

Red Hat Enterprise Virtualization 3.0 Administration Guide

**System Administration of Red Hat Enterprise Virtualization
Environments using the Administration Portal**



Red Hat Documentation Team

Kate Grainger

Cheryn Tan

Tim Hildred

Stephen Gordon

Zac Dover

Daniel Macpherson

Shikha Nansi

Susan Burgess

David Jorm

Red Hat Enterprise Virtualization 3.0 Administration Guide

System Administration of Red Hat Enterprise Virtualization Environments using the Administration Portal

Edition 2

Author	Red Hat Documentation Team
Author	Kate Grainger
Author	Cheryn Tan
Author	Tim Hildred
Author	Stephen Gordon
Author	Zac Dover
Author	Daniel Macpherson
Author	Shikha Nansi
Author	Susan Burgess
Author	David Jorm

Copyright © 2011 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

1801 Varsity Drive
Raleigh, NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701

	xi
1. About this Guide	xi
1.1. Documentation Suite	xi
1.2. Audience	xii
2. Document Conventions	xii
2.1. Typographic Conventions	xii
2.2. Pull-quote Conventions	xiii
2.3. Notes and Warnings	xiv
3. We Need Feedback!	xv
Introduction	xvii
1. Red Hat Enterprise Virtualization Architecture	xvii
1.1. System Components	xviii
2. Red Hat Enterprise Virtualization Resources	xix
3. Administration of the Red Hat Enterprise Virtualization Platform	xx
3.1. Maintaining the Red Hat Enterprise Virtualization Platform	xxi
I. The Red Hat Enterprise Virtualization Interface	1
1. Getting Started	3
1.1. Graphical User Interface	3
1.1.1. Tree Mode and Flat Mode	4
1.1.2. Using the Guide Me facility	6
1.2. Search	6
1.2.1. Search Syntax	7
1.2.2. Saving and Accessing Queries as Bookmarks	10
1.3. Tags	11
II. Managing System Components	13
2. Managing Data Centers	15
2.1. Data Centers	15
2.1.1. Data Center Properties	16
2.1.2. Data Centers Operations	17
2.1.3. Data Centers Related Entities	18
2.1.4. Data Centers Troubleshooting	20
2.1.5. System Permissions	23
2.2. Clusters	25
2.2.1. Cluster Properties	26
2.2.2. Cluster Operations	27
2.2.3. Cluster Related Entities	29
2.2.4. Cluster Troubleshooting	31
2.2.5. Cluster Permissions	34
2.3. Logical Networks	35
2.3.1. Logical Networks Properties	36
2.3.2. Logical Network Operations	36
2.3.3. Logical Networks Troubleshooting	39
2.3.4. Logical Networks System Permissions	39
3. Storage	43
3.1. Managing Storage	43
3.1.1. Storage Properties	44
3.1.2. Storage Operations	44
3.1.3. Storage Entities	62
3.1.4. Storage Permissions	63

3.1.5. Storage Troubleshooting	65
4. Red Hat Enterprise Virtualization Hosts	73
4.1. Managing Hosts	73
4.1.1. Hosts Properties	74
4.1.2. Host Operations	75
4.1.3. Hosts Entities	89
4.1.4. Host Troubleshooting	91
4.1.5. Hosts Permissions	94
5. Users	97
5.1. Authorization Model	97
5.2. User Properties	98
5.2.1. Roles	99
5.2.2. Permissions	101
5.3. Users Operations	101
5.3.1. Adding Users and Groups	102
5.4. Users Troubleshooting	103
5.4.1. Managing Event Notifiers	103
5.4.2. Removing Users	104
III. Managing Virtualization Infrastructure	107
6. Managing Virtual Resources	109
6.1. About Virtual Machines	109
6.1.1. Supported Virtual Machines	110
6.1.2. Virtual Machine Performance Parameters	110
6.1.3. Understanding Virtual Machine Storage	111
6.2. Creating New Virtual Machines	113
6.2.1. Creating Virtual Machines from Existing Templates	113
6.2.2. Creating New Virtual Machines without a Template	115
6.2.3. Cloning Virtual Machines from Existing Templates	118
6.3. Completing the Configuration of the Virtual Machine	119
6.4. Installing Operating Systems onto Blank Virtual Machines	122
6.5. Managing Permissions to Virtual Machines	125
6.5.1. Assigning Virtual Machines to Users	125
6.5.2. Assigning System Permissions to Virtual Machines	127
6.6. Logging in to Virtual Machines	128
6.6.1. Logging into Virtual Machines using SPICE	128
6.6.2. Logging in to Windows Virtual Machines with Remote Desktop (RDP)	129
6.6.3. Logging in to Virtual Machines with VNC	129
6.6.4. Console Window Menu Extension for Administrators	130
6.7. Managing Virtual Machines	131
6.7.1. Editing Virtual Machines	132
6.7.2. Powering Virtual Machines On	132
6.7.3. Shutting Down or Pausing Virtual Machines	133
6.7.4. Migrating Virtual Machines	134
6.7.5. Moving Virtual Machines within a Data Center	134
6.7.6. Removing Virtual Machines	135
6.8. Using Virtual Machine Snapshots	135
6.8.1. Creating Snapshots of Virtual Machines	136
6.8.2. Restoring Virtual Machines from Snapshots	136
6.8.3. Deleting Snapshots	138
6.9. Exporting and Importing Virtual Resources	138
6.9.1. Overview of the Export-Import Process	139

6.9.2. Exporting Virtual Machines	141
6.9.3. Importing Virtual Machines into the Destination Data Center	142
6.10. Backing Up Virtual Resources	143
7. Templates	145
7.1. Creating Templates from Existing Virtual Machines	145
7.1.1. Sealing a Windows Template with Sysprep	146
7.1.2. Sealing a Linux Template with sys-unconfig	147
7.2. Managing Templates	148
7.2.1. Editing Templates	148
7.2.2. Copying Templates to a Different Storage Domain	148
7.2.3. Deleting Templates	149
7.3. Exporting and Importing Templates	149
7.3.1. Exporting Templates	150
7.3.2. Importing Templates	150
7.3.3. Backing Up Templates	152
7.4. Managing Permissions to Templates	152
7.4.1. Assigning Templates to Users	152
7.4.2. Assigning System Permissions for Templates Usage	153
8. Pools	155
8.1. Creating Desktop Pools	155
8.2. Managing Desktop Pools	161
8.2.1. Assigning Users to a Desktop Pool	161
8.2.2. Editing Desktop Pools	163
8.2.3. Detaching Desktops from a Pool	164
8.3. Managing Permissions to Pools	165
8.4. Removing Desktop Pools	167
IV. Monitoring and Reporting	169
9. Monitoring Red Hat Enterprise Virtualization	171
9.1. Using the Monitoring Tools	171
9.1.1. Monitoring Storage	172
9.1.2. Monitoring Hosts	172
9.1.3. Monitoring Virtual Machines	172
9.1.4. Monitoring High Severity Events	174
9.1.5. Viewing the Event List	174
9.1.6. Viewing Alert Information	175
10. Red Hat Enterprise Virtualization Reports	177
10.1. Overview	177
10.1.1. JasperReports & JasperServer	177
10.1.2. Online Help	177
10.1.3. Red Hat Enterprise Virtualization History Database	177
10.2. System Requirements	178
10.3. Accessing Reports and Dashboards	178
10.3.1. Logging in	178
10.3.2. Navigating Reports	180
10.3.3. Managing Users	181
10.4. Running Reports	182
10.4.1. Report Parameters	182
10.4.2. Executive Reports	184
10.4.3. Inventory Reports	186
10.4.4. Service Level Reports	188

