Red Hat Virtualization 4.2

Administration Guide

Administration Tasks in Red Hat Virtualization

Red Hat Virtualization 4.2 Administration Guide

Administration Tasks in Red Hat Virtualization

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

Legal Notice

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

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 by clicking **Administration → Configure**, 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.

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.

Creating a New Role

1. Click **Administration → Configure** to open the **Configure** window. The **Roles** tab is selected by default, showing a list of default User and Administrator roles, and any custom roles.

2. Click **New**.

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 want to permit or deny for the role you are setting up.

7. Click **OK** to apply the changes. 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.

Editing or Copying a Role

1. Click **Administration → Configure** 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 or the VM Portal to manage these resources.

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 VM Portal. Because he has **UserVmManager** permissions, he can modify the virtual machine. 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 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 VM 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.

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.

Figure 1.1. 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, "Red Hat Virtualization Object Hierarchy". 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 use both the Administration Portal and the VM Portal.

1.2. SYSTEM PERMISSIONS

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

Figure 1.2. Permissions & Roles
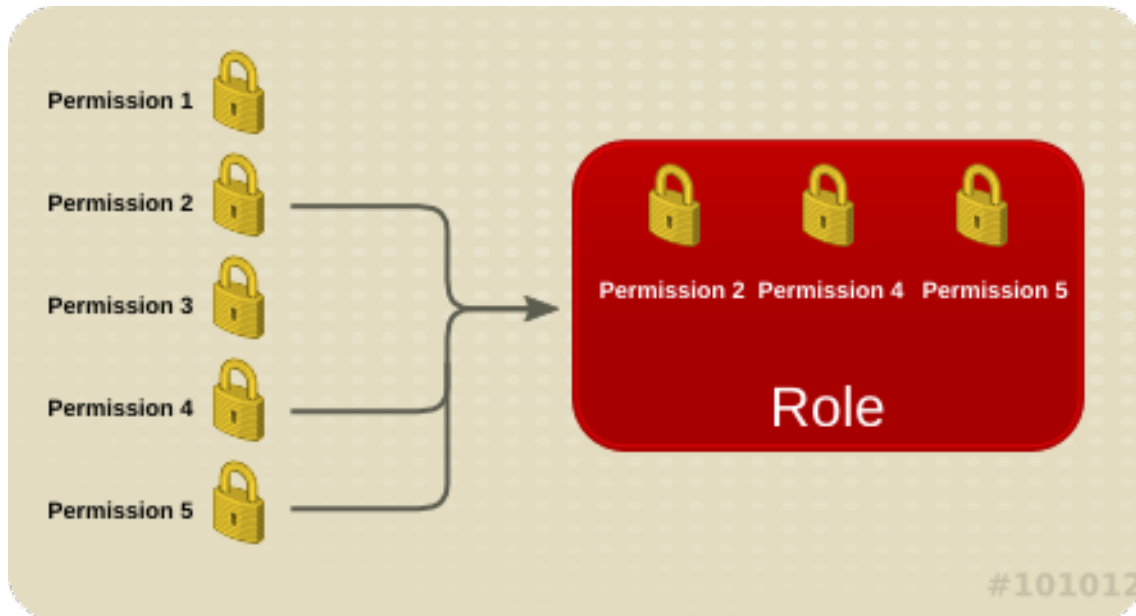
Any permissions that apply to a container object also apply to all members of that container. The following diagram depicts the hierarchy of objects in the system.

Figure 1.3. Red Hat Virtualization Object Hierarchy
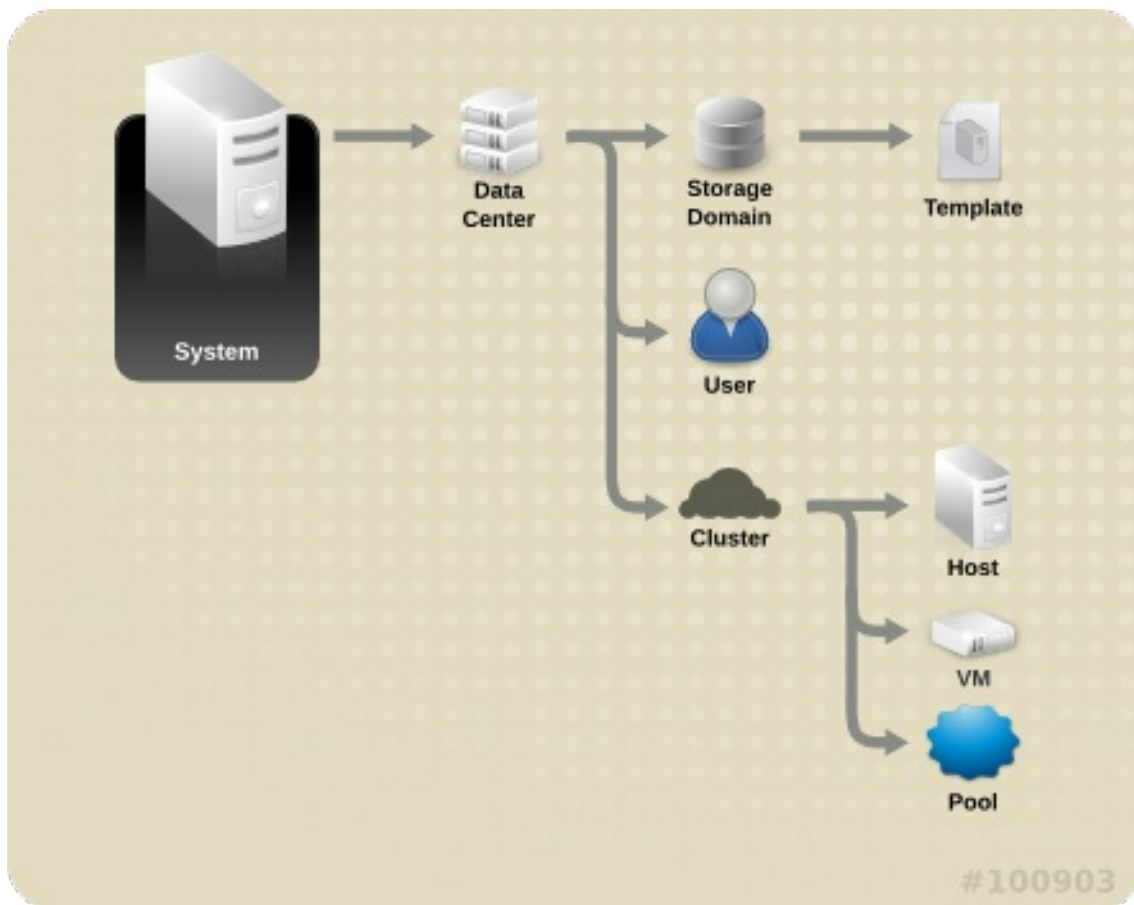
1.2.1. User Properties

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 VM Portal; however, it has no bearing on what a user can see in the VM Portal.

- User Role: Allows access to the **VM Portal** for managing and accessing virtual machines and templates. A user role determines what a user can see in the VM Portal. Permissions granted to a user with an administrator role are reflected in the actions available to that user in the VM Portal.

1.2.3. User Roles Explained

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

Role	Privileges	Notes
VmCreator	Can create virtual machines in the VM 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.
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.

Role	Privileges	Notes
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.

1.2.5. 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.

Assigning a Role to a Resource

1. Find and click the resource's name to open the details view.

2. Click the **Permissions** tab 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**.

The user now has the inherited permissions of that role enabled for that resource.

1.2.6. 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.

Removing a Role from a Resource

1. Find and click the resource's name to open the details view.

2. Click the **Permissions** tab 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**.

5. Click **OK**.

1.2.7. 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.

1.2.8. Data Center Administrator Roles Explained

Data Center Permission Roles

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

Table 1.5. 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.

Role	Privileges	Notes
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.

1.2.9. 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 cluster 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.

1.2.10. Cluster Administrator Roles Explained

Cluster Permission Roles

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

Table 1.6. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes

Role	Privileges	Notes
ClusterAdmin	Cluster Administrator	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. However, a **ClusterAdmin** does not have permissions to attach or detach networks from a cluster, to do so **NetworkAdmin** permissions are required.
NetworkAdmin	Network Administrator	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.

1.2.11. 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.

1.2.12. 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 1.7. Red Hat Virtualization Network Administrator and User Roles

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.

1.2.13. 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.

1.2.14. Host Administrator Roles Explained

Host Permission Roles

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

Table 1.8. 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.

1.2.15. 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.

1.2.16. Storage Administrator Roles Explained

Storage Domain Permission Roles

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

Table 1.9. 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.

1.2.17. 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.

1.2.18. Virtual Machine Pool Administrator Roles Explained

Pool Permission Roles

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

Table 1.10. 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.

1.2.19. 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 VM 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.

1.2.20. 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 VM Portal.

Table 1.11. 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.

1.2.21. Setting a Legacy SPICE Cipher

SPICE consoles use FIPS–compliant encryption by default, with a cipher string. The default SPICE cipher string is: **kECDHE+FIPS:kDHE+FIPS:kRSA+FIPS:!eNULL:!aNULL**

This string is generally sufficient. However, if you have a virtual machine with an older operating system or SPICE client, where either one or the other does not support FIPS–compliant encryption, you must use a weaker cipher string. Otherwise, a connection security error may occur if you install a new cluster or a new host in an existing cluster and try to connect to that virtual machine.

You can change the cipher string by using an Ansible playbook.

Changing the cipher string

1. On the Manager machine, create a file in the directory **/usr/share/ovirt-engine/playbooks**. For example:

    ```
    # vim /usr/share/ovirt-engine/playbooks/change-spice-cipher.yml
    ```

2. Enter the following in the file and save it:

    ```
    name: oVirt - setup weaker SPICE encryption for old clients
    hosts: hostname
    vars:
      host_deploy_spice_cipher_string: 'DEFAULT:-RC4:-3DES:-DES'
    roles:
      - ovirt-host-deploy-spice-encryption
    ```

3. Run the file you just created:

    ```
    # ansible-playbook -l hostname /usr/share/ovirt-engine/playbooks/change-spice-cipher.yml
    ```

Alternatively, you can reconfigure the host with the Ansible playbook **ovirt-host-deploy** using the **--extra-vars** option with the variable **host_deploy_spice_cipher_string**, as follows:

```
# ansible-playbook -l hostname \
  --extra-vars host_deploy_spice_cipher_string="DEFAULT:-RC4:-3DES:-DES" \
  /usr/share/ovirt-engine/playbooks/ovirt-host-deploy.yml
```

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 filter modules apply hard enforcement and filter out hosts that do not meet the conditions specified by that filter. The weights modules apply soft enforcement, and are used to control the relative priority of factors considered when determining the hosts in a cluster on which a virtual machine can run.

The Red Hat Virtualization Manager provides five default scheduling policies: **Evenly_Distributed**, **Cluster_Maintenance**, **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 5.2.5, "Scheduling Policy Settings Explained" for more information about the properties of each scheduling policy.

Figure 1.4. Evenly Distributed Scheduling Policy
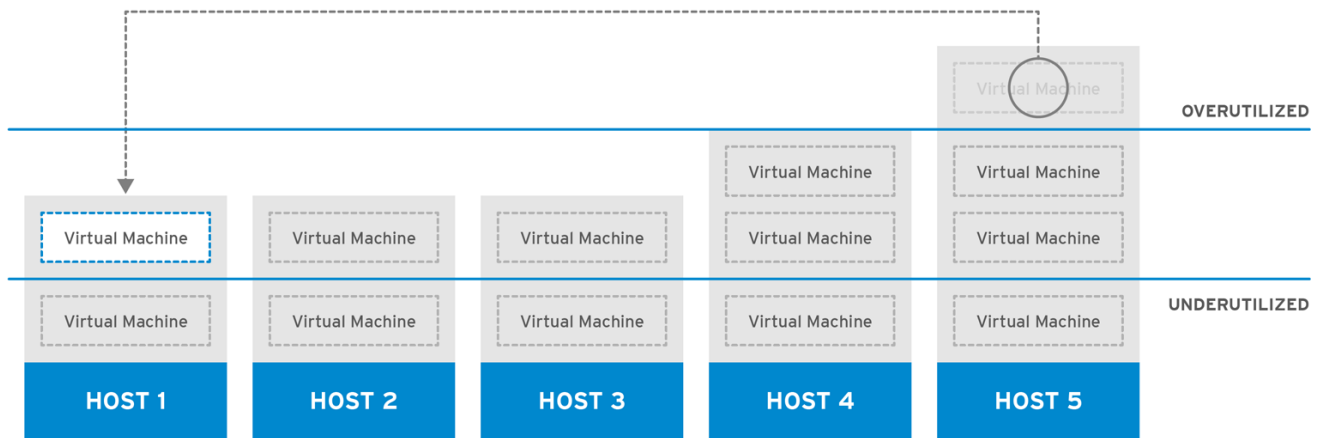
RHV_ 444396_0417

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**.

Figure 1.5. Power Saving Scheduling Policy
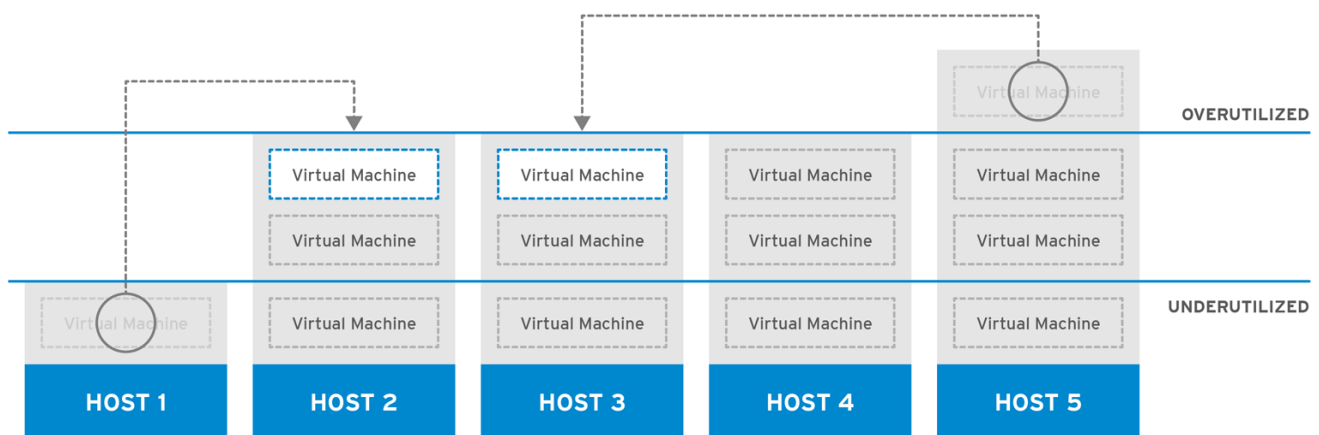
RHV_ 444396_0417

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 **Cluster_Maintenance** scheduling policy limits activity in a cluster during maintenance tasks. When the **Cluster_Maintenance** policy is set, no new virtual machines may be started, except highly available

virtual machines. If host failure occurs, highly available virtual machines will restart properly and any virtual machine can migrate.

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.

Creating a Scheduling Policy

1. Click **Administration → Configure**.

2. Click the **Scheduling Policies** tab.

3. Click **New**.

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 1.12. 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.
Filter Modules	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: • **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. • **VmToHostsAffinityGroups**: Group of 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 one of the hosts in a group or on a separate host that is excluded from the group. • **InClusterUpgrade**: Hosts that are running an earlier operating system than the host that the virtual machine currently runs on. • **HostDevice**: Hosts that do not support host devices required by the virtual machine. • **HA**: Forces the Manager virtual machine in a self-hosted engine environment to run only on hosts with a positive high availability score.

Field Name	Description
	- **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. - **HostedEnginesSpares**: Reserves space for the Manager virtual machine on a specified number of self-hosted engine nodes. - **Label**: Hosts that do not have the required affinity labels. - **Compatibility-Version**: Runs virtual machines only on hosts with the correct compatibility version support. - **CPUOverloaded**: Hosts that are CPU overloaded.
Weights Modules	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. - **InClusterUpgrade**: Weight hosts in accordance with their operating system version. The weight penalizes hosts with earlier operating systems more than hosts with the same operating system as the host that the virtual machine is currently running on. This ensures that priority is always given to hosts with later 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. - **VmToHostsAffinityGroups**: Weights hosts in accordance with the affinity groups defined for virtual machines. This weight module determines how likely virtual

Field Name	Description
	machines in an affinity group are to run on one of the hosts in a group or on a separate host that is excluded from the group. ● **OptimalForCPUPowerSaving**: Weights hosts in accordance with their CPU usage, giving priority to hosts with higher CPU usage. ● **OptimalForEvenCpuDistribution**: 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. ● **PreferredHosts**: Preferred hosts have priority during virtual machine setup. ● **OptimalForMemoryPowerSaving**: Weights hosts in accordance with their memory usage, giving priority to hosts with lower available memory. ● **OptimalForMemoryEvenDistribution**: Weights hosts in accordance with their memory usage, giving priority to hosts with higher available memory.
Load Balancer	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.
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 1.13. 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.

Fields in the **New Virtual Machine** and **Edit Virtual Machine** windows that are bound to an instance type 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.

Creating an Instance Type

1. Click **Administration → Configure**.

2. Click the **Instance Types** tab.

3. Click **New**.

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.

Editing Instance Type Properties

1. Click **Administration → Configure**.

2. Click the **Instance Types** tab.

3. Select the instance type to be edited.

4. Click **Edit**.

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

Removing an Instance Type

1. Click **Administration → Configure**.

2. Click the **Instance Types** tab.

3. Select the instance type to be removed.

4. Click **Remove**.

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** check box. 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(s) of MAC addresses allocated for each cluster. A MAC address pool is specified for each cluster. 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 cluster are within the range for the assigned MAC address pool.

The same MAC address pool can be shared by multiple clusters, but each cluster 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 clusters see Section 5.2.1, "Creating a New Cluster".
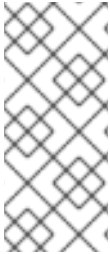
NOTE

If more than one Red Hat Virtualization cluster shares a network, do not rely solely on the default MAC address pool because the virtual machines of each cluster will try to use the same range of MAC addresses, leading to conflicts. To avoid MAC address conflicts, check the MAC address pool ranges to ensure that each cluster is assigned a unique MAC address range.

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.

Creating a MAC Address Pool

1. Click **Administration → Configure**.

2. Click the **MAC Address Pools** tab.

3. Click **Add**.

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.
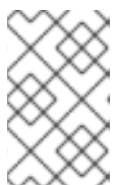
   NOTE

   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.

Editing MAC Address Pool Properties

1. Click **Administration → Configure**.

2. Click the **MAC Address Pools** tab.

3. Select the MAC address pool to be edited.

4. Click **Edit**.

5. Change the **Name**, **Description**, **Allow Duplicates**, and **MAC Address Ranges** fields as required.
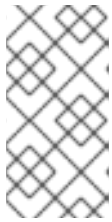
> **NOTE**
>
> When a MAC address range is updated, the MAC addresses of existing NICs are not reassigned. MAC addresses that were already assigned, but are outside of the new MAC address range, are added as user-specified MAC addresses and are still tracked by that MAC address pool.

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.

Editing MAC Address Pool Permissions

1. Click **Administration → Configure**.

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 a created MAC address pool if the pool is not associated with a cluster, but the default MAC address pool cannot be removed.

Removing a MAC Address Pool

1. Click **Administration → Configure**.

2. Click the **MAC Address Pools** tab.

3. Select the MAC address pool to be removed.

4. Click the **Remove**.

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 you to a problem and allows you to analyze the problem area.

The information in the dashboard is updated every 15 minutes by default from 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 page 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 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
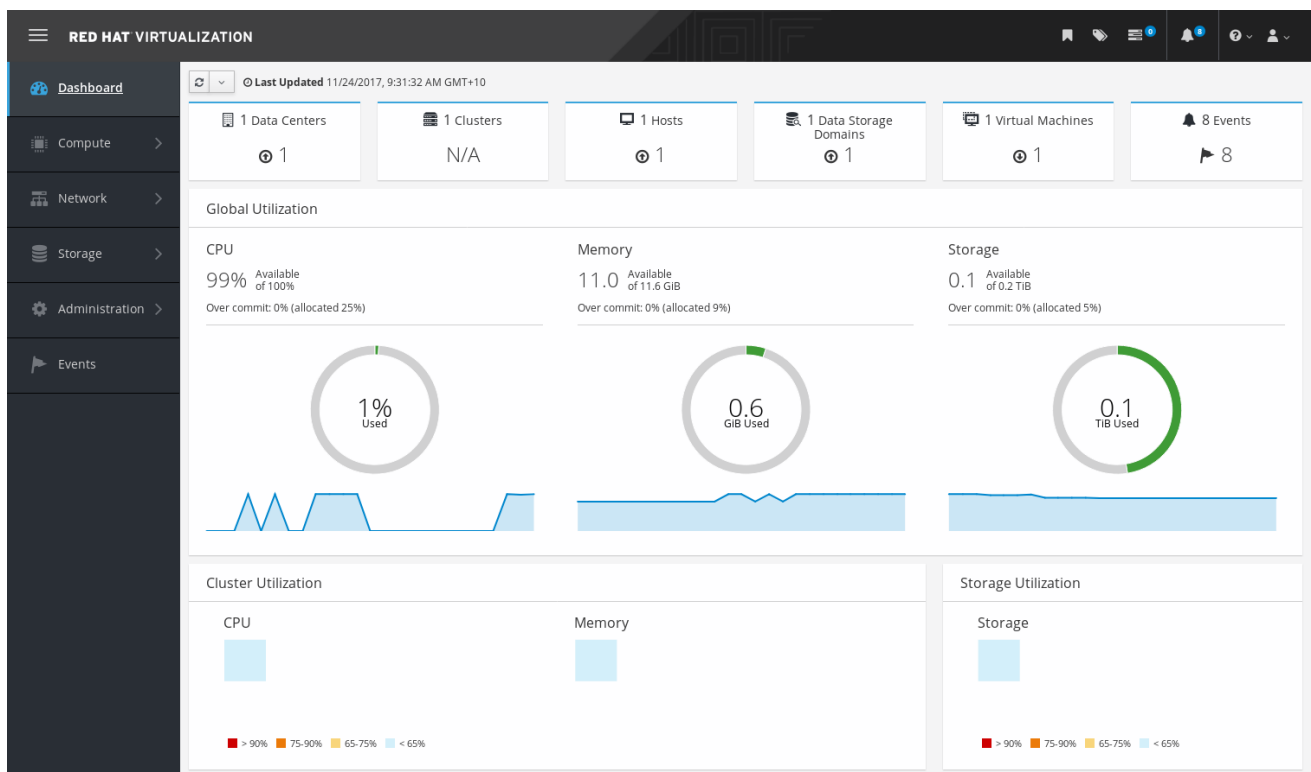
2.1. PREREQUISITES

The Dashboard requires that Data Warehouse is installed and configured. 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
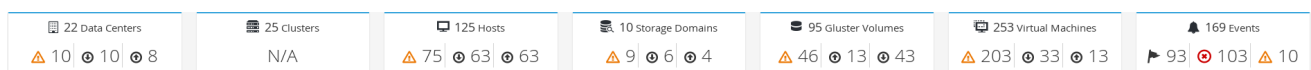
The title shows the number of a type of resource and their status is displayed below the title. Clicking on the resource title navigates to the related page in the Red Hat Virtualization Manager. The status for **Clusters** is always displayed as N/A.

Table 2.1. Resource Status

Icon	Status
✅	None of that resource added to Red Hat Virtualization.
⚠️	Shows the number of a resource with a warning status. Clicking on the icon navigates to the appropriate page with the search limited to that resource with a warning status. The search is limited differently for each resource: • **Data Centers**: The search is limited to data centers that are not operational or non-responsive. • **Gluster Volumes**: The search is limited to gluster volumes that are powering up, paused, migrating, waiting, suspended, or powering down. • **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 page with the search limited to resources that are up.

Icon	Status
⊕	Shows the number of a resource with a down status. Clicking on the icon navigates to the appropriate page with the search limited to resources with a down status. The search is limited differently for each resource: • **Data Centers**: The search is limited to data centers that are uninitialized, in maintenance mode, or with a down status. • **Gluster Volumes**: The search is limited to gluster volumes that are detached or inactive. • **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.
⚑	Shows the number of events with an alert status. Clicking on the icon navigates to **Events** 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 **Events** 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
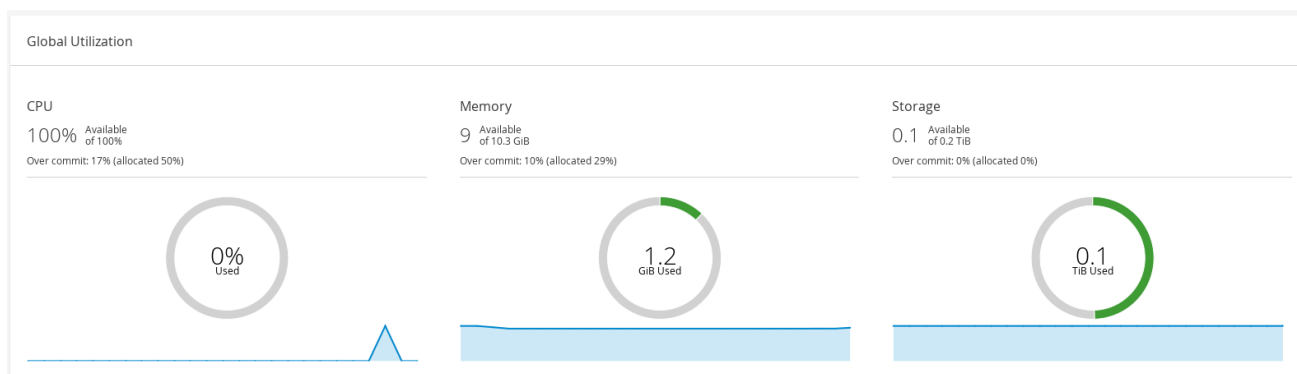
- The top section shows the percentage of the available CPU, memory or storage and the over commit ratio. For example, the over commit ratio for the CPU is calculated by dividing the number of virtual cores by the number of physical cores that are available for the running virtual machines based on the latest data in 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)
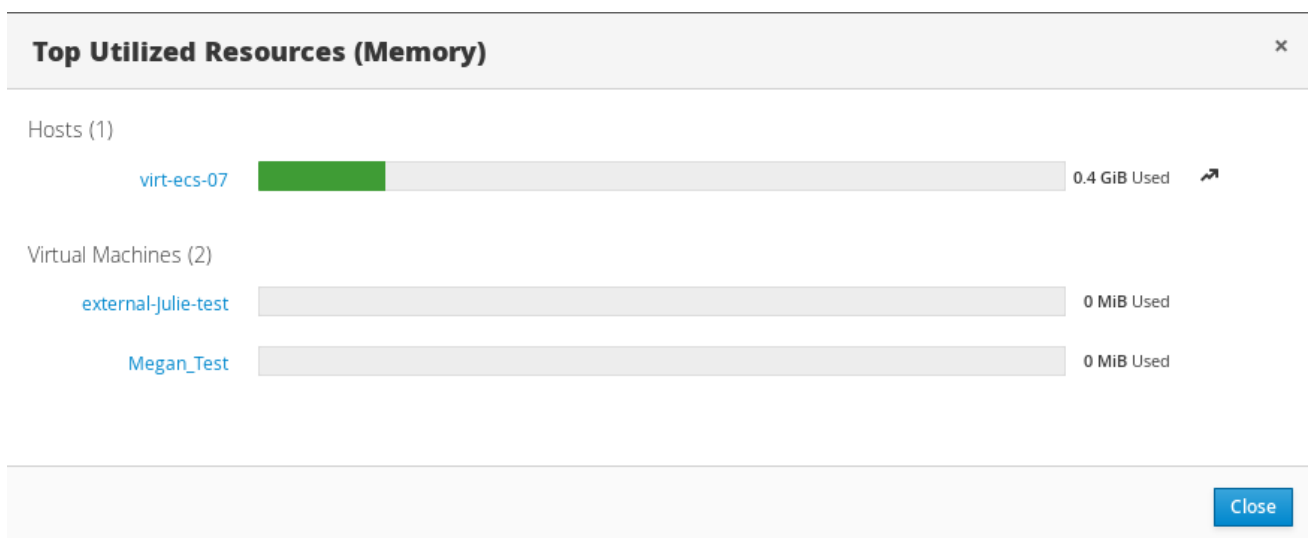
Clicking the donut in the global utilization section of the Dashboard will display a list of the top utilized resources for the CPU, memory or storage. For CPU and memory the pop-up shows a list of the ten hosts and virtual machines with the highest usage. For storage the pop-up shows a list of the top 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** section shows the cluster utilization for the CPU and memory in a heatmap.

Figure 2.5. Cluster Utilization
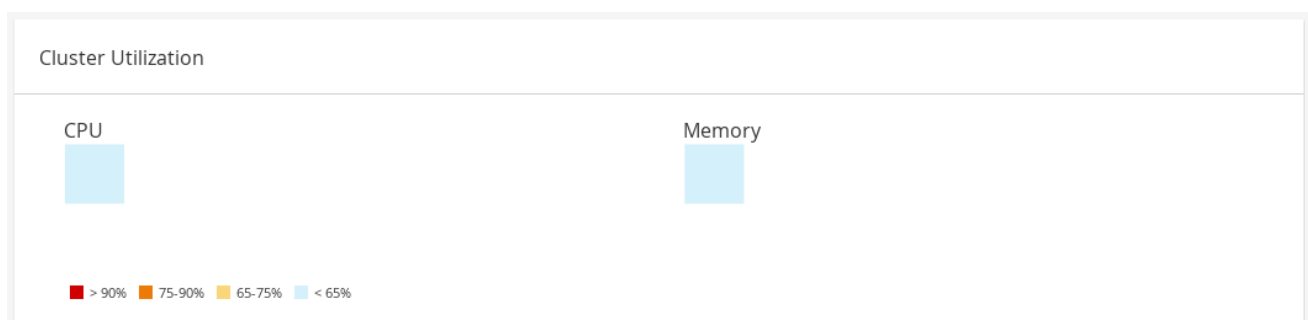
2.4.1. CPU

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 **Compute → Hosts** 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 **Compute → Hosts** 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** section shows the storage utilization in a heatmap.

