURING AN EXTERNAL LDAP PROVIDER	265
15.4. CONFIGURING LDAP AND KERBEROS FOR SINGLE SIGN-ON	276
15.5. USER AUTHORIZATION	282
15.6. ADMINISTERING USER TASKS FROM THE ADMINISTRATION PORTAL	282
15.7. ADMINISTERING USER TASKS FROM THE COMMAND LINE	285
15.8. CONFIGURING ADDITIONAL LOCAL DOMAINS	290
CHAPTER 16. QUOTAS AND SERVICE LEVEL AGREEMENT POLICY	291
16.1. INTRODUCTION TO QUOTA	291

16.2. SHARED QUOTA AND INDIVIDUALLY DEFINED QUOTA	292
16.3. QUOTA ACCOUNTING	292
16.4. ENABLING AND CHANGING A QUOTA MODE IN A DATA CENTER	293
16.5. CREATING A NEW QUOTA POLICY	293
16.6. EXPLANATION OF QUOTA THRESHOLD SETTINGS	294
16.7. ASSIGNING A QUOTA TO AN OBJECT	295
16.8. USING QUOTA TO LIMIT RESOURCES BY USER	296
16.9. EDITING QUOTAS	296
16.10. REMOVING QUOTAS	297
16.11. SERVICE LEVEL AGREEMENT POLICY ENFORCEMENT	297
CHAPTER 17. EVENT NOTIFICATIONS	299
17.1. CONFIGURING EVENT NOTIFICATIONS IN THE ADMINISTRATION PORTAL	299
17.2. CANCELING EVENT NOTIFICATIONS IN THE ADMINISTRATION PORTAL	301
17.3. PARAMETERS FOR EVENT NOTIFICATIONS IN OVIRT-ENGINE-NOTIFIER.CONF	301
17.4. CONFIGURING THE RED HAT VIRTUALIZATION MANAGER TO SEND SNMP TRAPS	305
CHAPTER 18. UTILITIES	308
18.1. THE OVIRT ENGINE RENAME TOOL	308
18.2. THE ENGINE CONFIGURATION TOOL	310
18.3. THE IMAGE UPLOADER TOOL	311
18.4. THE USB FILTER EDITOR	315
18.5. THE LOG COLLECTOR TOOL	321
18.6. THE ISO UPLOADER TOOL	325
18.7. THE ENGINE VACUUM TOOL	329
PART IV. GATHERING INFORMATION ABOUT THE ENVIRONMENT	331
CHAPTER 19. LOG FILES	332
19.1. RED HAT VIRTUALIZATION MANAGER INSTALLATION LOG FILES	332
19.2. RED HAT VIRTUALIZATION MANAGER LOG FILES	332
19.3. SPICE LOG FILES	333
19.4. RED HAT VIRTUALIZATION HOST LOG FILES	334
19.5. SETTING UP A VIRTUALIZATION HOST LOGGING SERVER	335
CHAPTER 20. PROXIES	337
20.1. SPICE PROXY	337
20.2. SQUID PROXY	339
20.3. WEBSOCKET PROXY	342
APPENDIX A. VDSM AND HOOKS	344
A.1. VDSM	344
A.2. VDSM HOOKS	344
A.3. EXTENDING VDSM WITH HOOKS	344
A.4. SUPPORTED VDSM EVENTS	344
A.5. THE VDSM HOOK ENVIRONMENT	347
A.6. THE VDSM HOOK DOMAIN XML OBJECT	347
A.7. DEFINING CUSTOM PROPERTIES	348
A.8. SETTING VIRTUAL MACHINE CUSTOM PROPERTIES	350
A.9. EVALUATING VIRTUAL MACHINE CUSTOM PROPERTIES IN A VDSM HOOK	350
A.10. USING THE VDSM HOOKING MODULE	350
A.11. VDSM HOOK EXECUTION	351
A.12. VDSM HOOK RETURN CODES	352
A.13. VDSM HOOK EXAMPLES	352

APPENDIX B. CUSTOM NETWORK PROPERTIES	355
B.1. EXPLANATION OF BRIDGE_OPTS PARAMETERS	355
B.2. HOW TO SET UP RED HAT VIRTUALIZATION MANAGER TO USE ETHTOOL	357
B.3. HOW TO SET UP RED HAT VIRTUALIZATION MANAGER TO USE FCOE	358
APPENDIX C. RED HAT VIRTUALIZATION USER INTERFACE PLUGINS	360
C.1. RED HAT VIRTUALIZATION USER INTERFACE PLUG-INS	360
C.2. RED HAT VIRTUALIZATION USER INTERFACE PLUGIN LIFECYCLE	360
C.3. USER INTERFACE PLUGIN-RELATED FILES AND THEIR LOCATIONS	362
C.4. EXAMPLE USER INTERFACE PLUG-IN DEPLOYMENT	362
C.5. USING RED HAT SUPPORT PLUG-IN	363
APPENDIX D. RED HAT VIRTUALIZATION AND SSL	368
D.1. REPLACING THE RED HAT VIRTUALIZATION MANAGER SSL/TLS CERTIFICATE	368
D.2. SETTING UP SSL OR TLS CONNECTIONS BETWEEN THE MANAGER AND AN LDAP SERVER	370
APPENDIX E. USING SEARCH, BOOKMARKS, AND TAGS	372
E.1. SEARCHES	372
E.2. BOOKMARKS	392
E.3. TAGS	394
APPENDIX F. BRANDING	396
F.1. BRANDING	396
APPENDIX G. SYSTEM ACCOUNTS	399
G.1. SYSTEM ACCOUNTS	399

PART I. ADMINISTERING AND MAINTAINING THE RED HAT VIRTUALIZATION ENVIRONMENT

The Red Hat Virtualization environment requires an administrator to keep it running. As an administrator, your tasks include:

- Managing physical and virtual resources such as hosts and virtual machines. This includes upgrading and adding hosts, importing domains, converting virtual machines created on foreign hypervisors, and managing virtual machine pools.
- Monitoring the overall system resources for potential problems such as extreme load on one of the hosts, insufficient memory or disk space, and taking any necessary actions (such as migrating virtual machines to other hosts to lessen the load or freeing resources by shutting down machines).
- Responding to the new requirements of virtual machines (for example, upgrading the operating system or allocating more memory).
- Managing customized object properties using tags.
- Managing searches saved as public bookmarks.
- Managing user setup and setting permission levels.
- Troubleshooting for specific users or virtual machines for overall system functionality.
- Generating general and specific reports.

CHAPTER 1. GLOBAL CONFIGURATION

Accessed from the header bar in the Administration Portal, the **Configure** window allows you to configure a number of global resources for your Red Hat Virtualization environment, such as users, roles, system permissions, scheduling policies, instance types, and MAC address pools. This window allows you to customize the way in which users interact with resources in the environment, and provides a central location for configuring options that can be applied to multiple clusters.




Figure 1.1. Accessing the Configure window

1.1. ROLES

Roles are predefined sets of privileges that can be configured from Red Hat Virtualization Manager. Roles provide access and management permissions to different levels of resources in the data center, and to specific physical and virtual resources.

With multilevel administration, any permissions which apply to a container object also apply to all individual objects within that container. For example, when a host administrator role is assigned to a user on a specific host, the user gains permissions to perform any of the available host operations, but only on the assigned host. However, if the host administrator role is assigned to a user on a data center, the user gains permissions to perform host operations on all hosts within the cluster of the data center.

1.1.1. Creating a New Role

If the role you require is not on Red Hat Virtualization's default list of roles, you can create a new role and customize it to suit your purposes.

Procedure 1.1. Creating a New Role

1. On the header bar, click the **Configure** button to open the **Configure** window. The window shows a list of default User and Administrator roles, and any custom roles.
2. Click **New**. The **New Role** dialog box displays.




Figure 1.2. The New Role Dialog

3. Enter the **Name** and **Description** of the new role.
4. Select either **Admin** or **User** as the **Account Type**.
5. Use the **Expand All** or **Collapse All** buttons to view more or fewer of the permissions for the listed objects in the **Check Boxes to Allow Action** list. You can also expand or collapse the options for each object.
6. For each of the objects, select or clear the actions you wish to permit or deny for the role you are setting up.
7. Click **OK** to apply the changes you have made. The new role displays on the list of roles.

1.1.2. Editing or Copying a Role

You can change the settings for roles you have created, but you cannot change default roles. To change default roles, clone and modify them to suit your requirements.

Procedure 1.2. Editing or Copying a Role

1. On the header bar, click the **Configure** button to open the **Configure** window. The window shows a list of default User and Administrator roles, and any custom roles.
2. Select the role you wish to change. Click **Edit** to open the **Edit Role** window, or click **Copy** to open the **Copy Role** window.
3. If necessary, edit the **Name** and **Description** of the role.
4. Use the **Expand All** or **Collapse All** buttons to view more or fewer of the permissions for the listed objects. You can also expand or collapse the options for each object.

5. For each of the objects, select or clear the actions you wish to permit or deny for the role you are editing.
6. Click **OK** to apply the changes you have made.

1.1.3. User Role and Authorization Examples

The following examples illustrate how to apply authorization controls for various scenarios, using the different features of the authorization system described in this chapter.

Example 1.1. Cluster Permissions

Sarah is the system administrator for the accounts department of a company. All the virtual resources for her department are organized under a Red Hat Virtualization **cluster** called **Accounts**. She is assigned the **ClusterAdmin** role on the accounts cluster. This enables her to manage all virtual machines in the cluster, since the virtual machines are child objects of the cluster. Managing the virtual machines includes editing, adding, or removing virtual resources such as disks, and taking snapshots. It does not allow her to manage any resources outside this cluster. Because **ClusterAdmin** is an administrator role, it allows her to use the Administration Portal to manage these resources, but does not give her any access via the User Portal.

Example 1.2. VM PowerUser Permissions

John is a software developer in the accounts department. He uses virtual machines to build and test his software. Sarah has created a virtual desktop called **johndesktop** for him. John is assigned the **UserVmManager** role on the **johndesktop** virtual machine. This allows him to access this single virtual machine using the User Portal. Because he has **UserVmManager** permissions, he can modify the virtual machine and add resources to it, such as new virtual disks. Because **UserVmManager** is a user role, it does not allow him to use the Administration Portal.

Example 1.3. Data Center Power User Role Permissions

Penelope is an office manager. In addition to her own responsibilities, she occasionally helps the HR manager with recruitment tasks, such as scheduling interviews and following up on reference checks. As per corporate policy, Penelope needs to use a particular application for recruitment tasks.

While Penelope has her own machine for office management tasks, she wants to create a separate virtual machine to run the recruitment application. She is assigned **PowerUserRole** permissions for the data center in which her new virtual machine will reside. This is because to create a new virtual machine, she needs to make changes to several components within the data center, including creating the virtual disk image in the storage domain.

Note that this is not the same as assigning **DataCenterAdmin** privileges to Penelope. As a PowerUser for a data center, Penelope can log in to the User Portal and perform virtual machine-specific actions on virtual machines within the data center. She cannot perform data center-level operations such as attaching hosts or storage to a data center.

Example 1.4. Network Administrator Permissions

Chris works as the network administrator in the IT department. Her day-to-day responsibilities include creating, manipulating, and removing networks in the department's Red Hat Virtualization environment. For her role, she requires administrative privileges on the resources and on the networks of each resource. For example, if Chris has **NetworkAdmin** privileges on the IT department's data center, she can add and remove networks in the data center, and attach and detach networks for all virtual machines belonging to the data center.

In addition to managing the networks of the company's virtualized infrastructure, Chris also has a junior network administrator reporting to her. The junior network administrator, Pat, is managing a smaller virtualized environment for the company's internal training department. Chris has assigned Pat **VnicProfileUser** permissions and **UserVmManager** permissions for the virtual machines used by the internal training department. With these permissions, Pat can perform simple administrative tasks such as adding network interfaces onto virtual machines in the **Extended** tab of the User Portal. However, he does not have permissions to alter the networks for the hosts on which the virtual machines run, or the networks on the data center to which the virtual machines belong.

Example 1.5. Custom Role Permissions

Rachel works in the IT department, and is responsible for managing user accounts in Red Hat Virtualization. She needs permission to add user accounts and assign them the appropriate roles and permissions. She does not use any virtual machines herself, and should not have access to administration of hosts, virtual machines, clusters or data centers. There is no built-in role which provides her with this specific set of permissions. A custom role must be created to define the set of permissions appropriate to Rachel's position.

The screenshot shows a 'New Role' dialog box with the following configuration:

- Name:** UserManager
- Description:** (empty)
- Account Type:** Admin (selected), User (unselected)
- Check Boxes to Allow Action:**
 - System:**
 - Configure System:
 - Manipulate Users:
 - Manipulate Permissions:
 - Add users and groups from directory while adding permissions:
 - Manipulate Roles:
 - Login Permissions:
 - Tag management Permissions:
 - Bookmark management Permissions:
 - Event notification management Permissions:
 - Audit Log management Permissions:
 - Generic Configuration:
 - Data Center:
 - Network:

Buttons at the bottom: OK, Reset, Cancel.

Figure 1.3. UserManager Custom Role

The **UserManager** custom role shown above allows manipulation of users, permissions and roles. These actions are organized under **System** - the top level object of the hierarchy shown in [Figure 1.3, "UserManager Custom Role"](#). This means they apply to all other objects in the system. The role is set to have an **Account Type** of **Admin**. This means that when she is assigned this role, Rachel can only use the Administration Portal, not the User Portal.

1.2. SYSTEM PERMISSIONS

Permissions enable users to perform actions on objects, where objects are either individual objects or container objects.




Figure 1.4. Permissions & Roles

Any permissions that apply to a container object also apply to all members of that container. The following diagram depicts the hierarchy of objects in the system.

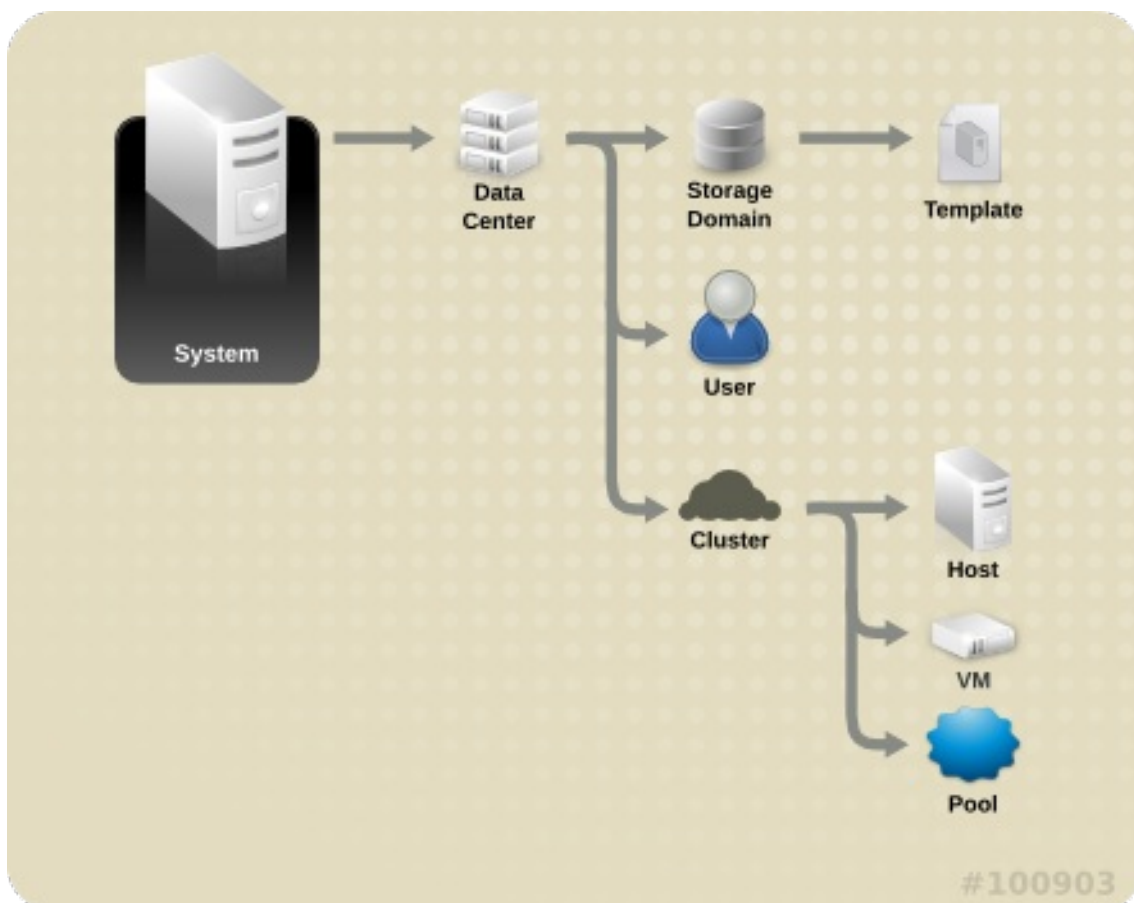


Figure 1.5. Red Hat Virtualization Object Hierarchy

1.2.1. User Properties

Roles and permissions are the properties of the user. Roles are predefined sets of privileges that permit access to different levels of physical and virtual resources. Multilevel

administration provides a finely grained hierarchy of permissions. For example, a data center administrator has permissions to manage all objects in the data center, while a host administrator has system administrator permissions to a single physical host. A user can have permissions to use a single virtual machine but not make any changes to the virtual machine configurations, while another user can be assigned system permissions to a virtual machine.

1.2.2. User and Administrator Roles

Red Hat Virtualization provides a range of pre-configured roles, from an administrator with system-wide permissions to an end user with access to a single virtual machine. While you cannot change or remove the default roles, you can clone and customize them, or create new roles according to your requirements. There are two types of roles:

- **Administrator Role:** Allows access to the *Administration Portal* for managing physical and virtual resources. An administrator role confers permissions for actions to be performed in the User Portal; however, it has no bearing on what a user can see in the User Portal.
- **User Role:** Allows access to the *User Portal* for managing and accessing virtual machines and templates. A user role determines what a user can see in the User Portal. Permissions granted to a user with an administrator role are reflected in the actions available to that user in the User Portal.

For example, if you have an **administrator** role on a cluster, you can manage all virtual machines in the cluster using the *Administration Portal*. However, you cannot access any of these virtual machines in the *User Portal*; this requires a **user** role.

1.2.3. User Roles Explained

The table below describes basic user roles which confer permissions to access and configure virtual machines in the User Portal.

Table 1.1. Red Hat Virtualization User Roles - Basic

Role	Privileges	Notes
UserRole	Can access and use virtual machines and pools.	Can log in to the User Portal, use assigned virtual machines and pools, view virtual machine state and details.
PowerUserRole	Can create and manage virtual machines and templates.	Apply this role to a user for the whole environment with the Configure window, or for specific data centers or clusters. For example, if a PowerUserRole is applied on a data center level, the PowerUser can create virtual machines and templates in the data center.

Role	Privileges	Notes
UserVmManager	System administrator of a virtual machine.	Can manage virtual machines and create and use snapshots. A user who creates a virtual machine in the User Portal is automatically assigned the UserVmManager role on the machine.

The table below describes advanced user roles which allow you to do more fine tuning of permissions for resources in the User Portal.

Table 1.2. Red Hat Virtualization User Roles - Advanced

Role	Privileges	Notes
UserTemplateBasedVm	Limited privileges to only use Templates.	Can use templates to create virtual machines.
DiskOperator	Virtual disk user.	Can use, view and edit virtual disks. Inherits permissions to use the virtual machine to which the virtual disk is attached.
VmCreator	Can create virtual machines in the User Portal.	This role is not applied to a specific virtual machine; apply this role to a user for the whole environment with the Configure window. Alternatively apply this role for specific data centers or clusters. When applying this role to a cluster, you must also apply the DiskCreator role on an entire data center, or on specific storage domains.
TemplateCreator	Can create, edit, manage and remove virtual machine templates within assigned resources.	This role is not applied to a specific template; apply this role to a user for the whole environment with the Configure window. Alternatively apply this role for specific data centers, clusters, or storage domains.

Role	Privileges	Notes
DiskCreator	Can create, edit, manage and remove virtual disks within assigned clusters or data centers.	This role is not applied to a specific virtual disk; apply this role to a user for the whole environment with the Configure window. Alternatively apply this role for specific data centers or storage domains.
TemplateOwner	Can edit and delete the template, assign and manage user permissions for the template.	This role is automatically assigned to the user who creates a template. Other users who do not have TemplateOwner permissions on a template cannot view or use the template.
VnicProfileUser	Logical network and network interface user for virtual machine and template.	Can attach or detach network interfaces from specific logical networks.

1.2.4. Administrator Roles Explained

The table below describes basic administrator roles which confer permissions to access and configure resources in the Administration Portal.

Table 1.3. Red Hat Virtualization System Administrator Roles - Basic

Role	Privileges	Notes
SuperUser	System Administrator of the Red Hat Virtualization environment.	Has full permissions across all objects and levels, can manage all objects across all data centers.
ClusterAdmin	Cluster Administrator.	Possesses administrative permissions for all objects underneath a specific cluster.
DataCenterAdmin	Data Center Administrator.	Possesses administrative permissions for all objects underneath a specific data center except for storage.



IMPORTANT

Do not use the administrative user for the directory server as the Red Hat Virtualization administrative user. Create a user in the directory server specifically for use as the Red Hat Virtualization administrative user.

The table below describes advanced administrator roles which allow you to do more fine tuning of permissions for resources in the Administration Portal.

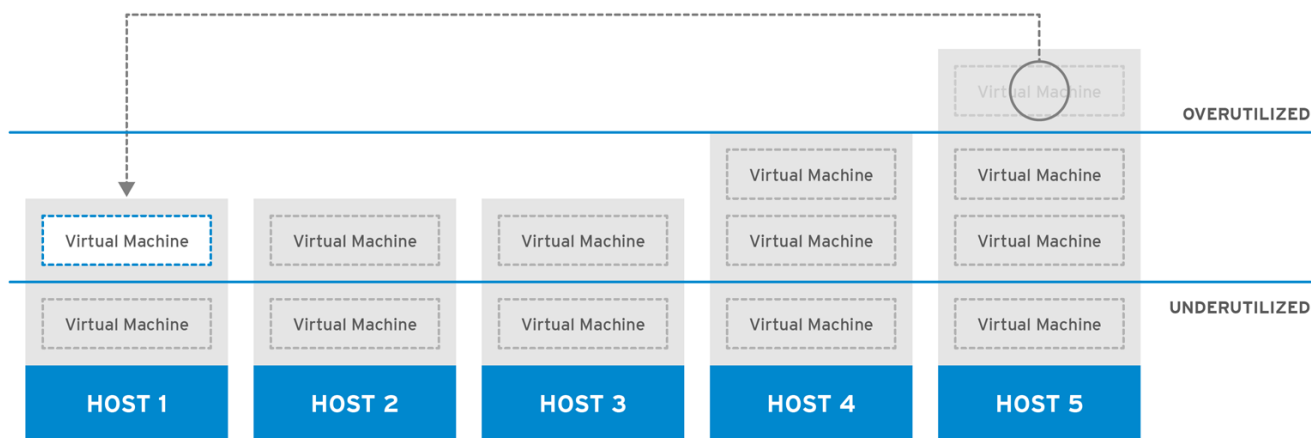
Table 1.4. Red Hat Virtualization System Administrator Roles - Advanced

Role	Privileges	Notes
TemplateAdmin	Administrator of a virtual machine template.	Can create, delete, and configure the storage domains and network details of templates, and move templates between domains.
StorageAdmin	Storage Administrator.	Can create, delete, configure, and manage an assigned storage domain.
HostAdmin	Host Administrator.	Can attach, remove, configure, and manage a specific host.
NetworkAdmin	Network Administrator.	Can configure and manage the network of a particular data center or cluster. A network administrator of a data center or cluster inherits network permissions for virtual pools within the cluster.
VmPoolAdmin	System Administrator of a virtual pool.	Can create, delete, and configure a virtual pool; assign and remove virtual pool users; and perform basic operations on a virtual machine in the pool.
GlusterAdmin	Gluster Storage Administrator.	Can create, delete, configure, and manage Gluster storage volumes.
VmImporterExporter	Import and export Administrator of a virtual machine.	Can import and export virtual machines. Able to view all virtual machines and templates exported by other users.

1.3. SCHEDULING POLICIES

A scheduling policy is a set of rules that defines the logic by which virtual machines are distributed amongst hosts in the cluster that scheduling policy is applied to. Scheduling policies determine this logic via a combination of filters, weightings, and a load balancing

policy. The Red Hat Virtualization Manager provides five default scheduling policies: **Evenly_Distributed**, **InClusterUpgrade**, **None**, **Power_Saving**, and **VM_Evenly_Distributed**. You can also define new scheduling policies that provide fine-grained control over the distribution of virtual machines. Regardless of the scheduling policy, a virtual machine will not start on a host with an overloaded CPU. By default, a host's CPU is considered overloaded if it has a load of more than 80% for 5 minutes, but these values can be changed using scheduling policies. See [Section 5.2.2.4, “Scheduling Policy Settings Explained”](#) for more information about the properties of each scheduling policy.

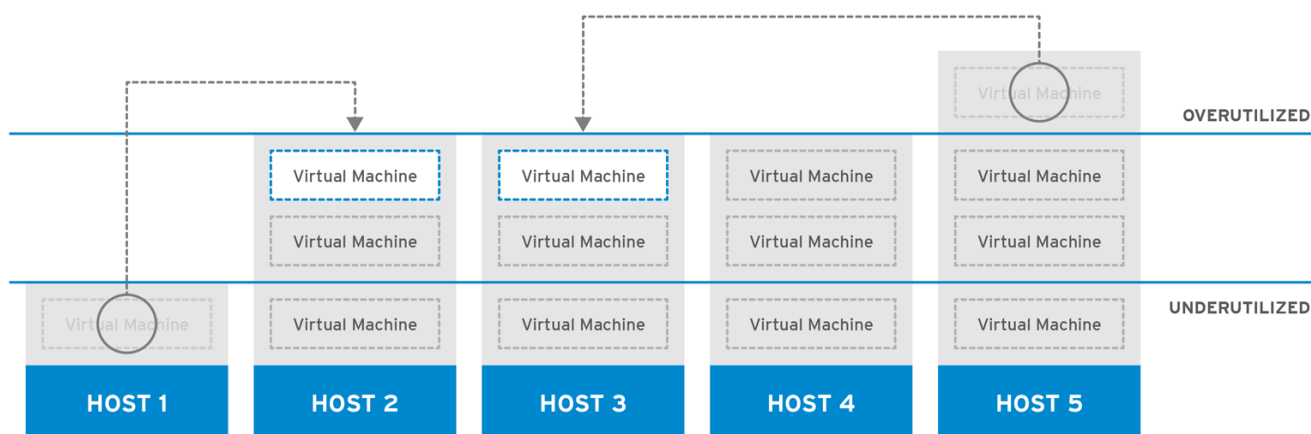


RHV_444396_0417

Figure 1.6. Evenly Distributed Scheduling Policy

The **Evenly_Distributed** scheduling policy distributes the memory and CPU processing load evenly across all hosts in the cluster. Additional virtual machines attached to a host will not start if that host has reached the defined **CpuOverCommitDurationMinutes**, **HighUtilization**, or **MaxFreeMemoryForOverUtilized**.

The **VM_Evenly_Distributed** scheduling policy virtual machines evenly between hosts based on a count of the virtual machines. The cluster is considered unbalanced if any host is running more virtual machines than the **HighVmCount** and there is at least one host with a virtual machine count that falls outside of the **MigrationThreshold**.



RHV_444396_0417


Figure 1.7. Power Saving Scheduling Policy

The **Power_Saving** scheduling policy distributes the memory and CPU processing load across a subset of available hosts to reduce power consumption on underutilized hosts. Hosts with a CPU load below the low utilization value for longer than the defined time

interval will migrate all virtual machines to other hosts so that it can be powered down. Additional virtual machines attached to a host will not start if that host has reached the defined high utilization value.

Set the **None** policy to have no load or power sharing between hosts for running virtual machines. This is the default mode. When a virtual machine is started, the memory and CPU processing load is spread evenly across all hosts in the cluster. Additional virtual machines attached to a host will not start if that host has reached the defined **CpuOverCommitDurationMinutes**, **HighUtilization**, or **MaxFreeMemoryForOverUtilized**.

The **InClusterUpgrade** scheduling policy distributes virtual machines based on host operating system version. Hosts with a newer operating system than the virtual machine currently runs on are given priority over hosts with the same operating system. Virtual machines that migrate to a host with a newer operating system will not migrate back to an older operating system. A virtual machine can restart on any host in the cluster. The policy allows hosts in the cluster to be upgraded by allowing the cluster to have mixed operating system versions. Preconditions must be met before the policy can be enabled. See [Upgrading Hosts in a Cluster from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7](#) in the *Red Hat Enterprise Virtualization 3.6 Upgrade Guide*



IMPORTANT

The InClusterUpgrade scheduling policy is used for upgrades between major versions only. For example, upgrading from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7.

1.3.1. Creating a Scheduling Policy

You can create new scheduling policies to control the logic by which virtual machines are distributed amongst a given cluster in your Red Hat Virtualization environment.

Procedure 1.3. Creating a Scheduling Policy

1. Click the **Configure** button in the header bar of the Administration Portal to open the **Configure** window.
2. Click **Scheduling Policies** to view the scheduling policies tab.
3. Click **New** to open the **New Scheduling Policy** window.

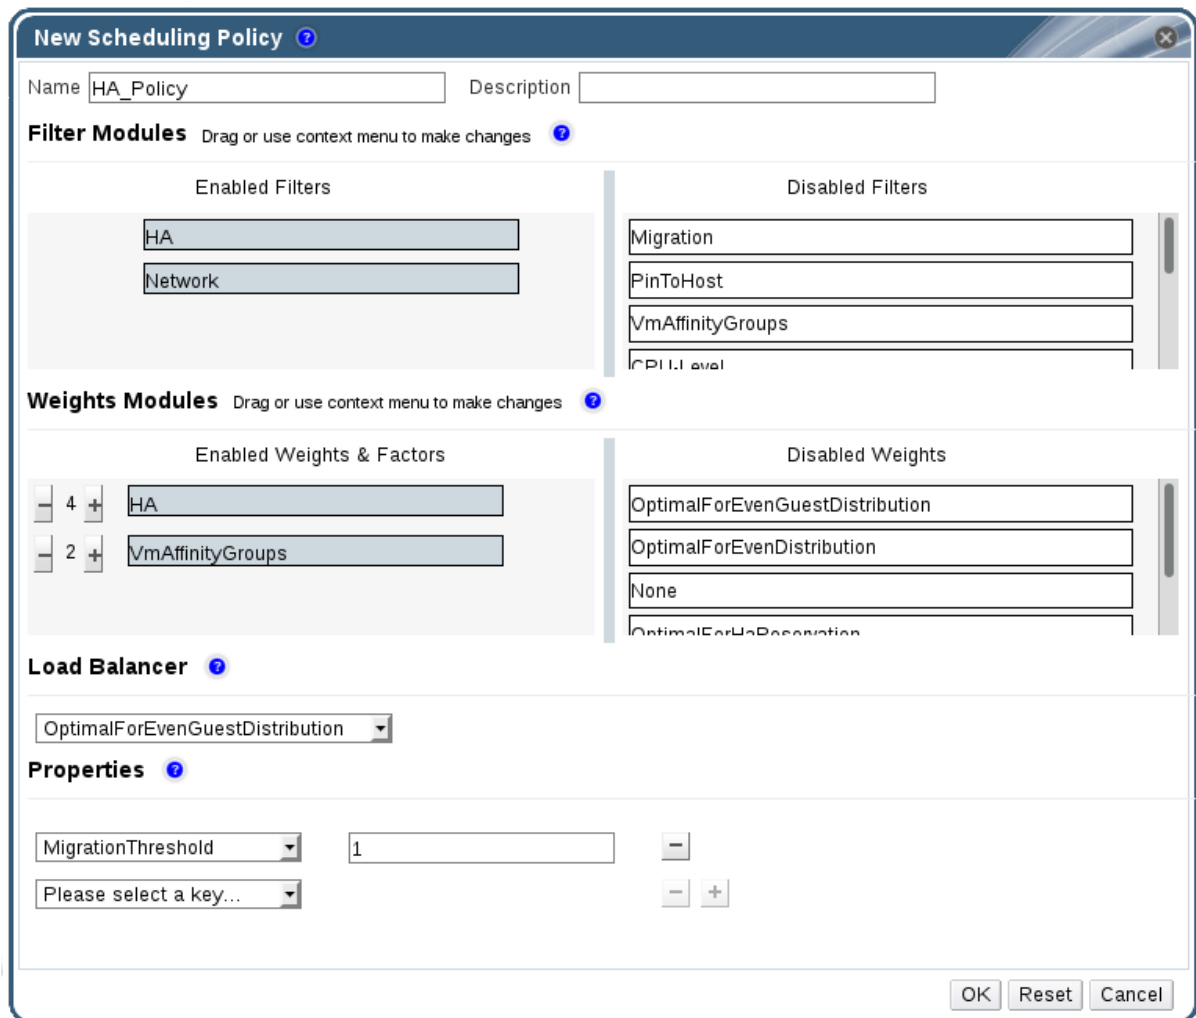


Figure 1.8. The New Scheduling Policy Window

4. Enter a **Name** and **Description** for the scheduling policy.
5. Configure filter modules:
 - a. In the **Filter Modules** section, drag and drop the preferred filter modules to apply to the scheduling policy from the **Disabled Filters** section into the **Enabled Filters** section.
 - b. Specific filter modules can also be set as the **First**, to be given highest priority, or **Last**, to be given lowest priority, for basic optimization.

To set the priority, right-click any filter module, hover the cursor over **Position** and select **First** or **Last**.
6. Configure weight modules:
 - a. In the **Weights Modules** section, drag and drop the preferred weights modules to apply to the scheduling policy from the **Disabled Weights** section into the **Enabled Weights & Factors** section.
 - b. Use the **+** and **-** buttons to the left of the enabled weight modules to increase or decrease the weight of those modules.
7. Specify a load balancing policy:

- a. From the drop-down menu in the **Load Balancer** section, select the load balancing policy to apply to the scheduling policy.
 - b. From the drop-down menu in the **Properties** section, select a load balancing property to apply to the scheduling policy and use the text field to the right of that property to specify a value.
 - c. Use the + and - buttons to add or remove additional properties.
8. Click **OK**.

1.3.2. Explanation of Settings in the New Scheduling Policy and Edit Scheduling Policy Window

The following table details the options available in the **New Scheduling Policy** and **Edit Scheduling Policy** windows.

Table 1.5. New Scheduling Policy and Edit Scheduling Policy Settings

Field Name	Description
Name	The name of the scheduling policy. This is the name used to refer to the scheduling policy in the Red Hat Virtualization Manager.
Description	A description of the scheduling policy. This field is recommended but not mandatory.

Field Name	Description
Filter Modules	<p>A set of filters for controlling the hosts on which a virtual machine in a cluster can run. Enabling a filter will filter out hosts that do not meet the conditions specified by that filter, as outlined below:</p> <ul style="list-style-type: none"> • CpuPinning: Hosts which do not satisfy the CPU pinning definition. • Migration: Prevent migration to the same host. • PinToHost: Hosts other than the host to which the virtual machine is pinned. • CPU-Level: Hosts that do not meet the CPU topology of the virtual machine. • CPU: Hosts with fewer CPUs than the number assigned to the virtual machine. • Memory: Hosts that do not have sufficient memory to run the virtual machine. • VmAffinityGroups: Hosts that do not meet the conditions specified for a virtual machine that is a member of an affinity group. For example, that virtual machines in an affinity group must run on the same host or on separate hosts. • InClusterUpgrade: Hosts which run an older operating system than the virtual machine currently runs on. • HostDevice: Hosts that do not support host devices required by the virtual machine. • HA: Forces the hosted engine virtual machine to only run on hosts with a positive high availability score. • Emulated-Machine: Hosts which do not have proper emulated machine support. • Network: Hosts on which networks required by the network interface controller of a virtual machine are not installed, or on which the cluster's display network is not installed.

Field Name	Description
Weights Modules	<p>A set of weightings for controlling the relative priority of factors considered when determining the hosts in a cluster on which a virtual machine can run.</p> <ul style="list-style-type: none"> • InClusterUpgrade: Weight hosts in accordance with their operating system version. The weight penalizes hosts with older operating systems more than hosts with the same operating system, giving priority to hosts with newer operating systems. • OptimalForHaReservation: Weights hosts in accordance with their high availability score. • None: Weights hosts in accordance with the even distribution module. • OptimalForEvenGuestDistribution: Weights hosts in accordance with the number of virtual machines running on those hosts. • VmAffinityGroups: Weights hosts in accordance with the affinity groups defined for virtual machines. This weight module determines how likely virtual machines in an affinity group are to run on the same host or on separate hosts in accordance with the parameters of that affinity group. • OptimalForPowerSaving: Weights hosts in accordance with their CPU usage, giving priority to hosts with higher CPU usage. • OptimalForEvenDistribution: Weights hosts in accordance with their CPU usage, giving priority to hosts with lower CPU usage. • HA: Weights hosts in accordance with their high availability score.
Load Balancer	<p>This drop-down menu allows you to select a load balancing module to apply. Load balancing modules determine the logic used to migrate virtual machines from hosts experiencing high usage to hosts experiencing lower usage.</p>

Field Name	Description
Properties	This drop-down menu allows you to add or remove properties for load balancing modules, and is only available when you have selected a load balancing module for the scheduling policy. No properties are defined by default, and the properties that are available are specific to the load balancing module that is selected. Use the + and - buttons to add or remove additional properties to or from the load balancing module.

1.4. INSTANCE TYPES

Instance types can be used to define the hardware configuration of a virtual machine. Selecting an instance type when creating or editing a virtual machine will automatically fill in the hardware configuration fields. This allows users to create multiple virtual machines with the same hardware configuration without having to manually fill in every field.

A set of predefined instance types are available by default, as outlined in the following table:

Table 1.6. Predefined Instance Types

Name	Memory	vCPUs
Tiny	512 MB	1
Small	2 GB	1
Medium	4 GB	2
Large	8 GB	2
XLarge	16 GB	4

Administrators can also create, edit, and remove instance types from the **Instance Types** tab of the **Configure** window.






Figure 1.9. The Instance Types Tab

Fields in the **New Virtual Machine** and **Edit Virtual Machine** windows that are bound to an instance type will have a chain link image next to them (). If the value of one of these fields is changed, the virtual machine will be detached from the instance type, changing to **Custom**, and the chain will appear broken (). However, if the value is changed back, the chain will relink and the instance type will move back to the selected one.

1.4.1. Creating Instance Types

Administrators can create new instance types, which can then be selected by users when creating or editing virtual machines.

Procedure 1.4. Creating an Instance Type

1. On the header bar, click **Configure**.
2. Click the **Instance Types** tab.
3. Click **New** to open the **New Instance Type** window.

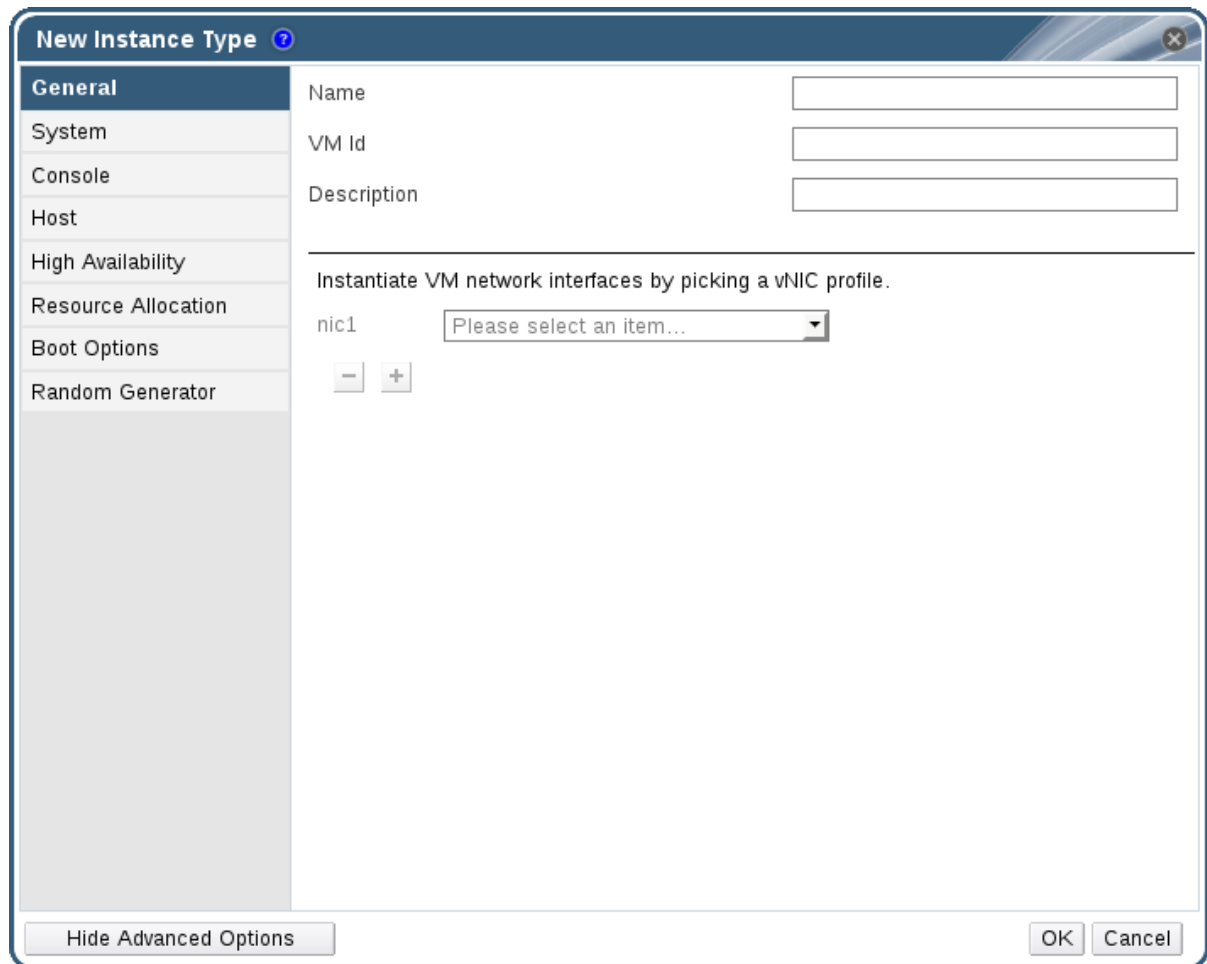


Figure 1.10. The New Instance Type Window

4. Enter a **Name** and **Description** for the instance type.
5. Click **Show Advanced Options** and configure the instance type's settings as required. The settings that appear in the **New Instance Type** window are identical to those in the **New Virtual Machine** window, but with the relevant fields only. See [Explanation of Settings in the New Virtual Machine and Edit Virtual Machine Windows](#) in the *Virtual Machine Management Guide*
6. Click **OK**.

The new instance type will appear in the **Instance Types** tab in the **Configure** window, and can be selected from the **Instance Type** drop-down list when creating or editing a virtual machine.

1.4.2. Editing Instance Types

Administrators can edit existing instance types from the **Configure** window.

Procedure 1.5. Editing Instance Type Properties

1. On the header bar, click **Configure**.
2. Click the **Instance Types** tab.
3. Select the instance type to be edited.

4. Click **Edit** to open the **Edit Instance Type** window.
5. Change the settings as required.
6. Click **OK**.

The configuration of the instance type is updated. When a new virtual machine based on this instance type is created, or when an existing virtual machine based on this instance type is updated, the new configuration is applied.

Existing virtual machines based on this instance type will display fields, marked with a chain icon, that will be updated. If the existing virtual machines were running when the instance type was changed, the orange Pending Changes icon will appear beside them and the fields with the chain icon will be updated at the next restart.

1.4.3. Removing Instance Types

Procedure 1.6. Removing an Instance Type

1. On the header bar, click **Configure**.
2. Click the **Instance Types** tab.
3. Select the instance type to be removed.
4. Click **Remove** to open the **Remove Instance Type** window.
5. If any virtual machines are based on the instance type to be removed, a warning window listing the attached virtual machines will appear. To continue removing the instance type, select the **Approve Operation** checkbox. Otherwise click **Cancel**.
6. Click **OK**.

The instance type is removed from the **Instance Types** list and can no longer be used when creating a new virtual machine. Any virtual machines that were attached to the removed instance type will now be attached to **Custom** (no instance type).

1.5. MAC ADDRESS POOLS

MAC address pools define the range of MAC addresses from which MAC addresses are allocated for each data center. A MAC address pool is specified for each data center. By using MAC address pools Red Hat Virtualization can automatically generate and assign MAC addresses to new virtual network devices, which helps to prevent MAC address duplication. MAC address pools are more memory efficient when all MAC addresses related to a data center are within the range for the assigned MAC address pool.

The same MAC address pool can be shared by multiple data centers, but each data center has a single MAC address pool assigned. A default MAC address pool is created by Red Hat Virtualization and is used if another MAC address pool is not assigned. For more information about assigning MAC address pools to data centers see [Section 4.5.1, “Creating a New Data Center”](#).

The MAC address pool assigns the next available MAC address following the last address that was returned to the pool. If there are no further addresses left in the range, the search starts again from the beginning of the range. If there are multiple MAC address ranges with

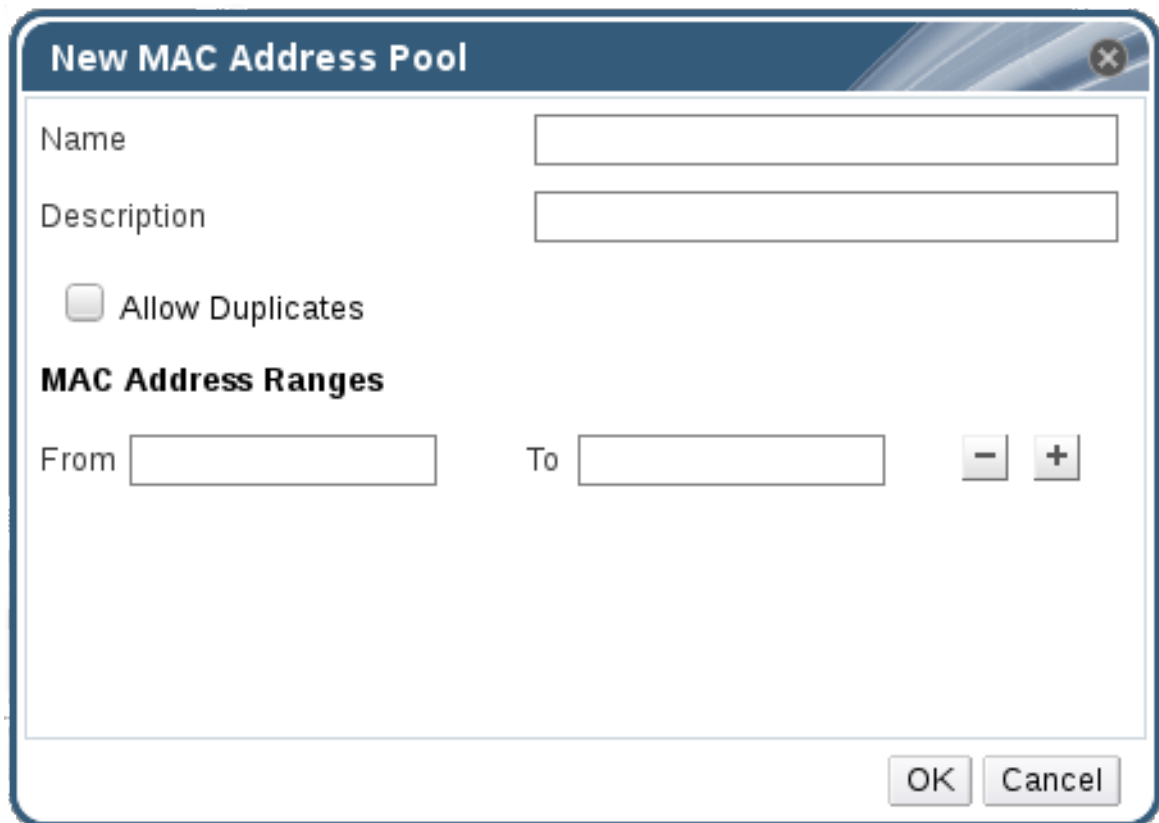
available MAC addresses defined in a single MAC address pool, the ranges take turns in serving incoming requests in the same way available MAC addresses are selected.

1.5.1. Creating MAC Address Pools

You can create new MAC address pools.

Procedure 1.7. Creating a MAC Address Pool

1. On the header bar, click the **Configure** button to open the **Configure** window.
2. Click the **MAC Address Pools** tab.
3. Click the **Add** button to open the **New MAC Address Pool** window.



The image shows a dialog box titled "New MAC Address Pool". It has a standard Windows-style title bar with a close button. The main area contains the following elements:

- A text input field for "Name".
- A text input field for "Description".
- A checkbox labeled "Allow Duplicates".
- A section header "MAC Address Ranges".
- Two text input fields for "From" and "To", with minus and plus buttons to the right of the "To" field.
- "OK" and "Cancel" buttons at the bottom right.

Figure 1.11. The New MAC Address Pool Window

4. Enter the **Name** and **Description** of the new MAC address pool.
5. Select the **Allow Duplicates** check box to allow a MAC address to be used multiple times in a pool. The MAC address pool will not automatically use a duplicate MAC address, but enabling the duplicates option means a user can manually use a duplicate MAC address.



NOTE

If one MAC address pool has duplicates disabled, and another has duplicates enabled, each MAC address can be used once in the pool with duplicates disabled but can be used multiple times in the pool with duplicates enabled.

6. Enter the required **MAC Address Ranges**. To enter multiple ranges click the plus button next to the **From** and **To** fields.
7. Click **OK**.

1.5.2. Editing MAC Address Pools

You can edit MAC address pools to change the details, including the range of MAC addresses available in the pool and whether duplicates are allowed.

Procedure 1.8. Editing MAC Address Pool Properties

1. On the header bar, click the **Configure** button to open the **Configure** window.
2. Click the **MAC Address Pools** tab.
3. Select the MAC address pool to be edited.
4. Click the **Edit** button to open the **Edit MAC Address Pool** window.
5. Change the **Name**, **Description**, **Allow Duplicates**, and **MAC Address Ranges** fields as required.



NOTE

When a MAC address range is updated, the MAC addresses of existing NICs are not reassigned. MAC addresses that were already assigned, but are outside of the new MAC address range, are added as user-specified MAC addresses and are still tracked by that MAC address pool.

6. Click **OK**.

1.5.3. Editing MAC Address Pool Permissions

After a MAC address pool has been created, you can edit its user permissions. The user permissions control which data centers can use the MAC address pool. See [Section 1.1, “Roles”](#) for more information on adding new user permissions.

Procedure 1.9. Editing MAC Address Pool Permissions

1. On the header bar, click the **Configure** button to open the **Configure** window.
2. Click the **MAC Address Pools** tab.
3. Select the required MAC address pool.
4. Edit the user permissions for the MAC address pool:
 - To add user permissions to a MAC address pool:
 - a. Click **Add** in the user permissions pane at the bottom of the **Configure** window.
 - b. Search for and select the required users.

- c. Select the required role from the **Role to Assign** drop-down list.
- d. Click **OK** to add the user permissions.
- To remove user permissions from a MAC address pool:
 - a. Select the user permission to be removed in the user permissions pane at the bottom of the **Configure** window.
 - b. Click **Remove** to remove the user permissions.

1.5.4. Removing MAC Address Pools

You can remove created MAC address pools, but the default MAC address pool cannot be removed.

Procedure 1.10. Removing a MAC Address Pool

1. On the header bar, click the **Configure** button to open the **Configure** window.
2. Click the **MAC Address Pools** tab.
3. Select the MAC address pool to be removed.
4. Click the **Remove** button to open the **Remove MAC Address Pool** window.
5. Click **OK**.

CHAPTER 2. DASHBOARD

The Dashboard provides an overview of the Red Hat Virtualization system status by displaying a summary of Red Hat Virtualization's resources and utilization. This summary can alert the user to a problem and allows them to analyse the problem area.

The information in the dashboard is updated every 15 minutes by default from the Data Warehouse, and every 15 seconds by default by the Manager API, or whenever the Dashboard is refreshed. The Dashboard is refreshed when the user changes back from another tab or when manually refreshed. The Dashboard does not automatically refresh. The inventory card information is supplied by the Manager API and the utilization information is supplied by the Data Warehouse. The Dashboard is implemented as a UI plugin component, which is automatically installed and upgraded alongside the Manager.

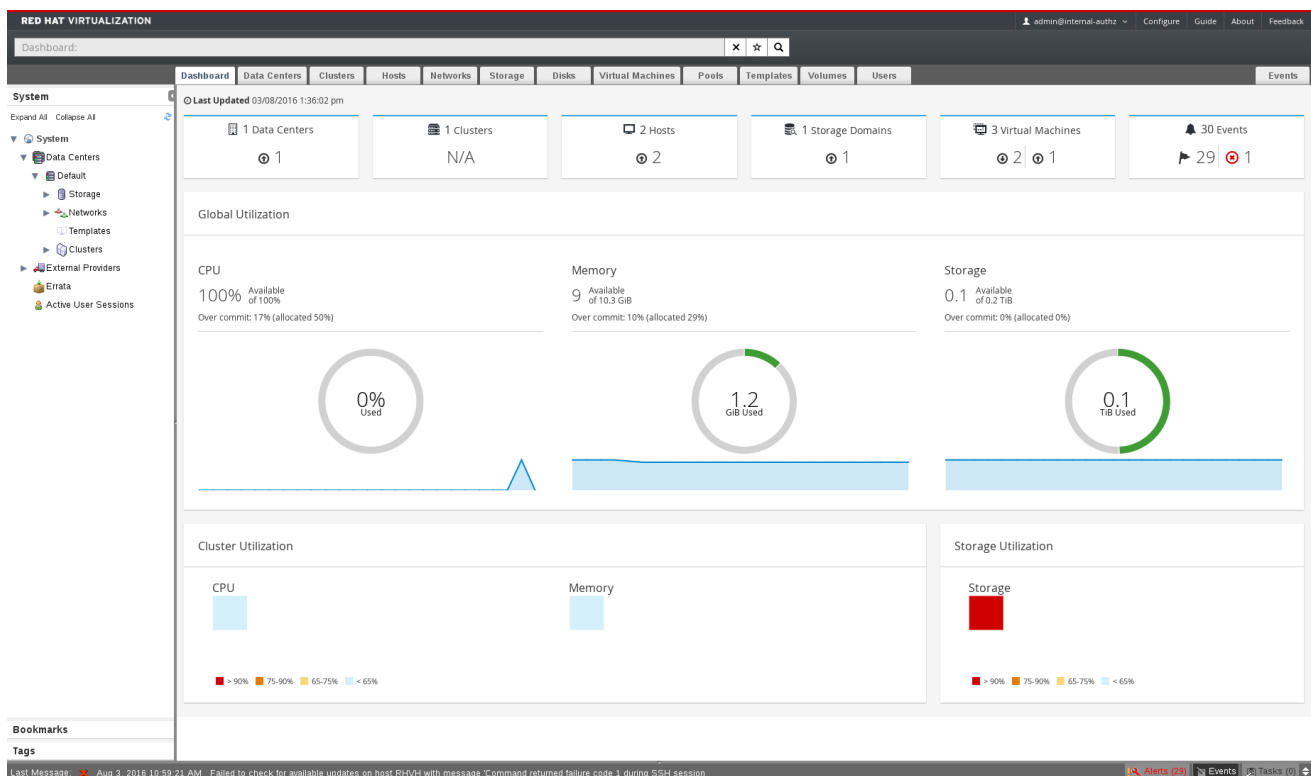


Figure 2.1. The Dashboard

2.1. PREREQUISITES

The Dashboard requires that the Data Warehouse is installed and configured. For more information on installing and configuring the Data Warehouse see [Installing and Configuring Data Warehouse](#) in the *Data Warehouse Guide*

2.2. GLOBAL INVENTORY

The top section of the Dashboard provides a global inventory of the Red Hat Virtualization resources and includes items for data centers, clusters, hosts, storage domains, virtual machines, and events. Icons show the status of each resource and numbers show the quantity of the each resource with that status.










Figure 2.2. Global Inventory

The title shows the number of a type of resource and their status is displayed below the title. Clicking on the resource title navigates to the related tab in the Red Hat Virtualization Manager. The status for **Clusters** is always displayed as N/A.

Table 2.1. Resource Status

Icon	Status
	None of that resource added to Red Hat Virtualization.
	<p>Shows the number of a resource with a warning status. Clicking on the icon navigates to the appropriate tab with the search limited to that resource with a warning status. The search is limited differently for each resource:</p> <ul style="list-style-type: none"> • Data Centers: The search is limited to data centers that are not operational or non-responsive. • Hosts: The search is limited to hosts that are unassigned, in maintenance mode, installing, rebooting, preparing for maintenance, pending approval, or connecting. • Storage Domains: The search is limited to storage domains that are uninitialized, unattached, inactive, in maintenance mode, preparing for maintenance, detaching, or activating. • Virtual Machines: The search is limited to virtual machines that are powering up, paused, migrating, waiting, suspended, or powering down. • Events: The search is limited to events with the severity of warning.
	Shows the number of a resource with an up status. Clicking on the icon navigates to the appropriate tab with the search limited to resources that are up.
	<p>Shows the number of a resource with a down status. Clicking on the icon navigates to the appropriate tab with the search limited to resources with a down status. The search is limited differently for each resource:</p> <ul style="list-style-type: none"> • Data Centers: The search is limited to data centers that are uninitialized, in maintenance mode, or with a down status. • Hosts: The search is limited to hosts that are non-responsive, have an error, have an installation error, non-operational, initializing, or down. • Storage Domains: The search is limited to storage domains that are detached or inactive. • Virtual Machines: The search is limited to virtual machines that are down, not responding, or rebooting.

Icon	Status
	Shows the number of events with an alert status. Clicking on the icon navigates to the Events tab with the search limited to events with the severity of alert.
	Shows the number of events with an error status. Clicking on the icon navigates to the Events tab with the search limited to events with the severity of error.

2.3. GLOBAL UTILIZATION

The **Global Utilization** section shows the system utilization of the CPU, Memory and Storage.

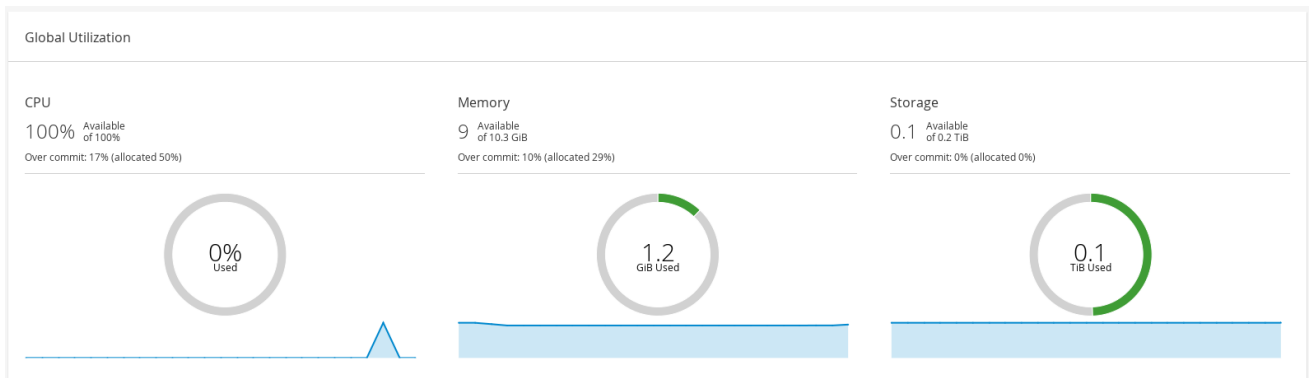


Figure 2.3. Global Utilization

- The top section shows the percentage of the available CPU, memory or storage and the over commit ratio. For example, the over commit ratio for the CPU is calculated by dividing the number of virtual cores by the number of physical cores that are available for the running virtual machines based on the latest data in the Data Warehouse.
- The donut displays the usage in percentage for the CPU, memory or storage and shows the average usage for all hosts based on the average usage in the last 5 minutes. Hovering over a section of the donut will display the value of the selected section.
- The line graph at the bottom displays the trend in the last 24 hours. Each data point shows the average usage for a specific hour. Hovering over a point on the graph displays the time and the percentage used for the CPU graph and the amount of usage for the memory and storage graphs.

2.3.1. Top Utilized Resources

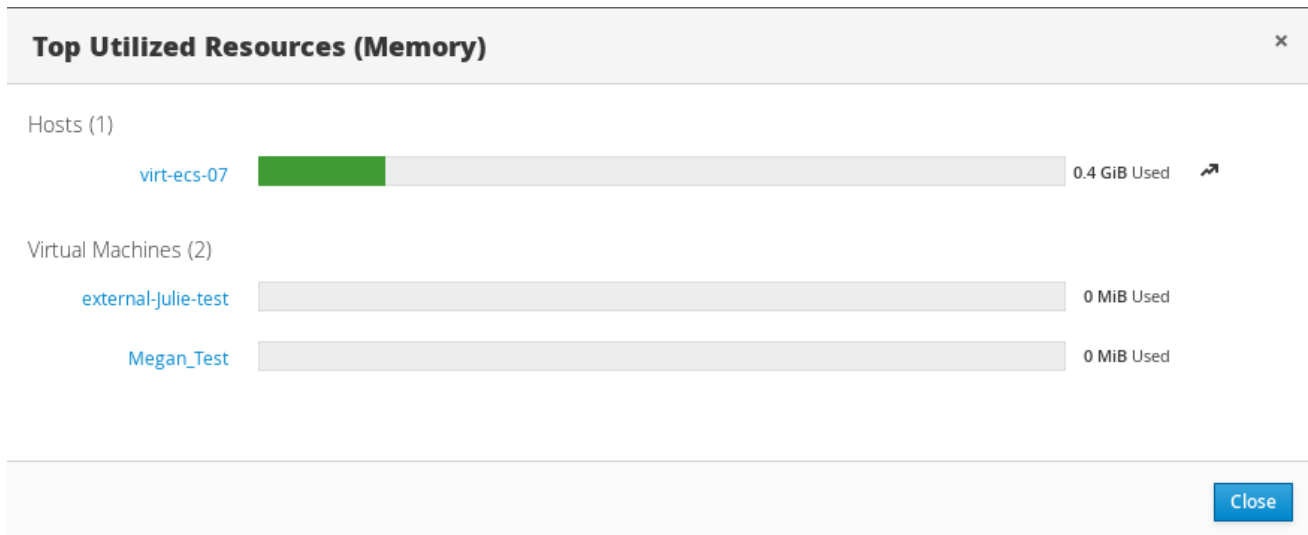


Figure 2.4. Top Utilized Resources (Memory)

Clicking the donut in the global utilization section of the Dashboard will display a list of the top utilized resources for the CPU, memory or storage. For CPU and memory the pop-up shows a list of the ten hosts and virtual machines with the highest usage. For storage the pop-up shows a list of the top ten utilized storage domains and virtual machines. The arrow to the right of the usage bar shows the trend of usage for that resource in the last minute.

2.4. CLUSTER UTILIZATION

The **Cluster Utilization** shows the cluster utilization for the CPU and memory in a heatmap.

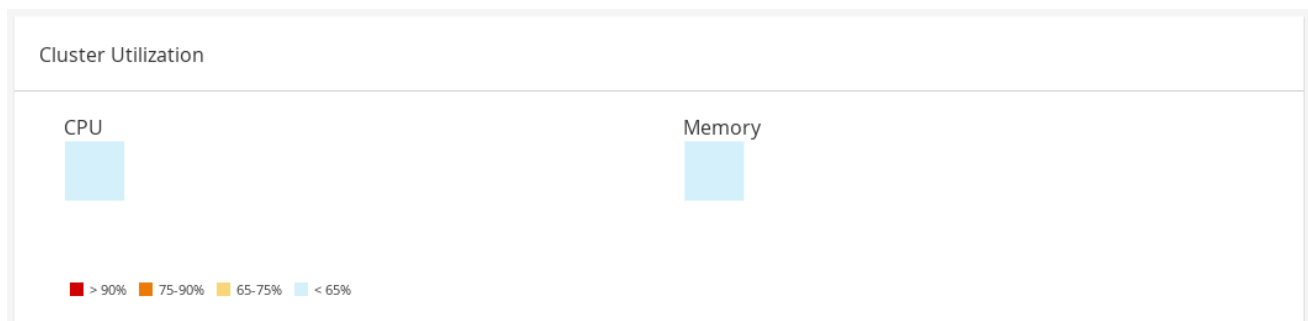


Figure 2.5. Cluster Utilization

2.4.1. CPU

The heatmap of the CPU utilization for a specific cluster that shows the average utilization of the CPU for the last 24 hours. Hovering over the heatmap displays the cluster name. Clicking on the heatmap navigates to the **Hosts** tab and displays the results of a search on a specific cluster sorted by CPU utilization. The formula used to calculate the usage of the CPU by the cluster is the average host CPU utilization in the cluster. This is calculated by using the average host CPU utilization for each host over the last 24 hours to find the total average usage of the CPU by the cluster.

2.4.2. Memory

The heatmap of the memory utilization for a specific cluster that shows the average utilization of the memory for the last 24 hours. Hovering over the heatmap displays the

cluster name. Clicking on the heatmap navigates to the **Hosts** tab and displays the results of a search on a specific cluster sorted by memory usage. The formula used to calculate the memory usage by the cluster is the total utilization of the memory in the cluster in GB. This is calculated by using the average host memory utilization for each host over the last 24 hours to find the total average usage of memory by the cluster.

2.5. STORAGE UTILIZATION

The **Storage Utilization** shows the storage utilization in a heatmap.

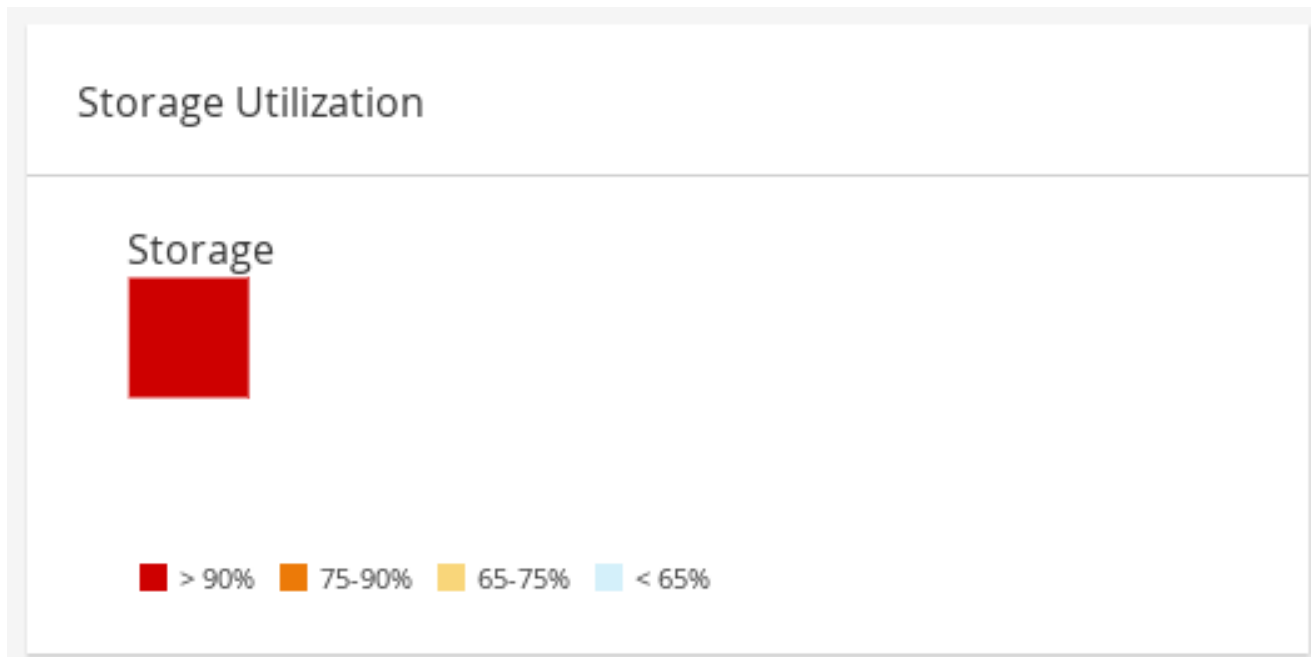


Figure 2.6. Storage Utilization

The heatmap shows the average utilization of the storage for the last 24 hours. The formula used to calculate the storage usage by the cluster is the total utilization of the storage in the cluster. This is calculated by using the average storage utilization for each host over the last 24 hours to find the total average usage of the storage by the cluster. Hovering over the heatmap displays the storage domain name. Clicking on the heatmap navigates to the **Storage** tab with the storage domains sorted by utilization.

PART II. ADMINISTERING THE RESOURCES

CHAPTER 3. QUALITY OF SERVICE

Red Hat Virtualization allows you to define quality of service entries that provide fine-grained control over the level of input and output, processing, and networking capabilities that resources in your environment can access. Quality of service entries are defined at the data center level and are assigned to profiles created under clusters and storage domains. These profiles are then assigned to individual resources in the clusters and storage domains where the profiles were created.

3.1. STORAGE QUALITY OF SERVICE

Storage quality of service defines the maximum level of throughput and the maximum level of input and output operations for a virtual disk in a storage domain. Assigning storage quality of service to a virtual disk allows you to fine tune the performance of storage domains and prevent the storage operations associated with one virtual disk from affecting the storage capabilities available to other virtual disks hosted in the same storage domain.

3.1.1. Creating a Storage Quality of Service Entry

Create a storage quality of service entry.

Procedure 3.1. Creating a Storage Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click **QoS** in the details pane.
3. Click **Storage**.
4. Click **New**.
5. Enter a name for the quality of service entry in the **QoS Name** field.
6. Enter a description for the quality of service entry in the **Description** field.
7. Specify the throughput quality of service:
 - a. Select the **Throughput** check box.
 - b. Enter the maximum permitted total throughput in the **Total** field.
 - c. Enter the maximum permitted throughput for read operations in the **Read** field.
 - d. Enter the maximum permitted throughput for write operations in the **Write** field.
8. Specify the input and output quality of service:
 - a. Select the **Iops** check box.
 - b. Enter the maximum permitted number of input and output operations per second in the **Total** field.
 - c. Enter the maximum permitted number of input operations per second in the **Read** field.

- d. Enter the maximum permitted number of output operations per second in the **Write** field.

9. Click **OK**.

You have created a storage quality of service entry, and can create disk profiles based on that entry in data storage domains that belong to the data center.

3.1.2. Removing a Storage Quality of Service Entry

Remove an existing storage quality of service entry.

Procedure 3.2. Removing a Storage Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click **QoS** in the details pane.
3. Click **Storage**.
4. Select the storage quality of service entry to remove.
5. Click **Remove**.
6. Click **OK** when prompted.

You have removed the storage quality of service entry, and that entry is no longer available. If any disk profiles were based on that entry, the storage quality of service entry for those profiles is automatically set to **[unlimited]**.

3.2. VIRTUAL MACHINE NETWORK QUALITY OF SERVICE

Virtual machine network quality of service is a feature that allows you to create profiles for limiting both the inbound and outbound traffic of individual virtual network interface controllers. With this feature, you can limit bandwidth in a number of layers, controlling the consumption of network resources.

3.2.1. Creating a Virtual Machine Network Quality of Service Entry

Create a virtual machine network quality of service entry to regulate network traffic when applied to a virtual network interface controller (vNIC) profile, also known as a virtual machine network interface profile.

Procedure 3.3. Creating a Virtual Machine Network Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click the **QoS** tab in the details pane.
3. Click **VM Network**.
4. Click **New**.
5. Enter a name for the virtual machine network quality of service entry in the **Name** field.

6. Enter the limits for the **Inbound** and **Outbound** network traffic.

7. Click **OK**.

You have created a virtual machine network quality of service entry that can be used in a virtual network interface controller.

3.2.2. Settings in the New Virtual Machine Network QoS and Edit Virtual Machine Network QoS Windows Explained

Virtual machine network quality of service settings allow you to configure bandwidth limits for both inbound and outbound traffic on three distinct levels.

Table 3.1. Virtual Machine Network QoS Settings

Field Name	Description
Data Center	The data center to which the virtual machine network QoS policy is to be added. This field is configured automatically according to the selected data center.
Name	A name to represent the virtual machine network QoS policy within the Manager.
Inbound	<p>The settings to be applied to inbound traffic. Select or clear the Inbound check box to enable or disable these settings.</p> <ul style="list-style-type: none"> • Average: The average speed of inbound traffic. • Peak: The speed of inbound traffic during peak times. • Burst: The speed of inbound traffic during bursts.
Outbound	<p>The settings to be applied to outbound traffic. Select or clear the Outbound check box to enable or disable these settings.</p> <ul style="list-style-type: none"> • Average: The average speed of outbound traffic. • Peak: The speed of outbound traffic during peak times. • Burst: The speed of outbound traffic during bursts.

3.2.3. Removing a Virtual Machine Network Quality of Service Entry

Remove an existing virtual machine network quality of service entry.

Procedure 3.4. Removing a Virtual Machine Network Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click the **QoS** tab in the details pane.
3. Click **VM Network**.
4. Select the virtual machine network quality of service entry to remove.
5. Click **Remove**.
6. Click **OK** when prompted.

You have removed the virtual machine network quality of service entry, and that entry is no longer available.

3.3. HOST NETWORK QUALITY OF SERVICE

Host network quality of service configures the networks on a host to enable the control of network traffic through the physical interfaces. Host network quality of service allows for the fine tuning of network performance by controlling the consumption of network resources on the same physical network interface controller. This helps to prevent situations where one network causes other networks attached to the same physical network interface controller to no longer function due to heavy traffic. By configuring host network quality of service, these networks can now function on the same physical network interface controller without congestion issues.

3.3.1. Creating a Host Network Quality of Service Entry

Create a host network quality of service entry.

Procedure 3.5. Creating a Host Network Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click **QoS** in the details pane.
3. Click **Host Network**.
4. Click **New**.
5. Enter a name for the quality of service entry in the **QoS Name** field.
6. Enter a description for the quality of service entry in the **Description** field.
7. Enter the desired values for **Weighted Share**, **Rate Limit [Mbps]**, and **Committed Rate [Mbps]**.
8. Click **OK**.

3.3.2. Settings in the New Host Network Quality of Service and Edit Host Network Quality of Service Windows Explained

Host network quality of service settings allow you to configure bandwidth limits for outbound traffic.

Table 3.2. Host Network QoS Settings

Field Name	Description
Data Center	The data center to which the host network QoS policy is to be added. This field is configured automatically according to the selected data center.
QoS Name	A name to represent the host network QoS policy within the Manager.
Description	A description of the host network QoS policy.
Outbound	<p>The settings to be applied to outbound traffic.</p> <ul style="list-style-type: none"> • Weighted Share: Signifies how much of the logical link's capacity a specific network should be allocated, relative to the other networks attached to the same logical link. The exact share depends on the sum of shares of all networks on that link. By default this is a number in the range 1-100. • Rate Limit [Mbps]: The maximum bandwidth to be used by a network. • Committed Rate [Mbps]: The minimum bandwidth required by a network. The Committed Rate requested is not guaranteed and will vary depending on the network infrastructure and the Committed Rate requested by other networks on the same logical link.

3.3.3. Removing a Host Network Quality of Service Entry

Remove an existing network quality of service entry.

Procedure 3.6. Removing a Host Network Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click the **QoS** tab in the details pane.
3. Click **Host Network**.
4. Select the network quality of service entry to remove.
5. Click **Remove**.
6. Click **OK** when prompted.

3.4. CPU QUALITY OF SERVICE

CPU quality of service defines the maximum amount of processing capability a virtual machine can access on the host on which it runs, expressed as a percent of the total processing capability available to that host. Assigning CPU quality of service to a virtual machine allows you to prevent the workload on one virtual machine in a cluster from affecting the processing resources available to other virtual machines in that cluster.

3.4.1. Creating a CPU Quality of Service Entry

Create a CPU quality of service entry.

Procedure 3.7. Creating a CPU Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click **QoS** in the details pane.
3. Click **CPU**.
4. Click **New**.
5. Enter a name for the quality of service entry in the **QoS Name** field.
6. Enter a description for the quality of service entry in the **Description** field.
7. Enter the maximum processing capability the quality of service entry permits in the **Limit** field, in percentage. Do not include the % symbol.
8. Click **OK**.

You have created a CPU quality of service entry, and can create CPU profiles based on that entry in clusters that belong to the data center.

3.4.2. Removing a CPU Quality of Service Entry

Remove an existing CPU quality of service entry.

Procedure 3.8. Removing a CPU Quality of Service Entry

1. Click the **Data Centers** resource tab and select a data center.
2. Click **QoS** in the details pane.
3. Click **CPU**.
4. Select the CPU quality of service entry to remove.
5. Click **Remove**.
6. Click **OK** when prompted.

You have removed the CPU quality of service entry, and that entry is no longer available. If any CPU profiles were based on that entry, the CPU quality of service entry for those profiles is automatically set to **[unlimited]**.

CHAPTER 4. DATA CENTERS

4.1. INTRODUCTION TO DATA CENTERS

A data center is a logical entity that defines the set of resources used in a specific environment. A data center is considered a container resource, in that it is comprised of logical resources, in the form of clusters and hosts; network resources, in the form of logical networks and physical NICs; and storage resources, in the form of storage domains.

A data center can contain multiple clusters, which can contain multiple hosts; it can have multiple storage domains associated to it; and it can support multiple virtual machines on each of its hosts. A Red Hat Virtualization environment can contain multiple data centers; the data center infrastructure allows you to keep these centers separate.

All data centers are managed from the single Administration Portal.

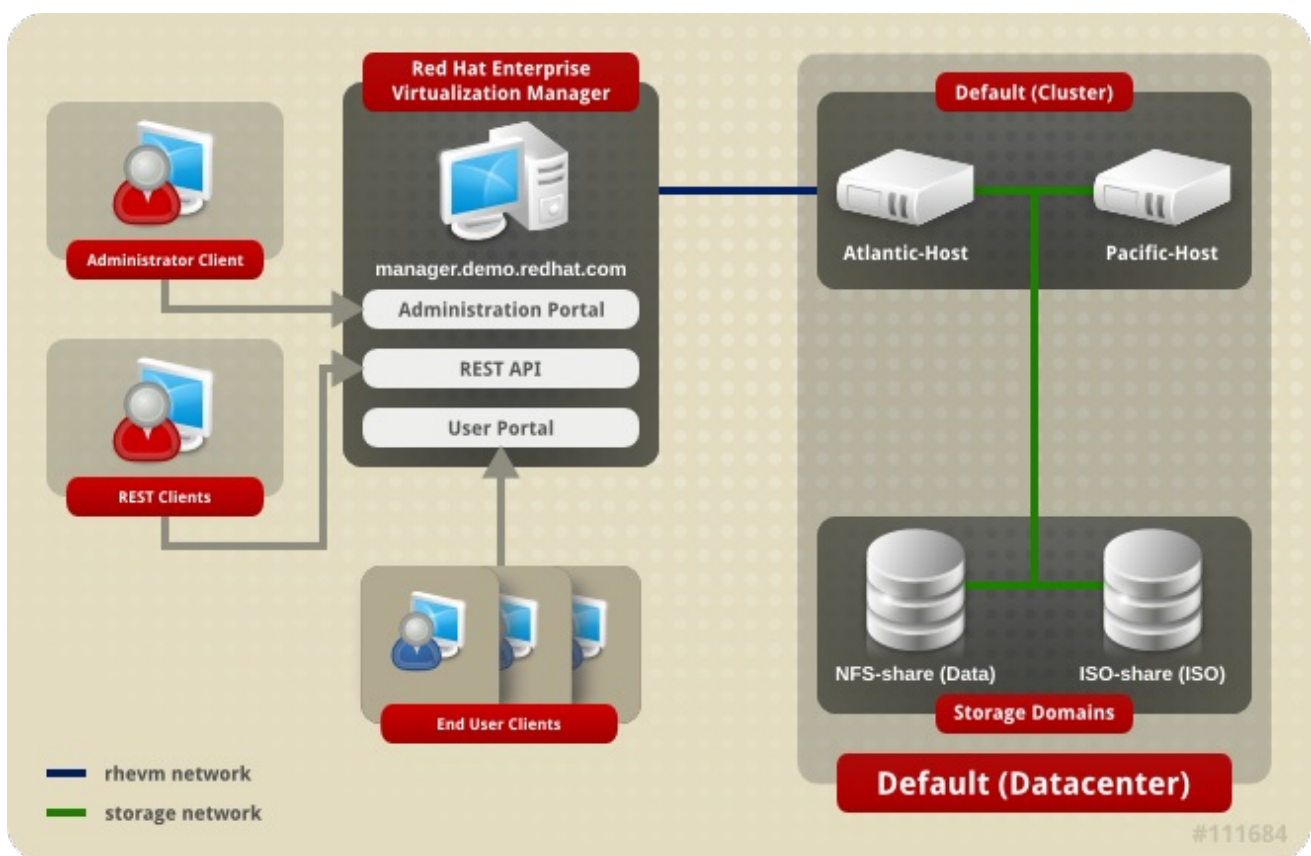


Figure 4.1. Data Centers

Red Hat Virtualization creates a default data center during installation. You can configure the default data center, or set up new appropriately named data centers.