10.4.5. Trend Reports	193
10.5. Dashboards	201
10.5.1. Data Center Inventory Dashboard	201
10.5.2. Data Center Trends Dashboard	202
10.5.3. Data Center Uptime Dashboard	203
10.5.4. System Overview Dashboard	203
10.6. Ad Hoc Reports	204
11. History Database Reports	205
11.1. Overview	205
11.1.1. Tracking Configuration History	205
11.1.2. Recording Statistical History	206
11.1.3. Tracking Tag History	206
11.2. Connecting to the History Database	206
11.3. Example Reports	207
11.3.1. Resource Utilization on a Single Host	207
11.3.2. Resource Utilization Across All Hosts	208
11.3.3. Tag Filter of Latest VM Configuration	209
11.3.4. List Current Virtual Machines' Names, Types, and Operating Systems	210
V. Managing Advanced Functionality	213
12. Live Migration	215
12.1. What is Live Migration?	215
12.2. Live Migration Prerequisites	215
12.3. Automatic Virtual Machine Migration	215
12.3.1. Moving a Host to Maintenance Mode	216
12.3.2. Cluster Policy	216
12.3.3. Preventing Automatic Migration of a Virtual Machine	218
12.4. Manually Migrating Virtual Machines	219
12.5. Setting Migration Priority	219
13. High Availability	221
13.1. What is High Availability?	221
13.1.1. Why Use High Availability?	221
13.2. High Availability Considerations	221
13.3. Host High Availability	222
13.3.1. Setting the Parameters for Fencing	222
13.3.2. Using Power Management Functions on a Fenced Host	224
13.3.3. Manually Fencing or Isolating a Host	225
13.4. Virtual Machine High Availability	226
13.4.1. Configuring a Highly Available Virtual Machine	227
13.4.2. Setting a Cluster Resilience Policy	227
14. Managing Multilevel Administration	229
14.1. Configuring Roles	229
14.1.1. Roles	229
14.1.2. Creating Custom Roles	230
14.1.3. Editing Roles	230
14.1.4. Cloning Roles	231
14.2. User Roles Examples	232
14.2.1. Setting Up an End User	232
14.2.2. Setting Up a Virtual Machine Administrator	233
14.2.3. Setting Up a Power User	233
14.3. Authorization Examples	234

15. Backing Up and Restoring the Red Hat Enterprise Virtualization Manager.	237
15.1. Backup and Restore the rhevm Postgres Database	237
15.1.1. Backing up Databases in Red Hat Enterprise Virtualization	237
15.1.2. Restoring Databases in Red Hat Enterprise Virtualization	238
15.2. Backing up and Restoring Manager Configuration Files	238
15.2.1. Red Hat Enterprise Virtualization Manager Configuration Files Requiring Backup	238
15.2.2. Restoring Red Hat Enterprise Virtualization Manager Configuration Files	240
16. Extending VDSM with Hooks	243
16.1. Environment	244
16.1.1. Domain XML	244
16.1.2. Custom Properties	245
16.1.3. Hooking module	248
16.2. Execution	248
16.3. Examples	249
A. Utilities	253
A.1. Domain Management Tool	253
A.1.1. Syntax	253
A.1.2. Examples	254
A.2. Configuration Tool	255
A.2.1. Syntax	256
A.2.2. Examples	257
A.3. Log Collector	257
A.3.1. Syntax	257
A.3.2. Examples	260
A.4. USB Filter Editor	261
A.4.1. Updating the USB Device Policy	262
A.4.2. Export a USB Policy	266
A.4.3. Import USB Policy	267
B. Configuring Red Hat Enterprise Linux 5.4 or Higher Virtual Machines to Use SPICE	269
C. Changing Passwords on the Red Hat Enterprise Virtualization Manager	271
C.1. Changing Password for the Administrative User	271
D. Reports Schema	273
D.1. Configuration History Views	273
D.1.1. v3_0_datacenter_configuration_view v3_0_latest_datacenter_configuration_view	273
D.1.2. v3_0_datacenter_storage_domain_map_view v3_0_latest_datacenter_configuration_view	274
D.1.3. v3_0_storage_domain_configuration_view v3_0_latest_storage_domain_configuration_view	274
D.1.4. v3_0_cluster_configuration_view/v3_0_latest_cluster_configuration_view	275
D.1.5. v3_0_host_configuration_view/v3_0_latest_host_configuration_view	275
D.1.6. v3_0_host_configuration_view/v3_0_latest_host_interface_configuration_view ..	277
D.1.7. v3_0_vm_configuration_view/v3_0_latest_vm_configuration_view	278
D.1.8. v3_0_vm_configuration_view/latest_vm_interface_configuration_view	279
D.1.9. v3_0_disks_vm_map_view/v3_0_latest_disks_vm_map_view	280
D.1.10. v3_0_vm_disk_configuration_view/v3_0_latest_vm_disk_configuration_view ...	281
D.2. Statistics History Views	282
D.2.1. v3_0_datacenter_samples_history_view/v3_0_datacenter_hourly_history_view v3_0_datacenter_daily_history_view	282

- D.2.2. v3_0_storage_domain_samples_history_view-
v3_0_storage_domain_hourly_history_view\v3_0_storage_domain_daily_history_view 283
- D.2.3. v3_0_host_samples_history_view\v3_0_host_hourly_history_view
v3_0_host_daily_history_view 283
- D.2.4. v3_0_host_interface_samples_history_view
v3_0_host_interface_hourly_history_view\v3_0_host_interface_daily_history_view 286
- D.2.5. v3_0_vm_samples_history_view\v3_0_vm_hourly_history_view
v3_0_vm_daily_history_view 287
- D.2.6. v3_0_vm_interface_samples_history_view
v3_0_vm_interface_hourly_history_view\v3_0_vm_interface_daily_history_view 289
- D.2.7. v3_0_vm_disk_daily_history_view\v3_0_vm_disk_hourly_history_view
v3_0_vm_disk_samples_history_view 290
- D.3. Tag History and ENUM Views 292
 - D.3.1. v3_0_tag_relations_history_view\v3_0_latest_tag_relations_history_view 292
 - D.3.2. v3_0_tag_details_view\v3_0_latest_tag_details_view 293
 - D.3.3. v3_0_enum_translator_view 293
- E. Search Parameters 295**
 - E.1. Searching for Resources 295
 - E.1.1. Searching for Data Centers 295
 - E.1.2. Searching for Clusters 295
 - E.1.3. Searching for Hosts 296
 - E.1.4. Searching for Storage 297
 - E.1.5. Searching for Virtual Machines 298
 - E.1.6. Searching for Pools 300
 - E.1.7. Searching for Templates 300
 - E.1.8. Searching for Users 302
 - E.1.9. Searching for Events 303
- F. SAP Monitoring 305**
- G. KVM Virtual Machine Timing Management 307**
- H. Additional References 311**
- I. Revision History 313**

Red Hat Enterprise Virtualization is a fully integrated virtualization solution for servers and desktops that provides management across the enterprise, and features live migration, high availability, system scheduling, power management, image management, snapshots, thin provisioning, import of foreign hypervisors, monitoring and reporting.