Figure 2.6. Storage Utilization
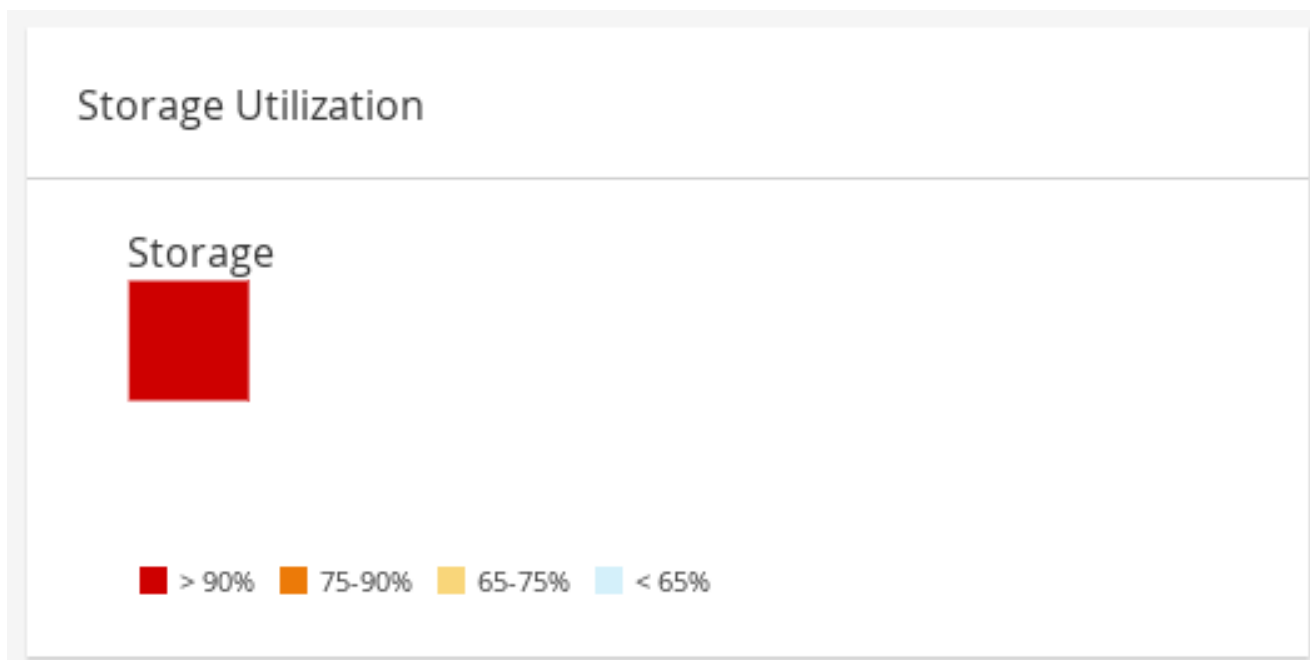
The heatmap shows the average utilization of the storage for the last 24 hours. The formula used to calculate the storage usage by the cluster is the total utilization of the storage in the cluster. This is calculated by using the average storage utilization for each host over the last 24 hours to find 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 **Storage → Domains** 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

Creating a Storage Quality of Service Entry

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **Storage**, click **New**.

5. Enter a **QoS Name** and a **Description** for the quality of service entry.

6. Specify the **Throughput** quality of service by clicking one of the radio buttons:

   - **None**

   - **Total** – Enter the maximum permitted total throughput in the **MB/s** field.

   - **Read/Write** – Enter the maximum permitted throughput for read operations in the left **MB/s** field, and the maximum permitted throughput for write operations in the right **MB/s** field.

7. Specify the input and output (**IOps**) quality of service by clicking one of the radio buttons:

   - **None**

   - **Total** – Enter the maximum permitted number of input and output operations per second in the **IOps** field.

   - **Read/Write** – Enter the maximum permitted number of input operations per second in the left **IOps** field, and the maximum permitted number of output operations per second in the right **IOps** field.

8. 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.

Removing a Storage Quality of Service Entry

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **Storage**, select a storage quality of service entry and click **Remove**.

5. Click **OK**.

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.

Creating a Virtual Machine Network Quality of Service Entry

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **VM Network**, click **New**.

5. Enter a **Name** for the virtual machine network quality of service entry.

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	The settings to be applied to inbound traffic. Select or clear the **Inbound** check box to enable or disable these settings. • **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	The settings to be applied to outbound traffic. Select or clear the **Outbound** check box to enable or disable these settings. • **Average**: The average speed of outbound traffic. • **Peak**: The speed of outbound traffic during peak times. • **Burst**: The speed of outbound traffic during bursts.

To change the maximum value allowed by the **Average**, **Peak**, or **Burst** fields, use the **engine-config** command to change the value of the **MaxAverageNetworkQoSValue**, **MaxPeakNetworkQoSValue**, or **MaxBurstNetworkQoSValue** configuration keys. You must restart the **ovirt-engine** service for any changes to take effect. For example:

```
# engine-config -s MaxAverageNetworkQoSValue=2048
# systemctl restart ovirt-engine
```

3.2.3. Removing a Virtual Machine Network Quality of Service Entry

Remove an existing virtual machine network quality of service entry.

Removing a Virtual Machine Network Quality of Service Entry

1. Click **Compute** → **Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **VM Network**, select a virtual machine network quality of service entry and click **Remove**.

5. Click **OK**.

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.

Creating a Host Network Quality of Service Entry

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **Host Network**, click **New**.

5. Enter a **Qos Name** and a description for the quality of service entry.

6. Enter the desired values for **Weighted Share**, **Rate Limit [Mbps]**, and **Committed Rate [Mbps]**.

7. 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.

Field Name	Description
Outbound	The settings to be applied to outbound traffic. • **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.

To change the maximum value allowed by the **Rate Limit [Mbps]** or **Committed Rate [Mbps]** fields, use the **engine-config** command to change the value of the **MaxAverageNetworkQoSValue** configuration key. You must restart the **ovirt-engine** service for the change to take effect. For example:

```
# engine-config -s MaxAverageNetworkQoSValue=2048
# systemctl restart ovirt-engine
```

3.3.3. Removing a Host Network Quality of Service Entry

Remove an existing network quality of service entry.

Removing a Host Network Quality of Service Entry

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **Host Network**, select a host network quality of service entry and click **Remove**.

5. 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.

Creating a CPU Quality of Service Entry

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **CPU**, click **New**.

5. Enter a **QoS Name** and a **Description** for the quality of service entry.

6. Enter the maximum processing capability the quality of service entry permits in the **Limit (%)** field. Do not include the **%** symbol.

7. 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.

Removing a CPU Quality of Service Entry

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **QoS** tab.

4. Under **CPU**, select a CPU quality of service entry and click **Remove**.

5. Click **OK**.

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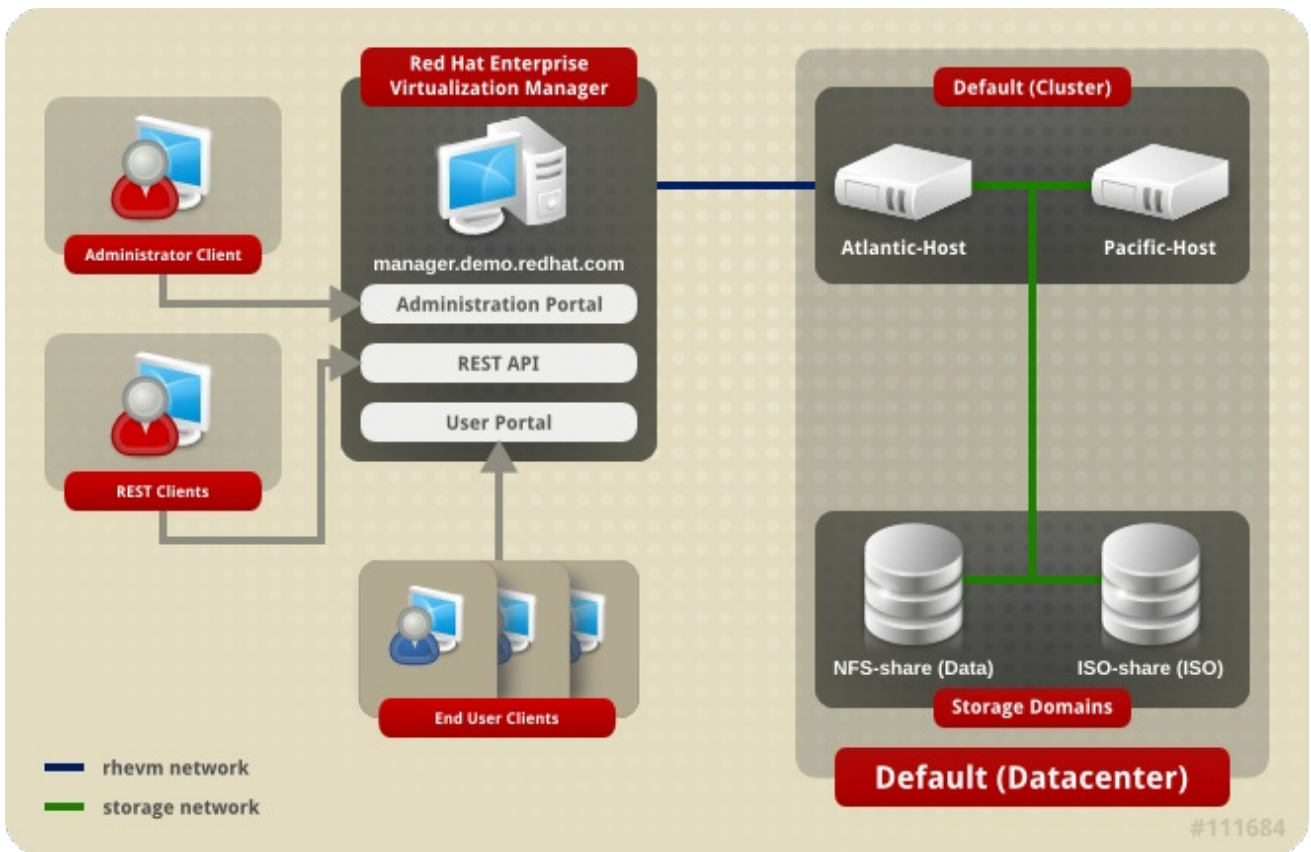
Red Hat Virtualization creates a default data center during installation. You can configure the default data center, or set up new appropriately named data centers.

Figure 4.2. Data Center Objects
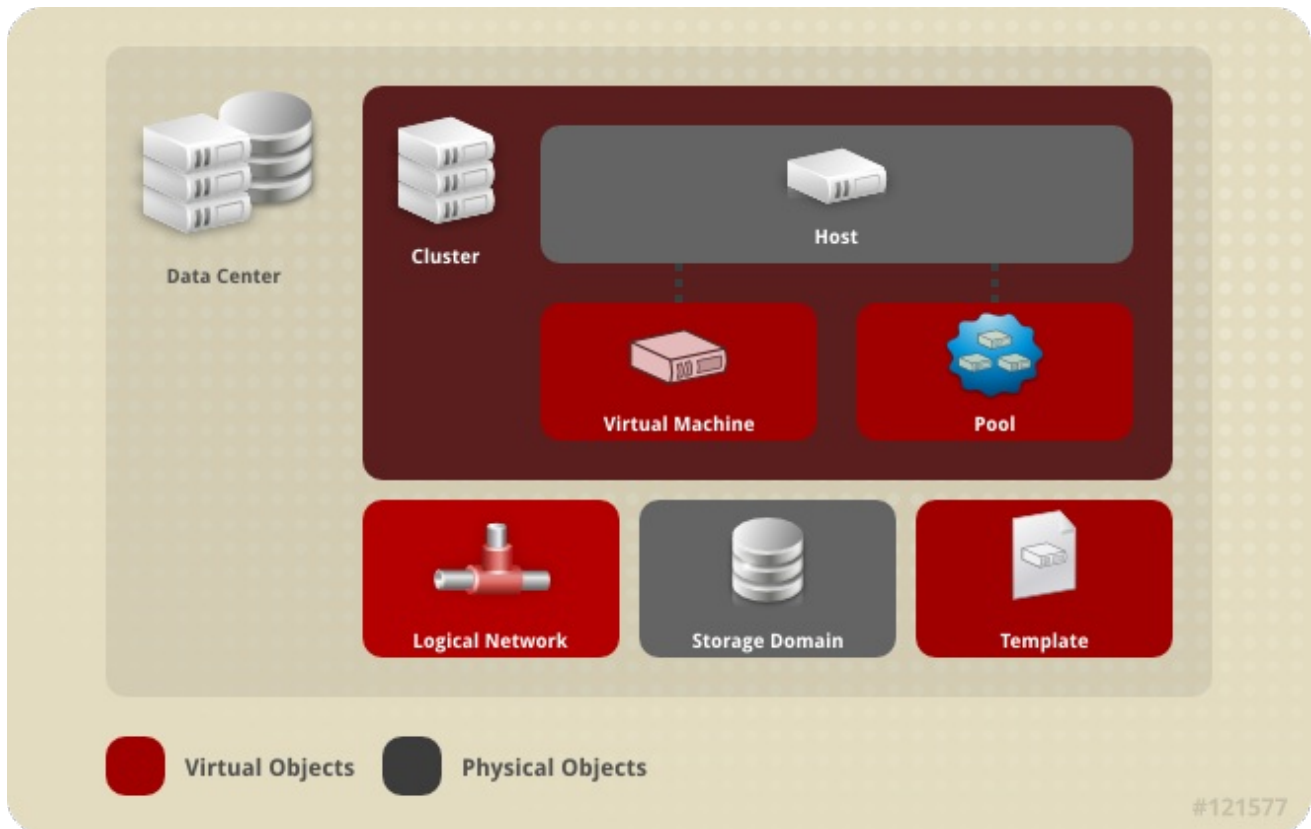
4.2. THE STORAGE POOL MANAGER

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 in the **SPM** tab in the **Edit Host** window.

4.4. DATA CENTER TASKS

4.4.1. 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.
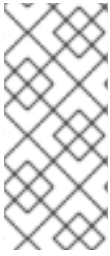
> **NOTE**
>
> Once the **Compatibility Version** is set, it cannot be lowered at a later time; version regression is not allowed.
>
> The option to specify a MAC pool range for a data center has been disabled, and is now done at the cluster level.

Creating a New Data Center

1. Click **Compute → Data Centers**.

2. Click **New**.

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. Click **OK** to create the data center and open the **Data Center - Guide Me** window.

6. 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 **More Actions → Guide Me**.

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

4.4.2. 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.

Field	Description/Action
Storage Type	Choose **Shared** or **Local** storage type. Different types of storage domains (iSCSI, NFS, FC, POSIX, and Gluster) can be added to the same data center. Local and shared domains, however, cannot be mixed. You can change the storage type after the data center is initialized. See Section 4.4.6, "Changing the Data Center Storage Type".
Compatibility Version	The version of Red Hat Virtualization. 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.
Quota Mode	Quota is a resource limitation tool provided with Red Hat Virtualization. Choose one of: • **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
Comment	Optionally add a plain text comment about the data center.

4.4.3. 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. You must re-initialize your master data domain if its data is corrupted. 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.

Re-Initializing a Data Center

1. Click **Compute → Data Centers** and select the data center.

2. Ensure that any storage domains attached to the data center are in maintenance mode.

3. Click **More Actions → Re-Initialize Data Center**.

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**.

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.

4.4.4. Removing a Data Center

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

Removing a Data Center

1. Ensure the storage domains attached to the data center are in maintenance mode.

2. Click **Compute → Data Centers** and select the data center to remove.

3. Click **Remove**.

4. Click **OK**.

4.4.5. 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.

Force Removing a Data Center

1. Click **Compute → Data Centers** and select the data center to remove.

2. Click **More Actions → Force Remove**.

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.

4.4.6. Changing the Data Center Storage Type

You can change the storage type of the data center after it has been initialized. This is useful for data domains that are used to move virtual machines or templates around.

Limitations

- Shared to Local – For a data center that does not contain more than one host and more than one cluster, since a local data center does not support it.

- Local to Shared – For a data center that does not contain a local storage domain.

Changing the Data Center Storage Type

1. Click **Compute → Data Centers** and select the data center to change.

2. Click **Edit**.

3. Change the **Storage Type** to the desired value.

4. Click **OK**.

4.4.7. 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.

> **IMPORTANT**
>
> 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

1. Click **Compute → Data Centers** and select the data center to change.

2. Click **Edit**.

3. Change the **Compatibility Version** to the desired value.

4. Click **OK** to open the **Change Data Center Compatibility Version** confirmation window.

5. Click **OK** to confirm.

4.5. DATA CENTERS AND STORAGE DOMAINS

4.5.1. 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.

Attaching an Existing Data Domain to a Data Center

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **Storage** tab to list the storage domains already attached to the data center.

4. Click **Attach Data**.

5. 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.

6. Click **OK**.

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

4.5.2. 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.

Attaching an Existing ISO Domain to a Data Center

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **Storage** tab to list the storage domains already attached to the data center.

4. Click **Attach ISO**.

5. Click the radio button for the appropriate ISO domain.

6. Click **OK**.

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

4.5.3. Attaching an Existing Export Domain to a Data Center
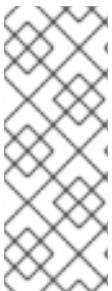
> **NOTE**
>
> The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disks, and templates can then be uploaded from the imported storage domain to the attached data center. See Section 8.7, "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.

Attaching an Existing Export Domain to a Data Center

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **Storage** tab to list the storage domains already attached to the data center.

4. Click **Attach Export**.

5. Click the radio button for the appropriate export domain.

6. Click **OK**.

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

4.5.4. Detaching a Storage Domain from a Data Center

Detaching a storage domain from a data center stops 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.

Detaching a Storage Domain from a Data Center

1. Click **Compute → Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **Storage** tab to list the storage domains attached to the data center.

4. Select the storage domain to detach. If the storage domain is **Active**, click **Maintenance**.

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

6. Click **Detach**.

7. Click **OK**.

It can take up to several minutes for the storage domain to disappear from the details view.

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 cluster and settings on the virtual machines. 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
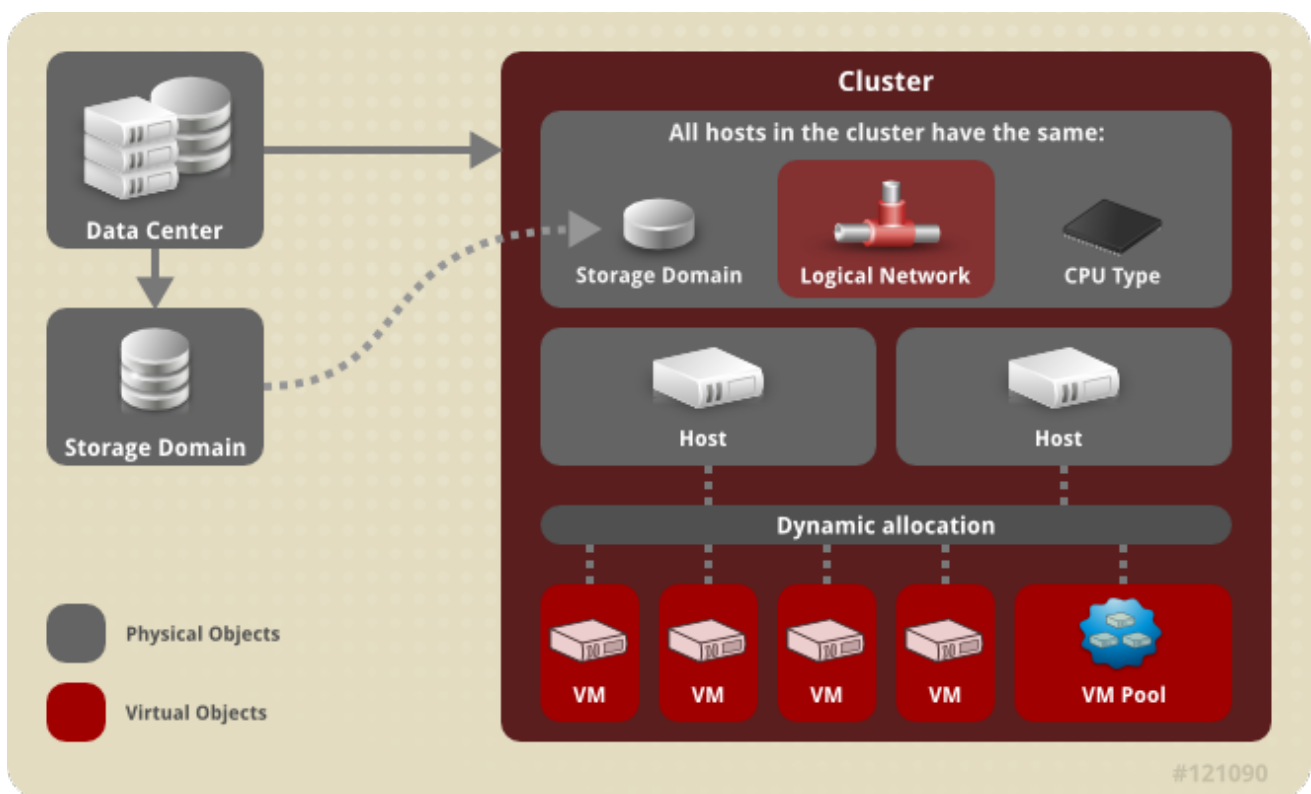
5.2. CLUSTER TASKS

> **NOTE**
>
> Some cluster options do not apply to Gluster clusters. For more information about using Red Hat Gluster Storage with Red Hat Virtualization, see *Configuring Red Hat Virtualization with Red Hat Gluster Storage*.

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.

Creating a New Cluster

1. Click **Compute → Clusters**.

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 the **Switch Type** from the drop-down list.

9. Select the **Firewall Type** for hosts in the cluster, either **iptables** or **firewalld**.

10. Select either the **Enable Virt Service** or **Enable Gluster Service** check box to define whether the cluster will be populated with virtual machine hosts or with Gluster-enabled nodes.

11. 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.

12. 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.

13. Optionally select the **/dev/hwrng source** (external hardware device) check box to specify the random number generator device that all hosts in the cluster will use. The **/dev/urandom source** (Linux-provided device) is enabled by default.

14. 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.

15. Click the **Migration Policy** tab to define the virtual machine migration policy for the cluster.

16. 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.

17. Click the **Console** tab to optionally override the global SPICE proxy, if any, and specify the address of a SPICE proxy for hosts in the cluster.

18. Click the **Fencing policy** tab to enable or disable fencing in the cluster, and select fencing options.

19. Click the **MAC Address Pool** tab to specify a MAC address pool other than the default pool for the cluster. For more options on creating, editing, or removing MAC address pools, see Section 1.5, "MAC Address Pools" .

20. Click **OK** to create the cluster and open the **Cluster - Guide Me** window.

21. 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 **More Actions → Guide Me**.

5.2.2. General Cluster Settings Explained

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 that will be assigned the management network role. The default is **ovirtmgmt**. This network will also be used for migrating virtual machines if the migration network is not properly attached to the source or the destination hosts. On existing clusters, the management network can only be changed using the **Manage Networks** button in the **Logical Networks** tab in the details view.

Field	Description/Action
CPU Architecture	The CPU architecture of the cluster. Different CPU types are available depending on which CPU architecture is selected. • **undefined**: All CPU types are available. • **x86_64**: All Intel and AMD CPU types are available. • **ppc64**: Only IBM POWER 8 is available.
CPU Type	The CPU type of the cluster. See CPU Requirements in the *Planning and Prerequisites Guide* for a list of supported CPU types. 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.
Compatibility Version	The version of Red Hat Virtualization. You will not be able to select a version earlier than the version specified for the data center.
Switch Type	The type of switch used by the cluster. **Linux Bridge** is the standard Red Hat Virtualization switch. **OVS** provides support for Open vSwitch networking features.
Firewall Type	Specifies the firewall type for hosts in the cluster, either **iptables** or **firewalld**. If you change an existing cluster's firewall type, you must reinstall all hosts in the cluster to apply the change.
Default Network Provider	Specifies the default external network provider that the cluster will use. If you select Open Virtual Network (OVN), the hosts added to the cluster are automatically configured to communicate with the OVN provider. If you change the default network provider, you must reinstall all hosts in the cluster to apply the change.
Enable Virt Service	If this radio button is selected, hosts in this cluster will be used to run virtual machines.

Field	Description/Action
Enable Gluster Service	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.
Import existing gluster configuration	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. The following options are required for each host in the cluster that is being imported: ● **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	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.
Enable to set Host maintenance reason	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.
Additional Random Number Generator source	If the check box is selected, all hosts in the cluster have the additional random number generator device available. This enables passthrough of entropy from the random number generator device to virtual machines.

5.2.3. Optimization Settings Explained

Memory Considerations

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 Considerations

- **For non-CPU-intensive workloads**, you can run virtual machines with a total number of processor cores greater than the number of cores in the host. Doing so enables the following:

  - You can run a greater number of virtual machines, which reduces hardware requirements.

  - You can configure virtual machines with CPU topologies that are otherwise not possible, such as when the number of virtual cores is between the number of host cores and the number of host threads.

- **For best performance, and especially for CPU-intensive workloads**, you should use the same topology in the virtual machine as in the host, so the host and the virtual machine expect the same cache usage. When the host has hyperthreading enabled, QEMU treats the host's hyperthreads as cores, so the virtual machine is not aware that it is running on a single core with multiple threads. This behavior might impact the performance of a virtual machine, because a virtual core that actually corresponds to a hyperthread in the host core might share a single cache with another hyperthread in the same host core, while the virtual machine treats it as a separate core.

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	**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	Selecting the **Count Threads As Cores** check box enables hosts to run virtual machines with a total number of processor cores greater than the number of cores in the host. When this check box is selected, the exposed host threads are treated as cores that virtual machines can use. 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.

Field	Description/Action
Memory Balloon	Selecting the **Enable Memory Balloon Optimization** check box enables memory overcommitment on virtual machines running on the hosts in this cluster. When this check box is selected, the Memory Overcommit Manager (MoM) starts ballooning where and when possible, with a limitation of the guaranteed memory size of every virtual machine. 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 5.2.9, "Updating the MoM Policy on Hosts in a Cluster". 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.
KSM control	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.

5.2.4. 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	Legacy behavior of 3.6 version. Overrides in **vdsm.conf** are still applied. The guest agent hook mechanism is disabled.

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.
Post-copy migration	This is a Technology Preview feature. Virtual machines should not experience any significant downtime similar to the minimal downtime policy. The migration will switch to post-copy if the virtual machine migration does not converge after a long time. The disadvantage of this policy is that in the post-copy phase, the virtual machine may slow down significantly as the missing parts of memory are transferred between the hosts. If anything goes wrong during the post-copy phase, such as a network failure between the hosts, then the running virtual machine instance will be lost. It is therefore not possible to abort a migration during the post-copy phase. 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.

Policy	Description
Custom	Defined by user (in Mbps). This value is divided by the number of concurrent migrations (default is 2, to account for ingoing and outgoing migration). Therefore, the user-defined bandwidth must be large enough to accommodate all concurrent migrations. For example, if the **Custom** bandwidth is defined as 600 Mbps, a virtual machine migration's maximum bandwidth is actually 300 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.
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

Property	Description
Auto Converge migrations	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. • 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.
Enable migration compression	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. • 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.

5.2.5. 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.

Table 5.7. Scheduling Policy Tab Properties

Field	Description/Action
Select Policy	Select a policy from the drop-down list. **none**: Set the policy value to **none** to have no load-balancing 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 **CpuOverCommitDurationMinutes**, **HighUtilization**, or **MaxFreeMemoryForOverUtilized**. **cluster_maintenance**: Limits activity in a cluster during maintenance tasks. No new virtual machines may be started, except highly available virtual machines. If host failure occurs, highly available virtual machines will restart properly and any virtual machine can migrate. **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	Description/Action
	The following properties appear depending on the selected policy, and can be edited if necessary: • **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 the SPM host can run in comparison to other hosts. 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 it again when load balancing requires or there are not enough free hosts in the cluster. The default value is **20**. • **ScaleDown**: Reduces the impact of the **HA Reservation** weight function, by dividing a host's score by the specified amount. This is

Field	Description/Action
	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 **MaxFreeMemoryForOverUtilized** and **MinFreeMemoryForUnderUtilized** to 0MB 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.

- **HeSparesCount**: Sets the number of additional self-hosted engine nodes that must reserve enough free memory to start the Manager virtual machine if it migrates or shuts down. Other virtual machines are prevented from starting on a self-hosted engine node if doing so would not leave enough free memory for the Manager virtual machine. This is an optional property that can be added to the **power_saving**,

Field	Description/Action
	vm_evenly_distributed, and _evenly_distributed policies. The default value is **0**.
Scheduler Optimization	Optimize scheduling for host weighing/ordering. • **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	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 9.9, "Trusted Compute Pools".
Enable HA Reservation	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.
Provide custom serial number policy	This check box allows you to specify a serial number policy for the virtual machines in the cluster. Select one of the following options: • **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.

5.2.6. 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 VM Portal) is outside of the network where the hypervisors reside.
Overridden SPICE proxy address	The proxy by which the SPICE client connects to virtual machines. The address must be in the following format: protocol://[host]:[port]

5.2.7. 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.
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**.

Field	Description/Action
Skip fencing if gluster bricks are up	This option is only available when Red Hat Gluster Storage functionality is enabled. If this check box is selected, fencing is skipped if bricks are running and can be reached from other peers. See Chapter 2. Configure High Availability using Fencing Policies and Appendix A. Fencing Policies for Red Hat Gluster Storage in *Maintaining Red Hat Hyperconverged Infrastructure* for more information.
Skip fencing if gluster quorum not met	This option is only available when Red Hat Gluster Storage functionality is enabled. If this check box is selected, fencing is skipped if bricks are running and shutting down the host will cause loss of quorum. See Chapter 2. Configure High Availability using Fencing Policies and Appendix A. Fencing Policies for Red Hat Gluster Storage in *Maintaining Red Hat Hyperconverged Infrastructure* for more information.

5.2.8. 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 5.2.5, "Scheduling Policy Settings Explained".

Setting Load and Power Management Policies for Hosts

1. Click **Compute → Clusters** and select a cluster.

2. Click **Edit**.

3. Click the **Scheduling Policy** tab.

4. 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.

     d. Optionally, in the **HeSparesCount** field, enter the number of additional self-hosted engine nodes on which to reserve enough free memory to start the Manager virtual

machine if it migrates or shuts down. See Configuring Memory Slots Reserved for the Self-Hosted Engine on Additional Hosts in the *Self-Hosted Engine Guide* for more information.

- **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**.

  e. Optionally, in the **HeSparesCount** field, enter the number of additional self-hosted engine nodes on which to reserve enough free memory to start the Manager virtual machine if it migrates or shuts down. See Configuring Memory Slots Reserved for the Self-Hosted Engine on Additional Hosts in the *Self-Hosted Engine Guide* for more information.

- **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**.

  f. Optionally, in the **HeSparesCount** field, enter the number of additional self-hosted engine nodes on which to reserve enough free memory to start the Manager virtual machine if it migrates or shuts down. See Configuring Memory Slots Reserved for the Self-Hosted Engine on Additional Hosts in the *Self-Hosted Engine Guide* for more information.

5. 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.

6. 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.

7. Optionally select the **Enable HA Reservation** check box to enable the Manager to monitor cluster capacity for highly available virtual machines.

8. 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.

9. Click **OK**.

5.2.9. 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.

Synchronizing MoM Policy on a Host

1. Click **Compute → Clusters**.

2. Click the cluster's name to open the details view.

3. Click the **Hosts** tab and select the host that requires an updated MoM policy.

4. Click **Sync MoM Policy**.

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

5.2.10. Creating a CPU Profile

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.

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

Creating a CPU Profile

1. Click **Compute → Clusters**.

2. Click the cluster's name to open the details view.

3. Click the **CPU Profiles** tab.

4. Click **New**.

5. Enter a **Name** and a **Description** for the CPU profile.

6. Select the quality of service to apply to the CPU profile from the **QoS** list.

7. Click **OK**.

5.2.11. Removing a CPU Profile

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

Removing a CPU Profile

1. Click **Compute** → **Clusters**.

2. Click the cluster's name to open the details view.

3. Click the **CPU Profiles** tab and select the CPU profile to remove.

4. Click **Remove**.

5. Click **OK**.

If the CPU profile was assigned to any virtual machines, those virtual machines are automatically assigned the **default** CPU profile.

5.2.12. 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.

Importing an Existing Red Hat Gluster Storage Cluster to the Red Hat Virtualization Manager

1. Click **Compute** → **Clusters**.

2. Click **New**.

3. Select the **Data Center** the cluster will belong to.

4. Enter the **Name** and **Description** of the cluster.

5. Select the **Enable Gluster Service** check box and the **Import existing gluster configuration** check box.
   The **Import existing gluster configuration** field is only displayed if the **Enable Gluster Service** is selected.

6. In the **Hostname** field, enter the host name or IP address of any server in the cluster.

The host **SSH 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 **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.

    Verify that 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.

5.2.13. 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.

5.2.14. Removing a Cluster

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.

Removing a Cluster

1. Click **Compute → Clusters** and select a cluster.

2. Ensure there are no hosts in the cluster.

3. Click **Remove**.

4. Click **OK**

5.2.15. 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.