Figure 4.2. Data Center Objects

4.2. THE STORAGE POOL MANAGER

The Storage Pool Manager (SPM) is a role given to one of the hosts in the data center enabling it to manage the storage domains of the data center. The SPM entity can be run on any host in the data center; the Red Hat Virtualization Manager grants the role to one of the hosts. The SPM does not preclude the host from its standard operation; a host running as SPM can still host virtual resources.

The SPM entity controls access to storage by coordinating the metadata across the storage domains. This includes creating, deleting, and manipulating virtual disks (images), snapshots, and templates, and allocating storage for sparse block devices (on SAN). This is an exclusive responsibility: only one host can be the SPM in the data center at one time to ensure metadata integrity.

The Red Hat Virtualization Manager ensures that the SPM is always available. The Manager moves the SPM role to a different host if the SPM host encounters problems accessing the storage. When the SPM starts, it ensures that it is the only host granted the role; therefore it will acquire a storage-centric lease. This process can take some time.

4.3. SPM PRIORITY

The SPM role uses some of a host's available resources. The SPM priority setting of a host alters the likelihood of the host being assigned the SPM role: a host with high SPM priority will be assigned the SPM role before a host with low SPM priority. Critical virtual machines on hosts with low SPM priority will not have to contend with SPM operations for host resources.

You can change a host's SPM priority by editing the host.

4.4. USING THE EVENTS TAB TO IDENTIFY PROBLEM OBJECTS IN DATA CENTERS

The **Events** tab for a data center displays all events associated with that data center; events include audits, warnings, and errors. The information displayed in the results list will enable you to identify problem objects in your Red Hat Virtualization environment.

The **Events** results list has two views: Basic and Advanced. Basic view displays the event icon, the time of the event, and the description of the events. Advanced view displays these also and includes, where applicable, the event ID; the associated user, host, virtual machine, template, data center, storage, and cluster; the Gluster volume, and the correlation ID.

4.5. DATA CENTER TASKS

4.5.1. Creating a New Data Center

This procedure creates a data center in your virtualization environment. The data center requires a functioning cluster, host, and storage domain to operate.



NOTE

The storage **Type** can be edited until the first storage domain is added to the data center. Once a storage domain has been added, the storage **Type** cannot be changed.

Once the **Compatibility Version** is set, it cannot be lowered at a later time; version regression is not allowed.

Procedure 4.1. Creating a New Data Center

1. Select the **Data Centers** resource tab to list all data centers in the results list.
2. Click **New** to open the **New Data Center** window.
3. Enter the **Name** and **Description** of the data center.
4. Select the storage **Type**, **Compatibility Version**, and **Quota Mode** of the data center from the drop-down menus.
5. Optionally, change the MAC address pool for the data center. The default MAC address pool is preselected by default. For more information on creating MAC address pools see [Section 1.5, "MAC Address Pools"](#).
 - a. Click the **MAC Address Pool** tab.
 - b. Select the required MAC address pool from the **MAC Address Pool** drop-down list.
6. Click **OK** to create the data center and open the **New Data Center - Guide Me** window.
7. The **Guide Me** window lists the entities that need to be configured for the data center. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the data center and

clicking the **Guide Me** button.

The new data center is added to the virtualization environment. It will remain **Uninitialized** until a cluster, host, and storage domain are configured for it; use **Guide Me** to configure these entities.

4.5.2. Explanation of Settings in the New Data Center and Edit Data Center Windows

The table below describes the settings of a data center as displayed in the **New Data Center** and **Edit Data Center** windows. Invalid entries are outlined in orange when you click **OK**, prohibiting the changes being accepted. In addition, field prompts indicate the expected values or range of values.

Table 4.1. Data Center Properties

Field	Description/Action
Name	The name of the data center. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Description	The description of the data center. This field is recommended but not mandatory.
Type	<p>The storage type. Choose one of the following:</p> <ul style="list-style-type: none"> • Shared • Local <p>The type of data domain dictates the type of the data center and cannot be changed after creation without significant disruption. Multiple types of storage domains (iSCSI, NFS, FC, POSIX, and Gluster) can be added to the same data center, though local and shared domains cannot be mixed.</p>
Compatibility Version	<p>The version of Red Hat Virtualization. Choose one of the following:</p> <ul style="list-style-type: none"> • 3.6 • 4.0 <p>After upgrading the Red Hat Virtualization Manager, the hosts, clusters and data centers may still be in the earlier version. Ensure that you have upgraded all the hosts, then the clusters, before you upgrade the Compatibility Level of the data center.</p>

Field	Description/Action
Quota Mode	<p>Quota is a resource limitation tool provided with Red Hat Virtualization. Choose one of:</p> <ul style="list-style-type: none"> • Disabled: Select if you do not want to implement Quota • Audit: Select if you want to edit the Quota settings • Enforced: Select to implement Quota
MAC Address Pool	<p>The MAC address pool of the data center. If no other MAC address pool is assigned the default MAC address pool is used. For more information on MAC address pools see Section 1.5, "MAC Address Pools"</p>

4.5.3. Re-Initializing a Data Center: Recovery Procedure

This recovery procedure replaces the master data domain of your data center with a new master data domain; necessary in the event of data corruption of your master data domain. Re-initializing a data center allows you to restore all other resources associated with the data center, including clusters, hosts, and non-problematic storage domains.

You can import any backup or exported virtual machines or templates into your new master data domain.

Procedure 4.2. Re-Initializing a Data Center

1. Click the **Data Centers** resource tab and select the data center to re-initialize.
2. Ensure that any storage domains attached to the data center are in maintenance mode.
3. Right-click the data center and select **Re-Initialize Data Center** from the drop-down menu to open the **Data Center Re-Initialize** window.
4. The **Data Center Re-Initialize** window lists all available (detached; in maintenance mode) storage domains. Click the radio button for the storage domain you are adding to the data center.
5. Select the **Approve operation** check box.
6. Click **OK** to close the window and re-initialize the data center.

The storage domain is attached to the data center as the master data domain and activated. You can now import any backup or exported virtual machines or templates into your new master data domain.

4.5.4. Removing a Data Center

An active host is required to remove a data center. Removing a data center will not remove the associated resources.

Procedure 4.3. Removing a Data Center

1. Ensure the storage domains attached to the data center is in maintenance mode.
2. Click the **Data Centers** resource tab and select the data center to remove.
3. Click **Remove** to open the **Remove Data Center(s)** confirmation window.
4. Click **OK**.

4.5.5. Force Removing a Data Center

A data center becomes **Non Responsive** if the attached storage domain is corrupt or if the host becomes **Non Responsive**. You cannot **Remove** the data center under either circumstance.

Force Remove does not require an active host. It also permanently removes the attached storage domain.

It may be necessary to **Destroy** a corrupted storage domain before you can **Force Remove** the data center.


Procedure 4.4. Force Removing a Data Center

1. Click the **Data Centers** resource tab and select the data center to remove.
2. Click **Force Remove** to open the **Force Remove Data Center** confirmation window.
3. Select the **Approve operation** check box.
4. Click **OK**

The data center and attached storage domain are permanently removed from the Red Hat Virtualization environment.

4.5.6. Changing the Data Center Compatibility Version

Red Hat Virtualization data centers have a compatibility version. The compatibility version indicates the version of Red Hat Virtualization that the data center is intended to be compatible with. All clusters in the data center must support the desired compatibility level.



NOTE


To change the data center compatibility version, you must have first updated all the clusters in your data center to a level that supports your desired compatibility level.

Procedure 4.5. Changing the Data Center Compatibility Version

1. From the Administration Portal, click the **Data Centers** tab.
2. Select the data center to change from the list displayed.

3. Click **Edit**.
4. Change the **Compatibility Version** to the desired value.
5. Click **OK** to open the **Change Data Center Compatibility Version** confirmation window.
6. Click **OK** to confirm.

You have updated the compatibility version of the data center.



IMPORTANT

Upgrading the compatibility will also upgrade all of the storage domains belonging to the data center.

4.6. DATA CENTERS AND STORAGE DOMAINS

4.6.1. Attaching an Existing Data Domain to a Data Center

Data domains that are **Unattached** can be attached to a data center. Shared storage domains of multiple types (iSCSI, NFS, FC, POSIX, and Gluster) can be added to the same data center.

Procedure 4.6. Attaching an Existing Data Domain to a Data Center

1. Click the **Data Centers** resource tab and select the appropriate data center.
2. Select the **Storage** tab in the details pane to list the storage domains already attached to the data center.
3. Click **Attach Data** to open the **Attach Storage** window.
4. Select the check box for the data domain to attach to the data center. You can select multiple check boxes to attach multiple data domains.
5. Click **OK**.

The data domain is attached to the data center and is automatically activated.

4.6.2. Attaching an Existing ISO domain to a Data Center

An ISO domain that is **Unattached** can be attached to a data center. The ISO domain must be of the same **Storage Type** as the data center.

Only one ISO domain can be attached to a data center.

Procedure 4.7. Attaching an Existing ISO Domain to a Data Center

1. Click the **Data Centers** resource tab and select the appropriate data center.
2. Select the **Storage** tab in the details pane to list the storage domains already attached to the data center.
3. Click **Attach ISO** to open the **Attach ISO Library** window.

4. Click the radio button for the appropriate ISO domain.
5. Click **OK**.

The ISO domain is attached to the data center and is automatically activated.

4.6.3. Attaching an Existing Export Domain to a Data Center



NOTE

The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disk images, and templates can then be uploaded from the imported storage domain to the attached data center. See [Section 8.6, “Importing Existing Storage Domains”](#) for information on importing storage domains.

An export domain that is **Unattached** can be attached to a data center. Only one export domain can be attached to a data center.

Procedure 4.8. Attaching an Existing Export Domain to a Data Center

1. Click the **Data Centers** resource tab and select the appropriate data center.
2. Select the **Storage** tab in the details pane to list the storage domains already attached to the data center.
3. Click **Attach Export** to open the **Attach Export Domain** window.
4. Click the radio button for the appropriate Export domain.
5. Click **OK**.

The export domain is attached to the data center and is automatically activated.

4.6.4. Detaching a Storage Domain from a Data Center

Detaching a storage domain from a data center will stop the data center from associating with that storage domain. The storage domain is not removed from the Red Hat Virtualization environment; it can be attached to another data center.

Data, such as virtual machines and templates, remains attached to the storage domain.



NOTE

The master storage, if it is the last available storage domain, cannot be removed.

Procedure 4.9. Detaching a Storage Domain from a Data Center

1. Click the **Data Centers** resource tab and select the appropriate data center.
2. Select the **Storage** tab in the details pane to list the storage domains attached to the data center.

3. Select the storage domain to detach. If the storage domain is **Active**, click **Maintenance** to open the **Maintenance Storage Domain(s)** confirmation window.
4. Click **OK** to initiate maintenance mode.
5. Click **Detach** to open the **Detach Storage** confirmation window.
6. Click **OK**.

You have detached the storage domain from the data center. It can take up to several minutes for the storage domain to disappear from the details pane.

4.7. DATA CENTERS AND PERMISSIONS

4.7.1. Managing System Permissions for a Data Center

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A data center administrator is a system administration role for a specific data center only. This is useful in virtualization environments with multiple data centers where each data center requires an administrator. The **DataCenterAdmin** role is a hierarchical model; a user assigned the data center administrator role for a data center can manage all objects in the data center with the exception of storage for that data center. Use the **Configure** button in the header bar to assign a data center administrator for all data centers in the environment.

The data center administrator role permits the following actions:

- Create and remove clusters associated with the data center.
- Add and remove hosts, virtual machines, and pools associated with the data center.
- Edit user permissions for virtual machines associated with the data center.



NOTE

You can only assign roles and permissions to existing users.

You can change the system administrator of a data center by removing the existing system administrator and adding the new system administrator.

4.7.2. Data Center Administrator Roles Explained

Data Center Permission Roles

The table below describes the administrator roles and privileges applicable to data center administration.

Table 4.2. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
DataCenterAdmin	Data Center Administrator	Can use, create, delete, manage all physical and virtual resources within a specific data center except for storage, including clusters, hosts, templates and virtual machines.
NetworkAdmin	Network Administrator	Can configure and manage the network of a particular data center. A network administrator of a data center inherits network permissions for virtual machines within the data center as well.

4.7.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 4.10. Assigning a Role to a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Click **Add**.
4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down list.
6. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

4.7.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 4.11. Removing a Role from a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's

role, and the inherited permissions for the selected resource.

3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

CHAPTER 5. CLUSTERS

5.1. INTRODUCTION TO CLUSTERS

A cluster is a logical grouping of hosts that share the same storage domains and have the same type of CPU (either Intel or AMD). If the hosts have different generations of CPU models, they use only the features present in all models.

Each cluster in the system must belong to a data center, and each host in the system must belong to a cluster. Virtual machines are dynamically allocated to any host in a cluster and can be migrated between them, according to policies defined on the **Clusters** tab and in the Configuration tool during runtime. The cluster is the highest level at which power and load-sharing policies can be defined.

The number of hosts and number of virtual machines that belong to a cluster are displayed in the results list under **Host Count** and **VM Count**, respectively.

Clusters run virtual machines or Red Hat Gluster Storage Servers. These two purposes are mutually exclusive: A single cluster cannot support virtualization and storage hosts together.

Red Hat Virtualization creates a default cluster in the default data center during installation.

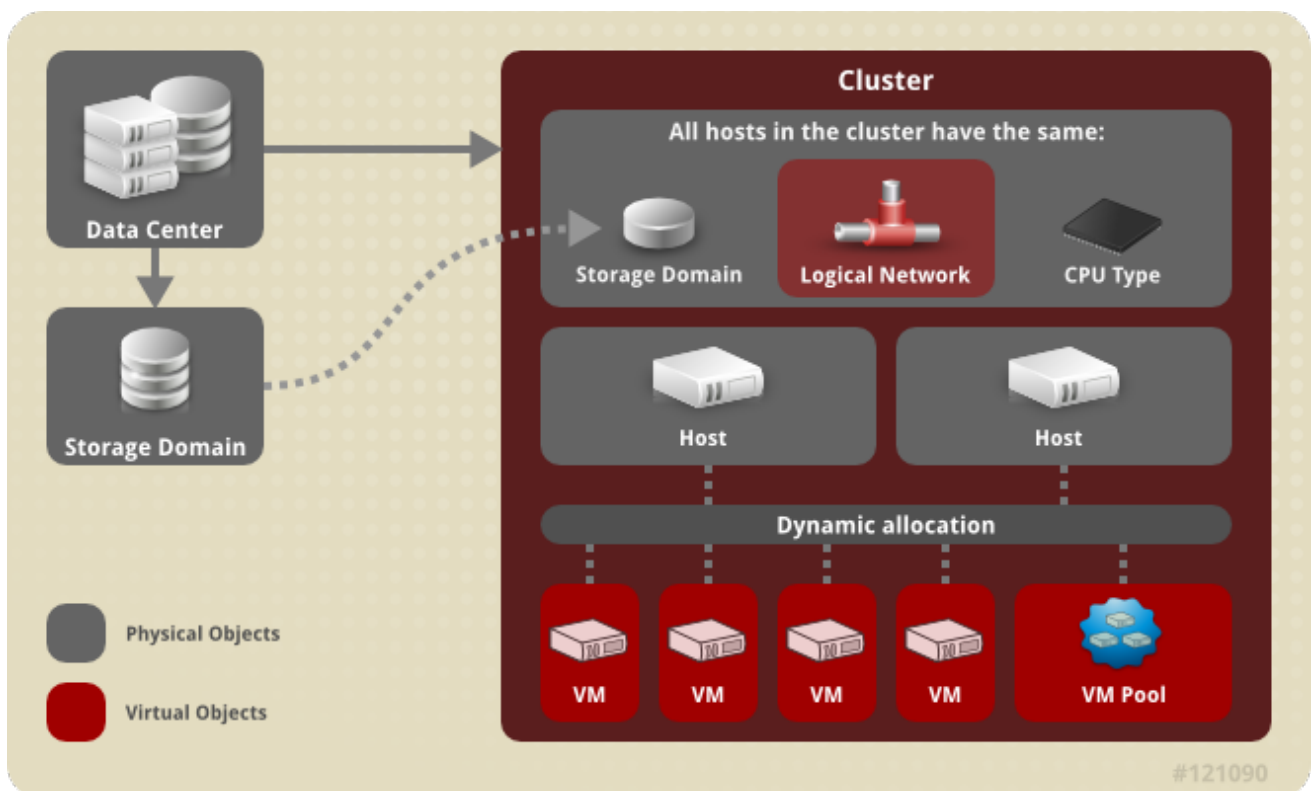


Figure 5.1. Cluster

5.2. CLUSTER TASKS

5.2.1. Creating a New Cluster

A data center can contain multiple clusters, and a cluster can contain multiple hosts. All hosts in a cluster must be of the same CPU type (Intel or AMD). It is recommended that you

create your hosts before you create your cluster to ensure CPU type optimization. However, you can configure the hosts at a later time using the **Guide Me** button.

Procedure 5.1. Creating a New Cluster

1. Select the **Clusters** resource tab.
2. Click **New**.
3. Select the **Data Center** the cluster will belong to from the drop-down list.
4. Enter the **Name** and **Description** of the cluster.
5. Select a network from the **Management Network** drop-down list to assign the management network role.
6. Select the **CPU Architecture** and **CPU Type** from the drop-down lists. It is important to match the CPU processor family with the minimum CPU processor type of the hosts you intend to attach to the cluster, otherwise the host will be non-operational.



NOTE

For both Intel and AMD CPU types, the listed CPU models are in logical order from the oldest to the newest. If your cluster includes hosts with different CPU models, select the oldest CPU model. For more information on each CPU model, see <https://access.redhat.com/solutions/634853>.

7. Select the **Compatibility Version** of the cluster from the drop-down list.
8. Select either the **Enable Virt Service** or **Enable Gluster Service** radio button to define whether the cluster will be populated with virtual machine hosts or with Gluster-enabled nodes. Note that you cannot add Red Hat Virtualization Host (RHVH) to a Gluster-enabled cluster.
9. Optionally select the **Enable to set VM maintenance reason** check box to enable an optional reason field when a virtual machine is shut down from the Manager, allowing the administrator to provide an explanation for the maintenance.
10. Optionally select the **Enable to set Host maintenance reason** check box to enable an optional reason field when a host is placed into maintenance mode from the Manager, allowing the administrator to provide an explanation for the maintenance.
11. Select either the **/dev/random source** (Linux-provided device) or **/dev/hwrng source** (external hardware device) check box to specify the random number generator device that all hosts in the cluster will use.
12. Click the **Optimization** tab to select the memory page sharing threshold for the cluster, and optionally enable CPU thread handling and memory ballooning on the hosts in the cluster.
13. Click the **Migration Policy** tab to define the virtual machine migration policy for the cluster.

14. Click the **Scheduling Policy** tab to optionally configure a scheduling policy, configure scheduler optimization settings, enable trusted service for hosts in the cluster, enable HA Reservation, and add a custom serial number policy.
15. Click the **Console** tab to optionally override the globalSPICE proxy, if any, and specify the address of a SPICE proxy for hosts in the cluster.
16. Click the **Fencing policy** tab to enable or disable fencing in the cluster, and select fencing options.
17. Click **OK** to create the cluster and open the **New Cluster - Guide Me** window.
18. The **Guide Me** window lists the entities that need to be configured for the cluster. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the cluster and clicking the **Guide Me** button.

The new cluster is added to the virtualization environment.

5.2.2. Explanation of Settings and Controls in the New Cluster and Edit Cluster Windows

5.2.2.1. General Cluster Settings Explained

Setting	Value
Data Center	Default
Name	
Description	
Comment	
Management Network	ovirtmgmt
CPU Architecture	undefined
CPU Type	
Compatibility Version	3.6
Enable Virt Service	<input checked="" type="radio"/>
Enable Gluster Service	<input type="radio"/>
Enable to set VM maintenance reason	<input type="checkbox"/>
Enable to set Host maintenance reason	<input type="checkbox"/>
Required Random Number Generator sources:	
/dev/random source	<input type="checkbox"/>
/dev/hwrng source	<input type="checkbox"/>

Figure 5.2. New Cluster window

The table below describes the settings for the **General** tab in the **New Cluster** and **Edit Cluster** windows. Invalid entries are outlined in orange when you click **OK**, prohibiting the changes being accepted. In addition, field prompts indicate the expected values or range of

values.

Table 5.1. General Cluster Settings

Field	Description/Action
Data Center	The data center that will contain the cluster. The data center must be created before adding a cluster.
Name	The name of the cluster. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Description / Comment	The description of the cluster or additional notes. These fields are recommended but not mandatory.
Management Network	The logical network which will be assigned the management network role. The default is ovirtmgmt . On existing clusters, the management network can only be changed via the Manage Networks button in the Logical Networks tab in the details pane.
CPU Architecture	<p>The CPU architecture of the cluster. Different CPU types are available depending on which CPU architecture is selected.</p> <ul style="list-style-type: none"> • undefined: All CPU types are available. • x86_64: All Intel and AMD CPU types are available. • ppc64: Only IBM POWER 8 is available.

Field	Description/Action
CPU Type	<p>The CPU type of the cluster. Choose one of:</p> <ul style="list-style-type: none"> • Intel Conroe Family • Intel Penryn Family • Intel Nehalem Family • Intel Westmere Family • Intel SandyBridge Family • Intel Haswell • AMD Opteron G1 • AMD Opteron G2 • AMD Opteron G3 • AMD Opteron G4 • AMD Opteron G5 • IBM POWER 8 <p>All hosts in a cluster must run either Intel, AMD, or IBM POWER 8 CPU type; this cannot be changed after creation without significant disruption. The CPU type should be set to the oldest CPU model in the cluster. Only features present in all models can be used. For both Intel and AMD CPU types, the listed CPU models are in logical order from the oldest to the newest.</p>
Compatibility Version	<p>The version of Red Hat Virtualization. Choose one of:</p> <ul style="list-style-type: none"> • 3.6 • 4.0 <p>You will not be able to select a version older than the version specified for the data center.</p>
Enable Virt Service	<p>If this radio button is selected, hosts in this cluster will be used to run virtual machines.</p>
Enable Gluster Service	<p>If this radio button is selected, hosts in this cluster will be used as Red Hat Gluster Storage Server nodes, and not for running virtual machines. You cannot add a Red Hat Virtualization Host to a cluster with this option enabled.</p>

Field	Description/Action
Import existing gluster configuration	<p>This check box is only available if the Enable Gluster Service radio button is selected. This option allows you to import an existing Gluster-enabled cluster and all its attached hosts to Red Hat Virtualization Manager.</p> <p>The following options are required for each host in the cluster that is being imported:</p> <ul style="list-style-type: none"> • Address: Enter the IP or fully qualified domain name of the Gluster host server. • Fingerprint: Red Hat Virtualization Manager fetches the host's fingerprint, to ensure you are connecting with the correct host. • Root Password: Enter the root password required for communicating with the host.
Enable to set VM maintenance reason	<p>If this check box is selected, an optional reason field will appear when a virtual machine in the cluster is shut down from the Manager. This allows you to provide an explanation for the maintenance, which will appear in the logs and when the virtual machine is powered on again.</p>
Enable to set Host maintenance reason	<p>If this check box is selected, an optional reason field will appear when a host in the cluster is moved into maintenance mode from the Manager. This allows you to provide an explanation for the maintenance, which will appear in the logs and when the host is activated again.</p>
Required Random Number Generator sources:	<p>If one of the following check boxes is selected, all hosts in the cluster must have that device available. This enables passthrough of entropy from the random number generator device to virtual machines.</p> <ul style="list-style-type: none"> • /dev/random source - The Linux-provided random number generator. • /dev/hwrng source - An external hardware generator.

5.2.2.2. Optimization Settings Explained

Memory page sharing allows virtual machines to use up to 200% of their allocated memory by utilizing unused memory in other virtual machines. This process is based on the

assumption that the virtual machines in your Red Hat Virtualization environment will not all be running at full capacity at the same time, allowing unused memory to be temporarily allocated to a particular virtual machine.

CPU Thread Handling allows hosts to run virtual machines with a total number of processor cores greater than number of cores in the host. This is useful for non-CPU-intensive workloads, where allowing a greater number of virtual machines to run can reduce hardware requirements. It also allows virtual machines to run with CPU topologies that would otherwise not be possible, specifically if the number of guest cores is between the number of host cores and number of host threads.

The table below describes the settings for the **Optimization** tab in the **New Cluster** and **Edit Cluster** windows.

Table 5.2. Optimization Settings

Field	Description/Action
Memory Optimization	<ul style="list-style-type: none"> • None - Disable memory overcommit: Disables memory page sharing. • For Server Load - Allow scheduling of 150% of physical memory: Sets the memory page sharing threshold to 150% of the system memory on each host. • For Desktop Load - Allow scheduling of 200% of physical memory: Sets the memory page sharing threshold to 200% of the system memory on each host.
CPU Threads	<p>Selecting the Count Threads As Cores check box allows hosts to run virtual machines with a total number of processor cores greater than the number of cores in the host.</p> <p>The exposed host threads would be treated as cores which can be utilized by virtual machines. For example, a 24-core system with 2 threads per core (48 threads total) can run virtual machines with up to 48 cores each, and the algorithms to calculate host CPU load would compare load against twice as many potential utilized cores.</p>

Field	Description/Action
Memory Balloon	<p>Selecting the Enable Memory Balloon Optimization check box enables memory overcommitment on virtual machines running on the hosts in this cluster. When this option is set, the Memory Overcommit Manager (MoM) will start ballooning where and when possible, with a limitation of the guaranteed memory size of every virtual machine.</p> <p>To have a balloon running, the virtual machine needs to have a balloon device with relevant drivers. Each virtual machine includes a balloon device unless specifically removed. Each host in this cluster receives a balloon policy update when its status changes to Up. If necessary, you can manually update the balloon policy on a host without having to change the status. See Section 5.2.5, “Updating the MoM Policy on Hosts in a Cluster”.</p> <p>It is important to understand that in some scenarios ballooning may collide with KSM. In such cases MoM will try to adjust the balloon size to minimize collisions. Additionally, in some scenarios ballooning may cause sub-optimal performance for a virtual machine. Administrators are advised to use ballooning optimization with caution.</p>
KSM control	<p>Selecting the Enable KSM check box enables MoM to run Kernel Same-page Merging (KSM) when necessary and when it can yield a memory saving benefit that outweighs its CPU cost.</p>

5.2.2.3. Migration Policy Settings Explained

A migration policy defines the conditions for live migrating virtual machines in the event of host failure. These conditions include the downtime of the virtual machine during migration, network bandwidth, and how the virtual machines are prioritized.

Table 5.3. Migration Policies Explained

Policy	Description
Legacy	<p>Legacy behavior of 3.6 version. Overrides in vdsm.conf are still applied. The guest agent hook mechanism is disabled.</p>

Policy	Description
Minimal downtime	A policy that lets virtual machines migrate in typical situations. Virtual machines should not experience any significant downtime. The migration will be aborted if the virtual machine migration does not converge after a long time (dependent on QEMU iterations, with a maximum of 500 milliseconds). The guest agent hook mechanism is enabled.
Suspend workload if needed	A policy that lets virtual machines migrate in most situations, including virtual machines running heavy workloads. Virtual machines may experience a more significant downtime. The migration may still be aborted for extreme workloads. The guest agent hook mechanism is enabled.

The bandwidth settings define the maximum bandwidth of both outgoing and incoming migrations per host.

Table 5.4. Bandwidth Explained

Policy	Description
Auto	Bandwidth is copied from the Rate Limit [Mbps] setting in the data center Host Network QoS . If the rate limit has not been defined, it is computed as a minimum of link speeds of sending and receiving network interfaces. If rate limit has not been set, and link speeds are not available, it is determined by local VDSM setting on sending host.
Hypervisor default	Bandwidth is controlled by local VDSM setting on sending Host.
Custom	Defined by user (in Mbps).

The resilience policy defines how the virtual machines are prioritized in the migration.

Table 5.5. Resilience Policy Settings

Field	Description/Action
Migrate Virtual Machines	Migrates all virtual machines in order of their defined priority.
Migrate only Highly Available Virtual Machines	Migrates only highly available virtual machines to prevent overloading other hosts.

Field	Description/Action
Do Not Migrate Virtual Machines	Prevents virtual machines from being migrated.

The **Additional Properties** are only applicable to the **Legacy** migration policy.

Table 5.6. Additional Properties Explained

Property	Description
Auto Converge migrations	<p>Allows you to set whether auto-convergence is used during live migration of virtual machines. Large virtual machines with high workloads can dirty memory more quickly than the transfer rate achieved during live migration, and prevent the migration from converging. Auto-convergence capabilities in QEMU allow you to force convergence of virtual machine migrations. QEMU automatically detects a lack of convergence and triggers a throttle-down of the vCPUs on the virtual machine. Auto-convergence is disabled globally by default.</p> <ul style="list-style-type: none"> • Select Inherit from global setting to use the auto-convergence setting that is set at the global level. This option is selected by default. • Select Auto Converge to override the global setting and allow auto-convergence for the virtual machine. • Select Don't Auto Converge to override the global setting and prevent auto-convergence for the virtual machine.

Property	Description
Enable migration compression	<p>The option allows you to set whether migration compression is used during live migration of the virtual machine. This feature uses Xor Binary Zero Run-Length-Encoding to reduce virtual machine downtime and total live migration time for virtual machines running memory write-intensive workloads or for any application with a sparse memory update pattern. Migration compression is disabled globally by default.</p> <ul style="list-style-type: none"> • Select Inherit from global setting to use the compression setting that is set at the global level. This option is selected by default. • Select Compress to override the global setting and allow compression for the virtual machine. • Select Don't compress to override the global setting and prevent compression for the virtual machine.

5.2.2.4. Scheduling Policy Settings Explained

Scheduling policies allow you to specify the usage and distribution of virtual machines between available hosts. Define the scheduling policy to enable automatic load balancing across the hosts in a cluster. Regardless of the scheduling policy, a virtual machine will not start on a host with an overloaded CPU. By default, a host's CPU is considered overloaded if it has a load of more than 80% for 5 minutes, but these values can be changed using scheduling policies. See [Section 1.3, "Scheduling Policies"](#) for more information about scheduling policies.

To add a scheduling policy to an existing cluster, click the **Clusters** tab and click the **Edit** button, then click the **Scheduling Policy** tab.

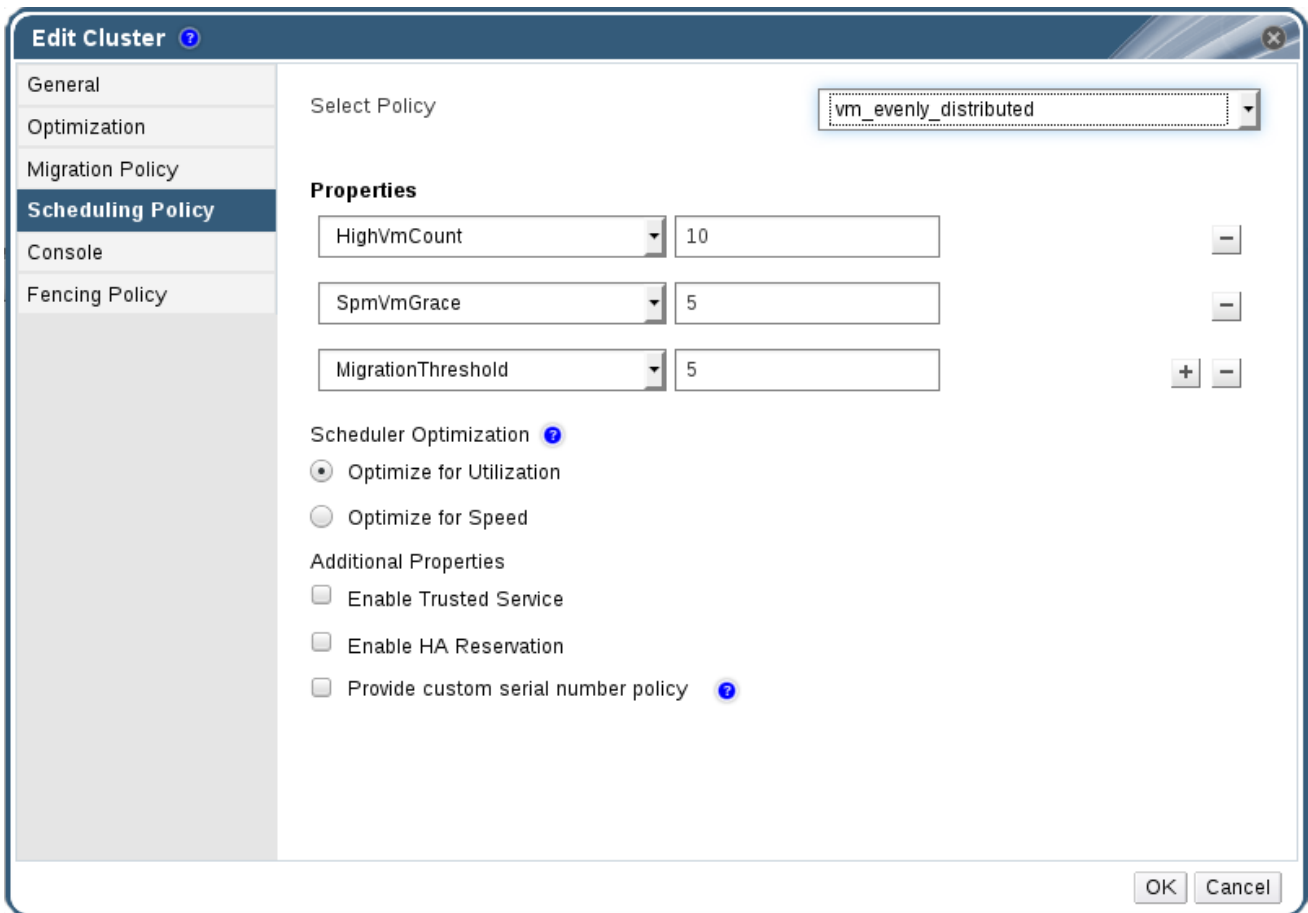



Figure 5.3. Scheduling Policy Settings: vm_evenly_distributed

The table below describes the settings for the **Scheduling Policy** tab.

Table 5.7. Scheduling Policy Tab Properties

Field	Description/Action
Select Policy	<p>Select a policy from the drop-down list.</p> <ul style="list-style-type: none"> none: Set the policy value to none to have no load or power sharing between hosts for already-running virtual machines. This is the default mode. When a virtual machine is started, the memory and CPU processing load is spread evenly across all hosts in the cluster. Additional virtual machines attached to a host will not start if that host has reached the defined CpuOverCommitDurationMinutes, HighUtilization, or MaxFreeMemoryForOverUtilized. evenly_distributed: Distributes the memory and CPU processing load evenly across all hosts in the cluster. Additional virtual machines attached to a host will not start if that host has reached the defined

Field	Description/Action CpuOverCommitDurationMinutes, HighUtilization, or MaxFreeMemoryForOverUtilized.
	<ul style="list-style-type: none"> InClusterUpgrade: Distributes virtual machines based on host operating system version. Hosts with a newer operating system than the virtual machine currently runs on are given priority over hosts with the same operating system. Virtual machines that migrate to a host with a newer operating system will not migrate back to an older operating system. A virtual machine can restart on any host in the cluster. The policy allows hosts in the cluster to be upgraded by allowing the cluster to have mixed operating system versions. Preconditions must be met before the policy can be enabled. See Upgrading Hosts in a Cluster from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7 in the <i>Red Hat Enterprise Virtualization 3.6 Upgrade Guide</i>. <div data-bbox="919 931 1024 1256" style="display: inline-block; vertical-align: top; margin-bottom: 10px;">  </div> <p data-bbox="1107 931 1315 965">IMPORTANT</p> <p data-bbox="1107 1010 1430 1249">The InClusterUpgrade scheduling policy is used for upgrades between major versions only. For example, upgrading from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7.</p> <ul style="list-style-type: none"> power_saving: Distributes the memory and CPU processing load across a subset of available hosts to reduce power consumption on underutilized hosts. Hosts with a CPU load below the low utilization value for longer than the defined time interval will migrate all virtual machines to other hosts so that it can be powered down. Additional virtual machines attached to a host will not start if that host has reached the defined high utilization value. vm_evenly_distributed: Distributes virtual machines evenly between hosts based on a count of the virtual machines. The cluster is considered unbalanced if any host is running more virtual machines than the HighVmCount and there is at least one host with a virtual machine count that falls outside of the MigrationThreshold.

Properties Field	The following properties appear depending on Description/Action, and can be edited if necessary:
	<ul style="list-style-type: none"> • HighVmCount: Sets the minimum number of virtual machines that must be running per host to enable load balancing. The default value is 10 running virtual machines on one host. Load balancing is only enabled when there is at least one host in the cluster that has at least HighVmCount running virtual machines. • MigrationThreshold: Defines a buffer before virtual machines are migrated from the host. It is the maximum inclusive difference in virtual machine count between the most highly-utilized host and the least-utilized host. The cluster is balanced when every host in the cluster has a virtual machine count that falls inside the migration threshold. The default value is 5. • SpmVmGrace: Defines the number of slots for virtual machines to be reserved on SPM hosts. The SPM host will have a lower load than other hosts, so this variable defines how many fewer virtual machines than other hosts it can run. The default value is 5. • CpuOverCommitDurationMinutes: Sets the time (in minutes) that a host can run a CPU load outside of the defined utilization values before the scheduling policy takes action. The defined time interval protects against temporary spikes in CPU load activating scheduling policies and instigating unnecessary virtual machine migration. Maximum two characters. The default value is 2. • HighUtilization: Expressed as a percentage. If the host runs with CPU usage at or above the high utilization value for the defined time interval, the Red Hat Virtualization Manager migrates virtual machines to other hosts in the cluster until the host's CPU load is below the maximum service threshold. The default value is 80. • LowUtilization: Expressed as a percentage. If the host runs with CPU usage below the low utilization value for the defined time interval, the Red Hat Virtualization Manager will migrate virtual machines to other hosts in the cluster. The Manager will power down the original host machine, and restart

Field	Description/Action
	<p>it again when load balancing requires not enough free hosts in the cluster. The default value is 20.</p> <ul style="list-style-type: none"> • ScaleDown: Reduces the impact of the HA Reservation weight function, by dividing a host's score by the specified amount. This is an optional property that can be added to any policy, including none. • HostsInReserve: Specifies a number of hosts to keep running even though there are no running virtual machines on them. This is an optional property that can be added to the power_saving policy. • EnableAutomaticHostPowerManagement: Enables automatic power management for all hosts in the cluster. This is an optional property that can be added to the power_saving policy. The default value is true. • MaxFreeMemoryForOverUtilized: Sets the minimum free memory required in MB for the minimum service level. If the host's available memory runs at, or below this value, the Red Hat Virtualization Manager migrates virtual machines to other hosts in the cluster while the host's available memory is below the minimum service threshold. Setting both MaxFreeMemoryForOverUtilized and MinFreeMemoryForUnderUtilized to 0 MB disables memory based balancing. If MaxFreeMemoryForOverUtilized is set, MinFreeMemoryForUnderUtilized must also be set to avoid unexpected behavior. This is an optional property that can be added to the power_saving and evenly_distributed policies. • MinFreeMemoryForUnderUtilized: Sets the minimum free memory required in MB before the host is considered underutilized. If the host's available memory runs above this value, the Red Hat Virtualization Manager migrates virtual machines to other hosts in the cluster and will automatically power down the host machine, and restart it again when load balancing requires or there are not enough free hosts in the cluster. Setting both

Field	Description/Action
	<p>MaxFreeMemoryForOverUtilized to OMB disables memory based balancing. If MinFreeMemoryForUnderUtilized is set, MaxFreeMemoryForOverUtilized must also be set to avoid unexpected behavior. This is an optional property that can be added to the power_saving and evenly_distributed policies.</p>
Scheduler Optimization	<p>Optimize scheduling for host weighing/ordering.</p> <ul style="list-style-type: none"> • Optimize for Utilization: Includes weight modules in scheduling to allow best selection. • Optimize for Speed: Skips host weighting in cases where there are more than ten pending requests.
Enable Trusted Service	<p>Enable integration with an OpenAttestation server. Before this can be enabled, use the engine-config tool to enter the OpenAttestation server's details. For more information, see Section 10.4, "Trusted Compute Pools".</p>
Enable HA Reservation	<p>Enable the Manager to monitor cluster capacity for highly available virtual machines. The Manager ensures that appropriate capacity exists within a cluster for virtual machines designated as highly available to migrate in the event that their existing host fails unexpectedly.</p>
Provide custom serial number policy	<p>This check box allows you to specify a serial number policy for the virtual machines in the cluster. Select one of the following options:</p> <ul style="list-style-type: none"> • Host ID: Sets the host's UUID as the virtual machine's serial number. • Vm ID: Sets the virtual machine's UUID as its serial number. • Custom serial number: Allows you to specify a custom serial number.

When a host's free memory drops below 20%, ballooning commands like **mom.Controllers.Balloon - INFO Ballooning guest:half1 from 1096400 to 1991580** are logged to **/var/log/vdsm/mom.log**. **/var/log/vdsm/mom.log** is the Memory

Overcommit Manager log file.

5.2.2.5. Cluster Console Settings Explained

The table below describes the settings for the **Console** tab in the **New Cluster** and **Edit Cluster** windows.

Table 5.8. Console Settings

Field	Description/Action
Define SPICE Proxy for Cluster	Select this check box to enable overriding the SPICE proxy defined in global configuration. This feature is useful in a case where the user (who is, for example, connecting via the User Portal) is outside of the network where the hypervisors reside.
Overridden SPICE proxy address	The proxy by which the SPICE client will connect to virtual machines. The address must be in the following format: protocol://[host]:[port]

5.2.2.6. Fencing Policy Settings Explained

The table below describes the settings for the **Fencing Policy** tab in the **New Cluster** and **Edit Cluster** windows.

Table 5.9. Fencing Policy Settings

Field	Description/Action
Enable fencing	Enables fencing on the cluster. Fencing is enabled by default, but can be disabled if required; for example, if temporary network issues are occurring or expected, administrators can disable fencing until diagnostics or maintenance activities are completed. Note that if fencing is disabled, highly available virtual machines running on non-responsive hosts will not be restarted elsewhere.
Skip fencing if host has live lease on storage	If this check box is selected, any hosts in the cluster that are Non Responsive and still connected to storage will not be fenced.

Field	Description/Action
Skip fencing on cluster connectivity issues	If this check box is selected, fencing will be temporarily disabled if the percentage of hosts in the cluster that are experiencing connectivity issues is greater than or equal to the defined Threshold . The Threshold value is selected from the drop-down list; available values are 25, 50, 75, and 100 .

5.2.3. Editing a Resource

Summary

Edit the properties of a resource.

Procedure 5.2. Editing a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click **Edit** to open the **Edit** window.
3. Change the necessary properties and click **OK**.

Result

The new properties are saved to the resource. The **Edit** window will not close if a property field is invalid.

5.2.4. Setting Load and Power Management Policies for Hosts in a Cluster

The **evenly_distributed** and **power_saving** scheduling policies allow you to specify acceptable memory and CPU usage values, and the point at which virtual machines must be migrated to or from a host. The **vm_evenly_distributed** scheduling policy distributes virtual machines evenly between hosts based on a count of the virtual machines. Define the scheduling policy to enable automatic load balancing across the hosts in a cluster. For a detailed explanation of each scheduling policy, see [Section 5.2.2.4, “Scheduling Policy Settings Explained”](#).

Procedure 5.3. Setting Load and Power Management Policies for Hosts

1. Use the resource tabs, tree mode, or the search function to find and select the cluster in the results list.
2. Click **Edit** to open the **Edit Cluster** window.

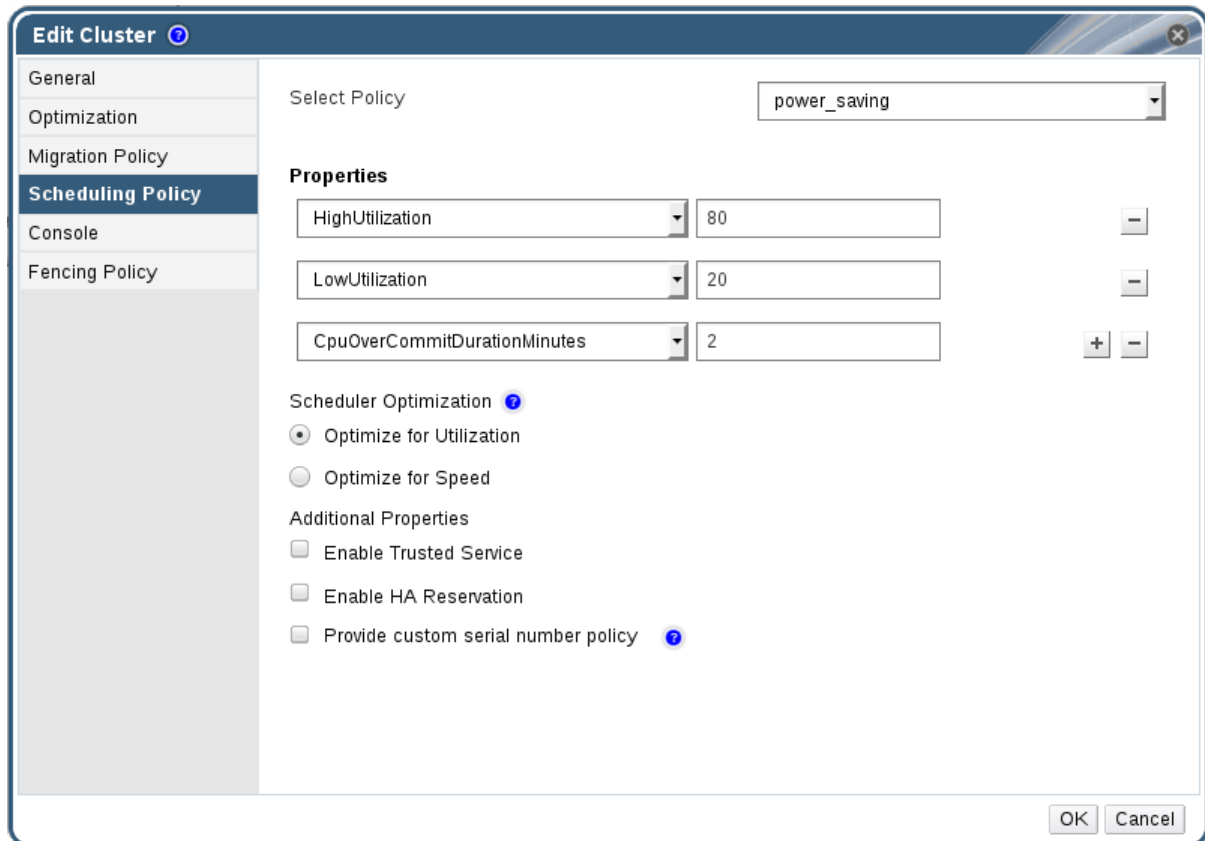


Figure 5.4. Edit Scheduling Policy

3. Select one of the following policies:

- **none**
- **vm_evenly_distributed**
 - a. Set the minimum number of virtual machines that must be running on at least one host to enable load balancing in the **HighVmCount** field.
 - b. Define the maximum acceptable difference between the number of virtual machines on the most highly-utilized host and the number of virtual machines on the least-utilized host in the **MigrationThreshold** field.
 - c. Define the number of slots for virtual machines to be reserved on SPM hosts in the **SpmVmGrace** field.
- **evenly_distributed**
 - a. Set the time (in minutes) that a host can run a CPU load outside of the defined utilization values before the scheduling policy takes action in the **CpuOverCommitDurationMinutes** field.
 - b. Enter the CPU utilization percentage at which virtual machines start migrating to other hosts in the **HighUtilization** field.
 - c. Enter the minimum required free memory in MB above which virtual machines start migrating to other hosts in the **MinFreeMemoryForUnderUtilized**.

- d. Enter the maximum required free memory in MB below which virtual machines start migrating to other hosts in the **MaxFreeMemoryForOverUtilized**.
- **power_saving**
 - a. Set the time (in minutes) that a host can run a CPU load outside of the defined utilization values before the scheduling policy takes action in the **CpuOverCommitDurationMinutes** field.
 - b. Enter the CPU utilization percentage below which the host will be considered under-utilized in the **LowUtilization** field.
 - c. Enter the CPU utilization percentage at which virtual machines start migrating to other hosts in the **HighUtilization** field.
 - d. Enter the minimum required free memory in MB above which virtual machines start migrating to other hosts in the **MinFreeMemoryForUnderUtilized**.
 - e. Enter the maximum required free memory in MB below which virtual machines start migrating to other hosts in the **MaxFreeMemoryForOverUtilized**.
4. Choose one of the following as the **Scheduler Optimization** for the cluster:
 - Select **Optimize for Utilization** to include weight modules in scheduling to allow best selection.
 - Select **Optimize for Speed** to skip host weighting in cases where there are more than ten pending requests.
5. If you are using an OpenAttestation server to verify your hosts, and have set up the server's details using the **engine-config** tool, select the **Enable Trusted Service** check box.
6. Optionally select the **Enable HA Reservation** check box to enable the Manager to monitor cluster capacity for highly available virtual machines.
7. Optionally select the **Provide custom serial number policy** check box to specify a serial number policy for the virtual machines in the cluster, and then select one of the following options:
 - Select **Host ID** to set the host's UUID as the virtual machine's serial number.
 - Select **Vm ID** to set the virtual machine's UUID as its serial number.
 - Select **Custom serial number**, and then specify a custom serial number in the text field.
8. Click **OK**.

5.2.5. Updating the MoM Policy on Hosts in a Cluster

The Memory Overcommit Manager handles memory balloon and KSM functions on a host. Changes to these functions at the cluster level are only passed to hosts the next time a host moves to a status of **Up** after being rebooted or in maintenance mode. However, if

necessary you can apply important changes to a host immediately by synchronizing the MoM policy while the host is **Up**. The following procedure must be performed on each host individually.

Procedure 5.4. Synchronizing MoM Policy on a Host

1. Click the **Clusters** tab and select the cluster to which the host belongs.
2. Click the **Hosts** tab in the details pane and select the host that requires an updated MoM policy.
3. Click **Sync MoM Policy**.

The MoM policy on the host is updated without having to move the host to maintenance mode and back **Up**.

5.2.6. CPU Profiles

CPU profiles define the maximum amount of processing capability a virtual machine in a cluster can access on the host on which it runs, expressed as a percent of the total processing capability available to that host. CPU profiles are created based on CPU profiles defined under data centers, and are not automatically applied to all virtual machines in a cluster; they must be manually assigned to individual virtual machines for the profile to take effect.

5.2.6.1. Creating a CPU Profile

Create a CPU profile. This procedure assumes you have already defined one or more CPU quality of service entries under the data center to which the cluster belongs.

Procedure 5.5. Creating a CPU Profile

1. Click the **Clusters** resource tab and select a cluster.
2. Click the **CPU Profiles** sub tab in the details pane.
3. Click **New**.
4. Enter a name for the CPU profile in the **Name** field.
5. Enter a description for the CPU profile in the **Description** field.
6. Select the quality of service to apply to the CPU profile from the **QoS** list.
7. Click **OK**.

You have created a CPU profile, and that CPU profile can be applied to virtual machines in the cluster.

5.2.6.2. Removing a CPU Profile

Remove an existing CPU profile from your Red Hat Virtualization environment.

Procedure 5.6. Removing a CPU Profile

1. Click the **Clusters** resource tab and select a cluster.
2. Click the **CPU Profiles** sub tab in the details pane.
3. Select the CPU profile to remove.
4. Click **Remove**.
5. Click **OK**.

You have removed a CPU profile, and that CPU profile is no longer available. If the CPU profile was assigned to any virtual machines, those virtual machines are automatically assigned the **default** CPU profile.

5.2.7. Importing an Existing Red Hat Gluster Storage Cluster

You can import a Red Hat Gluster Storage cluster and all hosts belonging to the cluster into Red Hat Virtualization Manager.

When you provide details such as the IP address or host name and password of any host in the cluster, the **gluster peer status** command is executed on that host through SSH, then displays a list of hosts that are a part of the cluster. You must manually verify the fingerprint of each host and provide passwords for them. You will not be able to import the cluster if one of the hosts in the cluster is down or unreachable. As the newly imported hosts do not have VDSM installed, the bootstrap script installs all the necessary VDSM packages on the hosts after they have been imported, and reboots them.

Procedure 5.7. Importing an Existing Red Hat Gluster Storage Cluster to Red Hat Virtualization Manager

1. Select the **Clusters** resource tab to list all clusters in the results list.
2. Click **New** to open the **New Cluster** window.
3. Select the **Data Center** the cluster will belong to from the drop-down menu.
4. Enter the **Name** and **Description** of the cluster.
5. Select the **Enable Gluster Service** radio button and the **Import existing gluster configuration** check box.

The **Import existing gluster configuration** field is displayed only if you select **Enable Gluster Service** radio button.

6. In the **Address** field, enter the hostname or IP address of any server in the cluster.

The host **Fingerprint** displays to ensure you are connecting with the correct host. If a host is unreachable or if there is a network error, an error **Error in fetching fingerprint** displays in the **Fingerprint** field.

7. Enter the **Root Password** for the server, and click **OK**.
8. The **Add Hosts** window opens, and a list of hosts that are a part of the cluster displays.
9. For each host, enter the **Name** and the **Root Password**.

10. If you wish to use the same password for all hosts, select the **Use a Common Password** check box to enter the password in the provided text field.

Click **Apply** to set the entered password all hosts.

Make sure the fingerprints are valid and submit your changes by clicking **OK**.

The bootstrap script installs all the necessary VDSM packages on the hosts after they have been imported, and reboots them. You have now successfully imported an existing Red Hat Gluster Storage cluster into Red Hat Virtualization Manager.

5.2.8. Explanation of Settings in the Add Hosts Window

The **Add Hosts** window allows you to specify the details of the hosts imported as part of a Gluster-enabled cluster. This window appears after you have selected the **Enable Gluster Service** check box in the **New Cluster** window and provided the necessary host details.

Table 5.10. Add Gluster Hosts Settings

Field	Description
Use a common password	Tick this check box to use the same password for all hosts belonging to the cluster. Enter the password in the Password field, then click the Apply button to set the password on all hosts.
Name	Enter the name of the host.
Hostname/IP	This field is automatically populated with the fully qualified domain name or IP of the host you provided in the New Cluster window.
Root Password	Enter a password in this field to use a different root password for each host. This field overrides the common password provided for all hosts in the cluster.
Fingerprint	The host fingerprint is displayed to ensure you are connecting with the correct host. This field is automatically populated with the fingerprint of the host you provided in the New Cluster window.

5.2.9. Removing a Cluster

Summary

Move all hosts out of a cluster before removing it.



NOTE

You cannot remove the **Default** cluster, as it holds the **Blank** template. You can however rename the **Default** cluster and add it to a new data center.

Procedure 5.8. Removing a Cluster

1. Use the resource tabs, tree mode, or the search function to find and select the cluster in the results list.
2. Ensure there are no hosts in the cluster.
3. Click **Remove** to open the **Remove Cluster(s)** confirmation window.
4. Click **OK**

Result

The cluster is removed.

5.2.10. Changing the Cluster Compatibility Version

Red Hat Virtualization clusters have a compatibility version. The cluster compatibility version indicates the features of Red Hat Virtualization supported by all of the hosts in the cluster. The cluster compatibility is set according to the version of the least capable host operating system in the cluster.



NOTE

To change the cluster compatibility version, you must have first updated all the hosts in your cluster to a level that supports your desired compatibility level.

After you update the cluster compatibility version of the cluster you need to update the cluster compatibility version of all running or suspended virtual machines to ensure that the changes become effective. This is achieved by restarting the virtual machines from within the Manager or REST API call instead of within the guest operating system. Virtual machines will continue to run in the previous cluster compatibility level until they are restarted. Those virtual machines that require a restart are marked with the **Next-Run** icon (triangle with an exclamation mark). You cannot change the cluster compatibility version of a virtual machine snapshot that is in preview, you need to first commit or undo the preview.

The self-hosted engine virtual machine does not need to be restarted, see [Maintenance and Upgrading Resources](#) in the *Self-Hosted Engine Guide* for more information about upgrading the Self-Hosted Engine environment.

Procedure 5.9. Changing the Cluster Compatibility Version

1. From the Administration Portal, click the **Clusters** tab.
2. Select the cluster to change from the list displayed.
3. Click **Edit**.
4. Change the **Compatibility Version** to the desired value.
5. Click **OK** to open the **Change Cluster Compatibility Version** confirmation window.
6. Click **OK** to confirm.

You have updated the compatibility version of the cluster. Once you have updated the compatibility version of all clusters in a data center, you can then change the compatibility version of the data center itself.



IMPORTANT

Upgrading the compatibility will also upgrade all of the storage domains belonging to the data center.

5.3. CLUSTERS AND PERMISSIONS

5.3.1. Managing System Permissions for a Cluster

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A cluster administrator is a system administration role for a specific data center only. This is useful in data centers with multiple clusters, where each cluster requires a system administrator. The **ClusterAdmin** role is a hierarchical model: a user assigned the cluster administrator role for a cluster can manage all objects in the cluster. Use the **Configure** button in the header bar to assign a cluster administrator for all clusters in the environment.

The cluster administrator role permits the following actions:

- Create and remove associated clusters.
- Add and remove hosts, virtual machines, and pools associated with the cluster.
- Edit user permissions for virtual machines associated with the cluster.



NOTE

You can only assign roles and permissions to existing users.

You can also change the system administrator of a cluster by removing the existing system administrator and adding the new system administrator.

5.3.2. Cluster Administrator Roles Explained

Cluster Permission Roles

The table below describes the administrator roles and privileges applicable to cluster administration.

Table 5.11. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
ClusterAdmin	Cluster Administrator	<p>Can use, create, delete, manage all physical and virtual resources in a specific cluster, including hosts, templates and virtual machines. Can configure network properties within the cluster such as designating display networks, or marking a network as required or non-required.</p> <p>However, a ClusterAdmin does not have permissions to attach or detach networks from a cluster, to do so NetworkAdmin permissions are required.</p>
NetworkAdmin	Network Administrator	<p>Can configure and manage the network of a particular cluster. A network administrator of a cluster inherits network permissions for virtual machines within the cluster as well.</p>

5.3.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 5.10. Assigning a Role to a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Click **Add**.
4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down list.
6. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

5.3.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 5.11. Removing a Role from a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

CHAPTER 6. LOGICAL NETWORKS

6.1. LOGICAL NETWORK TASKS

6.1.1. Using the Networks Tab

The **Networks** resource tab provides a central location for users to perform logical network-related operations and search for logical networks based on each network's property or association with other resources.

All logical networks in the Red Hat Virtualization environment display in the results list of the **Networks** tab. The **New**, **Edit** and **Remove** buttons allow you to create, change the properties of, and delete logical networks within data centers.

Click on each network name and use the **Clusters**, **Hosts**, **Virtual Machines**, **Templates**, and **Permissions** tabs in the details pane to perform functions including:

- Attaching or detaching the networks to clusters and hosts
- Removing network interfaces from virtual machines and templates
- Adding and removing permissions for users to access and manage networks

These functions are also accessible through each individual resource tab.



WARNING

Do not change networking in a data center or a cluster if any hosts are running as this risks making the host unreachable.

IMPORTANT

If you plan to use Red Hat Virtualization nodes to provide any services, remember that the services will stop if the Red Hat Virtualization environment stops operating.

This applies to all services, but you should be especially aware of the hazards of running the following on Red Hat Virtualization:

- Directory Services
- DNS
- Storage

6.1.2. Creating a New Logical Network in a Data Center or Cluster

Create a logical network and define its use in a data center, or in clusters in a data center.

Procedure 6.1. Creating a New Logical Network in a Data Center or Cluster

1. Click the **Data Centers** or **Clusters** resource tabs, and select a data center or cluster in the results list.
2. Click the **Logical Networks** tab of the details pane to list the existing logical networks.
3.
 - From the **Data Centers** details pane, click **New** to open the **New Logical Network** window.
 - From the **Clusters** details pane, click **Add Network** to open the **New Logical Network** window.
4. Enter a **Name**, **Description**, and **Comment** for the logical network.
5. Optionally select the **Create on external provider** check box. Select the **External Provider** from the drop-down list and provide the IP address of the **Physical Network**. The **External Provider** drop-down list will not list any external providers in read-only mode.

If **Create on external provider** is selected, the **Network Label**, **VM Network**, and **MTU** options are disabled.
6. Enter a new label or select an existing label for the logical network in the **Network Label** text field.
7. Optionally enable **Enable VLAN tagging**.
8. Optionally disable **VM Network**.
9. Set the **MTU** value to **Default (1500)** or **Custom**.
10. From the **Cluster** tab, select the clusters to which the network will be assigned. You can also specify whether the logical network will be a required network.
11. If **Create on external provider** is selected, the **Subnet** tab will be visible. From the **Subnet** tab, select the **Create subnet** and enter a **Name**, **CIDR**, and **Gateway** address, and select an **IP Version** for the subnet that the logical network will provide. You can also add DNS servers as required.
12. From the **vNIC Profiles** tab, add vNIC profiles to the logical network as required.
13. Click **OK**.

You have defined a logical network as a resource required by a cluster or clusters in the data center. If you entered a label for the logical network, it will be automatically added to all host network interfaces with that label.


**NOTE**

When creating a new logical network or making changes to an existing logical network that is used as a display network, any running virtual machines that use that network must be rebooted before the network becomes available or the changes are applied.

6.1.3. Editing a Logical Network

Edit the settings of a logical network.


Procedure 6.2. Editing a Logical Network



IMPORTANT

A logical network cannot be edited or moved to another interface if it is not synchronized with the network configuration on the host. See [Section 6.5.2, “Editing Host Network Interfaces and Assigning Logical Networks to Hosts”](#) on how to synchronize your networks.

1. Click the **Data Centers** resource tab, and select the data center of the logical network in the results list.
2. Click the **Logical Networks** tab in the details pane to list the logical networks in the data center.
3. Select a logical network and click **Edit** to open the **Edit Logical Network** window.
4. Edit the necessary settings.



NOTE

You can edit the name of a new or existing network, with the exception of the default network, without having to stop the virtual machines.

5. Click **OK** to save the changes.



NOTE

Multi-host network configuration automatically applies updated network settings to all of the hosts within the data center to which the network is assigned. Changes can only be applied when virtual machines using the network are down. You cannot rename a logical network that is already configured on a host. You cannot disable the **VM Network** option while virtual machines or templates using that network are running.

6.1.4. Removing a Logical Network

You can remove a logical network from the **Networks** resource tab or the **Data Centers** resource tab. The following procedure shows you how to remove logical networks associated to a data center. For a working Red Hat Virtualization environment, you must have at least one logical network used as the **ovirtmgmt** management network.

Procedure 6.3. Removing Logical Networks

1. Click the **Data Centers** resource tab, and select the data center of the logical network in the results list.
2. Click the **Logical Networks** tab in the details pane to list the logical networks in the data center.

3. Select a logical network and click **Remove** to open the **Remove Logical Network(s)** window.
4. Optionally, select the **Remove external network(s) from the provider(s) as well** check box to remove the logical network both from the Manager and from the external provider if the network is provided by an external provider. The check box is grayed out if the external provider is in read-only mode.
5. Click **OK**.

The logical network is removed from the Manager and is no longer available.

6.1.5. Viewing or Editing the Gateway for a Logical Network

Users can define the gateway, along with the IP address and subnet mask, for a logical network. This is necessary when multiple networks exist on a host and traffic should be routed through the specified network, rather than the default gateway.

If multiple networks exist on a host and the gateways are not defined, return traffic will be routed through the default gateway, which may not reach the intended destination. This would result in users being unable to ping the host.

Red Hat Virtualization handles multiple gateways automatically whenever an interface goes up or down.

Procedure 6.4. Viewing or Editing the Gateway for a Logical Network

1. Click the **Hosts** resource tab, and select the desired host.
2. Click the **Network Interfaces** tab in the details pane to list the network interfaces attached to the host and their configurations.
3. Click the **Setup Host Networks** button to open the **Setup Host Networks** window.
4. Hover your cursor over an assigned logical network and click the pencil icon to open the **Edit Management Network** window.

The **Edit Management Network** window displays the network name, the boot protocol, and the IP, subnet mask, and gateway addresses. The address information can be manually edited by selecting a **Static** boot protocol.

6.1.6. Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows

6.1.6.1. Logical Network General Settings Explained

The table below describes the settings for the **General** tab of the **New Logical Network** and **Edit Logical Network** window.

Table 6.1. New Logical Network and Edit Logical Network Settings

Field Name	Description
------------	-------------

Field Name	Description
Name	The name of the logical network. This text field has a 15-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Description	The description of the logical network. This text field has a 40-character limit.
Comment	A field for adding plain text, human-readable comments regarding the logical network.
Create on external provider	Allows you to create the logical network to an OpenStack Networking instance that has been added to the Manager as an external provider. External Provider - Allows you to select the external provider on which the logical network will be created.
Enable VLAN tagging	VLAN tagging is a security feature that gives all network traffic carried on the logical network a special characteristic. VLAN-tagged traffic cannot be read by interfaces that do not also have that characteristic. Use of VLANs on logical networks also allows a single network interface to be associated with multiple, differently VLAN-tagged logical networks. Enter a numeric value in the text entry field if VLAN tagging is enabled.
VM Network	Select this option if only virtual machines use this network. If the network is used for traffic that does not involve virtual machines, such as storage communications, do not select this check box.
MTU	Choose either Default , which sets the maximum transmission unit (MTU) to the value given in the parenthesis (), or Custom to set a custom MTU for the logical network. You can use this to match the MTU supported by your new logical network to the MTU supported by the hardware it interfaces with. Enter a numeric value in the text entry field if Custom is selected.
Network Label	Allows you to specify a new label for the network or select from existing labels already attached to host network interfaces. If you select an existing label, the logical network will be automatically assigned to all host network interfaces with that label.

6.1.6.2. Logical Network Cluster Settings Explained

The table below describes the settings for the **Cluster** tab of the **New Logical Network** window.

Table 6.2. New Logical Network Settings

Field Name	Description
Attach/Detach Network to/from Cluster(s)	<p>Allows you to attach or detach the logical network from clusters in the data center and specify whether the logical network will be a required network for individual clusters.</p> <p>Name - the name of the cluster to which the settings will apply. This value cannot be edited.</p> <p>Attach All - Allows you to attach or detach the logical network to or from all clusters in the data center. Alternatively, select or clear the Attach check box next to the name of each cluster to attach or detach the logical network to or from a given cluster.</p> <p>Required All - Allows you to specify whether the logical network is a required network on all clusters. Alternatively, select or clear the Required check box next to the name of each cluster to specify whether the logical network is a required network for a given cluster.</p>

6.1.6.3. Logical Network vNIC Profiles Settings Explained

The table below describes the settings for the **vNIC Profiles** tab of the **New Logical Network** window.

Table 6.3. New Logical Network Settings

Field Name	Description
vNIC Profiles	<p>Allows you to specify one or more vNIC profiles for the logical network. You can add or remove a vNIC profile to or from the logical network by clicking the plus or minus button next to the vNIC profile. The first field is for entering a name for the vNIC profile.</p> <p>Public - Allows you to specify whether the profile is available to all users.</p> <p>QoS - Allows you to specify a network quality of service (QoS) profile to the vNIC profile.</p>

6.1.7. Designate a Specific Traffic Type for a Logical Network with the Manage Networks Window

Specify the traffic type for the logical network to optimize the network traffic flow.

Procedure 6.5. Specifying Traffic Types for Logical Networks

1. Click the **Clusters** resource tab, and select a cluster from the results list.
2. Select the **Logical Networks** tab in the details pane to list the logical networks assigned to the cluster.
3. Click **Manage Networks** to open the **Manage Networks** window.

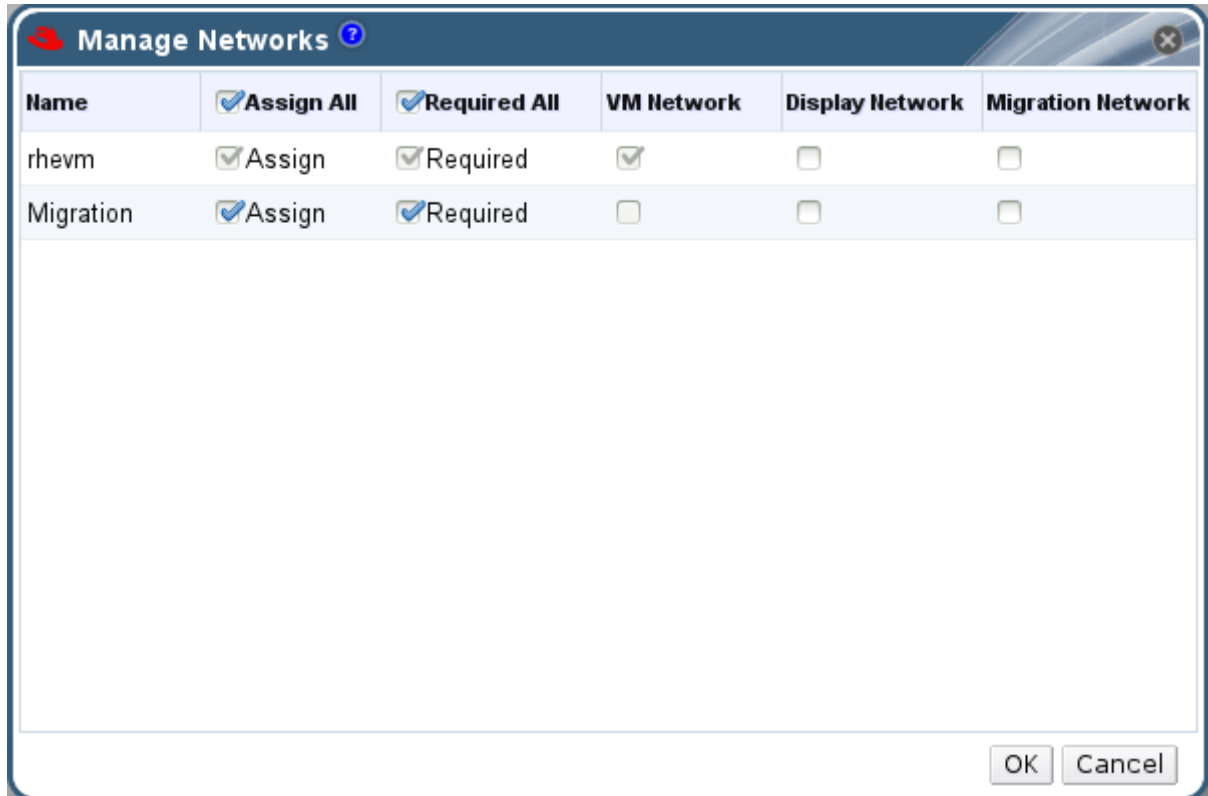


Figure 6.1. Manage Networks

4. Select appropriate check boxes.
5. Click **OK** to save the changes and close the window.

You have optimized the network traffic flow by assigning a specific type of traffic to be carried on a specific logical network.



NOTE

Logical networks offered by external providers must be used as virtual machine networks; they cannot be assigned special cluster roles such as display or migration.

6.1.8. Explanation of Settings in the Manage Networks Window

The table below describes the settings for the **Manage Networks** window.

Table 6.4. Manage Networks Settings

Field	Description/Action
Assign	Assigns the logical network to all hosts in the cluster.
Required	A Network marked "required" must remain operational in order for the hosts associated with it to function properly. If a required network ceases to function, any hosts associated with it become non-operational.
VM Network	A logical network marked "VM Network" carries network traffic relevant to the virtual machine network.
Display Network	A logical network marked "Display Network" carries network traffic relevant to SPICE and to the virtual network controller.
Migration Network	A logical network marked "Migration Network" carries virtual machine and storage migration traffic.

6.1.9. Editing the Virtual Function Configuration on a NIC


Single Root I/O Virtualization (SR-IOV) enables a single PCIe endpoint to be used as multiple separate devices. This is achieved through the introduction of two PCIe functions: physical functions (PFs) and virtual functions (VFs). A PCIe card can have between one and eight PFs, but each PF can support many more VFs (dependent on the device).

You can edit the configuration of SR-IOV-capable Network Interface Controllers (NICs) through the Red Hat Virtualization Manager, including the number of VFs on each NIC and to specify the virtual networks allowed to access the VFs.


Once VFs have been created, each can be treated as a standalone NIC. This includes having one or more logical networks assigned to them, creating bonded interfaces with them, and to directly assign vNICs to them for direct device passthrough.

A vNIC must have the passthrough property enabled in order to be directly attached to a VF. See [Section 6.2.4, “Enabling Passthrough on a vNIC Profile”](#).

Procedure 6.6. Editing the Virtual Function Configuration on a NIC

1. Select an SR-IOV-capable host and click the **Network Interfaces** tab in the details pane.
2. Click **Setup Host Networks** to open the **Setup Host Networks** window.
3. Select an SR-IOV-capable NIC, marked with a , and click the pencil icon to open the **Edit Virtual Functions (SR-IOV) configuration of NIC** window.

4. To edit the number of virtual functions, click the **Number of VFs setting** drop-down button and edit the **Number of VFs** text field.



IMPORTANT

Changing the number of VFs will delete all previous VFs on the network interface before creating new VFs. This includes any VFs that have virtual machines directly attached.

5. The **All Networks** check box is selected by default, allowing all networks to access the virtual functions. To specify the virtual networks allowed to access the virtual functions, select the **Specific networks** radio button to list all networks. You can then either select the check box for desired networks, or you can use the **Labels** text field to automatically select networks based on one or more network labels.
6. Click **OK** to close the window. Note that the configuration changes will not take effect until you click the **OK** button in the **Setup Host Networks** window.

6.2. VIRTUAL NETWORK INTERFACE CARDS

6.2.1. vNIC Profile Overview

A Virtual Network Interface Card (vNIC) profile is a collection of settings that can be applied to individual virtual network interface cards in the Manager. A vNIC profile allows you to apply Network QoS profiles to a vNIC, enable or disable port mirroring, and add or remove custom properties. A vNIC profile also offers an added layer of administrative flexibility in that permission to use (consume) these profiles can be granted to specific users. In this way, you can control the quality of service that different users receive from a given network.

6.2.2. Creating or Editing a vNIC Profile

Create or edit a Virtual Network Interface Controller (vNIC) profile to regulate network bandwidth for users and groups.



NOTE

If you are enabling or disabling port mirroring, all virtual machines using the associated profile must be in a down state before editing.

Procedure 6.7. Creating or editing a vNIC Profile

1. Click the **Networks** resource tab, and select a logical network in the results list.
2. Select the **vNIC Profiles** tab in the details pane. If you selected the logical network in tree mode, you can select the **vNIC Profiles** tab in the results list.
3. Click **New** or **Edit** to open the **VM Interface Profile** window.





Figure 6.2. The VM Interface Profile window

4. Enter the **Name** and **Description** of the profile.
5. Select the relevant Quality of Service policy from the **QoS** list.
6. Select a **Network Filter** from the drop-down list to manage the traffic of network packets to and from virtual machines. For more information on network filters, see [Applying network filtering](#) in the *Red Hat Enterprise Linux Virtualization Deployment and Administration Guide*.
7. Select the **Passthrough** check box to enable passthrough of the vNIC and allow direct device assignment of a virtual function. Enabling the passthrough property will disable QoS, network filtering, and port mirroring as these are not compatible. For more information on passthrough, see [Section 6.2.4, “Enabling Passthrough on a vNIC Profile”](#).
8. Use the **Port Mirroring** and **Allow all users to use this Profile** check boxes to toggle these options.
9. Select a custom property from the custom properties list, which displays **Please select a key...** by default. Use the + and - buttons to add or remove custom properties.
10. Click **OK**.

You have created a vNIC profile. Apply this profile to users and groups to regulate their network bandwidth. Note that if you edited a vNIC profile, you must either restart the virtual machine or hot unplug and then hot plug the vNIC.

**NOTE**

The guest operating system must support vNIC hot plug and hot unplug.

6.2.3. Explanation of Settings in the VM Interface Profile Window

Table 6.5. VM Interface Profile Window

Field Name	Description
Network	A drop-down list of the available networks to apply the vNIC profile to.
Name	The name of the vNIC profile. This must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores between 1 and 50 characters.
Description	The description of the vNIC profile. This field is recommended but not mandatory.
QoS	A drop-down list of the available Network Quality of Service policies to apply to the vNIC profile. QoS policies regulate inbound and outbound network traffic of the vNIC.
Network Filter	<p>A drop-down list of the available network filters to apply to the vNIC profile. Network filters improve network security by filtering the type of packets that can be sent to and from virtual machines. The default filter is vdsm-no-mac-spoofing, which is a combination of no-mac-spoofing and no-arp-mac-spoofing. For more information on the network filters provided by libvirt, see the Pre-existing network filters section of the <i>Red Hat Enterprise Linux Virtualization Deployment and Administration Guide</i>.</p> <p><No Network Filter> should be used for virtual machine VLANs and bonds. On trusted virtual machines, choosing not to use a network filter can improve performance.</p>
Passthrough	<p>A check box to toggle the passthrough property. Passthrough allows a vNIC to connect directly to a virtual function of a host NIC. The passthrough property cannot be edited if the vNIC profile is attached to a virtual machine.</p> <p>QoS, network filters, and port mirroring are disabled in the vNIC profile if passthrough is enabled.</p>

Field Name	Description
Port Mirroring	A check box to toggle port mirroring. Port mirroring copies layer 3 network traffic on the logical network to a virtual interface on a virtual machine. It is not selected by default. For further details, see Port Mirroring in the <i>Technical Reference</i> .
Device Custom Properties	A drop-down menu to select available custom properties to apply to the vNIC profile. Use the + and - buttons to add and remove properties respectively.
Allow all users to use this Profile	A check box to toggle the availability of the profile to all users in the environment. It is selected by default.

6.2.4. Enabling Passthrough on a vNIC Profile

The passthrough property of a vNIC profile enables a vNIC to be directly connected to a virtual function (VF) of an SR-IOV-enabled NIC. The vNIC will then bypass the software network virtualization and connect directly to the VF for direct device assignment.

The passthrough property cannot be enabled if the vNIC profile is already attached to a vNIC; this procedure creates a new profile to avoid this. If a vNIC profile has passthrough enabled, QoS and port mirroring are disabled for the profile.

For more information on SR-IOV, direct device assignment, and the hardware considerations for implementing these in Red Hat Virtualization, see [Hardware Considerations for Implementing SR-IOV](#).

Procedure 6.8. Enabling Passthrough

1. Select a logical network from the **Networks** results list and click the **vNIC Profiles** tab in the details pane to list all vNIC profiles for that logical network.
2. Click **New** to open the **VM Interface Profile** window.
3. Enter the **Name** and **Description** of the profile.
4. Select the **Passthrough** check box. This will disable **QoS** and **Port Mirroring**.
5. If necessary, select a custom property from the custom properties list, which displays **Please select a key...** by default. Use the **+** and **-** buttons to add or remove custom properties.
6. Click **OK** to save the profile and close the window.