Red Hat Enterprise Virtualization does not limit the amount of memory, cores, or any other feature of the physical hardware used by virtual machines and offers unmatched scalability in the management of large numbers of virtual machines.

1. About this Guide

This guide is intended for advanced users and it assumes that you have successfully installed the Red Hat Enterprise Virtualization Manager and hosts and have an understanding of your data center resources. It describes how to use the Administration Portal, manage system components and virtual infrastructure, and configure and use the advanced features of Red Hat Enterprise Virtualization.

1.1. Documentation Suite

The Red Hat Enterprise Virtualization documentation suite provides information on installation, development of applications, configuration and usage of the Red Hat Enterprise Virtualization platform and its related products.

- *Red Hat Enterprise Virtualization — Administration Guide* (the book you are reading) describes how to setup, configure and manage Red Hat Enterprise Virtualization. It assumes that you have successfully installed the Red Hat Enterprise Virtualization Manager and hosts.
- *Red Hat Enterprise Virtualization — Evaluation Guide* enables prospective customers to evaluate the features of Red Hat Enterprise Virtualization. Use this guide if you have an evaluation license.
- *Red Hat Enterprise Virtualization — Installation Guide* describes the installation prerequisites and procedures. Read this if you need to install Red Hat Enterprise Virtualization. The installation of hosts, Manager and storage are covered in this guide. You will need to refer to the *Red Hat Enterprise Virtualization Administration Guide* to configure the system before you can start using the platform.
- *Red Hat Enterprise Virtualization — Manager Release Notes* contain release specific information for Red Hat Enterprise Virtualization Managers.
- *Red Hat Enterprise Virtualization — Power User Portal Guide* describes how power users can create and manage virtual machines from the Red Hat Enterprise Virtualization user portal.
- *Red Hat Enterprise Virtualization — Quick Start Guide* provides quick and simple instructions for first time users to set up a basic Red Hat Enterprise Virtualization environment.
- *Red Hat Enterprise Virtualization — REST API Guide* describes how to use the REST API to set up and manage virtualization tasks. Use this guide if you wish to develop systems which integrate with Red Hat Enterprise Virtualization, using an open and platform independent API.
- *Red Hat Enterprise Virtualization — Technical Reference Guide* describes the technical architecture of Red Hat Enterprise Virtualization and its interactions with existing infrastructure.
- *Red Hat Enterprise Virtualization — User Portal Guide* describes how users of the Red Hat Enterprise Virtualization system can access and use virtual desktops from the user portal.
- *Red Hat Enterprise Linux — Hypervisor Deployment Guide* describes how to deploy and install the Hypervisor. Read this guide if you need advanced information about installing and deploying

Hypervisors. The basic installation of Hypervisor hosts is also described in the *Red Hat Enterprise Virtualization Installation Guide*.

- *Red Hat Enterprise Linux — V2V Guide* describes importing virtual machines from KVM, Xen and VMware ESX to Red Hat Enterprise Virtualization and KVM managed by libvirt.

1.2. Audience

This advanced guide is intended for Linux or Windows system administrators who need to manage a virtual environment using Red Hat Enterprise Virtualization platform. An advanced level of system administration, preferably including familiarity with virtual machine data center operations, is assumed. This document is not intended for beginners. If you want to install or evaluate the system, read the *Red Hat Enterprise Virtualization Installation Guide* and *Red Hat Enterprise Virtualization Evaluation Guide*.

2. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the *Liberation Fonts*¹ set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later includes the Liberation Fonts set by default.

2.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

Mono-spaced Bold

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keycaps and key combinations. For example:

To see the contents of the file **my_next_bestselling_novel** in your current working directory, enter the **cat my_next_bestselling_novel** command at the shell prompt and press **Enter** to execute the command.

The above includes a file name, a shell command and a keycap, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from keycaps by the hyphen connecting each part of a key combination. For example:

Press **Enter** to execute the command.

Press **Ctrl+Alt+F2** to switch to the first virtual terminal. Press **Ctrl+Alt+F1** to return to your X-Windows session.

The first paragraph highlights the particular keycap to press. The second highlights two key combinations (each a set of three keycaps with each set pressed simultaneously).

¹ <https://fedorahosted.org/liberation-fonts/>

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **mono-spaced bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog box text; labeled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

Choose **System** → **Preferences** → **Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, click the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

To insert a special character into a **gedit** file, choose **Applications** → **Accessories** → **Character Map** from the main menu bar. Next, choose **Search** → **Find...** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit** → **Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

Mono-spaced Bold Italic or *Proportional Bold Italic*

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh john@example.com**.

The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount /home**.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above — *username*, *domain.name*, *file-system*, *package*, *version* and *release*. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

Publican is a *DocBook* publishing system.

2.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in **mono-spaced roman** and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts svgs
```

Source-code listings are also set in **mono-spaced roman** but add syntax highlighting as follows:

```
package org.jboss.book.jca.ex1;

import javax.naming.InitialContext;

public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref    = iniCtx.lookup("EchoBean");
        EchoHome        home   = (EchoHome) ref;
        Echo             echo   = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

2.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled 'Important' will not cause data loss but may cause irritation and frustration.



Warning

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

3. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report by email to the author of the manual, Kate Grainger (kgrainge@redhat.com). When submitting a bug report, be sure to mention the manual's identifier: *Guides-Admin*.

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, include the section number and some of the surrounding text so we can find it easily.

Introduction

Red Hat Enterprise Virtualization provides IT departments with the tools to meet the challenges of managing complex environments. Red Hat Enterprise Virtualization's rich virtualization platform enables administrators to reduce the cost and complexity of large deployments. Red Hat Enterprise Virtualization platform includes:

- High availability to quickly configure virtual machines for fault tolerance.
- Live migration to move virtual machines between physical hosts without interruption.
- System scheduler to create policies to dynamically balance compute resources.
- Power saver to create policies to conserve power and cooling costs.
- Image manager to create, manage and provision virtual machines.
- Storage virtualization to enable consistent access of common storage from any server.
- Multi-level administration to enable administration of physical infrastructure as well as administration of virtual objects.
- Ability to convert existing virtual machines on foreign hypervisors to Red Hat Enterprise Virtualization platform. This is completely described in the *V2V Guide*.
- A range of reports either from the reports module based on JasperReports, or from the data warehouse. The reports enable administrators to monitor and analyze information on virtual machines, hosts and storage usage and performance.