> **IMPORTANT**
>
> 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. Check if there is an icon next to the host indicating an update is available. See Update the Hosts in the *Upgrade Guide* for more information on updating hosts.

Procedure

1. Click **Compute → Clusters** and select the cluster to change.

2. Click **Edit**.

3. Change the **Compatibility Version** to the desired value.

4. Click **OK** to open the **Change Cluster Compatibility Version** confirmation window.

5. Click **OK** to confirm.
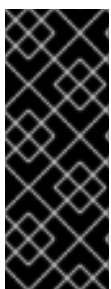
> **IMPORTANT**
>
> An error message may warn that some virtual machines and templates are incorrectly configured. To fix this error, edit each virtual machine manually. The **Edit Virtual Machine** window provides additional validations and warnings that show what to correct. Sometimes the issue is automatically corrected and the virtual machine's configuration just needs to be saved again. After editing each virtual machine, you will be able to change the cluster compatibility version.

After you update the cluster's compatibility version, you must update the cluster compatibility version of all running or suspended virtual machines by restarting them from within the Manager, or using the REST API, 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 pending changes icon ( ). You cannot change the cluster compatibility version of a virtual machine snapshot that is in preview; you must first commit or undo the preview.

The self-hosted engine virtual machine does not need to be restarted.

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.

CHAPTER 6. LOGICAL NETWORKS

6.1. LOGICAL NETWORK TASKS

6.1.1. Performing Networking Tasks

Network → Networks 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. 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 tabs in the details view 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.

> ⚠ **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.

Creating a New Logical Network in a Data Center or Cluster

1. Click **Compute → Data Centers** or **Compute → Clusters**.

2. Click the data center or cluster name to open the details view.

3. Click the **Logical Networks** tab.

4. Open the **New Logical Network** window:

   - In a data center details view, click **New**.

   - In a cluster details view, click **Add Network**.

5. Enter a **Name**, **Description**, and **Comment** for the logical network.

6. Optionally enable **Enable VLAN tagging**.

7. Optionally disable **VM Network**.

8. Optionally select the **Create on external provider** check box. This disables the **Network Label**, **VM Network**, and **MTU** options. See Chapter 11, *External Providers* for details.

9. Select the **External Provider**. The **External Provider** list does not include external providers that are in **read-only** mode.
   You can create an internal isolated network, by selecting **ovirt-provider-ovn** on the **External Provider** list and leaving **Connect to physical network** unselected.

10. Enter a new label or select an existing label for the logical network in the **Network Label** text field.

11. Set the **MTU** value to **Default (1500)** or **Custom**.

12. In 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.

13. 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.

14. In the **vNIC Profiles** tab, add vNIC profiles to the logical network as required.

15. Click **OK**.

If you entered a label for the logical network, it is 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
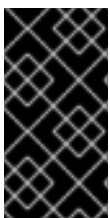
> **IMPORTANT**
>
> A logical network cannot be edited or moved to another interface if it is not synchronized with the network configuration on the host. See Section 6.4.2, "Editing Host Network Interfaces and Assigning Logical Networks to Hosts" on how to synchronize your networks.

Editing a Logical Network

1. Click **Compute → Data Centers**.

2. Click the data center's name to open the details view.

3. Click the **Logical Networks** tab and select a logical network.

4. Click **Edit**.

5. 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.

6. Click **OK**.

> **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.

> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon ⟳ appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>
>   - In the host's **Status** column in the **Compute → Hosts** window.
>
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

6.1.4. Removing a Logical Network

You can remove a logical network from **Network → Networks** or **Compute → Data Centers**. 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.

Removing Logical Networks

1. Click **Compute** → **Data Centers**.

2. Click a data center's name to open the details view.

3. Click the **Logical Networks** tab to list the logical networks in the data center.

4. Select a logical network and click **Remove**.

5. 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.

6. Click **OK**.

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

6.1.5. Configuring a Non-Management Logical Network as the Default Route

The default route used by hosts in a cluster is through the management network (**ovirtmgmt**). The following procedure provides instructions to configure a non-management logical network as the default route.

Prerequisite:

- If you are using the **default_route** custom property, you need to clear the custom property from all attached hosts and then follow this procedure.

Configuring the Default Route Role

1. Click **Network** → **Networks**.

2. Click the name of the non-management logical network to configure as the default route to access its details.

3. Click the **Clusters** tab.

4. Click **Manage Network** to open the **Manage Network** window.

5. Select the **Default Route** checkbox for the appropriate cluster(s).

6. Click **OK**.

When networks are attached to a host, the default route of the host will be set on the network of your choice. It is recommended to configure the default route role before any host is added to your cluster. If your cluster already contains hosts, they may become out-of-sync until you sync your change to them.

6.1.6. 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.

Viewing or Editing the Gateway for a Logical Network

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **Network Interfaces** tab to list the network interfaces attached to the host, and their configurations.

4. Click **Setup Host Networks**.

5. 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.

6.1.7. 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
Name	The name of the logical network. This text field must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores. Note that while the name of the logical network can be longer than 15 characters and can contain non-ASCII characters, the on-host identifier (*vdsm_name*) will differ from the name you defined. See Mapping VDSM Names to Logical Network Names for instructions on displaying a mapping of these names.
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.

Field Name	Description
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.

6.1.8. 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

Field Name	Description
Attach/Detach Network to/from Cluster(s)	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.
	Name – the name of the cluster to which the settings will apply. This value cannot be edited.
	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.
	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.

6.1.9. 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	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.
	Public – Allows you to specify whether the profile is available to all users.
	QoS – Allows you to specify a network quality of service (QoS) profile to the vNIC profile.

6.1.10. 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.

Specifying Traffic Types for Logical Networks

1. Click **Compute → Clusters**.

2. Click the cluster's name to open the details view.

3. Click the **Logical Networks** tab.

4. Click **Manage Networks**.

5. Select the appropriate check boxes and radio buttons.

6. Click **OK**.

> **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.

6.1.11. 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. If an outage occurs on this network, the management network (**ovirtmgmt** by default) will be used instead.

6.1.12. 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" .

Editing the Virtual Function Configuration on a NIC

1. Click **Compute → Hosts**.

2. Click the name of an SR-IOV-capable host to open the details view.

3. Click the **Network Interfaces** tab.

4. Click **Setup Host Networks**.

5. Select an SR-IOV-capable NIC, marked with a , and click the pencil icon.

6. 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.

7. 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.

8. Click **OK**.

9. In the **Setup Host Networks** window, click **OK**.

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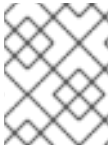
> **NOTE**
>
> If you are enabling or disabling port mirroring, all virtual machines using the associated profile must be in a down state before editing.

Creating or Editing a vNIC Profile

1. Click **Network → Networks**.

2. Click the logical network's name to open the details view.

3. Click the **vNIC Profiles** tab.

4. Click **New** or **Edit**.

5. Enter the **Name** and **Description** of the profile.

6. Select the relevant Quality of Service policy from the **QoS** list.

7. 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* .

8. 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" .

9. If **Passthrough** is selected, optionally deselect the **Migratable** check box to disable migration for vNICs using this profile. If you keep this check box selected, see Additional Prerequisites for Virtual Machines with SR-IOV-Enabled vNICs in the *Virtual Machine Management Guide* .

10. Use the **Port Mirroring** and **Allow all users to use this Profile** check boxes to toggle these options.

11. 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.

12. Click **OK**.

Apply this profile to users and groups to regulate their network bandwidth. If you edited a vNIC profile, you must either restart the virtual machine, or hot unplug and then hot plug the vNIC if the guest operating system supports 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	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*Red Hat Enterprise Linux Virtualization Deployment and Administration Guide*. Use **<No Network Filter>** for virtual machine VLANs and bonds. On trusted virtual machines, choosing not to use a network filter can improve performance. **NOTE** Red Hat no longer supports disabling filters by setting the **EnableMACAntiSpoofingFilterRules** parameter to **false** using the **engine-config** tool. Use the**<No Network Filter>** option instead.
Passthrough	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. QoS, network filters, and port mirroring are disabled in the vNIC profile if passthrough is enabled.

Field Name	Description
Migratable	A check box to toggle whether or not vNICs using this profile can be migrated. Migration is enabled by default on regular vNIC profiles; the check box is selected and cannot be changed. When the **Passthrough** check box is selected,**Migratable** becomes available and can be deselected, if required, to disable migration of passthrough vNICs.
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 it not selected by default. For further details, see Port Mirroring in the *Technical Reference*.
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, network filters, and port mirroring cannot be enabled on the same 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 .

Enabling Passthrough

1. Click **Network → Networks**.

2. Click the logical network's name to open the details view.

3. Click the **vNIC Profiles** tab to list all vNIC profiles for that logical network.

4. Click **New**.

5. Enter the **Name** and **Description** of the profile.

6. Select the **Passthrough** check box.

7. Optionally deselect the **Migratable** check box to disable migration for vNICs using this profile. If you keep this check box selected, see Additional Prerequisites for Virtual Machines with SR-IOV-Enabled vNICs in the *Virtual Machine Management Guide*.

8. 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.

9. Click **OK**.

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 **PCI Passthrough** vNIC on the desired virtual machine that uses the passthrough vNIC profile. For more information on these procedures respectively, see Section 6.4.2, "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.

Removing a vNIC Profile

1. Click **Network → Networks**.

2. Click the logical network's name to open the details view.

3. Click the **vNIC Profiles** tab to display available vNIC profiles.

4. Select one or more profiles and click **Remove**.

5. Click **OK**.

6.2.6. Assigning Security Groups to vNIC Profiles

> **NOTE**
>
> This feature is only available when OpenStack Networking (neutron) is added as an external network provider. Security groups cannot be created through the Red Hat Virtualization Manager. You must create security groups through OpenStack. For more information, see Project Security Management 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
> ```

Assigning Security Groups to vNIC Profiles

1. Click **Network → Networks**.

2. Click the logical network's name to open the details view.

3. Click the **vNIC Profiles** tab.

4. Click **New**, or select an existing vNIC profile and click **Edit**.

5. 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.

6. In the text field, enter the ID of the security group to attach to the vNIC profile.

7. 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.

User Permissions for vNIC Profiles

1. Click **Network → vNIC Profile**.

2. Click the vNIC profile's name to open the details view.

3. Click the **Permissions** tab to show the current user permissions for the profile.

4. Click **Add** or **Remove** 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 data center aspects such as computing, networking and storage resources.

The **vdsm-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 **vdsm-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.
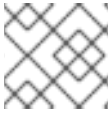
> **NOTE**
>
> A UCS port profile must be configured in Cisco UCS before configuring the vNIC profile.

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 **ovirt-engine** service.

   ```
   # 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.

Configuring a vNIC Profile for UCS Integration

1. Click **Network → Networks**.

2. Click the logical network's name to open the details view.

3. Click the **vNIC Profiles** tab.

4. Click **New**, or select a vNIC profile and click **Edit**.

5. Enter the **Name** and **Description** of the profile.

6. Select the **vmfex** custom property from the custom properties list and enter the UCS port profile name.

7. 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] or xref: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.

Importing a Network From an External Provider

1. Click **Network → Networks**.

2. Click **Import**.

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 to prevent that network from being available to all users.

8. Click **Import**.

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

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.

6.3.4. Adding Subnets to External Provider Logical Networks

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

Adding Subnets to External Provider Logical Networks

1. Click **Network → Networks**.

2. Click the logical network's name to open the details view.

3. Click the **Subnets** tab.

4. Click **New**.

5. Enter a **Name** and **CIDR** for the new subnet.

6. From the **IP Version** drop-down list, select either **IPv4** or **IPv6**.

7. Click **OK**.

6.3.5. Removing Subnets from External Provider Logical Networks

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

Removing Subnets from External Provider Logical Networks

1. Click **Network → Networks**.

2. Click the logical network's name to open the details view.

3. Click the **Subnets** tab.

4. Select a subnet and click **Remove**.

5. Click **OK**.

6.4. HOSTS AND NETWORKING

6.4.1. 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.

Refreshing Host Capabilities

1. Click **Compute → Hosts** and select a host.

2. Click **Management → Refresh Capabilities**.

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

> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon ⟳ appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>   - In the host's **Status** column in the **Compute → Hosts** window.
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

6.4.2. 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, see Section 6.4.4, "Editing a Host's VLAN Settings".
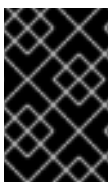
> **IMPORTANT**
>
> You cannot assign logical networks offered by external providers to physical host network interfaces; such networks are dynamically assigned to hosts as they are required by virtual machines.

NOTE

If the switch has been configured to provide Link Layer Discovery Protocol (LLDP) information, you can hover your cursor over a physical network interface to view the switch port's current configuration. This can help to prevent incorrect configuration. Red Hat recommends checking the following information prior to assigning logical networks:

- **Port Description (TLV type 4)** and **System Name (TLV type 5)** help to detect to which ports and on which switch the host's interfaces are patched.

- **Port VLAN ID** shows the native VLAN ID configured on the switch port for untagged ethernet frames. All VLANs configured on the switch port are shown as **VLAN Name** and **VLAN ID** combinations.

Editing Host Network Interfaces and Assigning Logical Networks to Hosts

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **Network Interfaces** tab.

4. Click **Setup Host Networks**.

5. Optionally, hover your cursor over host network interface to view configuration information provided by the switch.

6. 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.

7. 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

      If you change the host's management network IP address, you must reinstall the host for the new IP address to be configured.

      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.

      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.

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=** 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.
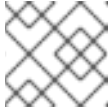
> **NOTE**
>
> A separate, dedicated logical network is recommended for use with FCoE.

g. To change the default network used by the host from the management network (ovirtmgmt) to a non-management network, configure the non-management network's default route. See Section 6.1.5, "Configuring a Non-Management Logical Network as the Default Route" for more information.

h. If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box. For more information about unsynchronized hosts and how to synchronize them, see Section 6.4.3, "Synchronizing Host Networks".

8. 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.

9. Select the **Save network configuration** check box to make the changes persistent when the environment is rebooted.

10. Click **OK**.

> **NOTE**
>
> If not all network interface cards for the host are displayed, click **Management → Refresh Capabilities** to update the list of network interface cards available for that host.
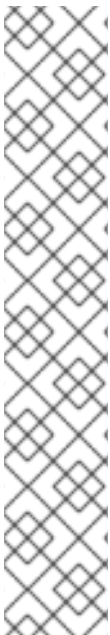
> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon ⟳ appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>   - In the host's **Status** column in the **Compute → Hosts** window.
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

6.4.3. Synchronizing Host Networks

The Manager defines a network interface as **out-of-sync** when the definition of the interface on the host differs from the definitions stored by the Manager. Out-of-sync networks appear with an Out-of-sync icon in the host's **Network Interfaces** tab and with this icon in the **Setup Host Networks** window.

When a host's network is out of sync, the only activities that you can perform on the unsynchronized network in the **Setup Host Networks** window are detaching the logical network from the network interface or synchronizing the network.

Understanding How a Host Becomes out-of-sync

A host will become out of sync if:

- You make configuration changes on the host rather than using the the **Edit Logical Networks** window, for example:

  - Changing the VLAN identifier on the physical host.

  - Changing the **Custom MTU** on the physical host.

- You move a host to a different data center with the same network name, but with different values/parameters.

- You change a network's **VM Network** property by manually removing the bridge from the host.

- You update definitions using the **Setup Host Networks** window, without selecting the **Save network configuration** check box when saving your changes. After rebooting the host, it may become unsynchronized.

Preventing Hosts from Becoming Unsynchronized

Following these best practices will prevent your host from becoming unsynchronized:

1. Ensure that the **Save network configuration** check box is selected when saving your changes in the **Setup Host Networks** window (it is selected by default).

2. Use the Administration Portal to make changes rather than making changes locally on the host.

3. Edit VLAN settings according to the instructions in Section 6.4.4, "Editing a Host's VLAN Settings".

Synchronizing Hosts

Synchronizing a host's network interface definitions involves using the definitions from the Manager and applying them to the host. If these are not the definitions that you require, after synchronizing your hosts update their definitions from the Administration Portal. You can synchronize a host's networks on three levels:

- Per logical network

- Per host

- Per cluster

> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon ⟳ appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>
>   - In the host's **Status** column in the **Compute → Hosts** window.
>
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

Synchronizing Host Networks on the Logical Network Level

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **Network Interfaces** tab.

4. Click **Setup Host Networks**.

5. Hover your cursor over the unsynchronized network and click the pencil icon to open the **Edit Network** window.

6. Select the **Sync network** check box.

7. Click **OK**.

8. Select the **Save network configuration** check box in the **Setup Host Networks** window to make the changes persistent when the environment is rebooted.

9. Click **OK**.

Synchronizing a Host's Networks on the Host level

- Click the **Sync All Networks** button in the host's **Network Interfaces** tab to synchronize all of the host's unsynchronized network interfaces.

Synchronizing a Host's Networks on the Cluster level

- Click the **Sync All Networks** button in the cluster's **Logical Networks** tab to synchronize all unsynchronized logical network definitions for the entire cluster.
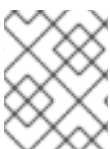
> **NOTE**
>
> You can also synchronize a host's networks via the REST API. See syncallnetworks in the *REST API Guide*.

6.4.4. Editing a Host's VLAN Settings

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:

1. Put the host in maintenance mode.

2. Manually remove the management network from the host. This will make the host reachable over the new VLAN.

3. 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

If you change the management network's VLAN ID, you must reinstall the host to apply the new VLAN ID.

6.4.5. 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.
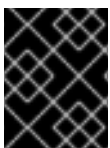
IMPORTANT

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.

Adding Multiple VLANs to a Network Interface using Logical Networks

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **Network Interfaces** tab.

4. Click **Setup Host Networks**.

5. 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.

6. Edit the logical networks:

   a. Hover your cursor over an assigned logical network and click the pencil icon.

a. Hover your cursor over an assigned logical network and click the pencil icon.

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

c. Select a **Boot Protocol**:

- None

- DHCP

- Static

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

e. Click **OK**.

7. 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.

8. Select the **Save network configuration** check box.

9. 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.

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.

> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon ⟳ appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>
>   - In the host's **Status** column in the **Compute → Hosts** window.
>
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

6.4.6. 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 **vdsm-hook-extra-ipv4-addrs** 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.

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 vdsm-hook-extra-ipv4-addrs
   ```

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

   ```
   # engine-config -s 'UserDefinedNetworkCustomProperties=ipv4_addrs=.*'
   ```

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

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

4. In the Administration Portal, click **Compute → Hosts**.

5. Click the host's name to open the details view.

6. Click the **Network Interfaces** tab and click **Setup Host Networks**.

7. Edit the host network interface by hovering the cursor over the assigned logical network and clicking the pencil icon.

8. 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.

9. Click **OK**.

10. Select the **Save network configuration** check box.

11. 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.

> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>
>   - In the host's **Status** column in the **Compute → Hosts** window.
>
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

6.4.7. 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.

There are two methods of adding labels to a host network interface:

- Manually, from the **Setup Host Networks** window.

- Automatically, using LLDP Labeler.

**Manually Adding Network Labels to Host Network Interfaces**

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **Network Interfaces** tab.

4. Click **Setup Host Networks**.

5. Click **Labels** and right-click **[New Label]**. Select a physical network interface to label.

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

7. Click **OK**.

Automatically Adding Network Labels to Host Network Interfaces

The LLDP Labeler service enables you to automate the process of assigning labels to host network interfaces in the configured list of clusters.

By default, LLDP Labeler is defined to run as an hourly service. This option is useful if you make hardware changes, for example, making changes to NICs, switches, cables, or if you change switch configurations.

Prerequisites

- The interfaces must be connected to a Juniper switch.

- The Juniper switch must be configured to provide the **Port VLAN** using LLDP.

Procedure

1. Configure the Manager's username and password by opening **/etc/ovirt-lldp-labeler/conf.d/ovirt-lldp-credentials.conf** in a text editor and updating the following values:

   - **username** – the username of the Manager's administrator. The default is **admin@internal**.

   - **password** – the password of the Manager's administrator. The default is **123456**.

2. Configure the LLDP Labeler service by opening **etc/ovirt-lldp-labeler/conf.d/ovirt-lldp-credentials.conf** in a text editor and updating the following values:

   - **clusters** – a comma-separated list of clusters on which the service should run. Wildcards are supported. For example, **Cluster*** defines LLDP Labeler to run on all clusters starting with word **Cluster**. To run the service on all clusters in the data center, type *****. The default is **Def***.

   - **api_url** – the full URL of the Manager's API. The default is **https://ovirt-engine/ovirt-engine/api**

   - **ca_file** – the path to the custom certificate file. Leave this value empty if you do not use custom certificates. The default is empty.

   - **auto_bonding** – enables LLDP Labeler's bonding capabilities. The default is **true**.

   - **auto_labeling** – enables LLDP Labeler's labeling capabilities. The default is **true**.

3. Optional. Configure the service to run at a different time interval, by editing **etc/ovirt-lldp-labeler/conf.d/ovirt-lldp-labeler.timer** in a text editor and changing the value of **OnUnitActiveSec**. The default is **1h**.

4. Enable the service by running:

   ```
   # systemctl enable --now ovirt-lldp-labeler
   ```

5. Optional. To invoke the service manually, run:

   ```
   # /usr/bin/python /usr/share/ovirt-lldp-labeler/ovirt_lldp_labeler_cli.py
   ```

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.

> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon ⟳ appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>
>   - In the host's **Status** column in the **Compute → Hosts** window.
>
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

6.4.8. Changing the FQDN of a Host

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

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 7.5.14, "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.

6.5. BOND DEVICES

6.5.1. Bonding Methods

You can bond compatible network devices together. This type of configuration can increase available bandwidth and reliability.

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.
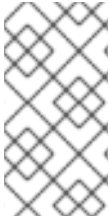
NOTE

For a bond in Mode 4, all slaves must be configured properly on the switch. Otherwise, the **ad_partner_mac** is 00:00:00:00:00:00 and the Manager displays a warning exclamation mark icon on the bond in the **Network Interfaces** tab. No warning is provided if any of the slaves are up and running.

There are two methods for creating bond devices:

- Manually, from the Administration Portal, on a specific host. See Section 6.5.2, "Creating a Bond Device Using the Administration Portal".

- Automatically, using the LLDP Labeler service, for non-bonded NICs for all cluster hosts. See Section 6.5.3, "Creating a Bond Device Automatically".

6.5.2. Creating a Bond Device Using the Administration Portal

Using the Administration Portal, 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-tagged traffic.
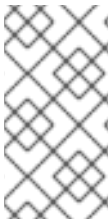
NOTE

If the switch has been configured to provide Link Layer Discovery Protocol (LLDP) information, you can hover your cursor over a physical network interface to view the switch port's current aggregation configuration. Red Hat recommends checking the configuration prior to creating a bond device.

Procedure

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **Network Interfaces** tab to list the physical network interfaces attached to the host.

4. Click **Setup Host Networks**.

5. Optionally, hover your cursor over host network interface to view configuration information provided by the switch.

6. 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. For information about bonding logic, see Section 6.5.4, "Bonding Logic in Red Hat Virtualization".

7. Select the **Bond Name** and **Bonding Mode** from the drop-down menus.
   You can select bonding modes 1, 2, 4, and 5. Any other mode can be configured using the **Custom** option. For more information about bond modes, see Section 6.5.5, "Bonding Modes".

8. Click **OK** to create the bond and close the **Create New Bond** window.

9. Assign a logical network to the newly created bond device.

10. Optionally, select **Verify connectivity between Host and Engine** and/or **Save network configuration**.

11. Click **OK**.

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 for the selected host.
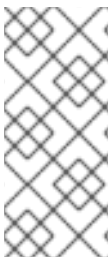
> **NOTE**
>
> While the host's configuration is being updated, an indication of this status appears as follows:
>
> - An Updating icon   appears below each network interface in the host's **Network Interfaces** tab.
>
> - The status **Networks updating** appears:
>   - In the host's **Status** column in the **Compute → Hosts** window.
>   - In the host's **Status** column in the **Hosts** tab that you access when selecting a cluster in the **Compute > Cluster** window.
>   - In the host's **Network Device Status** column in the **Hosts** tab that you access when selecting a network in the **Network > Networks** window.

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.

> **NOTE**
>
> For a bond in Mode 4, all slaves must be configured properly on the switch. If none of them is configured properly on the switch, the **ad_partner_mac** is reported as 00:00:00:00:00:00. The Manager will display a warning in the form of an exclamation mark icon on the bond in the **Network Interfaces** tab. No warning is provided if any of the slaves are up and running.

6.5.3. Creating a Bond Device Automatically

Red Hat enables you to automate the bonding process for non-bonded NICs, for one or more clusters, or for the entire data center using the LLDP Labeler. The bond is created using bonding mode 4. For more information about bond modes, see Section 6.5.5, "Bonding Modes".

Bonding Devices Automatically

By default, LLDP Labeler is defined to run as an hourly service. This option is useful if you make hardware changes, for example, making changes to NICs, switches, cables, or if you change switch configurations.

Prerequisites

- The interfaces must be connected to a Juniper switch.

- The Juniper switch must be configured for Link Aggregation Control Protocol (LACP) using LLDP.

Procedure

1. Configure the Manager's username and password by opening **/etc/ovirt-lldp-labeler/conf.d/ovirt-lldp-credentials.conf** in a text editor and updating the following values:

    - **username** – the username of the Manager's administrator. The default is **admin@internal**.

    - **password** – the password of the Manager's administrator. The default is **123456**.

2. Configure the LLDP Labeler service by opening **etc/ovirt-lldp-labeler/conf.d/ovirt-lldp-credentials.conf** in a text editor and updating the following values:

    - **clusters** – a comma-separated list of clusters on which the service should run. Wildcards are supported. For example, **Cluster*** defines LLDP Labeler to run on all clusters starting with word **Cluster**. To run the service on all clusters in the data center, type *****. The default is **Def***.

    - **api_url** – the full URL of the Manager's API. The default is **https://ovirt-engine/ovirt-engine/api**

    - **ca_file** – the path to the custom certificate file. Leave this value empty if you do not use custom certificates. The default is empty.

    - **auto_bonding** – enables LLDP Labeler's bonding capabilities. The default is **true**.

    - **auto_labeling** – enables LLDP Labeler's labeling capabilities. The default is **true**.

3. Optional. Configure the service to run at a different time interval, by editing **etc/ovirt-lldp-labeler/conf.d/ovirt-lldp-labeler.timer** in a text editor and changing the value of **OnUnitActiveSec**. The default is **1h**.

4. Enable the service by running:

    ```
    # systemctl enable --now ovirt-lldp-labeler
    ```

5. Optional. To invoke the service manually, run:

    ```
    # /usr/bin/python /usr/share/ovirt-lldp-labeler/ovirt_lldp_labeler_cli.py
    ```

If the devices are incompatible, the NICs that violate these rules will not be bonded. For information about bonding logic, see Section 6.5.4, "Bonding Logic in Red Hat Virtualization".

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 for the selected host. If the NICs were not already connected to logical networks, assign a logical network to the newly created bond device. See Section 6.4.2, "Editing Host Network Interfaces and Assigning Logical Networks to Hosts" for details.

6.5.4. Bonding Logic in Red Hat Virtualization

There are several distinct bond creation scenarios, each with its own logic.

Two factors that affect bonding logic are:

- Does either device already carry logical networks?

- Are the devices carrying compatible logical networks?

NOTE

If multiple logical networks are connected to a NIC, only one of the networks can be non-VLAN. All remaining logical networks must have unique VLANs.

In addition, the NICs must be connected to the same port on the switch.

NOTE

If your environment uses iSCSI storage and you want to implement redundancy, follow the instructions for configuring multipathing for iSCSI .

Table 6.6. Bonding Scenarios, Results, and Creation Method

Bonding Scenario	Result	Method
NIC + NIC	The **Create New Bond** window is displayed, and you can configure a new bond device. 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.	Adminstration Portal or LLDP Labeler
NIC + Bond	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. 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.	Administration Portal

Bonding Scenario	Result	Method
Bond + Bond	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. 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.	Administration Portal

6.5.5. Bonding Modes

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.

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 ends with the last available network interface card in the bond. All subsequent loops 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. Therefore, it is not compatible with virtual machine logical networks.