The vNIC profile is now passthrough-capable. To use this profile to directly attach a virtual machine to a NIC or PCI VF, attach the logical network to the NIC and create a new vNIC on the desired virtual machine that uses the passthrough vNIC profile. For more information on

these procedures respectively, see [Section 6.5.2, “Editing Host Network Interfaces and Assigning Logical Networks to Hosts”](#), and [Adding a New Network Interface](#) in the *Virtual Machine Management Guide*.


6.2.5. Removing a vNIC Profile

Remove a vNIC profile to delete it from your virtualized environment.

Procedure 6.9. Removing a vNIC Profile

1. Click the **Networks** resource tab, and select a logical network in the results list.
2. Select the **Profiles** tab in the details pane to display available vNIC profiles. If you selected the logical network in tree mode, you can select the **vNIC Profiles** tab in the results list.
3. Select one or more profiles and click **Remove** to open the **Remove VM Interface Profile(s)** window.
4. Click **OK** to remove the profile and close the window.

6.2.6. Assigning Security Groups to vNIC Profiles



NOTE

This feature is only available for users who are integrating with OpenStack Neutron. Security groups cannot be created with Red Hat Virtualization Manager. You must create security groups within OpenStack. For more information, see <https://access.redhat.com/documentation/en/red-hat-openstack-platform/9/single/users-and-identity-management-guide#project-security> in the *Red Hat OpenStack Platform Users and Identity Management Guide*.

You can assign security groups to the vNIC profile of networks that have been imported from an OpenStack Networking instance and that use the Open vSwitch plug-in. A security group is a collection of strictly enforced rules that allow you to filter inbound and outbound traffic over a network interface. The following procedure outlines how to attach a security group to a vNIC profile.



NOTE

A security group is identified using the ID of that security group as registered in the OpenStack Networking instance. You can find the IDs of security groups for a given tenant by running the following command on the system on which OpenStack Networking is installed:

```
# neutron security-group-list
```

Procedure 6.10. Assigning Security Groups to vNIC Profiles

1. Click the **Networks** tab and select a logical network from the results list.
2. Click the **vNIC Profiles** tab in the details pane.

3. Click **New**, or select an existing vNIC profile and click **Edit**, to open the **VM Interface Profile** window.
4. From the custom properties drop-down list, select **SecurityGroups**. Leaving the custom property drop-down blank applies the default security settings, which permit all outbound traffic and intercommunication but deny all inbound traffic from outside of the default security group. Note that removing the **SecurityGroups** property later will not affect the applied security group.
5. In the text field, enter the ID of the security group to attach to the vNIC profile.
6. Click **OK**.

You have attached a security group to the vNIC profile. All traffic through the logical network to which that profile is attached will be filtered in accordance with the rules defined for that security group.

6.2.7. User Permissions for vNIC Profiles

Configure user permissions to assign users to certain vNIC profiles. Assign the **VnicProfileUser** role to a user to enable them to use the profile. Restrict users from certain profiles by removing their permission for that profile.

Procedure 6.11. User Permissions for vNIC Profiles

1. Click the **Networks** tab and select a logical network from the results list.
2. Click the **vNIC Profiles** resource tab to display the vNIC profiles.
3. Click the **Permissions** tab in the details pane to show the current user permissions for the profile.
4. Click the **Add** button to open the **Add Permission to User** window, and the **Remove** button to open the **Remove Permission** window, to change user permissions for the vNIC profile.
5. In the **Add Permissions to User** window, click **My Groups** to display your user groups. You can use this option to grant permissions to other users in your groups.

You have configured user permissions for a vNIC profile.

6.2.8. Configuring vNIC Profiles for UCS Integration

Cisco's Unified Computing System (UCS) is used to manage datacenter aspects such as computing, networking and storage resources.

The **vds-hook-vmfex-dev** hook allows virtual machines to connect to Cisco's UCS-defined port profiles by configuring the vNIC profile. The UCS-defined port profiles contain the properties and settings used to configure virtual interfaces in UCS. The **vds-hook-vmfex-dev** hook is installed by default with VDSM. See [Appendix A, VDSM and Hooks](#) for more information.

When a virtual machine that uses the vNIC profile is created, it will use the Cisco vNIC.

The procedure to configure the vNIC profile for UCS integration involves first configuring a custom device property. When configuring the custom device property, any existing value

it contained is overwritten. When combining new and existing custom properties, include all of the custom properties in the command used to set the key's value. Multiple custom properties are separated by a semi-colon.

**NOTE**

A UCS port profile must be configured in Cisco UCS before configuring the vNIC profile.

Procedure 6.12. Configuring the Custom Device Property

1. On the Red Hat Virtualization Manager, configure the **vmfex** custom property and set the cluster compatibility level using **--cver**.

```
# engine-config -s CustomDeviceProperties='{type=interface;prop={vmfex=^[a-zA-Z0-9_.-]{2,32}$}}' --cver=3.6
```

2. Verify that the **vmfex** custom device property was added.

```
# engine-config -g CustomDeviceProperties
```

3. Restart the engine.

```
# systemctl restart ovirt-engine.service
```

The vNIC profile to configure can belong to a new or existing logical network. See [Section 6.1.2, “Creating a New Logical Network in a Data Center or Cluster”](#) for instructions to configure a new logical network.

Procedure 6.13. Configuring a vNIC Profile for UCS Integration

1. Click the **Networks** resource tab, and select a logical network in the results list.
2. Select the **vNIC Profiles** tab in the details pane. If you selected the logical network in tree mode, you can select the **vNIC Profiles** tab in the results list.
3. Click **New** or **Edit** to open the **VM Interface Profile** window.
4. Enter the **Name** and **Description** of the profile.
5. Select the **vmfex** custom property from the custom properties list and enter the UCS port profile name.
6. Click **OK**.

6.3. EXTERNAL PROVIDER NETWORKS

6.3.1. Importing Networks From External Providers

To use networks from an external network provider (OpenStack Networking or any third-party provider that implements the OpenStack Neutron REST API), register the provider with the Manager. See [Section 12.2.3, “Adding an OpenStack Networking \(Neutron\) Instance for Network Provisioning”](#) or [Section 12.2.6, “Adding an External Network](#)

[Provider](#)” for more information. Then, use the following procedure to import the networks provided by that provider into the Manager so the networks can be used by virtual machines.

Procedure 6.14. Importing a Network From an External Provider

1. Click the **Networks** tab.
2. Click the **Import** button to open the **Import Networks** window.

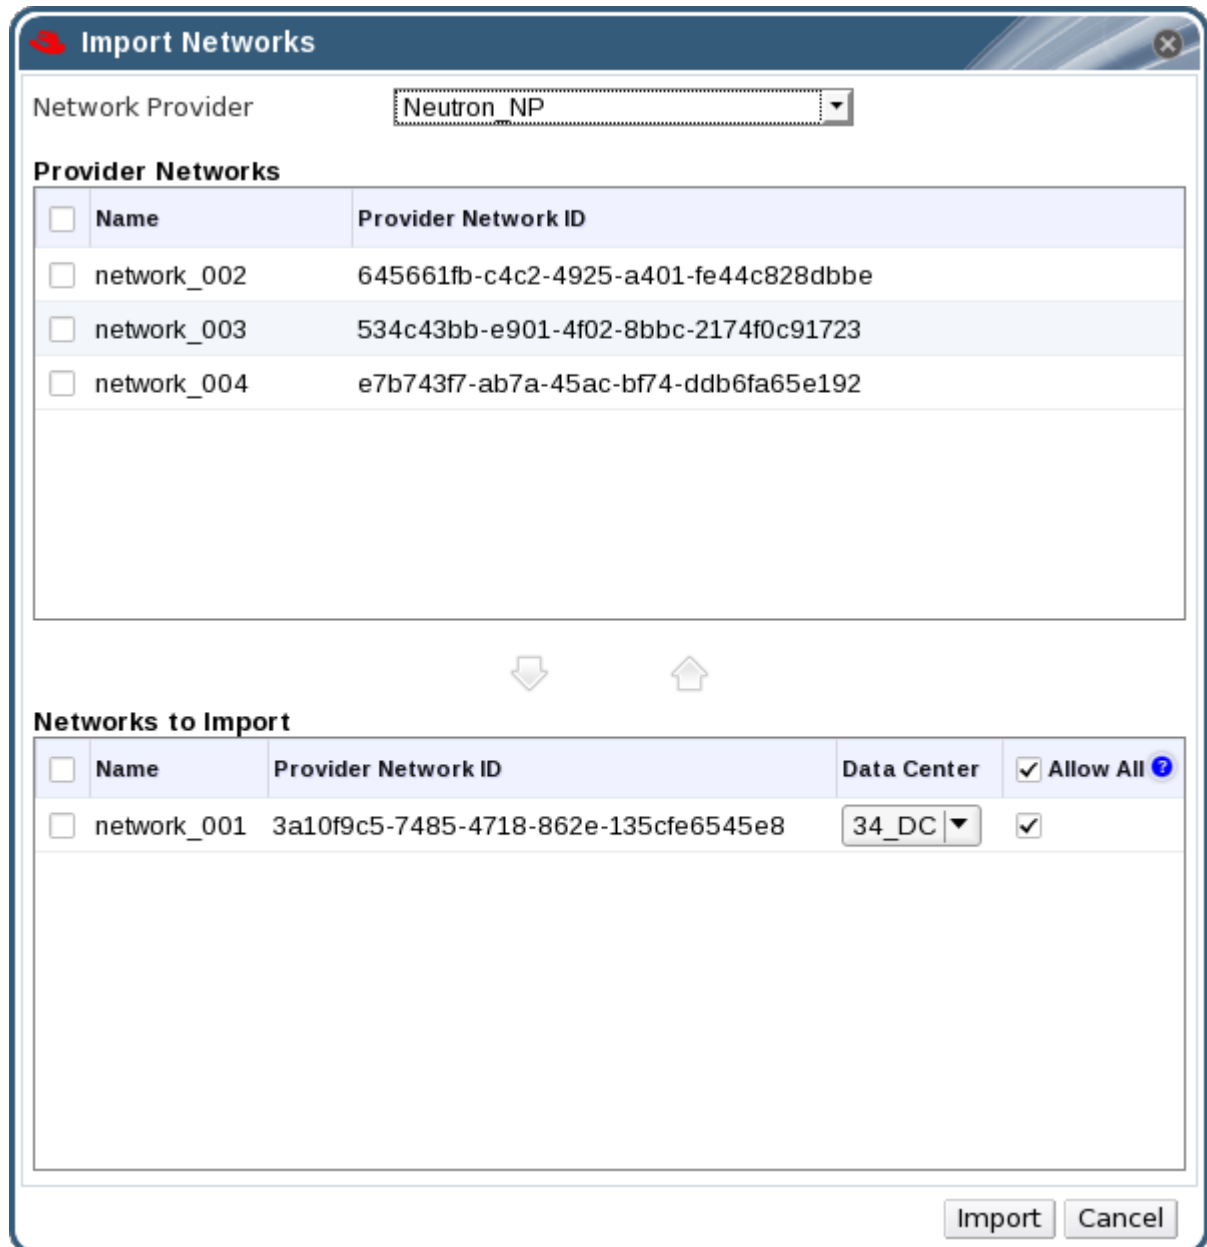


Figure 6.3. The Import Networks Window

3. From the **Network Provider** drop-down list, select an external provider. The networks offered by that provider are automatically discovered and listed in the **Provider Networks** list.
4. Using the check boxes, select the networks to import in the **Provider Networks** list and click the down arrow to move those networks into the **Networks to Import** list.

5. It is possible to customize the name of the network that you are importing. To customize the name, click on the network's name in the **Name** column, and change the text.
6. From the **Data Center** drop-down list, select the data center into which the networks will be imported.
7. Optionally, clear the **Allow All** check box for a network in the **Networks to Import** list to prevent that network from being available to all users.
8. Click the **Import** button.

The selected networks are imported into the target data center and can be attached to virtual machines. See [Adding a New Network Interface](#) in the *Virtual Machine Management Guide* for more information.

6.3.2. Limitations to Using External Provider Networks

The following limitations apply to using logical networks imported from an external provider in a Red Hat Virtualization environment.

- Logical networks offered by external providers must be used as virtual machine networks, and cannot be used as display networks.
- The same logical network can be imported more than once, but only to different data centers.
- You cannot edit logical networks offered by external providers in the Manager. To edit the details of a logical network offered by an external provider, you must edit the logical network directly from the external provider that provides that logical network.
- Port mirroring is not available for virtual network interface cards connected to logical networks offered by external providers.
- If a virtual machine uses a logical network offered by an external provider, that provider cannot be deleted from the Manager while the logical network is still in use by the virtual machine.
- Networks offered by external providers are non-required. As such, scheduling for clusters in which such logical networks have been imported will not take those logical networks into account during host selection. Moreover, it is the responsibility of the user to ensure the availability of the logical network on hosts in clusters in which such logical networks have been imported.

6.3.3. Configuring Subnets on External Provider Logical Networks

6.3.3.1. Configuring Subnets on External Provider Logical Networks

A logical network provided by an external provider can only assign IP addresses to virtual machines if one or more subnets have been defined on that logical network. If no subnets are defined, virtual machines will not be assigned IP addresses. If there is one subnet, virtual machines will be assigned an IP address from that subnet, and if there are multiple subnets, virtual machines will be assigned an IP address from any of the available subnets. The DHCP service provided by the external network provider on which the logical network is hosted is responsible for assigning these IP addresses.

While the Red Hat Virtualization Manager automatically discovers predefined subnets on imported logical networks, you can also add or remove subnets to or from logical networks from within the Manager.

6.3.3.2. Adding Subnets to External Provider Logical Networks

Create a subnet on a logical network provided by an external provider.

Procedure 6.15. Adding Subnets to External Provider Logical Networks

1. Click the **Networks** tab.
2. Click the logical network provided by an external provider to which the subnet will be added.
3. Click the **Subnets** tab in the details pane.
4. Click the **New** button to open the **New External Subnet** window.

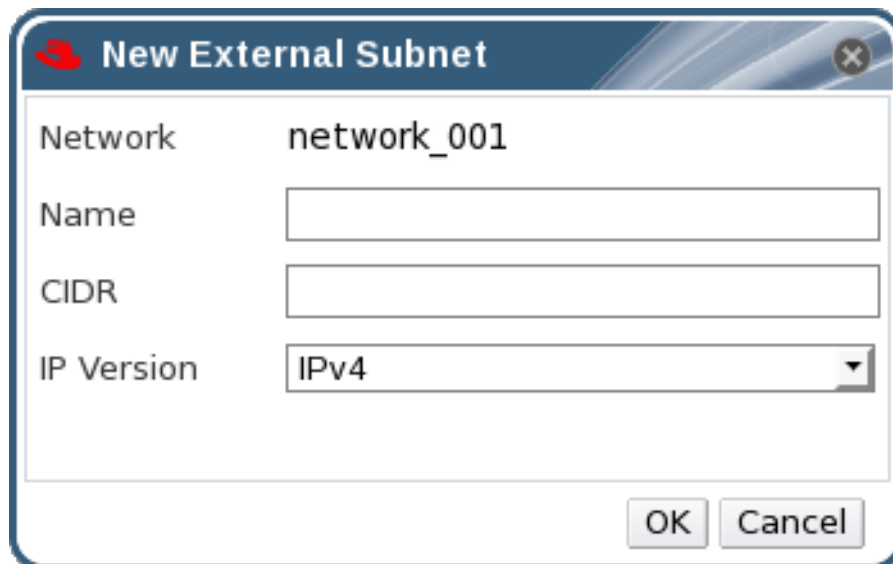


Figure 6.4. The New External Subnet Window

5. Enter a **Name** and **CIDR** for the new subnet.
6. From the **IP Version** drop-down menu, select either **IPv4** or **IPv6**.
7. Click **OK**.

6.3.3.3. Removing Subnets from External Provider Logical Networks

Remove a subnet from a logical network provided by an external provider.

Procedure 6.16. Removing Subnets from External Provider Logical Networks

1. Click the **Networks** tab.
2. Click the logical network provided by an external provider from which the subnet will be removed.
3. Click the **Subnets** tab in the details pane.

4. Click the subnet to remove.
5. Click the **Remove** button and click **OK** when prompted.

6.4. LOGICAL NETWORKS AND PERMISSIONS

6.4.1. Managing System Permissions for a Network

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A network administrator is a system administration role that can be applied for a specific network, or for all networks on a data center, cluster, host, virtual machine, or template. A network user can perform limited administration roles, such as viewing and attaching networks on a specific virtual machine or template. You can use the **Configure** button in the header bar to assign a network administrator for all networks in the environment.

The network administrator role permits the following actions:

- Create, edit and remove networks.
- Edit the configuration of the network, including configuring port mirroring.
- Attach and detach networks from resources including clusters and virtual machines.

The user who creates a network is automatically assigned **NetworkAdmin** permissions on the created network. You can also change the administrator of a network by removing the existing administrator and adding the new administrator.

6.4.2. Network Administrator and User Roles Explained

Network Permission Roles

The table below describes the administrator and user roles and privileges applicable to network administration.

Table 6.6. Red Hat Virtualization Network Administrator and User Roles

Role	Privileges	Notes
------	------------	-------

Role	Privileges	Notes
NetworkAdmin	Network Administrator for data center, cluster, host, virtual machine, or template. The user who creates a network is automatically assigned NetworkAdmin permissions on the created network.	Can configure and manage the network of a particular data center, cluster, host, virtual machine, or template. A network administrator of a data center or cluster inherits network permissions for virtual pools within the cluster. To configure port mirroring on a virtual machine network, apply the NetworkAdmin role on the network and the UserVmManager role on the virtual machine.
VnicProfileUser	Logical network and network interface user for virtual machine and template.	Can attach or detach network interfaces from specific logical networks.

6.4.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 6.17. Assigning a Role to a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Click **Add**.
4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down list.
6. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

6.4.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 6.18. Removing a Role from a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

6.5. HOSTS AND NETWORKING

6.5.1. Refreshing Host Capabilities

When a network interface card is added to a host, the capabilities of the host must be refreshed to display that network interface card in the Manager.

Procedure 6.19. To Refresh Host Capabilities

1. Use the resource tabs, tree mode, or the search function to find and select a host in the results list.
2. Click the **Refresh Capabilities** button.

The list of network interface cards in the **Network Interfaces** tab of the details pane for the selected host is updated. Any new network interface cards can now be used in the Manager.

6.5.2. Editing Host Network Interfaces and Assigning Logical Networks to Hosts

You can change the settings of physical host network interfaces, move the management network from one physical host network interface to another, and assign logical networks to physical host network interfaces. Bridge and ethtool custom properties are also supported.



WARNING

The only way to change the IP address of a host in Red Hat Virtualization is to remove the host and then to add it again.


To change the VLAN settings of a host, the host must be removed from the Manager, reconfigured, and re-added to the Manager.

To keep networking synchronized, do the following. Put the host in maintenance mode and manually remove the management network from the host. This will make the host reachable over the new VLAN. Add the host to the cluster. Virtual machines that are not connected directly to the management network can be migrated between hosts safely.

The following warning message appears when the VLAN ID of the management network is changed:

```
Changing certain properties (e.g. VLAN, MTU) of the
management network could lead to loss of connectivity to
hosts in the data center, if its underlying network
infrastructure isn't configured to accommodate the changes.
Are you sure you want to proceed?
```

Proceeding causes all of the hosts in the data center to lose connectivity to the Manager and causes the migration of hosts to the new management network to fail. The management network will be reported as "out-of-sync".



IMPORTANT

You cannot assign logical networks offered by external providers to physical host network interfaces; such networks are dynamically assigned to hosts as they are required by virtual machines.

Procedure 6.20. Editing Host Network Interfaces and Assigning Logical Networks to Hosts

1. Click the **Hosts** resource tab, and select the desired host.
2. Click the **Network Interfaces** tab in the details pane.
3. Click the **Setup Host Networks** button to open the **Setup Host Networks** window.
4. Attach a logical network to a physical host network interface by selecting and dragging the logical network into the **Assigned Logical Networks** area next to the physical host network interface.

Alternatively, right-click the logical network and select a network interface from the drop-down menu.

5. Configure the logical network:

- a. Hover your cursor over an assigned logical network and click the pencil icon to open the **Edit Management Network** window.
- b. From the **IPv4** tab, select a **Boot Protocol** from **None**, **DHCP**, or **Static**. If you selected **Static**, enter the **IP**, **Netmask / Routing Prefix**, and the **Gateway**.

**NOTE**

Each logical network can have a separate gateway defined from the management network gateway. This ensures traffic that arrives on the logical network will be forwarded using the logical network's gateway instead of the default gateway used by the management network.

**NOTE**

The **IPv6** tab should not be used as it is currently not supported.

- c. Use the **QoS** tab to override the default host network quality of service. Select **Override QoS** and enter the desired values in the following fields:
 - **Weighted Share**: Signifies how much of the logical link's capacity a specific network should be allocated, relative to the other networks attached to the same logical link. The exact share depends on the sum of shares of all networks on that link. By default this is a number in the range 1-100.
 - **Rate Limit [Mbps]**: The maximum bandwidth to be used by a network.
 - **Committed Rate [Mbps]**: The minimum bandwidth required by a network. The Committed Rate requested is not guaranteed and will vary depending on the network infrastructure and the Committed Rate requested by other networks on the same logical link.

For more information on configuring host network quality of service see [Section 3.3, "Host Network Quality of Service"](#)

- d. To configure a network bridge, click the **Custom Properties** tab and select **bridge_opts** from the drop-down list. Enter a valid key and value with the following syntax: *key=value*. Separate multiple entries with a whitespace character. The following keys are valid, with the values provided as examples. For more information on these parameters, see [Section B.1, "Explanation of bridge_opts Parameters"](#).

```
forward_delay=1500
gc_timer=3765
group_addr=1:80:c2:0:0:0
group_fwd_mask=0x0
hash_elasticity=4
hash_max=512
hello_time=200
hello_timer=70
max_age=2000
multicast_last_member_count=2
multicast_last_member_interval=100
multicast_membership_interval=26000
```

```

multicast_querier=0
multicast_querier_interval=25500
multicast_query_interval=13000
multicast_query_response_interval=1000
multicast_query_use_ifaddr=0
multicast_router=1
multicast_snooping=1
multicast_startup_query_count=2
multicast_startup_query_interval=3125

```

- e. To configure ethernet properties, click the **Custom Properties** tab and select **ethtool_opts** from the drop-down list. Enter a valid value using the format of the command-line arguments of ethtool. For example:

```

--coalesce em1 rx-usecs 14 sample-interval 3 --offload em2 rx on
lro on tso off --change em1 speed 1000 duplex half

```

This field can accept wildcards. For example, to apply the same option to all of this network's interfaces, use:

```

--coalesce * rx-usecs 14 sample-interval 3

```

The **ethtool_opts** option is not available by default; you need to add it using the engine configuration tool. See [Section B.2, “How to Set Up Red Hat Virtualization Manager to Use Ethtool”](#) for more information. For more information on ethtool properties, see the manual page by typing **man ethtool** in the command line.

- f. To configure Fibre Channel over Ethernet (FCoE), click the **Custom Properties** tab and select **fcoe** from the drop-down list. Enter a valid key and value with the following syntax: *key=value*. At least **enable=yes** is required. You can also add **dcb=[yes|no]** and **auto_vlan=[yes|no]**. Separate multiple entries with a whitespace character. The **fcoe** option is not available by default; you need to add it using the engine configuration tool. See [Section B.3, “How to Set Up Red Hat Virtualization Manager to Use FCoE”](#) for more information.



NOTE

A separate, dedicated logical network is recommended for use with FCoE.

- g. If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box. A logical network cannot be edited or moved to another interface until it is synchronized.

**NOTE**

Networks are not considered synchronized if they have one of the following conditions:

- The **VM Network** is different from the physical host network.
- The VLAN identifier is different from the physical host network.
- A **Custom MTU** is set on the logical network, and is different from the physical host network.

6. Select the **Verify connectivity between Host and Engine** check box to check network connectivity; this action will only work if the host is in maintenance mode.
7. Select the **Save network configuration** check box to make the changes persistent when the environment is rebooted.
8. Click **OK**.

**NOTE**

If not all network interface cards for the host are displayed, click the **Refresh Capabilities** button to update the list of network interface cards available for that host.

6.5.3. Adding Multiple VLANs to a Single Network Interface Using Logical Networks

Multiple VLANs can be added to a single network interface to separate traffic on the one host.

**IMPORTANT**

You must have created more than one logical network, all with the **Enable VLAN tagging** check box selected in the **New Logical Network** or **Edit Logical Network** windows.

Procedure 6.21. Adding Multiple VLANs to a Network Interface using Logical Networks

1. Click the **Hosts** resource tab, and select in the results list a host associated with the cluster to which your VLAN-tagged logical networks are assigned.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the data center.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Drag your VLAN-tagged logical networks into the **Assigned Logical Networks** area next to the physical network interface. The physical network interface can have multiple logical networks assigned due to the VLAN tagging.
5. Edit the logical networks by hovering your cursor over an assigned logical network and clicking the pencil icon to open the **Edit Network** window.

If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box.

Select a **Boot Protocol** from:

- **None,**
- **DHCP, or**
- **Static,**

Provide the **IP** and **Subnet Mask**.

Click **OK**.

6. Select the **Verify connectivity between Host and Engine** check box to run a network check; this will only work if the host is in maintenance mode.
7. Select the **Save network configuration** check box
8. Click **OK**.

Add the logical network to each host in the cluster by editing a NIC on each host in the cluster. After this is done, the network will become operational

You have added multiple VLAN-tagged logical networks to a single interface. This process can be repeated multiple times, selecting and editing the same network interface each time on each host to add logical networks with different VLAN tags to a single network interface.

6.5.4. Assigning Additional IPv4 Addresses to a Host Network

A host network, such as the **ovirtmgmt** management network, is created with only one IP address when initially set up. This means that if a NIC's configuration file (for example, **/etc/sysconfig/network-scripts/ifcfg-eth01**) is configured with multiple IP addresses, only the first listed IP address will be assigned to the host network. Additional IP addresses may be required if connecting to storage, or to a server on a separate private subnet using the same NIC.

The **vds-hook-extra-ipv4-addr**s hook allows you to configure additional IPv4 addresses for host networks. For more information about hooks, see [Appendix A, VDSM and Hooks](#).

In the following procedure, the host-specific tasks must be performed on each host for which you want to configure additional IP addresses.

Procedure 6.22. Assigning Additional IPv4 Addresses to a Host Network

1. On the host that you want to configure additional IPv4 addresses for, install the VDSM hook package. The package is available by default on Red Hat Virtualization Hosts but needs to be installed on Red Hat Enterprise Linux hosts.

```
# yum install vds-hook-extra-ipv4-addr
```

2. On the Manager, run the following command to add the key:

```
# engine-config -s
'UserDefinedNetworkCustomProperties=ipv4_addr=.*'
```

3. Restart the ovirt-engine service:

```
# systemctl restart ovirt-engine.service
```

4. In the Administration Portal, click the **Hosts** resource tab, and select the host for which additional IP addresses must be configured.
5. Click the **Network Interfaces** tab in the details pane and click the **Setup Host Networks** button to open the **Setup Host Networks** window.
6. Edit the host network interface by hovering the cursor over the assigned logical network and clicking the pencil icon to open the **Edit Management Network** window.
7. Select **ipv4_addr** from the **Custom Properties** drop-down list and add the additional IP address and prefix (for example 5.5.5.5/24). Multiple IP addresses must be comma-separated.
8. Click **OK**.
9. Select the **Save network configuration** check box.
10. Click **OK**.

The additional IP addresses will not be displayed in the Manager, but you can run the command **ip addr show** on the host to confirm that they have been added.

6.5.5. Adding Network Labels to Host Network Interfaces

Using network labels allows you to greatly simplify the administrative workload associated with assigning logical networks to host network interfaces.

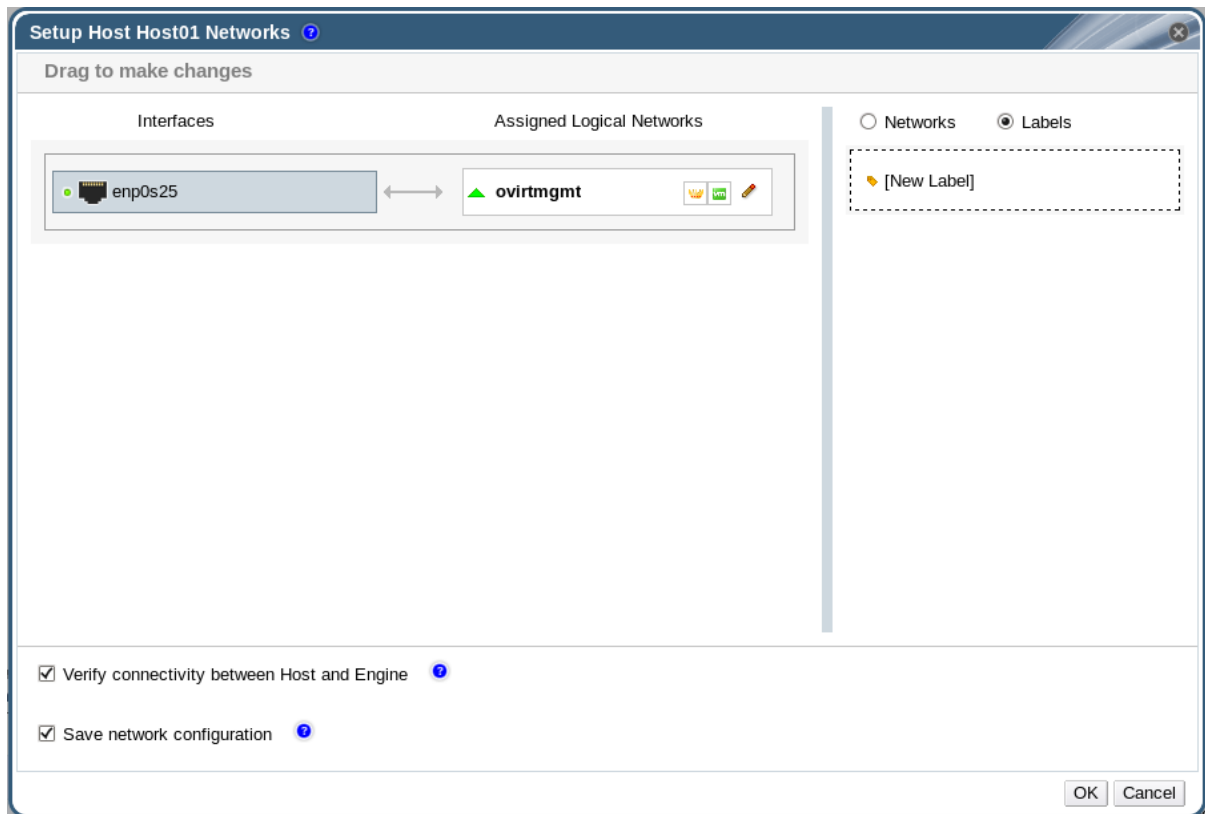


NOTE

Setting a label on a role network (for instance, a migration network or a display network) causes a mass deployment of that network on all hosts. Such mass additions of networks are achieved through the use of DHCP. This method of mass deployment was chosen over a method of typing in static addresses, because of the unscalable nature of the task of typing in many static IP addresses.

Procedure 6.23. Adding Network Labels to Host Network Interfaces

1. Click the **Hosts** resource tab, and select in the results list a host associated with the cluster to which your VLAN-tagged logical networks are assigned.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the data center.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Click **Labels**, and right-click [**New Label**]. Select a physical network interface to label.



5. Enter a name for the network label in the **Label** text field.
6. Click **OK**.

You have added a network label to a host network interface. Any newly created logical networks with the same label will be automatically assigned to all host network interfaces with that label. Also, removing a label from a logical network will automatically remove that logical network from all host network interfaces with that label.

6.5.6. Bonds

6.5.6.1. Bonding Logic in Red Hat Virtualization

The Red Hat Virtualization Manager Administration Portal allows you to create bond devices using a graphical interface. There are several distinct bond creation scenarios, each with its own logic.

Two factors that affect bonding logic are:

- Are either of the devices already carrying logical networks?
- Are the devices carrying compatible logical networks?

Table 6.7. Bonding Scenarios and Their Results


Bonding Scenario	Result
------------------	--------

Bonding Scenario	Result
NIC + NIC	<p>The Create New Bond window is displayed, and you can configure a new bond device.</p> <p>If the network interfaces carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.</p>
NIC + Bond	<p>The NIC is added to the bond device. Logical networks carried by the NIC and the bond are all added to the resultant bond device if they are compatible.</p> <p>If the bond devices carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.</p>
Bond + Bond	<p>If the bond devices are not attached to logical networks, or are attached to compatible logical networks, a new bond device is created. It contains all of the network interfaces, and carries all logical networks, of the component bond devices. The Create New Bond window is displayed, allowing you to configure your new bond.</p> <p>If the bond devices carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.</p>

6.5.6.2. Bonds

A *bond* is an aggregation of multiple network interface cards into a single software-defined device. Because bonded network interfaces combine the transmission capability of the network interface cards included in the bond to act as a single network interface, they can provide greater transmission speed than that of a single network interface card. Also, because all network interface cards in the bond must fail for the bond itself to fail, bonding provides increased fault tolerance. However, one limitation is that the network interface cards that form a bonded network interface must be of the same make and model to ensure that all network interface cards in the bond support the same options and modes.

The packet dispersal algorithm for a bond is determined by the bonding mode used.



IMPORTANT

Modes 1, 2, 3 and 4 support both virtual machine (bridged) and non-virtual machine (bridgeless) network types. Modes 0, 5 and 6 support non-virtual machine (bridgeless) networks only.

Bonding Modes

Red Hat Virtualization uses Mode 4 by default, but supports the following common bonding modes:

Mode 0 (round-robin policy)

Transmits packets through network interface cards in sequential order. Packets are transmitted in a loop that begins with the first available network interface card in the bond and end with the last available network interface card in the bond. All subsequent loops then start with the first available network interface card. Mode 0 offers fault tolerance and balances the load across all network interface cards in the bond. However, Mode 0 cannot be used in conjunction with bridges, and is therefore not compatible with virtual machine logical networks.

Mode 1 (active-backup policy)

Sets all network interface cards to a backup state while one network interface card remains active. In the event of failure in the active network interface card, one of the backup network interface cards replaces that network interface card as the only active network interface card in the bond. The MAC address of the bond in Mode 1 is visible on only one port to prevent any confusion that might otherwise be caused if the MAC address of the bond changed to reflect that of the active network interface card. Mode 1 provides fault tolerance and is supported in Red Hat Virtualization.

Mode 2 (XOR policy)

Selects the network interface card through which to transmit packets based on the result of an XOR operation on the source and destination MAC addresses modulo network interface card slave count. This calculation ensures that the same network interface card is selected for each destination MAC address used. Mode 2 provides fault tolerance and load balancing and is supported in Red Hat Virtualization.

Mode 3 (broadcast policy)

Transmits all packets to all network interface cards. Mode 3 provides fault tolerance and is supported in Red Hat Virtualization.

Mode 4 (IEEE 802.3ad policy)

Creates aggregation groups in which the interfaces share the same speed and duplex settings. Mode 4 uses all network interface cards in the active aggregation group in accordance with the IEEE 802.3ad specification and is supported in Red Hat Virtualization.

Mode 5 (adaptive transmit load balancing policy)

Ensures the distribution of outgoing traffic accounts for the load on each network interface card in the bond and that the current network interface card receives all incoming traffic. If the network interface card assigned to receive traffic fails, another network interface card is assigned to the role of receiving incoming traffic. Mode 5 cannot be used in conjunction with bridges, therefore it is not compatible with virtual machine logical networks.

Mode 6 (adaptive load balancing policy)

Combines Mode 5 (adaptive transmit load balancing policy) with receive load balancing for IPv4 traffic without any special switch requirements. ARP negotiation is used for balancing the receive load. Mode 6 cannot be used in conjunction with bridges, therefore it is not compatible with virtual machine logical networks.

6.5.6.3. Creating a Bond Device Using the Administration Portal

You can bond compatible network devices together. This type of configuration can increase available bandwidth and reliability. You can bond multiple network interfaces, pre-existing bond devices, and combinations of the two. A bond can carry both VLAN tagged and non-VLAN traffic.

Procedure 6.24. Creating a Bond Device using the Administration Portal

1. Click the **Hosts** resource tab, and select the host in the results list.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the host.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Select and drag one of the devices over the top of another device and drop it to open the **Create New Bond** window. Alternatively, right-click the device and select another device from the drop-down menu.

If the devices are incompatible, the bond operation fails and suggests how to correct the compatibility issue.

5. Select the **Bond Name** and **Bonding Mode** from the drop-down menus.

Bonding modes 1, 2, 4, and 5 can be selected. Any other mode can be configured using the **Custom** option.

6. Click **OK** to create the bond and close the **Create New Bond** window.
7. Assign a logical network to the newly created bond device.
8. Optionally choose to **Verify connectivity between Host and Engine** and **Save network configuration**.
9. Click **OK** accept the changes and close the **Setup Host Networks** window.

Your network devices are linked into a bond device and can be edited as a single interface. The bond device is listed in the **Network Interfaces** tab of the details pane for the selected host.

Bonding must be enabled for the ports of the switch used by the host. The process by which bonding is enabled is slightly different for each switch; consult the manual provided by your switch vendor for detailed information on how to enable bonding.

6.5.6.4. Example Uses of Custom Bonding Options with Host Interfaces

You can create customized bond devices by selecting **Custom** from the **Bonding Mode** of the **Create New Bond** window. The following examples should be adapted for your needs. For a comprehensive list of bonding options and their descriptions, see the [Linux Ethernet Bonding Driver HOWTO](#) on Kernel.org.

Example 6.1. xmit_hash_policy

This option defines the transmit load balancing policy for bonding modes 2 and 4. For example, if the majority of your traffic is between many different IP addresses, you may want to set a policy to balance by IP address. You can set this load-balancing policy by selecting a **Custom** bonding mode, and entering the following into the text field:

-

```
mode=4 xmit_hash_policy=layer2+3
```

Example 6.2. ARP Monitoring

ARP monitor is useful for systems which can't or don't report link-state properly via ethtool. Set an *arp_interval* on the bond device of the host by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=1 arp_interval=1 arp_ip_target=192.168.0.2
```

Example 6.3. Primary

You may want to designate a NIC with higher throughput as the primary interface in a bond device. Designate which NIC is primary by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=1 primary=eth0
```

6.5.7. Changing the FQDN of a Host

Use the following procedure to change the fully qualified domain name of hosts.

Procedure 6.25. Updating the FQDN of a Host

1. Place the host into maintenance mode so the virtual machines are live migrated to another host. See [Section 7.5.7, “Moving a Host to Maintenance Mode”](#) for more information. Alternatively, manually shut down or migrate all the virtual machines to another host. See [Manually Migrating Virtual Machines](#) in the *Virtual Machine Management Guide* for more information.
2. Click **Remove**, and click **OK** to remove the host from the Administration Portal.
3. Use the **hostnamectl** tool to update the host name. For more options, see [Configure Host Names](#) in the *Red Hat Enterprise Linux 7 Networking Guide*

```
# hostnamectl set-hostname NEW_FQDN
```

4. Reboot the host.
5. Re-register the host with the Manager. See [Section 7.5.1, “Adding a Host to the Red Hat Virtualization Manager”](#) for more information.

CHAPTER 7. HOSTS

7.1. INTRODUCTION TO 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 as individual Linux processes and threads on the host machine and are managed remotely by the Red Hat Virtualization Manager. A Red Hat Virtualization environment has one or more hosts attached to it.

Red Hat Virtualization supports two methods of installing hosts. You can use the Red Hat Virtualization Host (RHVH) installation media, or install hypervisor packages on a standard Red Hat Enterprise Linux installation.



NOTE

You can identify the host type of an individual host in the Red Hat Virtualization Manager by selecting the host, clicking **Software** under the **General** tab in the details pane, and checking the **OS Description**.

Hosts use **tuned** profiles, which provide virtualization optimizations. For more information on **tuned**, see the [Red Hat Enterprise Linux 7 Performance Tuning Guide](#)

The Red Hat Virtualization Host has security features enabled. Security Enhanced Linux (SELinux) and the iptables firewall are fully configured and on by default. The status of SELinux on a selected host is reported under **SELinux mode** in the **General** tab of the details pane. The Manager can open required ports on Red Hat Enterprise Linux hosts when it adds them to the environment.

A host is a physical 64-bit server with the Intel VT or AMD-V extensions running Red Hat Enterprise Linux 7 AMD64/Intel 64 version.

A physical host on the Red Hat Virtualization platform:

- Must belong to only one cluster in the system.
- Must have CPUs that support the AMD-V or Intel VT hardware virtualization extensions.
- Must have CPUs that support all functionality exposed by the virtual CPU type selected upon cluster creation.
- Has a minimum of 2 GB RAM.
- Can have an assigned system administrator with system permissions.

Administrators can receive the latest security advisories from the Red Hat Virtualization watch list. Subscribe to the Red Hat Virtualization watch list to receive new security advisories for Red Hat Virtualization products by email. Subscribe by completing this form:

<http://www.redhat.com/mailman/listinfo/rhev-watch-list/>

**WARNING**

Configuring networking through NetworkManager (including `nmcli`, `nmtui`, and the Cockpit user interface) is currently not supported. If additional network configuration is required before adding a host to the Manager, you must manually write `ifcfg` files. See the [Red Hat Enterprise Linux Networking Guide](#) for more information.

7.2. RED HAT VIRTUALIZATION HOST

Red Hat Virtualization Host (RHVH) is installed using a special build of Red Hat Enterprise Linux with only the packages required to host virtual machines. It uses an **Anaconda** installation interface based on the one used by Red Hat Enterprise Linux hosts, and can be updated through the Red Hat Virtualization Manager or via `yum`. However, installing additional packages is not currently supported; any additional packages that are installed must be reinstalled each time RHVH is updated.

RHVH features a Cockpit user interface for monitoring the host's resources and performing administrative tasks. Direct access to RHVH via SSH or console is not supported, so the Cockpit user interface provides a graphical user interface for tasks that are performed before the host is added to the Red Hat Virtualization Manager, such as deploying a self-hosted engine, and can also be used to run terminal commands via the **Tools > Terminal** sub-tab.

Access the Cockpit user interface at `https://HostFQDNorIP:9090` in your web browser. Cockpit for RHVH includes a custom **Virtualization** dashboard that displays the host's health status, SSH Host Key, self-hosted engine status, virtual machines, and virtual machine statistics.

**NOTE**

Custom boot kernel arguments can be added to Red Hat Virtualization Host using the `grubby` tool. The `grubby` tool makes persistent changes to the `grub.cfg` file. Navigate to the **Tools > Terminal** sub-tab in the host's Cockpit user interface to use `grubby` commands. See the [Red Hat Enterprise Linux System Administrator's Guide](#) for more information.

**WARNING**

Red Hat strongly recommends not creating untrusted users on RHVH, as this can lead to exploitation of local security vulnerabilities.

7.3. RED HAT ENTERPRISE LINUX HOSTS

You can use a Red Hat Enterprise Linux 7 installation on capable hardware as a host. Red

Hat Virtualization supports hosts running Red Hat Enterprise Linux 7 Server AMD64/Intel 64 version with Intel VT or AMD-V extensions. To use your Red Hat Enterprise Linux machine as a host, you must also attach the **Red Hat Enterprise Linux Server** entitlement and the **Red Hat Virtualization** entitlement.

Adding a host can take some time, as the following steps are completed by the platform: virtualization checks, installation of packages, creation of bridge, and a reboot of the host. Use the details pane to monitor the process as the host and management system establish a connection.



IMPORTANT

Third-party watchdogs should not be installed on Red Hat Enterprise Linux hosts, as they can interfere with the watchdog daemon provided by VDSM.

7.4. SATELLITE HOST PROVIDER HOSTS

Hosts provided by a Satellite host provider can also be used as virtualization hosts by the Red Hat Virtualization Manager. After a Satellite host provider has been added to the Manager as an external provider, any hosts that it provides can be added to and used in Red Hat Virtualization in the same way as Red Hat Virtualization Hosts (RHVH) and Red Hat Enterprise Linux hosts.

7.5. HOST TASKS

7.5.1. Adding a Host to the Red Hat Virtualization Manager

Adding a host to your Red Hat Virtualization environment can take some time, as the following steps are completed by the platform: virtualization checks, installation of packages, creation of bridge, and a reboot of the host. Use the details pane to monitor the process as the host and the Manager establish a connection.

Procedure 7.1. Adding a Host to the Red Hat Virtualization Manager

1. From the Administration Portal, click the **Hosts** resource tab.
2. Click **New**.
3. Use the drop-down list to select the **Data Center** and **Host Cluster** for the new host.
4. Enter the **Name** and the **Address** of the new host. The standard SSH port, port 22, is auto-filled in the **SSH Port** field.
5. Select an authentication method to use for the Manager to access the host.
 - Enter the root user's password to use password authentication.
 - Alternatively, copy the key displayed in the **SSH PublicKey** field to `/root/.ssh/authorized_keys` on the host to use public key authentication.
6. Click the **Advanced Parameters** button to expand the advanced host settings.
 - a. Optionally disable automatic firewall configuration.

- b. Optionally add a host SSH fingerprint to increase security. You can add it manually, or fetch it automatically.
7. Optionally configure **Power Management**, **SPM**, **Console**, **Network Provider**, and **Kernel**. See [Section 7.5.4, “Explanation of Settings and Controls in the New Host and Edit Host Windows”](#) for more information. **Hosted Engine** is used when deploying or undeploying a host for a self-hosted engine deployment.
8. Click **OK**.

The new host displays in the list of hosts with a status of **Installing**, and you can view the progress of the installation in the details pane. After a brief delay the host status changes to **Up**.

7.5.2. Adding a Satellite Host Provider Host

The process for adding a Satellite host provider host is almost identical to that of adding a Red Hat Enterprise Linux host except for the method by which the host is identified in the Manager. The following procedure outlines how to add a host provided by a Satellite host provider.

Procedure 7.2. Adding a Satellite Host Provider Host

1. Click the **Hosts** resource tab to list the hosts in the results list.
2. Click **New** to open the **New Host** window.
3. Use the drop-down menu to select the **Host Cluster** for the new host.
4. Select the **Foreman/Satellite** check box to display the options for adding a Satellite host provider host and select the provider from which the host is to be added.
5. Select either **Discovered Hosts** or **Provisioned Hosts**.
 - **Discovered Hosts** (default option): Select the host, host group, and compute resources from the drop-down lists.
 - **Provisioned Hosts**: Select a host from the **Providers Hosts** drop-down list.

Any details regarding the host that can be retrieved from the external provider are automatically set, and can be edited as desired.

6. Enter the **Name**, **Address**, and **SSH Port** (Provisioned Hosts only) of the new host.
7. Select an authentication method to use with the host.
 - Enter the root user's password to use password authentication.
 - Copy the key displayed in the **SSH PublicKey** field to `/root/.ssh/authorized_hosts` on the host to use public key authentication (Provisioned Hosts only).
8. You have now completed the mandatory steps to add a Red Hat Enterprise Linux host. Click the **Advanced Parameters** drop-down button to show the advanced host settings.

- a. Optionally disable automatic firewall configuration.
 - b. Optionally add a host SSH fingerprint to increase security. You can add it manually, or fetch it automatically.
9. You can configure the **Power Management**, **SPM**, **Console**, and **Network Provider** using the applicable tabs now; however, as these are not fundamental to adding a Red Hat Enterprise Linux host, they are not covered in this procedure.
 10. Click **OK** to add the host and close the window.

The new host displays in the list of hosts with a status of **Installing**, and you can view the progress of the installation in the details pane. After installation is complete, the status will update to **Reboot**. The host must be activated for the status to change to **Up**.

7.5.3. Configuring Satellite Errata Management for a Host

Red Hat Virtualization can be configured to view errata from Red Hat Satellite. This enables the host administrator to receive updates about available errata, and their importance, in the same dashboard used to manage host configuration. For more information about Red Hat Satellite see the [Red Hat Satellite User Guide](#)

Red Hat Virtualization 4.0 supports errata management with Red Hat Satellite 6.1.



IMPORTANT

Hosts are identified in the Satellite server by their FQDN. Hosts added using an IP address will not be able to report errata. This ensures that an external content host ID does not need to be maintained in Red Hat Virtualization.

The Satellite account used to manage the host must have Administrator permissions and a default organization set.

Procedure 7.3. Configuring Satellite Errata Management for a Host

1. Add the Satellite server as an external provider. See [Section 12.2.1, “Adding a Red Hat Satellite Instance for Host Provisioning”](#) for more information.
2. Associate the required host with the Satellite server.



NOTE

The host must be registered to the Satellite server and have the `katello-agent` package installed.

For more information on how to configure a host registration see [Configuring a Host for Registration](#) in the *Red Hat Satellite User Guide* and for more information on how to register a host and install the `katello-agent` package see [Registration](#) in the *Red Hat Satellite User Guide*

- a. In the **Hosts** tab, select the host in the results list.
- b. Click **Edit** to open the **Edit Host** window.

- c. Check the **Use Foreman/Satellite** checkbox.
- d. Select the required Satellite server from the drop-down list.
- e. Click **OK**.

The host is now configured to show the available errata, and their importance, in the same dashboard used to manage host configuration.

7.5.4. Explanation of Settings and Controls in the New Host and Edit Host Windows

7.5.4.1. Host General Settings Explained

These settings apply when editing the details of a host or adding new Red Hat Enterprise Linux hosts and Satellite host provider hosts.

The **General** settings table contains the information required on the **General** tab of the **New Host** or **Edit Host** window.

Table 7.1. General settings

Field Name	Description
Data Center	The data center to which the host belongs. Red Hat Virtualization Host (RHVH) cannot be added to Gluster-enabled clusters.
Host Cluster	The cluster to which the host belongs.

Field Name	Description
Use Foreman/Satellite	<p>Select or clear this check box to view or hide options for adding hosts provided by Satellite host providers. The following options are also available:</p> <p>Discovered Hosts</p> <ul style="list-style-type: none"> • Discovered Hosts - A drop-down list that is populated with the name of Satellite hosts discovered by the engine. • Host Groups -A drop-down list of host groups available. • Compute Resources - A drop-down list of hypervisors to provide compute resources. <p>Provisioned Hosts</p> <ul style="list-style-type: none"> • Providers Hosts - A drop-down list that is populated with the name of hosts provided by the selected external provider. The entries in this list are filtered in accordance with any search queries that have been input in the Provider search filter. • Provider search filter - A text field that allows you to search for hosts provided by the selected external provider. This option is provider-specific; see provider documentation for details on forming search queries for specific providers. Leave this field blank to view all available hosts.
Name	The name of the cluster. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Comment	A field for adding plain text, human-readable comments regarding the host.
Address	The IP address, or resolvable hostname of the host.
Password	The password of the host's root user. This can only be given when you add the host; it cannot be edited afterwards.

Field Name	Description
SSH PublicKey	Copy the contents in the text box to the <code>/root/.known_hosts</code> file on the host to use the Manager's ssh key instead of using a password to authenticate with the host.
Automatically configure host firewall	When adding a new host, the Manager can open the required ports on the host's firewall. This is enabled by default. This is an Advanced Parameter .
SSH Fingerprint	You can fetch the host's SSH fingerprint, and compare it with the fingerprint you expect the host to return, ensuring that they match. This is an Advanced Parameter .

7.5.4.2. Host Power Management Settings Explained

The **Power Management** settings table contains the information required on the **Power Management** tab of the **New Host** or **Edit Host** windows. You can configure power management if the host has a supported power management card.

Table 7.2. Power Management Settings

Field Name	Description
Enable Power Management	Enables power management on the host. Select this check box to enable the rest of the fields in the Power Management tab.
Kdump integration	Prevents the host from fencing while performing a kernel crash dump, so that the crash dump is not interrupted. From Red Hat Enterprise Linux 7.1 onwards, kdump is available by default. If kdump is available on the host, but its configuration is not valid (the kdump service cannot be started), enabling Kdump integration will cause the host (re)installation to fail. If this is the case, see Section 7.6.4, “fence_kdump Advanced Configuration” .
Disable policy control of power management	Power management is controlled by the Scheduling Policy of the host's cluster . If power management is enabled and the defined low utilization value is reached, the Manager will power down the host machine, and restart it again when load balancing requires or there are not enough free hosts in the cluster. Select this check box to disable policy control.

Field Name	Description
Agents by Sequential Order	<p>Lists the host's fence agents. Fence agents can be sequential, concurrent, or a mix of both.</p> <ul style="list-style-type: none"> • If fence agents are used sequentially, the primary agent is used first to stop or start a host, and if it fails, the secondary agent is used. • If fence agents are used concurrently, both fence agents have to respond to the Stop command for the host to be stopped; if one agent responds to the Start command, the host will go up. <p>Fence agents are sequential by default. Use the up and down buttons to change the sequence in which the fence agents are used.</p> <p>To make two fence agents concurrent, select one fence agent from the Concurrent with drop-down list next to the other fence agent. Additional fence agents can be added to the group of concurrent fence agents by selecting the group from the Concurrent with drop-down list next to the additional fence agent.</p>
Add Fence Agent	<p>Click the plus (+) button to add a new fence agent. The Edit fence agent window opens. See the table below for more information on the fields in this window.</p>
Power Management Proxy Preference	<p>By default, specifies that the Manager will search for a fencing proxy within the same cluster as the host, and if no fencing proxy is found, the Manager will search in the same dc (data center). Use the up and down buttons to change the sequence in which these resources are used. This field is available under Advanced Parameters.</p>

The following table contains the information required in the **Edit fence agent** window.

Table 7.3. Edit fence agent Settings

Field Name	Description
Address	<p>The address to access your host's power management device. Either a resolvable hostname or an IP address.</p>
User Name	<p>User account with which to access the power management device. You can set up a user on the device, or use the default user.</p>

Field Name	Description
Password	Password for the user accessing the power management device.
Type	<p>The type of power management device in your host.</p> <p>Choose one of the following:</p> <ul style="list-style-type: none"> • apc - APC MasterSwitch network power switch. Not for use with APC 5.x power switch devices. • apc_snmp - Use with APC 5.x power switch devices. • bladecenter - IBM Bladecenter Remote Supervisor Adapter. • cisco_ucs - Cisco Unified Computing System. • drac5 - Dell Remote Access Controller for Dell computers. • drac7 - Dell Remote Access Controller for Dell computers. • eps - ePowerSwitch 8M+ network power switch. • hpblade - HP BladeSystem. • ilo, ilo2, ilo3, ilo4 - HP Integrated Lights-Out. • ipmilan - Intelligent Platform Management Interface and Sun Integrated Lights Out Management devices. • rsa - IBM Remote Supervisor Adapter. • rsb - Fujitsu-Siemens RSB management interface. • wti - WTI Network Power Switch.
Port	The port number used by the power management device to communicate with the host.
Slot	The number used to identify the blade of the power management device.

Field Name	Description
Service Profile	The service profile name used to identify the blade of the power management device. This field appears instead of Slot when the device type is cisco_ucs .
Options	Power management device specific options. Enter these as 'key=value'. See the documentation of your host's power management device for the options available. For Red Hat Enterprise Linux 7 hosts, if you are using cisco_ucs as the power management device, you also need to append ssl_insecure=1 to the Options field.
Secure	Select this check box to allow the power management device to connect securely to the host. This can be done via ssh, ssl, or other authentication protocols depending on the power management agent.

7.5.4.3. SPM Priority Settings Explained

The **SPM** settings table details the information required on the **SPM** tab of the **New Host** or **Edit Host** window.

Table 7.4. SPM settings

Field Name	Description
SPM Priority	Defines the likelihood that the host will be given the role of Storage Pool Manager (SPM). The options are Low , Normal , and High priority. Low priority means that there is a reduced likelihood of the host being assigned the role of SPM, and High priority means there is an increased likelihood. The default setting is Normal.

7.5.4.4. Host Console Settings Explained

The **Console** settings table details the information required on the **Console** tab of the **New Host** or **Edit Host** window.

Table 7.5. Console settings

Field Name	Description
------------	-------------

Field Name	Description
Override display address	Select this check box to override the display addresses of the host. This feature is useful in a case where the hosts are defined by internal IP and are behind a NAT firewall. When a user connects to a virtual machine from outside of the internal network, instead of returning the private address of the host on which the virtual machine is running, the machine returns a public IP or FQDN (which is resolved in the external network to the public IP).
Display address	The display address specified here will be used for all virtual machines running on this host. The address must be in the format of a fully qualified domain name or IP.

7.5.4.5. Network Provider Settings Explained

The **Network Provider** settings table details the information required on the **Network Provider** tab of the **New Host** or **Edit Host** window.

Table 7.6. Network Provider settings

Field Name	Description
External Network Provider	If you have added an external network provider and want the host's network to be provisioned by the external network provider, select one from the list.

7.5.4.6. Kernel Settings Explained

The **Kernel** settings table details the information required on the **Kernel** tab of the **New Host** or **Edit Host** window. Common kernel boot parameter options are listed as check boxes so you can easily select them. For more complex changes, use the free text entry field next to **Kernel command line** to add in any additional parameters required.




IMPORTANT

If the host is attached to the Manager already, ensure you place the host into maintenance mode before applying any changes. You will need to reinstall the host by clicking **Reinstall**, and then reboot the host after the reinstallation is complete for the changes to take effect.

Table 7.7. Kernel Settings

Field Name	Description
Hostdev Passthrough & SR-IOV	Enables the IOMMU flag in the kernel to allow a host device to be used by a virtual machine as if the device is a device attached directly to the virtual machine itself. The host hardware and firmware must also support IOMMU. The virtualization extension and IOMMU extension must be enabled on the hardware. See Configuring a Host for PCI Passthrough in the <i>Installation Guide</i> . IBM POWER8 has IOMMU enabled by default.
Nested Virtualization	Enables the vmx or svm flag to allow you to run virtual machines within virtual machines. This option is only intended for evaluation purposes and not supported for production purposes. The vdsm-hook-nestedvt hook must be installed on the host.
Unsafe Interrupts	If IOMMU is enabled but the passthrough fails because the hardware does not support interrupt remapping, you can consider enabling this option. Note that you should only enable this option if the virtual machines on the host are trusted; having the option enabled potentially exposes the host to MSI attacks from the virtual machines. This option is only intended to be used as a workaround when using uncertified hardware for evaluation purposes.
PCI Reallocation	If your SR-IOV NIC is unable to allocate virtual functions because of memory issues, consider enabling this option. The host hardware and firmware must also support PCI reallocation. This option is only intended to be used as a workaround when using uncertified hardware for evaluation purposes.
Kernel command line	This field allows you to append more kernel parameters to the default parameters.

**NOTE**

If the kernel boot parameters are grayed out, click the **reset** button and the options will be available.

7.5.4.7. Hosted Engine Settings Explained

The **Hosted Engine** settings table details the information required on the **Hosted Engine** tab of the **New Host** or **Edit Host** window.

Table 7.8. Hosted Engine Settings

Field Name	Description
Choose hosted engine deployment action	<p>Three options are available:</p> <ul style="list-style-type: none"> • None - No actions required. • Deploy - Select this option to deploy the host as a self-hosted engine node. • Undeploy - For a self-hosted engine node, you can select this option to undeploy the host and remove self-hosted engine related configurations.

7.5.5. Configuring Host Power Management Settings

Configure your host power management device settings to perform host life-cycle operations (stop, start, restart) from the Administration Portal.

It is necessary to configure host power management in order to utilize host high availability and virtual machine high availability.




IMPORTANT

Ensure that your host is in **maintenance mode** before configuring power management settings. Otherwise, all running virtual machines on that host will be stopped ungracefully upon restarting the host, which can cause disruptions in production environments. A warning dialog will appear if you have not correctly set your host to **maintenance mode**.

Procedure 7.4. Configuring Power Management Settings

1. In the **Hosts** tab, select the host in the results list.
2. Click **Edit** to open the **Edit Host** window.
3. Click the **Power Management** tab to display the Power Management settings.
4. Select the **Enable Power Management** check box to enable the fields.
5. Select the **Kdump integration** check box to prevent the host from fencing while performing a kernel crash dump.



IMPORTANT

When you enable **Kdump integration** on an existing host, the host must be reinstalled for kdump to be configured. See [Section 7.5.11, “Reinstalling Hosts”](#).

6. Optionally, select the **Disable policy control of power management** check box if you do not want your host's power management to be controlled by the **Scheduling Policy** of the host's **cluster**.

7. Click the plus (+) button to add a new power management device. The **Edit fence agent** window opens.
8. Enter the **Address**, **User Name**, and **Password** of the power management device into the appropriate fields.
9. Select the power management device **Type** from the drop-down list.
10. Enter the **SSH Port** number used by the power management device to communicate with the host.
11. Enter the **Slot** number used to identify the blade of the power management device.
12. Enter the **Options** for the power management device. Use a comma-separated list of 'key=value' entries.
13. Select the **Secure** check box to enable the power management device to connect securely to the host.
14. Click **Test** to ensure the settings are correct. *Test Succeeded, Host Status is: on* will display upon successful verification.
15. Click **OK** to close the **Edit fence agent** window.
16. In the **Power Management** tab, optionally expand the **Advanced Parameters** and use the up and down buttons to specify the order in which the Manager will search the host's **cluster** and **dc** (datacenter) for a fencing proxy.
17. Click **OK**.

The **Power Management** drop-down menu is now enabled in the Administration Portal.

7.5.6. Configuring Host Storage Pool Manager Settings

The Storage Pool Manager (SPM) is a management role given to one of the hosts in a data center to maintain access control over the storage domains. The SPM must always be available, and the SPM role will be assigned to another host if the SPM host becomes unavailable. As the SPM role uses some of the host's available resources, it is important to prioritize hosts that can afford the resources.

The Storage Pool Manager (SPM) priority setting of a host alters the likelihood of the host being assigned the SPM role: a host with high SPM priority will be assigned the SPM role before a host with low SPM priority.

Procedure 7.5. Configuring SPM settings

1. Click the **Hosts** resource tab, and select a host from the results list.
2. Click **Edit** to open the **Edit Host** window.
3. Click the **SPM** tab to display the **SPM Priority** settings.
4. Use the radio buttons to select the appropriate SPM priority for the host.
5. Click **OK** to save the settings and close the window.

You have configured the SPM priority of the host.

7.5.7. Moving a Host to Maintenance Mode

Many common maintenance tasks, including network configuration and deployment of software updates, require that hosts be placed into maintenance mode. Hosts should be placed into maintenance mode before any event that might cause VDSM to stop working properly, such as a reboot, or issues with networking or storage.

When a host is placed into maintenance mode the Red Hat Virtualization Manager attempts to migrate all running virtual machines to alternative hosts. The standard prerequisites for live migration apply, in particular there must be at least one active host in the cluster with capacity to run the migrated virtual machines.

Procedure 7.6. Placing a Host into Maintenance Mode

1. Click the **Hosts** resource tab, and select the desired host.
2. Click **Maintenance** to open the **Maintenance Host(s)** confirmation window.
3. Optionally, enter a **Reason** for moving the host into maintenance mode in the **Maintenance Host(s)** confirmation window. This allows you to provide an explanation for the maintenance, which will appear in the logs and when the host is activated again.



NOTE

The host maintenance **Reason** field will only appear if it has been enabled in the cluster settings. See [Section 5.2.2.1, “General Cluster Settings Explained”](#) for more information.

4. Click **OK** to initiate maintenance mode.

All running virtual machines are migrated to alternative hosts. If the host is the Storage Pool Manager (SPM), the SPM role is migrated to another host. The **Status** field of the host changes to **Preparing for Maintenance**, and finally **Maintenance** when the operation completes successfully. VDSM does not stop while the host is in maintenance mode.



NOTE

If migration fails on any virtual machine, click **Activate** on the host to stop the operation placing it into maintenance mode, then click **Cancel Migration** on the virtual machine to stop the migration.

7.5.8. Activating a Host from Maintenance Mode

A host that has been placed into maintenance mode, or recently added to the environment, must be activated before it can be used. Activation may fail if the host is not ready; ensure that all tasks are complete before attempting to activate the host.

Procedure 7.7. Activating a Host from Maintenance Mode

1. Click the **Hosts** resources tab and select the host.
2. Click **Activate**.

The host status changes to **Unassigned**, and finally **Up** when the operation is complete. Virtual machines can now run on the host. Virtual machines that were migrated off the host when it was placed into maintenance mode are not automatically migrated back to the host when it is activated, but can be migrated manually. If the host was the Storage Pool Manager (SPM) before being placed into maintenance mode, the SPM role does not return automatically when the host is activated.

7.5.9. Removing a Host

Remove a host from your virtualized environment.

Procedure 7.8. Removing a host

1. In the Administration Portal, click the **Hosts** resource tab and select the host in the results list.
2. Place the host into maintenance mode.
3. Click **Remove** to open the **Remove Host(s)** confirmation window.
4. Select the **Force Remove** check box if the host is part of a Red Hat Gluster Storage cluster and has volume bricks on it, or if the host is non-responsive.
5. Click **OK**.

Your host has been removed from the environment and is no longer visible in the **Hosts** tab.

7.5.10. Updating a Host Between Minor Releases

See the following section in the Upgrade Guide for instructions on keeping your host current between minor releases: https://access.redhat.com/documentation/en/red-hat-virtualization/4.0/single/upgrade-guide/#chap-Updates_between_Minor_Releases.

7.5.11. Reinstalling Hosts

Reinstall Red Hat Virtualization Hosts (RHVH) and Red Hat Enterprise Linux hosts from the Administration Portal. The procedure includes stopping and restarting the host. If migration is enabled at cluster level, virtual machines are automatically migrated to another host in the cluster; as a result, it is recommended that host reinstalls are performed at a time when the host's usage is relatively low.

The cluster to which the host belongs must have sufficient memory reserve in order for its hosts to perform maintenance. Moving a host with live virtual machines to maintenance in a cluster that lacks sufficient memory causes the virtual machine migration operation to hang and then fail. You can reduce the memory usage of this operation by shutting down some or all virtual machines before moving the host to maintenance.




IMPORTANT

Ensure that the cluster contains more than one host before performing a reinstall. Do not attempt to reinstall all the hosts at the same time, as one host must remain available to perform Storage Pool Manager (SPM) tasks.

Procedure 7.9. Reinstalling Red Hat Virtualization Host or Red Hat Enterprise Linux hosts

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Maintenance**. If migration is enabled at cluster level, any virtual machines running on the host are migrated to other hosts. If the host is the SPM, this function is moved to another host. The status of the host changes as it enters maintenance mode.
3. Click **Reinstall** to open the **Install Host** window.
4. Click **OK** to reinstall the host.

Once successfully reinstalled, the host displays a status of **Up**. Any virtual machines that were migrated off the host, are at this point able to be migrated back to it.



IMPORTANT

After a Red Hat Virtualization Host is successfully registered to the Red Hat Virtualization Manager and then reinstalled, it may erroneously appear in the Administration Portal with the status of **Install Failed**. Click **Activate**, and the Host will change to an **Up** status and be ready for use.

7.5.12. Customizing Hosts with Tags

You can use tags to store information about your hosts. You can then search for hosts based on tags.

Procedure 7.10. Customizing hosts with tags

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Assign Tags** to open the **Assign Tags** window.

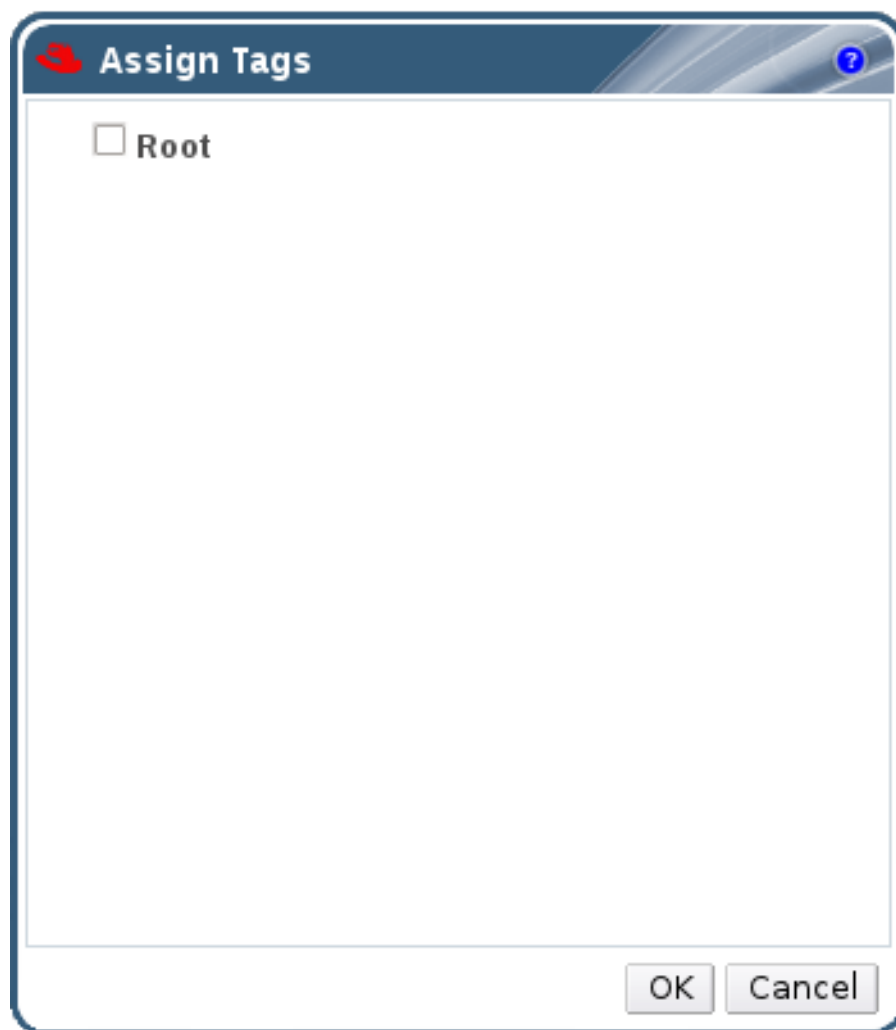


Figure 7.1. Assign Tags Window

3. The **Assign Tags** window lists all available tags. Select the check boxes of applicable tags.
4. Click **OK** to assign the tags and close the window.

You have added extra, searchable information about your host as tags.

7.5.13. Viewing Host Errata





Errata for each host can be viewed after the host has been configured to receive errata information from the Red Hat Satellite server. For more information on configuring a host to receive errata information see [Section 7.5.3, “Configuring Satellite Errata Management for a Host”](#)

Procedure 7.11. Viewing Host Errata

1. Click the **Hosts** resource tab, and select a host from the results list.
2. Click the **General** tab in the details pane.
3. Click the **Errata** sub-tab in the **General** tab.

7.5.14. Viewing the Health Status of a Host

Hosts have an external health status in addition to their regular **Status**. The external health status is reported by plug-ins or external systems, or set by an administrator, and appears to the left of the host's **Name** as one of the following icons:

- **OK:** No icon
- **Info:** 
- **Warning:** 
- **Error:** 
- **Failure:** 

To view further details about the host's health status, select the host and click the **Events** sub-tab.

The host's health status can also be viewed using the REST API. A **GET** request on a host will include the **external_status** element, which contains the health status.

You can set a host's health status in the REST API via the **events** collection. For more information, see [Adding Events](#) in the *REST API Guide*.

7.5.15. Viewing Host Devices

You can view the host devices for each host in the details pane. If the host has been configured for direct device assignment, these devices can be directly attached to virtual machines for improved performance.

For more information on the hardware requirements for direct device assignment, see [Additional Hardware Considerations for Using Device Assignment](#) in *Red Hat Virtualization Hardware Considerations for Implementing SR-IOV*.

For more information on configuring the host for direct device assignment, see [Configuring a Host for PCI Passthrough](#) in the *Installation Guide*.

For more information on attaching host devices to virtual machines, see [Host Devices](#) in the *Virtual Machine Management Guide*.

Procedure 7.12. Viewing Host Devices

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select a host from the results list.
2. Click the **Host Devices** tab in the details pane.

The details pane lists the details of the host devices, including whether the device is attached to a virtual machine, and currently in use by that virtual machine.


7.5.16. Preparing Host and Guest Systems for GPU Passthrough

The Graphics Processing Unit (GPU) device from a host can be directly assigned to a virtual machine. Before this can be achieved, both the host and the virtual machine require amendments to their **grub** configuration files. You can edit the host **grub** configuration file

using the **Kernel command line** free text entry field in the Administration Portal. Both the host machine and the virtual machine require reboot for the changes to take effect.

This procedure is relevant for hosts with either x86_64 or ppc64le architecture.

For more information on the hardware requirements for direct device assignment, see [PCI Device Requirements](#) in the *Installation Guide*.



IMPORTANT

If the host is attached to the Manager already, ensure you place the host into maintenance mode before applying any changes.

Procedure 7.13. Preparing a Host for GPU Passthrough

1. From the Administration Portal, select a host.
2. Click the **General** tab in the details pane, and click **Hardware**. Locate the GPU device *vendor ID:product ID*. In this example, the IDs are **10de:13ba** and **10de:0fbc**.
3. Right-click the host and select **Edit**. Click the **Kernel** tab.
4. In the **Kernel command line** free text entry field, enter the IDs located in the previous steps.

```
pci-stub.ids=10de:13ba,10de:0fbc
```

5. Blacklist the corresponding drivers on the host. For example, to blacklist nVidia's nouveau driver, next to *pci-stub.ids=xxxx:xxxx*, enter **rdblacklist=nouveau**.

```
pci-stub.ids=10de:13ba,10de:0fbc rdblacklist=nouveau
```

6. Click **OK** to save the changes.
7. Click **Reinstall** to commit the changes to the host.
8. Reboot the host after the reinstallation is complete.

**NOTE**

To confirm the device is bound to the **pci-stub** driver, run the **lspci** command:

```
# lspci -nnk
...
01:00.0 VGA compatible controller [0300]: NVIDIA Corporation
GM107GL [Quadro K2200] [10de:13ba] (rev a2)
    Subsystem: NVIDIA Corporation Device [10de:1097]
    Kernel driver in use: pci-stub
01:00.1 Audio device [0403]: NVIDIA Corporation Device
[10de:0fbc] (rev a1)
    Subsystem: NVIDIA Corporation Device [10de:1097]
    Kernel driver in use: pci-stub
...
```

For instructions on how to make the above changes by editing the **grub** configuration file manually, see [Preparing Host and Guest Systems for GPU Passthrough](#) in the *3.6 Administration Guide*.

Proceed to the next procedure to configure GPU passthrough on the guest system side.

Procedure 7.14. Preparing a Guest Virtual Machine for GPU Passthrough

- For Linux
 - a. Only proprietary GPU drivers are supported. Black list the corresponding open source driver in the **grub** configuration file. For example:

```
$ vi /etc/default/grub
...
GRUB_CMDLINE_LINUX="nofb splash=quiet console=tty0 ...
rdblacklist=nouveau"
...
```

- b. Locate the GPU BusID. In this example, is BusID is **00:09.0**.

```
# lspci | grep VGA
00:09.0 VGA compatible controller: NVIDIA Corporation GK106GL
[Quadro K4000] (rev a1)
```

- c. Edit the **/etc/X11/xorg.conf** file and append the following content:

```
Section "Device"
Identifier "Device0"
Driver "nvidia"
VendorName "NVIDIA Corporation"
BusID "PCI:0:9:0"
EndSection
```

- d. Restart the virtual machine.

- For Windows

- a. Download and install the corresponding drivers for the device. For example, for Nvidia drivers, go to [NVIDIA Driver Downloads](#).
- b. Restart the virtual machine.

The host GPU can now be directly assigned to the prepared virtual machine. For more information on assigning host devices to virtual machines, see [Host Devices](#) in the *Virtual Machine Management Guide*.

7.5.17. Accessing Cockpit from the Administration Portal

The Cockpit UI plug-in is an optional feature that can be installed in Red Hat Virtualization environments. The plug-in provides access to the Cockpit user interface, used for monitoring and administering host resources, through the Administration Portal. When a host with Cockpit installed is selected, the **Cockpit** sub-tab shows the Cockpit user interface directly in the details pane in the Administration Portal. Alternatively, the **Cockpit** button in the main **Hosts** menu opens the Cockpit user interface in a new browser tab.

The Cockpit user interface is available by default on Red Hat Virtualization Host (RHVH). It is currently not supported on Red Hat Enterprise Linux hosts.

Procedure 7.15. Accessing Cockpit from the Administration Portal

1. Install the Cockpit UI plug-in on the Manager machine:

```
# yum install cockpit-ovirt-uiplugin
```

2. Restart the **ovirt-engine** service:

```
# systemctl restart ovirt-engine.service
```

3. In the Administration Portal, click the **Hosts** tab and select a host.
4. Open the Cockpit user interface in a new tab, or view it directly through the Administration Portal:
 - Right-click the host and select **Cockpit** to open the Cockpit user interface in a new browser tab.
 - Click the **Cockpit** sub-tab to view the Cockpit user interface in the details pane of the **Hosts** tab.



NOTE

If Cockpit is not available on the selected host, the Cockpit sub-tab shows basic troubleshooting steps.

7.6. HOST RESILIENCE

7.6.1. Host High Availability

The Red Hat Virtualization Manager uses fencing to keep the hosts in a cluster responsive. A **Non Responsive** host is different from a **Non Operational** host. **Non Operational** hosts

can be communicated with by the Manager, but have an incorrect configuration, for example a missing logical network. **Non Responsive** hosts cannot be communicated with by the Manager.

If a host with a power management device loses communication with the Manager, it can be fenced (rebooted) from the Administration Portal. All the virtual machines running on that host are stopped, and highly available virtual machines are started on a different host.

All power management operations are done using a proxy host, as opposed to directly by the Red Hat Virtualization Manager. At least two hosts are required for power management operations.

Fencing allows a cluster to react to unexpected host failures as well as enforce power saving, load balancing, and virtual machine availability policies. You should configure the fencing parameters for your host's power management device and test their correctness from time to time.

Hosts can be fenced automatically using the power management parameters, or manually by right-clicking on a host and using the options on the menu. In a fencing operation, an unresponsive host is rebooted, and if the host does not return to an active status within a prescribed time, it remains unresponsive pending manual intervention and troubleshooting.

If the host is required to run virtual machines that are highly available, power management must be enabled and configured.

7.6.2. Power Management by Proxy in Red Hat Virtualization

The Red Hat Virtualization Manager does not communicate directly with fence agents. Instead, the Manager uses a proxy to send power management commands to a host power management device. The Manager uses VDSM to execute power management device actions, so another host in the environment is used as a fencing proxy.

You can select between:

- Any host in the same cluster as the host requiring fencing.
- Any host in the same data center as the host requiring fencing.

A viable fencing proxy host has a status of either *UP* or *Maintenance*.

7.6.3. Setting Fencing Parameters on a Host

The parameters for host fencing are set using the **Power Management** fields on the **New Host** or **Edit Host** windows. Power management enables the system to fence a troublesome host using an additional interface such as a Remote Access Card (RAC).

All power management operations are done using a proxy host, as opposed to directly by the Red Hat Virtualization Manager. At least two hosts are required for power management operations.

Procedure 7.16. Setting fencing parameters on a host

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Edit** to open the **Edit Host** window.

- Click the **Power Management** tab.