1. Red Hat Enterprise Virtualization Architecture

Red Hat Enterprise Virtualization platform consists of three components:

- **Red Hat Enterprise Virtualization Hypervisor** - based on Kernel Virtual Machine (KVM), is a thin virtualization layer deployed across the server's infrastructure. Because it is a core part of the Linux kernel, KVM is a highly efficient means of providing virtualization.
- **Agents and tools** include VDSM which runs in the hypervisor or host. These provide local management for virtual machines, networks and storage.
- **Red Hat Enterprise Virtualization platform management infrastructure** allows users to view and manage all the system components, machines and images from a single, powerful interface. The management system GUI provides a comprehensive range of features, including powerful search capabilities, resource management, live migrations, and provisioning.

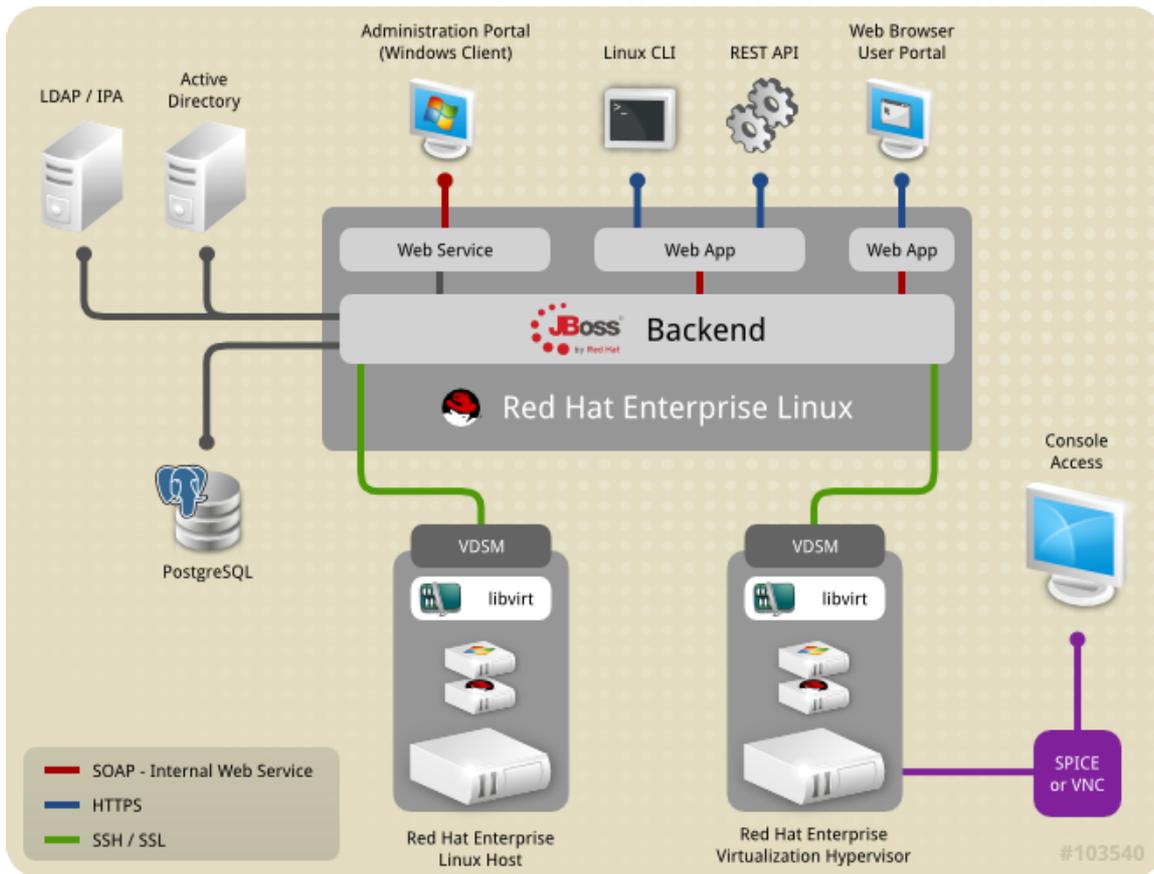


Figure 1. Red Hat Enterprise Virtualization Platform Overview

1.1. System Components

The various components work seamlessly together to enable the system administrator to set up, configure, and maintain the virtualized environment through an intuitive graphical interface.

1.1.1. About the Components

Red Hat Enterprise Virtualization platform consists of one or more hosts (either Red Hat Enterprise Virtualization Hypervisors or Red Hat Enterprise Linux 5.5 and higher systems) and at least one manager. The virtual machines are installed on the hosts. The system and all its components are managed through a centralized management system.

Hosts run the user's Windows XP, Windows 2000, Windows 7, or Linux virtual machines on Red Hat Enterprise Virtualization Hypervisor or Red Hat Enterprise Linux (5.5 and higher) and KVM virtualization technology (Kernel-based Virtual Machine). The hypervisor also includes a resource optimization layer that allows for better desktop interactivity and management.

The Red Hat Enterprise Virtualization Manager is a service running on a Red Hat Enterprise Linux 6 server that provides interfaces for controlling the Red Hat Enterprise Virtualization platform. It manages provisioning, connection protocols, user session logons and logoffs, virtual desktop pools, virtual machine images, and the high availability and clustering systems. Red Hat Enterprise Virtualization Manager requires Windows to remotely access the Administration Portal, but the Manager itself is hosted on a Red Hat Enterprise Linux server.

1.1.2. About the Virtual Machines

Red Hat Enterprise Virtualization platform enables you to create virtual machines that perform the same functions as physical machines. Using a standard Web browser, users can run virtual machines that behave like physical desktops. Multiple levels of permissions allow users with different roles to manage virtual machines to meet the requirements of the enterprise.

1.1.3. About SPICE

The SPICE protocol, which is installed on the client machine and runs on the client machine, allows you to experience PC-like graphics performance while connected to your virtual machine's computing desktop environment. It supplies video at more than 30 frames per second, bi-directional audio (for soft-phones/IP phones), bi-directional video (for video telephony/video conferencing) and USB redirection from the client's USB port into the virtual machine. SPICE also supports connection to multiple monitors with a single virtual machine. SPICE Client is installed on the client machine by ActiveX when the virtual machine is accessed through a browser from the Admin Portal or the User Portal.

2. Red Hat Enterprise Virtualization Resources

The Red Hat Enterprise Virtualization platform manages the following resources within the management infrastructure to create a powerful, scalable virtual environment. The components of the Red Hat Enterprise Virtualization platform fall into three categories: physical components, logical components, and resources. Physical components are the objects that occupy physical space in the world and make up a part of the Red Hat Enterprise Virtualization platform, and may be thought of as the "hardware" component of the platform. Logical components are the non-physical groupings and processes that enable Red Hat Enterprise Virtualization to provide its functionality, and may be thought of as the "software" component of the platform. Resources are administration tools and devices made available to the administrator by the Red Hat Enterprise Virtualization platform, and may be thought of as the "utilities" provided by the "software" of the logical components, which runs on the "hardware" of the physical components.