Mode 1 (active-backup policy)

One network interface card is active, while all the other network interface cards are in a backup state. If the active network interface card fails, 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

this bond is visible only on the network adapter port in order to prevent confusion that might occur if the MAC address of the bond were to change, reflecting the MAC address of the new 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 using the result of the following operation: (**XOR** the source MAC address with the destination MAC address) **modulo** network interface card count. This calculation ensures that the same network interface card is selected for each destination MAC address. 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 (dynamic link aggregation 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 that the outward traffic is distributed, based on the load, over all the network interface cards in the bond and that the incoming traffic is received by the active network interface card. If the network interface card receiving incoming traffic fails, another network interface card is assigned. 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 and has no 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. 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
```

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.
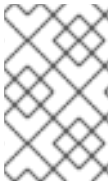
> **NOTE**
>
> You can identify the host type of an individual host in the Red Hat Virtualization Manager by selecting the host's name to open the details view, and checking the **OS Description** under **Software**.

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 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 view. 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:

https://www.redhat.com/mailman/listinfo/rhsa-announce

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**. Using the **yum** command is the only way to install additional packages and have them persist after an upgrade.

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 configuring networking and deploying a self-hosted engine, and can also be used to run terminal commands via the **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.

RHVH uses the Automatic Bug Reporting Tool (ABRT) to collect meaningful debug information about application crashes. For more information, see the *Red Hat Enterprise Linux System Administrator's Guide*.

> **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 **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** and **Red Hat Virtualization** subscriptions.

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

Optionally, you can install a Cockpit user interface for monitoring the host's resources and performing administrative tasks. 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 configuring networking and deploying a self-hosted engine, and can also be used to run terminal commands via the **Terminal** sub-tab.
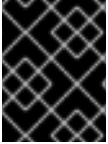
> **IMPORTANT**
>
> Third-party watchdogs should not be installed on Red Hat Enterprise Linux hosts, as they can interfere with the watchdog daemon provided by VDSM.

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, and creation of a bridge. Use the details view to monitor the process as the host and the Manager establish a connection.

Adding a Host to the Red Hat Virtualization Manager

1. Click **Compute → Hosts**.

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 **Hostname** 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 Public Key** 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 see the progress of the installation in the details view. After a brief delay the host status changes to **Up**.

> **IMPORTANT**
>
> Keep the environment up-to-date. See https://access.redhat.com/articles/2974891 for more information. Since bug fixes for known issues are frequently released, Red Hat recommends using scheduled tasks to update the hosts and the Manager.

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.

Adding a Satellite Host Provider Host

1. Click **Compute → Hosts**.

2. Click **New**.

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** 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 view. 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 4.2 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.

Configuring Satellite Errata Management for a Host

1. Add the Satellite server as an external provider. See Section 11.2.1, "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 host registration see Configuring a Host for Registration in the *Red Hat Satellite User Guide*. 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. Click **Compute → Hosts** and select the host.

    b. Click **Edit**.

    c. Select the **Use Foreman/Satellite** check box.

    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

7.5.5. 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
Host Cluster	The cluster and data center to which the host belongs.
Use Foreman/Satellite	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: **Discovered Hosts** • **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. **Provisioned Hosts** • **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 host. 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.
Hostname	The IP address or resolvable host name 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 Public Key	Copy the contents in the text box to the **/root/.ssh/authorized_hosts** file on the host to use the Manager's SSH key instead of a password to authenticate with a 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**.

7.5.6. 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. In Red Hat Enterprise Linux 7.1 and later, 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	Lists the host's fence agents. Fence agents can be sequential, concurrent, or a mix of both. • 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. Fence agents are sequential by default. Use the up and down buttons to change the sequence in which the fence agents are used. 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.
Add Fence Agent	Click the **+** 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.
Power Management Proxy Preference	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**.

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

Table 7.3. Edit fence agent Settings

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

Field Name	Description
Password	Password for the user accessing the power management device.
Type	The type of power management device in your host. Choose one of the following: • **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. For more information about power management devices, see Power Management in the *Technical Reference*.
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.
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**.

Field Name	Description
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.

7.5.7. 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.

7.5.8. 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.

7.5.9. 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.

7.5.10. 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. If you change any kernel command line parameters, you must reinstall the host.

> **IMPORTANT**
>
> If the host is attached to the Manager, you must place the host into maintenance mode before making changes. After making the changes, reinstall the host to apply the changes.

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 *Installation Guide*. 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.

7.5.11. 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	Three options are available: • **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.

7.5.12. 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.

You must configure host power management in order to utilize host high availability and virtual machine high availability. For more information about power management devices, see Power Management in the *Technical Reference*.

Configuring Power Management Settings

1. Click **Compute → Hosts** and select a host.

2. Click **Management → Maintenance**, and click **OK** to confirm.

3. When the host is in maintenance mode, click **Edit**.

4. Click the **Power Management** tab.

5. Select the **Enable Power Management** check box to enable the fields.

6. Select the **Kdump integration** check box to prevent the host from fencing while performing a kernel crash dump.
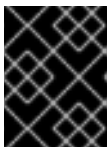
   > **IMPORTANT**
   >
   > If you enable or disable **Kdump integration** on an existing host, you must reinstall the host for kdump to be configured.

7. 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**.

8. Click the plus (**+**) button to add a new power management device. The **Edit fence agent** window opens.

9. Enter the **User Name** and **Password** of the power management device into the appropriate fields.

10. Select the power management device **Type** in the drop-down list.

11. Enter the IP address in the **Address** field.

12. Enter the **SSH Port** number used by the power management device to communicate with the host.

13. Enter the **Slot** number used to identify the blade of the power management device.

14. Enter the **Options** for the power management device. Use a comma-separated list of *'key=value'* entries.

    - If both IPv4 and IPv6 IP addresses can be used (default), leave the **Options** field blank.

    - If only IPv4 IP addresses can be used, enter **inet4_only=1**.

    - If only IPv6 IP addresses can be used, enter **inet6_only=1**.

15. Select the **Secure** check box to enable the power management device to connect securely to the host.

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

17. Click **OK** to close the **Edit fence agent** window.

18. 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.

19. Click **OK**.

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

7.5.13. 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.

Configuring SPM settings

1. Click **Compute → Hosts**.

2. Click **Edit**.

3. Click the **SPM** tab.

4. Use the radio buttons to select the appropriate SPM priority for the host.

5. Click **OK**.

7.5.14. 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.

> **NOTE**
>
> Virtual machines that are pinned to the host and cannot be migrated are shut down. You can check which virtual machines are pinned to the host by clicking **Pinned to Host** in the **Virtual Machines** tab of the host's details view.

Placing a Host into Maintenance Mode

1. Click **Compute → Hosts** and select the desired host.

2. Click **Management → Maintenance** to open the **Maintenance Host(s)** confirmation window.

3. Optionally, enter a **Reason** for moving the host into maintenance mode, which will appear in the logs and when the host is activated again.
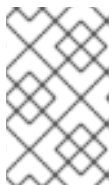
   > **NOTE**
   >
   > The host maintenance **Reason** field will only appear if it has been enabled in the cluster settings. See Section 5.2.2, "General Cluster Settings Explained" for more information.

4. Optionally, select the required options for hosts that support Gluster.
   Select the **Ignore Gluster Quorum and Self-Heal Validations** option to avoid the default checks. By default, the Manager checks that the Gluster quorum is not lost when the host is moved to maintenance mode. The Manager also checks that there is no self-heal activity that will be affected by moving the host to maintenance mode. If the Gluster quorum will be lost or if there is self-heal activity that will be affected, the Manager prevents the host from being placed into maintenance mode. Only use this option if there is no other way to place the host in maintenance mode.

   Select the **Stop Gluster Service** option to stop all Gluster services while moving the host to maintenance mode.
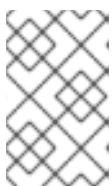
   > **NOTE**
   >
   > These fields will only appear in the host maintenance window when the selected host supports Gluster. See Replacing the Primary Gluster Storage Node in *Maintaining Red Hat Hyperconverged Infrastructure* for more information.

5. 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 **Management → 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.

7.5.15. 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.

Activating a Host from Maintenance Mode

1. Click **Compute → Hosts** and select the host.

2. Click **Management → 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.

7.5.16. Configuring Host Firewall Rules

You can configure the host firewall rules so that they are persistent, using Ansible. The cluster must be configured to use **firewalld**, not **iptables**.

Configuring Firewall Rules for Hosts

1. On the Manager machine, edit **ovirt-host-deploy-post-tasks.yml.example** to add a custom firewall port:

   ```
   # vi /etc/ovirt-engine/ansible/ovirt-host-deploy-post-tasks.yml.example
   ---
   #
   # Any additional tasks required to be executing during host deploy process can
   # be added below
   #
   - name: Enable additional port on firewalld
     firewalld:
       port: "12345/tcp"
       permanent: yes
       immediate: yes
       state: enabled
   ```

2. Save the file to another location as **ovirt-host-deploy-post-tasks.yml**.

New or reinstalled hosts are configured with the updated firewall rules.

Existing hosts must be reinstalled by clicking **Installation → Reinstall** and selecting **Automatically configure host firewall**.

7.5.17. Removing a Host

Remove a host from your virtualized environment.

Removing a host

1. Click **Compute → Hosts** and select the host.

2. Click **Management → Maintenance**.

3. when the host is in maintenance mode, 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**.

7.5.18. Updating a Host Between Minor Releases

7.5.18.1. Updating the Hosts

Use the host upgrade manager to update individual hosts directly from the Red Hat Virtualization Manager.

> **NOTE**
>
> The upgrade manager only checks hosts with a status of **Up** or **Non-operational**, but not **Maintenance**.
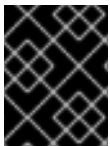
> **IMPORTANT**
>
> On RHVH, the update only preserves modified content in the **/etc** and **/var** directories. Modified data in other paths is overwritten during an update.

Prerequisites

- If migration is enabled at the cluster level, virtual machines are automatically migrated to another host in the cluster. Update a host when its usage is relatively low.

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

- Ensure that the cluster to which the host belongs has sufficient memory reserve for its hosts to perform maintenance. Otherwise, the virtual machine migration operation will hang and fail. You can reduce the memory usage of this operation by shutting down some or all virtual machines before updating the host.

- You cannot migrate a virtual machine using a vGPU to a different host. Virtual machines with vGPUs installed must be shut down before updating the host.

Procedure

1. Ensure that the correct repositories are enabled (to view a list of currently enabled repositories, type **yum repolist**):

   - For Red Hat Virtualization Hosts:

     ```
     # subscription-manager repos --enable=rhel-7-server-rhvh-4-rpms
     ```

   - For Red Hat Enterprise Linux hosts:

     ```
     # subscription-manager repos \
         --enable=rhel-7-server-rpms \
         --enable=rhel-7-server-rhv-4-mgmt-agent-rpms \
         --enable=rhel-7-server-ansible-2.9-rpms
     ```

2. In the Administration Portal, click **Compute → Hosts** and select the host to be updated.

3. Click **Installation → Check for Upgrade** and click **OK**.

   Click the **Events and alerts notification** icon (🔔) and expand the **Events** section to see the result.

4. If an update is available, click **Installation → Upgrade**.

5. Click **OK** to update the host. Running virtual machines will be migrated according to their migration policy. If migration is disabled for any virtual machines, you will be prompted to shut them down.
   The details of the host are updated in **Compute → Hosts** and the status transitions through these stages:

   - Maintenance

   - Installing

   - Reboot

   - Up
     If any virtual machines were migrated off the host, they are now migrated back.

   > **NOTE**
   >
   > If the update fails, the host's status changes to **Install Failed**. From **Install Failed** you can click **Installation → Upgrade** again.

Repeat this procedure for each host in the Red Hat Virtualization environment.

7.5.18.2. Manually Updating Hosts

You can use the **yum** command to update your hosts. Update your systems regularly, to ensure timely application of security and bug fixes.
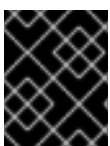
> **IMPORTANT**
>
> On RHVH, the update only preserves modified content in the **/etc** and **/var** directories. Modified data in other paths is overwritten during an update.

Prerequisites

- 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 updates are performed at a time when the host's usage is relatively low.

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

- Ensure that the cluster to which the host belongs has sufficient memory reserve in order for its hosts to perform maintenance. If a cluster lacks sufficient memory, the virtual machine migration operation will hang and then fail. You can reduce the memory usage of this operation by shutting down some or all virtual machines before updating the host.

- You cannot migrate a virtual machine using a vGPU to a different host. Virtual machines with vGPUs installed must be shut down before updating the host.

Procedure

1. Ensure the correct repositories are enabled. You can check which repositories are currently enabled by running **yum repolist**.

   - For Red Hat Virtualization Hosts:

     ```
     # subscription-manager repos --enable=rhel-7-server-rhvh-4-rpms
     ```

   - For Red Hat Enterprise Linux hosts:

     ```
     # subscription-manager repos \
         --enable=rhel-7-server-rpms \
         --enable=rhel-7-server-rhv-4-mgmt-agent-rpms \
         --enable=rhel-7-server-ansible-2.9-rpms
     ```

2. In the Administration Portal, click **Compute → Hosts** and select the host to be updated.

3. Click **Management → Maintenance**.

4. Update the host:

   ```
   # yum update
   ```

5. Reboot the host to ensure all updates are correctly applied.

   > **NOTE**
   >
   > Check the imgbased logs to see if any additional package updates have failed for a Red Hat Virtualization Host. If some packages were not successfully reinstalled after the update, check that the packages are listed in **/var/imgbased/persisted-rpms**. Add any missing packages then run **rpm -Uvh /var/imgbased/persisted-rpms/***.

Repeat this process for each host in the Red Hat Virtualization environment.

7.5.19. 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.

Prerequisites

- 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.

- Ensure that the cluster has sufficient memory reserve in order for its hosts to perform maintenance. If a cluster lacks sufficient memory, the virtual machine migration operation will 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.

- 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

1. Click **Compute → Hosts** and select the host.

2. Click **Management → Maintenance**.

3. Click **Installation → 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 can now 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 **Management → Activate**, and the host will change to an **Up** status and be ready for use.

7.5.20. Customizing Hosts with Tags

You can use tags to store information about your hosts. You can then search for hosts based on tags. For more information on searches, see Searching for Hosts in the *Introduction to the Administration Portal*.

Customizing hosts with tags

1. Click **Compute → Hosts** and select a host.

2. Click **More Actions → Assign Tags**.

3. Select the check boxes of applicable tags.

4. Click **OK**.

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

7.5.21. 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"

Viewing Host Errata

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **Errata** tab.

7.5.22. 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, click the host's name to open the details view, and click the **Events** 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*.

7.5.23. Viewing Host Devices

You can view the host devices for each host in the **Host Devices** tab in the details view. 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 *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*.

Viewing Host Devices

1. Click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click **Host Devices** tab.

This tab 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.

7.5.24. 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 *Planning and Prerequisites Guide* .

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

Preparing a Host for GPU Passthrough

1. In the Administration Portal, click **Compute → Hosts**.

2. Click the host's name to open the details view.

3. Click the **General** tab, and click **Hardware**. Locate the GPU device *vendor ID:product ID* . In this example, the IDs are **10de:13ba** and **10de:0fbc**.

4. Right-click the host and select **Edit**. Click the **Kernel** tab.

5. In the **Kernel command line** free text entry field, enter the IDs located in the previous steps.

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

6. 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
   ```

7. Click **OK**.

8. Click **Installation → Reinstall** to commit the changes to the host.

9. Reboot the host after the reinstallation is complete.
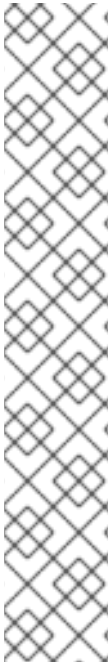
> **NOTE**
>
> To confirm the device is bound to the **pci-stub** driver, run the **lspci** command:
>
> ```
> # lspci -nnk
> ...
> 01:00.0 VGA compatible controller [0300]: NVIDIA Corporation GM107GL [Quadro
> K2200] [10de:13ba] (rev a2)
>         Subsystem: NVIDIA Corporation Device [10de:1097]
>         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.

Preparing a Guest Virtual Machine for GPU Passthrough

For Linux

1. 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"
   ...
   ```

2. 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)
   ```

3. 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
   ```

4. Restart the virtual machine.

For Windows

1. Download and install the corresponding drivers for the device. For example, for Nvidia drivers, go to NVIDIA Driver Downloads.

2. 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*.

7.5.25. Accessing Cockpit from the Administration Portal

Cockpit is available by default on Red Hat Virtualization Hosts (RHVH) and Red Hat Enterprise Linux hosts. You can access the Cockpit user interface by typing the address into a browser, or through the Administration Portal.

Accessing Cockpit from the Administration Portal

1. In the Administration Portal, click **Compute → Hosts** and select a host.

2. Click **Host Console**.

The Cockpit login page opens in a new browser window.

7.5.26. Setting a Legacy SPICE Cipher

SPICE consoles use FIPS–compliant encryption by default, with a cipher string. The default SPICE cipher string is: **kECDHE+FIPS:kDHE+FIPS:kRSA+FIPS:!eNULL:!aNULL**

This string is generally sufficient. However, if you have a virtual machine with an older operating system or SPICE client, where either one or the other does not support FIPS–compliant encryption, you must use a weaker cipher string. Otherwise, a connection security error may occur if you install a new cluster or a new host in an existing cluster and try to connect to that virtual machine.

You can change the cipher string by using an Ansible playbook.

Changing the cipher string

1. On the Manager machine, create a file in the directory **/usr/share/ovirt-engine/playbooks**. For example:

   ```
   # vim /usr/share/ovirt-engine/playbooks/change-spice-cipher.yml
   ```

2. Enter the following in the file and save it:

   ```
   name: oVirt - setup weaker SPICE encryption for old clients
   hosts: hostname
   vars:
     host_deploy_spice_cipher_string: 'DEFAULT:-RC4:-3DES:-DES'
   roles:
     - ovirt-host-deploy-spice-encryption
   ```

3. Run the file you just created:

   ```
   # ansible-playbook -l hostname /usr/share/ovirt-engine/playbooks/change-spice-cipher.yml
   ```

Alternatively, you can reconfigure the host with the Ansible playbook **ovirt-host-deploy** using the **--extra-vars** option with the variable **host_deploy_spice_cipher_string**, as follows:

```
# ansible-playbook -l hostname \
  --extra-vars host_deploy_spice_cipher_string="DEFAULT:-RC4:-3DES:-DES" \
  /usr/share/ovirt-engine/playbooks/ovirt-host-deploy.yml
```

7.6. HOST RESILIENCE

7.6.1. Host High Availability

The Red Hat Virtualization Manager uses fencing to keep 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.

Fencing allows a cluster to react to unexpected host failures and 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. In a fencing operation, a non-responsive host is rebooted, and if the host does not return to an active status within a prescribed time, it remains non-responsive pending manual intervention and troubleshooting.

NOTE

To automatically check the fencing parameters, you can configure the **PMHealthCheckEnabled** (false by default) and **PMHealthCheckIntervalInSec** (3600 sec by default) engine-config options.

When set to true, **PMHealthCheckEnabled** will check all host agents at the interval specified by **PMHealthCheckIntervalInSec**, and raise warnings if it detects issues. See Section 18.2.2, "Syntax for the engine-config Command" for more information about configuring engine-config options.

Power management operations can be performed by Red Hat Virtualization Manager after it reboots, by a proxy host, or manually in the Administration Portal. All the virtual machines running on the non-responsive host are stopped, and highly available virtual machines are started on a different host. At least two hosts are required for power management operations.

After the Manager starts up, it automatically attempts to fence non-responsive hosts that have power management enabled after the quiet time (5 minutes by default) has elapsed. The quiet time can be configured by updating the **DisableFenceAtStartupInSec** engine-config option.

NOTE

The **DisableFenceAtStartupInSec** engine-config option helps prevent a scenario where the Manager attempts to fence hosts while they boot up. This can occur after a data center outage because a host's boot process is normally longer than the Manager boot process.

Hosts can be fenced automatically by the proxy host using the power management parameters, or manually by right-clicking on a host and using the options on the menu.
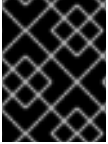
> **IMPORTANT**
>
> If a host runs virtual machines that are highly available, power management must be enabled and configured.

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.

Setting fencing parameters on a host

1. Click **Compute → Hosts** and select the host.

2. Click **Edit**.

3. Click the **Power Management** tab.

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.
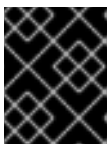
   > **IMPORTANT**
   >
   > If you enable or disable **Kdump integration** on an existing host, you must reinstall the host.

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 **+** 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.

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.

14. 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.

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**.

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

Click the name of a host to view the status of the kdump service in the **General** tab of the details view:

- **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 earlier 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]. For configuration of kdump on the Manager see xref: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.

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
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	–
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
```

Manually Configuring Kdump with engine-config

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

   ```
   # engine-config -s OPTION=value
   ```
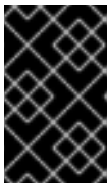
   > **IMPORTANT**
   >
   > The edited values must also be changed in the fence_kdump listener configuration file as outlined in the **Kdump Configuration Options** table. See Section 7.6.4.1, "fence_kdump listener Configuration".

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
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.

Variable	Description	Default	Note
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.
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, **vdsm restart** is executed via SSH.

4. If **vdsm 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

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.

Using Host Power Management Functions

1. Click **Compute → Hosts** and select the host.

2. Click the **Management** drop-down menu and select one of the following **Power Management** 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**.

     > **NOTE**
     >
     > If Power Management is not enabled, you can restart or stop the host by selecting it, clicking the **Management** drop-down menu, and selecting an **SSH Management** option, **Restart** or **Stop**.

> **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.

3. Click **OK**.

7.6.7. Manually Fencing or Isolating a Non Responsive Host

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 if 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.

Manually fencing or isolating a non-responsive host

1. In the Administration Portal, click **Compute → Hosts** and confirm the host's status is **Non Responsive**.

2. Manually reboot the host. This could mean physically entering the lab and rebooting the host.

3. In the Administration Portal, select the host and click **More Actions → Confirm 'Host has been Rebooted'**.

4. Select the **Approve Operation** check box and click **OK**.

5. If your hosts take an unusually long time to boot, you can set **ServerRebootTimeout** to specify how many seconds to wait before determining that the host is **Non Responsive**:

   ```
   # engine-config --set ServerRebootTimeout=integer
   ```

CHAPTER 8. STORAGE

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

- Network File System (NFS)

- GlusterFS exports

- 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.

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 disks, and templates can then be uploaded from the imported storage domain to the attached data center. See Section 8.7, "Importing Existing Storage Domains" for information on importing storage domains.
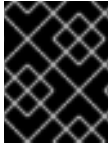
IMPORTANT

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 block devices (SAN – iSCSI or FCP) or a file system (NAS – NFS, GlusterFS, 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 sparse or preallocated. Snapshots are always sparse but can be taken for disks of either format.

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 storage domains on a Red Hat Enterprise Linux server.

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 disks, and templates can then be uploaded from the imported storage domain to the attached data center. See Section 8.7, "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.

Configuring the Required System User Accounts and System User Groups

1. Create the group **kvm**:

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

2. Create the user **vdsm** in the group **kvm**:

   ```
   # useradd vdsm -u 36 -g 36
   ```

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

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

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
   # chmod 0755 /exports/iso
   ```

For more information on the required system users and groups see Appendix F, *System Accounts*.

8.2.2. Attaching NFS Storage

Attach NFS storage domains to the data center in your Red Hat Virtualization environment. These storage domains provide storage for virtual disks (data domain) and ISO boot media (ISO domain). This procedure assumes that you have already exported shares. You must create the data domain before creating the ISO and export domains. Use the same procedure to create the ISO and export domains, selecting **ISO** or **Export** from the **Domain Function** list.

1. In the Administration Portal, click **Storage → Domains**.

2. Click **New Domain**.

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 *123.123.0.10:/data* (for IPv4), *[2001:0:0:0:0:0:0:5db1]:/data* (for IPv6), 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 has 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.

Increasing an Existing NFS Storage Domain

1. Click **Storage → Domains**.

2. Click the NFS storage domain's name to open the details view.

3. Click the **Data Center** tab and click **Maintenance** to place the storage domain into maintenance mode. This unmounts the existing share and makes it possible to resize the storage domain.

4. 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 .

5. In the details view, click the **Data Center** tab and click **Activate** 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 F, *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* for Red Hat Virtualization 4.0 for more details.

> **IMPORTANT**
>
> On Red Hat Virtualization Host (RHVH), local storage should always be defined on a file system that is separate from / (root). Red Hat recommends using a separate logical volume or disk, to prevent possible loss of data during upgrades.

Preparing Local Storage for Red Hat Enterprise Linux Hosts

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.

Preparing Local Storage for Red Hat Virtualization Hosts

Red Hat recommends creating the local storage on a logical volume as follows:

```
# mkdir /data
# lvcreate -L $SIZE rhvh -n data
# mkfs.ext4 /dev/mapper/rhvh-data
# echo "/dev/mapper/rhvh-data /data ext4 defaults,discard 1 2" >> /etc/fstab
```

Your local storage is ready to be added to the Red Hat Virtualization environment.

8.3.2. Adding Local Storage

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.

Adding Local Storage

1. Click **Compute → Hosts** and select the host.

2. Click **Management → Maintenance** and click **OK**.

3. Click **Management → Configure Local Storage**.

4. Click the **Edit** buttons next to the **Data Center**, **Cluster**, and **Storage** fields to configure and name the local storage domain.

5. Set the path to your local storage in the text entry field.

6. If applicable, click the **Optimization** tab to configure the memory optimization policy for the new local storage cluster.

7. Click **OK**.

Your host comes online in a data center of its own.

8.4. ADDING POSIX COMPLIANT FILE SYSTEM STORAGE

8.4.1. Attaching 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 file system used as a storage domain in Red Hat Virtualization must be a clustered file system, such as Global File System 2 (GFS2), and 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.

Attaching POSIX Compliant File System Storage

1. Click **Storage → Domains**.

2. Click **New Domain**.

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** from the **Domain Function** drop-down list, and **POSIX compliant FS** from the **Storage Type** drop-down list.
   If applicable, select the **Format** from the drop-down menu.

6. Select a host from the **Use Host** drop-down list.

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 in 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 in 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**.

8.5. ADDING BLOCK STORAGE

IMPORTANT

If you are using block storage and you intend to deploy virtual machines on raw devices or direct LUNs and to manage them with the Logical Volume Manager, you must create a filter to hide the guest logical volumes. This will prevent guest logical volumes from being activated when the host is booted, a situation that could lead to stale logical volumes and cause data corruption. See RHV: Hosts boot with Guest LVs activated for details.

IMPORTANT

Red Hat Virtualization currently does not support storage with a block size of 4K. You must configure block storage in legacy (512b block) mode.

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* .

Adding iSCSI Storage

1. Click **Storage → Domains**.

2. Click **New Domain**.

3. Enter the **Name** of the new storage domain.

4. Select a **Data Center** from the drop-down list.

5. Select the **Domain Function** and the **Storage Type** from the drop-down lists. 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 can 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.

   a. 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.

   b. Enter the fully qualified domain name or IP address of the iSCSI host in the **Address** field.

   c. Enter the port to connect to the host on when browsing for targets in the **Port** field. The default is **3260**.

   d. If the Challenge Handshake Authentication Protocol (CHAP) is being used to secure the storage, select the **User Authentication** check box. Enter the **CHAP user name** and **CHAP password**.

   > **NOTE**
   >
   > It is now possible to use the REST API to define specific credentials to each iSCSI target per host. See Adding a Storage Server Connection Extension in the *REST API Guide* for more information.

   e. Click **Discover**.

   f. 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 expands the entry and displays 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.

    e. Select the **Discard After Delete** check box to enable the discard after delete option. This option can be edited after the domain is created. This option is only available to block storage domains.

11. Click **OK**.

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.

If you want to migrate your current storage network to an iSCSI bond, see Section 8.5.3, "Migrating a Logical Network to an iSCSI Bond".

8.5.2. Configuring iSCSI Multipathing

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.

Configuring iSCSI Multipathing

1. Click **Compute → Data Centers**.

2. Click the data center's name to open the details view.

3. Click the **iSCSI Multipathing** tab.

4. Click **Add**.

5. In the **Add iSCSI Bond** window, enter a **Name** and a **Description** for the bond.

6. Select the networks to be used for the bond from the **Logical Networks** list. The networks must be **Non-Required** networks.
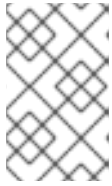
> **NOTE**
>
> To change a network's **Required** status to "Non-Required", in the Administration Portal, select a network, click the **Clusters** tab, click the **Manage Networks** button, and clear the **Required** check box.

7. Select the storage domain to be accessed via the chosen networks from the **Storage Targets** list. Ensure you select all paths to the same target.

8. Click **OK**.

All hosts in the data center are connected to the selected iSCSI target through the selected logical networks.

8.5.3. Migrating a Logical Network to an iSCSI Bond

If you have a logical network that you created for iSCSI traffic and configured on top of an existing network bond, you can migrate it to an iSCSI bond on the same subnet without disruption or downtime.

Procedure

1. Modify the current logical network:

   a. Click **Compute → Clusters**.

   b. Click the cluster name to open the details view.

   c. In the **Logical Networks** tab, select the current logical network (**net-1**) and click **Manage Networks**.

   d. Clear the **Require** check box and click **OK**.

2. Create a new logical network:

   a. Click **Add Network** to open the **New Logical Network** window.

   b. In the **General** tab, enter the **Name** (**net-2**) and clear the **VM network** check box.

   c. In the **Cluster** tab, clear the **Require** check box and click **OK**.

3. Remove the current network bond and reassign the logical networks:

   a. Click **Compute → Hosts**.

   b. Click the host name to open the details view.

   c. In the **Network Interfaces** tab, click **Setup Host Networks**.

   d. Drag **net-1** to the right to unassign it.

   e. Drag the current bond to the right to remove it.

f. Drag **net-1** and **net-2** to the left to assign them to physical interfaces.

g. Click the pencil icon of **net-2** to open the **Edit Network** window.

h. In the **IPV4** tab, select **Static**, enter the **IP** and **Netmask/Routing Prefix** of the subnet, and click **OK**.

4. Create the iSCSI bond:

a. Click **Compute → Data Centers**.

b. Click the data center name to open the details view.

c. In the **iSCSI Multipathing** tab, click **Add**.

d. In the **Add iSCSI Bond** window, enter a **Name**, select the networks, **net-1** and **net-2**, and click **OK**.

Your data center has an iSCSI bond containing the old and new logical networks.

8.5.4. 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*.

Adding FCP Storage

1. Click **Storage → Domains**.

2. Click **New Domain**.

3. Enter the **Name** of the storage domain.

4. Select an FCP **Data Center** from the drop-down list.
   If you do not yet have an appropriate FCP data center, select **(none)**.

5. Select the **Domain Function** and the **Storage Type** from the drop-down lists. 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 **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.

   e. Select the **Discard After Delete** check box to enable the discard after delete option. This option can be edited after the domain is created. This option is only available to block storage domains.

9. Click **OK**.

The new FCP data domain remains in a **Locked** status while it is being prepared for use. When ready, it is automatically attached to the data center.

8.5.5. Increasing iSCSI or FCP Storage

There are several ways to increase iSCSI or FCP storage size:

- Add an existing LUN to the current storage domain.

- Create a new storage domain with new LUNs and add it to an existing data center. See Section 8.5.1, "Adding iSCSI Storage".

- Expand the storage domain by resizing the underlying LUNs.

For information about creating, configuring, or resizing iSCSI storage on Red Hat Enterprise Linux 7 systems, see the *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.

Prerequisites

- The storage domain's status must be **UP**.

- The LUN must be accessible to all the hosts whose status is **UP**, or else the operation will fail and the LUN will not be added to the domain. The hosts themselves, however, will not be affected. If a newly added host, or a host that is coming out of maintenance or a **Non Operational** state, cannot access the LUN, the host's state will be **Non Operational**.

Increasing an Existing iSCSI or FCP Storage Domain

1. Click **Storage → Domains** and select an iSCSI or FCP domain.

2. Click **Manage Domain**.

3. Click **Targets > LUNs** and click the **Discover Targets** expansion button.

4. Enter the connection information for the storage server and click **Discover** to initiate the connection.

5. Click **LUNs > Targets** and select the check box of the newly available LUN.

6. 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 Administration Portal.

Refreshing the LUN Size

1. Click **Storage → Domains** and select an iSCSI or FCP domain.

2. Click **Manage Domain**.

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.

8.5.6. Reusing LUNs

LUNs cannot be reused, as is, to create a storage domain or virtual disk. If you try to reuse the LUNs, the Administration Portal displays the following error message:

> Physical device initialization failed. Please check that the device is empty and accessible by the host.

A self-hosted engine shows the following error during installation:

> [ERROR] Error creating Volume Group: Failed to initialize physical device: ("
> [u'/dev/mapper/00000000000000000000000000000000']",)
> [ERROR] Failed to execute stage 'Misc configuration': Failed to initialize physical device: ("
> [u'/dev/mapper/00000000000000000000000000000000']",)

Before the LUN can be reused, the old partitioning table must be cleared.

Clearing the Partition Table from a LUN
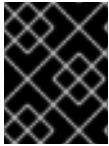
> **IMPORTANT**
>
> You must run this procedure on the correct LUN so that you do not inadvertently destroy data.

Run the **dd** command with the ID of the LUN that you want to reuse, the maximum number of bytes to read and write at a time, and the number of input blocks to copy:

```
# dd if=/dev/zero of=/dev/mapper/LUN_ID bs=1M count=200 oflag=direct
```

8.6. ADDING RED HAT GLUSTER STORAGE

To use Red Hat Gluster Storage with Red Hat Virtualization, see *Configuring Red Hat Virtualization with Red Hat Gluster Storage*.

For the Red Hat Gluster Storage versions that are supported with Red Hat Virtualization, see https://access.redhat.com/articles/2356261.
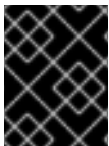
> **IMPORTANT**
>
> Red Hat Hyperconverged Infrastructure is not currently supported with Red Hat Virtualization 4.2.

8.7. IMPORTING EXISTING STORAGE DOMAINS

8.7.1. 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*.
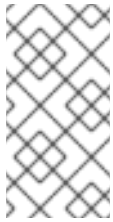
> **NOTE**
>
> The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disks, and templates can then be uploaded from the imported storage domain to the attached data center.

8.7.2. 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.

Importing a Storage Domain

1. Click **Storage → Domains**.

2. Click **Import Domain**.

3. Select the **Data Center** you want to import the storage domain to.

4. Enter a **Name** for the storage domain.

5. Select the **Domain Function** and **Storage Type** from the 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.
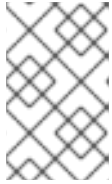
NOTE

The fields for specifying the details of the storage domain change depending on the values you select in the **Domain Function** and **Storage Type** lists. These fields are the same as those available for adding a new storage domain.

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**.

You can now import virtual machines and templates from the storage domain to the data center.

8.7.3. 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.

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 **Storage → Domains**.

3. Click the storage domain's name to open the details view.

4. Click the **Data Center** tab.

5. Click **Maintenance**, then click **OK**.

6. Click **Detach**, then click **OK**.

7. Click **Attach**.

8. 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.

8.7.4. 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.

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 **Storage → Domains**.

4. Click the storage domain's name to open the details view.

5. Click the **Data Center** tab.

6. Click **Maintenance**, then click **OK**.

7. Click **Detach**, then click **OK**.

8. Click **Remove**.

9. 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.

10. Click **OK** to remove the storage domain from the source environment.

11. Log in to the Administration Portal of the destination environment.

12. Click **Storage → Domains**.

13. Click **Import Domain**.

14. Select the destination data center from the **Data Center** drop-down list.

15. Enter a name for the storage domain.

16. Select the **Domain Function** and **Storage Type** from the appropriate drop-down lists.

17. Select a host from the **Use Host** drop-down list.

18. Enter the details of the storage domain.
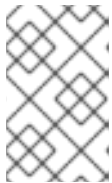
> **NOTE**
>
> The fields for specifying the details of the storage domain change depending on the value you select in the **Storage Type** drop-down list. These fields are the same as those available for adding a new storage domain.

19. Select the **Activate Domain in Data Center** check box to automatically activate the storage domain when it is attached.

20. 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.

8.7.5. Importing Virtual Machines from Imported Data Storage Domains

Import a virtual machine into one or more clusters 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.

Importing Virtual Machines from an Imported Data Storage Domain

1. Click **Storage → Domains**.

2. Click the imported storage domain's name to open the details view.

3. Click the **VM Import** tab.

4. Select one or more virtual machines to import.

5. Click **Import**.

6. For each virtual machine in the **Import Virtual Machine(s)** window, ensure the correct target cluster is selected in the **Cluster** list.

7. Map external virtual machine vNIC profiles to profiles that are present on the target cluster(s):

   a. Click **vNic Profiles Mapping**.

   b. Select the vNIC profile to use from the **Target vNic Profile** drop-down list.

   c. If multiple target clusters are selected in the **Import Virtual Machine(s)** window, select each target cluster in the **Target Cluster** drop-down list and ensure the mappings are correct.

   d. Click **OK**.

8. If a MAC address conflict is detected, an exclamation mark appears next to the name of the virtual machine. Mouse over the icon to view a tooltip displaying the type of error that occurred. Select the **Reassign Bad MACs** check box to reassign new MAC addresses to all problematic virtual machines. Alternatively, you can select the **Reassign** check box per virtual machine.

> **NOTE**
>
> If there are no available addresses to assign, the import operation will fail. However, in the case of MAC addresses that are outside the cluster's MAC address pool range, it is possible to import the virtual machine without reassigning a new MAC address.

9. Click **OK**.

The imported virtual machines no longer appear in the list under the **VM Import** tab.

8.7.6. 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.

Importing Templates from an Imported Data Storage Domain

1. Click **Storage → Domains**.

2. Click the imported storage domain's name to open the details view.

3. Click the **Template Import** tab.

4. Select one or more templates to import.

5. Click **Import**.

6. For each template in the **Import Templates(s)** window, ensure the correct target cluster is selected in the **Cluster** list.

7. Map external virtual machine vNIC profiles to profiles that are present on the target cluster(s):

   a. Click **vNic Profiles Mapping**.

   b. Select the vNIC profile to use from the **Target vNic Profile** drop-down list.

   c. If multiple target clusters are selected in the **Import Templates** window, select each target cluster in the **Target Cluster** drop-down list and ensure the mappings are correct.

   d. Click **OK**.

8. Click **OK**.

The imported templates no longer appear in the list under the **Template Import** tab.

8.8. STORAGE TASKS

8.8.1. Uploading Images to a Data Storage Domain

You can upload virtual disk images and ISO images to your data storage domain in the Administration Portal or with the REST API.
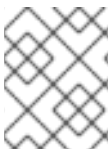
> **NOTE**
>
> To upload images with the REST API, see IMAGETRANSFERS and IMAGETRANSFER in the *REST API Guide*.

QEMU-compatible virtual disks can be attached to virtual machines. Virtual disk types must be either QCOW2 or raw. Disks created from a QCOW2 virtual disk cannot be shareable, and the QCOW2 virtual disk file must not have a backing file.

ISO images can be attached to virtual machines as CDROMs or used to boot virtual machines.

Prerequisites

The upload function uses HTML 5 APIs, which requires your environment to have the following:

- Image I/O Proxy (**ovirt-imageio-proxy**), configured with **engine-setup**. See Configuring the Red Hat Virtualization Manager in the *Installation Guide* for details.

- Certificate authority, imported into the web browser used to access the Administration Portal. To import the certificate authority, browse to **https://*engine_address*/ovirt-engine/services/pki-resource?resource=ca-certificate&format=X509-PEM-CA** and enable all the trust settings. Refer to the instructions to install the certificate authority in Firefox, Internet Explorer, or Google Chrome.

- Browser that supports HTML 5, such as Firefox 35, Internet Explorer 10, Chrome 13, or later.

Uploading an Image to a Data Storage Domain

1. Click **Storage → Disks**.

2. Select **Start** from the **Upload** menu.

3. Click **Choose File** and select the image to upload.

4. Fill in the **Disk Options** fields. See Section 10.6.2, "Explanation of Settings in the New Virtual Disk Window" for descriptions of the relevant fields.

5. Click **OK**.
   A progress bar indicates the status of the upload. You can pause, cancel, or resume uploads from the **Upload** menu.

Increasing the Upload Timeout Value

1. If the upload times out and you see the message, **Reason: timeout due to transfer inactivity**, increase the timeout value:

   ```
   # engine-config -s TransferImageClientInactivityTimeoutInSeconds=6000
   ```

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

   ```
   # systemctl restart ovirt-engine
   ```

8.8.2. Moving Storage Domains to Maintenance Mode

A storage domain must be in maintenance mode before it can be detached and removed. This is required to redesignate another data domain as the master data domain.
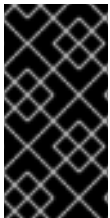
> **IMPORTANT**
>
> You cannot move a storage domain into maintenance mode if a virtual machine has a lease on the storage domain. The virtual machine needs to be shut down, or the lease needs to be to removed or moved to a different storage domain first. See the *Virtual Machine Management Guide* for information about virtual machine leases.

Expanding iSCSI domains by adding more LUNs can only be done when the domain is active.

Moving storage domains to maintenance mode

1. Shut down all the virtual machines running on the storage domain.

2. Click **Storage → Domains**.

3. Click the storage domain's name to open the details view.

4. Click the **Data Center** tab.

5. Click **Maintenance**.

   > **NOTE**
   >
   > The **Ignore OVF update failure** check box allows the storage domain to go into maintenance mode even if the OVF update fails.

6. Click **OK**.

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 in the details view of the data center it is associated with.

8.8.3. 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)**, **Wipe After Delete**, and **Discard 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 Storage Connections in the *REST API Guide*.

Editing an Active Storage Domain

1. Click **Storage → Domains** and select a storage domain.

2. Click **Manage Domain**.

3. Edit the available fields as required.

4. Click **OK**.

Editing an Inactive Storage Domain

1. Click **Storage → Domains**.

2. If the storage domain is active, move it to maintenance mode:

    a. Click the storage domain's name to open the details view.

    b. Click the **Data Center** tab.

    c. Click **Maintenance**.

    d. Click **OK**.

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. Activate the storage domain:

   a. Click the storage domain's name to open the details view.

   b. Click the **Data Center** tab.

   c. Click **Activate**.

8.8.4. Updating OVFs

By default, OVFs are updated every 60 minutes. However, if you have imported an important virtual machine or made a critical update, you can update OVFs manually.

Updating OVFs

1. Click **Storage → Domains**.

2. Select the storage domain and click **More Actions → Update OVFs**.
   The OVFs are updated and a message appears in **Events**.

8.8.5. 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 **Storage → Domains**.

2. Click an inactive storage domain's name to open the details view.

3. Click the **Data Centers** tab.

4. Click **Activate**.
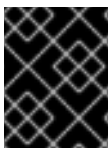
> **IMPORTANT**
>
> If you attempt to activate the ISO domain before activating the data domain, an error message displays and the domain is not activated.

8.8.6. Detaching a Storage Domain from a Data Center

Detach a storage domain from one data center to migrate it to another data center.

Detaching a Storage Domain from the Data Center

1. Click **Storage → Domains**.

2. Click the storage domain's name to open the details view.

3. Click the **Data Center** tab.

4. Click **Maintenance**.

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

6. Click **Detach**.

7. 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.

8.8.7. Attaching a Storage Domain to a Data Center

Attach a storage domain to a data center.

Attaching a Storage Domain to a Data Center

1. Click **Storage → Domains**.

2. Click the storage domain's name to open the details view.

3. Click the **Data Center** tab.

4. Click **Attach**.

5. Select the appropriate data center.

6. Click **OK**.

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

8.8.8. Removing a Storage Domain

You have a storage domain in your data center that you want to remove from the virtualized environment.

Procedure

1. Click **Storage → Domains**.

2. Move the storage domain to maintenance mode and detach it:

   a. Click the storage domain's name to open the details view.

   b. Click the **Data Center** tab.

   c. Click **Maintenance**, then click **OK**.

   d. Click **Detach**, then click **OK**.

3. Click **Remove**.

4. Optionally select the **Format Domain, i.e. Storage Content will be lost** check box to erase the content of the domain.

5. Click **OK**.

The storage domain is permanently removed from the environment.

8.8.9. Destroying a Storage Domain

A storage domain encountering errors may not be able to be removed through the normal procedure. Destroying a storage domain forcibly removes the storage domain from the virtualized environment.

Destroying a Storage Domain

1. Click **Storage → Domains**.

2. Select the storage domain and click **More Actions → Destroy**.

3. Select the **Approve operation** check box.

4. Click **OK**.

8.8.10. Creating a Disk Profile

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.

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.

Creating a Disk Profile

1. Click **Storage → Domains**.

2. Click the data storage domain's name to open the details view.

3. Click the **Disk Profiles** tab.

4. Click **New**.

5. Enter a **Name** and a **Description** for the disk profile.

6. Select the quality of service to apply to the disk profile from the **QoS** list.

7. Click **OK**.

8.8.11. Removing a Disk Profile

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

Removing a Disk Profile

1. Click **Storage → Domains**.

2. Click the data storage domain's name to open the details view.

3. Click the **Disk Profiles** tab.

4. Select the disk profile to remove.

5. Click **Remove**.

6. Click **OK**.

If the disk profile was assigned to any virtual disks, the disk profile is removed from those virtual disks.

8.8.12. 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, click the storage domain's name to open the details view, and click the **Events** 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*.

8.8.13. Setting Discard After Delete for a Storage Domain

When the **Discard After Delete** check box is selected, a **blkdiscard** command is called on a logical volume when it is removed and the underlying storage is notified that the blocks are free. The storage array can use the freed space and allocate it when requested. **Discard After Delete** only works on block storage. The flag is not available on the Red Hat Virtualization Manager for file storage, for example NFS.

Restrictions:

- **Discard After Delete** is only available on block storage domains, such as iSCSI or Fibre Channel.

- The underlying storage must support **Discard**.

Discard After Delete can be enabled both when creating a block storage domain or when editing a block storage domain. See] and xref:Editing_Storage_Domains[.

CHAPTER 9. POOLS

9.1. 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 and configuration changes are 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.

9.2. CREATING A VIRTUAL MACHINE POOL

You can create a virtual machine pool containing multiple virtual machines based on a common template. See Templates in the *Virtual Machine Management Guide* for information about sealing a virtual machine and creating a template.

Sysprep File Configuration Options for Windows Virtual Machines

Several **sysprep** file configuration options are available, depending on your requirements.

If your pool does not need to join a domain, you can use the default **sysprep** file, located in **/usr/share/ovirt-engine/conf/sysprep/**.

If your pool needs to join a domain, you can create a custom **sysprep** for each Windows operating system:

1. Copy the relevant sections for each operating system from **/usr/share/ovirt-engine/conf/osinfo-defaults.properties** to a new file and save as **99-defaults.properties**.

2. In **99-defaults.properties**, specify the Windows product activation key and the path of your new custom **sysprep** file:

```
os.operating_system.productKey.value=Windows_product_activation_key
...
os.operating_system.sysprepPath.value =
${ENGINE_USR}/conf/sysprep/sysprep.operating_system
```

3. Create a new **sysprep** file, specifying the domain, domain password, and domain administrator:

   ```
   <Credentials>
       <Domain>AD_Domain</Domain>
       <Password>Domain_Password</Password>
       <Username>Domain_Administrator</Username>
   </Credentials>
   ```

If you need to configure different **sysprep** settings for different pools of Windows virtual machines, you can create a custom **sysprep** file in the Administration Portal (see Creating a Virtual Machine Pool below). See Using Sysprep to Automate the Configuration of Virtual Machines in the *Virtual Machine Guide* for more information.

Creating a Virtual Machine Pool

1. Click **Compute → Pools**.

2. Click **New**.

3. Select a **Cluster** from the drop-down list.

4. Select a **Template** and version from the drop-down menu. A template provides standard settings for all the virtual machines in the pool.

5. Select an **Operating System** from the drop-down list.

6. Use the **Optimized for** drop-down list to optimize virtual machines for **Desktop** or **Server**.

   > **NOTE**
   >
   > **High Performance** optimization is not recommended for pools because a high performance virtual machine is pinned to a single host and concrete resources. A pool containing multiple virtual machines with such a configuration would not run well.

7. Enter a **Name** and, optionally, a **Description** and **Comment**.
   The **Name** of the pool is applied to each virtual machine in the pool, with a numeric suffix. You can customize the numbering of the virtual machines with **?** as a placeholder.

   > **Example 9.1. Pool Name and Virtual Machine Numbering Examples**
   >
   > - Pool: **MyPool**
   >   Virtual machines: **MyPool-1**, **MyPool-2**, ... **MyPool-10**
   >
   > - Pool: **MyPool-???**
   >   Virtual machines: **MyPool-001**, **MyPool-002**, ... **MyPool-010**

8. Enter the **Number of VMs** for the pool.

9. Enter the number of virtual machines to be prestarted in the **Prestarted** field.

10. Select the **Maximum number of VMs per user**that a single user is allowed to run in a session. The minimum is **1**.

11. Select the **Delete Protection** check box to enable delete protection.

12. If you are creating a pool of non-Windows virtual machines or if you are using the default **sysprep**, skip this step. If you are creating a custom **sysprep** file for a pool of Windows virtual machines:

    a. Click the **Show Advanced Options** button.

    b. Click the **Initial Run** tab and select the **Use Cloud-Init/Sysprep** check box.

    c. Click the **Authentication** arrow and enter the **User Name** and **Password** or select **Use already configured password**.

        > **NOTE**
        >
        > This **User Name** is the name of the local administrator. You can change its value from its default value (**user**) here in the **Authentication** section or in a custom **sysprep** file.

    d. Click the **Custom Script** arrow and paste the contents of the default **sysprep** file, located in **/usr/share/ovirt-engine/conf/sysprep/**, into the text box.

    e. You can modify the following values of the **sysprep** file:

        - **Key**. If you do not want to use the pre-defined Windows activation product key, replace **<![CDATA[$ProductKey$]]>** with a valid product key:

          ```
          <ProductKey>
              <Key><![CDATA[$ProductKey$]]></Key>
          </ProductKey>
          ```

          Example 9.2. Windows Product Key Example

          ```
          <ProductKey>
              <Key>0000-000-000-000</Key>
          </ProductKey>
          ```

        - **Domain** that the Windows virtual machines will join, the domain's **Password**, and the domain administrator's **Username**:

          ```
          <Credentials>
              <Domain>AD_Domain</Domain>
              <Password>Domain_Password</Password>
              <Username>Domain_Administrator</Username>
          </Credentials>
          ```

          Example 9.3. Domain Credentials Example

```
<Credentials>
    <Domain>addomain.local</Domain>
    <Password>12345678</Password>
    <Username>Sarah_Smith</Username>
</Credentials>
```

> **NOTE**
>
> The **Domain**, **Password**, and **Username** are required to join the domain. The **Key** is for activation. You do not necessarily need both.
>
> The domain and credentials cannot be modified in the **Initial Run** tab.

- **FullName** of the local administrator:

  ```
  <UserData>
  ...
      <FullName>Local_Administrator</FullName>
  ...
  </UserData>
  ```

- **DisplayName** and **Name** of the local administrator:

  ```
  <LocalAccounts>
      <LocalAccount wcm:action="add">
          <Password>
              <Value><![CDATA[$AdminPassword$]]></Value>
              <PlainText>true</PlainText>
          </Password>
          <DisplayName>Local_Administrator</DisplayName>
          <Group>administrators</Group>
          <Name>Local_Administrator</Name>
      </LocalAccount>
  </LocalAccounts>
  ```

  The remaining variables in the **sysprep** file can be filled in on the **Initial Run** tab.

13. Optional. Set a **Pool Type**:

    a. Click the **Type** tab and 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.

    b. Select the **Stateful Pool** check box to ensure that virtual machines are started in a stateful mode. This ensures that changes made by a previous user will persist on a virtual machine.

    c. Click **OK**.

14. Optional. Override the SPICE proxy:

    a. In the **Console** tab, select the **Override SPICE Proxy** check box.

b. In the **Overridden SPICE proxy address** text field, specify the address of a SPICE proxy to override the global SPICE proxy.

c. Click **OK**.

15. For a pool of Windows virtual machines, click **Compute → Virtual Machines**, select each virtual machine from the pool, and click **Run → Run Once**.

> **NOTE**
>
> If the virtual machine does not start and **Info [windeploy.exe] Found no unattend file** appears in **%WINDIR%\panther\UnattendGC\setupact.log**, add the **UnattendFile** key to the registry of the Windows virtual machine that was used to create the template for the pool:
>
> 1. Check that the Windows virtual machine has an attached floppy device with the unattend file, for example, **A:\Unattend.xml**.
>
> 2. Click **Start**, click **Run**, type **regedit** in the **Open** text box, and click **OK**.
>
> 3. In the left pane, go to **HKEY_LOCAL_MACHINE → SYSTEM → Setup**.
>
> 4. Right-click the right pane and select **New → String Value**.
>
> 5. Enter **UnattendFile** as the key name.
>
> 6. Double-click the new key and enter the **unattend** file name and path, for example, **A:\Unattend.xml**, as the key's value.
>
> 7. Save the registry, seal the Windows virtual machine, and create a new template. See Templates in the *Virtual Machine Management Guide* for details.

You have created and configured a virtual machine pool with the specified number of identical virtual machines. You can view these virtual machines in **Compute → Virtual Machines**, or by clicking the name of a pool to open its details view; a virtual machine in a pool is distinguished from independent virtual machines by its icon.

9.3. EXPLANATION OF SETTINGS AND CONTROLS IN THE NEW POOL AND EDIT POOL WINDOWS

9.3.1. 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 9.1. General settings

Field Name	Description

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 *Virtual Machine Management Guide*.
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.

9.3.2. New Pool 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 9.2. Type settings

Field Name	Description
Pool Type	This drop-down menu allows you to specify the type of the virtual machine pool. The following options are available: • **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	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.

9.3.3. 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 9.3. Console settings

Field Name	Description
Override SPICE proxy	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 VM Portal) is outside of the network where the hosts reside.
Overridden SPICE proxy address	The proxy by which the SPICE client connects 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: protocol://*host*:*port*

9.3.4. 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 9.4. Virtual Machine Pool: Host Settings

Field Name	Sub-element	Description
Start Running On		Defines the preferred host on which the virtual machine is to run. Select either: • **Any Host in Cluster** - The virtual machine can start and run on any available host in the cluster. • **Specific Host(s)** - 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.
Migration Options	Migration mode	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. • **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	Defines the migration convergence policy. If the check box is left unselected, the host determines the policy. • **Legacy** – Legacy behavior of 3.6 version. Overrides in **vdsm.conf** are still applied. The guest agent hook mechanism is disabled. • **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. • **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.
	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	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. • 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	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. • 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	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.
Configure NUMA	NUMA Node Count	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**.

Field Name	Sub-element	Description
	Tune Mode	The method used to allocate memory. ● **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	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.

9.3.5. New Pool and Edit Pool Resource Allocation Settings Explained

The following table details the information required on the **Resource Allocation** 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. See Virtual Machine Resource Allocation Settings Explained in the *Virtual Machine Management Guide* for more information.

Table 9.5. Resource Allocation settings

Field Name	Sub-element	Description
Disk Allocation	Auto select target	Select this check box to automatically select the storage domain that has the most free space. The **Target** and **Disk Profile** fields are disabled.

Field Name	Sub-element	Description
	Format	This field is read-only and always displays **QCOW2** unless the storage domain type is OpenStack Volume (Cinder), in which case the format is **Raw**.

9.4. 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.

Editing a Virtual Machine Pool

1. Click **Compute** → **Pools** and select a virtual machine pool.

2. Click **Edit**.

3. Edit the properties of the virtual machine pool.

4. Click **Ok**.

9.5. 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.

Prestarting Virtual Machines in a Pool

1. Click **Compute** → **Pools** and select the virtual machine pool.

2. Click **Edit**.

3. Enter the number of virtual machines to be prestarted in the **Prestarted VMs** field.

4. Click the **Type** tab. Ensure **Pool Type** is set to **Automatic**.

5. Click **OK**.

9.6. 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.

Adding Virtual Machines to a Virtual Machine Pool

1. Click **Compute → Pools** and select the virtual machine pool.

2. Click **Edit**.

3. Enter the number of additional virtual machines in the **Increase number of VMs in pool by** field.

4. Click **OK**.

9.7. 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.

Detaching Virtual Machines from a Virtual Machine Pool

1. Click **Compute → Pools**.

2. Click the pool's name to open the details view.

3. Click the **Virtual Machines** tab to list the virtual machines in the pool.

4. Ensure the virtual machine has a status of **Down**; you cannot detach a running virtual machine.

5. Select one or more virtual machines and click **Detach**.

6. Click **OK**.

> **NOTE**
>
> The virtual machine still exists in the environment and can be viewed and accessed from **Compute → Virtual Machines**. Note that the icon changes to denote that the detached virtual machine is an independent virtual machine.

9.8. 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.

Removing a Virtual Machine Pool

1. Click **Compute → Pools** and select the virtual machine pool.

2. Click **Remove**.

3. Click **OK**.

9.9. 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.

9.9.1. 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 9.6. 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.

Option	Default Value	Description
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.

9.9.2. 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.

Creating a Trusted Cluster

1. Click **Compute → Clusters**.

2. Click **New**.

3. Enter a **Name** for the cluster.

4. Select the **Enable Virt Service** check box.

5. Click the **Scheduling Policy** tab and select the **Enable Trusted Service** check box.

6. Click **OK**.

9.9.3. 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.

Adding a Trusted Host

1. Click **Compute → Hosts**.

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 **Hostname** 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 10. VIRTUAL DISKS

10.1. 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.

10.2. 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.

You can view a virtual disk's **ID** in **Storage → Disks**. The **ID** is used to identify a virtual disk because its device name (for example, **/dev/vda0**) can change, causing disk corruption. You can also view a virtual disk's ID in **/dev/disk/by-id**.

You can view the **Virtual Size** of a disk in **Storage → Disks** and in the **Disks** tab of the details view for storage domains, virtual machines, and templates. The **Virtual Size** 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 you create or edit a virtual disk.

You can view the **Actual Size** of a disk in the **Disks** tab of the details view for storage domains and templates. This is the amount of disk space that has been allocated to the virtual machine so far. Preallocated disks show the same value for **Virtual Size** and **Actual Size**. Sparse disks may show different values, depending on how much disk space has been allocated.
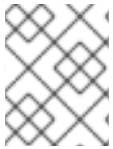
> **NOTE**
>
> When creating a Cinder virtual disk, the format and type of the disk are handled internally by Cinder and are not managed by Red Hat Virtualization.

The possible combinations of storage types and formats are described in the following table.

Table 10.1. Permitted Storage Combinations

Storage	Format	Type	Note
NFS	Raw	Preallocated	A file with an initial size that equals the amount of storage defined for the virtual disk, and has no formatting.
NFS	Raw	Sparse	A file with an initial size that is close to zero, and has no formatting.
NFS	QCOW2	Sparse	A file with an initial size that is close to zero, and has QCOW2 formatting. Subsequent layers will be QCOW2 formatted.
SAN	Raw	Preallocated	A block device with an initial size that equals the amount of storage defined for the virtual disk, and has no formatting.

Storage	Format	Type	Note
SAN	QCOW2	Sparse	A block device with an initial size that 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).

10.3. 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 reuse 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-config** tool on the Red Hat Virtualization Manager. Restart the **ovirt-engine** service for the setting change to take effect.
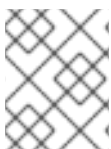
> **NOTE**
>
> Changing the **wipe_after_delete** flag's default setting will not affect the **Wipe After Delete** property of disks that already exist.

Setting SANWipeAfterDelete to Default to True Using the Engine Configuration Tool

1. Run the **engine-config** tool with the **--set** action:

   ```
   # engine-config --set SANWipeAfterDelete=true
   ```

2. Restart the **ovirt-engine** service 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': ImgsPar(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***.

10.4. 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.

10.5. 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 view 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**).

10.6. VIRTUAL DISK TASKS

10.6.1. Creating a Virtual Disk

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 11.2.4, "Adding an OpenStack Block Storage (Cinder) Instance for Storage Management" for more information.

You can create a virtual disk that is attached to a specific virtual machine. Additional options are available when creating an attached virtual disk, as specified in Section 10.6.2, "Explanation of Settings in the New Virtual Disk Window".

Creating a Virtual Disk Attached to a Virtual Machine

1. Click **Compute → Virtual Machines**.

2. Click the virtual machine's name to open the details view.

3. Click the **Disks** tab.

4. Click **New**.

5. Click the appropriate button to specify whether the virtual disk will be an **Image**, **Direct LUN**, or **Cinder** disk.

6. Select the options required for your virtual disk. The options change based on the disk type selected. See Section 10.6.2, "Explanation of Settings in the New Virtual Disk Window" for more details on each option for each disk type.

7. Click **OK**.

You can also create a floating virtual disk that does not belong to any virtual machines. You can attach this disk to a single virtual machine, or to multiple virtual machines if the disk is shareable. Some options are not available when creating a virtual disk, as specified in Section 10.6.2, "Explanation of Settings in the New Virtual Disk Window".

Creating a Floating Virtual Disk
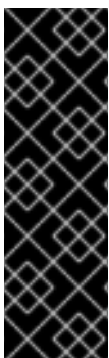
> **IMPORTANT**
>
> Creating floating virtual disks is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend to use them for production. These features 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/.

1. Click **Storage → Disks**.

2. Click **New**.

3. Click the appropriate button to specify whether the virtual disk will be an **Image**, **Direct LUN**, or **Cinder** disk.

4. Select the options required for your virtual disk. The options change based on the disk type selected. See Section 10.6.2, "Explanation of Settings in the New Virtual Disk Window" for more details on each option for each disk type.

5. Click **OK**.

10.6.2. Explanation of Settings in the New Virtual Disk Window

Because the New Virtual Disk windows for creating floating and attached virtual disks are very similar, their settings are described in a single section.

Table 10.2. New Virtual Disk and Edit Virtual Disk SettingsImage

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.
Interface	This field only appears when creating an attached disk. The virtual interface the disk presents to virtual machines. **VirtIO** is faster, but requires drivers. Red Hat Enterprise Linux 5 and later 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. The interface type can be updated after stopping all virtual machines that the disk is attached to.
Data Center	This field only appears when creating a floating disk. 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.

Field Name	Description
Allocation Policy	The provisioning policy for the new virtual disk. • **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 thin 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. Thin provisioned disks are faster to create than preallocated disks and allow for storage over-commitment. Thin provisioned virtual disks are recommended for desktops.
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.
Activate Disk(s)	This field only appears when creating an attached disk. Activate the virtual disk immediately after creation.
Wipe After Delete	Allows you to enable enhanced security for deletion of sensitive material when the virtual disk is deleted.
Bootable	This field only appears when creating an attached disk. 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
Read-Only	This field only appears when creating an attached disk. Allows you to set the disk as read-only. The same disk can be attached as read-only to one virtual machine, and as rewritable to another.
Enable Discard	This field only appears when creating an attached disk. Allows you to shrink a thin provisioned disk while the virtual machine is up. For block storage, the underlying storage device must support discard calls, and the option cannot be used with **Wipe After Delete** unless the underlying storage supports the discard_zeroes_data property. For file storage, the underlying file system and the block device must support discard calls. If all requirements are met, SCSI UNMAP commands issued from guest virtual machines is passed on by QEMU to the underlying storage to free up the unused space.

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.

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.

Table 10.3. New Virtual Disk and Edit Virtual Disk SettingsDirect LUN

Field Name	Description
Alias	The name of the virtual disk, limited to 40 characters.

Field Name	Description
Description	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. 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.
Interface	This field only appears when creating an attached disk. The virtual interface the disk presents to virtual machines. **VirtIO** is faster, but requires drivers. Red Hat Enterprise Linux 5 and later 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. The interface type can be updated after stopping all virtual machines that the disk is attached to.
Data Center	This field only appears when creating a floating disk. 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**.

Field Name	Description
Discover Targets	This section can be expanded when you are using iSCSI external LUNs and **Targets > LUNs** is selected. **Address** – The host name or IP address of the target server. **Port** – The port by which to attempt a connection to the target server. The default port is 3260. **User Authentication** – The iSCSI server requires User Authentication. The **User Authentication** field is visible when you are using iSCSI external LUNs. **CHAP user name** – 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. **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.
Activate Disk(s)	This field only appears when creating an attached disk. Activate the virtual disk immediately after creation.
Bootable	This field only appears when creating an attached disk. 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.
Read-Only	This field only appears when creating an attached disk. Allows you to set the disk as read-only. The same disk can be attached as read-only to one virtual machine, and as rewritable to another.
Enable Discard	This field only appears when creating an attached disk. Allows you to shrink a thin provisioned disk while the virtual machine is up. With this option enabled, SCSI UNMAP commands issued from guest virtual machines is passed on by QEMU to the underlying storage to free up the unused space.

Field Name	Description
Enable SCSI Pass-Through	This field only appears when creating an attached disk. 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. **Read-Only** is not supported when this check box is selected. When this check box is not selected, the virtual disk uses an emulated SCSI device. **Read-Only** is supported on emulated VirtIO-SCSI disks.
Allow Privileged SCSI I/O	This field only appears when creating an attached disk. 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.
Using SCSI Reservation	This field only appears when creating an attached disk. Available when the **Enable SCSI Pass-Through** and **Allow Privileged SCSI I/O** check boxes are selected. Selecting this check box disables migration for any virtual machine using this disk, to prevent virtual machines that are using SCSI reservation from losing access to the disk.

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 11.2.4, "Adding an OpenStack Block Storage (Cinder) Instance for Storage Management" for more information.

Table 10.4. New Virtual Disk and Edit Virtual Disk SettingsCinder

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.

Field Name	Description
Description	A description of the virtual disk. This field is recommended but not mandatory.
Interface	This field only appears when creating an attached disk. The virtual interface the disk presents to virtual machines. **VirtIO** is faster, but requires drivers. Red Hat Enterprise Linux 5 and later 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. The interface type can be updated after stopping all virtual machines that the disk is attached to.
Data Center	This field only appears when creating a floating disk. 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.
Activate Disk(s)	This field only appears when creating an attached disk. Activate the virtual disk immediately after creation.
Bootable	This field only appears when creating an attached disk. 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
Read-Only	This field only appears when creating an attached disk. Allows you to set the disk as read-only. The same disk can be attached as read-only to one virtual machine, and as rewritable to another.

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**).

10.6.3. 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.

10.6.4. 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.

Moving a Virtual Disk

1. Click **Storage → Disks** and select one or more virtual disks to move.

2. Click **Move**.

3. From the **Target** list, select the storage domain to which the virtual disk(s) will be moved.

4. From the **Disk Profile** list, select a profile for the disk(s), if applicable.

5. Click **OK**.

The virtual disks are moved to the target storage domain. During the move procedure, the **Status** column displays **Locked** and a progress bar indicating the progress of the move operation.

10.6.5. 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.

Changing a Disk Interface Type

1. Click **Compute → Virtual Machines** and stop the appropriate virtual machine(s).

2. Click the virtual machine's name to open the details view.

3. Click the **Disks** tab and select the disk.

4. Click **Edit**.

5. From the **Interface** list, select the new interface type and click **OK**.

You can attach a disk to a different virtual machine that requires a different interface type.

Attaching a Disk to a Different Virtual Machine using a Different Interface Type

1. Click **Compute → Virtual Machines** and stop the appropriate virtual machine(s).

2. Click the virtual machine's name to open the details view.

3. Click the **Disks** tab and select the disk.

4. Click **Remove**, then click **OK**.

5. Go back to **Virtual Machines** and click the name of the new virtual machine that the disk will be attached to.

6. Click the **Disks** tab, then click **Attach**.

7. Select the disk in the **Attach Virtual Disks** window and select the appropriate interface from the **Interface** drop-down.

8. Click **OK**.

10.6.6. Copying a Virtual Disk

You can copy a virtual disk from one storage domain to another. The copied disk can be attached to virtual machines.

Copying a Virtual Disk

1. Click **Storage → Disks** and select the virtual disk(s).

2. Click **Copy** .

3. Optionally, enter a new name in the **Alias** field.

4. From the **Target** list, select the storage domain to which the virtual disk(s) will be copied.

5. From the **Disk Profile** list, select a profile for the disk(s), if applicable.

6. Click **OK**.

The virtual disks have a status of **Locked** while being copied.

10.6.7. Uploading Images to a Data Storage Domain

You can upload virtual disk images and ISO images to your data storage domain in the Administration Portal or with the REST API. See Section 8.8.1, "Uploading Images to a Data Storage Domain" .

10.6.8. Importing a Disk Image from an Imported Storage Domain

Import floating virtual disks from an imported storage domain.
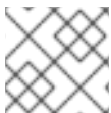
> **NOTE**
>
> Only QEMU-compatible disks can be imported into the Manager.

Importing a Disk Image

1. Click **Storage → Domains**.

2. Click the name of an imported storage domain to open the details view.

3. Click the **Disk Import** tab.

4. Select one or more disks and click **Import**.

5. Select the appropriate **Disk Profile** for each disk.

6. Click **OK**.

10.6.9. Importing an Unregistered Disk Image from an Imported Storage Domain

Import floating virtual disks from a storage domain. 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.
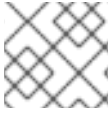
> **NOTE**
>
> Only QEMU-compatible disks can be imported into the Manager.

Importing a Disk Image

1. Click **Storage → Domains**.

2. Click the storage domain's name to open the details view.

3. Click **More Actions → Scan Disks** so that the Manager can identify unregistered disks.

4. Click the **Disk Import** tab.

5. Select one or more disk images and click **Import**.

6. Select the appropriate **Disk Profile** for each disk.

7. Click **OK**.

10.6.10. Importing a Virtual Disk from an OpenStack Image Service

Virtual disks 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 **Storage → Domains**.

2. Click the OpenStack Image Service domain's name to open the details view.

3. Click the **Images** tab and select an image.

4. Click **Import**.

5. Select the **Data Center** into which the image will be imported.

6. From the **Domain Name** drop-down list, select the storage domain in which the image will be stored.

7. Optionally, select a quota to apply to the image from the **Quota** drop-down list.

8. Click **OK**.

The disk can now be attached to a virtual machine.

10.6.11. Exporting a Virtual Disk to an OpenStack Image Service

Virtual disks can be exported to an OpenStack Image Service that has been added to the Manager as an external provider.
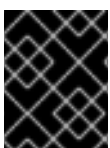
> **IMPORTANT**
>
> Virtual disks can only be exported if they do not have multiple volumes, are not thin provisioned, and do not have any snapshots.

1. Click **Storage → Disks** and select the disks to export.

2. Click **More Actions → Export**.

3. From the **Domain Name** drop-down list, select the OpenStack Image Service to which the disks will be exported.

4. From the **Quota** drop-down list, select a quota for the disks if a quota is to be applied.

5. Click **OK**.

10.6.12. Reclaiming Virtual Disk Space

Virtual disks that use thin provisioning do not automatically shrink after deleting files from them. For example, if the actual disk size is 100GB and you delete 50GB of files, the allocated disk size remains at 100GB, and the remaining 50GB is not returned to the host, and therefore cannot be used by other virtual machines. This unused disk space can be reclaimed by the host by performing a sparsify operation on the virtual machine's disks. This transfers the free space from the disk image to the host. You can sparsify multiple virtual disks in parallel.

Red Hat recommends performing this operation before cloning a virtual machine, creating a template based on a virtual machine, or cleaning up a storage domain's disk space.

Limitations

- NFS storage domains must use NFS version 4.2 or higher.

- You cannot sparsify a disk that uses a direct LUN or Cinder.

- You cannot sparsify a disk that uses a preallocated allocation policy. If you are creating a virtual machine from a template, you must select **Thin** from the **Storage Allocation** field, or if selecting **Clone**, ensure that the template is based on a virtual machine that has thin provisioning.

- You can only sparsify active snapshots.

Sparsifying a Disk

1. Click **Compute → Virtual Machines** and shut down the required virtual machine.

2. Click the virtual machine's name to open the details view.

3. Click the **Disks** tab. Ensure that the disk's status is **OK**.

4. Click **More Actions → Sparsify**.

5. Click **OK**.

A **Started to sparsify** event appears in the **Events** tab during the sparsify operation and the disk's status displays as **Locked**. When the operation is complete, a **Sparsified successfully** event appears in the **Events** tab and the disk's status displays as **OK**. The unused disk space has been returned to the host and is available for use by other virtual machines.

CHAPTER 11. EXTERNAL PROVIDERS

11.1. 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 disks.

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.

RHEL 5 Xen for Virtual Machine Provisioning

Virtual machines created in RHEL 5 Xen can be converted using V2V (**virt-v2v**) and imported into a Red Hat Virtualization environment. After you add a RHEL 5 Xen host 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.

KVM for Virtual Machine Provisioning

Virtual machines created in KVM can be imported into a Red Hat Virtualization environment. After you add a KVM host to the Manager, you can import the virtual machines it provides.

Open Virtual Network (OVN) for Network Provisioning

Open Virtual Network (OVN) is an Open vSwitch (OVS) extension that provides software-defined networks. After you add OVN to the Manager, you can import existing OVN networks, and create new OVN networks from the Manager. You can also automatically install OVN on the Manager using **engine-setup**.

External Network Provider for Network Provisioning

Supported external sofware-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.

11.2. ADDING EXTERNAL PROVIDERS

11.2.1. 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 4.2 is supported with Red Hat Satellite 6.1.

Adding a Satellite Instance for Host Provisioning

1. Click **Administration → Providers**.

2. Click **Add**.

3. Enter a **Name** and **Description**.

4. Select **Foreman/Satellite** from the **Type** drop-down list.

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. Select the **Requires Authentication** check box.

7. 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.

8. 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 to ensure the Manager can communicate with the instance.

9. Click **OK**.

11.2.2. Adding an OpenStack Image (Glance) Instance for Image Management

Add an OpenStack Image (Glance) instance for image management to the Red Hat Virtualization Manager.

Adding an OpenStack Image (Glance) Instance for Image Management

1. Click **Administration → Providers**.

2. Click **Add** and enter the details in the **General Settings** tab. For more information on these fields, see Section 11.2.10, "Add Provider General Settings Explained" .

3. Enter a **Name** and **Description**.

4. Select **OpenStack Image** from the **Type** drop-down list.

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** and **Password** for the OpenStack Image instance user registered in Keystone. You must also define the authentication URL of the Keystone server by defining the **Protocol** (must be **HTTP**), **Hostname**, and **API Port.**
   Enter the **Tenant** for the OpenStack Image instance.

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 to ensure the Manager can communicate with the instance.

8. Click **OK**.

11.2.3. 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 11.2.9, "Adding an External Network Provider" .

> **IMPORTANT**
>
> Red Hat Virtualization supports Red Hat OpenStack Platform versions 8, 9, 10, 11, and 12 as external network providers.

To use neutron networks, hosts must have the neutron agents configured. You can configure the agents manually, or use the Red Hat OpenStack Platform director to deploy the Networker role, before adding the network node to the Manager as a host. Using the director is recommended. Automatic deployment of the neutron agents through the **Network Provider** tab in the **New Host** window is not supported.

Although network nodes and regular hosts can be used in the same cluster, virtual machines using neutron networks can only run on network nodes.

Adding a Network Node as a Host

1. Use the Red Hat OpenStack Platform director to deploy the Networker role on the network node. See Creating a New Role and Networker in the *Red Hat OpenStack Platform Advanced Overcloud Customization Guide*.

2. Enable the Red Hat Virtualization repositories. See Enabling the Red Hat Enterprise Linux Host Repositories in the *Installation Guide*.

3. Install the Openstack Networking hook:

   ```
   # yum install vdsm-hook-openstacknet
   ```

4. Add the network node to the Manager as a host. See Section 7.5.1, "Adding a Host to the Red Hat Virtualization Manager".

   > **IMPORTANT**
   >
   > Do not select the OpenStack Networking provider from the **Network Provider** tab. This is currently not supported.

5. Remove the firewall rule that rejects ICMP traffic:

   ```
   # iptables -D INPUT -j REJECT --reject-with icmp-host-prohibited
   ```

Adding an OpenStack Networking (Neutron) Instance for Network Provisioning

1. Click **Administration → Providers**.

2. Click **Add** and enter the details in the **General Settings** tab. For more information on these fields, see Section 11.2.10, "Add Provider General Settings Explained" .

3. Enter a **Name** and **Description**.

4. Select **OpenStack Networking** from the **Type** drop-down list.

5. Ensure that **Open vSwitch** is selected in the **Networking Plugin** field.

6. Optionally, select the **Automatic Synchronization** check box. This enables automatic synchronization of the external network provider with existing networks.

7. 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.

8. Optionally, select the **Requires Authentication** check box and enter the **Username** and **Password** for the OpenStack Networking user registered in Keystone. You must also define the authentication URL of the Keystone server by defining the **Protocol, Hostname, API Port,** and **API Version**.
   For API version 2.0, enter the **Tenant** for the OpenStack Networking instance. For API version 3, enter the **User Domain Name**, **Project Name**, and **Project Domain Name**.

9. 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.

10. Click the **Agent Configuration** tab.

> ⚠️ **WARNING**
>
> The following steps are provided only as a Technology Preview. Red Hat Virtualization only supports preconfigured neutron hosts.

11. Enter a comma-separated list of interface mappings for the Open vSwitch agent in the **Interface Mappings** field.

12. Select the message broker type that the OpenStack Networking instance uses from the **Broker Type** list.

13. Enter the URL or fully qualified domain name of the host on which the message broker is hosted in the **Host** field.

14. 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.

15. Enter the **Username** and **Password** of the OpenStack Networking user registered in the message broker instance.

16. 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".

11.2.4. Adding an OpenStack Block Storage (Cinder) Instance for Storage Management
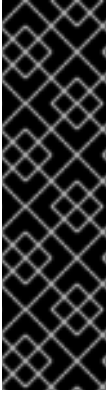
IMPORTANT

Using an OpenStack Block Storage (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 Block Storage (Cinder) instance for storage management to the Red Hat Virtualization Manager. The OpenStack Cinder volumes are provisioned by Ceph Storage.

Adding an OpenStack Block Storage (Cinder) Instance for Storage Management

1. Click **Administration → Providers**.

2. Click **Add** and enter the details in the **General Settings** tab. For more information on these fields, see Section 11.2.10, "Add Provider General Settings Explained" .

3. Enter a **Name** and **Description**.

4. Select **OpenStack Block Storage** from the **Type** drop-down list.

5. Select the **Data Center** to which OpenStack Block Storage volumes will be attached.

6. Enter the URL or fully qualified domain name of the machine on which the OpenStack Block Storage 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** and **Password** for the OpenStack Block Storage instance user registered in Keystone. Define the authentication URL of the Keystone server by defining the **Protocol** (must be **HTTP**), **Hostname**, and **API Port**.
Enter the **Tenant** for the OpenStack Block Storage instance.

8. Click **Test** to test whether you can authenticate successfully with the OpenStack Block Storage 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 Configuration 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** 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
```