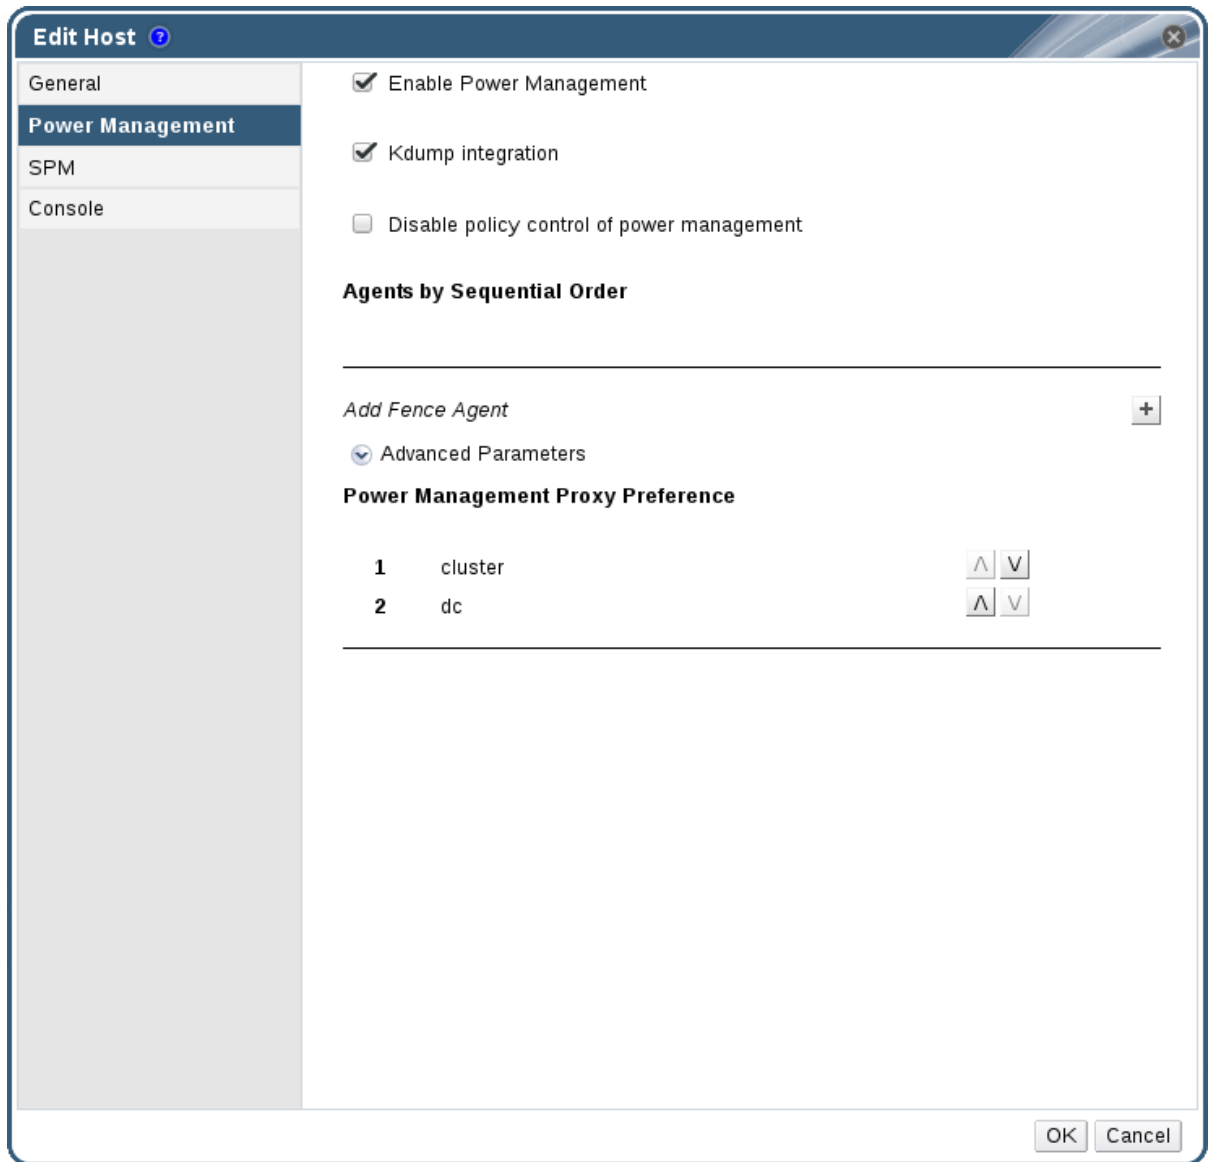


Figure 7.2. Power Management Settings

- Select the **Enable Power Management** check box to enable the fields.
- Select the **Kdump integration** check box to prevent the host from fencing while performing a kernel crash dump.



IMPORTANT

When you enable **Kdump integration** on an existing host, the host must be reinstalled for kdump to be configured. See [Section 7.5.11, “Reinstalling Hosts”](#).

- Optionally, select the **Disable policy control of power management** check box if you do not want your host's power management to be controlled by the **Scheduling Policy** of the host's cluster.
- Click the plus (+) button to add a new power management device. The **Edit fence agent** window opens.

Edit fence agent

Address

User Name

Password

Type

SSH Port

Slot

Options

Please use a comma-separated list of 'key=value'


Secure

Test

OK Cancel

Figure 7.3. Edit fence agent

8. Enter the **Address**, **User Name**, and **Password** of the power management device.
9. Select the power management device **Type** from the drop-down list.



NOTE

For more information on how to set up a custom power management device, see <https://access.redhat.com/articles/1238743>.

10. Enter the **SSH Port** number used by the power management device to communicate with the host.
11. Enter the **Slot** number used to identify the blade of the power management device.
12. Enter the **Options** for the power management device. Use a comma-separated list of 'key=value' entries.
13. Select the **Secure** check box to enable the power management device to connect securely to the host.

- Click the **Test** button to ensure the settings are correct. *Test Succeeded, Host Status is: on* will display upon successful verification.



WARNING

Power management parameters (userid, password, options, etc) are tested by Red Hat Virtualization Manager only during setup and manually after that. If you choose to ignore alerts about incorrect parameters, or if the parameters are changed on the power management hardware without the corresponding change in Red Hat Virtualization Manager, fencing is likely to fail when most needed.

- Click **OK** to close the **Edit fence agent** window.
- In the **Power Management** tab, optionally expand the **Advanced Parameters** and use the up and down buttons to specify the order in which the Manager will search the host's **cluster** and **dc** (datacenter) for a fencing proxy.
- Click **OK**.

You are returned to the list of hosts. Note that the exclamation mark next to the host's name has now disappeared, signifying that power management has been successfully configured.

7.6.4. fence_kdump Advanced Configuration

kdump

Select a host to view the status of the kdump service in the **General** tab of the details pane:

- Enabled:** kdump is configured properly and the kdump service is running.
- Disabled:** the kdump service is not running (in this case kdump integration will not work properly).
- Unknown:** happens only for hosts with an older VDSM version that does not report kdump status.

For more information on installing and using kdump, see the [Red Hat Enterprise Linux 7 Kernel Crash Dump Guide](#).

fence_kdump

Enabling **Kdump integration** in the **Power Management** tab of the **New Host** or **Edit Host** window configures a standard fence_kdump setup. If the environment's network configuration is simple and the Manager's FQDN is resolvable on all hosts, the default fence_kdump settings are sufficient for use.

However, there are some cases where advanced configuration of fence_kdump is necessary. Environments with more complex networking may require manual changes to

the configuration of the Manager, `fence_kdump` listener, or both. For example, if the Manager's FQDN is not resolvable on all hosts with **Kdump integration** enabled, you can set a proper host name or IP address using **engine-config**:

```
engine-config -s FenceKdumpDestinationAddress=A.B.C.D
```

The following example cases may also require configuration changes:

- The Manager has two NICs, where one of these is public-facing, and the second is the preferred destination for `fence_kdump` messages.
- You need to execute the `fence_kdump` listener on a different IP or port.
- You need to set a custom interval for `fence_kdump` notification messages, to prevent possible packet loss.


Customized `fence_kdump` detection settings are recommended for advanced users only, as changes to the default configuration are only necessary in more complex networking setups. For configuration options for the `fence_kdump` listener see [Section 7.6.4.1, “fence_kdump listener Configuration”](#). For configuration of kdump on the Manager see [Section 7.6.4.2, “Configuring fence_kdump on the Manager”](#).

7.6.4.1. fence_kdump listener Configuration

Edit the configuration of the `fence_kdump` listener. This is only necessary in cases where the default configuration is not sufficient.

Procedure 7.17. Manually Configuring the fence_kdump Listener

1. Create a new file (for example, `my-fence-kdump.conf`) in `/etc/ovirt-engine/ovirt-fence-kdump-listener.conf.d/`
2. Enter your customization with the syntax `OPTION=value` and save the file.



IMPORTANT

The edited values must also be changed in **engine-config** as outlined in the `fence_kdump` Listener Configuration Options table in [Section 7.6.4.2, “Configuring fence_kdump on the Manager”](#).

3. Restart the `fence_kdump` listener:

```
# systemctl restart ovirt-fence-kdump-listener.service
```

The following options can be customized if required:

Table 7.9. fence_kdump Listener Configuration Options

Variable	Description	Default	Note
----------	-------------	---------	------

Variable	Description	Default	Note
LISTENER_ADDRESS	Defines the IP address to receive fence_kdump messages on.	0.0.0.0	If the value of this parameter is changed, it must match the value of FenceKdumpDestinationAddress in engine-config .
LISTENER_PORT	Defines the port to receive fence_kdump messages on.	7410	If the value of this parameter is changed, it must match the value of FenceKdumpDestinationPort in engine-config .
HEARTBEAT_INTERVAL	Defines the interval in seconds of the listener's heartbeat updates.	30	If the value of this parameter is changed, it must be half the size or smaller than the value of FenceKdumpListenerTimeout in engine-config .
SESSION_SYNC_INTERVAL	Defines the interval in seconds to synchronize the listener's host kdumping sessions in memory to the database.	5	If the value of this parameter is changed, it must be half the size or smaller than the value of KdumpStartedTimeout in engine-config .
REOPEN_DB_CONNECTION_INTERVAL	Defines the interval in seconds to reopen the database connection which was previously unavailable.	30	-

Variable	Description	Default	Note
KDUMP_FINISHED_TIMEOUT	Defines the maximum timeout in seconds after the last received message from kdumping hosts after which the host kdump flow is marked as FINISHED.	60	If the value of this parameter is changed, it must be double the size or higher than the value of FenceKdumpMessageInterval in engine-config .

7.6.4.2. Configuring fence_kdump on the Manager


Edit the Manager's kdump configuration. This is only necessary in cases where the default configuration is not sufficient. The current configuration values can be found using:

```
# engine-config -g OPTION
```

Procedure 7.18. Manually Configuring Kdump with engine-config

1. Edit kdump's configuration using the **engine-config** command:

```
# engine-config -s OPTION=value
```



IMPORTANT

The edited values must also be changed in the fence_kdump listener configuration file as outlined in the **Kdump Configuration Options** table. See [Section 7.6.4.1, “fence_kdump listener Configuration”](#).

2. Restart the **ovirt-engine** service:

```
# systemctl restart ovirt-engine.service
```

3. Reinstall all hosts with **Kdump integration** enabled, if required (see the table below).

The following options can be configured using **engine-config**:

Table 7.10. Kdump Configuration Options

Variable	Description	Default	Note
----------	-------------	---------	------

Variable	Description	Default	Note
FenceKdumpDestinationAddress	Defines the hostname(s) or IP address(es) to send fence_kdump messages to. If empty, the Manager's FQDN is used.	Empty string (Manager FQDN is used)	If the value of this parameter is changed, it must match the value of LISTENER_ADDRESS in the fence_kdump listener configuration file, and all hosts with Kdump integration enabled must be reinstalled.
FenceKdumpDestinationPort	Defines the port to send fence_kdump messages to.	7410	If the value of this parameter is changed, it must match the value of LISTENER_PORT in the fence_kdump listener configuration file, and all hosts with Kdump integration enabled must be reinstalled.
FenceKdumpMessageInterval	Defines the interval in seconds between messages sent by fence_kdump.	5	If the value of this parameter is changed, it must be half the size or smaller than the value of KDUMP_FINISHED_TIMEOUT in the fence_kdump listener configuration file, and all hosts with Kdump integration enabled must be reinstalled.
FenceKdumpListenerTimeout	Defines the maximum timeout in seconds since the last heartbeat to consider the fence_kdump listener alive.	90	If the value of this parameter is changed, it must be double the size or higher than the value of HEARTBEAT_INTERVAL in the fence_kdump listener configuration file.

Variable	Description	Default	Note
KdumpStartedTimeout	Defines the maximum timeout in seconds to wait until the first message from the kdumping host is received (to detect that host kdump flow has started).	30	If the value of this parameter is changed, it must be double the size or higher than the value of SESSION_SYNC_INTERVAL in the fence_kdump listener configuration file, and FenceKdumpMessageInterval .


7.6.5. Soft-Fencing Hosts

Hosts can sometimes become non-responsive due to an unexpected problem, and though VDSM is unable to respond to requests, the virtual machines that depend upon VDSM remain alive and accessible. In these situations, restarting VDSM returns VDSM to a responsive state and resolves this issue.

"SSH Soft Fencing" is a process where the Manager attempts to restart VDSM via SSH on non-responsive hosts. If the Manager fails to restart VDSM via SSH, the responsibility for fencing falls to the external fencing agent if an external fencing agent has been configured.

Soft-fencing over SSH works as follows. Fencing must be configured and enabled on the host, and a valid proxy host (a second host, in an UP state, in the data center) must exist. When the connection between the Manager and the host times out, the following happens:

1. On the first network failure, the status of the host changes to "connecting".
2. The Manager then makes three attempts to ask VDSM for its status, or it waits for an interval determined by the load on the host. The formula for determining the length of the interval is configured by the configuration values `TimeoutToResetVdsInSeconds` (the default is 60 seconds) + `[DelayResetPerVmInSeconds (the default is 0.5 seconds)]*(the count of running virtual machines on host) + [DelayResetForSpmInSeconds (the default is 20 seconds)] * 1` (if host runs as SPM) or 0 (if the host does not run as SPM). To give VDSM the maximum amount of time to respond, the Manager chooses the longer of the two options mentioned above (three attempts to retrieve the status of VDSM or the interval determined by the above formula).
3. If the host does not respond when that interval has elapsed, **vdsmd restart** is executed via SSH.
4. If **vdsmd restart** does not succeed in re-establishing the connection between the host and the Manager, the status of the host changes to **Non Responsive** and, if power management is configured, fencing is handed off to the external fencing agent.

**NOTE**

Soft-fencing over SSH can be executed on hosts that have no power management configured. This is distinct from "fencing": fencing can be executed only on hosts that have power management configured.


7.6.6. Using Host Power Management Functions

Summary

When power management has been configured for a host, you can access a number of options from the Administration Portal interface. While each power management device has its own customizable options, they all support the basic options to start, stop, and restart a host.

Procedure 7.19. Using Host Power Management Functions

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Power Management** drop-down menu.
3. Select one of the following options:
 - **Restart**: This option stops the host and waits until the host's status changes to **Down**. When the agent has verified that the host is down, the highly available virtual machines are restarted on another host in the cluster. The agent then restarts this host. When the host is ready for use its status displays as **Up**.
 - **Start**: This option starts the host and lets it join a cluster. When it is ready for use its status displays as **Up**.
 - **Stop**: This option powers off the host. Before using this option, ensure that the virtual machines running on the host have been migrated to other hosts in the cluster. Otherwise the virtual machines will crash and only the highly available virtual machines will be restarted on another host. When the host has been stopped its status displays as **Non-Operational**.

**IMPORTANT**

When two fencing agents are defined on a host, they can be used concurrently or sequentially. For concurrent agents, both agents have to respond to the Stop command for the host to be stopped; and when one agent responds to the Start command, the host will go up. For sequential agents, to start or stop a host, the primary agent is used first; if it fails, the secondary agent is used.

4. Selecting one of the above options opens a confirmation window. Click **OK** to confirm and proceed.

Result

The selected action is performed.

7.6.7. Manually Fencing or Isolating a Non Responsive Host

Summary

If a host unpredictably goes into a non-responsive state, for example, due to a hardware failure; it can significantly affect the performance of the environment. If you do not have a power management device, or it is incorrectly configured, you can reboot the host manually.



WARNING

Do not use the **Confirm host has been rebooted** option unless you have manually rebooted the host. Using this option while the host is still running can lead to a virtual machine image corruption.

Procedure 7.20. Manually fencing or isolating a non-responsive host

1. On the **Hosts** tab, select the host. The status must display as **non-responsive**.
2. Manually reboot the host. This could mean physically entering the lab and rebooting the host.
3. On the Administration Portal, right-click the host entry and select the **Confirm Host has been rebooted** button.
4. A message displays prompting you to ensure that the host has been shut down or rebooted. Select the **Approve Operation** check box and click **OK**.

Result

You have manually rebooted your host, allowing highly available virtual machines to be started on active hosts. You confirmed your manual fencing action in the Administrator Portal, and the host is back online.

7.7. HOSTS AND PERMISSIONS

7.7.1. Managing System Permissions for a Host

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A host administrator is a system administration role for a specific host only. This is useful in clusters with multiple hosts, where each host requires a system administrator. You can use the **Configure** button in the header bar to assign a host administrator for all hosts in the environment.

The host administrator role permits the following actions:

- Edit the configuration of the host.

- Set up the logical networks.
- Remove the host.

You can also change the system administrator of a host by removing the existing system administrator and adding the new system administrator.

7.7.2. Host Administrator Roles Explained

Host Permission Roles

The table below describes the administrator roles and privileges applicable to host administration.

Table 7.11. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
HostAdmin	Host Administrator	Can configure, manage, and remove a specific host. Can also perform network-related operations on a specific host.

7.7.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 7.21. Assigning a Role to a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Click **Add**.
4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down list.
6. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

7.7.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 7.22. Removing a Role from a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

CHAPTER 8. STORAGE

Red Hat Virtualization uses a centralized storage system for virtual disk images, ISO files and snapshots. Storage networking can be implemented using:

- Network File System (NFS)
- GlusterFS exports
- CephFS
- Other POSIX compliant file systems
- Internet Small Computer System Interface (iSCSI)
- Local storage attached directly to the virtualization hosts
- Fibre Channel Protocol (FCP)
- Parallel NFS (pNFS)

Setting up storage is a prerequisite for a new data center because a data center cannot be initialized unless storage domains are attached and activated.

As a Red Hat Virtualization system administrator, you need to create, configure, attach and maintain storage for the virtualized enterprise. You should be familiar with the storage types and their use. Read your storage array vendor's guides, and see the [Red Hat Enterprise Linux Storage Administration Guide](#) for more information on the concepts, protocols, requirements or general usage of storage.

Red Hat Virtualization enables you to assign and manage storage using the Administration Portal's **Storage** tab. The **Storage** results list displays all the storage domains, and the details pane shows general information about the domain.

To add storage domains you must be able to successfully access the Administration Portal, and there must be at least one host connected with a status of **Up**.

Red Hat Virtualization has three types of storage domains:


- **Data Domain:** A data domain holds the virtual hard disks and OVF files of all the virtual machines and templates in a data center. In addition, snapshots of the virtual machines are also stored in the data domain.

The data domain cannot be shared across data centers. Data domains of multiple types (iSCSI, NFS, FC, POSIX, and Gluster) can be added to the same data center, provided they are all shared, rather than local, domains.

You must attach a data domain to a data center before you can attach domains of other types to it.


- **ISO Domain:** ISO domains store ISO files (or logical CDs) used to install and boot operating systems and applications for the virtual machines. An ISO domain removes the data center's need for physical media. An ISO domain can be shared across different data centers. ISO domains can only be NFS-based. Only one ISO domain can be added to a data center.
- **Export Domain:** Export domains are temporary storage repositories that are used

to copy and move images between data centers and Red Hat Virtualization environments. Export domains can be used to backup virtual machines. An export domain can be moved between data centers, however, it can only be active in one data center at a time. Export domains can only be NFS-based. Only one export domain can be added to a data center.



NOTE

The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disk images, and templates can then be uploaded from the imported storage domain to the attached data center. See [Section 8.6, “Importing Existing Storage Domains”](#) for information on importing storage domains.



IMPORTANT

Only commence configuring and attaching storage for your Red Hat Virtualization environment once you have determined the storage needs of your data center(s).

8.1. UNDERSTANDING STORAGE DOMAINS

A storage domain is a collection of images that have a common storage interface. A storage domain contains complete images of templates and virtual machines (including snapshots), or ISO files. A storage domain can be made of either block devices (SAN - iSCSI or FCP) or a file system (NAS - NFS, GlusterFS, CephFS, or other POSIX compliant file systems).

On NFS, all virtual disks, templates, and snapshots are files.

On SAN (iSCSI/FCP), each virtual disk, template or snapshot is a logical volume. Block devices are aggregated into a logical entity called a volume group, and then divided by LVM (Logical Volume Manager) into logical volumes for use as virtual hard disks. See *Red Hat Enterprise Linux Logical Volume Manager Administration Guide* for more information on LVM.

Virtual disks can have one of two formats, either QCOW2 or RAW. The type of storage can be either Sparse or Preallocated. Snapshots are always sparse but can be taken for disks created either as RAW or sparse.

Virtual machines that share the same storage domain can be migrated between hosts that belong to the same cluster.

8.2. PREPARING AND ADDING NFS STORAGE

8.2.1. Preparing NFS Storage

Set up NFS shares that will serve as a data domain on a Red Hat Enterprise Linux server. It is not necessary to create an ISO domain if one was created during the Red Hat Virtualization Manager installation procedure.

**NOTE**

The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disk images, and templates can then be uploaded from the imported storage domain to the attached data center. See [Section 8.6, “Importing Existing Storage Domains”](#) for information on importing storage domains.

For information on the setup and configuration of NFS on Red Hat Enterprise Linux, see [Network File System \(NFS\)](#) in the *Red Hat Enterprise Linux 6 Storage Administration Guide* or [Network File System \(NFS\)](#) in the *Red Hat Enterprise Linux 7 Storage Administration Guide*.

Specific system user accounts and system user groups are required by Red Hat Virtualization so the Manager can store data in the storage domains represented by the exported directories.

Procedure 8.1. Configuring the Required System User Accounts and System User Groups

1. Create the group **kvm**:

```
# groupadd kvm -g 36
```

2. Create the user **vds**m in the group **kvm**:

```
# useradd vds -u 36 -g 36
```

3. Set the ownership of your exported directories to 36:36, which gives vds:m:kvm ownership:

```
# chown -R 36:36 /exports/data
# chown -R 36:36 /exports/export
```

4. Change the mode of the directories so that read and write access is granted to the owner, and so that read and execute access is granted to the group and other users:

```
# chmod 0755 /exports/data
# chmod 0755 /exports/export
```

For more information on the required system users and groups see [Appendix G, System Accounts](#).

8.2.2. Attaching NFS Storage

Attach an NFS storage domain to the data center in your Red Hat Virtualization environment. This storage domain provides storage for virtualized guest images and ISO boot media. This procedure assumes that you have already exported shares. You must create the data domain before creating the export domain. Use the same procedure to create the export domain, selecting **Export / NFS** in the **Domain Function / Storage Type** list.

1. In the Red Hat Virtualization Manager Administration Portal, click the **Storage** resource tab.
2. Click **New Domain**.

Figure 8.1. The New Domain Window

3. Enter a **Name** for the storage domain.
4. Accept the default values for the **Data Center**, **Domain Function**, **Storage Type**, **Format**, and **Use Host** lists.
5. Enter the **Export Path** to be used for the storage domain.

The export path should be in the format of *192.168.0.10:/data* or *domain.example.com:/data*.

6. Optionally, you can configure the advanced parameters.
 - a. Click **Advanced Parameters**.
 - b. Enter a percentage value into the **Warning Low Space Indicator** field. If the free space available on the storage domain is below this percentage, warning messages are displayed to the user and logged.
 - c. Enter a GB value into the **Critical Space Action Blocker** field. If the free

space available on the storage domain is below this value, error messages are displayed to the user and logged, and any new action that consumes space, even temporarily, will be blocked.

- d. Select the **Wipe After Delete** check box to enable the wipe after delete option. This option can be edited after the domain is created, but doing so will not change the wipe after delete property of disks that already exist.

7. Click **OK**.

The new NFS data domain is displayed in the **Storage** tab with a status of **Locked** until the disk is prepared. The data domain is then automatically attached to the data center.

8.2.3. Increasing NFS Storage

To increase the amount of NFS storage, you can either create a new storage domain and add it to an existing data center, or increase the available free space on the NFS server. For the former option, see [Section 8.2.2, “Attaching NFS Storage”](#). The following procedure explains how to increase the available free space on the existing NFS server.

Procedure 8.2. Increasing an Existing NFS Storage Domain

1. Click the **Storage** resource tab and select an NFS storage domain.
2. In the details pane, click the **Data Center** tab and click the **Maintenance** button to place the storage domain into maintenance mode. This unmounts the existing share and makes it possible to resize the storage domain.
3. On the NFS server, resize the storage. For Red Hat Enterprise Linux 6 systems, see [Red Hat Enterprise Linux 6 Storage Administration Guide](#) For Red Hat Enterprise Linux 7 systems, see [Red Hat Enterprise Linux 7 Storage Administration Guide](#)
4. In the details pane, click the **Data Center** tab and click the **Activate** button to mount the storage domain.

8.3. PREPARING AND ADDING LOCAL STORAGE

8.3.1. Preparing Local Storage

A local storage domain can be set up on a host. When you set up a host to use local storage, the host automatically gets added to a new data center and cluster that no other hosts can be added to. Multiple host clusters require that all hosts have access to all storage domains, which is not possible with local storage. Virtual machines created in a single host cluster cannot be migrated, fenced or scheduled. For more information on the required system users and groups see [Appendix G, System Accounts](#).



NOTE

For information on preserving local storage domains when reinstalling Red Hat Virtualization Host (RHVH), see [Upgrading to RHVH While Preserving Local Storage](#) in the *Upgrade Guide*.



IMPORTANT

On Red Hat Virtualization Host (RHVH), the path used for local storage must be within the `/var` directory. For RHVH, prepend `/var` to the directories in the *Preparing Local Storage* procedure.

Local storage in the `/var` directory will be lost when Red Hat Virtualization Host is reinstalled. To avoid this, you can mount external storage to a host machine for use as a local storage domain. For more information on mounting storage, see the [Red Hat Enterprise Linux Storage Administration Guide](#)

Procedure 8.3. Preparing Local Storage

1. On the host, create the directory to be used for the local storage.

```
# mkdir -p /data/images
```

2. Ensure that the directory has permissions allowing read/write access to the `vdsm` user (UID 36) and `kvm` group (GID 36).

```
# chown 36:36 /data /data/images
```

```
# chmod 0755 /data /data/images
```

Your local storage is ready to be added to the Red Hat Virtualization environment.

8.3.2. Adding Local Storage

Storage local to your host has been prepared. Now use the Manager to add it to the host.

Adding local storage to a host in this manner causes the host to be put in a new data center and cluster. The local storage configuration window combines the creation of a data center, a cluster, and storage into a single process.

Procedure 8.4. Adding Local Storage

1. Click the **Hosts** resource tab, and select a host in the results list.
2. Click **Maintenance** to open the **Maintenance Host(s)** confirmation window.
3. Click **OK** to initiate maintenance mode.
4. Click **Configure Local Storage** to open the **Configure Local Storage** window.

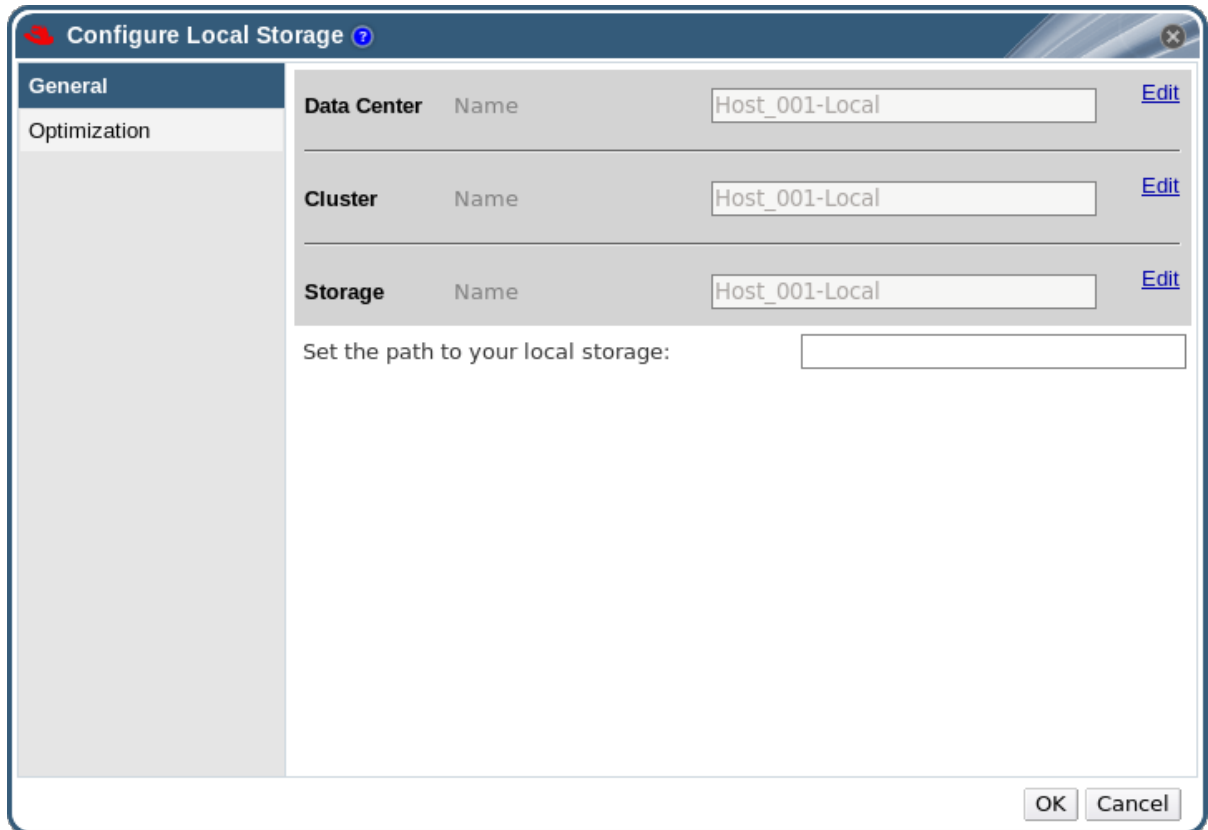


Figure 8.2. Configure Local Storage Window

5. Click the **Edit** buttons next to the **Data Center**, **Cluster**, and **Storage** fields to configure and name the local storage domain.
6. Set the path to your local storage in the text entry field.
7. If applicable, select the **Optimization** tab to configure the memory optimization policy for the new local storage cluster.
8. Click **OK** to save the settings and close the window.

Your host comes online in a data center of its own.

8.4. ADDING POSIX COMPLIANT FILE SYSTEM STORAGE

POSIX file system support allows you to mount file systems using the same mount options that you would normally use when mounting them manually from the command line. This functionality is intended to allow access to storage not exposed using NFS, iSCSI, or FCP.

Any POSIX compliant filesystem used as a storage domain in Red Hat Virtualization **MUST** support sparse files and direct I/O. The Common Internet File System (CIFS), for example, does not support direct I/O, making it incompatible with Red Hat Virtualization.



IMPORTANT

Do *not* mount NFS storage by creating a POSIX compliant file system Storage Domain. Always create an NFS Storage Domain instead.

8.4.1. Attaching POSIX Compliant File System Storage

You want to use a POSIX compliant file system that is not exposed using NFS, iSCSI, or FCP as a storage domain.

Procedure 8.5. Attaching POSIX Compliant File System Storage

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click **New Domain** to open the **New Domain** window.

Figure 8.3. POSIX Storage

3. Enter the **Name** for the storage domain.
4. Select the **Data Center** to be associated with the storage domain. The Data Center selected must be of type **POSIX (POSIX compliant FS)**. Alternatively, select **(none)**.
5. Select **Data / POSIX compliant FS** from the **Domain Function / Storage Type** drop-down menu.

If applicable, select the **Format** from the drop-down menu.

6. Select a host from the **Use Host** drop-down menu. Only hosts within the selected data center will be listed. The host that you select will be used to connect the storage domain.

7. Enter the **Path** to the POSIX file system, as you would normally provide it to the **mount** command.
8. Enter the **VFS Type**, as you would normally provide it to the **mount** command using the **-t** argument. See **man mount** for a list of valid VFS types.
9. Enter additional **Mount Options**, as you would normally provide them to the **mount** command using the **-o** argument. The mount options should be provided in a comma-separated list. See **man mount** for a list of valid mount options.
10. Optionally, you can configure the advanced parameters.
 - a. Click **Advanced Parameters**.
 - b. Enter a percentage value into the **Warning Low Space Indicator** field. If the free space available on the storage domain is below this percentage, warning messages are displayed to the user and logged.
 - c. Enter a GB value into the **Critical Space Action Blocker** field. If the free space available on the storage domain is below this value, error messages are displayed to the user and logged, and any new action that consumes space, even temporarily, will be blocked.
 - d. Select the **Wipe After Delete** check box to enable the wipe after delete option. This option can be edited after the domain is created, but doing so will not change the wipe after delete property of disks that already exist.
11. Click **OK** to attach the new Storage Domain and close the window.

8.5. ADDING BLOCK STORAGE

8.5.1. Adding iSCSI Storage

Red Hat Virtualization supports iSCSI storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.


For information on the setup and configuration of iSCSI on Red Hat Enterprise Linux, see [iSCSI Target Creation](#) in the *Red Hat Enterprise Linux 6 Storage Administration Guide* or [Online Storage Management](#) in the *Red Hat Enterprise Linux 7 Storage Administration Guide*.

Procedure 8.6. Adding iSCSI Storage

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click the **New Domain** button to open the **New Domain** window.
3. Enter the **Name** of the new storage domain.

Figure 8.4. New iSCSI Domain

4. Use the **Data Center** drop-down menu to select an data center.
5. Use the drop-down menus to select the **Domain Function** and the **Storage Type**. The storage domain types that are not compatible with the chosen domain function are not available.
6. Select an active host in the **Use Host** field. If this is not the first data domain in a data center, you must select the data center's SPM host.



IMPORTANT

All communication to the storage domain is through the selected host and not directly from the Red Hat Virtualization Manager. At least one active host must exist in the system and be attached to the chosen data center. All hosts must have access to the storage device before the storage domain can be configured.

7. The Red Hat Virtualization Manager is able to map either iSCSI targets to LUNs, or LUNs to iSCSI targets. The **New Domain** window automatically displays known targets with unused LUNs when iSCSI is selected as the storage type. If the target that you are adding storage from is not listed then you can use target discovery to find it, otherwise proceed to the next step.

iSCSI Target Discovery

1. Click **Discover Targets** to enable target discovery options. When targets have been discovered and logged in to, the **New Domain** window automatically displays targets with LUNs unused by the environment.




NOTE

LUNs used externally to the environment are also displayed.

You can use the **Discover Targets** options to add LUNs on many targets, or multiple paths to the same LUNs.

2. Enter the fully qualified domain name or IP address of the iSCSI host in the **Address** field.
3. Enter the port to connect to the host on when browsing for targets in the **Port** field. The default is **3260**.
4. If the Challenge Handshake Authentication Protocol (CHAP) is being used to secure the storage, select the **User Authentication** check box. Enter the **CHAP username** and **CHAP password**.




NOTE

It is now possible to use the REST API to define specific credentials to each iSCSI target per host. See [Defining Credentials to an iSCSI Target](#) in the *REST API Guide* for more information.

5. Click the **Discover** button.
6. Select the target to use from the discovery results and click the **Login** button.

Alternatively, click the **Login All** to log in to all of the discovered targets.



IMPORTANT

If more than one path access is required, ensure to discover and log in to the target through all the required paths. Modifying a storage domain to add additional paths is currently not supported.

8. Click the **+** button next to the desired target. This will expand the entry and display all unused LUNs attached to the target.
9. Select the check box for each LUN that you are using to create the storage domain.
10. Optionally, you can configure the advanced parameters.
 - a. Click **Advanced Parameters**.
 - b. Enter a percentage value into the **Warning Low Space Indicator** field. If the free space available on the storage domain is below this percentage, warning messages are displayed to the user and logged.

- c. Enter a GB value into the **Critical Space Action Blocker** field. If the free space available on the storage domain is below this value, error messages are displayed to the user and logged, and any new action that consumes space, even temporarily, will be blocked.
- d. Select the **Wipe After Delete** check box to enable the wipe after delete option. This option can be edited after the domain is created, but doing so will not change the wipe after delete property of disks that already exist.

11. Click **OK** to create the storage domain and close the window.

If you have configured multiple storage connection paths to the same target, follow the procedure in [Section 8.5.2, “Configuring iSCSI Multipathing”](#) to complete iSCSI bonding.

8.5.2. Configuring iSCSI Multipathing


The **iSCSI Multipathing** enables you to create and manage groups of logical networks and iSCSI storage connections. To prevent host downtime due to network path failure, configure multiple network paths between hosts and iSCSI storage. Once configured, the Manager connects each host in the data center to each bonded target via NICs/VLANs related to logical networks of the same iSCSI Bond. You can also specify which networks to use for storage traffic, instead of allowing hosts to route traffic through a default network. This option is only available in the Administration Portal after at least one iSCSI storage domain has been attached to a data center.

Prerequisites

- Ensure you have created an iSCSI storage domain and discovered and logged into all the paths to the iSCSI target(s).
- Ensure you have created **Non-Required** logical networks to bond with the iSCSI storage connections. You can configure multiple logical networks or bond networks to allow network failover.

Procedure 8.7. Configuring iSCSI Multipathing

1. Click the **Data Centers** tab and select a data center from the results list.
2. In the details pane, click the **iSCSI Multipathing** tab.
3. Click **Add**.
4. In the **Add iSCSI Bond** window, enter a **Name** and a **Description** for the bond.
5. Select the networks to be used for the bond from the **Logical Networks** list. The networks must be **Non-Required** networks.



NOTE

To change a network's **Required** designation, from the Administration Portal, select a network, click the **Cluster** tab, and click the **Manage Networks** button.

6. Select the storage domain to be accessed via the chosen networks from the **Storage Targets** list. Ensure to select all paths to the same target.

7. Click **OK**.

All hosts in the data center are connected to the selected iSCSI target through the selected logical networks.

8.5.3. Adding FCP Storage

Red Hat Virtualization platform supports SAN storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.

Red Hat Virtualization system administrators need a working knowledge of Storage Area Networks (SAN) concepts. SAN usually uses Fibre Channel Protocol (FCP) for traffic between hosts and shared external storage. For this reason, SAN may occasionally be referred to as FCP storage.

For information regarding the setup and configuration of FCP or multipathing on Red Hat Enterprise Linux, see the [Storage Administration Guide](#) and [DM Multipath Guide](#).

The following procedure shows you how to attach existing FCP storage to your Red Hat Virtualization environment as a data domain. For more information on other supported storage types, see [Chapter 8, Storage](#).

Procedure 8.8. Adding FCP Storage


1. Click the **Storage** resource tab to list all storage domains.
2. Click **New Domain** to open the **New Domain** window.
3. Enter the **Name** of the storage domain.

Figure 8.5. Adding FCP Storage

4. Use the **Data Center** drop-down menu to select an FCP data center.

If you do not yet have an appropriate FCP data center, select **(none)**.

5. Use the drop-down menus to select the **Domain Function** and the **Storage Type**. The storage domain types that are not compatible with the chosen data center are not available.
6. Select an active host in the **Use Host** field. If this is not the first data domain in a data center, you must select the data center's SPM host.



IMPORTANT

All communication to the storage domain is through the selected host and not directly from the Red Hat Virtualization Manager. At least one active host must exist in the system and be attached to the chosen data center. All hosts must have access to the storage device before the storage domain can be configured.

7. The **New Domain** window automatically displays known targets with unused LUNs when **Data / Fibre Channel** is selected as the storage type. Select the **LUN ID** check box to select all of the available LUNs.

8. Optionally, you can configure the advanced parameters.
 - a. Click **Advanced Parameters**.
 - b. Enter a percentage value into the **Warning Low Space Indicator** field. If the free space available on the storage domain is below this percentage, warning messages are displayed to the user and logged.
 - c. Enter a GB value into the **Critical Space Action Blocker** field. If the free space available on the storage domain is below this value, error messages are displayed to the user and logged, and any new action that consumes space, even temporarily, will be blocked.
 - d. Select the **Wipe After Delete** check box to enable the wipe after delete option. This option can be edited after the domain is created, but doing so will not change the wipe after delete property of disks that already exist.
9. Click **OK** to create the storage domain and close the window.

The new FCP data domain displays on the **Storage** tab. It will remain with a **Locked** status while it is being prepared for use. When ready, it is automatically attached to the data center.

8.5.4. Increasing iSCSI or FCP Storage

There are multiple ways to increase iSCSI or FCP storage size:

- Create a new storage domain with new LUNs and add it to an existing datacenter. See [Section 8.5.1, “Adding iSCSI Storage”](#).
- Create new LUNs and add them to an existing storage domain.
- Expand the storage domain by resizing the underlying LUNs.

For information about creating, configuring, or resizing iSCSI storage on Red Hat Enterprise Linux 6 systems, see [Red Hat Enterprise Linux 6 Storage Administration Guide](#) For Red Hat Enterprise Linux 7 systems, see [Red Hat Enterprise Linux 7 Storage Administration Guide](#)

The following procedure explains how to expand storage area network (SAN) storage by adding a new LUN to an existing storage domain.

Procedure 8.9. Increasing an Existing iSCSI or FCP Storage Domain

1. Create a new LUN on the SAN.
2. Click the **Storage** resource tab and select an iSCSI or FCP domain.
3. Click the **Manage Domain** button.
4. Click on **Targets > LUNs**, and click the **Discover Targets** expansion button.
5. Enter the connection information for the storage server and click the **Discover** button to initiate the connection.
6. Click on **LUNs > Targets** and select the check box of the newly available LUN.
7. Click **OK** to add the LUN to the selected storage domain.

This will increase the storage domain by the size of the added LUN.

When expanding the storage domain by resizing the underlying LUNs, the LUNs must also be refreshed in the Red Hat Virtualization Administration Portal.

Procedure 8.10. Refreshing the LUN Size

1. Click the **Storage** resource tab and select an iSCSI or FCP domain.
2. Click the **Manage Domain** button.
3. Click on **LUNs > Targets**.
4. In the **Additional Size** column, click the **Add *Additional_Storage_Size*** button of the LUN to refresh.
5. Click **OK** to refresh the LUN to indicate the new storage size.

8.5.5. Unusable LUNs in Red Hat Virtualization

In certain circumstances, the Red Hat Virtualization Manager will not allow you to use a LUN to create a storage domain or virtual machine hard disk.

- LUNs that are already part of the current Red Hat Virtualization environment are automatically prevented from being used.




Figure 8.6. Unusable LUNs in the Red Hat Virtualization Administration Portal

- LUNs that are already being used by the SPM host will also display as in use. You can choose to forcefully over ride the contents of these LUNs, but the operation is not guaranteed to succeed.

8.6. IMPORTING EXISTING STORAGE DOMAINS

8.6.1. Overview of Importing Existing Storage Domains


In addition to adding new storage domains that contain no data, you can also import existing storage domains and access the data they contain. The ability to import storage domains allows you to recover data in the event of a failure in the Manager database, and to migrate data from one data center or environment to another.

The following is an overview of importing each storage domain type:

Data

Importing an existing data storage domain allows you to access all of the virtual machines and templates that the data storage domain contains. After you import the storage domain, you must manually import virtual machines, floating disk images, and templates into the destination data center. The process for importing the virtual machines and templates that a data storage domain contains is similar to that for an export storage domain. However, because data storage domains contain all the virtual

machines and templates in a given data center, importing data storage domains is recommended for data recovery or large-scale migration of virtual machines between data centers or environments.



IMPORTANT

You can import existing data storage domains that were attached to data centers with a compatibility level of 3.5 or higher.

ISO

Importing an existing ISO storage domain allows you to access all of the ISO files and virtual diskettes that the ISO storage domain contains. No additional action is required after importing the storage domain to access these resources; you can attach them to virtual machines as required.

Export

Importing an existing export storage domain allows you to access all of the virtual machine images and templates that the export storage domain contains. Because export domains are designed for exporting and importing virtual machine images and templates, importing export storage domains is recommended method of migrating small numbers of virtual machines and templates inside an environment or between environments. For information on exporting and importing virtual machines and templates to and from export storage domains, see [Exporting and Importing Virtual Machines and Templates](#) in the *Virtual Machine Management Guide*.



NOTE

The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disk images, and templates can then be uploaded from the imported storage domain to the attached data center.

8.6.2. Importing Storage Domains

Import a storage domain that was previously attached to a data center in the same environment or in a different environment. This procedure assumes the storage domain is no longer attached to any data center in any environment, to avoid data corruption. To import and attach an existing data storage domain to a data center, the target data center must be initialized.

Procedure 8.11. Importing a Storage Domain

1. Click the **Storage** resource tab.
2. Click **Import Domain**.

Figure 8.7. The Import Pre-Configured Domain window

3. Select the data center to which to attach the storage domain from the **Data Center** drop-down list.
4. Enter a name for the storage domain.
5. Select the **Domain Function** and **Storage Type** from the appropriate drop-down lists.
6. Select a host from the **Use host** drop-down list.



IMPORTANT

All communication to the storage domain is through the selected host and not directly from the Red Hat Virtualization Manager. At least one active host must exist in the system and be attached to the chosen data center. All hosts must have access to the storage device before the storage domain can be configured.

7. Enter the details of the storage domain.

**NOTE**

The fields for specifying the details of the storage domain change in accordance with the value you select in the **Domain Function / Storage Type** list. These options are the same as those available for adding a new storage domain. For more information on these options, see [Section 8.1, “Understanding Storage Domains”](#).

8. Select the **Activate Domain in Data Center** check box to activate the storage domain after attaching it to the selected data center.
9. Click **OK**.

The storage domain is imported, and is displayed in the **Storage** tab. You can now import virtual machines and templates from the storage domain to the data center.

8.6.3. Migrating Storage Domains between Data Centers in the Same Environment

Migrate a storage domain from one data center to another in the same Red Hat Virtualization environment to allow the destination data center to access the data contained in the storage domain. This procedure involves detaching the storage domain from one data center, and attaching it to a different data center.

Procedure 8.12. Migrating a Storage Domain between Data Centers in the Same Environment

1. Shut down all virtual machines running on the required storage domain.
2. Click the **Storage** resource tab and select the storage domain from the results list.
3. Click the **Data Center** tab in the details pane.
4. Click **Maintenance**, then click **OK** to move the storage domain to maintenance mode.
5. Click **Detach**, then click **OK** to detach the storage domain from the source data center.
6. Click **Attach**.
7. Select the destination data center and click **OK**.

The storage domain is attached to the destination data center and is automatically activated. You can now import virtual machines and templates from the storage domain to the destination data center.

8.6.4. Migrating Storage Domains between Data Centers in Different Environments

Migrate a storage domain from one Red Hat Virtualization environment to another to allow the destination environment to access the data contained in the storage domain. This procedure involves removing the storage domain from one Red Hat Virtualization environment, and importing it into a different environment. To import and attach an existing data storage domain to a Red Hat Virtualization data center, the storage domain's source data center must have a compatibility level of 3.5 or higher.

Procedure 8.13. Migrating a Storage Domain between Data Centers in Different Environments

1. Log in to the Administration Portal of the source environment.
2. Shut down all virtual machines running on the required storage domain.
3. Click the **Storage** resource tab and select the storage domain from the results list.
4. Click the **Data Center** tab in the details pane.
5. Click **Maintenance**, then click **OK** to move the storage domain to maintenance mode.
6. Click **Detach**, then click **OK** to detach the storage domain from the source data center.
7. Click **Remove**.
8. In the **Remove Storage(s)** window, ensure the **Format Domain, i.e. Storage Content will be lost!** check box is not selected. This step preserves the data in the storage domain for later use.
9. Click **OK** to remove the storage domain from the source environment.
10. Log in to the Administration Portal of the destination environment.
11. Click the **Storage** resource tab.
12. Click **Import Domain**.

Figure 8.8. The Import Pre-Configured Domain window

13. Select the destination data center from the **Data Center** drop-down list.
14. Enter a name for the storage domain.
15. Select the **Domain Function** and **Storage Type** from the appropriate drop-down lists.
16. Select a host from the **Use Host** drop-down list.
17. Enter the details of the storage domain.



NOTE

The fields for specifying the details of the storage domain change in accordance with the value you select in the **Storage Type** drop-down list. These options are the same as those available for adding a new storage domain. For more information on these options, see [Section 8.1, “Understanding Storage Domains”](#).

18. Select the **Activate Domain in Data Center** check box to automatically activate the storage domain when it is attached.
19. Click **OK**.

The storage domain is attached to the destination data center in the new Red Hat Virtualization environment and is automatically activated. You can now import virtual machines and templates from the imported storage domain to the destination data center.

8.6.5. Importing Virtual Machines from Imported Data Storage Domains

Import a virtual machine from a data storage domain you have imported into your Red Hat Virtualization environment. This procedure assumes that the imported data storage domain has been attached to a data center and has been activated.

Procedure 8.14. Importing Virtual Machines from an Imported Data Storage Domain

1. Click the **Storage** resource tab.
2. Click the imported data storage domain.
3. Click the **VM Import** tab in the details pane.
4. Select one or more virtual machines to import.
5. Click **Import**.
6. Select the cluster into which the virtual machines are imported from the **Cluster** list.
7. Click **OK**.

You have imported one or more virtual machines into your environment. The imported virtual machines no longer appear in the list under the **VM Import** tab.

8.6.6. Importing Templates from Imported Data Storage Domains

Import a template from a data storage domain you have imported into your Red Hat Virtualization environment. This procedure assumes that the imported data storage domain has been attached to a data center and has been activated.

Procedure 8.15. Importing Templates from an Imported Data Storage Domain

1. Click the **Storage** resource tab.
2. Click the imported data storage domain.
3. Click the **Template Import** tab in the details pane.
4. Select one or more templates to import.
5. Click **Import**.
6. Select the cluster into which the templates are imported from the **Cluster** list.
7. Click **OK**.

You have imported one or more templates into your environment. The imported templates no longer appear in the list under the **Template Import** tab.

8.6.7. Importing a Disk Image from an Imported Storage Domain

Import floating virtual disks from an imported storage domain using the **Disk Import** tab of the details pane.



NOTE


Only QEMU-compatible disks can be imported into the Manager.

Procedure 8.16. Importing a Disk Image

1. Select a storage domain that has been imported into the data center.
2. In the details pane, click **Disk Import**.
3. Select one or more disk images and click **Import** to open the **Import Disk(s)** window.
4. Select the appropriate **Disk Profile** for each disk.
5. Click **OK** to import the selected disks.

8.6.8. Importing an Unregistered Disk Image from an Imported Storage Domain

Import floating virtual disks from a storage domain using the **Disk Import** tab of the details pane. Floating disks created outside of a Red Hat Virtualization environment are not registered with the Manager. Scan the storage domain to identify unregistered floating disks to be imported.



NOTE

Only QEMU-compatible disks can be imported into the Manager.

Procedure 8.17. Importing a Disk Image

1. Select a storage domain that has been imported into the data center.
2. Right-click the storage domain and select **Scan Disks** so that the Manager can identify unregistered disks.
3. In the details pane, click **Disk Import**.
4. Select one or more disk images and click **Import** to open the **Import Disk(s)** window.
5. Select the appropriate **Disk Profile** for each disk.
6. Click **OK** to import the selected disks.

8.7. STORAGE TASKS

8.7.1. Populating the ISO Storage Domain

An ISO storage domain is attached to a data center. ISO images must be uploaded to it. Red Hat Virtualization provides an ISO uploader tool that ensures that the images are uploaded into the correct directory path, with the correct user permissions.

The creation of ISO images from physical media is not described in this document. It is assumed that you have access to the images required for your environment.

Procedure 8.18. Populating the ISO Storage Domain

1. Copy the required ISO image to a temporary directory on the system running Red Hat Virtualization Manager.
2. Log in to the system running Red Hat Virtualization Manager as the **root** user.
3. Use the **engine-iso-uploader** command to upload the ISO image. This action will take some time. The amount of time varies depending on the size of the image being uploaded and available network bandwidth.

Example 8.1. ISO Uploader Usage

In this example the ISO image **RHEL6.iso** is uploaded to the ISO domain called **ISODomain** using NFS. The command will prompt for an administrative user name and password. The user name must be provided in the form *user name@domain*.

```
# engine-iso-uploader --iso-domain=ISODomain upload RHEL6.iso
```

The ISO image is uploaded and appears in the ISO storage domain specified. It is also available in the list of available boot media when creating virtual machines in the data center to which the storage domain is attached.

8.7.2. Moving Storage Domains to Maintenance Mode

Detaching and removing storage domains requires that they be in maintenance mode. This is required to redesignate another data domain as the master data domain.

Expanding iSCSI domains by adding more LUNs can only be done when the domain is active.

Procedure 8.19. Moving storage domains to maintenance mode

1. Shut down all the virtual machines running on the storage domain.
2. Click the **Storage** resource tab and select a storage domain.
3. Click the **Data Centers** tab in the details pane.
4. Click **Maintenance** to open the **Storage Domain maintenance** confirmation window.
5. Click **OK** to initiate maintenance mode. The storage domain is deactivated and has an **Inactive** status in the results list.

You can now edit, detach, remove, or reactivate the inactive storage domains from the data center.

**NOTE**

You can also activate, detach and place domains into maintenance mode using the Storage tab on the details pane of the data center it is associated with.

8.7.3. Editing Storage Domains

You can edit storage domain parameters through the Administration Portal. Depending on the state of the storage domain, either active or inactive, different fields are available for editing. Fields such as **Data Center**, **Domain Function**, **Storage Type**, and **Format** cannot be changed.

- **Active:** When the storage domain is in an active state, the **Name**, **Description**, **Comment**, **Warning Low Space Indicator (%)**, **Critical Space Action Blocker (GB)**, and **Wipe After Delete** fields can be edited. The **Name** field can only be edited while the storage domain is active. All other fields can also be edited while the storage domain is inactive.
- **Inactive:** When the storage domain is in maintenance mode or unattached, thus in an inactive state, you can edit all fields except **Name**, **Data Center**, **Domain Function**, **Storage Type**, and **Format**. The storage domain must be inactive to edit storage connections, mount options, and other advanced parameters. This is only supported for NFS, POSIX, and Local storage types.

**NOTE**

iSCSI storage connections cannot be edited via the Administration Portal, but can be edited via the REST API. See [Updating an iSCSI Storage Connection](#) in the *REST API Guide*.

Procedure 8.20. Editing an Active Storage Domain

1. Click the **Storage** resource tab and select a storage domain.
2. Click **Manage Domain**.
3. Edit the available fields as required.
4. Click **OK**.


Procedure 8.21. Editing an Inactive Storage Domain

1. Click the **Storage** resource tab and select a storage domain.
2. If the storage domain is active, click the **Data Center** tab in the details pane and click **Maintenance**.
3. Click **Manage Domain**.
4. Edit the storage path and other details as required. The new connection details must be of the same storage type as the original connection.
5. Click **OK**.
6. Click the **Data Center** tab in the details pane and click **Activate**.

8.7.4. Activating Storage Domains from Maintenance Mode

If you have been making changes to a data center's storage, you have to put storage domains into maintenance mode. Activate a storage domain to resume using it.

1. Click the **Storage** resource tab and select an inactive storage domain in the results list.
2. Click the **Data Centers** tab in the details pane.
3. Select the appropriate storage domain and click **Activate**.



IMPORTANT

If you attempt to activate the ISO domain before activating the data domain, an error message displays and the domain is not activated.

8.7.5. Removing a Storage Domain

You have a storage domain in your data center that you want to remove from the virtualized environment.

Procedure 8.22. Removing a Storage Domain

1. Click the **Storage** resource tab and select the appropriate storage domain in the results list.
2. Move the domain into maintenance mode to deactivate it.
3. Detach the domain from the data center.
4. Click **Remove** to open the **Remove Storage** confirmation window.
5. Select a host from the list.
6. Click **OK** to remove the storage domain and close the window.

The storage domain is permanently removed from the environment.

8.7.6. Destroying a Storage Domain

A storage domain encountering errors may not be able to be removed through the normal procedure. Destroying a storage domain will forcibly remove the storage domain from the virtualized environment without reference to the export directory.

When the storage domain is destroyed, you are required to manually fix the export directory of the storage domain before it can be used again.

Procedure 8.23. Destroying a Storage Domain

1. Use the **Storage** resource tab, tree mode, or the search function to find and select the appropriate storage domain in the results list.
2. Right-click the storage domain and select **Destroy** to open the **Destroy Storage Domain** confirmation window.

3. Select the **Approve operation** check box and click **OK** to destroy the storage domain and close the window.

The storage domain has been destroyed. Manually clean the export directory for the storage domain to recycle it.

8.7.7. Detaching a Storage Domain from a Data Center

Detach a storage domain from the data center to migrate virtual machines and templates to another data center.

Procedure 8.24. Detaching a Storage Domain from the Data Center

1. Click the **Storage** resource tab, and select a storage domain from the results list.
2. Click the **Data Centers** tab in the details pane and select the storage domain.
3. Click **Maintenance** to open the **Maintenance Storage Domain(s)** confirmation window.
4. Click **OK** to initiate maintenance mode.
5. Click **Detach** to open the **Detach Storage** confirmation window.
6. Click **OK** to detach the storage domain.

The storage domain has been detached from the data center, ready to be attached to another data center.

8.7.8. Attaching a Storage Domain to a Data Center

Attach a storage domain to a data center.

Procedure 8.25. Attaching a Storage Domain to a Data Center

1. Click the **Storage** resource tab, and select a storage domain from the results list.
2. Click the **Data Centers** tab in the details pane.
3. Click **Attach** to open the **Attach to Data Center** window.
4. Select the radio button of the appropriate data center.
5. Click **OK** to attach the storage domain.

The storage domain is attached to the data center and is automatically activated.

8.7.9. Disk Profiles

Disk profiles define the maximum level of throughput and the maximum level of input and output operations for a virtual disk in a storage domain. Disk profiles are created based on storage profiles defined under data centers, and must be manually assigned to individual virtual disks for the profile to take effect.

8.7.9.1. Creating a Disk Profile

Create a disk profile. This procedure assumes you have already defined one or more storage quality of service entries under the data center to which the storage domain belongs.

Procedure 8.26. Creating a Disk Profile

1. Click the **Storage** resource tab and select a data storage domain.
2. Click the **Disk Profiles** sub tab in the details pane.
3. Click **New**.
4. Enter a name for the disk profile in the **Name** field.
5. Enter a description for the disk profile in the **Description** field.
6. Select the quality of service to apply to the disk profile from the **QoS** list.
7. Click **OK**.

You have created a disk profile, and that disk profile can be applied to new virtual disks hosted in the data storage domain.

8.7.9.2. Removing a Disk Profile

Remove an existing disk profile from your Red Hat Virtualization environment.



Procedure 8.27. Removing a Disk Profile



1. Click the **Storage** resource tab and select a data storage domain.
2. Click the **Disk Profiles** sub tab in the details pane.
3. Select the disk profile to remove.
4. Click **Remove**.
5. Click **OK**.

You have removed a disk profile, and that disk profile is no longer available. If the disk profile was assigned to any virtual disks, the disk profile is removed from those virtual disks.

8.7.10. Viewing the Health Status of a Storage Domain

Storage domains have an external health status in addition to their regular **Status**. The external health status is reported by plug-ins or external systems, or set by an administrator, and appears to the left of the storage domain's **Name** as one of the following icons:

- **OK:** No icon
- **Info:** 
- **Warning:** 

- **Error:** 
- **Failure:** 

To view further details about the storage domain's health status, select the storage domain and click the **Events** sub-tab.

The storage domain's health status can also be viewed using the REST API. A **GET** request on a storage domain will include the **external_status** element, which contains the health status.

You can set a storage domain's health status in the REST API via the **events** collection. For more information, see [Adding Events](#) in the *REST API Guide*.

8.8. STORAGE AND PERMISSIONS


8.8.1. Managing System Permissions for a Storage Domain

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A storage administrator is a system administration role for a specific storage domain only. This is useful in data centers with multiple storage domains, where each storage domain requires a system administrator. Use the **Configure** button in the header bar to assign a storage administrator for all storage domains in the environment.

The storage domain administrator role permits the following actions:

- Edit the configuration of the storage domain.
- Move the storage domain into maintenance mode.
- Remove the storage domain.



NOTE

You can only assign roles and permissions to existing users.

You can also change the system administrator of a storage domain by removing the existing system administrator and adding the new system administrator.

8.8.2. Storage Administrator Roles Explained

Storage Domain Permission Roles

The table below describes the administrator roles and privileges applicable to storage domain administration.

Table 8.1. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
StorageAdmin	Storage Administrator	Can create, delete, configure and manage a specific storage domain.
GlusterAdmin	Gluster Storage Administrator	Can create, delete, configure and manage Gluster storage volumes.

8.8.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 8.28. Assigning a Role to a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Click **Add**.
4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down list.
6. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

8.8.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 8.29. Removing a Role from a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

CHAPTER 9. WORKING WITH RED HAT GLUSTER STORAGE

9.1. RED HAT GLUSTER STORAGE NODES

9.1.1. Adding Red Hat Gluster Storage Nodes

Add Red Hat Gluster Storage nodes to Gluster-enabled clusters and incorporate GlusterFS volumes and bricks into your Red Hat Virtualization environment.

This procedure presumes that you have a Gluster-enabled cluster of the appropriate **Compatibility Version** and a Red Hat Gluster Storage node already set up. For information on setting up a Red Hat Gluster Storage node, see the [Red Hat Gluster Storage Installation Guide](#). For more information on the compatibility matrix, see the [Red Hat Gluster Storage Version Compatibility and Support](#).

Procedure 9.1. Adding a Red Hat Gluster Storage Node

1. Click the **Hosts** resource tab to list the hosts in the results list.
2. Click **New** to open the **New Host** window.
3. Use the drop-down menus to select the **Data Center** and **Host Cluster** for the Red Hat Gluster Storage node.
4. Enter the **Name**, **Address**, and **SSH Port** of the Red Hat Gluster Storage node.
5. Select an authentication method to use with the Red Hat Gluster Storage node.
 - Enter the root user's password to use password authentication.
 - Copy the key displayed in the **SSH PublicKey** field to `/root/.ssh/authorized_keys` on the Red Hat Gluster Storage node to use public key authentication.
6. Click **OK** to add the node and close the window.

You have added a Red Hat Gluster Storage node to your Red Hat Virtualization environment. You can now use the volume and brick resources of the node in your environment.

9.1.2. Removing a Red Hat Gluster Storage Node

Remove a Red Hat Gluster Storage node from your Red Hat Virtualization environment.

Procedure 9.2. Removing a Red Hat Gluster Storage Node

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the Red Hat Gluster Storage node in the results list.
2. Click **Maintenance** to open the **Maintenance Host(s)** confirmation window.
3. Click **OK** to move the host to maintenance mode.

4. Click **Remove** to open the **Remove Host(s)** confirmation window.
5. Select the **Force Remove** check box if the node has volume bricks on it, or if the node is non-responsive.
6. Click **OK** to remove the node and close the window.

Your Red Hat Gluster Storage node has been removed from the environment and is no longer visible in the **Hosts** tab.

9.2. USING RED HAT GLUSTER STORAGE AS A STORAGE DOMAIN

9.2.1. Introduction to Red Hat Gluster Storage (GlusterFS) Volumes

Red Hat Gluster Storage volumes combine storage from more than one Red Hat Gluster Storage server into a single global namespace. A volume is a collection of bricks, where each brick is a mountpoint or directory on a Red Hat Gluster Storage Server in the trusted storage pool.

Most of the management operations of Red Hat Gluster Storage happen on the volume.

You can use the Administration Portal to create and start new volumes. You can monitor volumes in your Red Hat Gluster Storage cluster from the **Volumes** tab.

While volumes can be created and managed from the Administration Portal, bricks must be created on the individual Red Hat Gluster Storage nodes before they can be added to volumes using the Administration Portal

9.2.2. Gluster Storage Terminology

Table 9.1. Data Center Properties

Term	Definition
Brick	<p>A brick is the GlusterFS basic unit of storage, represented by an export directory on a server in the trusted storage pool. A Brick is expressed by combining a server with an export directory in the following format:</p> <p>SERVER:EXPORT</p> <p>For example:</p> <p>myhostname:/exports/myexportdir/</p>
Block Storage	<p>Block special files or block devices correspond to devices through which the system moves data in the form of blocks. These device nodes often represent addressable devices such as hard disks, CD-ROM drives, or memory-regions. Red Hat Gluster Storage supports XFS file system with extended attributes.</p>

Term	Definition
Cluster	A trusted pool of linked computers, working together closely thus in many respects forming a single computer. In Red Hat Gluster Storage terminology a cluster is called a trusted storage pool.
Client	The machine that mounts the volume (this may also be a server).
Distributed File System	A file system that allows multiple clients to concurrently access data spread across multiple servers/bricks in a trusted storage pool. Data sharing among multiple locations is fundamental to all distributed file systems.
Geo-Replication	Geo-replication provides a continuous, asynchronous, and incremental replication service from site to another over Local Area Networks (LAN), Wide Area Network (WAN), and across the Internet.
glusterd	The Gluster management daemon that needs to run on all servers in the trusted storage pool.
Metadata	Metadata is data providing information about one or more other pieces of data.
N-way Replication	Local synchronous data replication typically deployed across campus or Amazon Web Services Availability Zones.
Namespace	Namespace is an abstract container or environment created to hold a logical grouping of unique identifiers or symbols. Each Red Hat Gluster Storage trusted storage pool exposes a single namespace as a POSIX mount point that contains every file in the trusted storage pool.
POSIX	Portable Operating System Interface (for Unix) is the name of a family of related standards specified by the IEEE to define the application programming interface (API), along with shell and utilities interfaces for software compatible with variants of the UNIX operating system. Red Hat Gluster Storage exports a fully POSIX compatible file system.
RAID	Redundant Array of Inexpensive Disks (RAID) is a technology that provides increased storage reliability through redundancy, combining multiple low-cost, less-reliable disk drives components into a logical unit where all drives in the array are interdependent.

Term	Definition
RRDNS	Round Robin Domain Name Service (RRDNS) is a method to distribute load across application servers. RRDNS is implemented by creating multiple A records with the same name and different IP addresses in the zone file of a DNS server.
Server	The machine (virtual or bare-metal) which hosts the actual file system in which data will be stored.
Scale-Up Storage	Increases the capacity of the storage device, but only in a single dimension. An example might be adding additional disk capacity to a single computer in a trusted storage pool.
Scale-Out Storage	Increases the capability of a storage device in multiple dimensions. For example adding a server to a trusted storage pool increases CPU, disk capacity, and throughput for the trusted storage pool.
Subvolume	A subvolume is a brick after being processed by at least one translator.
Translator	A translator connects to one or more subvolumes, does something with them, and offers a subvolume connection.
Trusted Storage Pool	A storage pool is a trusted network of storage servers. When you start the first server, the storage pool consists of that server alone.
User Space	Applications running in user space do not directly interact with hardware, instead using the kernel to moderate access. User Space applications are generally more portable than applications in kernel space. Gluster is a user space application.
Virtual File System (VFS)	VFS is a kernel software layer that handles all system calls related to the standard Linux file system. It provides a common interface to several kinds of file systems.
Volume File	The volume file is a configuration file used by GlusterFS process. The volume file will usually be located at: <code>/var/lib/glusterd/vols/VOLNAME.</code>
Volume	A volume is a logical collection of bricks. Most of the Gluster management operations happen on the volume.

9.2.3. Attaching a Red Hat Gluster Storage Volume as a Storage Domain

Add a Red Hat Gluster Storage volume to the Red Hat Virtualization Manager to be used directly as a storage domain. This differs from adding a Red Hat Storage Gluster node, which enables control over the volumes and bricks of the node from within the Red Hat Virtualization Manager, and does not require a Gluster-enabled cluster.

The host requires the `glusterfs`, `glusterfs-fuse`, and `glusterfs-cli` packages to be installed in order to mount the volume. The `glusterfs-cli` package is available from the `rh-common-rpms` channel on the Customer Portal.

For information on setting up a Red Hat Gluster Storage node, see the [Red Hat Gluster Storage Installation Guide](#). For more information on preparing a host to be used with Red Hat Storage Gluster volumes, see the [Configuring Red Hat Virtualization with Red Hat Gluster Storage Guide](#). For more information on the compatibility matrix, see the [Red Hat Gluster Storage Version Compatibility and Support](#).

Procedure 9.3. Adding a Red Hat Gluster Storage Volume as a Storage Domain

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click **New Domain** to open the **New Domain** window.

New Domain ?

Data Center: Default (V3) | Name: RHGS_SD
 Domain Function: Data | Description: Red Hat Gluster Storage
 Storage Type: GlusterFS | Comment:
 Use Host: Host01

For data integrity make sure that the server is configured with Quorum (both client and server Quorum)

Path:
E.g.: myserver.mydomain.com:/myvolumename

VFS Type: glusterfs
 Mount Options:
 Advanced Parameters

Warning Low Space Indicator (%): 10
 Critical Space Action Blocker (GB): 5
 Format: V3
 Wipe After Delete:

OK Cancel

Figure 9.1. Red Hat Gluster Storage

3. Enter the **Name** for the storage domain.
4. Select the **Data Center** to be associated with the storage domain.
5. Select **Data** from the **Domain Function** drop-down list.
6. Select **GlusterFS** from the **Storage Type** drop-down list.
7. Select a host from the **Use Host** drop-down list. Only hosts within the selected data center will be listed. To mount the volume, the host that you select must have the `glusterfs` and `glusterfs-fuse` packages installed.
8. In the **Path** field, enter the IP address or FQDN of the Red Hat Gluster Storage server and the volume name separated by a colon.
9. Enter additional **Mount Options**, as you would normally provide them to the `mount` command using the `-o` argument. The mount options should be provided in a comma-separated list. See `man mount` for a list of valid mount options.
10. Optionally, you can configure the advanced parameters.
 - a. Click **Advanced Parameters**.
 - b. Enter a percentage value into the **Warning Low Space Indicator** field. If the free space available on the storage domain is below this percentage, warning messages are displayed to the user and logged.
 - c. Enter a GB value into the **Critical Space Action Blocker** field. If the free space available on the storage domain is below this value, error messages are displayed to the user and logged, and any new action that consumes space, even temporarily, will be blocked.
 - d. Select the **Wipe After Delete** check box to enable the wipe after delete option. This option can be edited after the domain is created, but doing so will not change the wipe after delete property of disks that already exist.
11. Click **OK** to mount the volume as a storage domain and close the window.

9.2.4. Creating a Storage Volume

You can create new volumes using the Administration Portal. When creating a new volume, you must specify the bricks that comprise the volume and specify whether the volume is to be distributed, replicated, or striped.

You must create brick directories or mountpoints before you can add them to volumes.



IMPORTANT

It is recommended that you use replicated volumes, where bricks exported from different hosts are combined into a volume. Replicated volumes create copies of files across multiple bricks in the volume, preventing data loss when a host is fenced.

Procedure 9.4. Creating A Storage Volume

1. Click the **Volumes** resource tab to list existing volumes in the results list.

2. Click **New** to open the **New Volume** window.
3. Use the drop-down menus to select the **Data Center** and **Volume Cluster**.
4. Enter the **Name** of the volume.
5. Use the drop-down menu to select the **Type** of the volume.
6. If active, select the appropriate **Transport Type** check box.
7. Click the **Add Bricks** button to select bricks to add to the volume. Bricks must be created externally on the Red Hat Gluster Storage nodes.
8. If active, use the **Gluster**, **NFS**, and **CIFS** check boxes to select the appropriate access protocols used for the volume.
9. Enter the volume access control as a comma-separated list of IP addresses or hostnames in the **Allow Access From** field.

You can use the * wildcard to specify ranges of IP addresses or hostnames.

10. Select the **Optimize for Virt Store** option to set the parameters to optimize your volume for virtual machine storage. Select this if you intend to use this volume as a storage domain.
11. Click **OK** to create the volume. The new volume is added and displays on the **Volume** tab.

You have added a Red Hat Gluster Storage volume. You can now use it for storage.

9.2.5. Adding Bricks to a Volume

Summary

You can expand your volumes by adding new bricks. You need to add at least one brick to a distributed volume, multiples of two bricks to replicated volumes, and multiples of four bricks to striped volumes when expanding your storage space.

Procedure 9.5. Adding Bricks to a Volume

1. On the **Volumes** tab on the navigation pane, select the volume to which you want to add bricks.
2. Click the **Bricks** tab from the Details pane.
3. Click **Add Bricks** to open the **Add Bricks** window.
4. Use the **Server** drop-down menu to select the server on which the brick resides.
5. Enter the path of the **Brick Directory**. The directory must already exist.
6. Click **Add**. The brick appears in the list of bricks in the volume, with server addresses and brick directory names.
7. Click **OK**.

Result

The new bricks are added to the volume and the bricks display in the volume's **Bricks** tab.

9.2.6. Explanation of Settings in the Add Bricks Window

Table 9.2. Add Bricks Tab Properties

Field Name	Description
Volume Type	Displays the type of volume. This field cannot be changed; it was set when you created the volume.
Server	The server where the bricks are hosted.
Brick Directory	The brick directory or mountpoint.

9.2.7. Optimizing Red Hat Gluster Storage Volumes to Store Virtual Machine Images

Optimize a Red Hat Gluster Storage volume to store virtual machine images using the Administration Portal.

To optimize a volume for storing virtual machines, the Manager sets a number of virtualization-specific parameters for the volume.

Volumes can be optimized to store virtual machines during creation by selecting the **Optimize for Virt Store** check box, or after creation using the **Optimize for Virt Store** button from the **Volumes** resource tab.

IMPORTANT

If a volume is replicated across three or more nodes, ensure the volume is optimized for virtual storage to avoid data inconsistencies across the nodes.

An alternate method is to access one of the Red Hat Gluster Storage nodes and set the volume group to **virt**. This sets the **cluster.quorum-type** parameter to **auto**, and the **cluster.server-quorum-type** parameter to **server**.

```
# gluster volume set VOLUME_NAME group virt
```

Verify the status of the volume by listing the volume information:

```
# gluster volume info VOLUME_NAME
```

9.2.8. Starting Volumes

Summary

After a volume has been created or an existing volume has been stopped, it needs to be started before it can be used.

Procedure 9.6. Starting Volumes

1. In the **Volumes** tab, select the volume to be started.

You can select multiple volumes to start by using **Shift** or **Ctrl** key.

2. Click the **Start** button.

The volume status changes to **Up**.

Result

You can now use your volume for virtual machine storage.

9.2.9. Tuning Volumes

Summary

Tuning volumes allows you to affect their performance. To tune volumes, you add options to them.

Procedure 9.7. Tuning Volumes

1. Click the **Volumes** tab.

A list of volumes displays.

2. Select the volume that you want to tune, and click the **Volume Options** tab from the Details pane.

The **Volume Options** tab displays a list of options set for the volume.

3. Click **Add** to set an option. The **Add Option** dialog box displays. Select the Option Key from the drop down list and enter the option value.

4. Click **OK**.

The option is set and displays in the **Volume Options** tab.

Result

You have tuned the options for your storage volume.

9.2.10. Editing Volume Options

Summary

You have tuned your volume by adding options to it. You can change the options for your storage volume.

Procedure 9.8. Editing Volume Options

1. Click the **Volumes** tab.

A list of volumes displays.

2. Select the volume that you want to edit, and click the **Volume Options** tab from the Details pane.

The **Volume Options** tab displays a list of options set for the volume.

3. Select the option you want to edit. Click **Edit**. The **Edit Option** dialog box displays. Enter a new value for the option.
4. Click **OK**.

The edited option displays in the **Volume Options** tab.

Result

You have changed the options on your volume.

9.2.11. Reset Volume Options

Summary

You can reset options to revert them to their default values.

1. Click the **Volumes** tab.

A list of volumes displays.

2. Select the volume and click the **Volume Options** tab from the Details pane.

The **Volume Options** tab displays a list of options set for the volume.

3. Select the option you want to reset. Click **Reset**. A dialog box displays, prompting to confirm the reset option.
4. Click **OK**.

The selected option is reset.



NOTE

You can reset all volume options by clicking **Reset All** button. A dialog box displays, prompting to confirm the reset option. Click **OK**. All volume options are reset for the selected volume.

Result

You have reset volume options to default.

9.2.12. Removing Bricks from a Volume

Summary

You can shrink volumes, as needed, while the cluster is online and available. For example, you might need to remove a brick that has become inaccessible in a distributed volume due to hardware or network failure.

Procedure 9.9. Removing Bricks from a Volume

1. On the **Volumes** tab on the navigation pane, select the volume from which you wish to remove bricks.

2. Click the **Bricks** tab from the Details pane.
3. Select the bricks you wish to remove. Click **Remove Bricks**.
4. A window opens, prompting to confirm the deletion. Click **OK** to confirm.

Result

The bricks are removed from the volume.

9.2.13. Stopping Red Hat Gluster Storage Volumes

After a volume has been started, it can be stopped.

Procedure 9.10. Stopping Volumes

1. In the **Volumes** tab, select the volume to be stopped.
You can select multiple volumes to stop by using **Shift** or **Ctrl** key.
2. Click **Stop**.

9.2.14. Deleting Red Hat Gluster Storage Volumes

You can delete a volume or multiple volumes from your cluster.

1. In the **Volumes** tab, select the volume to be deleted.
2. Click **Remove**. A dialog box displays, prompting to confirm the deletion. Click **OK**.

9.2.15. Rebalancing Volumes

Summary

If a volume has been expanded or shrunk by adding or removing bricks to or from that volume, the data on the volume must be rebalanced amongst the servers.

Procedure 9.11. Rebalancing a Volume

1. Click the **Volumes** tab.
A list of volumes displays.
2. Select the volume to rebalance.
3. Click **Rebalance**.

Result

The selected volume is rebalanced.

9.3. CLUSTERS AND GLUSTER HOOKS

9.3.1. Managing Gluster Hooks

Gluster hooks are volume life cycle extensions. You can manage Gluster hooks from the Manager. The content of the hook can be viewed if the hook content type is **Text**.

Through the Manager, you can perform the following:

- View a list of hooks available in the hosts.
- View the content and status of hooks.
- Enable or disable hooks.
- Resolve hook conflicts.

9.3.2. Listing Hooks

Summary

List the Gluster hooks in your environment.

Procedure 9.12. Listing a Hook

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Gluster Hooks** sub-tab to list the hooks in the details pane.

Result

You have listed the Gluster hooks in your environment.

9.3.3. Viewing the Content of Hooks

Summary

View the content of a Gluster hook in your environment.

Procedure 9.13. Viewing the Content of a Hook

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Gluster Hooks** sub-tab to list the hooks in the details pane.
3. Select a hook with content type **Text** and click the **View Content** button to open the **Hook Content** window.

Result

You have viewed the content of a hook in your environment.

9.3.4. Enabling or Disabling Hooks

Summary

Toggle the activity of a Gluster hook by enabling or disabling it.

Procedure 9.14. Enabling or Disabling a Hook

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Gluster Hooks** sub-tab to list the hooks in the details pane.
3. Select a hook and click one of the **Enable** or **Disable** buttons. The hook is enabled or disabled on all nodes of the cluster.

Result

You have toggled the activity of a Gluster hook in your environment.

9.3.5. Refreshing Hooks

Summary

By default, the Manager checks the status of installed hooks on the engine and on all servers in the cluster and detects new hooks by running a periodic job every hour. You can refresh hooks manually by clicking the **Sync** button.

Procedure 9.15. Refreshing a Hook

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Gluster Hooks** sub-tab to list the hooks in the details pane.
3. Click the **Sync** button.

Result

The hooks are synchronized and updated in the details pane.

9.3.6. Resolving Conflicts

The hooks are displayed in the **Gluster Hooks** sub-tab of the **Cluster** tab. Hooks causing a conflict are displayed with an exclamation mark. This denotes either that there is a conflict in the content or the status of the hook across the servers in the cluster, or that the hook script is missing in one or more servers. These conflicts can be resolved via the Manager. The hooks in the servers are periodically synchronized with engine database and the following conflicts can occur for the hooks:

- Content Conflict - the content of the hook is different across servers.
- Missing Conflict - one or more servers of the cluster do not have the hook.
- Status Conflict - the status of the hook is different across servers.
- Multiple Conflicts - a hook has a combination of two or more of the aforementioned conflicts.

9.3.7. Resolving Content Conflicts

Summary

A hook that is not consistent across the servers and engine will be flagged as having a conflict. To resolve the conflict, you must select a version of the hook to be copied across all servers and the engine.

Procedure 9.16. Resolving a Content Conflict

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Gluster Hooks** sub-tab to list the hooks in the details pane.
3. Select the conflicting hook and click the **Resolve Conflicts** button to open the **Resolve Conflicts** window.
4. Select the engine or a server from the list of sources to view the content of that hook and establish which version of the hook to copy.

**NOTE**

The content of the hook will be overwritten in all servers and in the engine.