- **Hosts** - A host is a physical server (a physical machine) that runs either Red Hat Enterprise Virtualization Hypervisor or Red Hat Enterprise Linux 5.5 and above, and hosts one or more virtual machines. Hosts are grouped into clusters. Virtual machines can be migrated from one host to another within a cluster.
- **Clusters** - A cluster is a set of physical hosts that are treated as a resource pool for a set of virtual machines. Hosts in a cluster share the same network infrastructure and the same storage. They form a migration domain within which virtual machines can be moved from host to host.
- **Data Center** - A data center is a logical entity that defines the set of resources used in a specific environment. It is a collection of a number of clusters of virtual machines, storage and networks.

The data center is the highest level container for all physical and logical resources within a managed virtual environment.

A data center relies on adequate and accessible physical storage. The storage pool provides an abstracted view of the physical storage assigned to a data center, that enables planners and administrators to easily monitor and manage storage requirements.

- **Storage Pool** - The storage pool is a logical entity that contains a standalone image repository of a certain type, either iSCSI, or Fiber Channel, or NFS. Each storage pool can contain several storage domains, for virtual machine disk images and for ISO images and for the import and export of virtual machine images.

- **Virtual Machines** - A virtual machine is a virtual desktop or virtual server containing an operating system and a set of applications. Multiple identical desktops can be created in a *Pool*. Virtual machines can be accessed and used by end users, and created, managed or deleted by power users.
- **Desktop Pools** - A desktop pool is a group of identical virtual desktops that are available on demand by each one of the group members (not concurrently). Desktop pools can be set up for different purposes. For example, one desktop pool may be for the Marketing department, another for Research and Development, and so on. Users get available desktops of the required type from the appropriate pool.
- **Templates** - A template is a model virtual machine with a unique configuration and settings. A virtual machine that is based on a particular template acquires the configurations and settings of the template. Templates are used to conveniently and efficiently create a set of identical virtual machines. Using templates is the quickest way of creating large number of virtual machines in a single step.
- **Snapshots** - A snapshot is a view of a virtual machine's operating system and all its applications at a given point in time. It can be used to save the settings of a virtual machine before an upgrade, or before new applications are installed. In case of problems, the parameters from the snapshot can be used to restore the virtual machine to the state before the upgrade or installation.
- **User Types** - Red Hat Enterprise Virtualization supports multiple levels of administrators and users with distinct levels of permissions. System administrators can manage and administrate objects of the physical infrastructure, such as data centers, hosts and storage. Red Hat Enterprise Virtualization Power Users are administrators who manage the end-users of the virtual machines, as well as act as administrators of the virtual machines. End users are the users who have access to a specified desktop, or an available virtual machine from a designated desktop pool.
- **Events and Monitors** - Alerts, warnings, and other notices about activities within the system help the administrator to monitor the performance and running of various resources. Monitoring details can be displayed in both graphic and textual fashion.
- **Reports** - A range of reports either from the reports module based on JasperReports, or from the data warehouse. Preconfigured or ad hoc reports can be generated from the reports module. Users can also generate reports using any query tool that supports SQL from a data warehouse that collects monitoring data for hosts, virtual machines and storage.

3. Administration of the Red Hat Enterprise Virtualization Platform

This section provides a high level overview of the tasks and responsibilities of a system administrator for the Red Hat Enterprise Virtualization platform. The tasks are divided into two general groups:

- Configuring a new logical data center is the most important task of the system administrator. Designing a new data center requires an understanding of capacity planning and definition of requirements. Typically this is determined by the solution architect, who provides the requirement to the system architect. Preparing to set up the virtualized environment is a significant part of the set up, and is usually part of the system administrator's role.
- Maintaining the data center, including performing updates and monitoring usage and performance to keep the data center responsive to changing needs and loads.

The procedures to complete these tasks are described in detail in later sections of this guide.

It is assumed that you have already read the material in *Red Hat Enterprise Virtualization 3.0 Installation Guide*.

3.1. Maintaining the Red Hat Enterprise Virtualization Platform

This section describes how to maintain a Red Hat Enterprise Virtualization platform.

The administrator's tasks include:

- Managing physical and virtual resources such as hosts and virtual machines. This includes upgrading or adding hosts, importing domains, and converting virtual machines created on foreign hypervisors, and the maintenance of pools of desktops.
- 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, freeing resources, for example, by shutting down machines).
- Responding to the new requirements of virtual machines (for example, upgrading the operating system for a set of virtual desktops or allocating more memory to a specific virtual server).
- Managing customized object properties (Tags).
- Managing searches saved as public bookmarks.
- Managing user setup and access and setting user and administrator permission levels. This includes assigning or customizing roles to suit the needs of the enterprise.
- Troubleshooting for specific users or virtual machines or overall system functionality.
- Generating general and specific reports.

These tasks are described in detail in later sections of this guide.

Part I. The Red Hat Enterprise Virtualization Interface

Getting Started

The Administration Portal allows you to monitor, create and maintain your Red Hat Enterprise Virtualization platform using an interactive graphical user interface (GUI). The GUI functions in two modes, tree or flat, allowing you to navigate to the system's resources, either hierarchically or directly. The powerful Search feature ensures that you can locate any resource in the enterprise, wherever it may be in the hierarchy; and you can use Tags and Bookmarks to help you to store the results of your searches for later reference.

It is assumed that you have correctly installed Red Hat Enterprise Virtualization, including Directory Services, and have logged into the Administration Portal. If you are attempting to set up Red Hat Enterprise Virtualization, it is recommended that you read the *Red Hat Enterprise Virtualization Quick Start Guide* and the *Red Hat Enterprise Virtualization Installation Guide 3.0*.

1.1. Graphical User Interface

After you have successfully logged into Red Hat Enterprise Virtualization, the Administration Portal displays. The graphical interface consists of a number of contextual panes and menus, and can be used in two modes, tree mode and flat mode. Tree mode allows you to browse the object hierarchy of a data center, and is the recommended manner of operation. The elements of the GUI are shown in the diagram below.