See Section 10.6.1, "Creating a Virtual Disk" for more information about creating a OpenStack Block Storage (Cinder) disk.

11.2.5. 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) and is installed on Red Hat Enterprise Linux hosts as a dependency of VDSM when added to the Red Hat Virtualization environment. Red Hat Enterprise Linux hosts must be Red Hat Enterprise Linux 7.2 or later.

> **NOTE**
>
> The **virt-v2v** package is not available on ppc64le architecture; these hosts cannot be used as proxy hosts.

Adding a VMware vCenter Instance as a Virtual Machine Provider

1. Click **Administration** → **Providers**.

2. Click **Add**.

3. Enter a **Name** and **Description**.

4. Select **VMware** from the **Type** drop-down list.

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.

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 the **Verify server's SSL certificate** check box selected to verify the ESXi host's certificate. If not, clear the check box.

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.

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 to ensure the Manager can communicate with the instance.

13. Click **OK**.

To import virtual machines from the VMware external provider, see Importing a Virtual Machine from a VMware Provider in the *Virtual Machine Management Guide*.

11.2.6. Adding a RHEL 5 Xen Host as a Virtual Machine Provider

Add a RHEL 5 Xen host to import virtual machines from Xen to Red Hat Virtualization.

Red Hat Virtualization uses V2V to convert RHEL 5 Xen 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) and is installed on Red Hat Enterprise Linux hosts as a dependency of VDSM when added to the Red Hat Virtualization environment. Red Hat Enterprise Linux hosts must be Red Hat Enterprise Linux 7.2 or later.

> **NOTE**
>
> The **virt-v2v** package is not available on ppc64le architecture; these hosts cannot be used as proxy hosts.

Adding a RHEL 5 Xen Instance as a Virtual Machine Provider

1. Enable public key authentication between the proxy host and the RHEL 5 Xen host:

    a. Log in to the proxy host and generate SSH keys for the **vdsm** user.

    ```
    # sudo -u vdsm ssh-keygen
    ```

    b. Copy the **vdsm** user's public key to the RHEL 5 Xen host. The proxy host's **known_hosts** file will also be updated to include the host key of the RHEL 5 Xen host.

    ```
    # sudo -u vdsm ssh-copy-id root@xenhost.example.com
    ```

    c. Log in to the RHEL 5 Xen host to verify that the login works correctly.

    ```
    # sudo -u vdsm ssh root@xenhost.example.com
    ```

2. Click **Administration** → **Providers**.

3. Click **Add**.

4. Enter a **Name** and **Description**.

5. Select **XEN** from the **Type** drop-down list.

6. Select the **Data Center** into which Xen virtual machines will be imported, or select **Any Data Center** to specify the destination data center during individual import operations.

7. Enter the URI of the RHEL 5 Xen host in the **URI** field.

8. 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 RHEL 5 Xen 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.

9. Click **Test** to test whether you can authenticate successfully with the RHEL 5 Xen host.

10. Click **OK**.

To import virtual machines from a RHEL 5 Xen external provider, see Importing a Virtual Machine from a RHEL 5 Xen Host in the *Virtual Machine Management Guide* .

11.2.7. Adding a KVM Host as a Virtual Machine Provider

Add a KVM host to import virtual machines from KVM to Red Hat Virtualization Manager.

Adding a KVM Host as a Virtual Machine Provider

1. Enable public key authentication between the proxy host and the KVM host:

   a. Log in to the proxy host and generate SSH keys for the **vdsm** user.

   ```
   # sudo -u vdsm ssh-keygen
   ```

   b. Copy the **vdsm** user's public key to the KVM host. The proxy host's **known_hosts** file will also be updated to include the host key of the KVM host.

   ```
   # sudo -u vdsm ssh-copy-id root@kvmhost.example.com
   ```

   c. Log in to the KVM host to verify that the login works correctly.

   ```
   # sudo -u vdsm ssh root@kvmhost.example.com
   ```

2. Click **Administration → Providers**.

3. Click **Add**.

4. Enter a **Name** and **Description**.

5. Select **KVM** from the **Type** drop-down list.

6. Select the **Data Center** into which KVM virtual machines will be imported, or select **Any Data Center** to specify the destination data center during individual import operations.

7. Enter the URI of the KVM host in the **URI** field.

8. Select a host in the chosen data center to serve as the **Proxy Host** during virtual machine import operations. This host must also be able to connect to the network of the KVM external provider. If you selected **Any Data Center** in the **Data Center** field above, you cannot choose the host here. The field is greyed out and shows **Any Host in Data Center**. Instead you can specify a host during individual import operations.

9. Optionally, select the **Requires Authentication** check box and enter the **Username** and **Password** for the KVM host. The user must have access to the KVM host on which the virtual machines reside.

10. Click **Test** to test whether you can authenticate successfully with the KVM host using the provided credentials.

11. Click **OK**.

To import virtual machines from a KVM external provider, see Importing a Virtual Machine from a KVM Host in the *Virtual Machine Management Guide* .

11.2.8. Adding Open Virtual Network (OVN) as an External Network Provider

Open Virtual Network (OVN) enables you to create networks without adding VLANs or changing the infrastructure. OVN is an Open vSwitch (OVS) extension that enables support for virtual networks by adding native OVS support for virtual L2 and L3 overlays.

You can either install a new OVN network provider or add an existing one .

You can also connect an OVN network to a native Red Hat Virtualization network. See Section 11.2.8.4, "Connecting an OVN Network to a Physical Network" for more information. This feature is available as a Technology Preview only.

A Neutron-like REST API is exposed by **ovirt-provider-ovn**, enabling you to create networks, subnets, ports, and routers (see the *OpenStack Networking API v2.0* for details). These overlay networks enable communication among the virtual machines.
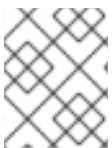
> **NOTE**
>
> OVN is supported as an external provider by CloudForms, using the OpenStack (Neutron) API. See Network Managers in *Red Hat CloudForms: Managing Providers* for details.

For more information on OVS and OVN, see the OVS documentation at http://docs.openvswitch.org/en/latest/ and http://openvswitch.org/support/dist-docs/.

11.2.8.1. Installing a New OVN Network Provider

> **WARNING**
>
> If the **openvswitch** package is already installed and if the version is 1:2.6.1 (version 2.6.1, epoch 1), the OVN installation will fail when it tries to install the latest **openvswitch** package. See the Doc Text in BZ#1505398 for the details and a workaround.

When you install OVN using **engine-setup**, the following steps are automated:

- Setting up an OVN central server on the Manager machine.

- Adding OVN to Red Hat Virtualization as an external network provider.

- Setting the **Default** cluster's default network provider to **ovirt-provider-ovn**.

- Configuring hosts to communicate with OVN when added to the cluster.

If you use a preconfigured answer file with **engine-setup**, you can add the following entry to install OVN:

OVESETUP_OVN/ovirtProviderOvn=bool:True

Installing a New OVN Network Provider

1. Install OVN on the Manager using engine-setup. During the installation, **engine-setup** asks the following questions:

   # Install ovirt-provider-ovn(Yes, No) [Yes]?:

   - If **Yes**, engine-setup installs **ovirt-provider-ovn**. If **engine-setup** is updating a system, this prompt only appears if **ovirt-provider-ovn** has not been installed previously.

   - If **No**, you will not be asked again on the next run of **engine-setup**. If you want to see this option, run **engine-setup --reconfigure-optional-components**.

     # Use default credentials (admin@internal) for ovirt-provider-ovn(Yes, No) [Yes]?:

   If **Yes**, **engine-setup** uses the default engine user and password specified earlier in the setup process. This option is only available during new installations.

   # oVirt OVN provider user[admin]:
   # oVirt OVN provider password[empty]:

   You can use the default values or specify the oVirt OVN provider user and password.

   NOTE

   To change the authentication method later, you can edit the **/etc/ovirt-provider-ovn/conf.d/10_engine_setup.conf** file, or create a new **/etc/ovirt-provider-ovn/conf.d/20_engine_setup.conf** file. Restart the **ovirt-provider-ovn** service for the change to take effect. See https://github.com/oVirt/ovirt-provider-ovn/blob/master/README.adoc for more information about OVN authentication.

2. Add hosts to the **Default** cluster. Hosts added to this cluster are automatically configured to communicate with OVN. To add new hosts, see Section 7.5.1, "Adding a Host to the Red Hat Virtualization Manager".
   To configure your hosts to use an existing, non-default network, see Section 11.2.8.3, "Configuring Hosts for an OVN Tunnel Network".

3. Add networks to the **Default** cluster; see Section 6.1.2, "Creating a New Logical Network in a Data Center or Cluster" and select the **Create on external provider** check box. **ovirt-provider-ovn** is selected by default.

4. To connect the OVN network to a native Red Hat Virtualization network, select the **Connect to physical network** check box and specify the Red Hat Virtualization network to use. See Section 11.2.8.4, "Connecting an OVN Network to a Physical Network" for more information and prerequisites.

5. Define whether the network should use Security Groups from the **Security Groups** drop-down. For more information on the available options see Section 6.1.7, "Logical Network General Settings Explained". You can now create virtual machines that use OVN networks.

11.2.8.2. Adding an Existing OVN Network Provider

Adding an existing OVN central server as an external network provider in Red Hat Virtualization involves the following key steps:

- Install the OVN provider, a proxy used by the Manager to interact with OVN. The OVN provider can be installed on any machine, but must be able to communicate with the OVN central server and the Manager.

- Add the OVN provider to Red Hat Virtualization as an external network provider.

- Create a new cluster that uses OVN as its default network provider. Hosts added to this cluster are automatically configured to communicate with OVN.

Prerequisites

The following packages are required by the OVN provider and must be available on the provider machine:

- openvswitch-ovn-central

- openvswitch

- openvswitch-ovn-common

- python-openvswitch

If these packages are not available from the repositories already enabled on the provider machine, they can be downloaded from the OVS website: http://openvswitch.org/download/.

Adding an Existing OVN Network Provider

1. Install and configure the OVN provider.

    a. Install the provider on the provider machine:

        ```
        # yum install ovirt-provider-ovn
        ```

    b. If you are not installing the provider on the same machine as the Manager, add the following entry to the /**etc/ovirt-provider-ovn/conf.d/10_engine_setup.conf** file (create this file if it does not already exist):

        ```
        [OVIRT]
        ovirt-host=https://Manager_host_name
        ```

This is used for authentication, if authentication is enabled.

c. If you are not installing the provider on the same machine as the OVN central server, add the following entry to the **/etc/ovirt-provider-ovn/conf.d/10_engine_setup.conf** file (create this file if it does not already exist):

> [OVN REMOTE]
> ovn-remote=tcp:*OVN_central_server_IP*:6641

d. Open ports 9696, 6641, and 6642 in the firewall to allow communication between the OVN provider, the OVN central server, and the Manager. This can be done either manually or by adding the **ovirt-provider-ovn** and **ovirt-provider-ovn-central** services to the appropriate zone:

> # firewall-cmd --zone=*ZoneName* --add-service=ovirt-provider-ovn --permanent
> # firewall-cmd --zone=*ZoneName* --add-service=ovirt-provider-ovn-central --permanent
> # firewall-cmd --reload

e. Start and enable the service:

> # systemctl start ovirt-provider-ovn
> # systemctl enable ovirt-provider-ovn

f. Configure the OVN central server to listen to requests from ports 6642 and 6641:

> # ovn-sbctl set-connection ptcp:6642
> # ovn-nbctl set-connection ptcp:6641

2. In the Administration Portal, click **Administration → Providers**.

3. Click **Add** and enter the details in the **General Settings** tab. For more information on these fields, see Section 11.2.10, "Add Provider General Settings Explained" .

4. Enter a **Name** and **Description**.

5. From the **Type** list, select **External Network Provider**.

6. Click the **Networking Plugin** text box and select **oVirt Network Provider for OVN** from the drop-down menu.

7. Optionally, select the **Automatic Synchronization** check box. This enables automatic synchronization of the external network provider with existing networks.
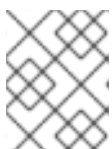
> **NOTE**
>
> Automatic synchronization is enabled by default on the **ovirt-provider-ovn** network provider created by the **engine-setup** tool.

8. Enter the URL or fully qualified domain name of the OVN provider in the **Provider URL** text field, followed by the port number. If the OVN provider and the OVN central server are on separate machines, this is the URL of the provider machine, not the central server. If the OVN provider is on the same machine as the Manager, the URL can remain the default **http://localhost:9696**.

9. Clear the **Read-Only** check box to allow creating new OVN networks from the Red Hat Virtualization Manager.

10. Optionally, select the **Requires Authentication** check box and enter the **Username** and **Password** for the for the external network provider user registered in Keystone. You must also define the authentication URL of the Keystone server by defining the **Protocol**, **Hostname**, and **API Port**.
    Optionally, enter the **Tenant** for the external network provider.

    The authentication method must be configured in the **/etc/ovirt-provider-ovn/conf.d/10_engine_setup.conf** file (create this file if it does not already exist). Restart the **ovirt-provider-ovn** service for the change to take effect. See https://github.com/oVirt/ovirt-provider-ovn/blob/master/README.adoc for more information about OVN authentication.

11. Test the credentials:

    a. Click **Test** to test whether you can authenticate successfully with OVN using the provided credentials.

    b. If the OVN instance uses SSL, the **Import provider certificates** window opens; click **OK** to import the certificate that the OVN instance provides to ensure the Manager can communicate with the instance.

12. Click **OK**.

13. Create a new cluster that uses OVN as its default network provider. See Section 5.2.1, "Creating a New Cluster" and select the OVN network provider from the **Default Network Provider** drop-down list.

14. Add hosts to the cluster. Hosts added to this cluster are automatically configured to communicate with OVN. To add new hosts, see Section 7.5.1, "Adding a Host to the Red Hat Virtualization Manager".

15. Import or add OVN networks to the new cluster. To import networks, see]. To create new networks using OVN, see xref:Creating_a_new_logical_network_in_a_data_center_or_cluster[, and select the **Create on external provider** check box. **ovirt-provider-ovn** is selected by default.
    To configure your hosts to use an existing, non-default network, see Section 11.2.8.3, "Configuring Hosts for an OVN Tunnel Network".

    To connect the OVN network to a native Red Hat Virtualization network, select the **Connect to physical network** check box and specify the Red Hat Virtualization network to use. See Section 11.2.8.4, "Connecting an OVN Network to a Physical Network" for more information and prerequisites.

You can now create virtual machines that use OVN networks.

11.2.8.3. Configuring Hosts for an OVN Tunnel Network

You can configure your hosts to use an existing network, other than the default **ovirtmgmt** network, with the **ovirt-provider-ovn-driver** Ansible playbook. The network must be accessible to all the hosts in the cluster.

> **NOTE**
>
> The **ovirt-provider-ovn-driver** Ansible playbook updates existing hosts. If you add new hosts to the cluster, you must run the playbook again.

Configuring Hosts for an OVN Tunnel Network

1. On the Manager machine, go to the **playbooks** directory:

   ```
   # cd /usr/share/ovirt-engine/playbooks
   ```

2. Run the **ansible-playbook** command with the following parameters:

   ```
   # ansible-playbook --private-key=/etc/pki/ovirt-engine/keys/engine_id_rsa -i /usr/share/ovirt-engine-metrics/bin/ovirt-engine-hosts-ansible-inventory --extra-vars
   " cluster_name=Cluster_Name ovn_central=OVN_Central_IP
   ovn_tunneling_interface=VDSM_Network_Name" ovirt-provider-ovn-driver.yml
   ```

   For example:

   ```
   # ansible-playbook --private-key=/etc/pki/ovirt-engine/keys/engine_id_rsa -i /usr/share/ovirt-engine-metrics/bin/ovirt-engine-hosts-ansible-inventory --extra-vars
   " cluster_name=MyCluster ovn_central=192.168.0.1 ovn_tunneling_interface=MyNetwork"
   ovirt-provider-ovn-driver.yml
   ```

> **NOTE**
>
> The *OVN_Central_IP* can be on the new network, but this is not a requirement. The *OVN_Central_IP* must be accessible to all hosts.
>
> The *VDSM_Network_Name* is limited to 15 characters. If you defined a logical network name that was longer than 15 characters or contained non–ASCII characters, a 15–character name is automatically generated. See Mapping VDSM Names to Logical Network Names for instructions on displaying a mapping of these names.

Updating the OVN Tunnel Network on a Single Host

You can update the OVN tunnel network on a single host with **vdsm-tool**:

```
# vdsm-tool ovn-config OVN_Central_IP Tunneling_IP_or_Network_Name
```

Example 11.1. Updating a Host with vdsm-tool

```
# vdsm-tool ovn-config 192.168.0.1 MyNetwork
```

11.2.8.4. Connecting an OVN Network to a Physical Network

> **IMPORTANT**
>
> This feature relies on Open vSwitch support, which is available only as a Technology Preview in Red Hat Virtualization. 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 using 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 an external provider network that overlays a native Red Hat Virtualization network so that the virtual machines on each appear to be sharing the same subnet.

> **IMPORTANT**
>
> If you created a subnet for the OVN network, a virtual machine using that network will receive an IP address from there. If you want the physical network to allocate the IP address, do not create a subnet for the OVN network.

Prerequisites

- The cluster must have **OVS** selected as the **Switch Type**. Hosts added to this cluster must not have any pre-existing Red Hat Virtualization networks configured, such as the **ovirtmgmt** bridge.

- The physical network must be available on the hosts. You can enforce this by setting the physical network as required for the cluster (in the **Manage Networks** window, or the **Cluster** tab of the **New Logical Network** window).

Creating a New External Network Connected to a Physical Network

1. Click **Compute → Clusters**.

2. Click the cluster's name to open the details view.

3. Click the **Logical Networks** tab and click **Add Network**.

4. Enter a **Name** for the network.

5. Select the **Create on external provider** check box. **ovirt-provider-ovn** is selected by default.

6. Select the **Connect to physical network** check box if it is not already selected by default.

7. Choose the physical network to connect the new network to:

   - Click the **Data Center Network** radio button and select the physical network from the drop-down list. This is the recommended option.

   - Click the **Custom** radio button and enter the name of the physical network. If the physical network has VLAN tagging enabled, you must also select the **Enable VLAN tagging** check box and enter the physical network's VLAN tag.

> **IMPORTANT**
>
> The physical network's name must not be longer than 15 characters, or contain special characters.

8. Click **OK**.

11.2.9. 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.

Adding an External Network Provider for Network Provisioning

1. Click **Administration → Providers**.

2. Click **Add** and enter the details in the **General Settings** tab. For more information on these fields, see Section 11.2.10, "Add Provider General Settings Explained" .

3. Enter a **Name** and **Description**.

4. Select **External Network Provider** from the **Type** drop-down list.

5. Optionally, click the **Networking Plugin** text box and select the appropriate driver from the drop-down menu.

6. Optionally, select the **Automatic Synchronization** check box. This enables automatic synchronization of the external network provider with existing networks. This feature is disabled by default when adding external network providers.
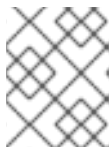
   > **NOTE**
   >
   > Automatic synchronization is enabled by default on the **ovirt-provider-ovn** network provider created by the **engine-setup** tool.

7. 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.
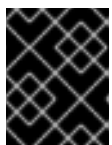
   > **IMPORTANT**
   >
   > You must leave the **Read-Only** check box selected for your setup to be supported by Red Hat.

8. Optionally, select the **Requires Authentication** check box and enter the **Username** and **Password** for the external network provider user registered in Keystone. You must also define the authentication URL of the Keystone server by defining the **Protocol**, **Hostname**, and **API Port**.
   Optionally, enter the **Tenant** for the external network provider.