5. Use the **Use content from** drop-down menu to select the preferred server or the engine.
6. Click **OK** to resolve the conflict and close the window.

Result

The hook from the selected server is copied across all servers and the engine to be consistent across the environment.

9.3.8. Resolving Missing Hook Conflicts**Summary**

A hook that is not present on all the servers and the engine will be flagged as having a conflict. To resolve the conflict, either select a version of the hook to be copied across all servers and the engine, or remove the missing hook entirely.

Procedure 9.17. Resolving a Missing Hook Conflict

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Gluster Hooks** sub-tab to list the hooks in the details pane.
3. Select the conflicting hook and click the **Resolve Conflicts** button to open the **Resolve Conflicts** window.
4. Select any source with a status of **Enabled** to view the content of the hook.
5. Select the appropriate radio button, either **Copy the hook to all the servers** or **Remove the missing hook**. The latter will remove the hook from the engine and all servers.
6. Click **OK** to resolve the conflict and close the window.

Result

Depending on your chosen resolution, the hook has either been removed from the environment entirely, or has been copied across all servers and the engine to be consistent across the environment.

9.3.9. Resolving Status Conflicts

Summary

A hook that does not have a consistent status across the servers and engine will be flagged as having a conflict. To resolve the conflict, select a status to be enforced across all servers in the environment.

Procedure 9.18. Resolving a Status Conflict

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Cluster Hooks** sub-tab to list the hooks in the details pane.
3. Select the conflicting hook and click the **Resolve Conflicts** button to open the **Resolve Conflicts** window.
4. Set **Hook Status** to **Enable** or **Disable**.
5. Click **OK** to resolve the conflict and close the window.

Result

The selected status for the hook is enforced across the engine and the servers to be consistent across the environment.

9.3.10. Resolving Multiple Conflicts

Summary

A hook may have a combination of two or more conflicts. These can all be resolved concurrently or independently through the **Resolve Conflicts** window. This procedure will resolve all conflicts for the hook so that it is consistent across the engine and all servers in the environment.

Procedure 9.19. Resolving Multiple Conflicts

1. Use the **Cluster** resource tab, tree mode, or the search function to find and select a cluster in the results list.
2. Select the **Cluster Hooks** sub-tab to list the hooks in the details pane.
3. Select the conflicting hook and click the **Resolve Conflicts** button to open the **Resolve Conflicts** window.
4. Choose a resolution to each of the affecting conflicts, as per the appropriate procedure.
5. Click **OK** to resolve the conflicts and close the window.

Result

You have resolved all of the conflicts so that the hook is consistent across the engine and all servers.

9.3.11. Managing Gluster Sync

The Gluster Sync feature periodically fetches the latest cluster configuration from GlusterFS and syncs the same with the engine DB. This process can be performed through the Manager. When a cluster is selected, the user is provided with the option to import hosts or detach existing hosts from the selected cluster. You can perform Gluster Sync if there is a host in the cluster.



NOTE

The Manager continuously monitors if hosts are added to or removed from the storage cluster. If the addition or removal of a host is detected, an action item is shown in the **General** tab for the cluster, where you can either to choose to **Import** the host into or **Detach** the host from the cluster.

CHAPTER 10. POOLS

10.1. INTRODUCTION TO VIRTUAL MACHINE POOLS

A virtual machine pool is a group of virtual machines that are all clones of the same template and that can be used on demand by any user in a given group. Virtual machine pools allow administrators to rapidly configure a set of generalized virtual machines for users.

Users access a virtual machine pool by taking a virtual machine from the pool. When a user takes a virtual machine from a pool, they are provided with any one of the virtual machines in the pool if any are available. That virtual machine will have the same operating system and configuration as that of the template on which the pool was based, but users may not receive the same member of the pool each time they take a virtual machine. Users can also take multiple virtual machines from the same virtual machine pool depending on the configuration of that pool.

Virtual machine pools are stateless by default, meaning that virtual machine data is not persistent across reboots. However, the pool can be configured to be stateful, allowing changes made by a previous user to persist. However, if a user configures console options for a virtual machine taken from a virtual machine pool, those options will be set as the default for that user for that virtual machine pool.



NOTE

Virtual machines taken from a pool are not stateless when accessed from the Administration Portal. This is because administrators need to be able to write changes to the disk if necessary.

In principle, virtual machines in a pool are started when taken by a user, and shut down when the user is finished. However, virtual machine pools can also contain pre-started virtual machines. Pre-started virtual machines are kept in an up state, and remain idle until they are taken by a user. This allows users to start using such virtual machines immediately, but these virtual machines will consume system resources even while not in use due to being idle.

10.2. VIRTUAL MACHINE POOL TASKS

10.2.1. Creating a Virtual Machine Pool

You can create a virtual machine pool that contains multiple virtual machines that have been created based on a common template.

Procedure 10.1. Creating a Virtual Machine Pool

1. Click the **Pools** tab.
2. Click the **New** button to open the **New Pool** window.
3. Use the drop down-list to select the **Cluster** or use the selected default.

4. Use the **Template** drop-down menu to select the required template and version or use the selected default. A template provides standard settings for all the virtual machines in the pool.
5. Use the **Operating System** drop-down list to select an **Operating System** or use the default provided by the template.
6. Use the **Optimized for** drop-down list to optimize virtual machines for either **Desktop** use or **Server** use.
7. Enter a **Name** and **Description**, any **Comments**, and the **Number of VMs** for the pool.
8. Enter the number of virtual machines to be prestarted in the **Prestarted VMs** field.
9. Select the **Maximum number of VMs per user** that a single user is allowed to run in a session. The minimum is one.
10. Select the **Delete Protection** check box to enable delete protection.
11. Optionally, click the **Show Advanced Options** button and perform the following steps:
 - a. Click the **Type** tab:
 - i. Select a **Pool Type**:
 - **Manual** - The administrator is responsible for explicitly returning the virtual machine to the pool.
 - **Automatic** - The virtual machine is automatically returned to the virtual machine pool.
 - ii. Select the **Stateful Pool** check box to ensure that virtual machines are started in a stateful mode. This means that changes made by a previous user will persist on a virtual machine.
 - b. Click the **Console** tab:
 - i. Select the **Override SPICE Proxy** check box.
 - ii. In the **Overridden SPICE proxy address** text field, specify the address of a SPICE proxy to override the global SPICE proxy.
12. Click **OK**.

You have created and configured a virtual machine pool with the specified number of identical virtual machines. You can view these virtual machines in the **Virtual Machines** resource tab, or in the details pane of the **Pools** resource tab; a virtual machine in a pool is distinguished from independent virtual machines by its icon.

10.2.2. Explanation of Settings and Controls in the New Pool and Edit Pool Windows

10.2.2.1. New Pool and Edit Pool General Settings Explained

The following table details the information required on the **General** tab of the **New Pool** and **Edit Pool** windows that are specific to virtual machine pools. All other settings are

identical to those in the **New Virtual Machine** window.

Table 10.1. General settings

Field Name	Description
Template	The template and template sub version on which the virtual machine pool is based. If you create a pool based on the latest sub version of a template, all virtual machines in the pool, when rebooted, will automatically receive the latest template version. For more information on configuring templates for virtual machines see Virtual Machine General Settings Explained and Explanation of Settings in the New Template and Edit Template Windows in the <i>Virtual Machine Management Guide</i> .
Description	A meaningful description of the virtual machine pool.
Comment	A field for adding plain text human-readable comments regarding the virtual machine pool.
Prestarted VMs	Allows you to specify the number of virtual machines in the virtual machine pool that will be started before they are taken and kept in that state to be taken by users. The value of this field must be between 0 and the total number of virtual machines in the virtual machine pool.
Number of VMs/Increase number of VMs in pool by	Allows you to specify the number of virtual machines to be created and made available in the virtual machine pool. In the edit window it allows you to increase the number of virtual machines in the virtual machine pool by the specified number. By default, the maximum number of virtual machines you can create in a pool is 1000. This value can be configured using the MaxVmsInPool key of the engine-config command.
Maximum number of VMs per user	Allows you to specify the maximum number of virtual machines a single user can take from the virtual machine pool at any one time. The value of this field must be between 1 and 32,767 .
Delete Protection	Allows you to prevent the virtual machines in the pool from being deleted.

10.2.2.2. New and Edit Pool Type Settings Explained

The following table details the information required on the **Type** tab of the **New Pool** and **Edit Pool** windows.

Table 10.2. Type settings

Field Name	Description
Pool Type	<p>This drop-down menu allows you to specify the type of the virtual machine pool. The following options are available:</p> <ul style="list-style-type: none"> • Automatic: After a user finishes using a virtual machine taken from a virtual machine pool, that virtual machine is automatically returned to the virtual machine pool. • Manual: After a user finishes using a virtual machine taken from a virtual machine pool, that virtual machine is only returned to the virtual machine pool when an administrator manually returns the virtual machine.
Stateful Pool	<p>Specify whether the state of virtual machines in the pool is preserved when a virtual machine is passed to a different user. This means that changes made by a previous user will persist on the virtual machine.</p>

10.2.2.3. New Pool and Edit Pool Console Settings Explained

The following table details the information required on the **Console** tab of the **New Pool** or **Edit Pool** window that is specific to virtual machine pools. All other settings are identical to those in the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table 10.3. Console settings

Field Name	Description
Override SPICE proxy	<p>Select this check box to enable overriding the SPICE proxy defined in global configuration. This feature is useful in a case where the user (who is, for example, connecting via the User Portal) is outside of the network where the hosts reside.</p>
Overridden SPICE proxy address	<p>The proxy by which the SPICE client will connect to virtual machines. This proxy overrides both the global SPICE proxy defined for the Red Hat Virtualization environment and the SPICE proxy defined for the cluster to which the virtual machine pool belongs, if any. The address must be in the following format:</p> <pre>protocol://[host]:[port]</pre>

10.2.2.4. Virtual Machine Pool Host Settings Explained

The following table details the options available on the **Host** tab of the **New Pool** and **Edit Pool** windows.

Table 10.4. Virtual Machine Pool: Host Settings

Field Name	Sub-element	Description
Start Running On		<p>Defines the preferred host on which the virtual machine is to run. Select either:</p> <ul style="list-style-type: none">• Any Host in Cluster - The virtual machine can start and run on any available host in the cluster.• Specific - The virtual machine will start running on a particular host in the cluster. However, the Manager or an administrator can migrate the virtual machine to a different host in the cluster depending on the migration and high-availability settings of the virtual machine. Select the specific host or group of hosts from the list of available hosts.

Field Name	Sub-element	Description
Migration Options	Migration mode	<p>Defines options to run and migrate the virtual machine. If the options here are not used, the virtual machine will run or migrate according to its cluster's policy.</p> <ul style="list-style-type: none">• Allow manual and automatic migration - The virtual machine can be automatically migrated from one host to another in accordance with the status of the environment, or manually by an administrator.• Allow manual migration only - The virtual machine can only be migrated from one host to another manually by an administrator.• Do not allow migration - The virtual machine cannot be migrated, either automatically or manually.

Field Name	Sub-element	Description
	Use custom migration policy	<p>Defines the migration convergence policy. If the check box is left unselected, the host determines the policy.</p> <ul style="list-style-type: none"><li data-bbox="1099 421 1430 667">• Legacy - Legacy behavior of 3.6 version. Overrides in vdsm.conf are still applied. The guest agent hook mechanism is disabled.<li data-bbox="1099 696 1430 1310">• Minimal downtime - Allows the virtual machine to migrate in typical situations. Virtual machines should not experience any significant downtime. The migration will be aborted if virtual machine migration does not converging after a long time (dependent on QEMU iterations, with a maximum of 500 milliseconds). The guest agent hook mechanism is enabled.<li data-bbox="1099 1339 1430 1868">• Suspend workload if needed - Allows the virtual machine to migrate in most situations, including when the virtual machine is running a heavy workload. Virtual machines may experience a more significant downtime. The migration may still be aborted for extreme workloads. The guest agent hook mechanism is enabled.

Field Name	Sub-element	Description
	Use custom migration downtime	This check box allows you to specify the maximum number of milliseconds the virtual machine can be down during live migration. Configure different maximum downtimes for each virtual machine according to its workload and SLA requirements. Enter 0 to use the VDSM default value.

Field Name	Sub-element	Description
	Auto Converge migrations	<p>Only activated with Legacy migration policy. Allows you to set whether auto-convergence is used during live migration of the virtual machine. Large virtual machines with high workloads can dirty memory more quickly than the transfer rate achieved during live migration, and prevent the migration from converging. Auto-convergence capabilities in QEMU allow you to force convergence of virtual machine migrations. QEMU automatically detects a lack of convergence and triggers a throttle-down of the vCPUs on the virtual machine. Auto-convergence is disabled globally by default.</p> <ul style="list-style-type: none">• Select Inherit from cluster setting to use the auto-convergence setting that is set at the cluster level. This option is selected by default.• Select Auto Converge to override the cluster setting or global setting and allow auto-convergence for the virtual machine.• Select Don't Auto Converge to override the cluster setting or global setting and prevent auto-convergence for the virtual machine.

Field Name	Sub-element	Description
	Enable migration compression	<p>Only activated with Legacy migration policy. The option allows you to set whether migration compression is used during live migration of the virtual machine. This feature uses Xor Binary Zero Run-Length-Encoding to reduce virtual machine downtime and total live migration time for virtual machines running memory write-intensive workloads or for any application with a sparse memory update pattern. Migration compression is disabled globally by default.</p> <ul style="list-style-type: none"> • Select Inherit from cluster setting to use the compression setting that is set at the cluster level. This option is selected by default. • Select Compress to override the cluster setting or global setting and allow compression for the virtual machine. • Select Don't compress to override the cluster setting or global setting and prevent compression for the virtual machine.
	Pass-Through Host CPU	<p>This check box allows virtual machines to take advantage of the features of the physical CPU of the host on which they are situated. This option can only be enabled when Do not allow migration is selected.</p>
Configure NUMA	NUMA Node Count	<p>The number of virtual NUMA nodes to assign to the virtual machine. If the Tune Mode is Preferred, this value must be set to 1.</p>

Field Name	Sub-element	Description
	Tune Mode	<p>The method used to allocate memory.</p> <ul style="list-style-type: none"> • Strict: Memory allocation will fail if the memory cannot be allocated on the target node. • Preferred: Memory is allocated from a single preferred node. If sufficient memory is not available, memory can be allocated from other nodes. • Interleave: Memory is allocated across nodes in a round-robin algorithm.
	NUMA Pinning	<p>Opens the NUMA Topology window. This window shows the host's total CPUs, memory, and NUMA nodes, and the virtual machine's virtual NUMA nodes. Pin virtual NUMA nodes to host NUMA nodes by clicking and dragging each vNUMA from the box on the right to a NUMA node on the left.</p>

10.2.3. Editing a Virtual Machine Pool

10.2.3.1. Editing a Virtual Machine Pool

After a virtual machine pool has been created, its properties can be edited. The properties available when editing a virtual machine pool are identical to those available when creating a new virtual machine pool except that the **Number of VMs** property is replaced by **Increase number of VMs in pool by**.



NOTE

When editing a virtual machine pool, the changes introduced affect only new virtual machines. Virtual machines that existed already at the time of the introduced changes remain unaffected.

Procedure 10.2. Editing a Virtual Machine Pool

1. Click the **Pools** resource tab, and select a virtual machine pool from the results list.

2. Click **Edit** to open the **Edit Pool** window.
3. Edit the properties of the virtual machine pool.
4. Click **Ok**.

10.2.3.2. Prestarting Virtual Machines in a Pool

The virtual machines in a virtual machine pool are powered down by default. When a user requests a virtual machine from a pool, a machine is powered up and assigned to the user. In contrast, a prestarted virtual machine is already running and waiting to be assigned to a user, decreasing the amount of time a user has to wait before being able to access a machine. When a prestarted virtual machine is shut down it is returned to the pool and restored to its original state. The maximum number of prestarted virtual machines is the number of virtual machines in the pool.

Prestarted virtual machines are suitable for environments in which users require immediate access to virtual machines which are not specifically assigned to them. Only automatic pools can have prestarted virtual machines.

Procedure 10.3. Prestarting Virtual Machines in a Pool

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Click **Edit** to open the **Edit Pool** window.
3. Enter the number of virtual machines to be prestarted in the **Prestarted VMs** field.
4. Select the **Pool** tab. Ensure **Pool Type** is set to **Automatic**.
5. Click **OK**.

You have set a number of prestarted virtual machines in a pool. The prestarted machines are running and available for use.

10.2.3.3. Adding Virtual Machines to a Virtual Machine Pool

If you require more virtual machines than originally provisioned in a virtual machine pool, add more machines to the pool.

Procedure 10.4. Adding Virtual Machines to a Virtual Machine Pool

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Click **Edit** to open the **Edit Pool** window.
3. Enter the number of additional virtual machines to add in the **Increase number of VMs in pool by** field.
4. Click **OK**.

You have added more virtual machines to the virtual machine pool.

10.2.3.4. Detaching Virtual Machines from a Virtual Machine Pool

You can detach virtual machines from a virtual machine pool. Detaching a virtual machine removes it from the pool to become an independent virtual machine.

Procedure 10.5. Detaching Virtual Machines from a Virtual Machine Pool

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Ensure the virtual machine has a status of **Down** because you cannot detach a running virtual machine.

Click the **Virtual Machines** tab in the details pane to list the virtual machines in the pool.

3. Select one or more virtual machines and click **Detach** to open the **Detach Virtual Machine(s)** confirmation window.
4. Click **OK** to detach the virtual machine from the pool.



NOTE

The virtual machine still exists in the environment and can be viewed and accessed from the **Virtual Machines** resource tab. Note that the icon changes to denote that the detached virtual machine is an independent virtual machine.

You have detached a virtual machine from the virtual machine pool.

10.2.4. Removing a Virtual Machine Pool

You can remove a virtual machine pool from a data center. You must first either delete or detach all of the virtual machines in the pool. Detaching virtual machines from the pool will preserve them as independent virtual machines.

Procedure 10.6. Removing a Virtual Machine Pool

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Click **Remove** to open the **Remove Pool(s)** confirmation window.
3. Click **OK** to remove the pool.

You have removed the pool from the data center.

10.3. POOLS AND PERMISSIONS

10.3.1. Managing System Permissions for a Virtual Machine Pool

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them

to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A virtual machine pool administrator is a system administration role for virtual machine pools in a data center. This role can be applied to specific virtual machine pools, to a data center, or to the whole virtualized environment; this is useful to allow different users to manage certain virtual machine pool resources.

The virtual machine pool administrator role permits the following actions:

- Create, edit, and remove pools.
- Add and detach virtual machines from the pool.



NOTE

You can only assign roles and permissions to existing users.

10.3.2. Virtual Machine Pool Administrator Roles Explained

Pool Permission Roles

The table below describes the administrator roles and privileges applicable to pool administration.

Table 10.5. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
VmPoolAdmin	System Administrator role of a virtual pool.	Can create, delete, and configure a virtual pool, assign and remove virtual pool users, and perform basic operations on a virtual machine.
ClusterAdmin	Cluster Administrator	Can use, create, delete, manage all virtual machine pools in a specific cluster.

10.3.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 10.7. Assigning a Role to a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.

3. Click **Add**.
4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down list.
6. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

10.3.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 10.8. Removing a Role from a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

10.4. TRUSTED COMPUTE POOLS

Trusted compute pools are secure clusters based on Intel Trusted Execution Technology (Intel TXT). Trusted clusters only allow hosts that are verified by Intel's OpenAttestation, which measures the integrity of the host's hardware and software against a White List database. Trusted hosts and the virtual machines running on them can be assigned tasks that require higher security. For more information on Intel TXT, trusted systems, and attestation, see <https://software.intel.com/en-us/articles/intel-trusted-execution-technology-intel-txt-enabling-guide>.

Creating a trusted compute pool involves the following steps:

- Configuring the Manager to communicate with an OpenAttestation server.
- Creating a trusted cluster that can only run trusted hosts.
- Adding trusted hosts to the trusted cluster. Hosts must be running the OpenAttestation agent to be verified as trusted by the OpenAttestation sever.

For information on installing an OpenAttestation server, installing the OpenAttestation agent on hosts, and creating a White List database, see <https://github.com/OpenAttestation/OpenAttestation/wiki>.

10.4.1. Connecting an OpenAttestation Server to the Manager

Before you can create a trusted cluster, the Red Hat Virtualization Manager must be configured to recognize the OpenAttestation server. Use **engine-config** to add the OpenAttestation server's FQDN or IP address:

```
# engine-config -s AttestationServer=attestationserver.example.com
```

The following settings can also be changed if required:

Table 10.6. OpenAttestation Settings for engine-config

Option	Default Value	Description
AttestationServer	oat-server	The FQDN or IP address of the OpenAttestation server. This must be set for the Manager to communicate with the OpenAttestation server.
AttestationPort	8443	The port used by the OpenAttestation server to communicate with the Manager.
AttestationTruststore	TrustStore.jks	The trust store used for securing communication with the OpenAttestation server.
AttestationTruststorePass	password	The password used to access the trust store.
AttestationFirstStageSize	10	Used for quick initialization. Changing this value without good reason is not recommended.
SecureConnectionWithOATServers	true	Enables or disables secure communication with OpenAttestation servers.
PollUri	AttestationService/resources/PollHosts	The URI used for accessing the OpenAttestation service.

10.4.2. Creating a Trusted Cluster

Trusted clusters communicate with an OpenAttestation server to assess the security of hosts. When a host is added to a trusted cluster, the OpenAttestation server measures the host's hardware and software against a White List database. Virtual machines can be migrated between trusted hosts in the trusted cluster, allowing for high availability in a secure environment.

Procedure 10.9. Creating a Trusted Cluster

1. Select the **Clusters** tab.

2. Click **New**.
3. Enter a **Name** for the cluster.
4. Select the **Enable Virt Service** radio button.
5. In the **Scheduling Policy** tab, select the **Enable Trusted Service** check box.
6. Click **OK**.

10.4.3. Adding a Trusted Host

Red Hat Enterprise Linux hosts can be added to trusted clusters and measured against a White List database by the OpenAttestation server. Hosts must meet the following requirements to be trusted by the OpenAttestation server:

- Intel TXT is enabled in the BIOS.
- The OpenAttestation agent is installed and running.
- Software running on the host matches the OpenAttestation server's White List database.

Procedure 10.10. Adding a Trusted Host

1. Select the **Hosts** tab.
2. Click **New**.
3. Select a trusted cluster from the **Host Cluster** drop-down list.
4. Enter a **Name** for the host.
5. Enter the **Address** of the host.
6. Enter the host's root **Password**.
7. Click **OK**.

After the host is added to the trusted cluster, it is assessed by the OpenAttestation server. If a host is not trusted by the OpenAttestation server, it will move to a **Non Operational** state and should be removed from the trusted cluster.

CHAPTER 11. VIRTUAL DISKS

11.1. UNDERSTANDING VIRTUAL MACHINE STORAGE

Red Hat Virtualization supports three storage types: NFS, iSCSI and FCP.

In each type, a host known as the Storage Pool Manager (SPM) manages access between hosts and storage. The SPM host is the only node that has full access within the storage pool; the SPM can modify the storage domain metadata, and the pool's metadata. All other hosts can only access virtual machine hard disk image data.

By default in an NFS, local, or POSIX compliant data center, the SPM creates the virtual disk using a thin provisioned format as a file in a file system.

In iSCSI and other block-based data centers, the SPM creates a volume group on top of the Logical Unit Numbers (LUNs) provided, and makes logical volumes to use as virtual disks. Virtual disks on block-based storage are preallocated by default.

If the virtual disk is preallocated, a logical volume of the specified size in GB is created. The virtual machine can be mounted on a Red Hat Enterprise Linux server using **kpartx**, **vgscan**, **vgchange** or **mount** to investigate the virtual machine's processes or problems.

If the virtual disk is thinly provisioned, a 1 GB logical volume is created. The logical volume is continuously monitored by the host on which the virtual machine is running. As soon as the usage nears a threshold the host notifies the SPM, and the SPM extends the logical volume by 1 GB. The host is responsible for resuming the virtual machine after the logical volume has been extended. If the virtual machine goes into a paused state it means that the SPM could not extend the disk in time. This occurs if the SPM is too busy or if there is not enough storage space.

A virtual disk with a preallocated (RAW) format has significantly faster write speeds than a virtual disk with a thin provisioning (QCOW2) format. Thin provisioning takes significantly less time to create a virtual disk. The thin provision format is suitable for non-I/O intensive virtual machines. The preallocated format is recommended for virtual machines with high I/O writes. If a virtual machine is able to write more than 1 GB every four seconds, use preallocated disks where possible.

11.2. UNDERSTANDING VIRTUAL DISKS

Red Hat Virtualization features **Preallocated** (thick provisioned) and **Sparse** (thin provisioned) storage options.

- Preallocated

A preallocated virtual disk allocates all the storage required for a virtual machine up front. For example, a 20 GB preallocated logical volume created for the data partition of a virtual machine will take up 20 GB of storage space immediately upon creation.

- Sparse

A sparse allocation allows an administrator to define the total storage to be assigned to the virtual machine, but the storage is only allocated when required.

For example, a 20 GB thin provisioned logical volume would take up 0 GB of storage

space when first created. When the operating system is installed it may take up the size of the installed file, and would continue to grow as data is added up to a maximum of 20 GB size.

The size of a disk is listed in the **Disks** sub-tab for each virtual machine and template. The **Virtual Size** of a disk is the total amount of disk space that the virtual machine can use; it is the number that you enter in the **Size(GB)** field when a disk is created or edited. The **Actual Size** of a disk is the amount of disk space that has been allocated to the virtual machine so far. Preallocated disks show the same value for both fields. Sparse disks may show a different value in the **Actual Size** field from the value in the **Virtual Size** field, depending on how much of the disk space has been allocated.



NOTE

When creating a Cinder virtual disk, the format and type of the disk are handled internally by Cinder and are not managed by Red Hat Virtualization.

The possible combinations of storage types and formats are described in the following table.

Table 11.1. Permitted Storage Combinations

Storage	Format	Type	Note
NFS or iSCSI/FCP	RAW or QCOW2	Sparse or Preallocated	
NFS	RAW	Preallocated	A file with an initial size which equals the amount of storage defined for the virtual disk, and has no formatting.
NFS	RAW	Sparse	A file with an initial size which is close to zero, and has no formatting.
NFS	QCOW2	Sparse	A file with an initial size which is close to zero, and has QCOW2 formatting. Subsequent layers will be QCOW2 formatted.
SAN	RAW	Preallocated	A block device with an initial size which equals the amount of storage defined for the virtual disk, and has no formatting.


Storage	Format	Type	Note
NFS or iSCSI/FCP	RAW or QCOW2	Sparse or Preallocated	
SAN	QCOW2	Sparse	A block device with an initial size which is much smaller than the size defined for the virtual disk (currently 1 GB), and has QCOW2 formatting for which space is allocated as needed (currently in 1 GB increments).

11.3. SETTINGS TO WIPE VIRTUAL DISKS AFTER DELETION

The `wipe_after_delete` flag, viewed in the Administration Portal as the **Wipe After Delete** check box will replace used data with zeros when a virtual disk is deleted. If it is set to false, which is the default, deleting the disk will open up those blocks for re-use but will not wipe the data. It is, therefore, possible for this data to be recovered because the blocks have not been returned to zero.

The `wipe_after_delete` flag only works on block storage. On file storage, for example NFS, the option does nothing because the file system will ensure that no data exists.

Enabling `wipe_after_delete` for virtual disks is more secure, and is recommended if the virtual disk has contained any sensitive data. This is a more intensive operation and users may experience degradation in performance and prolonged delete times.



NOTE

The wipe after delete functionality is not the same as secure delete, and cannot guarantee that the data is removed from the storage, just that new disks created on same storage will not expose data from old disks.

The `wipe_after_delete` flag default can be changed to `true` during the setup process (see [Configuring the Red Hat Virtualization Manager](#) in the *Installation Guide*), or by using the engine configuration tool on the Red Hat Virtualization Manager. Restart the engine for the setting change to take effect.



NOTE

Changing the `wipe_after_delete` flag default will not change the **Wipe After Delete** property of disks that already exist.

Procedure 11.1. Setting SANWipeAfterDelete to Default to True Using the Engine Configuration Tool

1. Run the engine configuration tool with the `--set` action:

```
# engine-config --set SANWipeAfterDelete=true
```

- Restart the engine for the change to take effect:

```
# systemctl restart ovirt-engine.service
```

The `/var/log/vdsm/vdsm.log` file located on the host can be checked to confirm that a virtual disk was successfully wiped and deleted.

For a successful wipe, the log file will contain the entry, ***storage_domain_id/volume_id was zeroed and will be deleted***. For example:

```
a9cb0625-d5dc-49ab-8ad1-72722e82b0bf/a49351a7-15d8-4932-8d67-512a369f9d61  
was zeroed and will be deleted
```

For a successful deletion, the log file will contain the entry, ***finished with VG:storage_domain_id LVs: list_of_volume_ids, img: image_id***. For example:

```
finished with VG:a9cb0625-d5dc-49ab-8ad1-72722e82b0bf LVs: {'a49351a7-  
15d8-4932-8d67-512a369f9d61': ImsPar(imgs=['11f8b3be-fa96-4f6a-bb83-  
14c9b12b6e0d'], parent='00000000-0000-0000-0000-000000000000')}, img:  
11f8b3be-fa96-4f6a-bb83-14c9b12b6e0d
```

An unsuccessful wipe will display a log message ***zeroing storage_domain_id/volume_id failed. Zero and remove this volume manually***, and an unsuccessful delete will display ***Remove failed for some of VG: storage_domain_id zeroed volumes: list_of_volume_ids***.

11.4. SHAREABLE DISKS IN RED HAT VIRTUALIZATION

Some applications require storage to be shared between servers. Red Hat Virtualization allows you to mark virtual machine hard disks as **Shareable** and attach those disks to virtual machines. That way a single virtual disk can be used by multiple cluster-aware guests.

Shared disks are not to be used in every situation. For applications like clustered database servers, and other highly available services, shared disks are appropriate. Attaching a shared disk to multiple guests that are not cluster-aware is likely to cause data corruption because their reads and writes to the disk are not coordinated.


You cannot take a snapshot of a shared disk. Virtual disks that have snapshots taken of them cannot later be marked shareable.

You can mark a disk shareable either when you create it, or by editing the disk later.

11.5. READ ONLY DISKS IN RED HAT VIRTUALIZATION

Some applications require administrators to share data with read-only rights. You can do this when creating or editing a disk attached to a virtual machine via the **Disks** tab in the details pane of the virtual machine and selecting the **Read Only** check box. That way, a single disk can be read by multiple cluster-aware guests, while an administrator maintains writing privileges.

You cannot change the read-only status of a disk while the virtual machine is running.




IMPORTANT

Mounting a journaled file system requires read-write access. Using the **Read Only** option is not appropriate for virtual disks that contain such file systems (e.g. **EXT3**, **EXT4**, or **XFS**).

11.6. VIRTUAL DISK TASKS

11.6.1. Creating Floating Virtual Disks



IMPORTANT

Creating floating virtual disks is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend to use them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information on Red Hat Technology Preview features support scope, see <https://access.redhat.com/support/offerings/techpreview/>.

You can create a virtual disk that does not belong to any virtual machines. You can then attach this disk to a single virtual machine, or to multiple virtual machines if the disk is shareable.

Image disk creation is managed entirely by the Manager. **Direct LUN** disks require externally prepared targets that already exist. **Cinder** disks require access to an instance of OpenStack Volume that has been added to the Red Hat Virtualization environment using the **External Providers** window; see [Section 12.2.4, “Adding an OpenStack Volume \(Cinder\) Instance for Storage Management”](#) for more information.

Procedure 11.2. Creating Floating Virtual Disks

1. Select the **Disks** resource tab.
2. Click **New**.

Figure 11.1. Add Virtual Disk Window

- Use the radio buttons to specify whether the virtual disk will be an **Image**, **Direct LUN**, or **Cinder** disk.
- Select the options required for your virtual disk. The options change based on the disk type selected. See [Section 11.6.2, “Explanation of Settings in the New Virtual Disk Window”](#) for more details on each option for each disk type.
- Click **OK**.

11.6.2. Explanation of Settings in the New Virtual Disk Window

Table 11.2. New Virtual Disk Settings: Image

Field Name	Description
Size(GB)	The size of the new virtual disk in GB.
Alias	The name of the virtual disk, limited to 40 characters.
Description	A description of the virtual disk. This field is recommended but not mandatory.

Field Name	Description
Interface	<p>The virtual interface the disk presents to virtual machines. VirtIO is faster, but requires drivers. Red Hat Enterprise Linux 5 and higher include these drivers. Windows does not include these drivers, but they can be installed from the guest tools ISO or virtual floppy disk. IDE devices do not require special drivers.</p> <p>The interface type can be updated after stopping all virtual machines that the disk is attached to.</p>
Data Center	<p>The data center in which the virtual disk will be available.</p>
Storage Domain	<p>The storage domain in which the virtual disk will be stored. The drop-down list shows all storage domains available in the given data center, and also shows the total space and currently available space in the storage domain.</p>
Allocation Policy	<p>The provisioning policy for the new virtual disk.</p> <ul style="list-style-type: none"> • Preallocated allocates the entire size of the disk on the storage domain at the time the virtual disk is created. The virtual size and the actual size of a preallocated disk are the same. Preallocated virtual disks take more time to create than thinly provisioned virtual disks, but have better read and write performance. Preallocated virtual disks are recommended for servers and other I/O intensive virtual machines. If a virtual machine is able to write more than 1 GB every four seconds, use preallocated disks where possible. • Thin Provision allocates 1 GB at the time the virtual disk is created and sets a maximum limit on the size to which the disk can grow. The virtual size of the disk is the maximum limit; the actual size of the disk is the space that has been allocated so far. Thinly provisioned disks are faster to create than preallocated disks and allow for storage over-commitment. Thinly provisioned virtual disks are recommended for desktops.

Field Name	Description
Disk Profile	The disk profile assigned to the virtual disk. Disk profiles define the maximum amount of throughput and the maximum level of input and output operations for a virtual disk in a storage domain. Disk profiles are defined on the storage domain level based on storage quality of service entries created for data centers.
Wipe After Delete	Allows you to enable enhanced security for deletion of sensitive material when the virtual disk is deleted.
Bootable	Allows you to enable the bootable flag on the virtual disk.
Shareable	Allows you to attach the virtual disk to more than one virtual machine at a time.

The **Direct LUN** settings can be displayed in either **Targets > LUNs** or **LUNs > Targets**. **Targets > LUNs** sorts available LUNs according to the host on which they are discovered, whereas **LUNs > Targets** displays a single list of LUNs.

Table 11.3. New Virtual Disk Settings: Direct LUN

Field Name	Description
Alias	The name of the virtual disk, limited to 40 characters.
Description	<p>A description of the virtual disk. This field is recommended but not mandatory. By default the last 4 characters of the LUN ID is inserted into the field.</p> <p>The default behavior can be configured by setting the PopulateDirectLUNDiskDescriptionWithLUNId configuration key to the appropriate value using the engine-config command. The configuration key can be set to -1 for the full LUN ID to be used, or 0 for this feature to be ignored. A positive integer populates the description with the corresponding number of characters of the LUN ID. See Section 18.2.2, "Syntax for the engine-config Command" for more information.</p>

Field Name	Description
Interface	<p>The virtual interface the disk presents to virtual machines. VirtIO is faster, but requires drivers. Red Hat Enterprise Linux 5 and higher include these drivers. Windows does not include these drivers, but they can be installed from the guest tools ISO or virtual floppy disk. IDE devices do not require special drivers.</p> <p>The interface type can be updated after stopping all virtual machines that the disk is attached to.</p>
Data Center	The data center in which the virtual disk will be available.
Use Host	The host on which the LUN will be mounted. You can select any host in the data center.
Storage Type	The type of external LUN to add. You can select from either iSCSI or Fibre Channel .
Discover Targets	<p>This section can be expanded when you are using iSCSI external LUNs and Targets > LUNs is selected.</p> <p>Address - The host name or IP address of the target server.</p> <p>Port - The port by which to attempt a connection to the target server. The default port is 3260.</p> <p>User Authentication - The iSCSI server requires User Authentication. The User Authentication field is visible when you are using iSCSI external LUNs.</p> <p>CHAP username - The user name of a user with permission to log in to LUNs. This field is accessible when the User Authentication check box is selected.</p> <p>CHAP password - The password of a user with permission to log in to LUNs. This field is accessible when the User Authentication check box is selected.</p>
Bootable	Allows you to enable the bootable flag on the virtual disk.
Shareable	Allows you to attach the virtual disk to more than one virtual machine at a time.

Field Name	Description
Enable SCSI Pass-Through	Available when the Interface is set to VirtIO-SCSI . Selecting this check box enables passthrough of a physical SCSI device to the virtual disk. A VirtIO-SCSI interface with SCSI passthrough enabled automatically includes SCSI discard support. When this check box is not selected, the virtual disk uses an emulated SCSI device.
Allow Privileged SCSI I/O	Available when the Enable SCSI Pass-Through check box is selected. Selecting this check box enables unfiltered SCSI Generic I/O (SG_IO) access, allowing privileged SG_IO commands on the disk. This is required for persistent reservations.

Fill in the fields in the **Discover Targets** section and click **Discover** to discover the target server. You can then click the **Login All** button to list the available LUNs on the target server and, using the radio buttons next to each LUN, select the LUN to add.

Using LUNs directly as virtual machine hard disk images removes a layer of abstraction between your virtual machines and their data.

The following considerations must be made when using a direct LUN as a virtual machine hard disk image:

- Live storage migration of direct LUN hard disk images is not supported.
- Direct LUN disks are not included in virtual machine exports.
- Direct LUN disks are not included in virtual machine snapshots.

The **Cinder** settings form will be disabled if there are no available OpenStack Volume storage domains on which you have permissions to create a disk in the relevant Data Center. **Cinder** disks require access to an instance of OpenStack Volume that has been added to the Red Hat Virtualization environment using the **External Providers** window; see [Section 12.2.4, “Adding an OpenStack Volume \(Cinder\) Instance for Storage Management”](#) for more information.

Table 11.4. New Virtual Disk Settings: Cinder

Field Name	Description
Size(GB)	The size of the new virtual disk in GB.
Alias	The name of the virtual disk, limited to 40 characters.
Description	A description of the virtual disk. This field is recommended but not mandatory.

Field Name	Description
Interface	<p>The virtual interface the disk presents to virtual machines. VirtIO is faster, but requires drivers. Red Hat Enterprise Linux 5 and higher include these drivers. Windows does not include these drivers, but they can be installed from the guest tools ISO or virtual floppy disk. IDE devices do not require special drivers.</p> <p>The interface type can be updated after stopping all virtual machines that the disk is attached to.</p>
Data Center	The data center in which the virtual disk will be available.
Storage Domain	The storage domain in which the virtual disk will be stored. The drop-down list shows all storage domains available in the given data center, and also shows the total space and currently available space in the storage domain.
Volume Type	The volume type of the virtual disk. The drop-down list shows all available volume types. The volume type will be managed and configured on OpenStack Cinder.
Bootable	Allows you to enable the bootable flag on the virtual disk.
Shareable	Allows you to attach the virtual disk to more than one virtual machine at a time.

11.6.3. Overview of Live Storage Migration

Virtual disks can be migrated from one storage domain to another while the virtual machine to which they are attached is running. This is referred to as live storage migration. When a disk attached to a running virtual machine is migrated, a snapshot of that disk's image chain is created in the source storage domain, and the entire image chain is replicated in the destination storage domain. As such, ensure that you have sufficient storage space in both the source storage domain and the destination storage domain to host both the disk image chain and the snapshot. A new snapshot is created on each live storage migration attempt, even when the migration fails.

Consider the following when using live storage migration:

- You can live migrate multiple disks at one time.
- Multiple disks for the same virtual machine can reside across more than one storage domain, but the image chain for each disk must reside on a single storage domain.
- You can live migrate disks between any two storage domains in the same data center.

- You cannot live migrate direct LUN hard disk images or disks marked as shareable.

11.6.4. Moving a Virtual Disk

Move a virtual disk that is attached to a virtual machine or acts as a floating virtual disk from one storage domain to another. You can move a virtual disk that is attached to a running virtual machine; this is referred to as live storage migration. Alternatively, shut down the virtual machine before continuing.

Consider the following when moving a disk:

- You can move multiple disks at the same time.
- You can move disks between any two storage domains in the same data center.
- If the virtual disk is attached to a virtual machine that was created based on a template and used the thin provisioning storage allocation option, you must copy the disks for the template on which the virtual machine was based to the same storage domain as the virtual disk.

Procedure 11.3. Moving a Virtual Disk

1. Select the **Disks** tab.
2. Select one or more virtual disks to move.
3. Click **Move** to open the **Move Disk(s)** window.
4. From the **Target** list, select the storage domain to which the virtual disk(s) will be moved.
5. From the **Disk Profile** list, select a profile for the disk(s), if applicable.
6. Click **OK**.

The virtual disks are moved to the target storage domain, and have a status of **Locked** while being moved.

11.6.5. Changing the Disk Interface Type

Users can change a disk's interface type after the disk has been created. This enables you to attach an existing disk to a virtual machine that requires a different interface type. For example, a disk using the **VirtIO** interface can be attached to a virtual machine requiring the **VirtIO-SCSI** or **IDE** interface. This provides flexibility to migrate disks for the purpose of backup and restore, or disaster recovery. The disk interface for shareable disks can also be updated per virtual machine. This means that each virtual machine that uses the shared disk can use a different interface type.

To update a disk interface type, all virtual machines using the disk must first be stopped.

Procedure 11.4. Changing a Disk Interface Type

1. Select the **Virtual Machines** tab and stop the appropriate virtual machine(s).
2. From the **Disks** sub-tab, select the disk and click **Edit**.

3. From the **Interface** list, select the new interface type and click **OK**.

The virtual machine will now use a different virtual interface for the disk.

The following procedure shows how to attach a disk to a different virtual machine that requires a different interface type.

Procedure 11.5. Attaching a Disk to a Different Virtual Machine using a Different Interface Type

1. Select the **Virtual Machines** tab and stop the appropriate virtual machine(s).
2. Select the virtual machine from which to detach the disk.
3. From the **Disks** sub-tab, select the disk and click **Remove**.
4. From the **Virtual Machines** tab, select the new virtual machine that the disk will be attached to.
5. Click **Attach**.
6. Select the disk in the **Attach Virtual Disks** window and select the appropriate interface from the **Interface** drop-down.
7. Click **OK**.

11.6.6. Copying a Virtual Disk

Summary

You can copy a virtual disk from one storage domain to another. The copied disk can be attached to virtual machines.

Procedure 11.6. Copying a Virtual Disk

1. Select the **Disks** tab.
2. Select the virtual disks to copy.
3. Click the **Copy** button to open the **Copy Disk(s)** window.
4. Optionally, enter an alias in the **Alias** text field.
5. Use the **Target** drop-down menus to select the storage domain to which the virtual disk will be copied.
6. Click **OK**.

Result

The virtual disks are copied to the target storage domain, and have a status of **Locked** while being copied.


11.6.7. Uploading a Disk Image to a Storage Domain

QEMU-compatible virtual disk images can be uploaded from your local machine to a Red Hat Virtualization storage domain and attached to virtual machines.

Virtual disk image types must be either QCOW2 or Raw. Disks created from a QCOW2 disk image cannot be shareable, and the QCOW2 disk image file must not have a backing file.

Prerequisites:

- You must configure the Image I/O Proxy when running **engine-setup**. See [Configuring the Red Hat Virtualization Manager](#) in the *Installation Guide* for more information.
- You must import the required certificate authority into the web browser used to access the Administration Portal.
- Internet Explorer 10, Firefox 35, or Chrome 13 or greater is required to perform this upload procedure. Previous browser versions do not support the required HTML5 APIs.

**NOTE**

To import the certificate authority, browse to `https://engine_address/ovirt-engine/services/pki-resource?resource=ca-certificate&format=X509-PEM-CA` and select all the trust settings. Refer to the instructions to install the certificate authority in [Firefox](#), [Internet Explorer](#), or [Google Chrome](#).

Procedure 11.7. Uploading a Disk Image to a Storage Domain

1. Open the Upload Image screen.
 - From the **Disks** tab, select **Start** from the **Upload** drop-down.
 - Alternatively, from the **Storage** tab select the storage domain, then select the **Disks** sub-tab, and select **Start** from the **Upload** drop-down.

Figure 11.2. The Upload Image Screen

2. In the **Upload Image** screen, click **Browse** and select the image on the local disk.
3. Set **Image Type** to **QCOW2** or **Raw**.
4. Fill in the **Disk Option** fields. See [Section 11.6.2, “Explanation of Settings in the New Virtual Disk Window”](#) for a description of the relevant fields.
5. Click **OK**.

A progress bar will indicate the status of the upload. You can also pause, cancel, or resume uploads from the **Upload** drop-down.

11.6.8. Importing a Disk Image from an Imported Storage Domain

Import floating virtual disks from an imported storage domain using the **Disk Import** tab of the details pane.



NOTE

Only QEMU-compatible disks can be imported into the Manager.


Procedure 11.8. Importing a Disk Image

1. Select a storage domain that has been imported into the data center.
2. In the details pane, click **Disk Import**.
3. Select one or more disk images and click **Import** to open the **Import Disk(s)** window.

4. Select the appropriate **Disk Profile** for each disk.
5. Click **OK** to import the selected disks.

11.6.9. Importing an Unregistered Disk Image from an Imported Storage Domain

Import floating virtual disks from a storage domain using the **Disk Import** tab of the details pane. Floating disks created outside of a Red Hat Virtualization environment are not registered with the Manager. Scan the storage domain to identify unregistered floating disks to be imported.



NOTE

Only QEMU-compatible disks can be imported into the Manager.

Procedure 11.9. Importing a Disk Image

1. Select a storage domain that has been imported into the data center.
2. Right-click the storage domain and select **Scan Disks** so that the Manager can identify unregistered disks.
3. In the details pane, click **Disk Import**.
4. Select one or more disk images and click **Import** to open the **Import Disk(s)** window.
5. Select the appropriate **Disk Profile** for each disk.
6. Click **OK** to import the selected disks.

11.6.10. Importing a Virtual Disk Image from an OpenStack Image Service

Summary

Virtual disk images managed by an OpenStack Image Service can be imported into the Red Hat Virtualization Manager if that OpenStack Image Service has been added to the Manager as an external provider.

1. Click the **Storage** resource tab and select the OpenStack Image Service domain from the results list.
2. Select the image to import in the **Images** tab of the details pane.
3. Click **Import** to open the **Import Image(s)** window.
4. From the **Data Center** drop-down menu, select the data center into which the virtual disk image will be imported.
5. From the **Domain Name** drop-down menu, select the storage domain in which the virtual disk image will be stored.
6. Optionally, select a quota from the **Quota** drop-down menu to apply a quota to the virtual disk image.

7. Click **OK** to import the image.

Result

The image is imported as a floating disk and is displayed in the results list of the **Disks** resource tab. It can now be attached to a virtual machine.

11.6.11. Exporting a Virtual Disk to an OpenStack Image Service


Summary

Virtual disks can be exported to an OpenStack Image Service that has been added to the Manager as an external provider.

1. Click the **Disks** resource tab.
2. Select the disks to export.
3. Click the **Export** button to open the **Export Image(s)** window.
4. From the **Domain Name** drop-down list, select the OpenStack Image Service to which the disks will be exported.
5. From the **Quota** drop-down list, select a quota for the disks if a quota is to be applied.
6. Click **OK**.

Result

The virtual disks are exported to the specified OpenStack Image Service where they are managed as virtual disk images.



IMPORTANT

Virtual disks can only be exported if they do not have multiple volumes, are not thinly provisioned, and do not have any snapshots.

11.7. VIRTUAL DISKS AND PERMISSIONS


11.7.1. Managing System Permissions for a Virtual Disk

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

Red Hat Virtualization Manager provides two default virtual disk user roles, but no default virtual disk administrator roles. One of these user roles, the **DiskCreator** role, enables the administration of virtual disks from the User Portal. This role can be applied to specific virtual machines, to a data center, to a specific storage domain, or to the whole virtualized environment; this is useful to allow different users to manage different virtual resources.

The virtual disk creator role permits the following actions:

- Create, edit, and remove virtual disks associated with a virtual machine or other resources.
- Edit user permissions for virtual disks.

**NOTE**

You can only assign roles and permissions to existing users.

11.7.2. Virtual Disk User Roles Explained

Virtual Disk User Permission Roles

The table below describes the user roles and privileges applicable to using and administrating virtual disks in the User Portal.

Table 11.5. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
DiskOperator	Virtual disk user.	Can use, view and edit virtual disks. Inherits permissions to use the virtual machine to which the virtual disk is attached.
DiskCreator	Can create, edit, manage and remove virtual disks within assigned clusters or data centers.	This role is not applied to a specific virtual disk; apply this role to a user for the whole environment with the Configure window. Alternatively apply this role for specific data centers, clusters, or storage domains.

11.7.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 11.10. Assigning a Role to a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Click **Add**.
4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down list.

6. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

11.7.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 11.11. Removing a Role from a Resource

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Permissions** tab in the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

CHAPTER 12. EXTERNAL PROVIDERS

12.1. INTRODUCTION TO EXTERNAL PROVIDERS IN RED HAT VIRTUALIZATION

In addition to resources managed by the Red Hat Virtualization Manager itself, Red Hat Virtualization can also take advantage of resources managed by external sources. The providers of these resources, known as external providers, can provide resources such as virtualization hosts, virtual machine images, and networks.

Red Hat Virtualization currently supports the following external providers:

Red Hat Satellite for Host Provisioning

Satellite is a tool for managing all aspects of the life cycle of both physical and virtual hosts. In Red Hat Virtualization, hosts managed by Satellite can be added to and used by the Red Hat Virtualization Manager as virtualization hosts. After you add a Satellite instance to the Manager, the hosts managed by the Satellite instance can be added by searching for available hosts on that Satellite instance when adding a new host. For more information on installing Red Hat Satellite and managing hosts using Red Hat Satellite, see the [Installation Guide](#) and [Host Configuration Guide](#).

OpenStack Image Service (Glance) for Image Management

OpenStack Image Service provides a catalog of virtual machine images. In Red Hat Virtualization, these images can be imported into the Red Hat Virtualization Manager and used as floating disks or attached to virtual machines and converted into templates. After you add an OpenStack Image Service to the Manager, it appears as a storage domain that is not attached to any data center. Virtual disks in a Red Hat Virtualization environment can also be exported to an OpenStack Image Service as virtual disk images.

OpenStack Networking (Neutron) for Network Provisioning

OpenStack Networking provides software-defined networks. In Red Hat Virtualization, networks provided by OpenStack Networking can be imported into the Red Hat Virtualization Manager and used to carry all types of traffic and create complicated network topologies. After you add OpenStack Networking to the Manager, you can access the networks provided by OpenStack Networking by manually importing them.

OpenStack Volume (Cinder) for Storage Management

OpenStack Volume provides persistent block storage management for virtual hard drives. The OpenStack Cinder volumes are provisioned by Ceph Storage. In Red Hat Virtualization, you can create disks on OpenStack Volume storage that can be used as floating disks or attached to virtual machines. After you add OpenStack Volume to the Manager, you can create a disk on the storage provided by OpenStack Volume.

VMware for Virtual Machine Provisioning

Virtual machines created in VMware can be converted using V2V (virt-v2v) and imported into a Red Hat Virtualization environment. After you add a VMware provider to the Manager, you can import the virtual machines it provides. V2V conversion is performed on a designated proxy host as part of the import operation.

External Network Provider for Network Provisioning

Supported external software-defined network providers include any provider that implements the OpenStack Neutron REST API. Unlike OpenStack Networking (Neutron), the Neutron agent is not used as the virtual interface driver implementation on the host. Instead, the virtual interface driver needs to be provided by the implementer of the external network provider.

All external resource providers are added using a single window that adapts to your input. You must add the resource provider before you can use the resources it provides in your Red Hat Virtualization environment.

12.2. ADDING EXTERNAL PROVIDERS

12.2.1. Adding a Red Hat Satellite Instance for Host Provisioning

Add a Satellite instance for host provisioning to the Red Hat Virtualization Manager. Red Hat Virtualization 4.0 is supported with Red Hat Satellite 6.1.

Procedure 12.1. Adding a Satellite Instance for Host Provisioning

1. Select the **External Providers** entry in the tree pane.
2. Click **Add** to open the **Add Provider** window.


The screenshot shows the 'Add Provider' dialog box with the following details:

- Title:** Add Provider
- Tab:** General
- Name:** Foreman_HP
- Description:** (empty)
- Type:** Foreman/Satellite
- Provider URL:** http://XX.XX.XX.XX
- Requires Authentication:**
- Username:** admin
- Password:** (masked with 10 dots)
- Test:** Test succeeded, managed to access provider.
- Buttons:** OK, Cancel

Figure 12.1. The Add Provider Window

3. Enter a **Name** and **Description**.
4. From the **Type** list, ensure that **Foreman/Satellite** is selected.

5. Enter the URL or fully qualified domain name of the machine on which the Satellite instance is installed in the **Provider URL** text field. You do not need to specify a port number.

**IMPORTANT**

IP addresses cannot be used to add a Satellite instance.

6. Enter the **Username** and **Password** for the Satellite instance. You must use the same user name and password as you would use to log in to the Satellite provisioning portal.
7. Test the credentials:
 - a. Click **Test** to test whether you can authenticate successfully with the Satellite instance using the provided credentials.
 - b. If the Satellite instance uses SSL, the **Import provider certificates** window opens; click **OK** to import the certificate that the Satellite instance provides.

**IMPORTANT**

You must import the certificate that the Satellite instance provides to ensure the Manager can communicate with the instance.

8. Click **OK**.

You have added the Satellite instance to the Red Hat Virtualization Manager, and can work with the hosts it provides.

12.2.2. Adding an OpenStack Image (Glance) Instance for Image Management

Add an OpenStack Image (Glance) instance for image management to the Red Hat Virtualization Manager.

Procedure 12.2. Adding an OpenStack Image (Glance) Instance for Image Management

1. Select the **External Providers** entry in the tree pane.
2. Click **Add** to open the **Add Provider** window.

Figure 12.2. The Add Provider Window

3. Enter a **Name** and **Description**.
4. From the **Type** list, select **OpenStack Image**.
5. Enter the URL or fully qualified domain name of the machine on which the OpenStack Image instance is installed in the **Provider URL** text field.
6. Optionally, select the **Requires Authentication** check box and enter the **Username**, **Password**, **Tenant Name**, and **Authentication URL** for the OpenStack Image instance. You must use the user name and password for the OpenStack Image user registered in Keystone, the tenant of which the OpenStack Image instance is a member, and the URL and port of the Keystone server.
7. Test the credentials:
 - a. Click **Test** to test whether you can authenticate successfully with the OpenStack Image instance using the provided credentials.
 - b. If the OpenStack Image instance uses SSL, the **Import provider certificates** window opens; click **OK** to import the certificate that the OpenStack Image instance provides.



IMPORTANT

You must import the certificate that the OpenStack Image instance provides to ensure the Manager can communicate with the instance.

- Click **OK**.

You have added the OpenStack Image instance to the Red Hat Virtualization Manager, and can work with the images it provides.

12.2.3. Adding an OpenStack Networking (Neutron) Instance for Network Provisioning

Add an OpenStack Networking (Neutron) instance for network provisioning to the Red Hat Virtualization Manager. To add other third-party network providers that implement the OpenStack Neutron REST API, see [Section 12.2.6, “Adding an External Network Provider”](#).



IMPORTANT

Red Hat Virtualization supports Red Hat OpenStack Platform 8, 9, and 10 as external network providers.


Procedure 12.3. Adding an OpenStack Networking (Neutron) Instance for Network Provisioning

- Select the **External Providers** entry in the tree pane.
- Click **Add** to open the **Add Provider** window.

Figure 12.3. The Add Provider Window

- Enter a **Name** and **Description**.
- From the **Type** list, select **OpenStack Networking**.

5. Ensure that **Open vSwitch** is selected in the **Networking Plugin** field.
6. Enter the URL or fully qualified domain name of the machine on which the OpenStack Networking instance is installed in the **Provider URL** text field, followed by the port number. The **Read Only** check box is selected by default. This prevents users from modifying the OpenStack Networking instance.



IMPORTANT

You must leave the **Read Only** check box selected for your setup to be supported by Red Hat.

7. Optionally, select the **Requires Authentication** check box and enter the **Username**, **Password**, **Tenant Name**, and **Authentication URL** for the OpenStack Networking instance. You must use the user name and password for the OpenStack Networking user registered in Keystone, the tenant of which the OpenStack Networking instance is a member, and the URL and port of the Keystone server.
8. Test the credentials:
 - a. Click **Test** to test whether you can authenticate successfully with the OpenStack Networking instance using the provided credentials.
 - b. If the OpenStack Networking instance uses SSL, the **Import provider certificates** window opens; click **OK** to import the certificate that the OpenStack Networking instance provides to ensure the Manager can communicate with the instance.



WARNING

The following steps are provided only as a Technology Preview. Red Hat Virtualization only supports preconfigured Neutron hosts.

9. Click the **Agent Configuration** tab.

Figure 12.4. The Agent Configuration Tab

10. Enter a comma-separated list of interface mappings for the Open vSwitch agent in the **Interface Mappings** field.
11. Select the message broker type that the OpenStack Networking instance uses from the **Broker Type** list.
12. Enter the URL or fully qualified domain name of the host on which the message broker is hosted in the **Host** field.
13. Enter the **Port** by which to connect to the message broker. This port number will be 5762 by default if the message broker is not configured to use SSL, and 5761 if it is configured to use SSL.
14. Enter the **Username** and **Password** of the OpenStack Networking user registered in the message broker instance.
15. Click **OK**.

You have added the OpenStack Networking instance to the Red Hat Virtualization Manager. Before you can use the networks it provides, import the networks into the Manager. See [Section 6.3.1, “Importing Networks From External Providers”](#).

12.2.4. Adding an OpenStack Volume (Cinder) Instance for Storage Management

IMPORTANT

Using an OpenStack Volume (Cinder) instance for storage management is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend to use them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information on Red Hat Technology Preview features support scope, see <https://access.redhat.com/support/offerings/techpreview/>.

Add an OpenStack Volume (Cinder) instance for storage management to the Red Hat Virtualization Manager. The OpenStack Cinder volumes are provisioned by Ceph Storage.

Procedure 12.4. Adding an OpenStack Volume (Cinder) Instance for Storage Management

1. Select the **External Providers** entry in the tree pane.
2. Click **Add** to open the **Add Provider** window.

The screenshot shows the 'Add Provider' dialog box with the following configuration:

- Name:** Cinder_VP
- Description:** (empty)
- Type:** OpenStack Volume
- Data Center:** Default
- Provider URL:** http://XX.XX.XX.XX:8776
- Requires Authentication:**
- Username:** cinder
- Password:** (masked)
- Tenant Name:** services
- Authentication URL:** http://XX.XX.XX.XX:35357/v2.0

A status message at the bottom left reads: ✔ Test succeeded, managed to access provider. Buttons for 'Test', 'OK', and 'Cancel' are located at the bottom right.

Figure 12.5. The Add Provider Window

3. Enter a **Name** and **Description**.
4. From the **Type** list, select **OpenStack Volume**.

5. Select the **Data Center** to which OpenStack Volume storage volumes will be attached.
6. Enter the URL or fully qualified domain name of the machine on which the OpenStack Volume instance is installed, followed by the port number, in the **Provider URL** text field.
7. Optionally, select the **Requires Authentication** check box and enter the **Username, Password, Tenant Name**, and **Authentication URL** for the OpenStack Volume instance. You must use the user name and password for the OpenStack Volume user registered in Keystone, the tenant of which the OpenStack Volume instance is a member, and the URL, port, and API version of the Keystone server.
8. Click **Test** to test whether you can authenticate successfully with the OpenStack Volume instance using the provided credentials.
9. Click **OK**.
10. If client Ceph authentication (**cephx**) is enabled, you must also complete the following steps. The **cephx** protocol is enabled by default.
 - a. On your Ceph server, create a new secret key for the **client.cinder** user using the **ceph auth get-or-create** command. See [Cephx Config Reference](#) for more information on **cephx**, and [Managing Users](#) for more information on creating keys for new users. If a key already exists for the **client.cinder** user, retrieve it using the same command.
 - b. In the Administration Portal, select the newly-created Cinder external provider from the **Providers** list.
 - c. Click the **Authentication Keys** sub-tab.
 - d. Click **New**.
 - e. Enter the secret key in the **Value** field.
 - f. Copy the automatically-generated **UUID**, or enter an existing UUID in the text field.
 - g. On your Cinder server, add the UUID from the previous step and the **cinder** user to **/etc/cinder/cinder.conf**:

```
rbd_secret_uuid = UUID
rbd_user = cinder
```

You have added the OpenStack Volume instance to the Red Hat Virtualization Manager, and can work with the storage volumes it provides. See [Section 11.6.1, “Creating Floating Virtual Disks”](#) for more information about creating a OpenStack Volume (Cinder) disk.

12.2.5. Adding a VMware Instance as a Virtual Machine Provider

Add a VMware vCenter instance to import virtual machines from VMware to the Red Hat Virtualization Manager.

Red Hat Virtualization uses V2V to convert VMware virtual machines to the correct format before they are imported. The `virt-v2v` package must be installed on at least one host. The

virt-v2v package is available by default on Red Hat Virtualization Hosts (RHVH) but must be installed on Red Hat Enterprise Linux hosts. Red Hat Enterprise Linux hosts must be Red Hat Enterprise Linux 7.2 or later.

Procedure 12.5. Adding a VMware vCenter Instance as a Virtual Machine Provider

1. Select the **External Providers** entry in the tree pane.
2. Click **Add** to open the **Add Provider** window.

Add Provider	
General	
Name	VMware_VM
Description	
Type	VMware
Data Center	Default
vCenter	myvcenter.example.com
ESXi	esxi.example.com
Data Center	VMwareDC1
Cluster	VMwareC1
Verify server's SSL certificate	<input checked="" type="checkbox"/>
Proxy Host	Any Host in Data Center
Username	admin
Password
	Test
	OK Cancel

Figure 12.6. The Add Provider Window

3. Enter a **Name** and **Description**.
4. From the **Type** list, select **VMware**.
5. Select the **Data Center** into which VMware virtual machines will be imported, or select **Any Data Center** to instead specify the destination data center during individual import operations (using the **Import** function in the **Virtual Machines** tab).
6. Enter the IP address or fully qualified domain name of the VMware vCenter instance in the **vCenter** field.
7. Enter the IP address or fully qualified domain name of the host from which the virtual machines will be imported in the **ESXi** field.
8. Enter the name of the data center in which the specified ESXi host resides in the **Data Center** field.

9. If you have exchanged the SSL certificate between the ESXi host and the Manager, leave **Verify server's SSL certificate** checked to verify the ESXi host's certificate. If not, uncheck the option.
10. Select a host in the chosen data center with virt-v2v installed to serve as the **Proxy Host** during virtual machine import operations. This host must also be able to connect to the network of the VMware vCenter external provider. If you selected **Any Data Center** above, you cannot choose the host here, but instead can specify a host during individual import operations (using the **Import** function in the **Virtual Machines** tab).
11. Enter the **Username** and **Password** for the VMware vCenter instance. The user must have access to the VMware data center and ESXi host on which the virtual machines reside.
12. Test the credentials:
 - a. Click **Test** to test whether you can authenticate successfully with the VMware vCenter instance using the provided credentials.
 - b. If the VMware vCenter instance uses SSL, the **Import provider certificates** window opens; click **OK** to import the certificate that the VMware vCenter instance provides.



IMPORTANT

You must import the certificate that the VMware vCenter instance provides to ensure the Manager can communicate with the instance.

13. Click **OK**.

You have added the VMware vCenter instance to the Red Hat Virtualization Manager, and can import the virtual machines it provides. See [Importing a Virtual Machine from a VMware Provider](#) in the *Virtual Machine Management Guide* for more information.

12.2.6. Adding an External Network Provider

Any network provider that implements the OpenStack Neutron REST API can be added to Red Hat Virtualization. The virtual interface driver needs to be provided by the implementer of the external network provider. A reference implementation of a network provider and a virtual interface driver are available at <https://github.com/mmirecki/ovirt-provider-mock> and https://github.com/mmirecki/ovirt-provider-mock/blob/master/docs/driver_instalation.

Procedure 12.6. Adding an External Network Provider for Network Provisioning

1. Select the **External Providers** entry in the tree pane.
2. Click **Add**.

The screenshot shows the 'Add Provider' dialog box. The 'General' tab is active. The 'Name' field contains 'Third_Party_Network_Provider'. The 'Description' field is empty. The 'Type' dropdown menu is set to 'External Network Provider'. The 'Provider URL' field contains 'http://XX.XX.XX.XX:9696'. The 'Read Only' checkbox is checked, while 'Requires Authentication' is unchecked. Below these are empty text boxes for 'Username', 'Password', 'Tenant Name', and 'Authentication URL'. A 'Test' button is located at the bottom right of the form area. At the very bottom of the dialog are 'OK' and 'Cancel' buttons.

Figure 12.7. The Add Provider Window

3. Enter a **Name** and **Description**.
4. From the **Type** list, select **External Network Provider**.
5. Enter the URL or fully qualified domain name of the machine on which the external network provider is installed in the **Provider URL** text field, followed by the port number. The **Read-Only** check box is selected by default. This prevents users from modifying the external network provider.



IMPORTANT

You must leave the **Read-Only** check box selected for your setup to be supported by Red Hat.

6. Optionally, select the **Requires Authentication** check box and enter the **Username**, **Password**, **Tenant Name**, and **Authentication URL** for the external network provider.
7. Test the credentials:
 - a. Click **Test** to test whether you can authenticate successfully with the external network provider using the provided credentials.
 - b. If the external network provider uses SSL, the **Import provider certificates** window opens; click **OK** to import the certificate that the external network provider provides to ensure the Manager can communicate with the instance.

You have added an external networking provider to the Red Hat Virtualization Manager.

Before you can use the networks it provides, you need to install the virtual interface driver on the hosts and import the networks. To import networks, see [Section 6.3.1, “Importing Networks From External Providers”](#).

12.2.7. Add Provider General Settings Explained

The **General** tab in the **Add Provider** window allows you to register the core details of the external provider.

Table 12.1. Add Provider: General Settings

Setting	Explanation
Name	A name to represent the provider in the Manager.
Description	A plain text, human-readable description of the provider.
Type	<p>The type of external provider. Changing this setting alters the available fields for configuring the provider.</p> <p>Foreman/Satellite</p> <ul style="list-style-type: none"> • Provider URL: The URL or fully qualified domain name of the machine that hosts the Satellite instance. You do not need to add the port number to the end of the URL or fully qualified domain name. • Requires Authentication: Allows you to specify whether authentication is required for the provider. Authentication is mandatory when Foreman/Satellite is selected. • Username: A user name for connecting to the Satellite instance. This user name must be the user name used to log in to the provisioning portal on the Satellite instance. By default, this user name is admin. • Password: The password against which the above user name is to be authenticated. This password must be the password used to log in to the provisioning portal on the Satellite instance. <p>OpenStack Image</p> <ul style="list-style-type: none"> • Provider URL: The URL or fully qualified domain name of the machine on which the OpenStack Image service is hosted. You must add the port number for the OpenStack Image

Setting	Explanation
	<p>service to the end of the URL or fully qualified domain name. By default, this port number is 9292.</p> <ul style="list-style-type: none"> • Requires Authentication: Allows you to specify whether authentication is required to access the OpenStack Image service. • Username: A user name for connecting to the OpenStack Image service. This user name must be the user name for the OpenStack Image service registered in the Keystone instance of which the OpenStack Image service is a member. By default, this user name is glance. • Password: The password against which the above user name is to be authenticated. This password must be the password for the OpenStack Image service registered in the Keystone instance of which the OpenStack Image service is a member. • Tenant Name: The name of the OpenStack tenant of which the OpenStack Image service is a member. By default, this is services. • Authentication URL: The URL and port of the Keystone server with which the OpenStack Image service authenticates. <p>OpenStack Networking</p> <ul style="list-style-type: none"> • Networking Plugin: The networking plugin with which to connect to the OpenStack Networking server. Open vSwitch is the only option, and is selected by default. • Provider URL: The URL or fully qualified domain name of the machine on which the OpenStack Networking instance is hosted. You must add the port number for the OpenStack Networking instance to the end of the URL or fully qualified domain name. By default, this port number is 9696. • Read Only: Allows you to specify whether the OpenStack Networking instance can be modified from the Administration Portal. • Requires Authentication: Allows you to specify whether authentication is required to access the OpenStack Networking service. • Username: A user name for connecting to the OpenStack

Setting	Explanation
	<p>Networking instance. This user name must be the user name for OpenStack Networking registered in the Keystone instance of which the OpenStack Networking instance is a member. By default, this user name is neutron.</p> <ul style="list-style-type: none"> • Password: The password against which the above user name is to be authenticated. This password must be the password for OpenStack Networking registered in the Keystone instance of which the OpenStack Networking instance is a member. • Tenant Name: The name of the OpenStack tenant of which the OpenStack Networking instance is a member. By default, this is services. • Authentication URL: The URL and port of the Keystone server with which the OpenStack Networking instance authenticates. <p>OpenStack Volume</p> <ul style="list-style-type: none"> • Data Center: The data center to which OpenStack Volume storage volumes will be attached. • Provider URL: The URL or fully qualified domain name of the machine on which the OpenStack Volume instance is hosted. You must add the port number for the OpenStack Volume instance to the end of the URL or fully qualified domain name. By default, this port number is 8776. • Requires Authentication: Allows you to specify whether authentication is required to access the OpenStack Volume service. • Username: A user name for connecting to the OpenStack Volume instance. This user name must be the user name for OpenStack Volume registered in the Keystone instance of which the OpenStack Volume instance is a member. By default, this user name is cinder. • Password: The password against which the above user name is to be authenticated. This password must be the password for OpenStack Volume registered in the Keystone instance of which the OpenStack Volume instance is a member. • Tenant Name: The name of the OpenStack tenant of which the OpenStack Volume instance is a

Setting	Explanation
	<p>member. By default, this is services.</p> <ul style="list-style-type: none"> • Authentication URL: The URL and port of the Keystone server with which the OpenStack Volume instance authenticates. <p>VMware</p> <ul style="list-style-type: none"> • Data Center: Specify the data center into which VMware virtual machines will be imported, or select Any Data Center to instead specify the destination data center during individual import operations (using the Import function in the Virtual Machines tab). • vCenter: The IP address or fully qualified domain name of the VMware vCenter instance. • ESXi: The IP address or fully qualified domain name of the host from which the virtual machines will be imported. • Data Center: The name of the data center in which the specified ESXi host resides. • Cluster: The name of the cluster in which the specified ESXi host resides. • Verify server's SSL certificate: Specify whether the ESXi host's certificate will be verified on connection. • Proxy Host: Select a host in the chosen data center with virt-v2v installed to serve as the host during virtual machine import operations. This host must also be able to connect to the network of the VMware vCenter external provider. If you selected Any Data Center, you cannot choose the host here, but instead can specify a host during individual import operations (using the Import function in the Virtual Machines tab). • Username: A user name for connecting to the VMware vCenter instance. The user must have access to the VMware data center and ESXi host on which the virtual machines reside. • Password: The password against which the above user name is to be authenticated. <p>External Network Provider</p>

Setting	Explanation
	<ul style="list-style-type: none"> • Provider URL: The URL or fully qualified domain name of the machine on which the external network provider is hosted. You must add the port number for the external network provider to the end of the URL or fully qualified domain name. By default, this port number is 9696. • Read Only: Allows you to specify whether the external network provider can be modified from the Administration Portal. • Requires Authentication: Allows you to specify whether authentication is required to access the external network provider. • Username: A user name for connecting to the external network provider. • Password: The password against which the above user name is to be authenticated. • Authentication URL: The URL and port of the authentication server with which the external network provider authenticates.
Test	Allows users to test the specified credentials. This button is available to all provider types.

12.2.8. Add Provider Agent Configuration Settings Explained

The **Agent Configuration** tab in the **Add Provider** window allows users to register details for networking plugins. This tab is only available for the **OpenStack Networking** provider type.

Table 12.2. Add Provider: General Settings

Setting	Explanation
Interface Mappings	A comma-separated list of mappings in the format of <i>label:interface</i> .
Broker Type	The message broker type that the OpenStack Networking instance uses. Select RabbitMQ or Qpid .
Host	The URL or fully qualified domain name of the machine on which the message broker is installed.

Setting	Explanation
Port	The remote port by which a connection with the above host is to be made. By default, this port is 5762 if SSL is not enabled on the host, and 5761 if SSL is enabled.
Username	A user name for authenticating the OpenStack Networking instance with the above message broker. By default, this user name is neutron .
Password	The password against which the above user name is to be authenticated.

12.3. EDITING EXTERNAL PROVIDERS

12.3.1. Editing an External Provider

Procedure 12.7. Editing an External Provider

1. Select the **External Providers** entry in the tree pane.
2. Select the external provider to edit.
3. Click the **Edit** button to open the **Edit Provider** window.
4. Change the current values for the provider to the preferred values.
5. Click **OK**.

12.4. REMOVING EXTERNAL PROVIDERS

12.4.1. Removing an External Provider

Procedure 12.8. Removing an External Provider

1. Select the **External Providers** entry in the tree pane.
2. Select the external provider to remove.
3. Click **Remove**.
4. Click **OK** in the **Remove Provider(s)** window to confirm the removal of this provider.

PART III. ADMINISTERING THE ENVIRONMENT

CHAPTER 13. BACKUPS AND MIGRATION

13.1. BACKING UP AND RESTORING THE RED HAT VIRTUALIZATION MANAGER

13.1.1. Backing up Red Hat Virtualization Manager - Overview

Use the **engine-backup** tool to take regular backups of the Red Hat Virtualization Manager. The tool backs up the engine database and configuration files into a single file and can be run without interrupting the **ovirt-engine** service.

13.1.2. Syntax for the engine-backup Command

The **engine-backup** command works in one of two basic modes:

```
# engine-backup --mode=backup
```

```
# engine-backup --mode=restore
```

These two modes are further extended by a set of parameters that allow you to specify the scope of the backup and different credentials for the engine database. Run **engine-backup --help** for a full list of parameters and their function.

Basic Options

--mode

Specifies whether the command will perform a backup operation or a restore operation. Two options are available - **backup**, and **restore**. This is a required parameter.

--file

Specifies the path and name of a file into which backups are to be taken in backup mode, and the path and name of a file from which to read backup data in restore mode. This is a required parameter in both backup mode and restore mode.

--log

Specifies the path and name of a file into which logs of the backup or restore operation are to be written. This parameter is required in both backup mode and restore mode.

--scope

Specifies the scope of the backup or restore operation. There are four options: **all**, which backs up or restores all databases and configuration data; **files**, which backs up or restores only files on the system; **db**, which backs up or restores only the Manager database; and **dwhdb**, which backs up or restores only the Data Warehouse database. The default scope is **all**.

The **--scope** parameter can be specified multiple times in the same **engine-backup** command.

Manager Database Options

The following options are only available when using the **engine-backup** command in **restore** mode. The option syntax below applies to restoring the Manager database. The same options exist for restoring the Data Warehouse database. See **engine-backup --help** for the Data Warehouse option syntax.

--provision-db


Creates a PostgreSQL database for the Manager database backup to be restored to. This is a required parameter when restoring a backup on a remote host or fresh installation that does not have a PostgreSQL database already configured.

--change-db-credentials

Allows you to specify alternate credentials for restoring the Manager database using credentials other than those stored in the backup itself. See **engine-backup --help** for the additional parameters required by this parameter.

--restore-permissions or --no-restore-permissions

Restores (or does not restore) the permissions of database users. One of these parameters is required when restoring a backup.



NOTE

If a backup contains grants for extra database users, restoring the backup with the **--restore-permissions** and **--provision-db** (or **--provision-dwh-db**) options will create the extra users with random passwords. You must change these passwords manually if the extra users require access to the restored system. See <https://access.redhat.com/articles/2686731>.

13.1.3. Creating a Backup with the engine-backup Command

The Red Hat Virtualization Manager can be backed up using the **engine-backup** command while the Manager is active. Append one of the following options to **--scope** to specify which backup to perform:

- **all**: A full backup of all databases and configuration files on the Manager
- **files**: A backup of only the files on the system
- **db**: A backup of only the Manager database
- **dwhdb**: A backup of only the Data Warehouse database



IMPORTANT

To restore a database to a fresh installation of Red Hat Virtualization Manager, a database backup alone is not sufficient; the Manager also requires access to the configuration files. Any backup that specifies a scope other than the default, **all**, must be accompanied by the **files** scope, or a filesystem backup.

Procedure 13.1. Example Usage of the engine-backup Command

1. Log on to the machine running the Red Hat Virtualization Manager.

2. Create a backup:

Example 13.1. Creating a Full Backup

```
# engine-backup --scope=all --mode=backup --file=file_name --
log=log_file_name
```

Example 13.2. Creating a Manager Database Backup

```
# engine-backup --scope=files --scope=db --mode=backup --
file=file_name --log=log_file_name
```

Replace the **db** option with **dwhdb** to back up the Data Warehouse database.

A **tar** file containing a backup is created using the path and file name provided.

The **tar** files containing the backups can now be used to restore the environment.

13.1.4. Restoring a Backup with the engine-backup Command

Restoring a backup using the `engine-backup` command involves more steps than creating a backup does, depending on the restoration destination. For example, the **engine-backup** command can be used to restore backups to fresh installations of Red Hat Virtualization, on top of existing installations of Red Hat Virtualization, and using local or remote databases.



IMPORTANT

Backups can only be restored to environments of the same major release as that of the backup. For example, a backup of a Red Hat Virtualization version 4.0 environment can only be restored to another Red Hat Virtualization version 4.0 environment. To view the version of Red Hat Virtualization contained in a backup file, unpack the backup file and read the value in the **version** file located in the root directory of the unpacked files.

13.1.5. Restoring a Backup to a Fresh Installation

The **engine-backup** command can be used to restore a backup to a fresh installation of the Red Hat Virtualization Manager. The following procedure must be performed on a machine on which the base operating system has been installed and the required packages for the Red Hat Virtualization Manager have been installed, but the **engine-setup** command has not yet been run. This procedure assumes that the backup file or files can be accessed from the machine on which the backup is to be restored.

Procedure 13.2. Restoring a Backup to a Fresh Installation

1. Log on to the Manager machine. If you are restoring the engine database to a remote host, you will need to log on to and perform the relevant actions on that host. Likewise, if also restoring the Data Warehouse to a remote host, you will need to log on to and perform the relevant actions on that host.
2. Restore a complete backup or a database-only backup.

- Restore a complete backup:

```
# engine-backup --mode=restore --file=file_name --  
log=log_file_name --provision-db --restore-permissions
```

If Data Warehouse is also being restored as part of the complete backup, provision the additional database:

```
engine-backup --mode=restore --file=file_name --log=log_file_name  
--provision-db --provision-dwh-db --restore-permissions
```

- Restore a database-only backup by restoring the configuration files and database backup:

```
# engine-backup --mode=restore --scope=files --scope=db --  
file=file_name --log=log_file_name --provision-db --restore-  
permissions
```

The example above restores a backup of the Manager database.

```
# engine-backup --mode=restore --scope=files --scope=dwhdb --  
file=file_name --log=log_file_name --provision-dwh-db --restore-  
permissions
```

The example above restores a backup of the Data Warehouse database.

If successful, the following output displays:

```
You should now run engine-setup.  
Done.
```

3. Run the following command and follow the prompts to configure the restored Manager:

```
# engine-setup
```

The Red Hat Virtualization Manager has been restored to the version preserved in the backup. To change the fully qualified domain name of the new Red Hat Virtualization system, see [Section 18.1.1, “The oVirt Engine Rename Tool”](#).

13.1.6. Restoring a Backup to Overwrite an Existing Installation

The **engine-backup** command can restore a backup to a machine on which the Red Hat Virtualization Manager has already been installed and set up. This is useful when you have taken a backup up of an installation, performed changes on that installation, and then want to restore the installation from the backup.



IMPORTANT

When restoring a backup to overwrite an existing installation, you must run the **engine-cleanup** command to clean up the existing installation before using the **engine-backup** command. Because the **engine-cleanup** command only cleans the engine database, and does not drop the database or delete the user that owns that database, you do not need to create a new database or specify the database credentials because the user and database already exist.