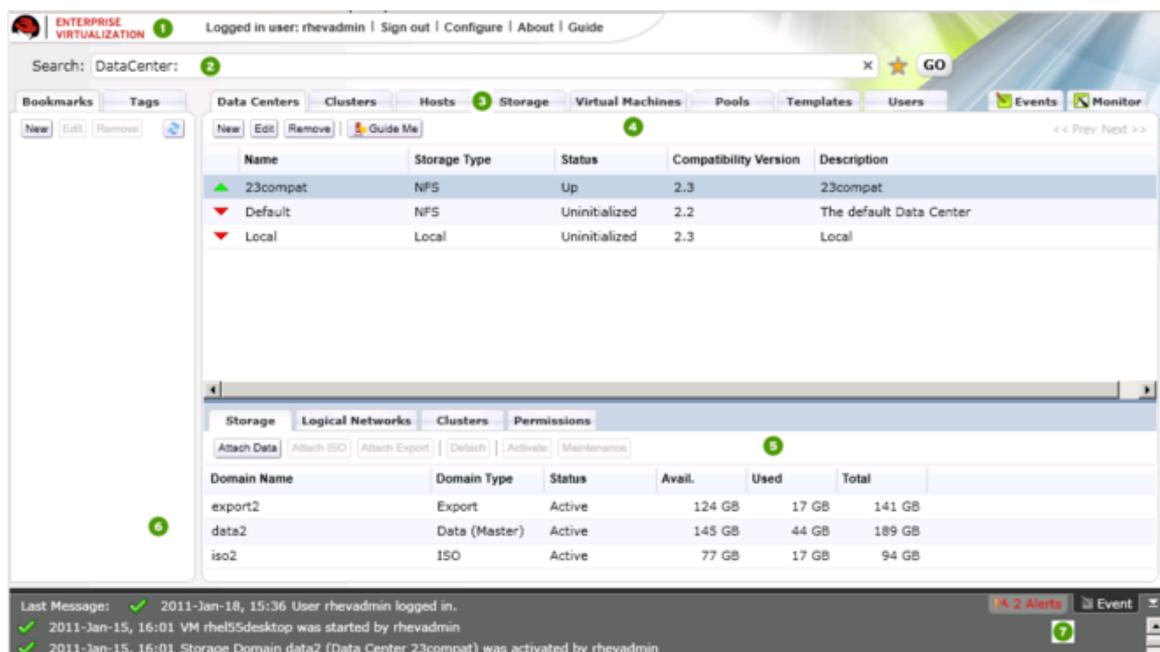


Figure 1.1. User Interface Elements of the Administration Portal

User Interface Elements

- **1 Header**

The **Header** Bar contains the name of the current logged in user, the **Sign out** button, the **About** button, and the **Configure** button. The **About** button provides access to version information. The **Configure** button allows you to configure user roles. The **Guide** provides a shortcut to the book you are reading now.

- **2 Search Bar**

The **Search** bar allows you to quickly search for resources such as hosts and virtual machines. You can build queries to find the resources that you need. Queries can be as simple as a list of all the

hosts in the system, or much more complex. As you type each part of the search query, you will be offered choices to assist you in building the search. The star icon can be used to save the search as a bookmark.

- **3 Resource Tabs**

All **Resources**, such as Hosts and Clusters, can be managed using the appropriate tab. Additionally, the Events and Monitor tabs allow you to manage and view events across the entire system.

Clicking a tab displays the results of the most recent search query on the selected object. For example, if you recently searched for all virtual machines starting with "M", clicking the virtual machines tab displays a list of all virtual machines starting with "M".

The Administration Portal uses the following tabs: Data Centers, Clusters, Hosts, Storage, Virtual Machines, Pools, Templates, Users, Events, and Monitor.

- **4 Results list**

Perform a task on an individual item, multiple items, or all the items in the results list, by selecting the item(s) and then clicking the relevant action button. If multiple selection is not possible, the button is disabled.

Details of a selected item display in the details pane.

- **5 Details Pane**

The **Details** pane displays detailed information about a selected item in the Results Grid. If multiple items are selected, the Details pane displays information on the first selected item only.

- **6 Bookmarks Pane**

Bookmarks are used to save frequently-used or complicated searches for repeated use. Bookmarks can be added, edited, or removed.

- **7 Alerts/Events Pane**

The **Alerts** pane lists all events with a severity of Error or Warning. The system records all events, which are listed as audits in the Alerts section. Like Events, Alerts can also be viewed in the lowermost panel of both the Monitor and Events tab, by re-sizing the panel and clicking the Alert tab. This tabbed panel also appears in other tabs, such as the Hosts tab.



Minimum Supported Browser Resolution for Working with the Administration Portal

The minimum supported resolution viewing the Administration Portal in a web browser is 1024x768. When viewed at a lower resolution, the Administration Portal will not render correctly.

1.1.1. Tree Mode and Flat Mode

The Administration Portal provides two different modes for managing your resources, tree mode and flat mode. Tree mode displays resources in a hierarchical view per data center, from the highest level of the data center, down to the individual virtual machine. Tree mode provides a visual representation of the virtualization system. Working in tree mode is highly recommended for most operations.

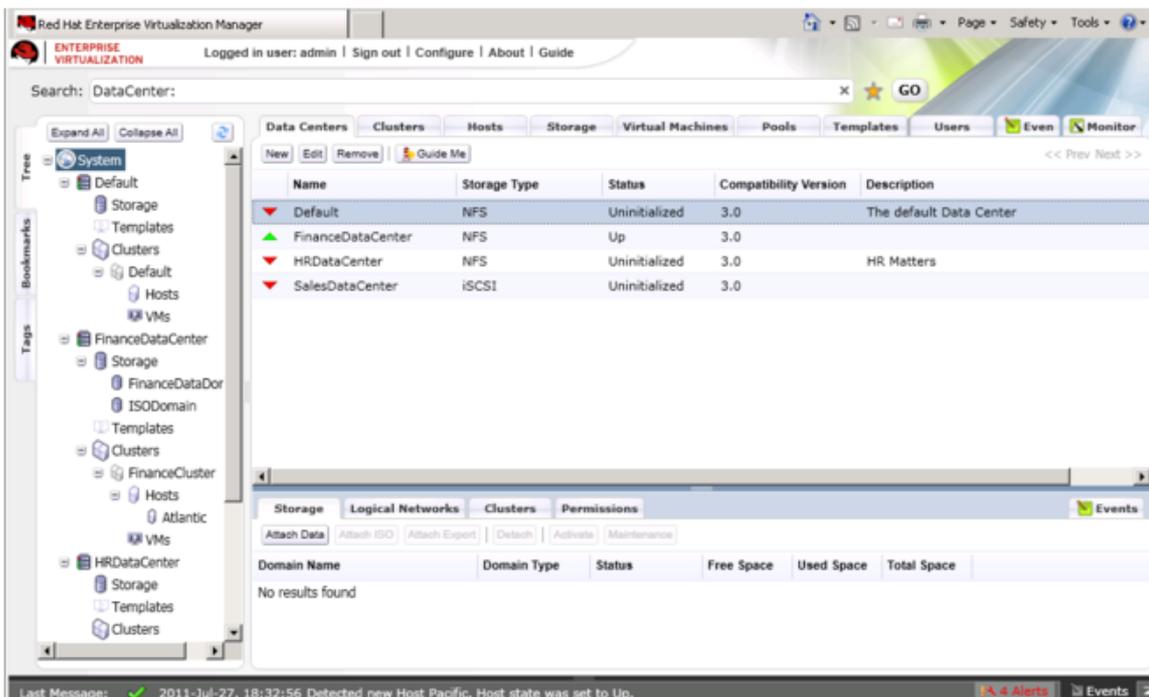


Figure 1.2. Tree Mode

Flat mode offers powerful search functionality and allows you to customize how you manage your system. It gives you access to any resource, regardless of its position in the enterprise. In this mode the full power of the search feature can be used. Flat mode does not limit you to viewing the resources of a single hierarchy, but allows you to search across data centers, or storage domains. For example, you may need to find all virtual machines that are using more than 80% CPU across clusters and data centers, or locate all hosts that have the highest utilization. Flat mode makes this possible. In addition, certain objects are not in the data centers hierarchy, so they will not appear in tree mode. Pools and users are not parts of the data centers hierarchy, and can be accessed only in flat mode.