9. 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.

10. Click **OK**.

Before you can use networks from this provider, you must 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".

11.2.10. 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 11.1. 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	The type of external provider. Changing this setting alters the available fields for configuring the provider.
	Foreman/Satellite
	● **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.
	● **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.
	OpenStack Image

Setting	Explanation
	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 service to the end of the URL or fully qualified domain name. By default, this port number is 9292.
	Requires Authentication: Allows you to specify whether authentication is required to access the OpenStack Image service.
	Username: A user name for connecting to the Keystone server. 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.
	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.
	Protocol: The protocol used to communicate with the Keystone server. This must be set to **HTTP**.
	Hostname: The IP address or hostname of the Keystone server.
	API port: The API port number of the Keystone server.
	API Version: The version of the Keystone service. The value is **v2.0** and the field is disabled.
	Tenant Name: The name of the OpenStack tenant of which the OpenStack Image service is a member.
	OpenStack Networking
	Networking Plugin: The networking plugin with which to connect to the OpenStack Networking server. For OpenStack Networking, **Open vSwitch** is the only option, and is selected by default.
	Automatic Synchronization: Allows you to specify whether the provider will be automatically synchronized with existing networks.
	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.

Setting	Explanation
	- **Read Only**: Allows you to specify whether 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 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. - **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. - **Protocol**: The protocol used to communicate with the Keystone server. The default is **HTTPS**. - **Hostname**: The IP address or hostname of the Keystone server. - **API port**: The API port number of the Keystone server. - **API Version**: The version of the Keystone server. This appears in the URL. If v2.0 appears, select **v2.0**. If v3 appears select **v3**. The following fields appear when you select **v3** from the **API Version** field: - **User Domain Name**: The name of the user defined in the domain. With Keystone API v3, domains are used to determine administrative boundaries of service entities in OpenStack. Domains allow you to group users together for various purposes, such as setting domain-specific configuration or security options. For more information, see OpenStack Identity (keystone) in the Red Hat OpenStack Platform *Architecture Guide*. - **Project Name**: Defines the project name for OpenStack Identity API v3. - **Project Domain Name**: Defines the project's domain name for OpenStack Identity API v3. The following field appears when you select **v2.0** from the **API Version** field: - **Tenant Name**: Appears only when v2 is selected from the **API Version** field. The name of the OpenStack tenant of which the

Setting	Explanation
	OpenStack Networking instance is a member.

OpenStack Volume

- **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 Keystone server. 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.

- **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.

- **Protocol**: The protocol used to communicate with the Keystone server. This must be set to **HTTP**.

- **Hostname**: The IP address or hostname of the Keystone server.

- **API port**: The API port number of the Keystone server.

- **API Version**: The version of the Keystone server. The value is **v2.0** and the field is disabled.

- **Tenant Name**: The name of the OpenStack tenant of which the OpenStack Volume instance is a member.

VMware

- **Data Center**: Specify the data center into which VMware virtual machines will be imported, or select **Any Data Center** to specify the destination data center during individual import operations (using the **Import** function in the **Virtual Machines** tab).

Setting	Explanation
	- **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 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.
	RHEL 5 Xen
	- **Data Center**: Specify the data center into which Xen 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).
	- **URI**: The URI of the RHEL 5 Xen host.
	- **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 RHEL 5 Xen 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).
	KVM

Setting	Explanation
	• **Data Center**: Specify the data center into which KVM 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).

• **URI**: The URI of the KVM host.

• **Proxy Host**: Select a host in the chosen data center to serve as the host during virtual machine import operations. This host must also be able to connect to the network of the KVM 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).

• **Requires Authentication**: Allows you to specify whether authentication is required to access the KVM host.

• **Username**: A user name for connecting to the KVM host.

• **Password**: The password against which the above user name is to be authenticated.

External Network Provider

• **Networking Plugin**: Determines which implementation of the driver will be used on the host to handle NIC operations. If an external network provider with the **oVirt Network Provider for OVN** plugin is added as the default network provider for a cluster, this also determines which driver will be installed on hosts added to the cluster.

• **Automatic Synchronization**: Allows you to specify whether the provider will be automatically synchronized with existing networks.

• **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. If you are

Setting	Explanation
	authenticating with Active Directory, the user name must be in the format of *username@domain@auth_profile* instead of the default *username@domain*. ● **Password**: The password against which the above user name is to be authenticated. ● **Protocol**: The protocol used to communicate with the Keystone server. The default is **HTTPS**. ● **Hostname**: The IP address or hostname of the Keystone server. ● **API port**: The API port number of the Keystone server. ● **API Version**: The version of the Keystone server. The value is **v2.0** and the field is disabled. ● **Tenant Name**: Optional. The name of the tenant of which the external network provider is a member.
Test	Allows users to test the specified credentials. This button is available to all provider types.

11.2.11. 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 11.2. Add Provider: Agent Configuration Settings

Setting	Explanation
Interface Mappings	A comma-separated list of mappings in the format of *label:interface*.
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.
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.

Setting	Explanation
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.

11.3. EDITING AN EXTERNAL PROVIDER

Editing an External Provider

1. Click **Administration → Providers** and select the external provider to edit.

2. Click **Edit**.

3. Change the current values for the provider to the preferred values.

4. Click **OK**.

11.4. REMOVING AN EXTERNAL PROVIDER

Removing an External Provider

1. Click **Administration → Providers** and select the external provider to remove.

2. Click **Remove**.

3. Click **OK**.

PART III. ADMINISTERING THE ENVIRONMENT

CHAPTER 12. BACKUPS AND MIGRATION

12.1. BACKING UP AND RESTORING THE RED HAT VIRTUALIZATION MANAGER

12.1.1. 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.

12.1.2. 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.

12.1.3. 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
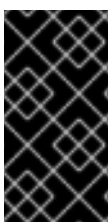
> **IMPORTANT**
>
> To restore a database to a fresh installation of Red Hat Virtualization Manager, a database backup alone is not sufficient; the Manager also requires access to the configuration files. Any backup that specifies a scope other than the default, **all**, must be accompanied by the **files** scope, or a filesystem backup.

Example Usage of the engine-backup Command

1. Log on to the machine running the Red Hat Virtualization Manager.

2. Create a backup:

   Example 12.1. Creating a Full Backup

   ```
   # engine-backup --scope=all --mode=backup --file=file_name --log=log_file_name
   ```

   Example 12.2. 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.

12.1.4. 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.
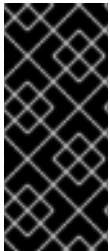
> **IMPORTANT**
>
> Backups can only be restored to environments of the same major release as that of the backup. For example, a backup of a Red Hat Virtualization version 4.2 environment can only be restored to another Red Hat Virtualization version 4.2 environment. To view the version of Red Hat Virtualization 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.

12.1.5. 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.

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".

12.1.6. 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.
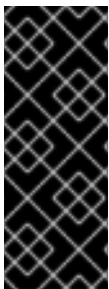
> **IMPORTANT**
>
> When restoring a backup to overwrite an existing installation, you must run the **engine-cleanup** command to clean up the existing installation before using the **engine-backup** command. Because the **engine-cleanup** command only cleans the engine database, and does not drop the 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.

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:

4. Restore a full backup:

```
# engine-backup --mode=restore --file=file_name --log=log_file_name --restore-permissions
```

5. 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.

6. Run the following command and follow the prompts to reconfigure the firewall and ensure the **ovirt-engine** service is correctly configured:

> # engine-setup

12.1.7. 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.
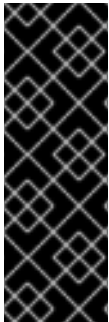
> **IMPORTANT**
>
> When restoring a backup to overwrite an existing installation, you must run the **engine-cleanup** command to clean up the existing installation before using the **engine-backup** command. Because the **engine-cleanup** command only cleans the engine database, and does not drop the 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.

Restoring a Backup with Different Credentials

1. Log on to the Red Hat Virtualization Manager machine.

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:

   a. Enter the postgresql command line:

```
# su - postgres -c 'scl enable rh-postgresql95 -- psql'
```

b. 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**.
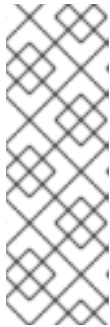
> **NOTE**
>
> The following examples use a **--*password** option for each database without specifying a password, which will prompt for a password for each database. Passwords can be supplied for these options in the command itself, however this is not recommended as the password will then be stored 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
   ```

12.1.8. 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 Enabling the Red Hat Virtualization Manager Repositories in the *Installation Guide*.

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
   ```

12.2. BACKING UP AND RESTORING VIRTUAL MACHINES USING THE BACKUP AND RESTORE API

12.2.1. 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 .

12.2.2. 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.

Backing Up a Virtual Machine

1. Using the REST API, create a snapshot of the virtual machine to back up:

```
POST /api/vms/{vm:id}/snapshots/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<snapshot>
    <description>BACKUP</description>
</snapshot>
```

> **NOTE**
>
> - Here, replace **{vm:id}** with the VM ID of the virtual machine whose snapshot you are making. This ID is available from the **General** tab of the **New Virtual Machine** and **Edit Virtual Machine**windows in the **Administration Portal** and **VM Portal**.
>
> - Taking a snapshot of a virtual machine stores its current configuration data in the **data** attribute of the **configuration** attribute in **initialization** under the snapshot.

> **IMPORTANT**
>
> You cannot take snapshots of disks marked as shareable or based on direct LUN disks.

2. Retrieve the configuration data of the virtual machine from the **data** attribute under the snapshot:

```
GET /api/vms/{vm:id}/snapshots/{snapshot:id} HTTP/1.1
All-Content: true
Accept: application/xml
Content-type: application/xml
```

> **NOTE**
>
> - Here, replace *{vm:id}* with the ID of the virtual machine whose snapshot you made earlier. Replace *{snapshot:id}* with the snapshot ID.
>
> - Add the **All-Content: true** header to retrieve additional OVF data in the response. The OVF data in the XML response is located within the VM configuration element, **\<initialization\>\<configuration\>**. Later, you will use this data to restore the virtual machine.

3. Get the snapshot ID:

   ```
   GET /api/vms/{vm:id}/snapshots/ HTTP/1.1
   Accept: application/xml
   Content-type: application/xml
   ```

4. Identify the disk ID of the snapshot:

   ```
   GET /api/vms/{vm:id}/snapshots/{snapshot:id}/disks HTTP/1.1
   Accept: application/xml
   Content-type: application/xml
   ```

5. Attach the snapshot to a backup virtual machine as an active disk attachment, with the correct interface type (for example, **virtio_scsi**):

   ```
   POST /api/vms/{vm:id}/diskattachments/ HTTP/1.1
   Accept: application/xml
   Content-type: application/xml

   <disk_attachment>
    <active>true</active>
    <interface>_virtio_scsi_</interface>
    <disk id="{disk:id}">
    <snapshot id="{snapshot:id}"/>
    </disk>
   </disk_attachment>
   ```

> **NOTE**
>
> Here, replace *{vm:id}* with the ID of the *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace *{disk:id}* with the disk ID. Replace *{snapshot:id}* with the snapshot ID.

6. Use the backup software on the backup virtual machine to back up the data on the snapshot disk.

7. Remove the snapshot disk attachment from the backup virtual machine:

   ```
   DELETE /api/vms/{vm:id}/diskattachments/{snapshot:id} HTTP/1.1
   Accept: application/xml
   Content-type: application/xml
   ```

NOTE

Here, replace *{vm:id}* with the ID of the *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace *{snapshot:id}* with the snapshot ID.

8. Optionally, delete the snapshot:

```
DELETE /api/vms/{vm:id}/snapshots/{snapshot:id} HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

NOTE

Here, replace *{vm:id}* with the ID of the virtual machine whose snapshot you made earlier. Replace *{snapshot:id}* with the snapshot ID.

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.

12.2.3. 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.

Restoring a Virtual Machine

1. In the Administration Portal, create a floating disk on which to restore the backup. See Section 10.6.1, "Creating a Virtual Disk" for details on how to create a floating disk.

2. Attach the disk to the backup virtual machine:

```
POST /api/vms/{vm:id}/disks/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<disk id="{disk:id}">
</disk>
```

NOTE

Here, replace *{vm:id}* with the ID of this *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace *{disk:id}* with the disk ID you got while backing up the virtual machine.

3. Use the backup software to restore the backup to the disk.

4. Detach the disk from the backup virtual machine:

```
DELETE /api/vms/{vm:id}/disks/{disk:id} HTTP/1.1
Accept: application/xml
Content-type: application/xml
```

```
<action>
    <detach>true</detach>
</action>
```

+ NOTE: Here, replace *{vm:id}* with the ID of this *backup* virtual machine, not the virtual machine whose snapshot you made earlier. Replace *{disk:id}* with the disk ID.

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>
   <initialization>
   <configuration>
 <data>
 < -- omitting long ovf data -->
 </data>
    <type>ovf</type>
    </configuration>
    </initialization>

   ...
</vm>
```
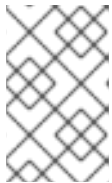
> **NOTE**
>
> To override any of the values in the ovf while creating the virtual machine, redefine the element *before* or *after* the **initialization** element. Not within the initialization element.

6. Attach the disk to the new virtual machine:

```
POST /api/vms/{vm:id}/disks/ HTTP/1.1
Accept: application/xml
Content-type: application/xml

<disk id="{disk:id}">
</disk>
```
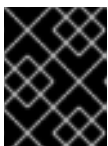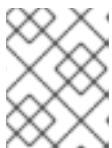
> **NOTE**
>
> Here, replace *{vm:id}* with the ID of the *new* virtual machine, not the virtual machine whose snapshot you made earlier. Replace *{disk:id}* with the disk ID.

You have restored a virtual machine using a backup that was created using the backup and restore API.

CHAPTER 13. 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 4.2 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.

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 11.2.1, "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.

Viewing Red Hat Virtualization Manager Errata

1. Click **Administration → Errata**.

2. Select the **Security**, **Bugs**, or **Enhancements** check boxes to view only those errata types.

For more information on viewing available errata for hosts see Section 7.5.21, "Viewing Host Errata" and for virtual machines see Viewing Red Hat Satellite Errata for a Virtual Machine in the *Virtual Machine Management Guide*.

CHAPTER 14. AUTOMATING CONFIGURATION TASKS USING ANSIBLE

Ansible is an automation tool used to configure systems, deploy software, and perform rolling updates. Ansible includes support for Red Hat Virtualization, and Ansible modules are available to allow you to automate post-installation tasks such as data center setup and configuration, managing users, or virtual machine operations.

Ansible provides an easier method of automating Red Hat Virtualization configuration compared to REST APIs and SDKs, and allows you to integrate with other Ansible modules. For more information about the Ansible modules available for Red Hat Virtualization, see the Ovirt modules in the Ansible documentation.

> **NOTE**
>
> Ansible Tower is a graphically enabled framework accessible through a web interface and REST APIs for Ansible. If you want support for Ansible Tower, then you must have an Ansible Tower license, which is not part of the Red Hat Virtualization subscription.

Ansible is shipped with Red Hat Virtualization. To install Ansible, ensure that you have enabled the required repositories. See Enabling the Red Hat Virtualization Manager Repositories in the *Installation Guide*, and run the following command:

```
# yum install ansible
```

See the Ansible Documentation for alternate installation instructions, and information about using Ansible.

> **NOTE**
>
> To permanently increase the verbose level for the Manager when running Ansible playbooks, create a configuration file in **/etc/ovirt-engine/engine.conf.d/** with following line:
>
> **ANSIBLE_PLAYBOOK_VERBOSE_LEVEL=4**
>
> You must restart the Manager after creating the file by running **systemctl restart ovirt-engine**.

14.1. ANSIBLE ROLES

Multiple Ansible roles are available to help configure and manage various parts of the Red Hat Virtualization infrastructure. Ansible roles provide a method of modularizing Ansible code by breaking up large playbooks into smaller, reusable files that can be shared with other users.

The Ansible roles available for Red Hat Virtualization are categorized by the various infrustructure components. For more information about the Ansible roles, see the oVirt Ansible Roles documentation. For the documentation installed with Ansible roles, see Section 14.1.1, "Installing Ansible Roles".

14.1.1. Installing Ansible Roles

You can install Ansible roles for Red Hat Virtualization from the 'rhel-7-server-rhv-4.2-manager-rpms' repository. See Enabling the Red Hat Virtualization Manager Repositories in the *Installation Guide* for more information.

Use the following command to install the Ansible roles:

```
# yum install ovirt-ansible-roles
```

By default the roles are installed to **/usr/share/ansible/roles**. The structure of the **ovirt-ansible-roles** package is as follows:

- **/usr/share/ansible/roles** – stores the roles.

- **/usr/share/doc/ovirt-ansible-roles/** – stores the examples, a basic overview, and the licence.

- **/usr/share/doc/ansible/roles/***role_name* – stores the documentation specific to the role.

14.1.2. Using Ansible Roles to Configure Red Hat Virtualization

The following procedure guides you through creating and running a playbook that uses Ansible roles to configure Red Hat Virtualization. This example uses Ansible to connect to the Manager on the local machine and create a new data center.

Prerequisites

- Ensure the **roles_path** option in **/etc/ansible/ansible.cfg** points to the location of your Ansible roles (**/usr/share/ansible/roles**).

- Ensure that you have the Python SDK installed on the machine running the playbook.

Configuring Red Hat Virtualization using Ansible Roles

1. Create a file in your working directory to store the Red Hat Virtualization Manager user password:

   ```
   # cat passwords.yml
    ---
   engine_password: youruserpassword
   ```

2. Encrypt the user password. You will be asked for a Vault password.

   ```
   # ansible-vault encrypt passwords.yml
   New Vault password:
   Confirm New Vault password:
   ```

3. Create a file that stores the Manager details such as the URL, certificate location, and user.

   ```
   # cat engine_vars.yml
   ---
   engine_url: https://example.engine.redhat.com/ovirt-engine/api
   engine_user: admin@internal
   engine_cafile: /etc/pki/ovirt-engine/ca.pem
   ```

> **NOTE**
>
> If you prefer, these variables can be added directly to the playbook instead.

4. Create your playbook. To simplify this you can copy and modify an example in **/usr/share/doc/ovirt-ansible-roles/examples**.

```
# cat rhv_infra.yml
---
- name: RHV infrastructure
  hosts: localhost
  connection: local
  gather_facts: false

  vars_files:
    # Contains variables to connect to the Manager
    - engine_vars.yml
    # Contains encrypted engine_password variable using ansible-vault
    - passwords.yml

  pre_tasks:
    - name: Login to RHV
      ovirt_auth:
        url: "{{ engine_url }}"
        username: "{{ engine_user }}"
        password: "{{ engine_password }}"
        ca_file: "{{ engine_cafile | default(omit) }}"
        insecure: "{{ engine_insecure | default(true) }}"
      tags:
        - always

  vars:
    data_center_name: mydatacenter
    data_center_description: mydatacenter
    data_center_local: false
    compatibility_version: 4.1

  roles:
    - ovirt-datacenters

  post_tasks:
    - name: Logout from RHV
      ovirt_auth:
        state: absent
        ovirt_auth: "{{ ovirt_auth }}"
      tags:
        - always
```

5. Run the playbook.

```
# ansible-playbook --ask-vault-pass rhv_infra.yml
```

You have successfully used the **ovirt-datacenters** Ansible role to create a data center named **mydatacenter**.

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 VM 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

- IBM Security Directory Server

- IBM Security Directory Server RFC-2307 Schema

- FreeIPA

- 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 (**rhevm**) and IdM (**ipa-server**) on the same system. IdM is incompatible with the **mod_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 must know the domain name of the DNS or the LDAP server.

- To set up secure connection between the LDAP server and the Manager, ensure that a PEM-encoded CA certificate has been prepared.

- Have at least one set of account name and password ready to perform search and login queries to the LDAP server.

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. 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 .

   > Available LDAP implementations:
   > 1 - 389ds
   > 2 - 389ds RFC-2307 Schema
   > 3 - Active Directory
   > 4 - IBM Security Directory Server
   > 5 - IBM Security Directory Server RFC-2307 Schema
   > 6 - IPA
   > 7 - Novell eDirectory RFC-2307 Schema
   > 8 - OpenLDAP RFC-2307 Schema
   > 9 - OpenLDAP Standard Schema
   > 10 - Oracle Unified Directory RFC-2307 Schema
   > 11 - RFC-2307 Schema (Generic)
   > 12 - RHDS
   > 13 - RHDS RFC-2307 Schema
   > 14 - iPlanet
   > Please select:

4. 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]:

5. Select a DNS policy method:

   - For option 1, the DNS servers listed in **/etc/resolv.conf** are used to resolve the IP address. Check that the **/etc/resolv.conf** file is updated with the correct DNS servers.

   - For option 2, 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

     Example: **dig _ldap._tcp.redhat.com SRV**.

   - 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.

     > 1 - Single server
     > 2 - DNS domain LDAP SRV record
     > 3 - Round-robin between multiple hosts

> 4 - Failover between multiple hosts
> Please select:

6. Select the secure connection method your LDAP server supports and specify the method to obtain a PEM-encoded CA certificate:

   - **File** allows you to provide the full path to the certificate.

   - **URL** allows you to specify a URL for the certificate.

   - **Inline** allows you to paste the content of the certificate in the terminal.

   - **System** allows you to specify the default location for all CA files.

   - **Insecure** skips certificate validation, but the connection is still encrypted using TLS.

     > 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):
     > Please enter the password:

     > **NOTE**
     >
     > LDAPS stands for Lightweight Directory Access Protocol Over Secure Socket Links. For SSL connections, select the **ldaps** option.

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 specified in LDAP annotation. If anonymous search is allowed, press **Enter** without any input.

   > Enter search user DN (for example uid=username,dc=example,dc=com or leave empty for anonymous): uid=user1,ou=Users,ou=department-1,dc=example,dc=com
   > Enter search user password:

8. Enter the base DN:

   > Please enter base DN (dc=redhat,dc=com) [dc=redhat,dc=com]: *ou=department-1,dc=redhat,dc=com*

9. Select **Yes** if you intend to configure single sign-on for virtual machines. Note that the feature cannot be used with single sign-on to the Administration Portal feature. The script reminds you that the profile name must match the domain name. You will still need to follow the instructions in Configuring Single Sign-On for Virtual Machines in the *Virtual Machine Management Guide*.

   > Are you going to use Single Sign-On for Virtual Machines (Yes, No) [Yes]:

10. Specify a profile name. The profile name is visible to users on the login page. This example uses **redhat.com**.
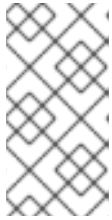
> **NOTE**
>
> To rename the profile after the domain has been configured, edit the **ovirt.engine.aaa.authn.profile.name** attribute in the **/etc/ovirt-engine/extensions.d/***redhat.com***-authn.properties** file. Restart the **ovirt-engine** service for the changes to take effect.

> Please specify profile name that will be visible to users: *redhat.com*

Figure 15.1. The Administration Portal Login Page
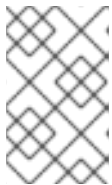
> **NOTE**
>
> Users must select the profile from the drop-down list when logging in for the first time. The information is stored in browser cookies and preselected the next time the user logs in.

11. Test the login function to ensure your LDAP server is connected to your Red Hat Virtualization environment properly. For the login query, enter your **user name** and **password**:

    > NOTE:
    > It is highly recommended to test drive the configuration before applying it into engine.
    > Login sequence is executed automatically, but it is recommended to also execute Search sequence manually after successful Login sequence.
    >
    > Please provide credentials to test login flow:
    > Enter user name:
    > Enter user password:
    > [INFO] Executing login sequence...
    > ...
    > [INFO] Login sequence executed successfully

12. Check that the user details are correct. If the user details are incorrect, select **Abort**:

    > Please make sure that user details are correct and group membership meets expectations

(search for PrincipalRecord and GroupRecord titles).
Abort if output is incorrect.
Select test sequence to execute (Done, Abort, Login, Search) [Abort]:

13. Manually testing the Search function is recommended. For the search query, select **Principal** for user accounts or **Group** for group accounts. Select **Yes** to **Resolve Groups** if you want the group account information for the user account to be returned. Three configuration files are created and displayed in the screen output.

    Select test sequence to execute (Done, Abort, Login, Search) [Search]: *Search*
    Select entity to search (Principal, Group) [Principal]:
    Term to search, trailing '*' is allowed: *testuser1*
    Resolve Groups (Yes, No) [No]:

14. Select **Done** to complete the setup:

    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.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-20171004101225-mmneib.log:
    [INFO] Stage: Pre-termination
    [INFO] Stage: Termination

15. Restart the **ovirt-engine** service. The profile you have created is now available on the Administration Portal and the VM Portal login pages. To assign the user accounts on the LDAP server appropriate roles and permissions, for example, to log in to the VM 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.

> **NOTE**
>
> Examples of the most common Active Directory configurations, which cannot be configured using the ovirt-engine-extension-aaa-ldap-setup tool, are provided in **/usr/share/ovirt-engine-extension-aaa-ldap/examples/README.md**.

- 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 Encrypted Communication 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.

- You must have at least one account name and password ready to perform search and login queries to the Active Directory.

- If your Active Directory deployment spans multiple domains, be aware of the limitation described in the **/usr/share/ovirt-engine-extension-aaa-ldap/profiles/ad.properties** file.

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. Select an LDAP type by entering the corresponding number. The LDAP-related questions after this step are different for different LDAP types.

   ```
   Available LDAP implementations:
    1 - 389ds
    2 - 389ds RFC-2307 Schema
    3 - Active Directory
    4 - IBM Security Directory Server
    5 - IBM Security Directory Server RFC-2307 Schema
    6 - IPA
    7 - Novell eDirectory RFC-2307 Schema
    8 - OpenLDAP RFC-2307 Schema
    9 - OpenLDAP Standard Schema
   10 - Oracle Unified Directory RFC-2307 Schema
   11 - RFC-2307 Schema (Generic)
   12 - RHDS
   13 - RHDS RFC-2307 Schema
   14 - iPlanet
   Please select: 3
   ```

4. 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*

5. 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 Encrypted Communication between the Manager and an LDAP Server".

6. 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):
> *cn=user1,ou=Users,dc=test,dc=redhat,dc=com*
> Enter search user password:

7. Specify whether to use single sign-on for virtual machines. This feature is enabled by default, but cannot be used if single sign-on to the Administration Portal is enabled. The script reminds you that the profile name must match the domain name. You will still need to follow the instructions in Configuring Single Sign-On for Virtual Machines in the *Virtual Machine Management Guide*.

> Are you going to use Single Sign-On for Virtual Machines (Yes, No) [Yes]:

8. 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.

9. 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.
Login sequence is executed automatically, but it is recommended to also execute Search
sequence manually after successful Login 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
```

10. The profile you have created is now available on the Administration Portal and the VM Portal login pages. To assign the user accounts on the LDAP server appropriate roles and permissions, for example, to log in to the VM Portal, see Section 15.6, "Administering User Tasks From the Administration Portal".
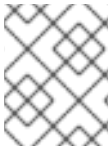
> **NOTE**
>
> For more information, see the LDAP authentication and authorization extension README file at **/usr/share/doc/ovirt-engine-extension-aaa-ldap-*version***.

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 VM 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.

Configuring an External LDAP Provider Manually

1. On the Red Hat Virtualization Manager, install the LDAP extension package:

   ```
   # yum install ovirt-engine-extension-aaa-ldap
   ```

2. 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
   ```

3. Rename the configuration files to match the profile name you want visible to users on the Administration Portal and the VM 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 Encrypted Communication 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
    ```

5. Review the authentication configuration file. The profile name visible to users on the

Administration Portal and the VM 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
```

6. 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
   ```

7. 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
   ```

8. Restart the engine service:

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

9. The *example* profile you have created is now available on the Administration Portal and the VM
   Portal login pages. To give the user accounts on the LDAP server appropriate permissions, for
   example, to log in to the VM 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.

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 VM 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 VM 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 VM Portal is enabled, single sign-on to virtual machines will not be possible. With single sign-on to the VM Portal enabled, the VM 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.

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
   ```

Configuring Single Sign-on to the VM 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>
```

6. Rename the configuration files to match the profile name you want visible to users on the Administration Portal and the VM 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
```

7. 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 Encrypted Communication 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
```

8. Review the authentication configuration file. The profile name visible to users on the Administration Portal and the VM 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
```

9. 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
ovirt.engine.extension.provides = org.ovirt.engine.api.extensions.aaa.Authz
config.profile.file.1 = ../aaa/example.properties
```

10. Review the authentication mapping configuration file. The configuration profile location must match the LDAP configuration file location. The configuration profile extension name must match the **ovirt.engine.aaa.authn.mapping.plugin** value in the authentication configuration file. All fields can be left as default.

```
# vi /etc/ovirt-engine/extensions.d/example-http-mapping.properties
```

Example 15.10. Example authentication mapping configuration file

```
ovirt.engine.extension.name = example-http-mapping
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.mapping.MappingExtension
ovirt.engine.extension.provides = org.ovirt.engine.api.extensions.aaa.Mapping
config.mapAuthRecord.type = regex
config.mapAuthRecord.regex.mustMatch = true
config.mapAuthRecord.regex.pattern = ^(?<user>.?)((\\\\(?<at>@)(?<suffix>.?)@.)|(?
<realm>@.))$
config.mapAuthRecord.regex.replacement = ${user}${at}${suffix}
```

11. Ensure that the ownership and permissions of the configuration files are appropriate:

```
# chown ovirt:ovirt /etc/ovirt-engine/aaa/example.properties
```

```
# chown ovirt:ovirt /etc/ovirt-engine/extensions.d/example-http-authn.properties
```

```
# chown ovirt:ovirt /etc/ovirt-engine/extensions.d/example-http-mapping.properties
```

```
# chown ovirt:ovirt /etc/ovirt-engine/extensions.d/example-authz.properties
```

```
# chmod 600 /etc/ovirt-engine/aaa/example.properties
```

```
# chmod 640 /etc/ovirt-engine/extensions.d/example-http-authn.properties
```

```
# chmod 640 /etc/ovirt-engine/extensions.d/example-http-mapping.properties
```

```
# 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 has a corresponding **permission**.

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 VM 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 VM Portal and to start creating virtual machines. The procedure also applies to group accounts.

Adding Users and Assigning VM Portal Permissions

1. On the header bar, click **Administration → Configure** to open the **Configure** window.

2. Click **System Permissions**.

3. Click **Add** to open the **Add System Permission to User** window.

4. 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.

5. Select the check boxes for the appropriate users or groups.

6. Select an appropriate role to assign under **Role to Assign**. The **UserRole** role gives the user account the permission to log in to the VM Portal.

7. Click **OK**.

Log in to the VM Portal to verify that the user account has the permissions to log in.

15.6.2. Viewing User Information

Viewing User Information

1. Click **Administration → Users** to display the list of authorized users.

2. Click the user's name to open the details view, usually with the **General** tab displaying general information, such as the domain name, email and status of the user.

3. 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.

Viewing User Permissions on Resources

1. Find and click the resource's name to open the details view.

2. Click the **Permissions** tab 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.

Removing Users

1. Click **Administration → Users** 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 **Remove**, then 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 **Administration** → **Active User Sessions** to view the **Session DB ID**, **User Name**, **Authorization provider**, **User id**, **Source IP**, **Session Start Time**, and **Session Last Active Time** for each logged-in user.

15.6.6. Terminating a User Session

You can terminate the session of a user who is currently logged in.

Terminating a User Session

1. Click **Administration** → **Active User Sessions**.

2. Select the user session to be terminated.

3. Click **Terminate Session**.

4. Click **OK**.

15.7. ADMINISTERING USER TASKS FROM THE COMMAND LINE

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.
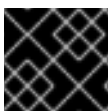
> **IMPORTANT**
>
> You must be logged into the Manager machine.