Procedure 13.3. Restoring a Backup to Overwrite an Existing Installation

1. Log on to the Red Hat Virtualization Manager machine.
2. Remove the configuration files and clean the database associated with the Manager:

```
# engine-cleanup
```

3. Restore a full backup or a database-only backup:

- Restore a full backup:

```
# engine-backup --mode=restore --file=file_name --
log=log_file_name --restore-permissions
```

- Restore a database-only backup by restoring the configuration files and the database backup:

```
# engine-backup --mode=restore --scope=files --scope=db --
file=file_name --log=log_file_name --restore-permissions
```

The example above restores a backup of the Manager database. If necessary, also restore the Data Warehouse database:

```
# engine-backup --mode=restore --scope=dwhdb --file=file_name --
log=log_file_name --restore-permissions
```

If successful, the following output displays:

```
You should now run engine-setup.
Done.
```

4. Run the following command and follow the prompts to reconfigure the firewall and ensure the **ovirt-engine** service is correctly configured:

```
# engine-setup
```

13.1.7. Restoring a Backup with Different Credentials

The **engine-backup** command can restore a backup to a machine on which the Red Hat Virtualization Manager has already been installed and set up, but the credentials of the database in the backup are different to those of the database on the machine on which the

backup is to be restored. This is useful when you have taken a backup of an installation and want to restore the installation from the backup to a different system.



IMPORTANT

When restoring a backup to overwrite an existing installation, you must run the **engine-cleanup** command to clean up the existing installation before using the **engine-backup** command. Because the **engine-cleanup** command only cleans the engine database, and does not drop the database or delete the user that owns that database, you do not need to create a new database or specify the database credentials because the user and database already exist. However, if the credentials for the owner of the engine database are not known, you must change them before you can restore the backup.

Procedure 13.4. Restoring a Backup with Different Credentials

1. Log on to the machine on which the Red Hat Virtualization Manager is installed.
2. Run the following command and follow the prompts to remove the configuration files for and clean the database associated with the Manager:

```
# engine-cleanup
```

3. Change the password for the owner of the engine database if the credentials of that user are not known:

1. Enter the postgresql command line:

```
# su postgres
$ psql
```

2. Change the password of the user that owns the **engine** database:

```
postgres=# alter role user_name encrypted password
'new_password';
```

Repeat this for the user that owns the **ovirt_engine_dwh** database if necessary.

4. Restore a complete backup or a database-only backup with the **--change-db-credentials** parameter to pass the credentials of the new database. The *database_location* for a database local to the Manager is **localhost**.



NOTE

The following examples use a **--*password** option for each database without specifying a password, which will prompt for a password for each database. Passwords can be supplied for these options in the command itself, however this is not recommended as the password will then be stored in the shell history. Alternatively, **--*passfile=password_file** options can be used for each database to securely pass the passwords to the **engine-backup** tool without the need for interactive prompts.

- Restore a complete backup:

```
# engine-backup --mode=restore --file=file_name --
log=log_file_name --change-db-credentials --db-
host=database_location --db-name=database_name --db-user=engine -
-db-password --no-restore-permissions
```

If Data Warehouse is also being restored as part of the complete backup, include the revised credentials for the additional database:

```
engine-backup --mode=restore --file=file_name --log=log_file_name
--change-db-credentials --db-host=database_location --db-
name=database_name --db-user=engine --db-password --change-dwh-
db-credentials --dwh-db-host=database_location --dwh-db-
name=database_name --dwh-db-user=ovirt_engine_history --dwh-db-
password --no-restore-permissions
```

- Restore a database-only backup by restoring the configuration files and the database backup:

```
# engine-backup --mode=restore --scope=files --scope=db --
file=file_name --log=log_file_name --change-db-credentials --db-
host=database_location --db-name=database_name --db-user=engine -
-db-password --no-restore-permissions
```

The example above restores a backup of the Manager database.

```
# engine-backup --mode=restore --scope=files --scope=dwhdb --
file=file_name --log=log_file_name --change-dwh-db-credentials --
dwh-db-host=database_location --dwh-db-name=database_name --dwh-
db-user=ovirt_engine_history --dwh-db-password --no-restore-
permissions
```

The example above restores a backup of the Data Warehouse database.

If successful, the following output displays:

```
You should now run engine-setup.
Done.
```

5. Run the following command and follow the prompts to reconfigure the firewall and ensure the **ovirt-engine** service is correctly configured:

```
# engine-setup
```

13.1.8. Migrating the Engine Database to a Remote Server Database

You can migrate the **engine** database to a remote database server after the Red Hat Virtualization Manager has been initially configured. Use **engine-backup** to create a database backup and restore it on the new database server. This procedure assumes that the new database server has Red Hat Enterprise Linux 7 installed and the appropriate subscriptions configured. See [Subscribing to the Required Entitlements](#) in the *Installation Guide*.

Procedure 13.5. Migrating the Database

1. Log in to the Red Hat Virtualization Manager machine and stop the **ovirt-engine** service so that it does not interfere with the engine backup:

```
# systemctl stop ovirt-engine.service
```

2. Create the **engine** database backup:

```
# engine-backup --scope=files --scope=db --mode=backup --  
file=file_name --log=log_file_name
```

3. Copy the backup file to the new database server:

```
# scp /tmp/engine.dump root@new.database.server.com:/tmp
```

4. Log in to the new database server and install **engine-backup**:

```
# yum install ovirt-engine-tools-backup
```

5. Restore the database on the new database server. *file_name* is the backup file copied from the Manager.

```
# engine-backup --mode=restore --scope=files --scope=db --  
file=file_name --log=log_file_name --provision-db --no-restore-  
permissions
```

6. Now that the database has been migrated, start the **ovirt-engine** service:


```
# systemctl start ovirt-engine.service
```

13.2. BACKING UP AND RESTORING VIRTUAL MACHINES USING THE BACKUP AND RESTORE API

13.2.1. The Backup and Restore API

The backup and restore API is a collection of functions that allows you to perform full or file-level backup and restoration of virtual machines. The API combines several components of Red Hat Virtualization, such as live snapshots and the REST API, to create and work with temporary volumes that can be attached to a virtual machine containing backup software provided by an independent software provider.

For supported third-party backup vendors, consult the Red Hat Virtualization Ecosystem at [Red Hat Marketplace](#).



NOTE

For information on how to work with the REST API, see [The Backup and Restore API](#) in the *REST API Guide*.

13.2.2. Backing Up a Virtual Machine

Use the backup and restore API to back up a virtual machine. This procedure assumes you have two virtual machines: the virtual machine to back up, and a virtual machine on which the software for managing the backup is installed.

Procedure 13.6. Backing Up a Virtual Machine

1. Using the REST API, create a snapshot of the virtual machine to back up:

```
POST /api/vms/11111111-1111-1111-1111-111111111111/snapshots/
HTTP/1.1
Accept: application/xml
Content-type: application/xml

<snapshot>
  <description>BACKUP</description>
</snapshot>
```



NOTE

When you take a snapshot of a virtual machine, a copy of the configuration data of the virtual machine as at the time the snapshot was taken is stored in the **data** attribute of the **configuration** attribute in **initialization** under the snapshot.



IMPORTANT

You cannot take snapshots of disks that are marked as shareable or that are based on direct LUN disks.

2. Retrieve the configuration data of the virtual machine from the **data** attribute under the snapshot:

```
GET /api/vms/11111111-1111-1111-1111-
111111111111/snapshots/11111111-1111-1111-1111-111111111111 HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

3. Identify the disk ID and snapshot ID of the snapshot:

```
GET /api/vms/11111111-1111-1111-1111-
111111111111/snapshots/11111111-1111-1111-1111-111111111111/disks
HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

4. Attach the snapshot to the backup virtual machine and activate the disk:

```
POST /api/vms/22222222-2222-2222-2222-222222222222/disks/ HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

```
<disk id="11111111-1111-1111-1111-111111111111">
  <snapshot id="11111111-1111-1111-1111-111111111111"/>
  <active>true</active>
</disk>
```

5. Use the backup software on the backup virtual machine to back up the data on the snapshot disk.
6. Detach the snapshot disk from the backup virtual machine:

```
DELETE /api/vms/22222222-2222-2222-2222-222222222222/disks/11111111-
1111-1111-1111-111111111111 HTTP/1.1
Accept: application/xml
Content-type: application/xml

<action>
  <detach>true</detach>
</action>
```

7. Optionally, delete the snapshot:

```
DELETE /api/vms/11111111-1111-1111-1111-
111111111111/snapshots/11111111-1111-1111-1111-111111111111 HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

You have backed up the state of a virtual machine at a fixed point in time using backup software installed on a separate virtual machine.

13.2.3. Restoring a Virtual Machine

Restore a virtual machine that has been backed up using the backup and restore API. This procedure assumes you have a backup virtual machine on which the software used to manage the previous backup is installed.

Procedure 13.7. Restoring a Virtual Machine

1. In the Administration Portal, create a floating disk on which to restore the backup. See [Section 11.6.1, "Creating Floating Virtual Disks"](#) for details on how to create a floating disk.
2. Attach the disk to the backup virtual machine:

```
POST /api/vms/22222222-2222-2222-2222-222222222222/disks/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<disk id="11111111-1111-1111-1111-111111111111">
</disk>
```

3. Use the backup software to restore the backup to the disk.
4. Detach the disk from the backup virtual machine:

■


```
DELETE /api/vms/22222222-2222-2222-2222-222222222222/disks/11111111-
1111-1111-1111-111111111111 HTTP/1.1
Accept: application/xml
Content-type: application/xml

<action>
  <detach>true</detach>
</action>
```

5. Create a new virtual machine using the configuration data of the virtual machine being restored:

```
POST /api/vms/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<vm>
  <cluster>
    <name>cluster_name</name>
  </cluster>
  <name>NAME</name>
  ...
</vm>
```

6. Attach the disk to the new virtual machine:

```
POST /api/vms/33333333-3333-3333-3333-333333333333/disks/ HTTP/1.1
Accept: application/xml
Content-type: application/xml


<disk id="11111111-1111-1111-1111-111111111111">
</disk>
```

You have restored a virtual machine using a backup that was created using the backup and restore API.

CHAPTER 14. ERRATA MANAGEMENT WITH RED HAT SATELLITE

Red Hat Virtualization can be configured to view errata from Red Hat Satellite in the Red Hat Virtualization Manager. This enables the administrator to receive updates about available errata, and their importance, for hosts, virtual machines, and the Manager once they have been associated with a Red Hat Satellite provider. Administrators can then choose to apply the updates by running an update on the required host, virtual machine, or on the Manager. For more information about Red Hat Satellite see the [Red Hat Satellite User Guide](#).

Red Hat Virtualization 4.0 supports errata management with Red Hat Satellite 6.1.



IMPORTANT


The Manager, hosts, and virtual machines are identified in the Satellite server by their FQDN. This ensures that external content host IDs do not need to be maintained in Red Hat Virtualization.

The Satellite account used to manage the Manager, hosts and virtual machines must have Administrator permissions and a default organization set.

Procedure 14.1. Configuring Red Hat Virtualization Errata

To associate a Manager, host, and virtual machine with a Red Hat Satellite provider first the Manager must be associated with a provider. Then the host is associated with the same provider and configured. Finally, the virtual machine is associated with the same provider and configured.

1. Associate the Manager by adding the required Satellite server as an external provider. See [Section 12.2.1, “Adding a Red Hat Satellite Instance for Host Provisioning”](#) for more information.



NOTE

The Manager must be registered to the Satellite server as a content host and have the katello-agent package installed.

For more information on how to configure a host registration see [Configuring a Host for Registration](#) in the *Red Hat Satellite User Guide* and for more information on how to register a host and install the katello-agent package see [Registration](#) in the *Red Hat Satellite User Guide*

2. Optionally, configure the required hosts to display available errata. See [Section 7.5.3, “Configuring Satellite Errata Management for a Host”](#) for more information.
3. Optionally, configure the required virtual machines to display available errata. The associated host needs to be configured prior to configuring the required virtual machines. See [Configuring Red Hat Satellite Errata Management for a Virtual Machine](#) in the *Virtual Machine Management Guide* for more information.

Procedure 14.2. Viewing Red Hat Virtualization Manager Errata

1. Select the **Errata** entry in the tree pane.
2. Click the **Security**, **Bugs**, or **Enhancements** checkboxes to view only those errata types.

For more information on viewing available errata for hosts see [Section 7.5.13, “Viewing Host Errata”](#) and for virtual machines see [Viewing Red Hat Satellite Errata for a Virtual Machine](#) in the *Virtual Machine Management Guide*.

CHAPTER 15. USERS AND ROLES

15.1. INTRODUCTION TO USERS

In Red Hat Virtualization, there are two types of user domains: local domain and external domain. A default local domain called the **internal** domain and a default user **admin** is created during the the Manager installation process.

You can create additional users on the **internal** domain using **ovirt-aaa-jdbc-tool**. User accounts created on local domains are known as local users. You can also attach external directory servers such as Red Hat Directory Server, Active Directory, OpenLDAP, and many other supported options to your Red Hat Virtualization environment and use them as external domains. User accounts created on external domains are known as directory users.

Both local users and directory users need to be assigned with appropriate roles and permissions through the Administration Portal before they can function in the environment. There are two main types of user roles: end user and administrator. An end user role uses and manages virtual resources from the User Portal. An administrator role maintains the system infrastructure using the Administration Portal. The roles can be assigned to the users for individual resources like virtual machines and hosts, or on a hierarchy of objects like clusters and data centers.

15.2. INTRODUCTION TO DIRECTORY SERVERS

During installation, Red Hat Virtualization Manager creates an **admin** user on the **internal** domain. The user is also referred to as **admin@internal**. This account is intended for use when initially configuring the environment and for troubleshooting. After you have attached an external directory server, added the directory users, and assigned them with appropriate roles and permissions, the **admin@internal** user can be disabled if it is not required. The directory servers supported are:

- 389ds
- 389ds RFC-2307 Schema
- Active Directory
- FreeIPA
- Red Hat Identity Management (IdM)
- Novell eDirectory RFC-2307 Schema
- OpenLDAP RFC-2307 Schema
- OpenLDAP Standard Schema
- Oracle Unified Directory RFC-2307 Schema
- RFC-2307 Schema (Generic)
- Red Hat Directory Server (RHDS)
- Red Hat Directory Server (RHDS) RFC-2307 Schema

- iPlanet



IMPORTANT

It is not possible to install Red Hat Virtualization Manager (rhev) and IdM (ipa-server) on the same system. IdM is incompatible with `themod_ssl` package, which is required by Red Hat Virtualization Manager.



IMPORTANT

If you are using Active Directory as your directory server, and you want to use **sysprep** in the creation of templates and virtual machines, then the Red Hat Virtualization administrative user must be delegated control over the Domain to:

- **Join a computer to the domain**
- **Modify the membership of a group**

For information on creation of user accounts in Active Directory, see <http://technet.microsoft.com/en-us/library/cc732336.aspx>.

For information on delegation of control in Active Directory, see <http://technet.microsoft.com/en-us/library/cc732524.aspx>.

15.3. CONFIGURING AN EXTERNAL LDAP PROVIDER

15.3.1. Configuring an External LDAP Provider (Interactive Setup)

The `ovirt-engine-extension-aaa-ldap` extension allows users to customize their external directory setup easily. The `ovirt-engine-extension-aaa-ldap` extension supports many different LDAP server types, and an interactive setup script is provided to assist you with the setup for most LDAP types.

If the LDAP server type is not listed in the interactive setup script, or you want to do more customization, you can manually edit the configuration files. See [Section 15.3.3, “Configuring an External LDAP Provider \(Manual Method\)”](#) for more information.

For an Active Directory example, see [Section 15.3.2, “Attaching an Active Directory”](#).

Prerequisites:

- You need to know the domain name of the DNS or the LDAP server. Round-robin and failover policies are also supported.
- To set up secure connection between the LDAP server and the Manager, ensure a PEM-encoded CA certificate has been prepared. See [Section D.2, “Setting Up SSL or TLS Connections between the Manager and an LDAP Server”](#) for more information.
- Unless anonymous search is supported, a user with permissions to browse all users and groups must be available on the directory server to be used as the search user. Note down the search user's distinguished name (DN). Do not use the administrative user for the directory server.

- Have at least one set of account name and password ready to perform search and login queries to the LDAP server.

Procedure 15.1. Configuring an External LDAP Provider


1. On the Red Hat Virtualization Manager, install the LDAP extension package:

```
# yum install ovirt-engine-extension-aaa-ldap-setup
```

2. Run **ovirt-engine-extension-aaa-ldap-setup** to start the interactive setup:

```
# ovirt-engine-extension-aaa-ldap-setup
```

3. Specify a profile name. The profile name is visible to users on the login page. This example uses **redhat.com**.



NOTE

To rename the profile after the domain has been configured, edit the **ovirt.engine.aaa.authn.profile.name** attribute in the **/etc/ovirt-engine/extensions.d/redhat.com-authn.properties** file. Restart the engine service for the changes to take effect.

Please specify profile name that will be visible to users:*redhat.com*

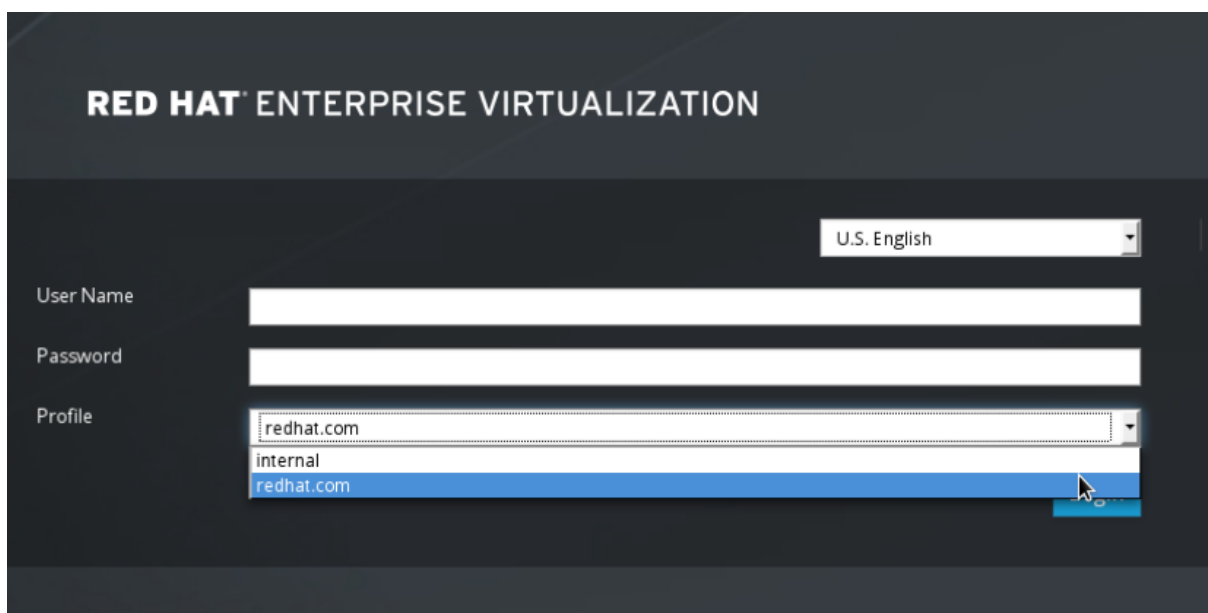


Figure 15.1. The Administration Portal Login Page



NOTE

Users need to select the desired profile from the drop-down list when logging in for the first time. The information is then stored in browser cookies and preselected the next time the user logs in.

4. Select an LDAP type by entering the corresponding number. If you are not sure

which schema your LDAP server is, select the standard schema of your LDAP server type. For Active Directory, follow the procedure at [Section 15.3.2, “Attaching an Active Directory”](#).

```
Available LDAP implementations:
1 - 389ds
2 - 389ds RFC-2307 Schema
3 - Active Directory
4 - IPA
5 - Novell eDirectory RFC-2307 Schema
6 - OpenLDAP RFC-2307 Schema
7 - OpenLDAP Standard Schema
8 - Oracle Unified Directory RFC-2307 Schema
9 - RFC-2307 Schema (Generic)
10 - RHDS
11 - RHDS RFC-2307 Schema
12 - iPlanet
Please select: 10
```

5. Press **Enter** to accept the default and configure domain name resolution for your LDAP server name:

```
It is highly recommended to use DNS resolution for LDAP server.
If for some reason you intend to use hosts or plain address disable
DNS usage.
Use DNS (Yes, No) [Yes]:
```

6. Select a DNS policy method by entering the corresponding number:

```
1 - Single server
2 - DNS domain LDAP SRV record
3 - Round-robin between multiple hosts
4 - Failover between multiple hosts
Please select:
```

- For option 1, the DNS servers listed in `/etc/resolv.conf` is used to resolve the IP address. Ensure the `/etc/resolv.conf` file is updated with the correct DNS servers.

Enter the fully qualified domain name (FQDN) or the IP address of the LDAP server. You can use the **dig** command with the SRV record to find out the domain name. An SRV record takes the following format:
`_service._protocol.domain name`. For example: **dig _ldap._tcp.redhat.com SRV**.

- For option 2, enter the domain name of the DNS server. A DNS search is performed to look up the SRV records to find out the domain name of the LDAP server.
- For option 3, enter a space-separated list of LDAP servers. Use either the FQDN or IP address of the servers. This policy provides load-balancing between the LDAP servers. Queries are distributed among all LDAP servers according to the round-robin algorithm.
- For option 4, enter a space-separated list of LDAP servers. Use either the FQDN

or IP address of the servers. This policy defines the first LDAP server to be the default LDAP server to respond to queries. If the first server is not available, the query will go to the next LDAP server on the list.

7. Select the secure connection method your LDAP server supports and specify the method to obtain a PEM-encoded CA certificate. The file option allows you to provide the full path to the certificate. The URL option allows you to specify a URL to the certificate. Use the inline option to paste the content of the certificate in the terminal. The system option allows you to specify the default location for all CA files. If the insecure mode is selected, the connection is still encrypted using TLS but certificate validation is skipped.

NOTE:

It is highly recommended to use secure protocol to access the LDAP server.

Protocol `startTLS` is the standard recommended method to do so.


Only in cases in which the `startTLS` is not supported, fallback to non standard `ldaps` protocol.

Use `plain` for test environments only.

Please select protocol to use (`startTLS`, `ldaps`, `plain`) [`startTLS`]:
`startTLS`

Please select method to obtain PEM encoded CA certificate (File, URL, Inline, System, Insecure): `File`

Please enter the password:



NOTE

LDAPS stands for Lightweight Directory Access Protocol Over Secure Socket Links. For SSL connections, select the **ldaps** option.

For more information on creating a PEM-encoded CA certificate, see [Section D.2, “Setting Up SSL or TLS Connections between the Manager and an LDAP Server”](#).

8. Enter the search user distinguished name (DN). The user must have permissions to browse all users and groups on the directory server. The search user must be specified in LDAP annotation. If anonymous search is allowed, press **Enter** without any input.

Enter search user DN (empty for anonymous):

`uid=user1,ou=Users,dc=test,dc=redhat,dc=com`

Enter search user password:

9. Test the search and login function to ensure the your LDAP server is connected to your Red Hat Virtualization environment properly. For the login query, enter the account name and password. For the search query, select **Principal** for user accounts, and select **Group** for group accounts. Enter **Yes** to **Resolve Groups** if you want the group account information for the user account to be returned. Select **Done** to complete the setup. Three configuration files are created and displayed in the screen output.

NOTE:

It is highly recommended to test drive the configuration before applying it into engine.


```

Perform at least one Login sequence and one Search sequence.
Select test sequence to execute (Done, Abort, Login, Search)
[Abort]: Login
Enter search user name: testuser1
Enter search user password:
[ INFO ] Executing login sequence...
...
Select test sequence to execute (Done, Abort, Login, Search)
[Abort]: Search
Select entity to search (Principal, Group) [Principal]:
Term to search, trailing '*' is allowed: testuser1
Resolve Groups (Yes, No) [No]:
[ INFO ] Executing login sequence...
...
Select test sequence to execute (Done, Abort, Login, Search)
[Abort]: Done
[ INFO ] Stage: Transaction setup
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Package installation
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Transaction commit
[ INFO ] Stage: Closing up
CONFIGURATION SUMMARY
Profile name is: redhat.com
The following files were created:
    /etc/ovirt-engine/aaa/redhat.com.properties
    /etc/ovirt-engine/extensions.d/redhat.com-
authz.properties
    /etc/ovirt-engine/extensions.d/redhat.com-
authn.properties
[ INFO ] Stage: Clean up
    Log file is available at /tmp/ovirt-engine-extension-aaa-
ldap-setup-20160114064955-1yar9i.log:
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination

```

- Restart the engine service. The profile you have created is now available on the Administration Portal and the User Portal login pages. To assign the user accounts on the LDAP server appropriate roles and permissions, for example to log in to the User Portal, see [Section 15.6, “Administering User Tasks From the Administration Portal”](#).

```
# systemctl restart ovirt-engine.service
```



NOTE

For more information, see the LDAP authentication and authorization extension README file at `/usr/share/doc/ovirt-engine-extension-aaa-ldap-version`.

15.3.2. Attaching an Active Directory

Prerequisites:

- You need to know the Active Directory forest name. The forest name is also known as the root domain name.
- You need to either add the DNS server that can resolve the Active Directory forest name to the `/etc/resolv.conf` file on the Manager, or note down the Active Directory DNS servers and enter them when prompted by the interactive setup script.
- To set up secure connection between the LDAP server and the Manager, ensure a PEM-encoded CA certificate has been prepared. See [Section D.2, “Setting Up SSL or TLS Connections between the Manager and an LDAP Server”](#) for more information.
- Unless anonymous search is supported, a user with permissions to browse all users and groups must be available on the Active Directory to be used as the search user. Note down the search user's distinguished name (DN). Do not use the administrative user for the Active Directory.
- Have at least one set of account name and password ready to perform search and login queries to the Active Directory.

Procedure 15.2. Configuring an External LDAP Provider

1. On the Red Hat Virtualization Manager, install the LDAP extension package:

```
# yum install ovirt-engine-extension-aaa-ldap-setup
```

2. Run `ovirt-engine-extension-aaa-ldap-setup` to start the interactive setup:

```
# ovirt-engine-extension-aaa-ldap-setup
```

3. Specify a profile name. The profile name is visible to users on the login page. This example uses `redhat.com`.

```
Please specify profile name that will be visible to users:redhat.com
```

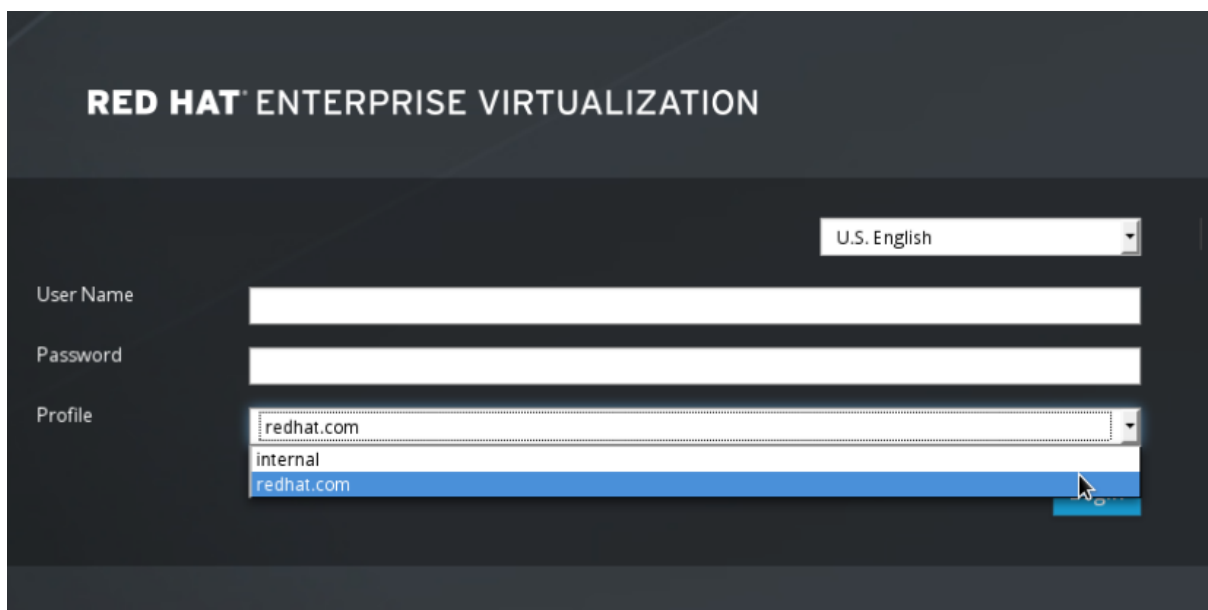



Figure 15.2. The Administration Portal Login Page

**NOTE**

Users need to select the desired profile from the drop-down list when logging in for the first time. The information is then stored in browser cookies and preselected the next time the user logs in.

- Select an LDAP type by entering the corresponding number. The LDAP related questions after this step is different for different LDAP types.


```
Available LDAP implementations:
1 - 389ds
2 - 389ds RFC-2307 Schema
3 - Active Directory
4 - IPA
5 - Novell eDirectory RFC-2307 Schema
6 - OpenLDAP RFC-2307 Schema
7 - OpenLDAP Standard Schema
8 - Oracle Unified Directory RFC-2307 Schema
9 - RFC-2307 Schema (Generic)
10 - RHDS
11 - RHDS RFC-2307 Schema
12 - iPlanet
Please select: 3
```

- Enter the Active Directory forest name. If the forest name is not resolvable by your Manager's DNS, the script prompts you to enter a space-separated list of Active Directory DNS server names.

```
Please enter Active Directory Forest name: ad-example.redhat.com
[ INFO ] Resolving Global Catalog SRV record for ad-example.redhat.com
[ INFO ] Resolving LDAP SRV record for ad-example.redhat.com
```

- Select the secure connection method your LDAP server supports and specify the method to obtain a PEM-encoded CA certificate. The file option allows you to provide the full path to the certificate. The URL option allows you to specify a URL to the certificate. Use the inline option to paste the content of the certificate in the terminal. The system option allows you to specify the location for all CA files. The insecure option allows you to use startTLS in insecure mode.

```
NOTE:
It is highly recommended to use secure protocol to access the LDAP
server.
Protocol startTLS is the standard recommended method to do so.
Only in cases in which the startTLS is not supported, fallback to
non standard ldaps protocol.
Use plain for test environments only.
Please select protocol to use (startTLS, ldaps, plain) [startTLS]:
startTLS
Please select method to obtain PEM encoded CA certificate (File,
URL, Inline, System, Insecure): File
Please enter the password:
```

**NOTE**

LDAPS stands for Lightweight Directory Access Protocol Over Secure Socket Links. For SSL connections, select the **ldaps** option.

For more information on creating a PEM-encoded CA certificate, see [Section D.2, “Setting Up SSL or TLS Connections between the Manager and an LDAP Server”](#).

7. Enter the search user distinguished name (DN). The user must have permissions to browse all users and groups on the directory server. The search user must be of LDAP annotation. If anonymous search is allowed, press **Enter** without any input.

```
Enter search user DN (empty for anonymous):
uid=user1,ou=Users,dc=test,dc=redhat,dc=com
Enter search user password:
```

8. Test the search and login function to ensure your LDAP server is connected to your Red Hat Virtualization environment properly. For the login query, enter the account name and password. For the search query, select **Principal** for user accounts, and select **Group** for group accounts. Enter **Yes** to **Resolve Groups** if you want the group account information for the user account to be returned. Select **Done** to complete the setup. Three configuration files are created and displayed in the screen output.

NOTE:

It is highly recommended to test drive the configuration before applying it into engine.

Perform at least one Login sequence and one Search sequence.

Select test sequence to execute (Done, Abort, Login, Search)

[Abort]: Login

Enter search user name: *testuser1*

Enter search user password:

[INFO] Executing login sequence...

...

Select test sequence to execute (Done, Abort, Login, Search)

[Abort]: Search

Select entity to search (Principal, Group) [Principal]:

Term to search, trailing '*' is allowed: *testuser1*

Resolve Groups (Yes, No) [No]:

[INFO] Executing login sequence...

...

Select test sequence to execute (Done, Abort, Login, Search)

[Abort]: Done

[INFO] Stage: Transaction setup

[INFO] Stage: Misc configuration

[INFO] Stage: Package installation

[INFO] Stage: Misc configuration

[INFO] Stage: Transaction commit

[INFO] Stage: Closing up

CONFIGURATION SUMMARY

Profile name is: *redhat.com*

The following files were created:

/etc/ovirt-engine/aaa/redhat.com.properties

/etc/ovirt-engine/extensions.d/redhat.com-

authz.properties

```

        /etc/ovirt-engine/extensions.d/redhat.com-
authn.properties
[ INFO ] Stage: Clean up
        Log file is available at /tmp/ovirt-engine-extension-aaa-
ldap-setup-20160114064955-1yar9i.log:
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination

```

- The profile you have created is now available on the Administration Portal and the User Portal login pages. To assign the user accounts on the LDAP server appropriate roles and permissions, for example to log in to the User Portal, see [Section 15.6, “Administering User Tasks From the Administration Portal”](#).



NOTE

For more information, see the LDAP authentication and authorization extension README file at `/usr/share/doc/ovirt-engine-extension-aaa-ldap-version`.

15.3.3. Configuring an External LDAP Provider (Manual Method)

The `ovirt-engine-extension-aaa-ldap` extension uses the LDAP protocol to access directory servers and is fully customizable. Kerberos authentication is not required unless you want to enable the single sign-on to the User Portal or the Administration Portal feature.

If the interactive setup method in the previous section does not cover your use case, you can manually modify the configuration files to attach your LDAP server. The following procedure uses generic details. Specific values depend on your setup.

Procedure 15.3. Configuring an External LDAP Provider Manually

- On the Red Hat Virtualization Manager, install the LDAP extension package:

```
# yum install ovirt-engine-extension-aaa-ldap
```

- Copy the LDAP configuration template file into the `/etc/ovirt-engine` directory. Template files are available for active directories (**ad**) and other directory types (**simple**). This example uses the simple configuration template.

```
# cp -r /usr/share/ovirt-engine-extension-aaa-ldap/examples/simple/.
/etc/ovirt-engine
```

- Rename the configuration files to match the profile name you want visible to users on the Administration Portal and the User Portal login pages:

```
# mv /etc/ovirt-engine/aaa/profile1.properties /etc/ovirt-
engine/aaa/example.properties
# mv /etc/ovirt-engine/extensions.d/profile1-authn.properties
/etc/ovirt-engine/extensions.d/example-authn.properties
# mv /etc/ovirt-engine/extensions.d/profile1-authz.properties
/etc/ovirt-engine/extensions.d/example-authz.properties
```

4. Edit the LDAP property configuration file by uncommenting an LDAP server type and updating the domain and passwords fields:

```
# vi /etc/ovirt-engine/aaa/example.properties
```

Example 15.1. Example profile: LDAP server section


```
# Select one
#
include = <openldap.properties>
#include = <389ds.properties>
#include = <rhds.properties>
#include = <ipa.properties>
#include = <iplanet.properties>
#include = <rfc2307-389ds.properties>
#include = <rfc2307-rhds.properties>
#include = <rfc2307-openldap.properties>
#include = <rfc2307-edir.properties>
#include = <rfc2307-generic.properties>

# Server
#
vars.server = ldap1.company.com

# Search user and its password.
#
vars.user = uid=search,cn=users,cn=accounts,dc=company,dc=com
vars.password = 123456

pool.default.serverset.single.server = ${global:vars.server}
pool.default.auth.simple.bindDN = ${global:vars.user}
pool.default.auth.simple.password = ${global:vars.password}
```

To use TLS or SSL protocol to interact with the LDAP server, obtain the root CA certificate for the LDAP server and use it to create a public keystore file. Uncomment the following lines and specify the full path to the public keystore file and the password to access the file.



NOTE

For more information on creating a public keystore file, see [Section D.2, “Setting Up SSL or TLS Connections between the Manager and an LDAP Server”](#).

Example 15.2. Example profile: keystore section

```
# Create keystore, import certificate chain and uncomment
# if using tls.
pool.default.ssl.startTLS = true
pool.default.ssl.truststore.file = /full/path/to/myrootca.jks
pool.default.ssl.truststore.password = password
```

- Review the authentication configuration file. The profile name visible to users on the Administration Portal and the User Portal login pages is defined by **ovirt.engine.aaa.authn.profile.name**. The configuration profile location must match the LDAP configuration file location. All fields can be left as default.

```
# vi /etc/ovirt-engine/extensions.d/example-authn.properties
```

Example 15.3. Example authentication configuration file

```
ovirt.engine.extension.name = example-authn
ovirt.engine.extension.bindings.method = jbossmodule
ovirt.engine.extension.binding.jbossmodule.module =
org.ovirt.engine-extensions.aaa.ldap
ovirt.engine.extension.binding.jbossmodule.class =
org.ovirt.engineextensions.aaa.ldap.AuthnExtension
ovirt.engine.extension.provides =
org.ovirt.engine.api.extensions.aaa.Authn
ovirt.engine.aaa.authn.profile.name = example
ovirt.engine.aaa.authn.authz.plugin = example-authz
config.profile.file.1 = ../aaa/example.properties
```

- Review the authorization configuration file. The configuration profile location must match the LDAP configuration file location. All fields can be left as default.

```
# vi /etc/ovirt-engine/extensions.d/example-authz.properties
```

Example 15.4. Example authorization configuration file

```
ovirt.engine.extension.name = example-authz
ovirt.engine.extension.bindings.method = jbossmodule
ovirt.engine.extension.binding.jbossmodule.module =
org.ovirt.engine-extensions.aaa.ldap
ovirt.engine.extension.binding.jbossmodule.class =
org.ovirt.engineextensions.aaa.ldap.AuthzExtension
ovirt.engine.extension.provides =
org.ovirt.engine.api.extensions.aaa.Authz
config.profile.file.1 = ../aaa/example.properties
```

- Ensure that the ownership and permissions of the configuration profile are appropriate:

```
# chown ovirt:ovirt /etc/ovirt-engine/aaa/example.properties
```

```
# chmod 600 /etc/ovirt-engine/aaa/example.properties
```

- Restart the engine service:

```
# systemctl restart ovirt-engine.service
```

- The *example* profile you have created is now available on the Administration Portal

and the User Portal login pages. To give the user accounts on the LDAP server appropriate permissions, for example to log in to the User Portal, see [Section 15.6, “Administering User Tasks From the Administration Portal”](#).

**NOTE**

For more information, see the LDAP authentication and authorization extension README file at `/usr/share/doc/ovirt-engine-extension-aaa-ldap-version`.

15.3.4. Removing an External LDAP Provider

This procedure shows you how to remove an external configured LDAP provider and its users.

Procedure 15.4. Removing an External LDAP Provider

1. Remove the LDAP provider configuration files, replacing the default name **profile1**:

```
# rm /etc/ovirt-engine/extensions.d/profile1-authn.properties
# rm /etc/ovirt-engine/extensions.d/profile1-authz.properties
# rm /etc/ovirt-engine/aaa/profile1.properties
```

2. Restart the **ovirt-engine** service:

```
# systemctl restart ovirt-engine
```

3. In the Administration Portal, in the **Users** resource tab, select the users of this provider (those whose **Authorization provider** is **profile1-authz**) and click **Remove**.

15.4. CONFIGURING LDAP AND KERBEROS FOR SINGLE SIGN-ON

Single sign-on allows users to log in to the User Portal or the Administration Portal without re-typing their passwords. Authentication credentials are obtained from the Kerberos server. To configure single sign-on to the Administration Portal and the User Portal, you need to configure two extensions: **ovirt-engine-extension-aaa-misc** and **ovirt-engine-extension-aaa-ldap**; and two Apache modules: **mod_auth_gssapi** and **mod_session**. You can configure single sign-on that does not involve Kerberos, however this is outside the scope of this documentation.

**NOTE**

If single sign-on to the User Portal is enabled, single sign-on to virtual machines will not be possible. With single sign-on to the User Portal enabled, the User Portal does not need to accept a password, thus the password cannot be delegated to sign in to virtual machines.

This example assumes the following:

- The existing Key Distribution Center (KDC) server uses the MIT version of Kerberos 5.
- You have administrative rights to the KDC server.
- The Kerberos client is installed on the Red Hat Virtualization Manager and user machines.
- The **kadmin** utility is used to create Kerberos service principals and **keytab** files.

This procedure involves the following components:

On the KDC server

- Create a service principal and a **keytab** file for the Apache service on the Red Hat Virtualization Manager.

On the Red Hat Virtualization Manager

- Install the authentication and authorization extension packages and the Apache Kerberos authentication module.
- Configure the extension files.

Procedure 15.5. Configuring Kerberos for the Apache Service

1. On the KDC server, use the **kadmin** utility to create a service principal for the Apache service on the Red Hat Virtualization Manager. The service principal is a reference ID to the KDC for the Apache service.

```
# kadmin
kadmin> addprinc -randkey HTTP/fqdn-of-rhevm@REALM.COM
```

2. Generate a **keytab** file for the Apache service. The **keytab** file stores the shared secret key.

```
kadmin> ktadd -k /tmp/http.keytab HTTP/fqdn-of-rhevm@REALM.COM
```

```
kadmin> quit
```

3. Copy the **keytab** file from the KDC server to the Red Hat Virtualization Manager:

```
# scp /tmp/http.keytab root@rhevm.example.com:/etc/httpd
```

Procedure 15.6. Configuring Single Sign-on to the User Portal or Administration Portal

1. On the Red Hat Virtualization Manager, ensure that the ownership and permissions for the keytab are appropriate:

```
# chown apache /etc/httpd/http.keytab
```

```
# chmod 400 /etc/httpd/http.keytab
```

- 2. Install the authentication extension package, LDAP extension package, and the `mod_auth_gssapi` and `mod_session` Apache modules:

```
# yum install ovirt-engine-extension-aaa-misc ovirt-engine-
extension-aaa-ldap mod_auth_gssapi mod_session
```

- 3. Copy the SSO configuration template file into the `/etc/ovirt-engine` directory. Template files are available for Active Directory (`ad-sso`) and other directory types (`simple-sso`). This example uses the simple SSO configuration template.

```
# cp -r /usr/share/ovirt-engine-extension-aaa-ldap/examples/simple-
sso/. /etc/ovirt-engine
```

- 4. Move `ovirt-sso.conf` into the Apache configuration directory:

```
# mv /etc/ovirt-engine/aaa/ovirt-sso.conf /etc/httpd/conf.d
```

- 5. Review the authentication method file. You do not need to edit this file, as the realm is automatically fetched from the `keytab` file.

```
# vi /etc/httpd/conf.d/ovirt-sso.conf
```

Example 15.5. Example authentication method file

```
<LocationMatch ^/ovirt-engine/sso/(interactive-login-
negotiate|oauth/token-http-auth)|^/ovirt-engine/api>
  <If "req('Authorization') !~ /^^(Bearer|Basic)/i">
    RewriteEngine on
    RewriteCond %{LA-U:REMOTE_USER} ^(.*)$
    RewriteRule ^(.*)$ - [L,NS,P,E=REMOTE_USER:%1]
    RequestHeader set X-Remote-User %{REMOTE_USER}s

    AuthType GSSAPI
    AuthName "Kerberos Login"

    # Modify to match installation
    GssapiCredStore keytab:/etc/httpd/http.keytab
    GssapiUseSessions On
    Session On
    SessionCookieName ovirt_gssapi_session
    path=/private;httponly;secure;

    Require valid-user
    ErrorDocument 401 "<html><meta http-equiv=\"refresh\"
content=\"0; url=/ovirt-engine/sso/login-unauthorized\"/><body><a
href=\"/ovirt-engine/sso/login-unauthorized\">Here</a></body>
</html>"
  </If>
</LocationMatch>
```

- Rename the configuration files to match the profile name you want visible to users on the Administration Portal and the User Portal login pages:

```
# mv /etc/ovirt-engine/aaa/profile1.properties /etc/ovirt-engine/aaa/example.properties

# mv /etc/ovirt-engine/extensions.d/profile1-http-authn.properties /etc/ovirt-engine/extensions.d/example-http-authn.properties

# mv /etc/ovirt-engine/extensions.d/profile1-http-mapping.properties /etc/ovirt-engine/extensions.d/example-http-mapping.properties

# mv /etc/ovirt-engine/extensions.d/profile1-authz.properties /etc/ovirt-engine/extensions.d/example-authz.properties
```

- Edit the LDAP property configuration file by uncommenting an LDAP server type and updating the domain and passwords fields:

```
# vi /etc/ovirt-engine/aaa/example.properties
```

Example 15.6. Example profile: LDAP server section


```
# Select one
include = <openldap.properties>
#include = <389ds.properties>
#include = <rhds.properties>
#include = <ipa.properties>
#include = <iplanet.properties>
#include = <rfc2307-389ds.properties>
#include = <rfc2307-rhds.properties>
#include = <rfc2307-openldap.properties>
#include = <rfc2307-edir.properties>
#include = <rfc2307-generic.properties>

# Server
#
vars.server = ldap1.company.com

# Search user and its password.
#
vars.user = uid=search,cn=users,cn=accounts,dc=company,dc=com
vars.password = 123456

pool.default.serverset.single.server = ${global:vars.server}
pool.default.auth.simple.bindDN = ${global:vars.user}
pool.default.auth.simple.password = ${global:vars.password}
```

To use TLS or SSL protocol to interact with the LDAP server, obtain the root CA certificate for the LDAP server and use it to create a public keystore file. Uncomment the following lines and specify the full path to the public keystore file and the password to access the file.

**NOTE**

For more information on creating a public keystore file, see [Section D.2, “Setting Up SSL or TLS Connections between the Manager and an LDAP Server”](#).

Example 15.7. Example profile: keystore section

```
# Create keystore, import certificate chain and uncomment
# if using ssl/tls.
pool.default.ssl.startTLS = true
pool.default.ssl.truststore.file = /full/path/to/myrootca.jks
pool.default.ssl.truststore.password = password
```

- Review the authentication configuration file. The profile name visible to users on the Administration Portal and the User Portal login pages is defined by **ovirt.engine.aaa.authn.profile.name**. The configuration profile location must match the LDAP configuration file location. All fields can be left as default.

```
# vi /etc/ovirt-engine/extensions.d/example-http-authn.properties
```

Example 15.8. Example authentication configuration file

```
ovirt.engine.extension.name = example-http-authn
ovirt.engine.extension.bindings.method = jbossmodule
ovirt.engine.extension.binding.jbossmodule.module =
org.ovirt.engine-extensions.aaa.misc
ovirt.engine.extension.binding.jbossmodule.class =
org.ovirt.engineextensions.aaa.misc.http.AuthnExtension
ovirt.engine.extension.provides =
org.ovirt.engine.api.extensions.aaa.Authn
ovirt.engine.aaa.authn.profile.name = example-http
ovirt.engine.aaa.authn.authz.plugin = example-authz
ovirt.engine.aaa.authn.mapping.plugin = example-http-mapping
config.artifact.name = HEADER
config.artifact.arg = X-Remote-User
```

- Review the authorization configuration file. The configuration profile location must match the LDAP configuration file location. All fields can be left as default.

```
# vi /etc/ovirt-engine/extensions.d/example-authz.properties
```

Example 15.9. Example authorization configuration file

```
ovirt.engine.extension.name = example-authz
ovirt.engine.extension.bindings.method = jbossmodule
ovirt.engine.extension.binding.jbossmodule.module =
org.ovirt.engine-extensions.aaa.ldap
ovirt.engine.extension.binding.jbossmodule.class =
org.ovirt.engineextensions.aaa.ldap.AuthzExtension
```



```
# chmod 640 /etc/ovirt-engine/extensions.d/example-authz.properties
```

12. Restart the Apache service and the **ovirt-engine** service:

```
# systemctl restart httpd.service
```

```
# systemctl restart ovirt-engine.service
```

15.5. USER AUTHORIZATION

15.5.1. User Authorization Model

Red Hat Virtualization applies authorization controls based on the combination of the three components:

- The user performing the action
- The type of action being performed
- The object on which the action is being performed

15.5.2. User Actions

For an action to be successfully performed, the **user** must have the appropriate **permission** for the **object** being acted upon. Each type of action corresponds to a **permission**. There are many different permissions in the system, so for simplicity:





Figure 15.3. Actions



IMPORTANT

Some actions are performed on more than one object. For example, copying a template to another storage domain will impact both the template and the destination storage domain. The user performing an action must have appropriate permissions for all objects the action impacts.

15.6. ADMINISTERING USER TASKS FROM THE ADMINISTRATION PORTAL

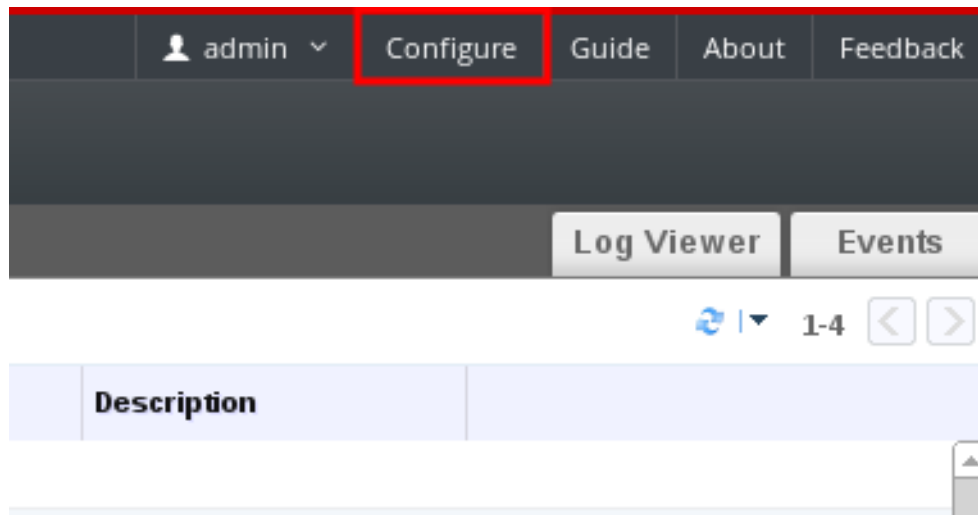
15.6.1. Adding Users and Assigning User Portal Permissions

Users must be created already before they can be added and assigned roles and

permissions. The roles and permissions assigned in this procedure give the user the permission to log in to the User Portal and to start creating virtual machines. The procedure also applies to group accounts.

Procedure 15.7. Adding Users and Assigning User Portal Permissions

1. On the header bar, click **Configure** to open the **Configure** window. Click **System Permissions**.



2. Click **Add** to open the **Add System Permission to User** window.
3. Select a profile under **Search**. The profile is the domain you want to search. Enter a name or part of a name in the search text field, and click **GO**. Alternatively, click **GO** to view a list of all users and groups.
4. Select the check boxes for the appropriate users or groups.
5. Select an appropriate role to assign under **Role to Assign**. The **UserRole** role gives the user account the permission to log in to the User Portal.
6. Click **OK**.

Log in to the User Portal to verify that the user account has the permissions to log in.

15.6.2. Viewing User Information

You can view detailed information about each user in the **Users** tab.

Procedure 15.8. Viewing User Information

1. Click the **Users** tab to display the list of authorized users.
2. Select the user, or perform a search if the user is not visible on the results list.
3. The details pane displays for the selected user, usually with the **General** tab displaying general information, such as the domain name, email and status of the user.
4. The other tabs allow you to view groups, permissions, quotas, and events for the user.

For example, to view the groups to which the user belongs, click the **Directory Groups** tab.

15.6.3. Viewing User Permissions on Resources

Users can be assigned permissions on specific resources or a hierarchy of resources. You can view the assigned users and their permissions on each resource.

Procedure 15.9. Viewing User Permissions on Resources

1. Click the resource tabs, and select a resource in the results list.
2. Click the **Permissions** tab of the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.

15.6.4. Removing Users

When a user account is no longer required, remove it from Red Hat Virtualization.

Procedure 15.10. Removing Users

1. Click the **Users** tab to display the list of authorized users.
2. Select the user to be removed. Ensure the user is not running any virtual machines.
3. Click the **Remove** button. A message displays prompting you to confirm the removal. Click **OK**.

The user is removed from Red Hat Virtualization, but not from the external directory.

15.6.5. Viewing Logged-In Users

You can view the users who are currently logged in, along with session times and other details. Click the **Active User Sessions** entry in the tree pane to view the details of the session for each logged-in user.

The **Active User Sessions** tab displays the **Session DB ID**, **User Name**, **Authorization provider**, **User id**, **Source IP**, **Session Start Time**, and **Session Last Active Time**.

15.6.6. Terminating a User Session

You can terminate the session of a user who is currently logged in.

Procedure 15.11. Terminating a User Session

1. Click the **Active User Sessions** entry in the tree pane.
2. Select the user session to be terminated.
3. Click **Terminate Session**.
4. Click **OK**.

15.7. ADMINISTERING USER TASKS FROM THE COMMAND LINE

15.7.1. Managing Users

You can use the `ovirt-aaa-jdbc-tool` tool to manage user accounts on the internal domain. Changes made using the tool take effect immediately and do not require you to restart the `ovirt-engine` service. For a full list of user options, run `ovirt-aaa-jdbc-tool user --help`. Common examples are provided in this section.

Procedure 15.12. Creating a User

This procedure shows you how to create a user, set the user password, and add it to your Red Hat Virtualization environment.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2. Create a new user account. Optionally use `--attribute` to specify account details. For a full list of options, run `ovirt-aaa-jdbc-tool user add --help`.

```
# ovirt-aaa-jdbc-tool user add test1 --attribute=firstName=John --
attribute=lastName=Doe
adding user test1...
user added successfully
```

3. Set a password. You must set a value for `--password-valid-to`, otherwise the password expiry time defaults to the current time. The date format is `yyyy-MM-dd HH:mm:ssX`. In this example, `-0800` stands for GMT minus 8 hours. For more options, run `ovirt-aaa-jdbc-tool user password-reset --help`.

```
# ovirt-aaa-jdbc-tool user password-reset test1 --password-valid-
to="2025-08-01 12:00:00-0800"
Password:
updating user test1...
user updated successfully
```



NOTE

By default, the password policy for user accounts on the internal domain has the following restrictions:

- A minimum of 6 characters.
- Three previous passwords used cannot be set again during the password change.

For more information on the password policy and other default settings, run `ovirt-aaa-jdbc-tool settings show`.

4. Add the newly created user in the Administration Portal and assign the user appropriate roles and permissions. See [Section 15.6.1, “Adding Users and Assigning User Portal Permissions”](#) for more information.

Procedure 15.13. Viewing User Information

This procedure shows you how to view user account information. More information is displayed than in the Administration Portal, **Users** tab.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2. Run the following command:

```
# ovirt-aaa-jdbc-tool user show test1
```

Procedure 15.14. Editing User Information

This procedure shows you how to update user account information.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2. Run the following command to edit user information. This example updates the email address.

```
# ovirt-aaa-jdbc-tool user edit test1 --  
attribute=email=jdoe@example.com
```

Procedure 15.15. Removing a User

This procedure shows you how to delete a user account

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2. Delete the user:

```
# ovirt-aaa-jdbc-tool user delete test1
```

3. Remove the user from the Administration Portal. See [Section 15.6.4, “Removing Users”](#) for more information.

15.7.2. Changing the Password for the Internal Administrative User

To reset the password of the internal administrative user (**admin@internal**), use the **ovirt-aaa-jdbc-tool** tool. You do not need to restart the **ovirt-engine** service for the password change to take effect.

By default, the password policy for user accounts on the internal domain has the following restrictions:

- A minimum of 6 characters.
- Three previous passwords used cannot be set again during the password change.

For more information on the password policy and other default settings, run **ovirt-aaa-jdbc-tool settings show**.

Procedure 15.16. Resetting the Password for the Internal Administrative User

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.

- To change the password in interactive mode, run the following command. You must set a value for `--password-valid-to`, otherwise the password expiry time defaults to the current time. The date format is `yyyy-MM-dd HH:mm:ssX`. In this example, `Z` stands for UTC time. For more options, run `ovirt-aaa-jdbc-tool user password-reset --help`.

```
# ovirt-aaa-jdbc-tool user password-reset admin --password-valid-to="2025-08-01 12:00:00Z"
```


15.7.3. Disabling the Internal Administrative User

You can disable users on the local domains including the `admin@internal` user created during `engine-setup`. Make sure you have at least one user in the environment with full administrative permissions before disabling the default `admin` user.

Procedure 15.17. Disabling the Internal Administrative User

- Log in to the machine on which the Red Hat Virtualization Manager is installed.
- Make sure another user with the `SuperUser` role has been added to the environment. See [Section 15.6.1, “Adding Users and Assigning User Portal Permissions”](#) for more information.
- Disable the default `admin` user:

```
# ovirt-aaa-jdbc-tool user edit admin --flag=+disabled
```



NOTE

To enable a disabled user, run `ovirt-aaa-jdbc-tool user edit username --flag=-disabled`

15.7.4. Managing Groups

You can use the `ovirt-aaa-jdbc-tool` tool to manage group accounts on your internal domain. Managing group accounts is similar to managing user accounts. For a full list of group options, run `ovirt-aaa-jdbc-tool group --help`. Common examples are provided in this section.

Procedure 15.18. Creating a Group

This procedure shows you how to create a group account, add users to the group, and view the details of the group.

- Log in to the machine on which the Red Hat Virtualization Manager is installed.
- Create a new group:

```
# ovirt-aaa-jdbc-tool group add group1
```

- Add users to the group. The users must be created already.

```
# ovirt-aaa-jdbc-tool group-manage useradd group1 --user=test1
```

**NOTE**

For a full list of the group-manage options, run **ovirt-aaa-jdbc-tool group-manage --help**.

4. View group account details:

```
# ovirt-aaa-jdbc-tool group show group1
```

5. Add the newly created group in the Administration Portal and assign the group appropriate roles and permissions. The users in the group inherit the roles and permissions of the group. See [Section 15.6.1, “Adding Users and Assigning User Portal Permissions”](#) for more information.

Procedure 15.19. Creating Nested Groups

This procedure shows you how to create groups within groups.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2. Create the first group:

```
# ovirt-aaa-jdbc-tool group add group1
```

3. Create the second group:

```
# ovirt-aaa-jdbc-tool group add group1-1
```

4. Add the second group to the first group:

```
# ovirt-aaa-jdbc-tool group-manage groupadd group1 --group=group1-1
```

5. Add the first group in the Administration Portal and assign the group appropriate roles and permissions. See [Section 15.6.1, “Adding Users and Assigning User Portal Permissions”](#) for more information.

15.7.5. Querying Users and Groups

The **query** module allows you to query user and group information. For a full list of options, run **ovirt-aaa-jdbc-tool query --help**.

Procedure 15.20. Listing All User or Group Account Details

This procedure shows you how to list all account information.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2.
 - List all user account details:

```
# ovirt-aaa-jdbc-tool query --what=user
```

- List all group account details:

```
# ovirt-aaa-jdbc-tool query --what=group
```

Procedure 15.21. Listing Filtered Account Details

This procedure shows you how to apply filters when listing account information.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2.
 - List user account details with names that start with the character j.

```
# ovirt-aaa-jdbc-tool query --what=user --pattern="name=j*"
```

- List groups that have the department attribute set to *marketing*:

```
# ovirt-aaa-jdbc-tool query --what=group --
pattern="department=marketing"
```

15.7.6. Managing Account Settings

To change the default account settings, use the **ovirt-aaa-jdbc-tool settings** module.

Procedure 15.22. Updating Account Settings

This procedure shows you how to update the default account settings.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.
2. Run the following command to show all the settings available:

```
# ovirt-aaa-jdbc-tool setting show
```

3. Change the desired settings:

- This example updates the default log in session time to 60 minutes for all user accounts. The default value is 10080 minutes.

```
# ovirt-aaa-jdbc-tool setting set --name=MAX_LOGIN_MINUTES --
value=60
```

- This example updates the number of failed login attempts a user can perform before the user account is locked. The default value is 5.

```
# ovirt-aaa-jdbc-tool setting set --
name=MAX_FAILURES_SINCE_SUCCESS --value=3
```

**NOTE**

To unlock a locked user account, run **`ovirt-aaa-jdbc-tool user unlock test1`**.

15.8. CONFIGURING ADDITIONAL LOCAL DOMAINS

Creating additional local domains other than the default **internal** domain is also supported. This can be done using the **`ovirt-engine-extension-aaa-jdbc`** extension and allows you to create multiple domains without attaching external directory servers, though the use case may not be common for enterprise environments.

Additionally created local domains will not get upgraded automatically during standard Red Hat Virtualization upgrades and need to be upgraded manually for each future release. For more information on creating additional local domains and how to upgrade the domains, see the README file at **`/usr/share/doc/ovirt-engine-extension-aaa-jdbc-version/README.admin`**.

CHAPTER 16. QUOTAS AND SERVICE LEVEL AGREEMENT POLICY

16.1. INTRODUCTION TO QUOTA

Quota is a resource limitation tool provided with Red Hat Virtualization. Quota may be thought of as a layer of limitations on top of the layer of limitations set by User Permissions.

Quota is a data center object.

Quota allows administrators of Red Hat Virtualization environments to limit user access to memory, CPU, and storage. Quota defines the memory resources and storage resources an administrator can assign users. As a result users may draw on only the resources assigned to them. When the quota resources are exhausted, Red Hat Virtualization does not permit further user actions.

There are two different kinds of Quota:

Table 16.1. The Two Different Kinds of Quota

Quota type	Definition
Run-time Quota	This quota limits the consumption of runtime resources, like CPU and memory.
Storage Quota	This quota limits the amount of storage available.

Quota, like SELinux, has three modes:

Table 16.2. Quota Modes

Quota Mode	Function
Enforced	This mode puts into effect the Quota that you have set in audit mode, limiting resources to the group or user affected by the quota.
Audit	This mode allows you to change Quota settings. Choose this mode to increase or decrease the amount of runtime quota and the amount of storage quota available to users affected by it.
Disabled	This mode turns off the runtime and storage limitations defined by the quota.

When a user attempts to run a virtual machine, the specifications of the virtual machine are compared to the storage allowance and the runtime allowance set in the applicable quota.

If starting a virtual machine causes the aggregated resources of all running virtual machines covered by a quota to exceed the allowance defined in the quota, then the Manager refuses to run the virtual machine.

When a user creates a new disk, the requested disk size is added to the aggregated disk usage of all the other disks covered by the applicable quota. If the new disk takes the total aggregated disk usage above the amount allowed by the quota, disk creation fails.

Quota allows for resource sharing of the same hardware. It supports hard and soft thresholds. Administrators can use a quota to set thresholds on resources. These thresholds appear, from the user's point of view, as 100% usage of that resource. To prevent failures when the customer unexpectedly exceeds this threshold, the interface supports a "grace" amount by which the threshold can be briefly exceeded. Exceeding the threshold results in a warning sent to the customer.



IMPORTANT

Quota imposes limitations upon the running of virtual machines. Ignoring these limitations is likely to result in a situation in which you cannot use your virtual machines and virtual disks.

When quota is running in enforced mode, virtual machines and disks that do not have quotas assigned cannot be used.

To power on a virtual machine, a quota must be assigned to that virtual machine.

To create a snapshot of a virtual machine, the disk associated with the virtual machine must have a quota assigned.

When creating a template from a virtual machine, you are prompted to select the quota that you want the template to consume. This allows you to set the template (and all future machines created from the template) to consume a different quota than the virtual machine and disk from which the template is generated.

16.2. SHARED QUOTA AND INDIVIDUALLY DEFINED QUOTA

Users with SuperUser permissions can create quotas for individual users or quotas for groups.

Group quotas can be set for Active Directory users. If a group of ten users are given a quota of 1 TB of storage and one of the ten users fills the entire terabyte, then the entire group will be in excess of the quota and none of the ten users will be able to use any of the storage associated with their group.

An individual user's quota is set for only the individual. Once the individual user has used up all of his or her storage or runtime quota, the user will be in excess of the quota and the user will no longer be able to use the storage associated with his or her quota.

16.3. QUOTA ACCOUNTING

When a quota is assigned to a consumer or a resource, each action by that consumer or on the resource involving storage, vCPU, or memory results in quota consumption or quota release.

Since the quota acts as an upper bound that limits the user's access to resources, the quota calculations may differ from the actual current use of the user. The quota is calculated for the max growth potential and not the current usage.

Example 16.1. Accounting example

A user runs a virtual machine with 1 vCPU and 1024 MB memory. The action consumes 1 vCPU and 1024 MB of the quota assigned to that user. When the virtual machine is stopped 1 vCPU and 1024 MB of RAM are released back to the quota assigned to that user. Run-time quota consumption is accounted for only during the actual run-time of the consumer.

A user creates a virtual thin provision disk of 10 GB. The actual disk usage may indicate only 3 GB of that disk are actually in use. The quota consumption, however, would be 10 GB, the max growth potential of that disk.

16.4. ENABLING AND CHANGING A QUOTA MODE IN A DATA CENTER

This procedure enables or changes the quota mode in a data center. You must select a quota mode before you can define quotas. You must be logged in to the Administration Portal to follow the steps of this procedure.

Use **Audit** mode to test your quota to make sure it works as you expect it to. You do not need to have your quota in **Audit** mode to create or change a quota.

Procedure 16.1. Enabling and Changing Quota in a Data Center

1. Click the **Data Centers** tab in the Navigation Pane.
2. From the list of data centers displayed in the Navigation Pane, choose the data center whose quota policy you plan to edit.
3. Click **Edit** in the top left of the Navigation Pane.

An **Edit Data Center** window opens.

4. In the **Quota Mode** drop-down, change the quota mode to **Enforced**.
5. Click **OK**.

You have now enabled a quota mode at the Data Center level. If you set the quota mode to **Audit** during testing, then you must change it to **Enforced** in order for the quota settings to take effect.

16.5. CREATING A NEW QUOTA POLICY

You have enabled quota mode, either in Audit or Enforcing mode. You want to define a quota policy to manage resource usage in your data center.

Procedure 16.2. Creating a New Quota Policy

1. In tree mode, select the data center. The **Quota** tab appears in the Navigation Pane.

2. Click the **Quota** tab in the Navigation Pane.
3. Click **Add** in the Navigation Pane. The **New Quota** window opens.
4. Fill in the **Name** field with a meaningful name.
Fill in the **Description** field with a meaningful name.
5. In the **Memory & CPU** section of the **New Quota** window, use the green slider to set **Cluster Threshold**.
6. In the **Memory & CPU** section of the **New Quota** window, use the blue slider to set **Cluster Grace**.
7. Select the **All Clusters** or the **Specific Clusters** radio button. If you select **Specific Clusters**, select the check box of the clusters that you want to add a quota policy to.
8. Click **Edit** to open the **Edit Quota** window.
9. Under the **Memory** field, select either the **Unlimited** radio button (to allow limitless use of Memory resources in the cluster), or select the **limit to** radio button to set the amount of memory set by this quota. If you select the **limit to** radio button, input a memory quota in megabytes (MB) in the **MB** field.
10. Under the **CPU** field, select either the **Unlimited** radio button or the **limit to** radio button to set the amount of CPU set by this quota. If you select the **limit to** radio button, input a number of vCPUs in the **vCpus** field.
11. Click **OK** in the **Edit Quota** window.
12. In the **Storage** section of the **New Quota** window, use the green slider to set **Storage Threshold**.
13. In the **Storage** section of the **New Quota** window, use the blue slider to set **Storage Grace**.
14. Select the **All Storage Domains** or the **Specific Storage Domains** radio button. If you select **Specific Storage Domains**, select the check box of the storage domains that you want to add a quota policy to.
15. Click **Edit** to open the **Edit Quota** window.
16. Under the **Storage Quota** field, select either the **Unlimited** radio button (to allow limitless use of Storage) or the **limit to** radio button to set the amount of storage to which quota will limit users. If you select the **limit to** radio button, input a storage quota size in gigabytes (GB) in the **GB** field.
17. Click **OK** in the **Edit Quota** window. You are returned to the **New Quota** window.
18. Click **OK** in the **New Quota** window.

Result

You have created a new quota policy.

16.6. EXPLANATION OF QUOTA THRESHOLD SETTINGS

Table 16.3. Quota thresholds and grace

Setting	Definition
Cluster Threshold	The amount of cluster resources available per data center.
Cluster Grace	The amount of the cluster available for the data center after exhausting the data center's Cluster Threshold.
Storage Threshold	The amount of storage resources available per data center.
Storage Grace	The amount of storage available for the data center after exhausting the data center's Storage Threshold.

If a quota is set to 100 GB with 20% Grace, then consumers are blocked from using storage after they use 120 GB of storage. If the same quota has a Threshold set at 70%, then consumers receive a warning when they exceed 70 GB of storage consumption (but they remain able to consume storage until they reach 120 GB of storage consumption.) Both "Threshold" and "Grace" are set relative to the quota. "Threshold" may be thought of as the "soft limit", and exceeding it generates a warning. "Grace" may be thought of as the "hard limit", and exceeding it makes it impossible to consume any more storage resources.

16.7. ASSIGNING A QUOTA TO AN OBJECT

Summary

This procedure explains how to associate a virtual machine with a quota.

Procedure 16.3. Assigning a Quota to a Virtual Machine

1. In the navigation pane, select the Virtual Machine to which you plan to add a quota.
2. Click **Edit**. The **Edit Virtual Machine** window appears.
3. Select the quota you want the virtual machine to consume. Use the **Quota** drop-down to do this.
4. Click **OK**.

Result

You have designated a quota for the virtual machine you selected.

Summary

This procedure explains how to associate a virtual disk with a quota.

Procedure 16.4. Assigning a Quota to a Virtual Disk

1. In the navigation pane, select the Virtual Machine whose disk(s) you plan to add a quota.

2. In the details pane, select the disk you plan to associate with a quota.
3. Click **Edit**. The **Edit Virtual Disk** window appears.
4. Select the quota you want the virtual disk to consume.
5. Click **OK**.

Result

You have designated a quota for the virtual disk you selected.



IMPORTANT

Quota must be selected for all objects associated with a virtual machine, in order for that virtual machine to work. If you fail to select a quota for the objects associated with a virtual machine, the virtual machine will not work. The error that the Manager throws in this situation is generic, which makes it difficult to know if the error was thrown because you did not associate a quota with all of the objects associated with the virtual machine. It is not possible to take snapshots of virtual machines that do not have an assigned quota. It is not possible to create templates of virtual machines whose virtual disks do not have assigned quotas.

16.8. USING QUOTA TO LIMIT RESOURCES BY USER

Summary

This procedure describes how to use quotas to limit the resources a user has access to.

Procedure 16.5. Assigning a User to a Quota

1. In the tree, click the Data Center with the quota you want to associate with a User.
2. Click the **Quota** tab in the navigation pane.
3. Select the target quota in the list in the navigation pane.
4. Click the **Consumers** tab in the details pane.
5. Click **Add** at the top of the details pane.
6. In the **Search** field, type the name of the user you want to associate with the quota.
7. Click **GO**.
8. Select the check box at the left side of the row containing the name of the target user.
9. Click **OK** in the bottom right of the **Assign Users and Groups to Quota** window.

Result

After a short time, the user will appear in the **Consumers** tab of the details pane.

16.9. EDITING QUOTAS

Summary

This procedure describes how to change existing quotas.

Procedure 16.6. Editing Quotas

1. On the tree pane, click on the data center whose quota you want to edit.
2. Click on the **Quota** tab in the Navigation Pane.
3. Click the name of the quota you want to edit.
4. Click **Edit** in the Navigation pane.
5. An **Edit Quota** window opens. If required, enter a meaningful name in the **Name** field.
6. If required, you can enter a meaningful description in the **Description** field.
7. Select either the **All Clusters** radio button or the **Specific Clusters** radio button. Move the **Cluster Threshold** and **Cluster Grace** sliders to the desired positions on the **Memory & CPU** slider.
8. Select either the **All Storage Domains** radio button or the **Specific Storage Domains** radio button. Move the **Storage Threshold** and **Storage Grace** sliders to the desired positions on the **Storage** slider.
9. Click **OK** in the **Edit Quota** window to confirm the new quota settings.

Result

You have changed an existing quota.

16.10. REMOVING QUOTAS

Summary

This procedure describes how to remove quotas.

Procedure 16.7. Removing Quotas

1. On the tree pane, click on the data center whose quota you want to edit.
2. Click on the **Quota** tab in the Navigation Pane.
3. Click the name of the quota you want to remove.
4. Click **Remove** at the top of the Navigation pane, under the row of tabs.
5. Click **OK** in the **Remove Quota(s)** window to confirm the removal of this quota.

Result

You have removed a quota.

16.11. SERVICE LEVEL AGREEMENT POLICY ENFORCEMENT

Summary

This procedure describes how to set service level agreement CPU features.

Procedure 16.8. Setting a Service Level Agreement CPU Policy

1. Select **New VM** in the Navigation Pane.
2. Select **Show Advanced Options**.
3. Select the **Resource Allocation** tab.

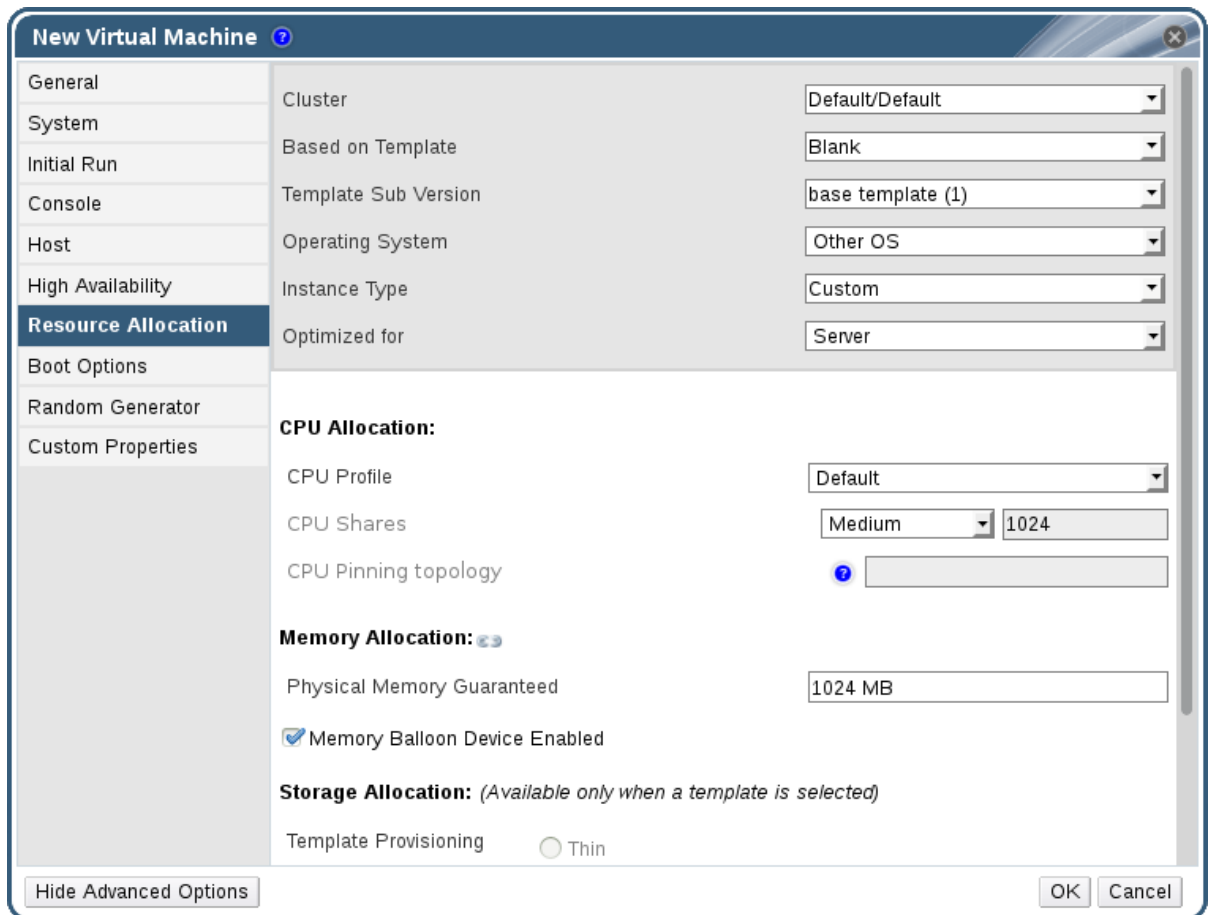


Figure 16.1. Service Level Agreement Policy Enforcement - CPU Allocation Menu

4. Specify **CPU Shares**. Possible options are **Low**, **Medium**, **High**, **Custom**, and **Disabled**. Virtual machines set to **High** receive twice as many shares as **Medium**, and virtual machines set to **Medium** receive twice as many shares as virtual machines set to **Low**. **Disabled** instructs VDSM to use an older algorithm for determining share dispensation; usually the number of shares dispensed under these conditions is 1020.

Result

You have set a service level agreement CPU policy. The CPU consumption of users is now governed by the policy you have set.

CHAPTER 17. EVENT NOTIFICATIONS

17.1. CONFIGURING EVENT NOTIFICATIONS IN THE ADMINISTRATION PORTAL

Summary

The Red Hat Virtualization Manager can notify designated users via email when specific events occur in the environment that the Red Hat Virtualization Manager manages. To use this functionality, you must set up a mail transfer agent to deliver messages. Only email notifications can be configured through the Administration Portal. SNMP traps must be configured on the Manager machine.

Procedure 17.1. Configuring Event Notifications

1. Ensure you have set up the mail transfer agent with the appropriate variables.
2. Use the **Users** resource tab, tree mode, or the search function to find and select the user to which event notifications will be sent.
3. Click the **Event Notifier** tab in the details pane to list the events for which the user will be notified. This list is blank if you have not configured any event notifications for that user.
4. Click **Manage Events** to open the **Add Event Notification** window.