To access flat mode, click on the "System" item in the pane on the left-hand side of the screen. You will know that you are in flat mode if the "Pools" and "Users" resource tabs appear.

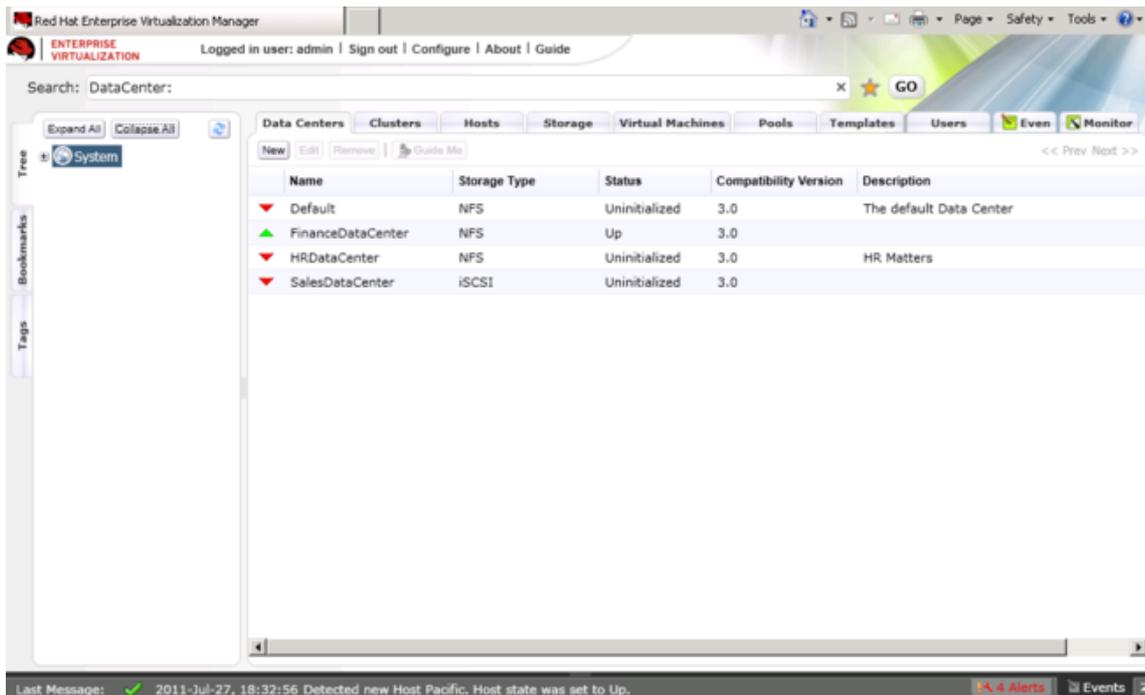


Figure 1.3. Flat Mode

1.1.2. Using the Guide Me facility

At set up or configuration time, a number of tasks must be completed in sequence. The Administration Portal provides prompts in the form of context sensitive **Guide Me** dialog box, buttons and icons. The **Guide Me** dialog box allows you to directly perform the required tasks. The **Guide Me** dialog box is context sensitive, and only displays the actions that are appropriate to the resource that is being configured. The **Guide Me** dialog box can be accessed at any time by clicking the **Guide Me** button on the resource toolbar. You can use the **Guide Me** facility to set up or configure data centers, storage, and clusters, and its use in specific contexts is described in later sections of this document.

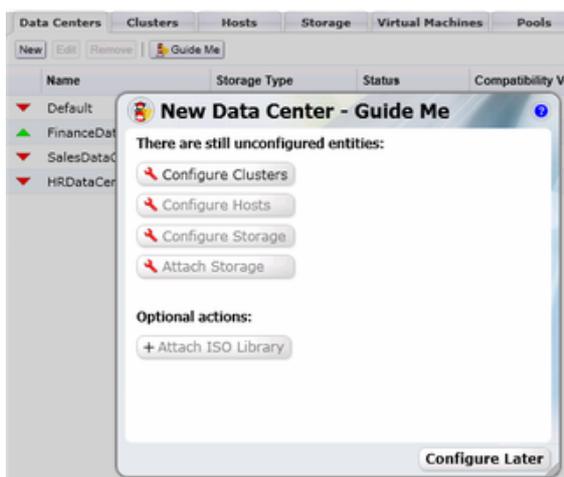


Figure 1.4. New Data Center Guide Me Dialog Box

1.2. Search

The Administration Portal environment is designed to enable the management of thousands of resources, such as virtual machines, hosts, users, and more. When managing the virtual desktop environment, it is recommended that large lists of resources, such as virtual machines, are reduced to

a manageable number (for example, 10). This allows tasks to be performed on the smaller list, or to select specific items on the list on which to perform a given task.

To perform a search, enter the search query (free-text or syntax-based) in the Search Bar at the top of the Administration Portal. Search queries can be saved as a Bookmarks for future reuse ([Section 1.2.2, "Saving and Accessing Queries as Bookmarks"](#)). This eliminates the need to reenter a search query each time the specific search results are needed.

1.2.1. Search Syntax

The syntax of the search queries for Red Hat Enterprise Virtualization resources is as follows:

result-type: {criteria} [sortby sort_spec]

Syntax Examples

The following examples describe how the search query is used and help you to understand how Red Hat Enterprise Virtualization assists with building search queries.

Table 1.1. Example Search Queries

Example	Result
Hosts: Vms.status = up	Displays a list of all hosts running virtual machines that are up.
Vms: domain = qa.company.com	Displays a list of all virtual machines running on the specified domain.
Vms: users.name = Mary	Displays a list of all virtual machines belonging to users with the username Mary.
Events: severity > normal sortby time	Displays the list of all Events whose severity is higher than Normal, sorted by time.

1.2.1.1. Auto-Completion

The Administration Portal provides auto-completion to help you create valid and powerful search queries. As you type each part of a search query, a drop-down list of choices for the next part of the search opens below the Search Bar. You can either select from the list and then continue typing/ selecting the next part of the search, or ignore the options and continue entering your query manually.

The following table specifies by example how the Administration Portal auto-completion assists in constructing a query:

Hosts: Vms.status = down

Table 1.2. Example Search Queries Using Auto-Completion

Input	List Items Displayed	Action
h	Hosts (1 option only)	Select Hosts or; Type Hosts
Hosts:	All host properties	Type v
Hosts: v	host properties starting with a v	Select Vms or type Vms
Hosts: Vms	All virtual machine properties	Type s
Hosts: Vms.s	All virtual machine properties beginning with s	Select status or type status

Input	List Items Displayed	Action
Hosts: Vms.status	= =!	Select or type =
Hosts: Vms.status =	All status values	Select or type down

1.2.1.2. Result-Type Options

The Result-type allows you to search for resources of any of the following types:

- **Vms** for a list of virtual machines
- **Host** for a list of hosts
- **Pools** for a list of pools
- **Template** for a list of templates
- **Event** for a list of events
- **Users** for a list of users
- **Cluster** for a list of clusters
- **Datacenter** for a list of data centers
- **Storage** for a list of storage domains

As each type of resource has a unique set of properties and a set of other resource types that it is associated with, each search type has a set of valid syntax combinations. These are specified in [Appendix E, Search Parameters](#). However, using the auto-complete feature will help you to create valid queries easily.