15.7.1. Creating a New User

You can create a new user account. The optional **--attribute** command specifies 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
```

You can 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 VM Portal Permissions" for more information.

15.7.2. Setting a User Password

You can create 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**.

15.7.3. Setting User Timeout

You can set the user timeout period:

```
# engine-config --set UserSessionTimeOutInterval=integer
```

15.7.4. Pre-encrypting a User Password

You can create a pre-encrypted user password using the **ovirt-engine-crypto-tool** script. This option is useful if you are adding users and passwords to the database with a script.

> **NOTE**
>
> Passwords are stored in the Manager database in encrypted form. The **ovirt-engine-crypto-tool** script is used because all passwords must be encrypted with the same algorithm.
>
> If the password is pre-encrypted, password validity tests cannot be performed. The password will be accepted even if it does not comply with the password validation policy.

1. Run the following command:

   ```
   # /usr/share/ovirt-engine/bin/ovirt-engine-crypto-tool.sh pbe-encode
   ```

   The script will prompt you to enter the password.

   Alternatively, you can use the **--password=file:*file*** option to encrypt a single password that appears as the first line of a file. This option is useful for automation. In the following example, **_file_** is a text file containing a single password for encryption:

   ```
   # /usr/share/ovirt-engine/bin/ovirt-engine-crypto-tool.sh pbe-encode --password=file:file
   ```

2. Set the new password with the **ovirt-aaa-jdbc-tool** script, using the **--encrypted** option:

```
# ovirt-aaa-jdbc-tool user password-reset test1 --password-valid-to="2025-08-01 12:00:00-0800" --encrypted
```

3. Enter and confirm the encrypted password:

```
Password:
Reenter password:
updating user test1...
user updated successfully
```

15.7.5. Viewing User Information

You can view detailed user account information:

```
# ovirt-aaa-jdbc-tool user show test1
```

This command displays more information than in the Administration Portal's **Administration → Users** screen.

15.7.6. Editing User Information

You can update user information, such as the email address:

```
# ovirt-aaa-jdbc-tool user edit test1 --attribute=email=jdoe@example.com
```

15.7.7. Removing a User

You can remove a user account:

```
# ovirt-aaa-jdbc-tool user delete test1
```

Remove the user from the Administration Portal. See Section 15.6.4, "Removing Users" for more information.

15.7.8. 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 envrionment with full administrative permissions before disabling the default **admin** user.

Disabling the Internal Administrative User

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.

2. Make sure another user with the **SuperUser** role has been added to the environment. See Section 15.6.1, "Adding Users and Assigning VM Portal Permissions" for more information.

3. Disable the default **admin** user:

```
# ovirt-aaa-jdbc-tool user edit admin --flag=+disabled
```
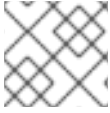
NOTE

To enable a disabled user, run **ovirt-aaa-jdbc-tool user edit** *username* **--flag=-disabled**

15.7.9. 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.

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.

1. Log in to the machine on which the Red Hat Virtualization Manager is installed.

2. Create a new group:

   # ovirt-aaa-jdbc-tool group add *group1*

3. Add users to the group. The users must be created already.

   # ovirt-aaa-jdbc-tool group-manage useradd *group1* --user=*test1*
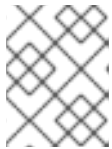
   NOTE

   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 VM Portal Permissions" for more information.

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 VM Portal Permissions" for more information.

15.7.10. 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**.

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 the account details.

   - All user account details:

     ```
     # ovirt-aaa-jdbc-tool query --what=user
     ```

   - All group account details:

     ```
     # ovirt-aaa-jdbc-tool query --what=group
     ```

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. Filter account details using the **--pattern** parameter.

   - 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"
     ```

15.7.11. Managing Account Settings

To change the default account settings, use the **ovirt-aaa-jdbc-tool settings** module.

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 settings 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 settings 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 settings 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 autonmatically 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 logs quota violations without blocking users and can be used to test quotas. In Audit mode, you can 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 verify that it works as you expect it to. You do not need to have your quota in **Audit** mode to create or change a quota.

Enabling and Changing Quota in a Data Center

1. Click **Compute → Data Centers** and select a data center.

2. Click **Edit**.

3. In the **Quota Mode** drop-down list, change the quota mode to **Enforced**.

4. Click **OK**.

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.

Creating a New Quota Policy

1. Click **Administration → Quota**.

2. Click **Add**.

3. Fill in the **Name** and **Description** fields.

4. Select a **Data Center**.

5. In the **Memory & CPU** section, use the green slider to set **Cluster Threshold**.

6. In the **Memory & CPU** section, use the blue slider to set **Cluster Grace**.

7. Click 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.

   a. 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.

b. 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.

c. Click **OK** in the **Edit Quota** window.

9. In the **Storage** section, use the green slider to set **Storage Threshold**.

10. In the **Storage** section, use the blue slider to set **Storage Grace**.

11. Click 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.

12. Click **Edit** to open the **Edit Quota** window.

a. 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.

b. Click **OK** in the **Edit Quota** window.

13. Click **OK** in the **New Quota** window.

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

Assigning a Quota to a Virtual Machine

1. Click **Compute → Virtual Machines** and select a virtual machine.

2. Click **Edit**.

3. Select the quota you want the virtual machine to consume from the **Quota** drop-down list.

4. Click **OK**.

Assigning a Quota to a Virtual Disk

1. Click **Compute → Virtual Machines**.

2. Click a virtual machine's name to open the details view.

3. Click the **Disks** tab and select the disk you plan to associate with a quota.

4. Click **Edit**.

5. Select the quota you want the virtual disk to consume from the **Quota** drop-down list.

6. Click **OK**.
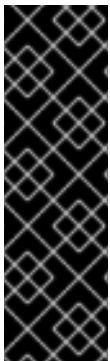
> **IMPORTANT**
>
> Quota must be selected for all objects associated with a virtual machine, in order for that virtual machine to work. If you fail to select a quota for the objects associated with a virtual machine, the virtual machine will not work. The error that the Manager throws in this situation 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

This procedure describes how to use quotas to limit the resources a user has access to.

Assigning a User to a Quota

1. Click **Administration → Quota**.

2. Click the name of the target quota to open the details view.

3. Click the **Consumers** tab.

4. Click **Add**.

5. In the **Search** field, type the name of the user you want to associate with the quota.

6. Click **GO**.

7. Select the check box next to the user's name.

8. Click **OK**.

After a short time, the user will appear in the **Consumers** tab in the details view.

16.9. EDITING QUOTAS

This procedure describes how to change existing quotas.

Editing Quotas

1. Click **Administration → Quota** and select a quota.

2. Click **Edit**.

3. Edit the fields as required.

4. Click **OK**.

16.10. REMOVING QUOTAS

This procedure describes how to remove quotas.

Removing Quotas

1. Click **Administration → Quota** and select a quota.

2. Click **Remove**.

3. Click **OK**.

16.11. SERVICE LEVEL AGREEMENT POLICY ENFORCEMENT

This procedure describes how to set service level agreement CPU features.

Setting a Service Level Agreement CPU Policy

1. Click **Compute → Virtual Machines**.

2. Click **New**, or select a virtual machine and click **Edit**.

3. Click the **Resource Allocation** tab.

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.

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

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.

Configuring Event Notifications

1. Ensure that you have access to an email server that can accept automated messages from RHVM and deliver them to a distribution list.

2. Click **Administration → Users** and select a user.

3. Click the user's **User Name** to go to the details page.

4. In the **Event Notifier** tab, click **Manage Events**.

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.
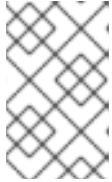
> **NOTE**
>
> The email address can be a text message email address (for example, ***1234567890@carrierdomainname.com***) or an email group address that includes email addresses and text message email addresses.

8. Click **OK**.

9. On the Manager machine, copy **ovirt-engine-notifier.conf** to a new file called **90-email-notify.conf**:

   ```
   # cp /usr/share/ovirt-engine/services/ovirt-engine-notifier/ovirt-engine-notifier.conf /etc/ovirt-engine/notifier/notifier.conf.d/90-email-notify.conf
   ```

10. Edit **90-email-notify.conf**, deleting everything except the **EMAIL Notifications** section.

11. Enter the correct email variables, as in the example below. This file overrides the values in the original **ovirt-engine-notifier.conf** file.

    ```
    --------------------
    # EMAIL Notifications #
    --------------------

    # The SMTP mail server address. Required.
    MAIL_SERVER=myemailserver.example.com

    # The SMTP port (usually 25 for plain SMTP, 465 for SMTP with SSL, 587 for SMTP with
    ```

```
TLS)
MAIL_PORT=25

# Required if SSL or TLS enabled to authenticate the user. Used also to specify 'from' user
address if mail server
# supports, when MAIL_FROM is not set. Address is in RFC822 format
MAIL_USER=

# Required to authenticate the user if mail server requires authentication or if SSL or TLS is
enabled
SENSITIVE_KEYS="${SENSITIVE_KEYS},MAIL_PASSWORD"
MAIL_PASSWORD=

# Indicates type of encryption (none, ssl or tls) should be used to communicate with mail
server.
MAIL_SMTP_ENCRYPTION=none

# If set to true, sends a message in HTML format.
HTML_MESSAGE_FORMAT=false

# Specifies 'from' address on sent mail in RFC822 format, if supported by mail server.
MAIL_FROM=rhevm2017@example.com

# Specifies 'reply-to' address on sent mail in RFC822 format.
MAIL_REPLY_TO=

# Interval to send smtp messages per # of IDLE_INTERVAL
MAIL_SEND_INTERVAL=1

# Amount of times to attempt sending an email before failing.
MAIL_RETRIES=4
```
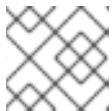
> **NOTE**
>
> See **/etc/ovirt-engine/notifier/notifier.conf.d/README** for more options.

12. Enable and restart the **ovirt-engine-notifier** service to activate the changes you have made:

    ```
    # systemctl daemon-reload
    # systemctl enable ovirt-engine-notifier.service
    # systemctl restart ovirt-engine-notifier.service
    ```

The specified user now receives emails based on events in the Red Hat Virtualization environment. The
selected events are displayed on the **Event Notifier** tab for that user.

17.2. CANCELING EVENT NOTIFICATIONS IN THE ADMINISTRATION PORTAL

A user has configured some unnecessary email notifications and wants them canceled.

Canceling Event Notifications

1. Click **Administration → Users**.

2. Click the user's **User Name** to open the details view.

3. Click the **Event Notifier** tab to list events for which the user receives email notifications.

4. Click **Manage Events**.

5. Use the **Expand All** button, or the subject-specific expansion buttons, to view the events.

6. Clear the appropriate check boxes to remove notification for that event.

7. Click **OK**.

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.
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.

Variable Name	Default	Remarks
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.
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.

Variable Name	Default	Remarks
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**.
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.

Variable Name	Default	Remarks
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.
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.

Variable Name	Default	Remarks
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 after changes such as system upgrades.

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 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/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 Manager 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 existing answer file to automate the rename operation.

--generate-answer=*[file]*

Allows you to specify the path and file name of the file in which your answers and the values changed by the **ovirt-engine-rename** command are recorded.

18.1.3. Renaming the Manager with the oVirt Engine Rename Tool

You can use the **ovirt-engine-rename** command to update records of the fully qualified domain name of the Manager.

> **IMPORTANT**
>
> The **ovirt-engine-rename** command does not update SSL certificates, such as **imageio-proxy** or **websocket-proxy**. These must be updated manually, after running **ovirt-engine-rename**. See Updating SSL Certificates below.

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.

Using the oVirt Engine Rename Tool

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]:
   ```

6. When prompted, enter the new fully qualified domain name for the Manager:

   ```
   New fully qualified server name:_new-name_
   ```

The **ovirt-engine-rename** command updates records of the fully qualified domain name of the Manager.

Updating SSL Certificates

Run the following commands after the **ovirt-engine-rename** command to update the SSL certificates:

```
1. # names="websocket-proxy imageio-proxy"
```

```
2. # subject="$(\
     openssl x509 \
     -in /etc/pki/ovirt-engine/certs/apache.cer \
     -noout \
     -subject | \
       sed \
         's;subject= \(.*\);\1;'
   )"
```

```
3. # . /usr/share/ovirt-engine/bin/engine-prolog.sh
```

```
4. # for name in $names; do
     /usr/share/ovirt-engine/bin/pki-enroll-pkcs12.sh \
       --name="${name}" \
       --password=mypass \
       --subject="${subject}" \
       --keep-key \
       --san=DNS:"${ENGINE_FQDN}"
   done
```

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 USB FILTER EDITOR

18.3.1. 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 **USB Filter Editor** installer from "Installers and Images for Red Hat Virtualization Manager".

Installing the USB Filter Editor

1. On a Windows machine, run the **.msi** file you downloaded for the **USB Filter Editor**.

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.

18.3.2. The USB Filter Editor Interface

Double-click the USB Filter Editor shortcut icon on your desktop.

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.

18.3.3. Adding a USB Policy

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

Adding a USB Policy

1. Click **Add**.

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 18.1. 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.
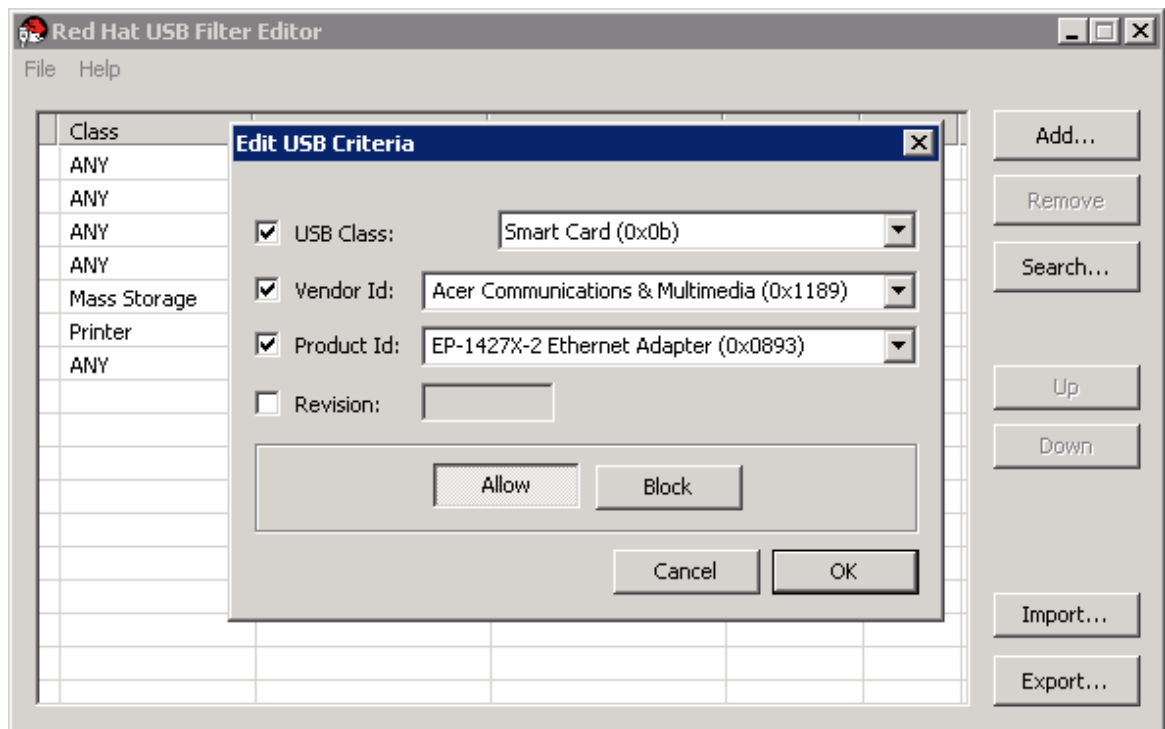
3. Click **File → Save** to save the changes.

You have added a USB policy to the USB Filter Editor. USB filter policies must be exported to the Red Hat Virtualization Manager to take effect.

18.3.4. Removing a USB Policy

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

Removing a USB Policy

1. Select the policy to be removed.

2. Click **Remove**. A message displays prompting you to confirm that you want to remove the policy.

3. Click **Yes** to confirm that you want to remove the policy.

4. Click **File → Save** to save the changes.

You have removed a USB policy from the USB Filter Editor. USB filter policies must be exported to the Red Hat Virtualization Manager to take effect.

18.3.5. Searching for USB Device Policies

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.

Searching for USB Device Policies

1. Click **Search**. The **Attached USB Devices** window displays a list of all the attached 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.

You have searched the attached USB devices. USB filter policies need to be exported to the Red Hat Virtualization Manager to take effect.

18.3.6. Exporting a USB Policy

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.

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
   ```

18.3.7. Importing a USB Policy

An existing USB device policy must be downloaded and imported into the USB Filter Editor before you can edit it.

Importing a USB Policy

1. Using a Secure Copy client, such as WinSCP, download the **usbfilter.txt** file from the server running Red Hat Virtualization Manager. The file can be found 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.

18.4. THE LOG COLLECTOR TOOL

18.4.1. 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.

18.4.2. 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.

--time-only

Displays only information about time differences between hosts, without generating a full SOS report.

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-dbname=*DBNAME*

Sets *DBNAME* as the database name to use for connections with the database server. The default is **rhevm**.

--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.

18.4.3. 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 18.2. 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

18.5. THE ISO UPLOADER TOOL

18.5.1. The ISO Uploader Tool

> **NOTE**
>
> The ISO Uploader tool has been deprecated. Red Hat recommends uploading ISO images to the data domain with the Administration Portal or with the REST API. See Section 8.8.1, "Uploading Images to a Data Storage Domain" for details.

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.

18.5.2. 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 to 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 used 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.

18.5.3. Specifying an NFS Server

> **Example 18.3. Uploading to an NFS Server**
>
> > # engine-iso-uploader --nfs-server=*storage.demo.redhat.com:/iso/path* upload *RHEL6.0.iso*

18.5.4. 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 18.4. 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):
> > ```

18.5.5. Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain

The **virtio-win** ISO and Virtual Floppy Drive (VFD) images, which contain the VirtIO drivers for Windows virtual machines, and the **rhv-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 **rhv-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/rhv-guest-tools-iso/rhv-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.

The example below demonstrates the command to upload the **virtio-win.iso**, **virtio-win_x86.vfd**, **virtio-win_amd64.vfd**, and **rhv-tools-setup.iso** image files to the **ISODomain**.

> **Example 18.5. 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/rhv-guest-tools-iso/rhv-tools-setup.iso

18.6. THE ENGINE VACUUM TOOL

18.6.1. 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 must 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/9.5/static/sql-vacuum.html
> (Yes, No) [No]:

The **Yes** option runs the Engine Vacuum tool in full vacuum verbose mode.

18.6.2. Engine Vacuum Modes

Engine Vacuum has two modes:

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.

18.6.3. Syntax for the engine-vacuum Command

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
```

18.7. THE VDSM TO NETWORK NAME MAPPING TOOL

18.7.1. Mapping VDSM Names to Logical Network Names

If the name of a logical network is longer than 15 characters or contains non-ASCII characters, the system automatically generates an on-host identifier (*vdsm_name*) name; it comprises the letters *on* and the first 13 characters of the network's unique identifier, for example, **ona1b2c3d4e5f6g**. It is this name that appears in the host's log files. To view a list of logical network names and their auto-generated network name, use the VDSM-to-Network-Name Mapping tool located in **/usr/share/ovirt-engine/bin/**.

Procedure

1. The first time you run the tool, define a PASSWORD environment variable, which is the password of a database user with read access to the Manager database. For example, run:

   ```
   # export PASSWORD=DatabaseUserPassword
   ```

2. Run the VDSM-to-Network-Name Mapping tool:

   ```
   # vdsm_to_network_name_map  --user USER
   ```

   where *USER* is the database user with read access to the Manager database, whose password is assigned to the PASSWORD environment variable.

The tool displays a list of logical network names that are mapped to their equivalent on-host identifiers.

Additional Flags

You can run the tool with the following flags:

--host is the hostname/IP address of the database server. The default value is **localhost**.

--port is the port number of the database server. The default value is **5432**. **--database** is the name of the database. The default value is **engine**, which is the Manager database.

--secure enables a secure connection with the database. By default the tool is run without a secure connection.

PART IV. GATHERING INFORMATION ABOUT THE ENVIRONMENT

CHAPTER 19. LOG FILES

19.1. MANAGER INSTALLATION LOG FILES

Table 19.1. Installation

Log File	Description
/var/log/ovirt-engine/engine-cleanup_*yyyy_mm_dd_hh_mm_ss*.log	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.
/var/log/ovirt-engine/engine-db-install-*yyyy_mm_dd_hh_mm_ss*.log	Log from the **engine-setup** command detailing the creation and configuration of the **engine** database.
/var/log/ovirt-engine/ovirt-engine-dwh-setup-*yyyy_mm_dd_hh_mm_ss*.log	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.
/var/log/ovirt-engine/setup/ovirt-engine-setup-*yyyymmddhhmmss*.log	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
/var/log/ovirt-engine/engine.log	Reflects all Red Hat Virtualization Manager GUI crashes, Active Directory lookups, Database issues, and other events.
/var/log/ovirt-engine/host-deploy	Log files from hosts deployed from the Red Hat Virtualization Manager.
/var/lib/ovirt-engine/setup-history.txt	Tracks the installation and upgrade of packages associated with the Red Hat Virtualization Manager.

Log File	Description
/var/log/httpd/ovirt-requests-log	Logs files from requests made to the Red Hat Virtualization Manager via HTTPS, including how long each request took. A **Correlation-Id** header is included to allow you to compare requests when comparing a log file with **/var/log/ovirt-engine/engine.log**.
/var/log/ovn-provider/ovirt-provider-ovn.log	Logs the activities of the OVN provider. For information about Open vSwitch logs, see the Open vSwitch documentation.

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
Red Hat Enterprise Linux Guest	Use **journalctl** as the root user.	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"** To run **spice-vdagent** in debug mode, from the command line: $ killall - u $USER spice-vdagent $ spice-vdagent -x -d [-d] [|& tee spice-vdagent.log]

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 later, **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. HOST LOG FILES

Log File	Description
/var/log/messages	The log file used by **libvirt**. Use **journalctl** to view the log. You must be a member of the *adm*, *systemd-journal*, or *wheel* groups to view the log.
/var/log/vdsm/spm-lock.log	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.
/var/log/vdsm/vdsm.log	Log file for VDSM, the Manager's agent on the host(s).
/tmp/ovirt-host-deploy-*Date*.log	A host deployment log that is copied to the Manager as **/var/log/ovirt-engine/host-deploy/ovirt-*Date-Host-Correlation_ID***.log after the host has been successfully deployed.
/var/log/vdsm/import/import-*UUID-Date*.log	Log file detailing virtual machine imports from a KVM host, a VMWare provider, or a RHEL 5 Xen host, including import failure information. *UUID* is the UUID of the virtual machine that was imported and *Date* is the date and time that the import began.
/var/log/vdsm/supervdsm.log	Logs VDSM tasks that were executed with superuser permissions.
/var/log/vdsm/upgrade.log	VDSM uses this log file during host upgrades to log configuration changes.
/var/log/vdsm/mom.log	Logs the activities of the VDSM's memory overcommitment manager.

19.5. SETTING UP A 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.

Setting up a 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 the firewall to allow proxy traffic. 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 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.

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 a SPICE Proxy

This procedure explains how to activate (or turn on) the SPICE proxy.

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 later, support HTTPS proxies. Earlier clients only support HTTP. If HTTPS is specified for earlier 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.

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

This section explains how to install and configure a Squid proxy to the VM Portal. A Squid proxy server is used as a content accelerator. It caches frequently-viewed content, reducing bandwidth and improving response times.

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
   ```
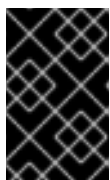
   > **IMPORTANT**
   >
   > You must include the quotes around the distinguished name for the certificate. The **-nodes** option ensures that the private key is not encrypted; this means that you do not need to enter the password to start the proxy server.

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.

4. 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/.

5. 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.

6. 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 .

7. Ensure both **proxy.key** and **proxy.cer** are present on the proxy machine:

   # ls -l proxy.key proxy.cer

8. Install the Squid proxy server package on the proxy machine:

   # yum install squid

9. 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/.

10. 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
```

> **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 a noVNC console. 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" .

20.3.2. Migrating the Websocket Proxy to a Separate Machine
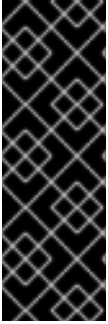
> **IMPORTANT**
>
> The websocket proxy and noVNC are Technology Preview features only. Technology Preview features are not supported with Red Hat production service-level agreements (SLAs) and might not be functionally complete, and Red Hat does not recommend using them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information see Red Hat Technology Preview Features Support Scope.

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.

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.

Name	Description
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.
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.

Name	Description
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.
before_device_destroy	Before destroying a device that supports custom properties.
after_device_destroy	After destroying a device that supports custom properties.

Name	Description
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 **vdsm** 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_*
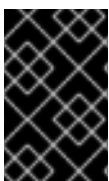
> **IMPORTANT**
>
> The **before_migration_destination** and **before_dehibernation** hooks currently receive the XML of the domain from the source host. The XML of the domain at the destination will have various differences.

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 theinterface **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
tobool	string	Converts a string "true" or "false" to a Boolean value
read_domxml	-	Reads the virtual machine's libvirt XML into a DOM object

Name	Argument	Description
write_domxml	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.
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.

Parameter	Description
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.
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 multicast_query_interval. If a multicast query for this membership has been sent within the last multicast_query_interval, it is not sent again.

Parameter	Description
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 multicast_query_interval.
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.

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 vdsm-hook-ethtool-options
```

The **ethtool_opts** key is now available in the Administration Portal. See Section 6.4.2, "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* .

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 6.4.2, "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:

- Plug-in discovery.

- Plug-in loading.

- 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:

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)	**/usr/share/ovirt-engine/ui-plugins/my-plugin.json**	
Plug-in user configuration files	**/etc/ovirt-engine/ui-plugins/my-plugin-config.json**	
Plug-in resource files	**/usr/share/ovirt-engine/ui-plugins/<resourcePath>/PluginHostPage.html**	**<resourcePath>** 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.

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 theHello World! Plug-in
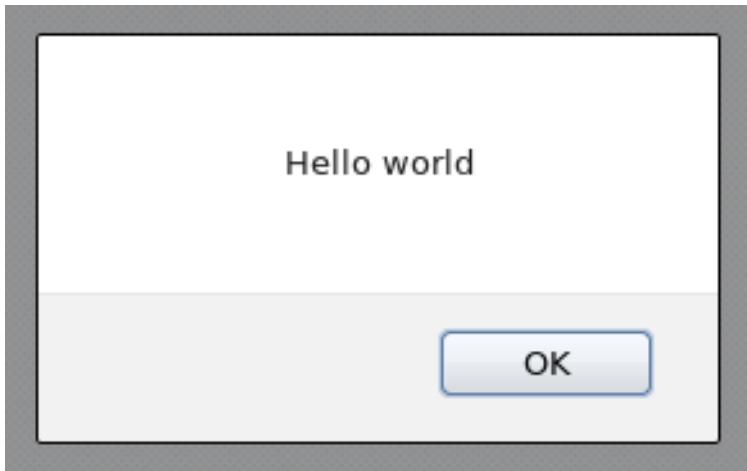
APPENDIX D. RED HAT VIRTUALIZATION AND ENCRYPTED COMMUNICATION

D.1. REPLACING THE RED HAT VIRTUALIZATION MANAGER CA 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**.

You can configure 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 .

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.

Replacing the Red Hat Virtualization Manager Apache CA 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, **/etc/ovirt-engine/engine.conf.d/99-custom-truststore.conf**, with the following parameters:

   ```
   ENGINE_HTTPS_PKI_TRUST_STORE="/etc/pki/java/cacerts"
   ENGINE_HTTPS_PKI_TRUST_STORE_PASSWORD=""
   ```

9. Edit the **/etc/ovirt-engine/ovirt-websocket-proxy.conf.d/10-setup.conf** file, adding the following parameters:

   ```
   SSL_CERTIFICATE=/etc/pki/ovirt-engine/apache.cer
   SSL_KEY=/etc/pki/ovirt-engine/keys/apache.key.nopass
   ```

10. Restart the **ovirt-provider-ovn** service:

    ```
    # systemctl restart ovirt-provider-ovn.service
    ```

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

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

Your users can now connect to the Administration Portal and VM Portal, without seeing a warning about the authenticity of the certificate used to encrypt HTTPS traffic.

D.2. SETTING UP ENCRYPTED COMMUNICATION BETWEEN THE MANAGER AND AN LDAP SERVER

To set up encrypted communication between the Red Hat 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.
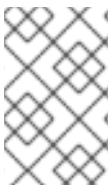
> **NOTE**
>
> For more information on creating a PEM-encoded CA certificate and importing certificates, see the **X.509 CERTIFICATE TRUST STORE** section of the README file at **/usr/share/doc/ovirt-engine-extension-aaa-ldap-***version*.

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]. To continue configuring LDAP and Kerberos for Single Sign-on, see xref:Configuring_LDAP_and_Kerberos_for_Single_Sign-on[.

D.3. MANUALLY SETTING UP ENCRYPTED COMMUNICATION FOR VDSM

You can manually set up encrypted communication for VDSM with the Manager and with other VDSM instances.

Only hosts in clusters with cluster level 3.6, 4.0, and 4.1 require manual configuration. Hosts in clusters with level 4.2 are automatically reconfigured for strong encryption during host reinstallation.

> **NOTE**
>
> RHVH 3.6, 4.0, and 4.1 hosts do not support strong encryption. RHVH 4.2 and RHEL hosts do support it.
>
> If you have 3.6, 4.0, or 4.1 clusters with RHVH 4.2 hosts, you can use strong encryption.

Procedure

1. Click **Compute → Hosts** and select the host.

2. Click **Management → Maintenance** to open the **Maintenance Host(s)** confirmation window.

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

4. On the host, create **/etc/vdsm/vdsm.conf.d/99-custom-ciphers.conf** with the following setting:

    ```
    [vars]
    ssl_ciphers = HIGH
    ```

■

See OpenSSL Cipher Strings for more information.

5. Restart VDSM:

```
# systemctl restart vdsm
```

6. Click **Compute → Hosts** and select the host.

7. Click **Management → Activate** to reactivate the host.

APPENDIX E. BRANDING

E.1. BRANDING

E.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.

E.1.2. Login Screen

The login screen is the login screen used by both the Administration Portal and VM 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**.

E.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**.

E.1.4. VM Portal Screen

The VM Portal screen is the screen that is shown when you log into the VM Portal. The elements of the VM 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 VM Portal screen are located in **user_portal.css**.

E.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**.

E.1.6. Tabs

Many pop-up windows in the Administration Portal 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**.

E.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.

E.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 F. SYSTEM ACCOUNTS

F.1. SYSTEM ACCOUNTS

F.1.1. Red Hat Virtualization Manager User Accounts

A number of system user accounts are created to support Red Hat Virtualization when the **rhevm** package is installed. Each system user has a default user identifier (UID). The system user accounts created are:

- The **vdsm** 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.

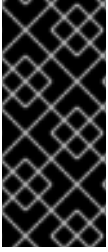### F.1.2. Red Hat Virtualization Manager Groups

A number of system user groups are created to support Red Hat Virtualization when the **rhevm** 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 **vdsm** 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.

F.1.3. Virtualization Host User Accounts

A number of system user accounts are created on the virtualization host when the **vdsm** and **qemu-kvm-rhev** packages are installed. Each system user has a default user identifier (UID). The system user accounts created are:

- The **vdsm** 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 **vdsm** 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 **vdsm** and **qemu-kvm-rhev** packages.

F.1.4. Virtualization Host Groups

A number of system user groups are created on the virtualization host when the **vdsm** 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 **vdsm** 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 **vdsm** 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 **vdsm** and **qemu-kvm-rhev** packages.