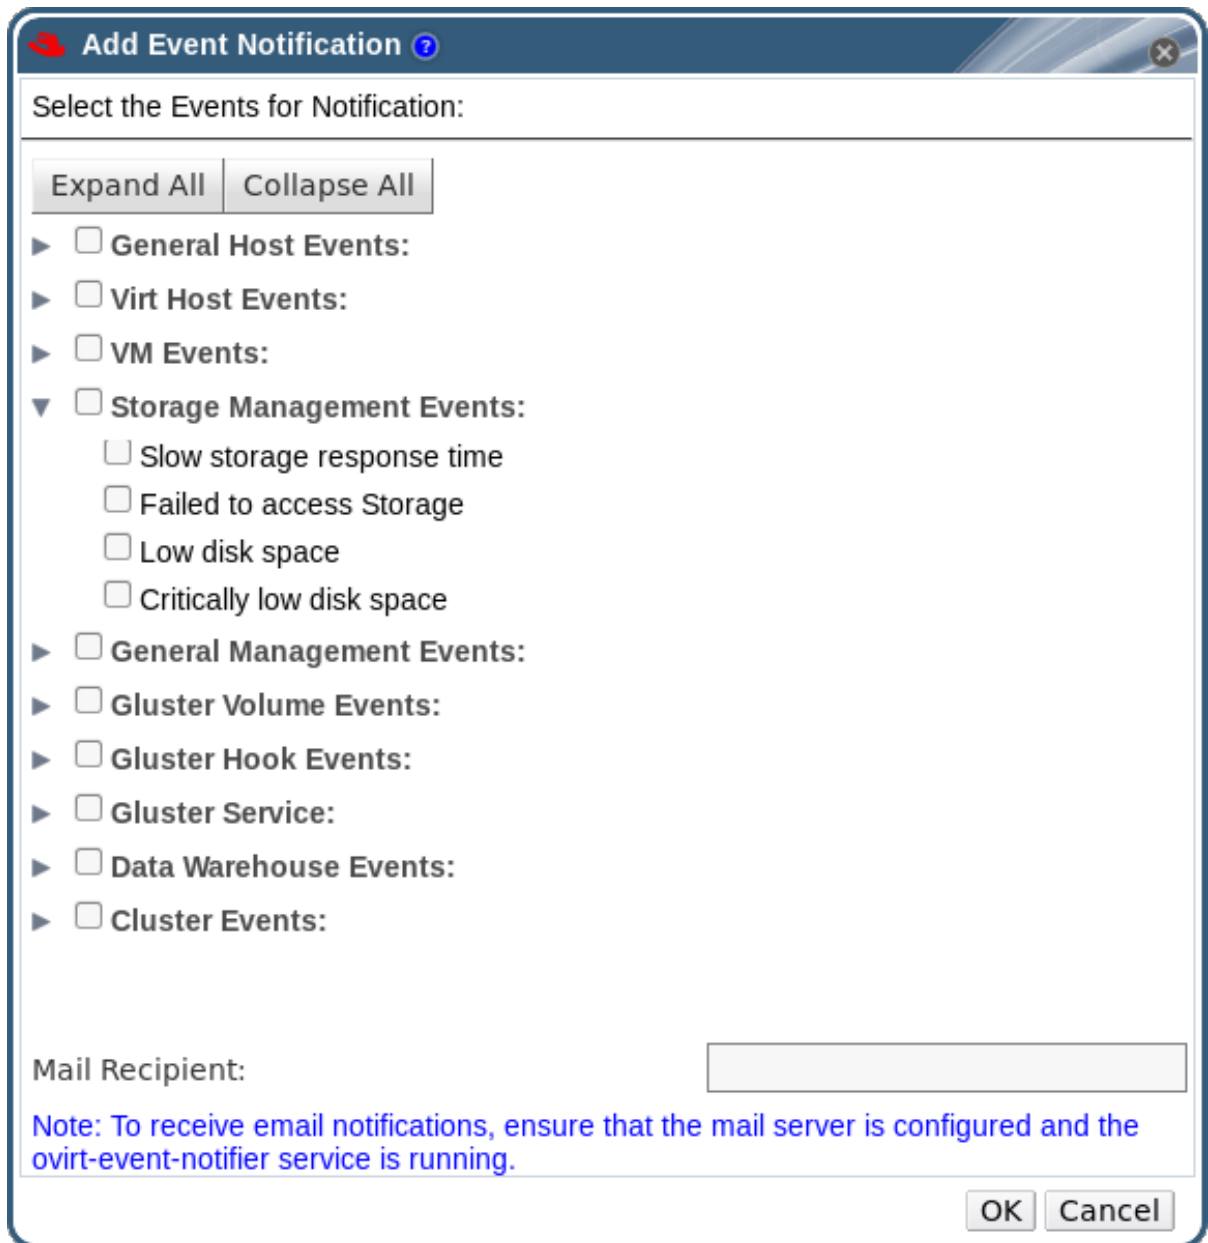


Figure 17.1. The Add Events Notification Window

5. Use the **Expand All** button or the subject-specific expansion buttons to view the events.
6. Select the appropriate check boxes.
7. Enter an email address in the **Mail Recipient** field.
8. Click **OK** to save changes and close the window.
9. Add and start the **ovirt-engine-notifier** service on the Red Hat Virtualization Manager. This activates the changes you have made:

```
# systemctl daemon-reload
# systemctl enable ovirt-engine-notifier.service
# systemctl restart ovirt-engine-notifier.service
```

Result

The specified user now receives emails based on events in the Red Hat Virtualization environment. The selected events display on the **Event Notifier** tab for that user.

17.2. CANCELING EVENT NOTIFICATIONS IN THE ADMINISTRATION PORTAL

Summary

A user has configured some unnecessary email notifications and wants them canceled.

Procedure 17.2. Canceling Event Notifications

1. In the **Users** tab, select the user or the user group.
2. Select the **Event Notifier** tab in the details pane to list events for which the user receives email notifications.
3. Click **Manage Events** to open the **Add Event Notification** window.
4. Use the **Expand All** button, or the subject-specific expansion buttons, to view the events.
5. Clear the appropriate check boxes to remove notification for that event.
6. Click **OK** to save changes and close the window.

Result

You have canceled unnecessary event notifications for the user.

17.3. PARAMETERS FOR EVENT NOTIFICATIONS IN OVIRT-ENGINE-NOTIFIER.CONF

The event notifier configuration file can be found in `/usr/share/ovirt-engine/services/ovirt-engine-notifier/ovirt-engine-notifier.conf`.

Table 17.1. ovirt-engine-notifier.conf variables

Variable Name	Default	Remarks
SENSITIVE_KEYS	none	A comma-separated list of keys that will not be logged.
JBOSS_HOME	/opt/rh/eap7/root /usr/share/wildfly	The location of the JBoss application server used by the Manager.
ENGINE_ETC	/etc/ovirt-engine	The location of the etc directory used by the Manager.
ENGINE_LOG	/var/log/ovirt-engine	The location of the logs directory used by the Manager.

Variable Name	Default	Remarks
ENGINE_USR	/usr/share/ovirt-engine	The location of the usr directory used by the Manager.
ENGINE_JAVA_MODULEPATH	\${ENGINE_USR}/modules	The file path to which the JBoss modules are appended.
NOTIFIER_DEBUG_ADDRESS	none	The address of a machine that can be used to perform remote debugging of the Java virtual machine that the notifier uses.
NOTIFIER_STOP_TIME	30	The time, in seconds, after which the service will time out.
NOTIFIER_STOP_INTERVAL	1	The time, in seconds, by which the timeout counter will be incremented.
INTERVAL_IN_SECONDS	120	The interval in seconds between instances of dispatching messages to subscribers.
IDLE_INTERVAL	30	The interval, in seconds, between which low-priority tasks will be performed.
DAYS_TO_KEEP_HISTORY	0	This variable sets the number of days dispatched events will be preserved in the history table. If this variable is not set, events remain on the history table indefinitely.
FAILED_QUERIES_NOTIFICATION_THRESHOLD	30	The number of failed queries after which a notification email is sent. A notification email is sent after the first failure to fetch notifications, and then once every time the number of failures specified by this variable is reached. If you specify a value of 0 or 1 , an email will be sent with each failure.
FAILED_QUERIES_NOTIFICATION_RECIPIENTS	none	The email addresses of the recipients to which notification emails will be sent. Email addresses must be separated by a comma. This entry has been deprecated by the FILTER variable.

Variable Name	Default	Remarks
DAYS_TO_SEND_ON_STARTUP	0	The number of days of old events that will be processed and sent when the notifier starts.
FILTER	exclude:*	The algorithm used to determine the triggers for and recipients of email notifications. The value for this variable comprises a combination of include or exclude , the event, and the recipient. For example, include:VDC_START(smtp:mail@example.com) \${FILTER}
MAIL_SERVER	none	The SMTP mail server address. Required.
MAIL_PORT	25	The port used for communication. Possible values include 25 for plain SMTP, 465 for SMTP with SSL, and 587 for SMTP with TLS.
MAIL_USER	none	If SSL is enabled to authenticate the user, then this variable must be set. This variable is also used to specify the "from" user address when the MAIL_FROM variable is not set. Some mail servers do not support this functionality. The address is in RFC822 format.
SENSITIVE_KEYS	\${SENSITIVE_KEYS},MAIL_PASSWORD	Required to authenticate the user if the mail server requires authentication or if SSL or TLS is enabled.
MAIL_PASSWORD	none	Required to authenticate the user if the mail server requires authentication or if SSL or TLS is enabled.
MAIL_SMTP_ENCRYPTION	none	The type of encryption to be used in communication. Possible values are none, ssl, tls .
HTML_MESSAGE_FORMAT	false	The mail server sends messages in HTML format if this variable is set to true .

Variable Name	Default	Remarks
MAIL_FROM	none	This variable specifies a sender address in RFC822 format, if supported by the mail server.
MAIL_REPLY_TO	none	This variable specifies reply-to addresses in RFC822 format on sent mail, if supported by the mail server.
MAIL_SEND_INTERVAL	1	The number of SMTP messages to be sent for each IDLE_INTERVAL
MAIL_RETRIES	4	The number of times to attempt to send an email before failing.
SNMP_MANAGER	none	The IP addresses or fully qualified domain names of machines that will act as the SNMP managers. Entries must be separated by a space and can contain a port number. For example, manager1.example.com manager2.example.com:164
SNMP_COMMUNITY	public	The default SNMP community.
SNMP_OID	1.3.6.1.4.1.2312.13.1.1	The default trap object identifiers for alerts. All trap types are sent, appended with event information, to the SNMP manager when this OID is defined. Note that changing the default trap prevents generated traps from complying with the Manager's management information base.
ENGINE_INTERVAL_IN_SECONDS	300	The interval, in seconds, between monitoring the machine on which the Manager is installed. The interval is measured from the time the monitoring is complete.
ENGINE_MONITOR_RETRIES	3	The number of times the notifier attempts to monitor the status of the machine on which the Manager is installed in a given interval after a failure.

Variable Name	Default	Remarks
ENGINE_TIMEOUT_IN_SECONDS	30	The time, in seconds, to wait before the notifier attempts to monitor the status of the machine on which the Manager is installed in a given interval after a failure.
IS_HTTPS_PROTOCOL	false	This entry must be set to true if JBoss is being run in secured mode.
SSL_PROTOCOL	TLS	The protocol used by JBoss configuration connector when SSL is enabled.
SSL_IGNORE_CERTIFICATE_ERRORS	false	This value must be set to true if JBoss is running in secure mode and SSL errors is to be ignored.
SSL_IGNORE_HOST_VERIFICATION	false	This value must be set to true if JBoss is running in secure mode and host name verification is to be ignored.
REPEAT_NON_RESPONSIVE_NOTIFICATION	false	This variable specifies whether repeated failure messages will be sent to subscribers if the machine on which the Manager is installed is non-responsive.
ENGINE_PID	/var/lib/ovirt-engine/ovirt-engine.pid	The path and file name of the PID of the Manager.

17.4. CONFIGURING THE RED HAT VIRTUALIZATION MANAGER TO SEND SNMP TRAPS

Configure your Red Hat Virtualization Manager to send Simple Network Management Protocol traps to one or more external SNMP managers. SNMP traps contain system event information; they are used to monitor your Red Hat Virtualization environment. The number and type of traps sent to the SNMP manager can be defined within the Red Hat Virtualization Manager.

This procedure assumes that you have configured one or more external SNMP managers to receive traps, and that you have the following details:

- The IP addresses or fully qualified domain names of machines that will act as SNMP managers. Optionally, determine the port through which the manager receives trap notifications; by default, this is UDP port 162.

- The SNMP community. Multiple SNMP managers can belong to a single community. Management systems and agents can communicate only if they are within the same community. The default community is **public**.
- The trap object identifier for alerts. The Red Hat Virtualization Manager provides a default OID of 1.3.6.1.4.1.2312.13.1.1. All trap types are sent, appended with event information, to the SNMP manager when this OID is defined. Note that changing the default trap prevents generated traps from complying with the Manager's management information base.



NOTE

The Red Hat Virtualization Manager provides management information bases at `/usr/share/doc/ovirt-engine/mibs/OVIRT-MIB.txt` and `/usr/share/doc/ovirt-engine/mibs/REDHAT-MIB.txt`. Load the MIBs in your SNMP manager before proceeding.

Default SNMP configuration values exist on the Manager in the events notification daemon configuration file `/usr/share/ovirt-engine/services/ovirt-engine-notifier/ovirt-engine-notifier.conf`. The values outlined in the following procedure are based on the default or example values provided in that file. It is recommended that you define an override file, rather than edit the `ovirt-engine-notifier.conf` file, to persist your configuration options across system changes, like upgrades.

Procedure 17.3. Configuring SNMP Traps on the Manager

1. On the Manager, create the SNMP configuration file:

```
# vi /etc/ovirt-engine/notifier/notifier.conf.d/20-snmp.conf
```

2. Specify the SNMP manager(s), the SNMP community, and the OID in the following format:

```
SNMP_MANAGERS="manager1.example.com manager2.example.com:162"
SNMP_COMMUNITY=public
SNMP_OID=1.3.6.1.4.1.2312.13.1.1
```

3. Define which events to send to the SNMP manager:

Example 17.1. Event Examples

Send all events to the default SNMP profile:

```
FILTER="include:*(snmp:) ${FILTER}"
```

Send all events with the severity **ERROR** or **ALERT** to the default SNMP profile:

```
FILTER="include:*ERROR(snmp:) ${FILTER}"
```

```
FILTER="include:*ALERT(snmp:) ${FILTER}"
```

Send events for `VDC_START` to the specified email address:

-

```
FILTER="include:VDC_START(snmp:mail@example.com) ${FILTER}"
```

Send events for everything but `VDC_START` to the default SNMP profile:

```
FILTER="exclude:VDC_START include:*(snmp:) ${FILTER}"
```

This is the default filter defined in `ovirt-engine-notifier.conf`; if you do not disable this filter or apply overriding filters, no notifications will be sent:

```
FILTER="exclude:*"
```

`VDC_START` is an example of the audit log messages available. A full list of audit log messages can be found in `/usr/share/doc/ovirt-engine/AuditLogMessages.properties`. Alternatively, filter results within your SNMP manager.

4. Save the file.
5. Start the `ovirt-engine-notifier` service, and ensure that this service starts on boot:

```
# systemctl start ovirt-engine-notifier.service
# systemctl enable ovirt-engine-notifier.service
```

Check your SNMP manager to ensure that traps are being received.



NOTE

`SNMP_MANAGERS`, `MAIL_SERVER`, or both must be properly defined in `/usr/share/ovirt-engine/services/ovirt-engine-notifier/ovirt-engine-notifier.conf` or in an override file in order for the notifier service to run.

CHAPTER 18. UTILITIES

18.1. THE OVIRT ENGINE RENAME TOOL

18.1.1. The oVirt Engine Rename Tool

When the **engine-setup** command is run in a clean environment, the command generates a number of certificates and keys that use the fully qualified domain name of the Manager supplied during the setup process. If the fully qualified domain name of the Manager must be changed later on (for example, due to migration of the machine hosting the Manager to a different domain), the records of the fully qualified domain name must be updated to reflect the new name. The **ovirt-engine-rename** command automates this task.

The **ovirt-engine-rename** command updates records of the fully qualified domain name of the Manager in the following locations:

- /etc/ovirt-engine/engine.conf.d/10-setup-protocols.conf
- /etc/ovirt-engine/imageuploader.conf.d/10-engine-setup.conf
- /etc/ovirt-engine/isouploader.conf.d/10-engine-setup.conf
- /etc/ovirt-engine/logcollector.conf.d/10-engine-setup.conf
- /etc/pki/ovirt-engine/cert.conf
- /etc/pki/ovirt-engine/cert.template
- /etc/pki/ovirt-engine/certs/apache.cer
- /etc/pki/ovirt-engine/keys/apache.key.nopass
- /etc/pki/ovirt-engine/keys/apache.p12



WARNING

While the **ovirt-engine-rename** command creates a new certificate for the web server on which the Manager runs, it does not affect the certificate for the engine or the certificate authority. Due to this, there is some risk involved in using the **ovirt-engine-rename** command, particularly in environments that have been upgraded from Red Hat Enterprise Virtualization 3.2 and earlier. Therefore, changing the fully qualified domain name of the Manager by running **engine-cleanup** and **engine-setup** is recommended where possible.

18.1.2. Syntax for the oVirt Engine Rename Command

The basic syntax for the **ovirt-engine-rename** command is:

```
# /usr/share/ovirt-engine/setup/bin/ovirt-engine-rename
```


■

The command also accepts the following options:

`--newname=[new name]`

Allows you to specify the new fully qualified domain name for the Manager without user interaction.

`--log=[file]`

Allows you to specify the path and name of a file into which logs of the rename operation are to be written.

`--config=[file]`

Allows you to specify the path and file name of a configuration file to load into the rename operation.

`--config-append=[file]`

Allows you to specify the path and file name of a configuration file to append to the rename operation. This option can be used to specify the path and file name of an answer file.

`--generate-answer=[file]`

Allows you to specify the path and file name of a file into which your answers to and the values changed by the **`ovirt-engine-rename`** command are recorded.

18.1.3. Using the oVirt Engine Rename Tool

Summary

You can use the **`ovirt-engine-rename`** command to update records of the fully qualified domain name of the Manager.

The tool checks whether the Manager provides a local ISO or Data storage domain. If it does, the tool prompts the user to eject, shut down, or place into maintenance mode any virtual machine or storage domain connected to the storage before continuing with the operation. This ensures that virtual machines do not lose connectivity with their virtual disks, and prevents ISO storage domains from losing connectivity during the renaming process.

Procedure 18.1. Renaming the Red Hat Virtualization Manager

1. Prepare all DNS and other relevant records for the new fully qualified domain name.
2. Update the DHCP server configuration if DHCP is used.
3. Update the host name on the Manager.
4. Run the following command:

```
# /usr/share/ovirt-engine/setup/bin/ovirt-engine-rename
```

5. When prompted, press **Enter** to stop the engine service:

During execution engine service will be stopped (OK, Cancel) [OK]:

- When prompted, enter the new fully qualified domain name for the Manager:

New fully qualified server name: *[new name]*

Result

The **ovirt-engine-rename** command updates records of the fully qualified domain name of the Manager.

18.2. THE ENGINE CONFIGURATION TOOL

18.2.1. The Engine Configuration Tool

The engine configuration tool is a command-line utility for configuring global settings for your Red Hat Virtualization environment. The tool interacts with a list of key-value mappings that are stored in the engine database, and allows you to retrieve and set the value of individual keys, and retrieve a list of all available configuration keys and values. Furthermore, different values can be stored for each configuration level in your Red Hat Virtualization environment.



NOTE

Neither the Red Hat Virtualization Manager nor Red Hat JBoss Enterprise Application Platform need to be running to retrieve or set the value of a configuration key. Because the configuration key value-key mappings are stored in the engine database, they can be updated while the **postgresql** service is running. Changes are then applied when the **ovirt-engine** service is restarted.

18.2.2. Syntax for the engine-config Command

You can run the engine configuration tool from the machine on which the Red Hat Virtualization Manager is installed. For detailed information on usage, print the help output for the command:

```
# engine-config --help
```

Common tasks

List available configuration keys

```
# engine-config --list
```

List available configuration values

```
# engine-config --all
```

Retrieve value of configuration key

```
# engine-config --get [KEY_NAME]
```

Replace `[KEY_NAME]` with the name of the preferred key to retrieve the value for the given version of the key. Use the `--cver` parameter to specify the configuration version of the value to be retrieved. If no version is provided, values for all existing versions are returned.

Set value of configuration key

```
# engine-config --set [KEY_NAME]=[KEY_VALUE] --cver=[VERSION]
```

Replace `[KEY_NAME]` with the name of the specific key to set, and replace `[KEY_VALUE]` with the value to be set. You must specify the `[VERSION]` in environments with more than one configuration version.

Restart the ovirt-engine service to load changes

The `ovirt-engine` service needs to be restarted for your changes to take effect.

```
# systemctl restart ovirt-engine.service
```

18.3. THE IMAGE UPLOADER TOOL

18.3.1. The Image Uploader Tool



NOTE

The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disk images, and templates can then be uploaded from the imported storage domain to the attached data center. See [Section 8.6, “Importing Existing Storage Domains”](#) for information on importing storage domains.

The `engine-image-uploader` command allows you to list export storage domains and upload virtual machine images in OVF or OVA format to an export storage domain and have them automatically recognized in the Red Hat Virtualization Manager.

An OVA is a `tar` archive of the OVF files.



NOTE

The image uploader only supports gzip-compressed OVF files, or OVA files, created by Red Hat Virtualization.

The OVF contains images and master directories in the following format:

```
|-- images
|   |-- [Image Group UUID]
|       |-- [Image UUID (this is the disk image)]
|       |-- [Image UUID (this is the disk image)].meta
```

```
| -- master  
|   |-- vms  
|       |-- [UUID]  
|           |-- [UUID].ovf
```

18.3.2. Syntax for the engine-image-uploader Command

The basic syntax for the image uploader command is:

```
engine-image-uploader [options] list  
engine-image-uploader [options] upload [file].[file]...[file]
```

The image uploader command supports two actions - **list**, and **upload**.

- The **list** action lists the export storage domains to which images can be uploaded.
- The **upload** action uploads images to the specified export storage domain.

You must specify one of the above actions when you use the image uploader command. Moreover, you must specify at least one local file to use the **upload** action.

There are several parameters to further refine the **engine-image-uploader** command. You can set defaults for any of these parameters in the **/etc/ovirt-engine/imageuploader.conf** file.

General Options

-h, --help

Displays information on how to use the image uploader command.

--conf-file=[PATH]

Sets **[PATH]** as the configuration file the command will use. The default is **etc/ovirt-engine/imageuploader.conf**.

--log-file=[PATH]

Sets **[PATH]** as the specific file name the command will use to write log output. The default is **/var/log/ovirt-engine/ovirt-image-uploader/ovirt-image-uploader-[date].log**.

--cert-file=[PATH]

Sets **[PATH]** as the certificate for validating the engine. The default is **etc/pki/ovirt-engine/ca.pem**.

--insecure

Specifies that no attempt will be made to verify the engine.

--quiet

Sets quiet mode, reducing console output to a minimum.

-v, --verbose

Sets verbose mode, providing more console output.

-f, --force

Force mode is necessary when the source file being uploaded has the same file name as an existing file in the destination export domain. This option forces the existing file to be overwritten.

Red Hat Virtualization Manager Options**-u [USER], --user=[USER]**

Specifies the user whose credentials will be used to execute the command. The *[USER]* is specified in the format *[username]@[domain]*. The user must exist in the specified domain and be known to the Red Hat Virtualization Manager.

-r [FQDN], --engine=[FQDN]

Specifies the IP address or fully qualified domain name of the Red Hat Virtualization Manager from which the images will be uploaded. It is assumed that the image uploader is being run from the same machine on which the Red Hat Virtualization Manager is installed. The default value is **localhost:443**.

Export Storage Domain Options

The following options specify the export domain to which the images will be uploaded. These options cannot be used together; you must use either the **-e** option or the **-n** option.

-e [EXPORT_DOMAIN], --export-domain=[EXPORT_DOMAIN]

Sets the storage domain *EXPORT_DOMAIN* as the destination for uploads.

-n [NFSSERVER], --nfs-server=[NFSSERVER]

Sets the NFS path *[NFSSERVER]* as the destination for uploads.

Import Options

The following options allow you to customize which attributes of the images being uploaded are included when the image is uploaded to the export domain.

-i, --ovf-id

Specifies that the UUID of the image will not be updated. By default, the command generates a new UUID for images that are uploaded. This ensures there is no conflict between the id of the image being uploaded and the images already in the environment.

-d, --disk-instance-id

Specifies that the instance ID for each disk in the image will not be renamed. By default, the command generates new UUIDs for disks in images that are uploaded. This ensures there are no conflicts between the disks on the image being uploaded and the disks already in the environment.

-m, --mac-address

Specifies that network components in the image will not be removed from the image. By

default, the command removes network interface cards from image being uploaded to prevent conflicts with network cards on other virtual machines already in the environment. If you do not use this option, you can use the Administration Portal to add network interface cards to newly imported images and the Manager will ensure there are no MAC address conflicts.

-N [NEW_IMAGE_NAME], --name=[NEW_IMAGE_NAME]

Specifies a new name for the image being uploaded.

18.3.3. Creating an OVF Archive That is Compatible With the Image Uploader

Summary

You can create files that can be uploaded using the **engine-image-uploader** tool.

Procedure 18.2. Creating an OVF Archive That is Compatible With the Image Uploader

1. Use the Manager to create an empty export domain. An empty export domain makes it easy to see which directory contains your virtual machine.
2. Export your virtual machine to the empty export domain you just created.
3. Log in to the storage server that serves as the export domain, find the root of the NFS share and change to the subdirectory under that mount point. You started with a new export domain, there is only one directory under the exported directory. It contains the **images/** and **master/** directories.
4. Run the **tar -zcvf my.ovf images/ master/** command to create the tar/gzip OVF archive.
5. Anyone you give the resulting OVF file to (in this example, called **my.ovf**) can import it to Red Hat Virtualization Manager using the **engine-image-uploader** command.

Result

You have created a compressed OVF image file that can be distributed. Anyone you give it to can use the **engine-image-uploader** command to upload your image into their Red Hat Virtualization environment.

18.3.4. Basic engine-image-uploader Usage Examples

The following is an example of how to use the engine uploader command to list export storage domains:

Example 18.1. Listing export storage domains using the image uploader

```
# engine-image-uploader list
Please provide the REST API password for the admin@internal oVirt Engine
user (CTRL+D to abort):
Export Storage Domain Name | Datacenter | Export Domain Status
myexportdom                | Myowndc   | active
```

The following is an example of how to upload an Open Virtualization Format (OVF) file:

Example 18.2. Uploading a file using the image uploader

```
# engine-image-uploader -e myexportdom upload myrhel6.ovf
Please provide the REST API password for the admin@internal oVirt Engine
user (CTRL+D to abort):
```

18.4. THE USB FILTER EDITOR

18.4.1. Installing the USB Filter Editor

The USB Filter Editor is a Windows tool used to configure the **usbfilter.txt** policy file. The policy rules defined in this file allow or deny automatic pass-through of specific USB devices from client machines to virtual machines managed using the Red Hat Virtualization Manager. The policy file resides on the Red Hat Virtualization Manager in the following location:


```
/etc/ovirt-engine/usbfilter.txt
```

Changes to USB filter policies do not take effect unless the **ovirt-engine** service on the Red Hat Virtualization Manager server is restarted.

Download the **USBFilterEditor.msi** file from the Content Delivery Network:
<https://rhn.redhat.com/rhn/software/channel/downloads/Download.do?cid=20703>.

Procedure 18.3. Installing the USB Filter Editor

1. On a Windows machine, launch the **USBFilterEditor.msi** installer obtained from the Content Delivery Network.
2. Follow the steps of the installation wizard. Unless otherwise specified, the USB Filter Editor will be installed by default in either **C:\Program Files\RedHat\USB Filter Editor** or **C:\Program Files(x86)\RedHat\USB Filter Editor** depending on your version of Windows.
3. A USB Filter Editor shortcut icon is created on your desktop.



IMPORTANT

Use a Secure Copy (SCP) client to import and export filter policies from the Red Hat Virtualization Manager. A Secure Copy tool for Windows machines is WinSCP (<http://winscp.net>).

The default USB device policy provides virtual machines with basic access to USB devices; update the policy to allow the use of additional USB devices.

18.4.2. The USB Filter Editor Interface

- Double-click the USB Filter Editor shortcut icon on your desktop.

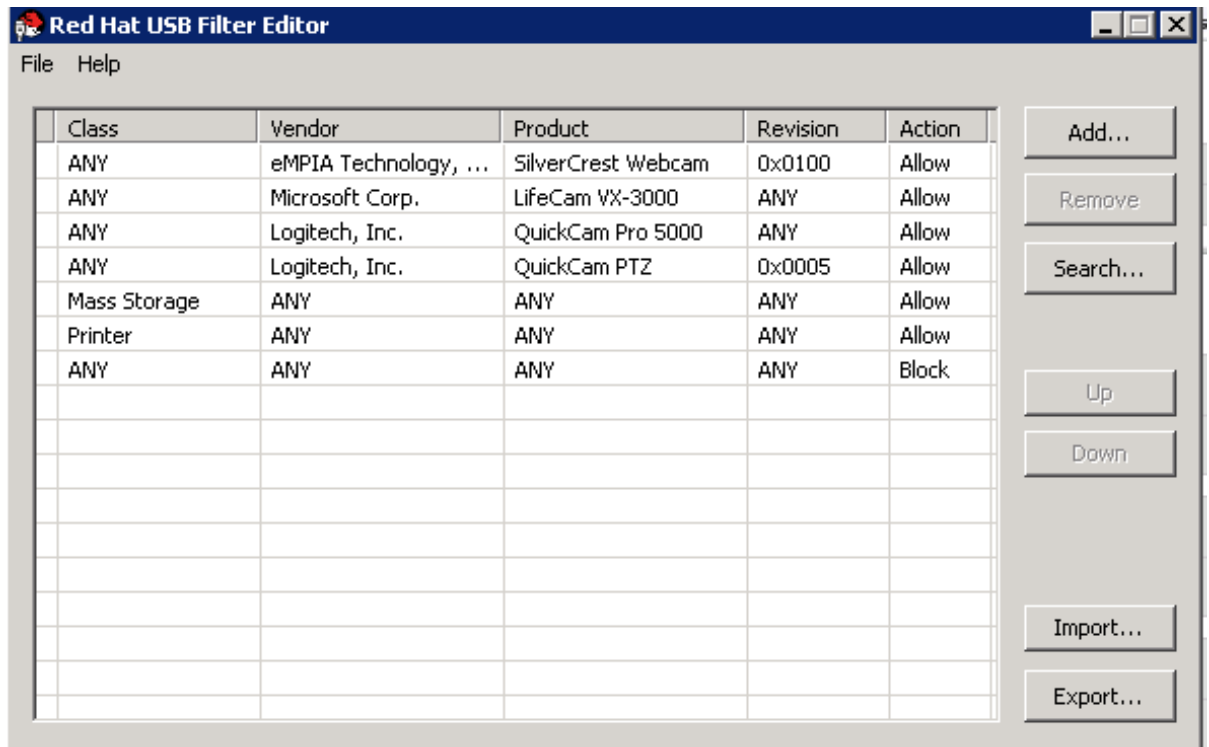


Figure 18.1. Red Hat USB Filter Editor

The **Red Hat USB Filter Editor** interface displays the **Class**, **Vendor**, **Product**, **Revision**, and **Action** for each USB device. Permitted USB devices are set to **Allow** in the **Action** column; prohibited devices are set to **Block**.

Table 18.1. USB Editor Fields

Name	Description
Class	Type of USB device; for example, printers, mass storage controllers.
Vendor	The manufacturer of the selected type of device.
Product	The specific USB device model.
Revision	The revision of the product.
Action	Allow or block the specified device.

The USB device policy rules are processed in their listed order. Use the **Up** and **Down** buttons to move rules higher or lower in the list. The universal **Block** rule needs to remain as the lowest entry to ensure all USB devices are denied unless explicitly allowed in the USB Filter Editor.

18.4.3. Adding a USB Policy

Summary

Add a USB policy to the USB Filter Editor.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

Procedure 18.4. Adding a USB Policy

1. Click the **Add** button. The **Edit USB Criteria** window opens:

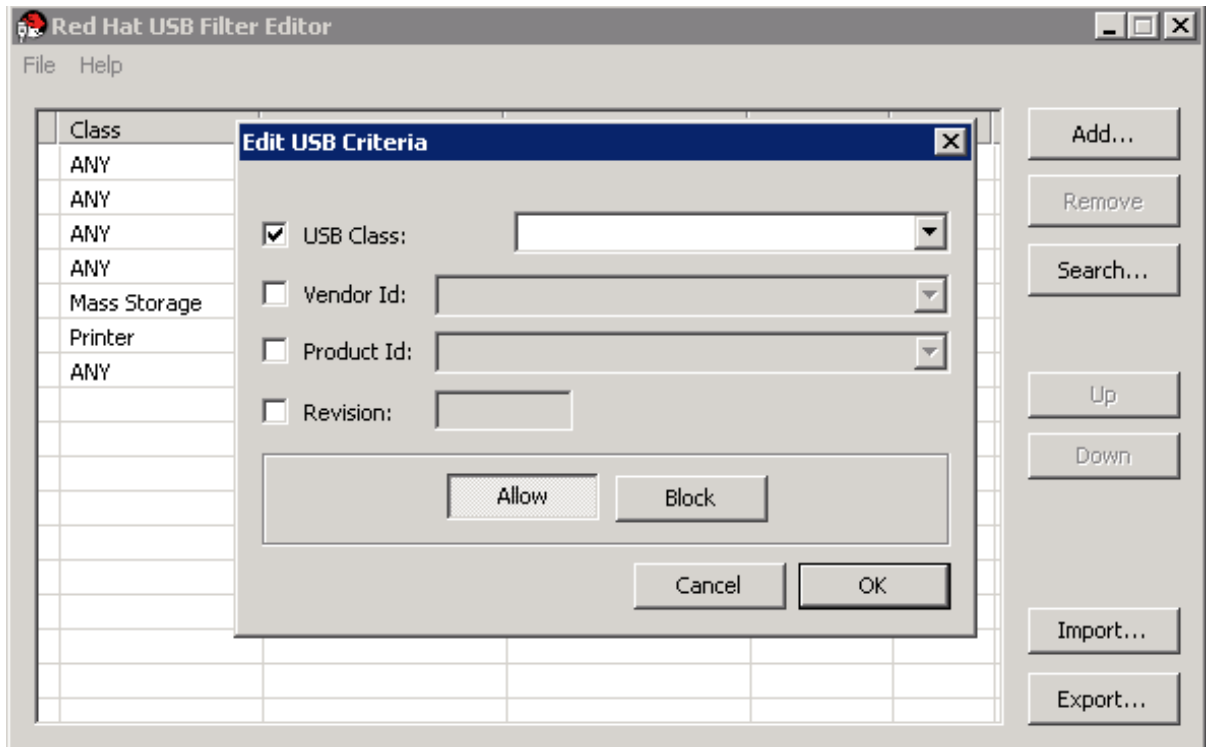


Figure 18.2. Edit USB Criteria

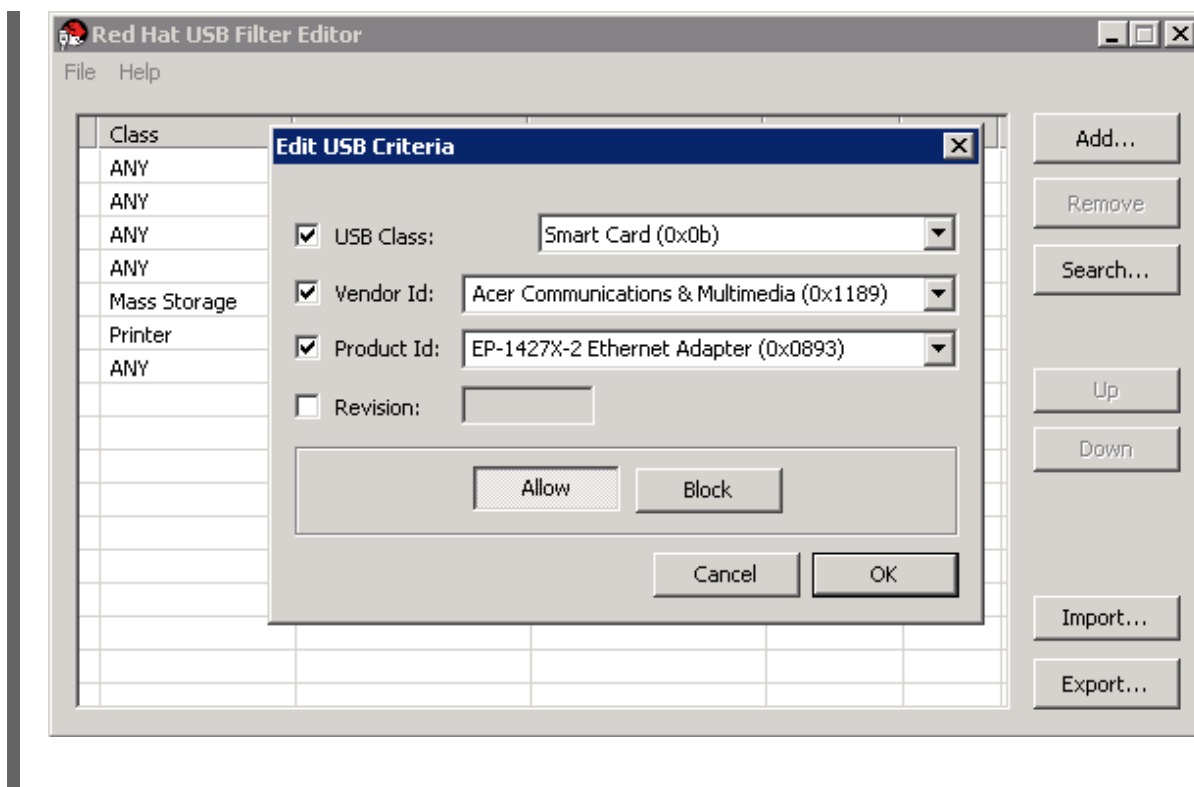
2. Use the **USB Class**, **Vendor ID**, **Product ID**, and **Revision** check boxes and lists to specify the device.

Click the **Allow** button to permit virtual machines use of the USB device; click the **Block** button to prohibit the USB device from virtual machines.

Click **OK** to add the selected filter rule to the list and close the window.

Example 18.3. Adding a Device

The following is an example of how to add USB Class **Smartcard**, device **EP-1427X-2 Ethernet Adapter**, from manufacturer **Acer Communications & Multimedia** to the list of allowed devices.



3. Click **File** → **Save** to save the changes.

Result

You have added a USB policy to the USB Filter Editor. USB filter policies need to be exported to the Red Hat Virtualization Manager to take effect.

18.4.4. Removing a USB Policy

Summary

Remove a USB policy from the USB Filter Editor.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

Procedure 18.5. Removing a USB Policy

1. Select the policy to be removed.




Figure 18.3. Select USB Policy

2. Click **Remove**. A message displays prompting you to confirm that you want to remove the policy.

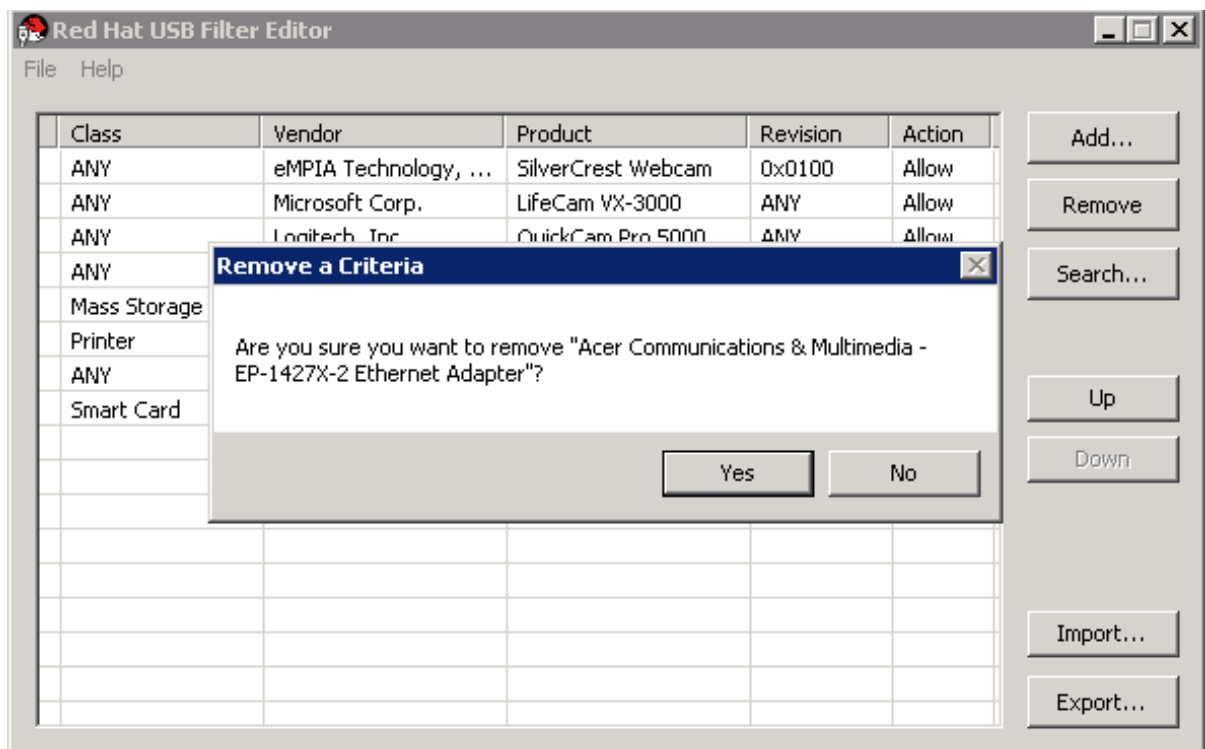


Figure 18.4. Edit USB Criteria

3. Click **Yes** to confirm that you want to remove the policy.
4. Click **File** → **Save** to save the changes.

Result

You have removed a USB policy from the USB Filter Editor. USB filter policies need to be exported to the Red Hat Virtualization Manager to take effect.

18.4.5. Searching for USB Device Policies

Summary

Search for attached USB devices to either allow or block them in the USB Filter Editor.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

Procedure 18.6. Searching for USB Device Policies

1. Click **Search**. The **Attached USB Devices** window displays a list of all the attached devices.

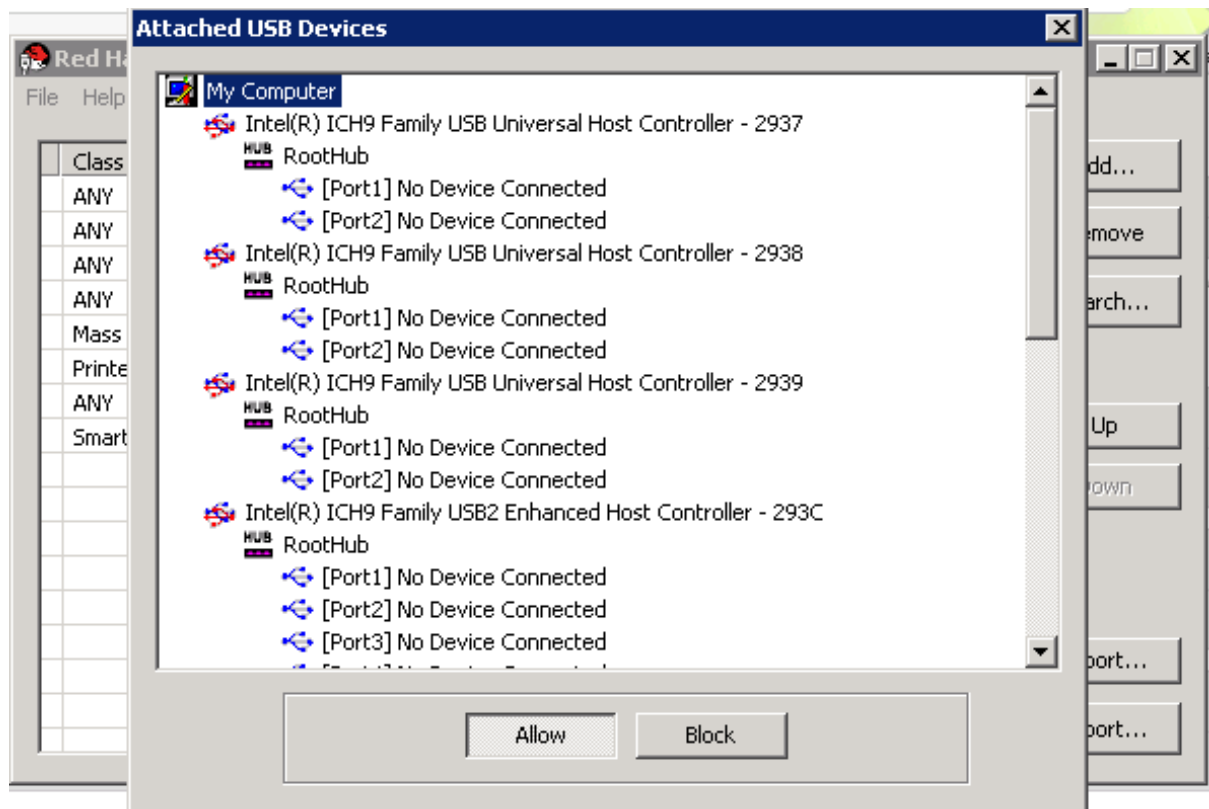


Figure 18.5. Attached USB Devices

2. Select the device and click **Allow** or **Block** as appropriate. Double-click the selected device to close the window. A policy rule for the device is added to the list.
3. Use the **Up** and **Down** buttons to change the position of the new policy rule in the list.
4. Click **File** → **Save** to save the changes.

Result

You have searched the attached USB devices. USB filter policies need to be exported to the Red Hat Virtualization Manager to take effect.

18.4.6. Exporting a USB Policy

Summary

USB device policy changes need to be exported and uploaded to the Red Hat Virtualization Manager for the updated policy to take effect. Upload the policy and restart the **ovirt-engine** service.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

Procedure 18.7. Exporting a USB Policy

1. Click **Export**; the **Save As** window opens.
2. Save the file with a file name of **usbfilter.txt**.
3. Using a Secure Copy client, such as WinSCP, upload the **usbfilter.txt** file to the server running Red Hat Virtualization Manager. The file must be placed in the following directory on the server:

```
/etc/ovirt-engine/
```

4. As the **root** user on the server running Red Hat Virtualization Manager, restart the **ovirt-engine** service.

```
# systemctl restart ovirt-engine.service
```

Result

The USB device policy will now be implemented on virtual machines running in the Red Hat Virtualization environment.

18.4.7. Importing a USB Policy

Summary

An existing USB device policy must be downloaded and imported into the USB Filter Editor before you can edit it.

Procedure 18.8. Importing a USB Policy

1. Using a Secure Copy client, such as WinSCP, upload the **usbfilter.txt** file to the server running Red Hat Virtualization Manager. The file must be placed in the following directory on the server:

```
/etc/ovirt-engine/
```

2. Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.
3. Click **Import** to open the **Open** window.
4. Open the **usbfilter.txt** file that was downloaded from the server.

Result

You are able to edit the USB device policy in the USB Filter Editor.

18.5. THE LOG COLLECTOR TOOL

18.5.1. Log Collector

A log collection tool is included in the Red Hat Virtualization Manager. This allows you to easily collect relevant logs from across the Red Hat Virtualization environment when requesting support.

The log collection command is **ovirt-log-collector**. You are required to log in as the **root** user and provide the administration credentials for the Red Hat Virtualization environment. The **ovirt-log-collector -h** command displays usage information, including a list of all valid options for the **ovirt-log-collector** command.

18.5.2. Syntax for the **ovirt-log-collector** Command

The basic syntax for the log collector command is:

```
ovirt-log-collector [options] list [all, clusters, datacenters]
ovirt-log-collector [options] collect
```

The two supported modes of operation are **list** and **collect**.

- The **list** parameter lists either the hosts, clusters, or data centers attached to the Red Hat Virtualization Manager. You are able to filter the log collection based on the listed objects.
- The **collect** parameter performs log collection from the Red Hat Virtualization Manager. The collected logs are placed in an archive file under the **/tmp/logcollector** directory. The **ovirt-log-collector** command assigns each log a specific file name.

Unless another parameter is specified, the default action is to list the available hosts together with the data center and cluster to which they belong. You will be prompted to enter user names and passwords to retrieve certain logs.

There are numerous parameters to further refine the **ovirt-log-collector** command.

General options

--version

Displays the version number of the command in use and returns to prompt.

-h, --help

Displays command usage information and returns to prompt.

--conf-file=PATH

Sets *PATH* as the configuration file the tool is to use.

--local-tmp=PATH

Sets *PATH* as the directory in which logs are saved. The default directory is **/tmp/logcollector**.

--ticket-number=TICKET

Sets *TICKET* as the ticket, or case number, to associate with the SOS report.

--upload=FTP_SERVER

Sets `FTP_SERVER` as the destination for retrieved logs to be sent using FTP. Do not use this option unless advised to by a Red Hat support representative.

--log-file=PATH

Sets `PATH` as the specific file name the command should use for the log output.

--quiet

Sets quiet mode, reducing console output to a minimum. Quiet mode is off by default.

-v, --verbose

Sets verbose mode, providing more console output. Verbose mode is off by default.

Red Hat Virtualization Manager Options

These options filter the log collection and specify authentication details for the Red Hat Virtualization Manager.

These parameters can be combined for specific commands. For example, `ovirt-log-collector --user=admin@internal --cluster ClusterA,ClusterB --hosts "SalesHost"*` specifies the user as `admin@internal` and limits the log collection to only `SalesHost` hosts in clusters `A` and `B`.

--no-hypervisors

Omits virtualization hosts from the log collection.

--one-hypervisor-per-cluster

Collects the logs of one host (the SPM, if there is one) from each cluster.

-u USER, --user=USER

Sets the user name for login. The `USER` is specified in the format `user@domain`, where `user` is the user name and `domain` is the directory services domain in use. The user must exist in directory services and be known to the Red Hat Virtualization Manager.

-r FQDN, --rhevm=FQDN

Sets the fully qualified domain name of the Red Hat Virtualization Manager server from which to collect logs, where `FQDN` is replaced by the fully qualified domain name of the Manager. It is assumed that the log collector is being run on the same local host as the Red Hat Virtualization Manager; the default value is `localhost`.

-c CLUSTER, --cluster=CLUSTER

Collects logs from the virtualization hosts in the nominated `CLUSTER` in addition to logs from the Red Hat Virtualization Manager. The cluster(s) for inclusion must be specified in a comma-separated list of cluster names or match patterns.

-d DATACENTER, --data-center=DATACENTER

Collects logs from the virtualization hosts in the nominated `DATACENTER` in addition to logs from the Red Hat Virtualization Manager. The data center(s) for inclusion must be specified in a comma-separated list of data center names or match patterns.

-H HOSTS_LIST, --hosts=HOSTS_LIST

Collects logs from the virtualization hosts in the nominated *HOSTS_LIST* in addition to logs from the Red Hat Virtualization Manager. The hosts for inclusion must be specified in a comma-separated list of host names, fully qualified domain names, or IP addresses. Match patterns are also valid.

SSH Configuration

--ssh-port=PORT

Sets *PORT* as the port to use for SSH connections with virtualization hosts.

-k KEYFILE, --key-file=KEYFILE

Sets *KEYFILE* as the public SSH key to be used for accessing the virtualization hosts.

--max-connections=MAX_CONNECTIONS

Sets *MAX_CONNECTIONS* as the maximum concurrent SSH connections for logs from virtualization hosts. The default is **10**.

PostgreSQL Database Options

The database user name and database name must be specified, using the *pg-user* and *dbname* parameters, if they have been changed from the default values.

Use the *pg-dbhost* parameter if the database is not on the local host. Use the optional *pg-host-key* parameter to collect remote logs. The PostgreSQL SOS plugin must be installed on the database server for remote log collection to be successful.

--no-postgresql

Disables collection of database. The log collector will connect to the Red Hat Virtualization Manager PostgreSQL database and include the data in the log report unless the *--no-postgresql* parameter is specified.

--pg-user=USER

Sets *USER* as the user name to use for connections with the database server. The default is **postgres**.

--pg-database=DATABASE

Sets *DATABASE* as the database name to use for connections with the database server. The default is **rhev**.

--pg-dbhost=DBHOST

Sets *DBHOST* as the host name for the database server. The default is **localhost**.

--pg-host-key=KEYFILE

Sets *KEYFILE* as the public identity file (private key) for the database server. This value is not set by default; it is required only where the database does not exist on the local host.

18.5.3. Basic Log Collector Usage

When the **ovirt-log-collector** command is run without specifying any additional parameters, its default behavior is to collect all logs from the Red Hat Virtualization Manager and its attached hosts. It will also collect database logs unless the **--no-postgresql** parameter is added. In the following example, log collector is run to collect all logs from the Red Hat Virtualization Manager and three attached hosts.

Example 18.4. Log Collector Usage

```
# ovirt-log-collector
INFO: Gathering oVirt Engine information...
INFO: Gathering PostgreSQL the oVirt Engine database and log files from
localhost...
Please provide REST API password for the admin@internal oVirt Engine
user (CTRL+D to abort):
About to collect information from 3 hypervisors. Continue? (Y/n):
INFO: Gathering information from selected hypervisors...
INFO: collecting information from 192.168.122.250
INFO: collecting information from 192.168.122.251
INFO: collecting information from 192.168.122.252
INFO: finished collecting information from 192.168.122.250
INFO: finished collecting information from 192.168.122.251
INFO: finished collecting information from 192.168.122.252
Creating compressed archive...
INFO Log files have been collected and placed in
/tmp/logcollector/sosreport-rhn-account-20110804121320-ce2a.tar.xz.
The MD5 for this file is 6d741b78925998caff29020df2b2ce2a and its size
is 26.7M
```

18.6. THE ISO UPLOADER TOOL

18.6.1. The ISO Uploader Tool

The ISO uploader is a tool for uploading ISO images to the ISO storage domain. It is installed as part of the Red Hat Virtualization Manager.

The ISO uploader command is **engine-iso-uploader**. You must log in as the **root** user and provide the administration credentials for the Red Hat Virtualization environment to use this command. The **engine-iso-uploader -h** command displays usage information, including a list of all valid options for the **engine-iso-uploader** command.

18.6.2. Syntax for the engine-iso-uploader Command

The basic syntax for the ISO uploader command is:

```
engine-iso-uploader [options] list
engine-iso-uploader [options] upload [file].[file]...[file]
```

The ISO uploader command supports two actions - **list**, and **upload**.

- The **list** action lists the ISO storage domains to which ISO files can be uploaded. The Red Hat Virtualization Manager creates this list on the machine on which the Manager is installed during the installation process.

- The **upload** action uploads a single ISO file or multiple ISO files separated by spaces to the specified ISO storage domain. NFS is used by default, but SSH is also available.

You must specify one of the above actions when you use the ISO uploader command. Moreover, you must specify at least one local file to use the **upload** action.

There are several parameters to further refine the **engine-iso-uploader** command.

General Options

--version

Displays the version of the ISO uploader command.

-h, --help

Displays information on how to use the ISO uploader command.

--conf-file=[PATH]

Sets *[PATH]* as the configuration file the command will use. The default is **/etc/ovirt-engine/isouploader.conf**.

--log-file=[PATH]

Sets *[PATH]* as the specific file name the command will use to write log output. The default is **/var/log/ovirt-engine/ovirt-iso-uploader/ovirt-iso-uploader[*date*].log**.

--cert-file=[PATH]

Sets *[PATH]* as the certificate for validating the engine. The default is **/etc/pki/ovirt-engine/ca.pem**.

--insecure

Specifies that no attempt will be made to verify the engine.

--nossl

Specifies that SSL will not be used to connect to the engine.

--quiet

Sets quiet mode, reducing console output to a minimum.

-v, --verbose

Sets verbose mode, providing more console output.

-f, --force

Force mode is necessary when the source file being uploaded has the same file name as an existing file in the destination ISO domain. This option forces the existing file to be overwritten.

Red Hat Virtualization Manager Options

-u [USER], --user=[USER]

Specifies the user whose credentials will be used to execute the command. The *[USER]* is specified in the format *[username]@[domain]*. The user must exist in the specified domain and be known to the Red Hat Virtualization Manager.

-r [FQDN], --engine=[FQDN]

Specifies the IP address or fully qualified domain name of the Red Hat Virtualization Manager from which the images will be uploaded. It is assumed that the image uploader is being run from the same machine on which the Red Hat Virtualization Manager is installed. The default value is **localhost:443**.

ISO Storage Domain Options

The following options specify the ISO domain to which the images will be uploaded. These options cannot be used together; you must use either the **-i** option or the **-n** option.

-i, --iso-domain=[ISODOMAIN]

Sets the storage domain *[ISODOMAIN]* as the destination for uploads.

-n, --nfs-server=[NFSSERVER]

Sets the NFS path *[NFSSERVER]* as the destination for uploads.

Connection Options

The ISO uploader uses NFS as default to upload files. These options specify SSH file transfer instead.

--ssh-user=[USER]

Sets *[USER]* as the SSH user name to use for the upload. The default is **root**.

--ssh-port=[PORT]

Sets *[PORT]* as the port to use when connecting to SSH.

-k [KEYFILE], --key-file=[KEYFILE]

Sets *[KEYFILE]* as the public key to use for SSH authentication. You will be prompted to enter the password of the user specified with **--ssh-user=[USER]** if no key is set.

18.6.3. Specifying an NFS Server**Example 18.5. Uploading to an NFS Server**

```
# engine-iso-uploader --nfs-server=storage.demo.redhat.com:/iso/path
upload RHEL6.0.iso
```

18.6.4. Basic ISO Uploader Usage

The example below demonstrates the ISO uploader and the list parameter. The first

command lists the available ISO storage domains; the **admin@internal** user is used because no user was specified in the command. The second command uploads an ISO file over NFS to the specified ISO domain.

Example 18.6. List Domains and Upload Image

```
# engine-iso-uploader list
Please provide the REST API password for the admin@internal oVirt Engine
user (CTRL+D to abort):
ISO Storage Domain Name | Datacenter          | ISO Domain Status
ISODomain                | Default             | active

# engine-iso-uploader --iso-domain=[ISODomain] upload [RHEL6.iso]
Please provide the REST API password for the admin@internal oVirt Engine
user (CTRL+D to abort):
```

18.6.5. Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain

The example below demonstrates the command to upload the **virtio-win.iso**, **virtio-win_x86.vfd**, **virtio-win_amd64.vfd**, and **rhev-tools-setup.iso** image files to the **ISODomain**.

Example 18.7. Uploading the VirtIO and Guest Tool Image Files

```
# engine-iso-uploader --iso-domain=[ISODomain] upload
/usr/share/virtio-win/virtio-win.iso /usr/share/virtio-win/virtio-
win_x86.vfd /usr/share/virtio-win/virtio-win_amd64.vfd /usr/share/rhev-
guest-tools-iso/rhev-tools-setup.iso
```

18.6.6. VirtIO and Guest Tool Image Files

The virtio-win ISO and Virtual Floppy Drive (VFD) images, which contain the VirtIO drivers for Windows virtual machines, and the rhev-tools-setup ISO, which contains the Red Hat Virtualization Guest Tools for Windows virtual machines, are copied to an ISO storage domain upon installation and configuration of the domain.

These image files provide software that can be installed on virtual machines to improve performance and usability. The most recent virtio-win and rhev-tools-setup files can be accessed via the following symbolic links on the file system of the Red Hat Virtualization Manager:

- **/usr/share/virtio-win/virtio-win.iso**
- **/usr/share/virtio-win/virtio-win_x86.vfd**
- **/usr/share/virtio-win/virtio-win_amd64.vfd**
- **/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso**

These image files must be manually uploaded to ISO storage domains that were not

created locally by the installation process. Use the **engine-iso-uploader** command to upload these images to your ISO storage domain. Once uploaded, the image files can be attached to and used by virtual machines.

18.7. THE ENGINE VACUUM TOOL

18.7.1. The Engine Vacuum Tool

The Engine Vacuum tool maintains PostgreSQL databases by updating tables and removing dead rows, allowing disk space to be reused. See the [PostgreSQL documentation](#) for information about the **VACUUM** command and its parameters.

The Engine Vacuum command is **engine-vacuum**. You are required to log in as the **root** user and provide the administration credentials for the Red Hat Virtualization environment.

Alternatively, the Engine Vacuum tool can be run while using the **engine-setup** command to customize an existing installation:

```
$ engine-setup
...
[ INFO ] Stage: Environment customization
...
Perform full vacuum on the engine database engine@localhost?
This operation may take a while depending on this setup health and the
configuration of the db vacuum process.
See https://www.postgresql.org/docs/9.2/static/sql-vacuum.html
(Yes, No) [No]:
```

The **Yes** option runs the Engine Vacuum tool in full vacuum verbose mode.

18.7.2. Engine Vacuum Modes

Engine Vacuum runs in two modes, Standard and Full.

Standard Vacuum

Frequent standard vacuuming is recommended.

Standard vacuum removes dead row versions in tables and indexes and marks the space as available for future reuse. Frequently updated tables should be vacuumed on a regular basis. However, standard vacuum does not return the space to the operating system.

Standard vacuum, with no parameters, processes every table in the current database.

Full Vacuum

Full vacuum is not recommended for routine use, but should only be run when a significant amount of space needs to be reclaimed from within the table.

Full vacuum compacts the tables by writing a new copy of the table file with no dead space, thereby enabling the operating system to reclaim the space. Full vacuum can take a long time.

Full vacuum requires extra disk space for the new copy of the table, until the operation completes and the old copy is deleted. Because full vacuum requires an exclusive lock on the table, it cannot be run in parallel with other uses of the table.

18.7.3. Syntax for the Engine Vacuum

The basic syntax for the Engine Vacuum command is:

```
engine-vacuum
```

```
engine-vacuum [option]
```

Running the **engine-vacuum** command with no options performs a standard vacuum.

There are several parameters to further refine the **engine-vacuum** command.

General Options

-h --help

Displays information on how to use the **engine-vacuum** command.

-a

Runs a standard vacuum, analyzes the database, and updates the optimizer statistics.

-A

Analyzes the database and updates the optimizer statistics, without vacuuming.

-f

Runs a full vacuum.

-v

Runs in verbose mode, providing more console output.

-t *[table_name]*

Vacuums a specific table or tables.

```
engine-vacuum -f -v -t vm_dynamic -t vds_dynamic
```

PART IV. GATHERING INFORMATION ABOUT THE ENVIRONMENT

CHAPTER 19. LOG FILES

19.1. RED HAT VIRTUALIZATION MANAGER INSTALLATION LOG FILES

Table 19.1. Installation

Log File	Description
<code>/var/log/ovirt-engine/engine-cleanup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the engine-cleanup command. This is the command used to reset a Red Hat Virtualization Manager installation. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist.
<code>/var/log/ovirt-engine/engine-db-install-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the engine-setup command detailing the creation and configuration of the rhev database.
<code>/var/log/ovirt-engine/ovirt-engine-dwh-setup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the ovirt-engine-dwh-setup command. This is the command used to create the ovirt_engine_history database for reporting. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
<code>/var/log/ovirt-engine/setup/ovirt-engine-setup-yyyymmddhhmmss.log</code>	Log from the engine-setup command. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.

19.2. RED HAT VIRTUALIZATION MANAGER LOG FILES

Table 19.2. Service Activity

Log File	Description
<code>/var/log/ovirt-engine/engine.log</code>	Reflects all Red Hat Virtualization Manager GUI crashes, Active Directory lookups, Database issues, and other events.
<code>/var/log/ovirt-engine/host-deploy</code>	Log files from hosts deployed from the Red Hat Virtualization Manager.
<code>/var/lib/ovirt-engine/setup-history.txt</code>	Tracks the installation and upgrade of packages associated with the Red Hat Virtualization Manager.

19.3. SPICE LOG FILES

SPICE log files are useful when troubleshooting SPICE connection issues. To start SPICE debugging, change the log level to **debugging**. Then, identify the log location.

Both the clients used to access the guest machines and the guest machines themselves have SPICE log files. For client side logs, if a SPICE client was launched using the native client, for which a **console.vv** file is downloaded, use the **remote-viewer** command to enable debugging and generate log output.

19.3.1. SPICE Logs for Hypervisor SPICE Servers

Table 19.3. SPICE Logs for Hypervisor SPICE Servers

Log Type	Log Location	To Change Log Level:
Host/Hypervisor SPICE Server	/var/log/libvirt/qemu/(guest_name).log	Run export SPICE_DEBUG_LEVEL=5 on the host/hypervisor prior to launching the guest. This variable is parsed by QEMU, and if run system-wide will print the debugging information of all virtual machines on the system. This command must be run on each host in the cluster. This command works only on a per-host/hypervisor basis, not a per-cluster basis.

19.3.2. SPICE Logs for Guest Machines

Table 19.4. spice-vdagent Logs for Guest Machines

Log Type	Log Location	To Change Log Level:
Windows Guest	C:\Windows\Temp\vdagent.log C:\Windows\Temp\vdservice.log	Not applicable

Log Type	Log Location	To Change Log Level:
Red Hat Enterprise Linux Guest	Use journalctl as the root user.	<p>To run the spice-vdagentd service in debug mode, as the root user create a /etc/sysconfig/spice-vdagentd file with this entry: SPICE_VDAGENTD_EXTRA_ARGS="-d -d"</p> <p>To run spice-vdagent in debug mode, from the command line:</p> <pre>\$ killall -u \$USER spice-vdagent \$ spice-vdagent -x -d [-d] [& tee spice-vdagent.log]</pre>

19.3.3. SPICE Logs for SPICE Clients Launched Using console.vv Files

For Linux client machines:

1. Enable SPICE debugging by running the **remote-viewer** command with the **--spice-debug** option. When prompted, enter the connection URL, for example, `spice://[virtual_machine_IP]:[port]`.

```
# remote-viewer --spice-debug
```

2. To run SPICE client with the debug parameter and to pass a .vv file to it, download the **console.vv** file and run the **remote-viewer** command with the **--spice-debug** option and specify the full path to the **console.vv** file.

```
# remote-viewer --spice-debug /path/to/console.vv
```

For Windows client machines:

1. In versions of **virt-viewer** 2.0-11.el7ev and higher, **virt-viewer.msi** installs **virt-viewer** and **debug-viewer.exe**.
2. Run the **remote-viewer** command with the **spice-debug** argument and direct the command at the path to the console:

```
remote-viewer --spice-debug path\to\console.vv
```

3. To view logs, connect to the virtual machine, and you will see a command prompt running GDB that prints standard output and standard error of **remote-viewer**.

19.4. RED HAT VIRTUALIZATION HOST LOG FILES

Table 19.5.

Log File	Description
<code>/var/log/vdsm/libvirt.log</code>	Log file for libvirt .
<code>/var/log/vdsm/spm-lock.log</code>	Log file detailing the host's ability to obtain a lease on the Storage Pool Manager role. The log details when the host has acquired, released, renewed, or failed to renew the lease.
<code>/var/log/vdsm/vdsm.log</code>	Log file for VDSM, the Manager's agent on the virtualization host(s).
<code>/tmp/ovirt-host-deploy-@DATE@.log</code>	Host deployment log, copied to engine as <code>/var/log/ovirt-engine/host-deploy/ovirt-@DATE@-@HOST@-@CORRELATION_ID@.log</code> after the host has been successfully deployed.

19.5. SETTING UP A VIRTUALIZATION HOST LOGGING SERVER

Hosts generate and update log files, recording their actions and problems. Collecting these log files centrally simplifies debugging.

This procedure should be used on your centralized log server. You could use a separate logging server, or use this procedure to enable host logging on the Red Hat Virtualization Manager.

Procedure 19.1. Setting up a Virtualization Host Logging Server

1. Configure SELinux to allow **rsyslog** traffic.

```
# semanage port -a -t syslogd_port_t -p udp 514
```

2. Edit `/etc/rsyslog.conf` and add the following lines:

```
$template TmplAuth, "/var/log/%fromhost%/secure"
$template TmplMsg, "/var/log/%fromhost%/messages"

$RuleSet remote
authpriv.* ?TmplAuth
*.info,mail.none;authpriv.none,cron.none ?TmplMsg
$RuleSet RSYSLOG_DefaultRuleset
$InputUDPServerBindRuleset remote
```

Uncomment the following:

```
#$ModLoad imudp
#$UDPServerRun 514
```

3. Restart the rsyslog service:

```
┃ # systemctl restart rsyslog.service
```

Your centralized log server is now configured to receive and store the **messages** and **secure** logs from your virtualization hosts.

CHAPTER 20. PROXIES

20.1. SPICE PROXY

20.1.1. SPICE Proxy Overview

The SPICE Proxy is a tool used to connect SPICE Clients to virtual machines when the SPICE Clients are outside the network that connects the hypervisors. Setting up a SPICE Proxy consists of installing **Squid** on a machine and configuring **iptables** to allow proxy traffic through the firewall. Turning a SPICE Proxy on consists of using **engine-config** on the Manager to set the key **SpiceProxyDefault** to a value consisting of the name and port of the proxy. Turning a SPICE Proxy off consists of using **engine-config** on the Manager to remove the value to which the key **SpiceProxyDefault** has been set.



IMPORTANT

The SPICE Proxy can only be used in conjunction with the standalone SPICE client, and cannot be used to connect to virtual machines using SPICE HTML5 or noVNC.

20.1.2. SPICE Proxy Machine Setup

This procedure explains how to set up a machine as a SPICE Proxy. A SPICE Proxy makes it possible to connect to the Red Hat Virtualization network from outside the network. We use **Squid** in this procedure to provide proxy services.

Procedure 20.1. Installing Squid on Red Hat Enterprise Linux

1. Install **Squid** on the Proxy machine:

```
# yum install squid
```

2. Open `/etc/squid/squid.conf`. Change:

```
http_access deny CONNECT !SSL_ports
```

to:

```
http_access deny CONNECT !Safe_ports
```

3. Start the proxy:

```
# systemctl start squid.service
```

4. Open the default squid port:

```
# iptables -A INPUT -p tcp --dport 3128 -j ACCEPT
```

5. Make this iptables rule persistent:

```
# service iptables save
```

You have now set up a machine as a SPICE proxy. Before connecting to the Red Hat Virtualization network from outside the network, activate the SPICE proxy.

20.1.3. Turning on SPICE Proxy

This procedure explains how to activate (or turn on) the SPICE proxy.

Procedure 20.2. Activating SPICE Proxy

1. On the Manager, use the engine-config tool to set a proxy:

```
# engine-config -s SpiceProxyDefault=someProxy
```

2. Restart the **ovirt-engine** service:

```
# systemctl restart ovirt-engine.service
```

The proxy must have this form:

```
protocol://[host]:[port]
```



NOTE

Only SPICE clients shipped with Red Hat Enterprise Linux 6.7, Red Hat Enterprise Linux 7.2, or newer support HTTPS proxies. Older clients only support HTTP. If HTTPS is specified for older clients, the client will ignore the proxy setting and attempt a direct connection to the host.

SPICE Proxy is now activated (turned on). It is now possible to connect to the Red Hat Virtualization network through the SPICE proxy.

20.1.4. Turning Off a SPICE Proxy

This procedure explains how to turn off (deactivate) a SPICE proxy.

Procedure 20.3. Turning Off a SPICE Proxy

1. Log in to the Manager:

```
$ ssh root@[IP of Manager]
```

2. Run the following command to clear the SPICE proxy:

```
# engine-config -s SpiceProxyDefault=""
```

3. Restart the Manager:

```
# systemctl restart ovirt-engine.service
```

SPICE proxy is now deactivated (turned off). It is no longer possible to connect to the Red Hat Virtualization network through the SPICE proxy.

20.2. SQUID PROXY

20.2.1. Installing and Configuring a Squid Proxy

Summary

This section explains how to install and configure a Squid proxy to the User Portal. A Squid proxy server is used as a content accelerator. It caches frequently-viewed content, reducing bandwidth and improving response times.

Procedure 20.4. Configuring a Squid Proxy

1. Obtain a keypair and certificate for the HTTPS port of the Squid proxy server. You can obtain this keypair the same way that you would obtain a keypair for another SSL/TLS service. The keypair is in the form of two PEM files which contain the private key and the signed certificate. For this procedure, we assume that they are named **proxy.key** and **proxy.cer**.



NOTE

The keypair and certificate can also be generated using the certificate authority of the engine. If you already have the private key and certificate for the proxy and do not want to generate it with the engine certificate authority, skip to the next step.

2. Choose a host name for the proxy. Then, choose the other components of the distinguished name of the certificate for the proxy.



NOTE

It is good practice to use the same country and same organization name used by the engine itself. Find this information by logging in to the machine where the Manager is installed and running the following command:

```
# openssl x509 -in /etc/pki/ovirt-engine/ca.pem -noout -subject
```

This command outputs something like this:

```
subject= /C=US/O=Example Inc./CN=engine.example.com.81108
```


The relevant part here is **/C=US/O=Example Inc..** Use this to build the complete distinguished name for the certificate for the proxy:

```
/C=US/O=Example Inc./CN=proxy.example.com
```

3. Log in to the proxy machine and generate a certificate signing request:



```
# openssl req -newkey rsa:2048 -subj '/C=US/O=Example  
Inc./CN=proxy.example.com' -nodes -keyout proxy.key -out proxy.req
```



IMPORTANT

You must include the quotes around the distinguished name for the certificate. The **-nodes** option ensures that the private key is not encrypted; this means that you do not need to enter the password to start the proxy server.

The command generates two files: **proxy.key** and **proxy.req**. **proxy.key** is the private key. Keep this file safe. **proxy.req** is the certificate signing request. **proxy.req** does not require any special protection.

- To generate the signed certificate, copy the certificate signing request file from the proxy machine to the Manager machine:

```
# scp proxy.req engine.example.com:/etc/pki/ovirt-engine/requests/.
```

- Log in to the Manager machine and sign the certificate:

```
# /usr/share/ovirt-engine/bin/pki-enroll-request.sh --name=proxy --  
days=3650 --subject='/C=US/O=Example Inc./CN=proxy.example.com'
```

This signs the certificate and makes it valid for 10 years (3650 days). Set the certificate to expire earlier, if you prefer.

- The generated certificate file is available in the directory **/etc/pki/ovirt-engine/certs** and should be named **proxy.cer**. On the proxy machine, copy this file from the Manager machine to your current directory:

```
# scp engine.example.com:/etc/pki/ovirt-engine/certs/proxy.cer .
```

- Ensure both **proxy.key** and **proxy.cer** are present on the proxy machine:

```
# ls -l proxy.key proxy.cer
```

- Install the Squid proxy server package on the proxy machine:

```
# yum install squid
```

- Move the private key and signed certificate to a place where the proxy can access them, for example to the **/etc/squid** directory:


```
# cp proxy.key proxy.cer /etc/squid/.
```

- Set permissions so that the **squid** user can read these files:

```
# chgrp squid /etc/squid/proxy.*  
# chmod 640 /etc/squid/proxy.*
```


11. The Squid proxy must verify the certificate used by the engine. Copy the Manager certificate to the proxy machine. This example uses the file path **/etc/squid/**:

```
# scp engine.example.com:/etc/pki/ovirt-engine/ca.pem /etc/squid/.
```



NOTE

The default CA certificate is located in **/etc/pki/ovirt-engine/ca.pem** on the Manager machine.

12. Set permissions so that the **squid** user can read the certificate file:

```
# chgrp squid /etc/squid/ca.pem
# chmod 640 /etc/squid/ca.pem
```

13. If SELinux is in enforcing mode, change the context of port 443 using the **semanage** tool to permit Squid to use port 443:

```
# yum install policycoreutils-python
# semanage port -m -p tcp -t http_cache_port_t 443
```

14. Replace the existing Squid configuration file with the following:


```
https_port 443 key=/etc/squid/proxy.key cert=/etc/squid/proxy.cer
ssl-bump defaultsite=engine.example.com
cache_peer engine.example.com parent 443 0 no-query originserver ssl
sslcafile=/etc/squid/ca.pem name=engine
cache_peer_access engine allow all
ssl_bump allow all
http_access allow all
```

15. Restart the Squid proxy server:

```
# systemctl restart squid.service
```

16. Connect to the User Portal using the complete URL, for instance:

```
https://proxy.example.com/UserPortal/org.ovirt.engine.ui.userportal.
UserPortal/UserPortal.html
```



NOTE

Shorter URLs, for example **https://proxy.example.com/UserPortal**, will not work. These shorter URLs are redirected to the long URL by the application server, using the 302 response code and the Location header. The version of **Squid** in Red Hat Enterprise Linux does not support rewriting these headers.

**NOTE**

Squid Proxy in the default configuration will terminate its connection after 15 idle minutes. To increase the amount of time before Squid Proxy terminates its idle connection, adjust the `read_timeout` option in `squid.conf` (for instance `read_timeout 10 hours`).

20.3. WEBSOCKET PROXY

20.3.1. Websocket Proxy Overview

The websocket proxy allows users to connect to virtual machines via noVNC and SPICE HTML5 consoles. Previously, the websocket proxy could only run on the Red Hat Virtualization Manager machine, but now the proxy can run on any machine that has access to the network.

The websocket proxy can be installed and configured on the Red Hat Virtualization Manager machine during the initial configuration (see [Configuring the Red Hat Virtualization Manager](#) in the *Installation Guide*), or on a separate machine (see [Installing a Websocket Proxy on a Separate Machine](#) in the *Installation Guide*).

The websocket proxy can also be migrated from the Manager machine to a separate machine. See [Section 20.3.2, “Migrating the Websocket Proxy to a Separate Machine”](#).

**NOTE**

SPICE HTML5 support is a Technology Preview feature. Technology Preview features are not fully supported under Red Hat Subscription Service Level Agreements (SLAs), may not be functionally complete, and are not intended for production use. However, these features provide early access to upcoming product innovations, enabling customers to test functionality and provide feedback during the development process.

20.3.2. Migrating the Websocket Proxy to a Separate Machine

For security or performance reasons the websocket proxy can run on a separate machine that does not run Red Hat Virtualization Manager. The procedure to migrate the websocket proxy from the Manager machine to a separate machine involves removing the websocket proxy configuration from the Manager machine, then installing the proxy on the separate machine.

The `engine-cleanup` command can be used to remove the websocket proxy from the Manager machine.

Procedure 20.5. Migrating the Websocket Proxy to a Separate Machine

1. On the Manager machine, run `engine-cleanup` to remove the required configuration.

```
# engine-cleanup
```

2. Type **No** when asked to remove all components and press **Enter**.

Do you want to remove all components? (Yes, No) [Yes]: No

3. Type **No** when asked to remove the engine and press **Enter**.

Do you want to remove the engine? (Yes, No) [Yes]: No

4. Type **Yes** when asked to remove the websocket proxy and press **Enter**.

Do you want to remove the WebSocket proxy? (Yes, No) [No]: Yes

Select **No** if asked to remove any other components.

5. Install and configure the proxy on the separate machine. See [Installing a WebSocket Proxy on a Separate Machine](#) in the *Installation Guide* for instructions.

APPENDIX A. VDSM AND HOOKS

A.1. VDSM

The VDSM service is used by the Red Hat Virtualization Manager to manage Red Hat Virtualization Hosts (RHVH) and Red Hat Enterprise Linux hosts. VDSM manages and monitors the host's storage, memory and network resources. It also co-ordinates virtual machine creation, statistics gathering, log collection and other host administration tasks. VDSM is run as a daemon on each host managed by Red Hat Virtualization Manager. It answers XML-RPC calls from clients. The Red Hat Virtualization Manager functions as a VDSM client.

A.2. VDSM HOOKS

VDSM is extensible via hooks. Hooks are scripts executed on the host when key events occur. When a supported event occurs VDSM runs any executable hook scripts in `/usr/libexec/vdsm/hooks/nn_event-name` on the host in alphanumeric order. By convention each hook script is assigned a two digit number, included at the front of the file name, to ensure that the order in which the scripts will be run in is clear. You are able to create hook scripts in any programming language, Python will however be used for the examples contained in this chapter.

Note that all scripts defined on the host for the event are executed. If you require that a given hook is only executed for a subset of the virtual machines which run on the host then you must ensure that the hook script itself handles this requirement by evaluating the **Custom Properties** associated with the virtual machine.



WARNING

VDSM hooks can interfere with the operation of Red Hat Virtualization. A bug in a VDSM hook has the potential to cause virtual machine crashes and loss of data. VDSM hooks should be implemented with caution and tested rigorously. The Hooks API is new and subject to significant change in the future.

A.3. EXTENDING VDSM WITH HOOKS

This chapter describes how to extend VDSM with event-driven hooks. Extending VDSM with hooks is an experimental technology, and this chapter is intended for experienced developers. By setting custom properties on virtual machines it is possible to pass additional parameters, specific to a given virtual machine, to the hook scripts.

A.4. SUPPORTED VDSM EVENTS

Table A.1. Supported VDSM Events

Name	Description
before_vm_start	Before virtual machine starts.
after_vm_start	After virtual machine starts.
before_vm_cont	Before virtual machine continues.
after_vm_cont	After virtual machine continues.
before_vm_pause	Before virtual machine pauses.
after_vm_pause	After virtual machine pauses.
before_vm_hibernate	Before virtual machine hibernates.
after_vm_hibernate	After virtual machine hibernates.
before_vm_dehibernate	Before virtual machine dehibernates.
after_vm_dehibernate	After virtual machine dehibernates.
before_vm_migrate_source	Before virtual machine migration, run on the source host from which the migration is occurring.
after_vm_migrate_source	After virtual machine migration, run on the source host from which the migration is occurring.
before_vm_migrate_destination	Before virtual machine migration, run on the destination host to which the migration is occurring.
after_vm_migrate_destination	After virtual machine migration, run on the destination host to which the migration is occurring.
after_vm_destroy	After virtual machine destruction.
before_vdsm_start	Before VDSM is started on the host. before_vdsm_start hooks are executed as the user root, and do not inherit the environment of the VDSM process.
after_vdsm_stop	After VDSM is stopped on the host. after_vdsm_stop hooks are executed as the user root, and do not inherit the environment of the VDSM process.

Name	Description
before_nic_hotplug	Before the NIC is hot plugged into the virtual machine.
after_nic_hotplug	After the NIC is hot plugged into the virtual machine.
before_nic_hotunplug	Before the NIC is hot unplugged from the virtual machine
after_nic_hotunplug	After the NIC is hot unplugged from the virtual machine.
after_nic_hotplug_fail	After hot plugging the NIC to the virtual machine fails.
after_nic_hotunplug_fail	After hot unplugging the NIC from the virtual machine fails.
before_disk_hotplug	Before the disk is hot plugged into the virtual machine.
after_disk_hotplug	After the disk is hot plugged into the virtual machine.
before_disk_hotunplug	Before the disk is hot unplugged from the virtual machine
after_disk_hotunplug	After the disk is hot unplugged from the virtual machine.
after_disk_hotplug_fail	After hot plugging the disk to the virtual machine fails.
after_disk_hotunplug_fail	After hot unplugging the disk from the virtual machine fails.
before_device_create	Before creating a device that supports custom properties.
after_device_create	After creating a device that supports custom properties.
before_update_device	Before updating a device that supports custom properties.
after_update_device	After updating a device that supports custom properties.

Name	Description
before_device_destroy	Before destroying a device that supports custom properties.
after_device_destroy	After destroying a device that supports custom properties.
before_device_migrate_destination	Before device migration, run on the destination host to which the migration is occurring.
after_device_migrate_destination	After device migration, run on the destination host to which the migration is occurring.
before_device_migrate_source	Before device migration, run on the source host from which the migration is occurring.
after_device_migrate_source	After device migration, run on the source host from which the migration is occurring.
after_network_setup	After setting up the network when starting a host machine.
before_network_setup	Before setting up the network when starting a host machine.

A.5. THE VDSM HOOK ENVIRONMENT

Most hook scripts are run as the **vds**m user and inherit the environment of the VDSM process. The exceptions are hook scripts triggered by the **before_vdsm_start** and **after_vdsm_stop** events. Hook scripts triggered by these events run as the **root** user and do not inherit the environment of the VDSM process.

A.6. THE VDSM HOOK DOMAIN XML OBJECT

When hook scripts are started, the **_hook_domxml** variable is appended to the environment. This variable contains the path of the libvirt domain XML representation of the relevant virtual machine. Several hooks are an exception to this rule, as outlined below.

The **_hook_domxml** variable of the following hooks contains the XML representation of the NIC and not the virtual machine.

- ***_nic_hotplug_***
- ***_nic_hotunplug_***
- ***_update_device**
- ***_device_create**
- ***_device_migrate_***



IMPORTANT

The `before_migration_destination` and `before_dehibernation` hooks currently receive the XML of the domain from the source host. The XML of the domain at the destination will have various differences.

The libvirt domain XML format is used by VDSM to define virtual machines. Details on the libvirt domain XML format can be found at <http://libvirt.org/formatdomain.html>. The UUID of the virtual machine may be deduced from the domain XML, but it is also available as the environment variable `vmId`.

A.7. DEFINING CUSTOM PROPERTIES

The custom properties that are accepted by the Red Hat Virtualization Manager - and in turn passed to custom hooks - are defined using the `engine-config` command. Run this command as the `root` user on the host where Red Hat Virtualization Manager is installed.

The `UserDefinedVMProperties` and `CustomDeviceProperties` configuration keys are used to store the names of the custom properties supported. Regular expressions defining the valid values for each named custom property are also contained in these configuration keys.

Multiple custom properties are separated by a semi-colon. Note that when setting the configuration key, any existing value it contained is overwritten. When combining new and existing custom properties, all of the custom properties in the command used to set the key's value must be included.

Once the configuration key has been updated, the `ovirt-engine` service must be restarted for the new values to take effect.

Example A.1. Virtual Machine Properties - Defining the *smartcard* Custom Property

1. Check the existing custom properties defined by the `UserDefinedVMProperties` configuration key using the following command:

```
# engine-config -g UserDefinedVMProperties
```

As shown by the output below, the custom property `memory` is already defined. The regular expression `^[0-9]+$` ensures that the custom property will only ever contain numeric characters.

```
# engine-config -g UserDefinedVMProperties
UserDefinedVMProperties:  version: 3.6
UserDefinedVMProperties:  version: 4.0
UserDefinedVMProperties : memory=^[0-9]+$ version: 4.0
```

2. Because the `memory` custom property is already defined in the `UserDefinedVMProperties` configuration key, the new custom property must be appended to it. The additional custom property, `smartcard`, is added to the configuration key's value. The new custom property is able to hold a value of `true` or `false`.


```
# engine-config -s UserDefinedVMProperties='memory=^[0-9]+$;smartcard=^(true|false)$' --cver=4.0
```

3. Verify that the custom properties defined by the **UserDefinedVMProperties** configuration key have been updated correctly.

```
# engine-config -g UserDefinedVMProperties
UserDefinedVMProperties:  version: 3.6
UserDefinedVMProperties:  version: 4.0
UserDefinedVMProperties : memory=^[0-9]+$;smartcard=^(true|false)$
version: 4.0
```

4. Finally, the **ovirt-engine** service must be restarted for the configuration change to take effect.

```
# systemctl restart ovirt-engine.service
```

Example A.2. Device Properties - Defining the *interface* Custom Property

1. Check the existing custom properties defined by the **CustomDeviceProperties** configuration key using the following command:

```
# engine-config -g CustomDeviceProperties
```

As shown by the output below, no custom properties have yet been defined.

```
# engine-config -g CustomDeviceProperties
CustomDeviceProperties:  version: 3.6
CustomDeviceProperties:  version: 4.0
```

2. The **interface** custom property does not already exist, so it can be appended as is. In this example, the value of the **speed** sub-property is set to a range of 0 to 99999, and the value of the **duplex** sub-property is set to a selection of either **full** or **half**.

```
# engine-config -s CustomDeviceProperties="{type=interface;prop={speed=^[0-9]{1,5}};$;duplex=^(full|half)$}" --cver=4.0
```

3. Verify that the custom properties defined by the **CustomDeviceProperties** configuration key have been updated correctly.

```
# engine-config -g CustomDeviceProperties
UserDefinedVMProperties:  version: 3.6
UserDefinedVMProperties:  version: 4.0
UserDefinedVMProperties : {type=interface;prop={speed=^[0-9]{1,5}};$;duplex=^(full|half)$} version: 4.0
```

4. Finally, the **ovirt-engine** service must be restarted for the configuration change to take effect.

```
# systemctl restart ovirt-engine.service
```

A.8. SETTING VIRTUAL MACHINE CUSTOM PROPERTIES

Once custom properties are defined in the Red Hat Virtualization Manager, you can begin setting them on virtual machines. Custom properties are set on the **Custom Properties** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows in the Administration Portal.

You can also set custom properties from the **Run Virtual Machine(s)** dialog box. Custom properties set from the **Run Virtual Machine(s)** dialog box will only apply to the virtual machine until it is next shutdown.

The **Custom Properties** tab provides a facility for you to select from the list of defined custom properties. Once you select a custom property key an additional field will display allowing you to enter a value for that key. Add additional key/value pairs by clicking the + button and remove them by clicking the - button.

A.9. EVALUATING VIRTUAL MACHINE CUSTOM PROPERTIES IN A VDSM HOOK

Each key set in the **Custom Properties** field for a virtual machine is appended as an environment variable when calling hook scripts. Although the regular expressions used to validate the **Custom Properties** field provide some protection you should ensure that your scripts also validate that the inputs provided match their expectations.

Example A.3. Evaluating Custom Properties

This short Python example checks for the existence of the custom property **key1**. If the custom property is set then the value is printed to standard error. If the custom property is not set then no action is taken.

```
#!/usr/bin/python

import os
import sys

if os.environ.has_key('key1'):
    sys.stderr.write('key1 value was : %s\n' % os.environ['key1'])
else:
    sys.exit(0)
```

A.10. USING THE VDSM HOOKING MODULE

VDSM ships with a Python hooking module, providing helper functions for VDSM hook scripts. This module is provided as an example, and is only relevant to VDSM hooks written in Python.

The hooking module supports reading of a virtual machine's libvirt XML into a DOM object. Hook scripts can then use Python's built in **xml.dom** library (<http://docs.python.org/release/2.6/library/xml.dom.html>) to manipulate the object.

The modified object can then be saved back to libvirt XML using the hooking module. The hooking module provides the following functions to support hook development:

Table A.2. Hooking module functions

Name	Argument	Description
<code>tobool</code>	string	Converts a string "true" or "false" to a Boolean value
<code>read_domxml</code>	-	Reads the virtual machine's libvirt XML into a DOM object
<code>write_domxml</code>	DOM object	Writes the virtual machine's libvirt XML from a DOM object

A.11. VDSM HOOK EXECUTION

`before_vm_start` scripts can edit the domain XML in order to change VDSM's definition of a virtual machine before it reaches libvirt. Caution must be exercised in doing so. Hook scripts have the potential to disrupt the operation of VDSM, and buggy scripts can result in outages to the Red Hat Virtualization environment. In particular, ensure you never change the UUID of the domain, and do not attempt to remove a device from the domain without sufficient background knowledge.

Both `before_vdsm_start` and `after_vdsm_stop` hook scripts are run as the `root` user. Other hook scripts that require `root` access to the system must be written to use the `sudo` command for privilege escalation. To support this the `/etc/sudoers` must be updated to allow the `vdsm` user to use `sudo` without reentering a password. This is required as hook scripts are executed non-interactively.

Example A.4. Configuring sudo for VDSM Hooks

In this example the `sudo` command will be configured to allow the `vdsm` user to run the `/bin/chown` command as `root`.

1. Log into the virtualization host as `root`.
2. Open the `/etc/sudoers` file in a text editor.
3. Add this line to the file:

```
vdsm ALL=(ALL) NOPASSWD: /bin/chown
```

This specifies that the `vdsm` user has the ability to run the `/bin/chown` command as the `root` user. The `NOPASSWD` parameter indicates that the user will not be prompted to enter their password when calling `sudo`.

Once this configuration change has been made VDSM hooks are able to use the `sudo` command to run `/bin/chown` as `root`. This Python code uses `sudo` to execute `/bin/chown` as `root` on the file `/my_file`.

```
retcode = subprocess.call( ["/usr/bin/sudo", "/bin/chown", "root",
"/my_file"] )
```

The standard error stream of hook scripts is collected in VDSM's log. This information is used to debug hook scripts.

A.12. VDSM HOOK RETURN CODES

Hook scripts must return one of the return codes shown in [Table A.3, “Hook Return Codes”](#). The return code will determine whether further hook scripts are processed by VDSM.

Table A.3. Hook Return Codes

Code	Description
0	The hook script ended successfully
1	The hook script failed, other hooks should be processed
2	The hook script failed, no further hooks should be processed
>2	Reserved

A.13. VDSM HOOK EXAMPLES

The example hook scripts provided in this section are strictly not supported by Red Hat. You must ensure that any and all hook scripts that you install to your system, regardless of source, are thoroughly tested for your environment.

Example A.5. NUMA Node Tuning

Purpose:

This hook script allows for tuning the allocation of memory on a NUMA host based on the **numaset** custom property. Where the custom property is not set no action is taken.

Configuration String:

```
numaset=^(interleave|strict|preferred):[\^]?\d+(-\d+)?(,[\^]?\d+(-\d+)?)*$
```

The regular expression used allows the **numaset** custom property for a given virtual machine to specify both the allocation mode (**interleave**, **strict**, **preferred**) and the node to use. The two values are separated by a colon (:). The regular expression allows specification of the **nodeset** as:

- that a specific node (**numaset=strict:1**, specifies that only node 1 be used), or

- that a range of nodes be used (**numaset=strict:1-4**, specifies that nodes 1 through 4 be used), or
- that a specific node not be used (**numaset=strict:^3**, specifies that node 3 not be used), or
- any comma-separated combination of the above (**numaset=strict:1-4,6**, specifies that nodes 1 to 4, and 6 be used).

Script:

```
/usr/libexec/vdsm/hooks/before_vm_start/50_numa
```

```
#!/usr/bin/python

import os
import sys
import hooking
import traceback

...
numa hook
=====
add numa support for domain xml:

<numatune>
  <memory mode="strict" nodeset="1-4,^3" />
</numatune>

memory=interleave|strict|preferred

numaset="1" (use one NUMA node)
numaset="1-4" (use 1-4 NUMA nodes)
numaset="^3" (don't use NUMA node 3)
numaset="1-4,^3,6" (or combinations)

syntax:
  numa=strict:1-4
  ...

if os.environ.has_key('numa'):
    try:
        mode, nodeset = os.environ['numa'].split(':')

        domxml = hooking.read_domxml()

        domain = domxml.getElementsByTagName('domain')[0]
        numas = domxml.getElementsByTagName('numatune')

        if not len(numas) > 0:
            numatune = domxml.createElement('numatune')
            domain.appendChild(numatune)

            memory = domxml.createElement('memory')
            memory.setAttribute('mode', mode)
            memory.setAttribute('nodeset', nodeset)
            numatune.appendChild(memory)
```

```
        hooking.write_domxml(domxml)
    else:
        sys.stderr.write('numa: numa already exists in domain xml')
        sys.exit(2)
    except:
        sys.stderr.write('numa: [unexpected error]: %s\n' %
            traceback.format_exc())
        sys.exit(2)
```

APPENDIX B. CUSTOM NETWORK PROPERTIES

B.1. EXPLANATION OF BRIDGE_OPTS PARAMETERS

Table B.1. bridge_opts parameters

Parameter	Description
forward_delay	Sets the time, in deciseconds, a bridge will spend in the listening and learning states. If no switching loop is discovered in this time, the bridge will enter forwarding state. This allows time to inspect the traffic and layout of the network before normal network operation.
gc_timer	Sets the garbage collection time, in deciseconds, after which the forwarding database is checked and cleared of timed-out entries.
group_addr	Set to zero when sending a general query. Set to the IP multicast address when sending a group-specific query, or group-and-source-specific query.
group_fwd_mask	Enables bridge to forward link local group addresses. Changing this value from the default will allow non-standard bridging behavior.
hash_elasticity	The maximum chain length permitted in the hash table. Does not take effect until the next new multicast group is added. If this cannot be satisfied after rehashing, a hash collision occurs and snooping is disabled.
hash_max	The maximum amount of buckets in the hash table. This takes effect immediately and cannot be set to a value less than the current number of multicast group entries. Value must be a power of two.
hello_time	Sets the time interval, in deciseconds, between sending 'hello' messages, announcing bridge position in the network topology. Applies only if this bridge is the Spanning Tree root bridge.
hello_timer	Time, in deciseconds, since last 'hello' message was sent.

Parameter	Description
max_age	Sets the maximum time, in deciseconds, to receive a 'hello' message from another root bridge before that bridge is considered dead and takeover begins.
multicast_last_member_count	Sets the number of 'last member' queries sent to the multicast group after receiving a 'leave group' message from a host.
multicast_last_member_interval	Sets the time, in deciseconds, between 'last member' queries.
multicast_membership_interval	Sets the time, in deciseconds, that a bridge will wait to hear from a member of a multicast group before it stops sending multicast traffic to the host.
multicast_querier	Sets whether the bridge actively runs a multicast querier or not. When a bridge receives a 'multicast host membership' query from another network host, that host is tracked based on the time that the query was received plus the multicast query interval time. If the bridge later attempts to forward traffic for that multicast membership, or is communicating with a querying multicast router, this timer confirms the validity of the querier. If valid, the multicast traffic is delivered via the bridge's existing multicast membership table; if no longer valid, the traffic is sent via all bridge ports. Broadcast domains with, or expecting, multicast memberships should run at least one multicast querier for improved performance.
multicast_querier_interval	Sets the maximum time, in deciseconds, between last 'multicast host membership' query received from a host to ensure it is still valid.
multicast_query_use_ifaddr	Boolean. Defaults to '0', in which case the querier uses 0.0.0.0 as source address for IPv4 messages. Changing this sets the bridge IP as the source address.

Parameter	Description
multicast_query_interval	Sets the time, in deciseconds, between query messages sent by the bridge to ensure validity of multicast memberships. At this time, or if the bridge is asked to send a multicast query for that membership, the bridge checks its own multicast querier state based on the time that a check was requested plus <code>multicast_query_interval</code> . If a multicast query for this membership has been sent within the last <code>multicast_query_interval</code> , it is not sent again.
multicast_query_response_interval	Length of time, in deciseconds, a host is allowed to respond to a query once it has been sent. Must be less than or equal to the value of the <code>multicast_query_interval</code> .
multicast_router	Allows you to enable or disable ports as having multicast routers attached. A port with one or more multicast routers will receive all multicast traffic. A value of 0 disables completely, a value of 1 enables the system to automatically detect the presence of routers based on queries, and a value of 2 enables ports to always receive all multicast traffic.
multicast_snooping	Toggles whether snooping is enabled or disabled. Snooping allows the bridge to listen to the network traffic between routers and hosts to maintain a map to filter multicast traffic to the appropriate links. This option allows the user to re-enable snooping if it was automatically disabled due to hash collisions, however snooping will not be re-enabled if the hash collision has not been resolved.
multicast_startup_query_count	Sets the number of queries sent out at startup to determine membership information.
multicast_startup_query_interval	Sets the time, in deciseconds, between queries sent out at startup to determine membership information.

B.2. HOW TO SET UP RED HAT VIRTUALIZATION MANAGER TO USE ETHTOOL

You can configure `ethtool` properties for host network interface cards from the Administration Portal. The `ethtool_opts` key is not available by default and needs to be added to the Manager using the engine configuration tool. You also need to install the required VDSM hook package on the hosts.

Procedure B.1. Adding the `ethtool_opts` Key to the Manager

1. On the Manager, run the following command to add the key:

```
# engine-config -s
UserDefinedNetworkCustomProperties=ethtool_opts=. * --cver=4.0
```

2. Restart the **ovirt-engine** service:

```
# systemctl restart ovirt-engine.service
```

3. On the hosts that you want to configure `ethtool` properties, install the VDSM hook package. The package is available by default on Red Hat Virtualization Host but needs to be installed on Red Hat Enterprise Linux hosts.

```
# yum install vdsmd-hook-ethtool-options
```

The **ethtool_opts** key is now available in the Administration Portal. See [Section 6.5.2, “Editing Host Network Interfaces and Assigning Logical Networks to Hosts”](#) to apply `ethtool` properties to logical networks.

B.3. HOW TO SET UP RED HAT VIRTUALIZATION MANAGER TO USE FCOE

You can configure Fibre Channel over Ethernet (FCoE) properties for host network interface cards from the Administration Portal. The **fcoe** key is not available by default and needs to be added to the Manager using the engine configuration tool. You can check whether **fcoe** has already been enabled by running the following command:

```
# engine-config -g UserDefinedNetworkCustomProperties
```

You also need to install the required VDSM hook package on the hosts. Depending on the FCoE card on the hosts, special configuration may also be needed; see [Configuring a Fibre Channel over Ethernet Interface](#) in the *Red Hat Enterprise Linux Storage Administration Guide*.

Procedure B.2. Adding the `fcoe` Key to the Manager

1. On the Manager, run the following command to add the key:

```
# engine-config -s
UserDefinedNetworkCustomProperties='fcoe=^((enable|dcb|auto_vlan)=
(yes|no),?)*$'
```

2. Restart the **ovirt-engine** service:

```
# systemctl restart ovirt-engine.service
```

3. Install the VDSM hook package on each of the Red Hat Enterprise Linux hosts on which you want to configure FCoE properties. The package is available by default on Red Hat Virtualization Host (RHVH).

```
# yum install vdsm-hook-fcoe
```

The **fcoe** key is now available in the Administration Portal. See [Section 6.5.2, “Editing Host Network Interfaces and Assigning Logical Networks to Hosts”](#) to apply FCoE properties to logical networks.

APPENDIX C. RED HAT VIRTUALIZATION USER INTERFACE PLUGINS

C.1. RED HAT VIRTUALIZATION USER INTERFACE PLUG-INS

Red Hat Virtualization supports plug-ins that expose non-standard features. This makes it easier to use the Red Hat Virtualization Administration Portal to integrate with other systems. Each interface plug-in represents a set of user interface extensions that can be packaged and distributed for use with Red Hat Virtualization.

Red Hat Virtualization's User Interface plug-ins integrate with the Administration Portal directly on the client using the JavaScript programming language. Plug-ins are invoked by the Administration Portal and executed in the web browser's JavaScript runtime. User Interface plug-ins can use the JavaScript language and its libraries.

At key events during runtime, the Administration Portal invokes individual plug-ins via event handler functions representing Administration-Portal-to-plug-in communication. Although the Administration Portal supports multiple event-handler functions, a plug-in declares functions which are of interest only to its implementation. Each plug-in must register relevant event handler functions as part of the plug-in bootstrap sequence before the plug-in is put to use by the administration portal.

To facilitate the plug-in-to-Administration-Portal communication that drives the User Interface extension, the Administration Portal exposes the plug-in API as a global (top-level) `pluginApi` JavaScript object that individual plug-ins can consume. Each plug-in obtains a separate `pluginApi` instance, allowing the Administration Portal to control plug-in API-function invocation for each plug-in with respect to the plug-in's life cycle.

C.2. RED HAT VIRTUALIZATION USER INTERFACE PLUGIN LIFECYCLE

C.2.1. Red Hat Virtualization User Interface Plug-in Life cycle

The basic life cycle of a User Interface Plug-in divides into three stages:

1. Plug-in discovery.
2. Plug-in loading.
3. Plug-in bootstrapping.

C.2.2. Red Hat Virtualization User Interface Plug-in Discovery

Creating plug-in descriptors is the first step in the plug-in discovery process. Plug-in descriptors contain important plug-in metadata and optional default plug-in-specific configurations.

As part of handling administration portal HTML page requests (**HTTP GET**), User Interface plug-in infrastructure attempts to discover and load plug-in descriptors from your local file system. For each plug-in descriptor, the infrastructure also attempts to load corresponding plug-in user configurations used to override default plug-in-specific configurations (if any exist) and tweak plug-in runtime behavior. Plug-in user configuration is optional. After

loading descriptors and corresponding user configuration files, oVirt Engine aggregates User Interface plug-in data and embeds it into the administration portal HTML page for runtime evaluation.

By default, plug-in descriptors reside in `$ENGINE_USR/ui-plug-ins`, with a default mapping of `ENGINE_USR=/usr/share/ovirt-engine` as defined by oVirt Engine local configuration. Plug-in descriptors are expected to comply with JSON format specifications, but plug-in descriptors allow Java/C++ style comments (of both `/*` and `//` varieties) in addition to the JSON format specifications.

By default, plug-in user configuration files reside in `$ENGINE_ETC/ui-plug-ins`, with a default mapping of `ENGINE_ETC=/etc/ovirt-engine` as defined by oVirt Engine local configuration. Plug-in user configuration files are expected to comply with same content format rules as plug-in descriptors.



NOTE

Plug-in user configuration files generally follow the `<descriptorFileName>-config.json` naming convention.

C.2.3. Red Hat Virtualization User Interface Plug-in Loading

After a plug-in has been discovered and its data is embedded into the administration portal HTML page, administration portal tries to load the plug-in as part of application startup (unless you have configured it not to load as part of application startup).

For each plug-in that has been discovered, the administration portal creates an HTML iframe element that is used to load its host page. The plug-in host page is necessary to begin the plug-in bootstrap process, which (the bootstrap process) is used to evaluate the plug-in code in the context of the plug-in's iframe element. User interface plug-in infrastructure supports serving plug-in resource files (such as the plug-in host page) from the local file system. The plug-in host page is loaded into the iframe element and the plug-in code is evaluated. After the plug-in code is evaluated, the plug-in communicates with the administration portal by means of the plug-in API.

C.2.4. Red Hat Virtualization User Interface Plug-in Bootstrapping

A typical plug-in bootstrap sequence consists of following steps:

Procedure C.1. Plug-in Bootstrap Sequence

1. Obtain `pluginApi` instance for the given plug-in
2. Obtain runtime plug-in configuration object (optional)
3. Register relevant event handler functions
4. Notify UI plug-in infrastructure to proceed with plug-in initialization

The following code snippet illustrates the above mentioned steps in practice:

```
// Access plug-in API using 'parent' due to this code being evaluated
// within the context of an iframe element.
// As 'parent.pluginApi' is subject to Same-Origin Policy, this will only
// work when WebAdmin HTML page and plug-in
```

```

// host page are served from same origin. WebAdmin HTML page and plug-in
// host page will always be on same origin
// when using UI plug-in infrastructure support to serve plug-in resource
// files.
var api = parent.pluginApi('MyPlugin');

// Runtime configuration object associated with the plug-in (or an empty
// object).
var config = api.configObject();

// Register event handler function(s) for later invocation by UI plug-in
// infrastructure.
api.register({
  // UiInit event handler function.
  UiInit: function() {
    // Handle UiInit event.
    window.alert('Favorite music band is ' + config.band);
  }
});

// Notify UI plug-in infrastructure to proceed with plug-in
// initialization.
api.ready();

```

C.3. USER INTERFACE PLUGIN-RELATED FILES AND THEIR LOCATIONS

Table C.1. UI Plugin-related Files and their Locations

File	Location	Remarks
Plug-in descriptor files (meta-data)	<code>/usr/share/ovirt-engine/ui-plugins/my-plugin.json</code>	
Plug-in user configuration files	<code>/etc/ovirt-engine/ui-plugins/my-plugin-config.json</code>	
Plug-in resource files	<code>/usr/share/ovirt-engine/ui-plugins/<resourcePath>/PluginHostPage.html</code>	<code><resourcePath></code> is defined by the corresponding attribute in the plug-in descriptor.

C.4. EXAMPLE USER INTERFACE PLUG-IN DEPLOYMENT

Follow these instructions to create a user interface plug-in that runs a **Hello World!** program when you sign in to the Red Hat Virtualization Manager administration portal.

Procedure C.2. Deploying a Hello World! Plug-in

1. Create a plug-in descriptor by creating the following file in the Manager at `/usr/share/ovirt-engine/ui-plugins/helloWorld.json`:

```
{
  "name": "HelloWorld",
  "url": "/ovirt-engine/webadmin/plugin/HelloWorld/start.html",
  "resourcePath": "hello-files"
}
```

2. Create the plug-in host page by creating the following file in the Manager at `/usr/share/ovirt-engine/ui-plugins/hello-files/start.html`:

```
<!DOCTYPE html><html><head>
<script>
  var api = parent.pluginApi('HelloWorld');
  api.register({
    UiInit: function() { window.alert('Hello world'); }
  });
  api.ready();
</script>
</head><body></body></html>
```

If you have successfully implemented the **Hello World!** plug-in, you will see this screen when you sign in to the administration portal:





Figure C.1. A Successful Implementation of the Hello World! Plug-in

C.5. USING RED HAT SUPPORT PLUG-IN

The Red Hat Access Plug-in allows you to use Red Hat access services from the Red Hat Virtualization Administration Portal. You must log in using your Red Hat login credentials. The Red Hat Access Plug-in detects when you are not logged in; if you are not logged in, a login window opens.



NOTE

Red Hat Virtualization Administration Portal credentials are not the same as a user's Red Hat login.

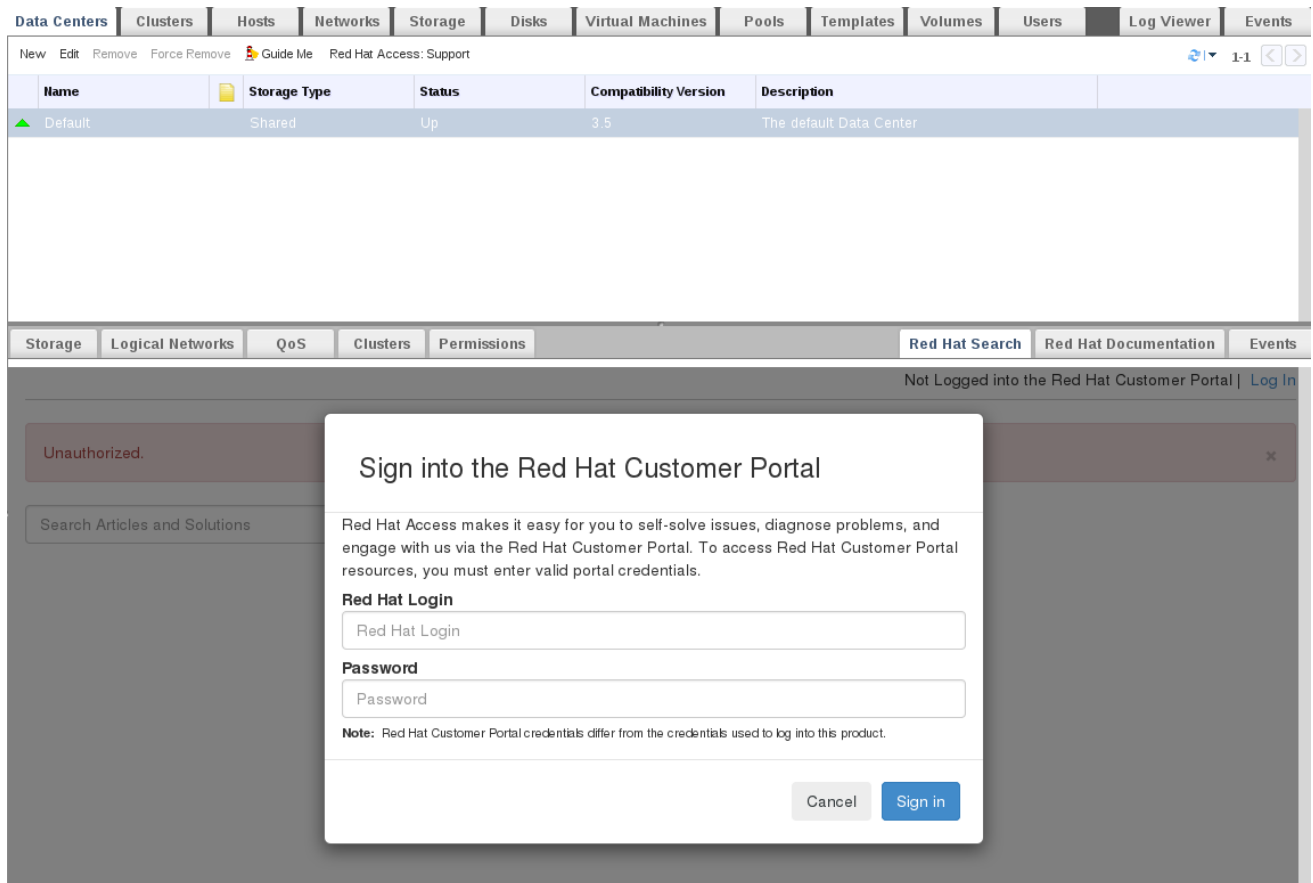


Figure C.2. Red Hat Support Plug-in - Login Window

After logging in, you will be able to access the Red Hat Customer Portal. Red Hat Support Plug-in is available in the details pane as well as in several context menus in the Red Hat Virtualization Administration Portal. Search the Red Hat Access database using the Search bar. Search results display in the left-hand navigation list in the details pane.

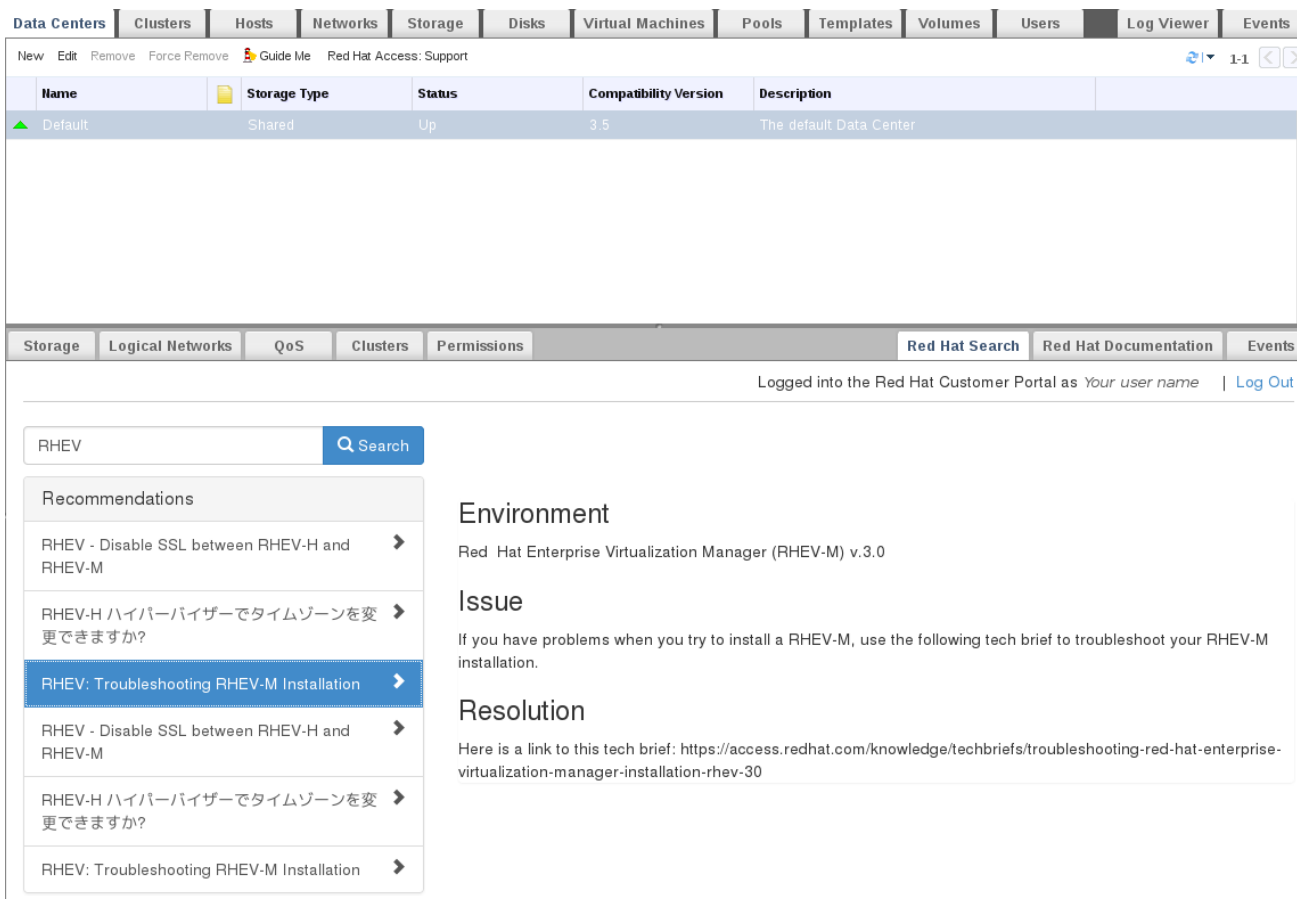


Figure C.3. Red Hat Support Plug-in - Query Results in the Left-Hand Navigation List

Right-click on context menus in the Red Hat Virtualization Administrator Portal to access the Red Hat Support Plug-in.

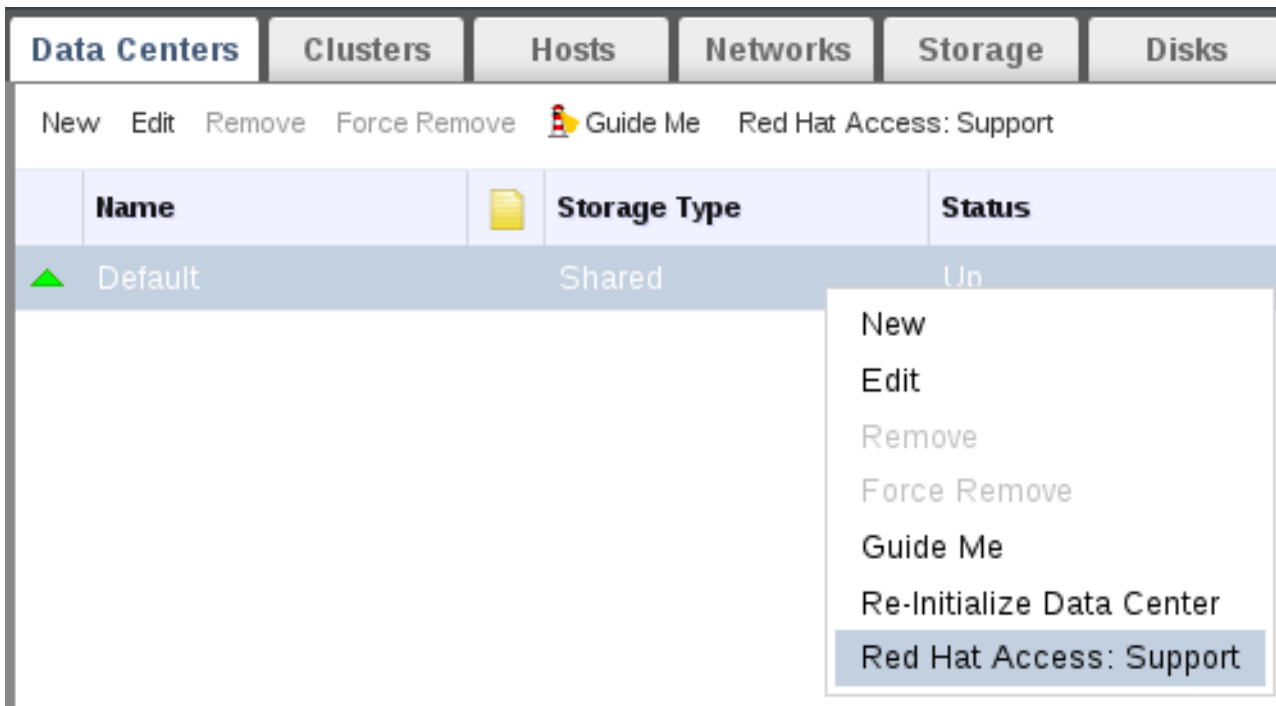
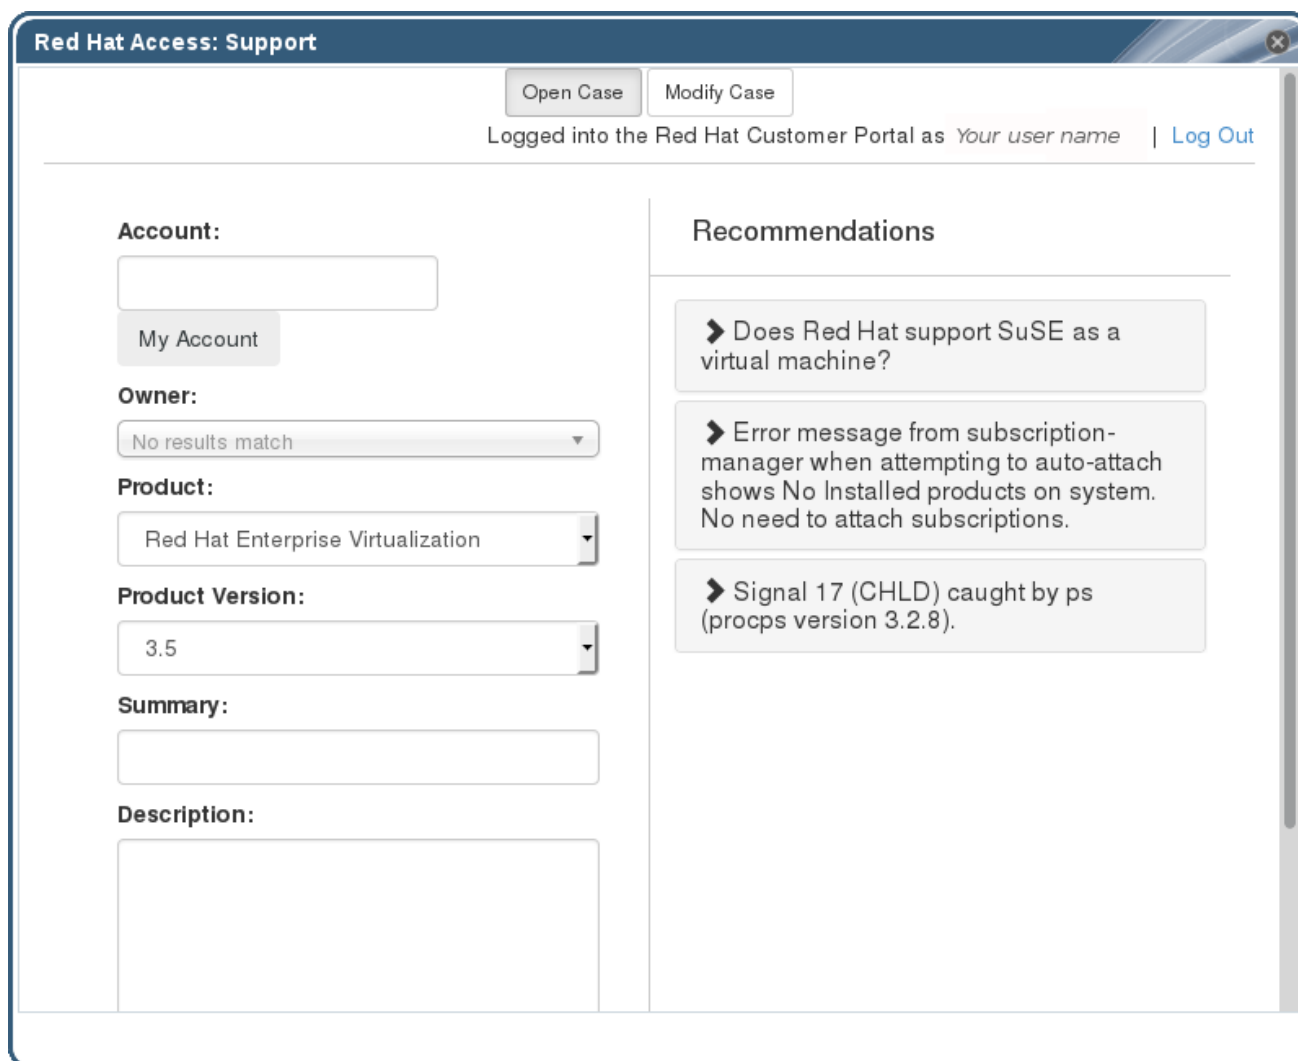


Figure C.4. Right-clicking on a Context Menu to Access Red Hat Support Plug-in

Open a new support case or modify an existing case by selecting the **Open New Support Case** or **Modify Existing Case** buttons.



The screenshot displays the 'Red Hat Access: Support' window. At the top, there are two buttons: 'Open Case' and 'Modify Case'. Below these, a status bar indicates 'Logged into the Red Hat Customer Portal as *Your user name* | [Log Out](#)'. The main content area is divided into two columns. The left column contains form fields for 'Account:' (with a text input and a 'My Account' button), 'Owner:' (with a dropdown menu showing 'No results match'), 'Product:' (with a dropdown menu showing 'Red Hat Enterprise Virtualization'), and 'Product Version:' (with a dropdown menu showing '3.5'). Below these are 'Summary:' and 'Description:' text input areas. The right column is titled 'Recommendations' and contains three items, each with a right-pointing arrow icon: 'Does Red Hat support SuSE as a virtual machine?', 'Error message from subscription-manager when attempting to auto-attach shows No Installed products on system. No need to attach subscriptions.', and 'Signal 17 (CHLD) caught by ps (procps version 3.2.8)'.

Figure C.5. Red Hat Support Plug-in - Opening a New Support Case

Select the **Red Hat Documentation** tab to open the documentation relevant to the part of the Administration Portal currently on the screen.

Data Centers	Clusters	Hosts	Networks	Storage	Disks	Virtual Machines	Pools	Templates	Volumes	Users	Log Viewer	Events
New Edit Remove Force Remove Guide Me Red Hat Access: Support												
Name	Storage Type	Status	Compatibility Version	Description								
Default	Shared	Up	3.5	The default Data Center								



Storage	Logical Networks	QoS	Clusters	Permissions	Red Hat Search	Red Hat Documentation	Events
				◀ Prev Next ▶		<h2 style="color: red;">Administration Guide</h2>	
<h3 style="color: red;">Chapter 4. Data Centers</h3> <ul style="list-style-type: none"> 4.1. Introduction to Data Centers 4.2. The Storage Pool Manager 4.3. SPM Priority 4.4. Using the Events Tab to Identify Problem Objects in Data Centers 4.5. Data Center Tasks 4.6. Data Centers and Storage Domains 4.7. Data Centers and Permissions <h4 style="color: red;">4.1. Introduction to Data Centers</h4> <p>A data center is a logical entity that defines the set of resources used in a specific environment. A data center is considered a container resource, in that it is comprised of logical resources, in the form of clusters and hosts; network resources, in the form of logical networks and physical NICs; and storage resources, in the form of storage domains.</p>							

Figure C.6. Red Hat Support Plug-in - Accessing Documentation

APPENDIX D. RED HAT VIRTUALIZATION AND SSL


D.1. REPLACING THE RED HAT VIRTUALIZATION MANAGER SSL/TLS CERTIFICATE



WARNING

Do not change the permissions and ownerships for the `/etc/pki` directory or any subdirectories. The permission for the `/etc/pki` and the `/etc/pki/ovirt-engine` directory must remain as the default 755.

Use the following procedure(s) if you want to use your organization's third-party CA certificate to identify the Red Hat Virtualization Manager to users connecting over HTTPS.



NOTE

Using a third-party CA certificate for HTTPS connections does not affect the certificate used for authentication between the Manager and hosts. They will continue to use the self-signed certificate generated by the Manager.

Prerequisites

- A third-party CA certificate. This is the certificate of the CA (Certificate Authority) that issued the certificate you want to use. It is provided as a *PEM* file. The certificate chain must be complete up to the root certificate. The chain's order is critical and must be from the last intermediate certificate to the root certificate. This procedure assumes that the third-party CA certificate is provided in `/tmp/3rd-party-ca-cert.pem`.
- The private key that you want to use for Apache httpd. It must not have a password. This procedure assumes that it is located in `/tmp/apache.key`.
- The certificate issued by the CA. This procedure assumes that it is located in `/tmp/apache.cer`.

If you received the private key and certificate from your CA in a P12 file, use the following procedure to extract them. For other file formats, contact your CA. After extracting the private key and certificate, proceed to [Replacing the Red Hat Virtualization Manager Apache SSL Certificate](#).

Procedure D.1. Extracting the Certificate and Private Key from a P12 Bundle

The internal CA stores the internally generated key and certificate in a *P12* file, in `/etc/pki/ovirt-engine/keys/apache.p12`. Red Hat recommends storing your new file in the same location. The following procedure assumes that the new *P12* file is in `/tmp/apache.p12`.

1. Back up the current `apache.p12` file:

```
# cp -p /etc/pki/ovirt-engine/keys/apache.p12 /etc/pki/ovirt-
engine/keys/apache.p12.bck
```

2. Replace the current file with the new file:

```
# cp /tmp/apache.p12 /etc/pki/ovirt-engine/keys/apache.p12
```

3. Extract the private key and certificate to the required locations. If the file is password protected, you must add **-passin pass:password**, replacing *password* with the required password.

```
# openssl pkcs12 -in /etc/pki/ovirt-engine/keys/apache.p12 -nocerts
-nodes > /tmp/apache.key
# openssl pkcs12 -in /etc/pki/ovirt-engine/keys/apache.p12 -nokeys >
/tmp/apache.cer
```



IMPORTANT

For new Red Hat Virtualization installations, you must complete all of the steps in this procedure. If you upgraded from a Red Hat Enterprise Virtualization 3.6 environment with a commercially signed certificate already configured, only steps 1, 8, and 9 are required.

Procedure D.2. Replacing the Red Hat Virtualization Manager Apache SSL Certificate

1. Add your CA certificate to the host-wide trust store:

```
# cp /tmp/3rd-party-ca-cert.pem /etc/pki/ca-trust/source/anchors
```

```
# update-ca-trust
```

2. The Manager has been configured to use **/etc/pki/ovirt-engine/apache-ca.pem**, which is symbolically linked to **/etc/pki/ovirt-engine/ca.pem**. Remove the symbolic link:

```
# rm /etc/pki/ovirt-engine/apache-ca.pem
```

3. Save your CA certificate as **/etc/pki/ovirt-engine/apache-ca.pem**:

```
# cp /tmp/3rd-party-ca-cert.pem /etc/pki/ovirt-engine/apache-ca.pem
```

4. Back up the existing private key and certificate:

```
# cp /etc/pki/ovirt-engine/keys/apache.key.nopass /etc/pki/ovirt-
engine/keys/apache.key.nopass.bck
# cp /etc/pki/ovirt-engine/certs/apache.cer /etc/pki/ovirt-
engine/certs/apache.cer.bck
```

5. Copy the private key to the required location:

■

```
# cp /tmp/apache.key /etc/pki/ovirt-engine/keys/apache.key.nopass
```

6. Copy the certificate to the required location:

```
# cp /tmp/apache.cer /etc/pki/ovirt-engine/certs/apache.cer
```

7. Restart the Apache server:

```
# systemctl restart httpd.service
```

8. Create a new trust store configuration file:

```
# vi /etc/ovirt-engine/engine.conf.d/99-custom-truststore.conf
```

Add the following content and save the file:

```
ENGINE_HTTPS_PKI_TRUST_STORE="/etc/pki/java/cacerts"  
ENGINE_HTTPS_PKI_TRUST_STORE_PASSWORD=""
```

9. Restart the **ovirt-engine** service:

```
# systemctl restart ovirt-engine.service
```

10. Replacing the certificate can cause the log collector to fail. To prevent this, create a new log collector configuration file:

```
# vi /etc/ovirt-engine/logcollector.conf.d/99-custom-ca-cert.conf
```


Add the following content and save the file:

```
[LogCollector]  
cert-file=/etc/pki/ovirt-engine/apache-ca.pem
```

Your users can now connect to the Administration and User portals without being warned about the authenticity of the certificate used to encrypt HTTPS traffic.

D.2. SETTING UP SSL OR TLS CONNECTIONS BETWEEN THE MANAGER AND AN LDAP SERVER

To set up a secure connection between the Red Hat Enterprise Virtualization Manager and an LDAP server, obtain the root CA certificate of the LDAP server, copy the root CA certificate to the Manager, and create a PEM-encoded CA certificate. The keystore type can be any Java-supported type. The following procedure uses the Java KeyStore (JKS) format.



NOTE

For more information on creating a PEM-encoded CA certificate and importing certificates, see the **X.509 CERTIFICATE TRUST STORE** section of the README file at `/usr/share/doc/ovirt-engine-extension-aaa-ldap-version`.

Procedure D.3. Creating a PEM-encoded CA certificate

1. On the Red Hat Virtualization Manager, copy the root CA certificate of the LDAP server to the `/tmp` directory and import the root CA certificate using `keytool` to create a PEM-encoded CA certificate. The following command imports the root CA certificate at `/tmp/myrootca.pem` and creates a PEM-encoded CA certificate `myrootca.jks` under `/etc/ovirt-engine/aaa/`. Note down the certificate's location and password. If you are using the interactive setup tool, this is all the information you need. If you are configuring the LDAP server manually, follow the rest of the procedure to update the configuration files.

```
$ keytool -importcert -noprompt -trustcacerts -alias myrootca -file
/tmp/myrootca.pem -keystore /etc/ovirt-engine/aaa/myrootca.jks -
storepass password
```

2. Update the `/etc/ovirt-engine/aaa/profile1.properties` file with the certificate information:



NOTE

`${local:_basedir}` is the directory where the LDAP property configuration file resides and points to the `/etc/ovirt-engine/aaa` directory. If you created the PEM-encoded CA certificate in a different directory, replace `${local:_basedir}` with the full path to the certificate.

- To use startTLS (recommended):

```
# Create keystore, import certificate chain and uncomment
pool.default.ssl.startTLS = true
pool.default.ssl.truststore.file = ${local:_basedir}/myrootca.jks
pool.default.ssl.truststore.password = password
```

- To use SSL:

```
# Create keystore, import certificate chain and uncomment
pool.default.serverset.single.port = 636
pool.default.ssl.enable = true
pool.default.ssl.truststore.file = ${local:_basedir}/myrootca.jks
pool.default.ssl.truststore.password = password
```

To continue configuring an external LDAP provider, see [Section 15.3.1, “Configuring an External LDAP Provider \(Interactive Setup\)”](#). To continue configuring LDAP and Kerberos for Single Sign-on, see [Section 15.4, “Configuring LDAP and Kerberos for Single Sign-on”](#).

APPENDIX E. USING SEARCH, BOOKMARKS, AND TAGS

E.1. SEARCHES

E.1.1. Performing Searches in Red Hat Virtualization

The Administration Portal enables the management of thousands of resources, such as virtual machines, hosts, users, and more. To perform a search, enter the search query (free-text or syntax-based) into the search bar. Search queries can be saved as bookmarks for future reuse, so you do not have to reenter a search query each time the specific search results are required. Searches are not case sensitive.

E.1.2. Search Syntax and Examples

The syntax of the search queries for Red Hat Virtualization resources is as follows:

result type: {criteria} [sortby sort_spec]

Syntax Examples

The following examples describe how the search query is used and help you to understand how Red Hat Virtualization assists with building search queries.

Table E.1. Example Search Queries

Example	Result
Hosts: Vms.status = up	Displays a list of all hosts running virtual machines that are up.
Vms: domain = qa.company.com	Displays a list of all virtual machines running on the specified domain.
Vms: users.name = Mary	Displays a list of all virtual machines belonging to users with the user name Mary.
Events: severity > normal sortby time	Displays the list of all Events whose severity is higher than Normal, sorted by time.

E.1.3. Search Auto-Completion

The Administration Portal provides auto-completion to help you create valid and powerful search queries. As you type each part of a search query, a drop-down list of choices for the next part of the search opens below the Search Bar. You can either select from the list and then continue typing/selecting the next part of the search, or ignore the options and continue entering your query manually.

The following table specifies by example how the Administration Portal auto-completion assists in constructing a query:

Hosts: Vms.status = down

Table E.2. Example Search Queries Using Auto-Completion

Input	List Items Displayed	Action
h	Hosts (1 option only)	Select Hosts or; Type Hosts
Hosts:	All host properties	Type v
Hosts: v	host properties starting with a v	Select Vms or type Vms
Hosts: Vms	All virtual machine properties	Type s
Hosts: Vms.s	All virtual machine properties beginning with s	Select status or type status
Hosts: Vms.status	= =!	Select or type =
Hosts: Vms.status =	All status values	Select or type down

E.1.4. Search Result Type Options

The result type allows you to search for resources of any of the following types:

- **Vms** for a list of virtual machines
- **Host** for a list of hosts
- **Pools** for a list of pools
- **Template** for a list of templates
- **Event** for a list of events
- **Users** for a list of users
- **Cluster** for a list of clusters
- **Datacenter** for a list of data centers
- **Storage** for a list of storage domains

As each type of resource has a unique set of properties and a set of other resource types that it is associated with, each search type has a set of valid syntax combinations. You can also use the auto-complete feature to create valid queries easily.

E.1.5. Search Criteria

You can specify the search criteria after the colon in the query. The syntax of **{criteria}** is as follows:

<prop><operator><value>

or

<obj-type><prop><operator><value>

Examples

The following table describes the parts of the syntax:

Table E.3. Example Search Criteria

Part	Description	Values	Example	Note
prop	The property of the searched-for resource. Can also be the property of a resource type (see obj-type), or <i>tag</i> (custom tag).	Limit your search to objects with a certain property. For example, search for objects with a <i>status</i> property.	Status	N/A
obj-type	A resource type that can be associated with the searched-for resource.	These are system objects, like data centers and virtual machines.	Users	N/A
operator	Comparison operators.	= != (not equal) > < >= <=	N/A	Value options depend on obj-type.

Part	Description	Values	Example	Note
Value	What the expression is being compared to.	String Integer Ranking Date (formatted according to Regional Settings)	Jones 256 normal	<ul style="list-style-type: none"> • Wildcards can be used within strings. • "" (two sets of quotation marks with no space between them) can be used to represent an uninitialized (empty) string. • Double quotes should be used around a string or date containing spaces

E.1.6. Search: Multiple Criteria and Wildcards

Wildcards can be used in the `<value>` part of the syntax for strings. For example, to find all users beginning with `m`, enter `m*`.

You can perform a search having two criteria by using the Boolean operators **AND** and **OR**. For example:

```
Vms: users.name = m* AND status = Up
```

This query returns all running virtual machines for users whose names begin with "m".

```
Vms: users.name = m* AND tag = "paris-loc"
```

This query returns all virtual machines tagged with "paris-loc" for users whose names begin with "m".

When two criteria are specified without **AND** or **OR**, **AND** is implied. **AND** precedes **OR**, and **OR** precedes implied **AND**.

E.1.7. Search: Determining Search Order

You can determine the sort order of the returned information by using **sortby**. Sort direction (**asc** for ascending, **desc** for descending) can be included.

For example:

```
events: severity > normal sortby time desc
```

This query returns all Events whose severity is higher than Normal, sorted by time (descending order).

E.1.8. Searching for Data Centers

The following table describes all search options for Data Centers.

Table E.4. Searching for Data Centers

Property (of resource or resource-type)	Type	Description (Reference)
<i>Clusters.clusters-prop</i>	Depends on property type	The property of the clusters associated with the data center.
name	String	The name of the data center.
description	String	A description of the data center.
type	String	The type of data center.
status	List	The availability of the data center.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

```
Datcenter: type = nfs and status != up
```

This example returns a list of data centers with:

- A storage type of NFS and status other than up

E.1.9. Searching for Clusters

The following table describes all search options for clusters.

Table E.5. Searching Clusters

Property (of resource or resource-type)	Type	Description (Reference)
<i>Datacenter.datacenter-prop</i>	Depends on property type	The property of the data center associated with the cluster.
Datacenter	String	The data center to which the cluster belongs.
name	String	The unique name that identifies the clusters on the network.
description	String	The description of the cluster.
initialized	String	True or False indicating the status of the cluster.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Clusters: initialized = true or name = Default

This example returns a list of clusters which are:

- initialized; or
- named Default

E.1.10. Searching for Hosts

The following table describes all search options for hosts.

Table E.6. Searching for Hosts

Property (of resource or resource-type)	Type	Description (Reference)
<i>Vms.Vms-prop</i>	Depends on property type	The property of the virtual machines associated with the host.
<i>Templates.templates-prop</i>	Depends on property type	The property of the templates associated with the host.

Property (of resource or resource-type)	Type	Description (Reference)
Events. <i>events-prop</i>	Depends on property type	The property of the events associated with the host.
Users. <i>users-prop</i>	Depends on property type	The property of the users associated with the host.
name	String	The name of the host.
status	List	The availability of the host.
external_status	String	The health status of the host as reported by external systems and plug-ins.
cluster	String	The cluster to which the host belongs.
address	String	The unique name that identifies the host on the network.
cpu_usage	Integer	The percent of processing power used.
mem_usage	Integer	The percentage of memory used.
network_usage	Integer	The percentage of network usage.
load	Integer	Jobs waiting to be executed in the <i>run-queue</i> per processor, in a given time slice.
version	Integer	The version number of the operating system.
cpus	Integer	The number of CPUs on the host.
memory	Integer	The amount of memory available.
cpu_speed	Integer	The processing speed of the CPU.

Property (of resource or resource-type)	Type	Description (Reference)
cpu_model	String	The type of CPU.
active_vms	Integer	The number of virtual machines currently running.
migrating_vms	Integer	The number of virtual machines currently being migrated.
committed_mem	Integer	The percentage of committed memory.
tag	String	The tag assigned to the host.
type	String	The type of host.
datacenter	String	The data center to which the host belongs.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Hosts: cluster = Default and Vms.os = rhel6

This example returns a list of hosts which:

- Are part of the Default cluster and host virtual machines running the Red Hat Enterprise Linux 6 operating system.

E.1.11. Searching for Networks

The following table describes all search options for networks.

Table E.7. Searching for Networks

Property (of resource or resource-type)	Type	Description (Reference)
Cluster_network.clusternetwor k-prop	Depends on property type	The property of the cluster associated with the network.

Property (of resource or resource-type)	Type	Description (Reference)
Host_Network. <i>hostnetwork-prop</i>	Depends on property type	The property of the host associated with the network.
name	String	The human readable name that identifies the network.
description	String	Keywords or text describing the network, optionally used when creating the network.
vlanid	Integer	The VLAN ID of the network.
stp	String	Whether Spanning Tree Protocol (STP) is enabled or disabled for the network.
mtu	Integer	The maximum transmission unit for the logical network.
vmnetwork	String	Whether the network is only used for virtual machine traffic.
datacenter	String	The data center to which the network is attached.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Network: `mtu > 1500` and `vmnetwork = true`

This example returns a list of networks:

- with a maximum transmission unit greater than 1500 bytes
- which are set up for use by only virtual machines.

E.1.12. Searching for Storage

The following table describes all search options for storage.

Table E.8. Searching for Storage

Property (of resource or resource-type)	Type	Description (Reference)
Hosts. <i>hosts-prop</i>	Depends on property type	The property of the hosts associated with the storage.
Clusters. <i>clusters-prop</i>	Depends on property type	The property of the clusters associated with the storage.
name	String	The unique name that identifies the storage on the network.
status	String	The status of the storage domain.
external_status	String	The health status of the storage domain as reported by external systems and plugins.
datacenter	String	The data center to which the storage belongs.
type	String	The type of the storage.
size	Integer	The size of the storage.
used	Integer	The amount of the storage that is used.
committed	Integer	The amount of the storage that is committed.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Storage: size > 200 or used < 50

This example returns a list of storage with:

- total storage space greater than 200 GB; or
- used storage space less than 50 GB.

E.1.13. Searching for Disks

The following table describes all search options for disks.

Table E.9. Searching for Disks

Property (of resource or resource-type)	Type	Description (Reference)
<code>Datacenters.datacenters-prop</code>	Depends on property type	The property of the data centers associated with the disk.
<code>Storages.storages-prop</code>	Depends on property type	The property of the storage associated with the disk.
alias	String	The human readable name that identifies the storage on the network.
description	String	Keywords or text describing the disk, optionally used when creating the disk.
provisioned_size	Integer	The virtual size of the disk.
size	Integer	The size of the disk.
actual_size	Integer	The actual size allocated to the disk.
creation_date	Integer	The date the disk was created.
bootable	String	Whether the disk can or cannot be booted. Valid values are one of 0 , 1 , yes , or no
shareable	String	Whether the disk can or cannot be attached to more than one virtual machine at a time. Valid values are one of 0 , 1 , yes , or no
format	String	The format of the disk. Can be one of unused , unassigned , cow , or raw .

Property (of resource or resource-type)	Type	Description (Reference)
status	String	The status of the disk. Can be one of unassigned , ok , locked , invalid , or illegal .
disk_type	String	The type of the disk. Can be one of image or lun .
number_of_vms	Integer	The number of virtual machine(s) to which the disk is attached.
vm_names	String	The name(s) of the virtual machine(s) to which the disk is attached.
quota	String	The name of the quota enforced on the virtual disk.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Disks: format = cow and provisioned_size > 8

This example returns a list of virtual disks with:

- QCOW, also known as thin provisioning, format; and
- an allocated disk size greater than 8 GB.

E.1.14. Searching for Volumes

The following table describes all search options for volumes.

Table E.10. Searching for Volumes

Property (of resource or resource-type)	Type	Description (Reference)
Volume. <i>cluster-prop</i>	Depends on property type	The property of the clusters associated with the volume.

Property (of resource or resource-type)	Type	Description (Reference)
Cluster	String	The name of the cluster associated with the volume.
name	String	The human readable name that identifies the volume.
type	String	Can be one of distribute, replicate, distributed_replicate, stripe, or distributed_stripe.
transport_type	Integer	Can be one of TCP or RDMA.
replica_count	Integer	Number of replica.
stripe_count	Integer	Number of stripes.
status	String	The status of the volume. Can be one of Up or Down.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Volume: transport_type = rdma and stripe_count >= 2

This example returns a list of volumes with:

- Transport type set to RDMA; and
- with 2 or more stripes.

E.1.15. Searching for Virtual Machines

The following table describes all search options for virtual machines.

Table E.11. Searching for Virtual Machines

Property (of resource or resource-type)	Type	Description (Reference)
---	------	-------------------------

Property (of resource or resource-type)	Type	Description (Reference)
Hosts. <i>hosts-prop</i>	Depends on property type	The property of the hosts associated with the virtual machine.
Templates. <i>templates-prop</i>	Depends on property type	The property of the templates associated with the virtual machine.
Events. <i>events-prop</i>	Depends on property type	The property of the events associated with the virtual machine.
Users. <i>users-prop</i>	Depends on property type	The property of the users associated with the virtual machine.
Storage. <i>storage-prop</i>	Depends on the property type	The property of storage devices associated with the virtual machine.
Vnic. <i>mac-prop</i>	Depends on the property type	The property of the MAC address associated with the virtual machine.
name	String	The name of the virtual machine.
status	List	The availability of the virtual machine.
ip	Integer	The IP address of the virtual machine.
uptime	Integer	The number of minutes that the virtual machine has been running.
domain	String	The domain (usually Active Directory domain) that groups these machines.
os	String	The operating system selected when the virtual machine was created.
creationdate	Date	The date on which the virtual machine was created.

Property (of resource or resource-type)	Type	Description (Reference)
address	String	The unique name that identifies the virtual machine on the network.
cpu_usage	Integer	The percent of processing power used.
mem_usage	Integer	The percentage of memory used.
network_usage	Integer	The percentage of network used.
memory	Integer	The maximum memory defined.
apps	String	The applications currently installed on the virtual machine.
cluster	List	The cluster to which the virtual machine belongs.
pool	List	The virtual machine pool to which the virtual machine belongs.
loggedinuser	String	The name of the user currently logged in to the virtual machine.
tag	List	The tags to which the virtual machine belongs.
datacenter	String	The data center to which the virtual machine belongs.
type	List	The virtual machine type (server or desktop).
quota	String	The name of the quota associated with the virtual machine.

Property (of resource or resource-type)	Type	Description (Reference)
description	String	Keywords or text describing the virtual machine, optionally used when creating the virtual machine.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.
next_run_configuration_exists	Boolean	The virtual machine has pending configuration changes.

Example

Vms: template.name = Win* and user.name = ""

This example returns a list of virtual machines, where:

- The template on which the virtual machine is based begins with Win and the virtual machine is assigned to any user.

Example

Vms: cluster = Default and os = windows7

This example returns a list of virtual machines, where:

- The cluster to which the virtual machine belongs is named Default and the virtual machine is running the Windows 7 operating system.

E.1.16. Searching for Pools

The following table describes all search options for Pools.

Table E.12. Searching for Pools

Property (of resource or resource-type)	Type	Description (Reference)
name	String	The name of the pool.
description	String	The description of the pool.
type	List	The type of pool.

Property (of resource or resource-type)	Type	Description (Reference)
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Pools: type = automatic

This example returns a list of pools with:

- Type of automatic

E.1.17. Searching for Templates

The following table describes all search options for templates.

Table E.13. Searching for Templates

Property (of resource or resource-type)	Type	Description (Reference)
<i>Vms.Vms-prop</i>	String	The property of the virtual machines associated with the template.
<i>Hosts.hosts-prop</i>	String	The property of the hosts associated with the template.
<i>Events.events-prop</i>	String	The property of the events associated with the template.
<i>Users.users-prop</i>	String	The property of the users associated with the template.
name	String	The name of the template.
domain	String	The domain of the template.
os	String	The type of operating system.
creationdate	Integer	The date on which the template was created. Date format is <i>mm/dd/yy</i> .

Property (of resource or resource-type)	Type	Description (Reference)
childcount	Integer	The number of virtual machines created from the template.
mem	Integer	Defined memory.
description	String	The description of the template.
status	String	The status of the template.
cluster	String	The cluster associated with the template.
datacenter	String	The data center associated with the template.
quota	String	The quota associated with the template.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Template: `Events.severity >= normal and Vms.uptime > 0`

This example returns a list of templates, where:

- Events of normal or greater severity have occurred on virtual machines derived from the template, and the virtual machines are still running.

E.1.18. Searching for Users

The following table describes all search options for users.

Table E.14. Searching for Users

Property (of resource or resource-type)	Type	Description (Reference)
<code>Vms.Vms-prop</code>	Depends on property type	The property of the virtual machines associated with the user.

Property (of resource or resource-type)	Type	Description (Reference)
Hosts. <i>hosts-prop</i>	Depends on property type	The property of the hosts associated with the user.
Templates. <i>templates-prop</i>	Depends on property type	The property of the templates associated with the user.
Events. <i>events-prop</i>	Depends on property type	The property of the events associated with the user.
name	String	The name of the user.
lastname	String	The last name of the user.
username	String	The unique name of the user.
department	String	The department to which the user belongs.
group	String	The group to which the user belongs.
title	String	The title of the user.
status	String	The status of the user.
role	String	The role of the user.
tag	String	The tag to which the user belongs.
pool	String	The pool to which the user belongs.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Users: Events.severity > normal and Vms.status = up or Vms.status = pause

This example returns a list of users where:

- Events of greater than normal severity have occurred on their virtual machines AND the virtual machines are still running; or
- The users' virtual machines are paused.

E.1.19. Searching for Events

The following table describes all search options you can use to search for events. Auto-completion is offered for many options as appropriate.

Table E.15. Searching for Events

Property (of resource or resource-type)	Type	Description (Reference)
<i>Vms.Vms-prop</i>	Depends on property type	The property of the virtual machines associated with the event.
<i>Hosts.hosts-prop</i>	Depends on property type	The property of the hosts associated with the event.
<i>Templates.templates-prop</i>	Depends on property type	The property of the templates associated with the event.
<i>Users.users-prop</i>	Depends on property type	The property of the users associated with the event.
<i>Clusters.clusters-prop</i>	Depends on property type	The property of the clusters associated with the event.
<i>Volumes.Volumes-prop</i>	Depends on property type	The property of the volumes associated with the event.
type	List	Type of the event.
severity	List	The severity of the event: Warning/Error/Normal.
message	String	Description of the event type.
time	List	Day the event occurred.
username	String	The user name associated with the event.

Property (of resource or resource-type)	Type	Description (Reference)
event_host	String	The host associated with the event.
event_vm	String	The virtual machine associated with the event.
event_template	String	The template associated with the event.
event_storage	String	The storage associated with the event.
event_datacenter	String	The data center associated with the event.
event_volume	String	The volume associated with the event.
correlation_id	Integer	The identification number of the event.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Events: Vms.name = testdesktop and Hosts.name = gonzo.example.com

This example returns a list of events, where:

- The event occurred on the virtual machine named **testdesktop** while it was running on the host **gonzo.example.com**.

E.2. BOOKMARKS

E.2.1. Saving a Query String as a Bookmark

A bookmark can be used to remember a search query, and shared with other users.

Procedure E.1. Saving a Query String as a Bookmark

1. Enter the desired search query in the search bar and perform the search.

2. Click the star-shaped **Bookmark** button to the right of the search bar to open the **New Bookmark** window.




Figure E.1. Bookmark Icon

3. Enter the **Name** of the bookmark.
4. Edit the **Search string** field (if applicable).
5. Click **OK** to save the query as a bookmark and close the window.
6. The search query is saved and displays in the **Bookmarks** pane.

You have saved a search query as a bookmark for future reuse. Use the **Bookmark** pane to find and select the bookmark.

E.2.2. Editing a Bookmark

You can modify the name and search string of a bookmark.

Procedure E.2. Editing a Bookmark

1. Click the **Bookmarks** tab on the far left side of the screen.
2. Select the bookmark you wish to edit.
3. Click the **Edit** button to open the **Edit Bookmark** window.
4. Change the **Name** and **Search string** fields as necessary.
5. Click **OK** to save the edited bookmark.

You have edited a bookmarked search query.

E.2.3. Deleting a Bookmark

When a bookmark is no longer needed, remove it.

Procedure E.3. Deleting a Bookmark

1. Click the **Bookmarks** tab on the far left side of the screen.
2. Select the bookmark you wish to remove.

3. Click the **Remove** button to open the **Remove Bookmark** window.
4. Click **OK** to remove the selected bookmark.

You have removed a bookmarked search query.

E.3. TAGS

E.3.1. Using Tags to Customize Interactions with Red Hat Virtualization

After your Red Hat Virtualization platform is set up and configured to your requirements, you can customize the way you work with it using tags. Tags provide one key advantage to system administrators: they allow system resources to be arranged into groups or categories. This is useful when many objects exist in the virtualization environment and the administrator would like to concentrate on a specific set of them.

This section describes how to create and edit tags, assign them to hosts or virtual machines and search using the tags as criteria. Tags can be arranged in a hierarchy that matches a structure, to fit the needs of the enterprise.

Administration Portal Tags can be created, modified, and removed using the **Tags** pane.

E.3.2. Creating a Tag

Create tags so you can filter search results using tags.

Procedure E.4. Creating a Tag

1. Click the **Tags** tab on the left side of the screen.
2. Select the node under which you wish to create the tag. For example, to create it at the highest level, click the **root** node.
3. Click the **New** button to open the **New Tag** window.
4. Enter the **Name** and **Description** of the new tag.
5. Click **OK** to create the tag.

The new tag is created and displays on the **Tags** tab.

E.3.3. Modifying a Tag

You can edit the name and description of a tag.

Procedure E.5. Modifying a Tag

1. Click the **Tags** tab on the left side of the screen.
2. Select the tag you wish to modify.
3. Click **Edit** to open the **Edit Tag** window.
4. Change the **Name** and **Description** fields as necessary.

5. Click **OK** to save the edited tag.

You have modified the properties of a tag.

E.3.4. Deleting a Tag

When a tag is no longer needed, remove it.

Procedure E.6. Deleting a Tag

1. Click the **Tags** tab on the left side of the screen.
2. Select the tag you wish to delete.
3. Click **Remove** to open the **Remove Tag(s)** window. The message warns you that removing the tag will also remove all descendants of the tag.
4. Click **OK** to delete the selected tag.

You have removed the tag and all its descendants. The tag is also removed from all the objects that it was attached to.

E.3.5. Adding and Removing Tags to and from Objects

You can assign tags to and remove tags from hosts, virtual machines, and users.

Procedure E.7. Adding and Removing Tags to and from Objects

1. Click the resource tab, and select the object(s) you wish to tag or untag.
2. Click the **Assign Tags** button to open the **Assign Tags** window.
3. Select the check box to assign a tag to the object, or clear the check box to detach the tag from the object.
4. Click **OK**.

The specified tag is now added or removed as a custom property of the selected object(s).

E.3.6. Searching for Objects Using Tags

- Enter a search query using **tag** as the property and the desired value or set of values as criteria for the search.

The objects tagged with the specified criteria are listed in the results list.

APPENDIX F. BRANDING

F.1. BRANDING

F.1.1. Re-Branding the Manager

Various aspects of the Red Hat Virtualization Manager can be customized, such as the icons used by and text displayed in pop-up windows and the links shown on the Welcome Page. This allows you to re-brand the Manager and gives you fine-grained control over the end look and feel presented to administrators and users.

The files required to customize the Manager are located in the `/etc/ovirt-engine/branding/` directory on the system on which the Manager is installed. The files comprise a set of cascading style sheet files that are used to style various aspects of the graphical user interface and a set of properties files that contain messages and links that are incorporated into various components of the Manager.

To customize a component, edit the file for that component and save the changes. The next time you open or refresh that component, the changes will be applied.

F.1.2. Login Screen

The login screen is the login screen used by both the Administration Portal and User Portal. The elements of the login screen that can be customized are as follows:

- The border
- The header image on the left
- The header image on the right
- The header text

The classes for the login screen are located in `common.css`.

F.1.3. Administration Portal Screen

The administration portal screen is the main screen that is shown when you log into the Administration Portal. The elements of the administration portal screen that can be customized are as follows:

- The logo
- The left background image
- The center background image
- The right background image
- The text to the right of the logo

The classes for the administration portal screen are located in `web_admin.css`.

F.1.4. User Portal Screen

The user portal screen is the screen that is shown when you log into the User Portal. The elements of the user portal screen that can be customized are as follows:

- The logo
- The center background image
- The right background image
- The border around the main grid
- The text above the **Logged in user** label

The classes for the user portal screen are located in **user_portal.css**.

F.1.5. Pop-Up Windows

Pop-up windows are all windows in the Manager that allow you to create, edit or update an entity such as a host or virtual machine. The elements of pop-up windows that can be customized are as follows:

- The border
- The header image on the left
- The header center image (repeated)

The classes for pop-up windows are located in **common.css**.

F.1.6. Tabs

There are two types of tab elements in the User Portal - the main tabs for switching between the Basic view and the Extended view, and the tabs on the left side of the screen when the Extended view is selected. Many pop-up windows in the Administration Portal also include tabs. The elements of these tabs that can be customized are as follows:

- Active
- Inactive

The classes for tabs are located in **common.css** and **user_portal.css**.

F.1.7. The Welcome Page

The Welcome Page is the page that is initially displayed when you visit the homepage of the Manager. In addition to customizing the overall look and feel, you can also make other changes such as adding links to the page for additional documentation or internal websites by editing a template file. The elements of the Welcome Page that can be customized are as follows:

- The page title
- The header (left, center and right)
- The error message

- The link to forward and the associated message for that link

The classes for the Welcome Page are located in **welcome_style.css**.

The Template File

The template file for the Welcome Page is a regular HTML file of the name **welcome_page.template** that does not contain **HTML**, **HEAD** or **BODY** tags. This file is inserted directly into the Welcome Page itself, and acts as a container for the content that is displayed in the Welcome Page. As such, you must edit this file to add new links or change the content itself. Another feature of the template file is that it contains placeholder text such as **{user_portal}** that is replaced by corresponding text in the **messages.properties** file when the Welcome Page is processed.

F.1.8. The Page Not Found Page

The Page Not Found page is a page that is displayed when you open a link to a page that cannot be found in the Red Hat Virtualization Manager. The elements of the Page Not Found page that can be customized are as follows:

- The page title
- The header (left, center and right)
- The error message
- The link to forward and the associated message for that link

The classes for the Page Not Found page are located in **welcome_style.css**.

APPENDIX G. SYSTEM ACCOUNTS

G.1. SYSTEM ACCOUNTS

G.1.1. Red Hat Virtualization Manager User Accounts

A number of system user accounts are created to support Red Hat Virtualization when the `rhev` package is installed. Each system user has a default user identifier (UID). The system user accounts created are:

- The **vds**m user (UID **36**). Required for support tools that mount and access NFS storage domains.
- The **ovirt** user (UID **108**). Owner of the **ovirt-engine** Red Hat JBoss Enterprise Application Platform instance.
- The **ovirt-vmconsole** user (UID **498**). Required for the guest serial console.

G.1.2. Red Hat Virtualization Manager Groups

A number of system user groups are created to support Red Hat Virtualization when the `rhev` package is installed. Each system user group has a default group identifier (GID). The system user groups created are:

- The **kvm** group (GID **36**). Group members include:
 - The **vds**m user.
- The **ovirt** group (GID **108**). Group members include:
 - The **ovirt** user.
- The **ovirt-vmconsole** group (GID **498**). Group members include:
 - The **ovirt-vmconsole** user.

G.1.3. Virtualization Host User Accounts

A number of system user accounts are created on the virtualization host when the `vds`m and `qemu-kvm-rhev` packages are installed. Each system user has a default user identifier (UID). The system user accounts created are:

- The **vds**m user (UID **36**).
- The **qemu** user (UID **107**).
- The **sanlock** user (UID **179**).
- The **ovirt-vmconsole** user (UID **498**).

**IMPORTANT**


The user identifiers (UIDs) and group identifiers (GIDs) allocated may vary between systems. The **vds**m user is fixed to a UID of **36** and the **kvm** group is fixed to a GID of **36**.

If UID **36** or GID **36** is already used by another account on the system a conflict will arise during installation of the **vds**m and **qemu-kvm-rhev** packages.

G.1.4. Virtualization Host Groups

A number of system user groups are created on the virtualization host when the **vds**m and **qemu-kvm-rhev** packages are installed. Each system user group has a default group identifier (GID). The system user groups created are:

- The **kvm** group (GID **36**). Group members include:
 - The **qemu** user.
 - The **sanlock** user.
- The **qemu** group (GID **107**). Group members include:
 - The **vds**m user.
 - The **sanlock** user.
- The **ovirt-vmconsole** group (GID **498**). Group members include:
 - The **ovirt-vmconsole** user.

**IMPORTANT**

The user identifiers (UIDs) and group identifiers (GIDs) allocated may vary between systems. The **vds**m user is fixed to a UID of **36** and the **kvm** group is fixed to a GID of **36**.

If UID **36** or GID **36** is already used by another account on the system a conflict will arise during installation of the **vds**m and **qemu-kvm-rhev** packages.