1.2.1.3. Search Criteria

You can specify the search criteria after the colon in the query. The syntax of **{criteria}** is as follows:

<prop> <operator> <value>

or

<obj-type> <prop> <operator> <value>

Examples

The following table describes the parts of the syntax:

Table 1.3. Example Search Criteria

Part	Description	Values	Example	Note
prop	The property of the searched-for resource. Can also be the property of a resource type	See the information for each of the search types in Section 1.2.1.3.1 ,	Status	--

Part	Description	Values	Example	Note
	(see obj - type), or <i>tag</i> (custom tag).	<i>"Wildcards and Multiple Criteria"</i>		
obj-type	A resource type that can be associated with the searched-for resource.	See the explanation of each of the search types in Section 1.2.1.3.1 , <i>"Wildcards and Multiple Criteria"</i>	Users	--
operator	Comparison operators.	= != (not equal) > < >=	--	Value options depend on obj-type.
Value	What the expression is being compared to.	<= String Integer Ranking Date (formatted according to Regional Settings)	Jones 256 normal	<ul style="list-style-type: none"> Wildcards can be used within strings. "" (two sets of quotation marks with no space between them) can be used to represent an un-initialized (empty) string. Double quotes should be used around a string or date containing spaces

1.2.1.3.1. Wildcards and Multiple Criteria

Wildcards can be used in the **<value>** part of the syntax for strings. For example, to find all users beginning with **m**, enter **m***.

You can perform a search having two criteria by using the Boolean operators **AND** and **OR**. For example:

Vms: users.name = m* AND Vms.status = Up

This query returns all running virtual machines for users whose names begin with "m".

Vms: users.name = m* AND Vms.tag = "paris-loc"

This query returns all virtual machines tagged with "paris-loc" for users whose names begin with "m".

When two criteria are specified without **AND** or **OR**, **AND** is implied. **AND** precedes **OR**, and **OR** precedes implied **AND**.

1.2.1.4. Determining Sort Order

You can determine the sort order of the returned information by using **sortby**. Sort direction (**asc** for ascending, **desc** for descending) can be included.

For example:

```
events: severity > normal sortby time desc
```

This query returns all Events whose severity is higher than Normal, sorted by time (descending order).

1.2.2. Saving and Accessing Queries as Bookmarks

Search queries can be saved as **Bookmarks**. This allows you to sort and display results lists with a single click. You can save, edit and remove bookmarks with the Bookmarks pane.

1.2.2.1. Creating Bookmarks

Bookmarks can be created for any type of available search, using a number of criteria.

To save a query string as a **Bookmark**:

1. In the Search Bar, enter the desired search query (see [Section 1.2.1, "Search Syntax"](#)).
2. Click the star-shaped **Bookmark** button to the right of the Search Bar.

The **New Bookmark** dialog box displays. The query displays in the **Search String** field. You can edit it if required.

3. In **Name**, specify a descriptive name for the search query.
4. Click **OK** to save the query as a bookmark.
5. The search query is saved and displays in the **Bookmarks** pane.

1.2.2.2. Editing Bookmarks

Bookmarks can be edited for any type of available search, using an existing bookmark.

To edit a bookmark:

1. Select the **Bookmark** pane by clicking on the **Bookmarks** tab on the far left side of the screen.
2. Select a bookmark from the **Bookmark** pane.
3. The results list displays the items according to the criteria. Click the **Edit** button on the **Bookmark** pane.

The **Edit Bookmark** dialog box displays. The query displays in the **Search String** field. Edit the search string to your requirements.

4. Change the **Name** and **Search String** as necessary.
5. Click **OK** to save the edited bookmark.

1.2.2.3. Deleting Bookmarks

Bookmarks can be deleted.

To delete a bookmark:

1. Select one or more bookmark from the **Bookmarks** pane.
2. The results list displays the items according to the criteria. Click the **Remove** button at the top of the **Bookmark** pane.

The **Remove Bookmark** dialog box displays, prompting you to confirm your decision to remove the bookmark.

3. Click **OK** to remove the selected bookmarks.

1.3. Tags

After your Red Hat Enterprise Virtualization platform is set up and configured to your requirements, you can customize the way you work with it using tags. Tags provide one key advantage to system administrators: they allow system resources to be arranged into groups or categories. This is useful when many objects exist in the virtualization environment and the administrator would like to concentrate on a specific set of them.

This section describes how to create and edit tags, assign them to hosts or virtual machines and search using the tags as criteria. Tags can be arranged in a hierarchy that matches a structure, to fit the needs of the enterprise.

Administration Portal Tags can be created, modified, and removed using the **Tags** pane.

To create a Tag:

1. In tree mode or flat mode, click the Resource tab for which you wish to create a tag, for example, **Hosts**.
2. Click the **Tags** tab. Select the node under which you wish to create the tag. For example, to create it at the highest level, click the root node. The **New** button is enabled.
3. Click **New** at the top of the **Tags** pane. The **New Tag** dialog box displays.
4. Enter the **Name** and **Description** of the new tag.
5. Click **OK**. The new tag is created and displays on the **Tags** tab.

To modify a Tag:

1. Click the **Tags** tab. Select the tag that you wish to modify. The buttons on the **Tags** tab are enabled.
2. Click **Edit** on the **Tags** pane. The **Edit Tag** dialog box displays.
3. You can change the **Name** and **Description** of the tag.
4. Click **OK**. The changes in the tag display on the **Tags** tab.

To delete a Tag:

1. Click the **Tags** tab. The list of tags will display.
2. Select the Tag(s) to be deleted. The **Remove Tag(s)** dialog box displays.

3. The tag, or list of tags are displayed in the dialog box. Check that you are sure about the removal. The message warns you that removing the tag will also remove all descendants of the tag.
4. Click **OK**. The new tag is removed and no longer displays on the **Tags** tab. The tag is also removed from all the objects that it was attached to.

Tags can be attached to Hosts and Virtual Machines only.

To add or remove a Tag to/from one or more object instances:

1. Search for the object(s) that you wish to tag/untag so that they are among the objects displayed in the results list.
2. Select one or more objects on the results list.
3. Click the **Assign Tags** button on the tool bar or right-click menu option.
4. A dialog box provides a list of Tags. Select the check box to assign a tag to the object. Or, deselect the check box to detach the tag from the object.
5. Click **OK**. The specified tag is now added/removed as a custom property of the selected object(s).

A user-defined tag can be a property of any object (for example, a virtual server or a host), and a search can be conducted to find it.

To search for objects using Tags:

- Follow the search instructions in [Section 1.2, "Search"](#), and enter a search query using "tag" as the property and the desired value or set of values as criteria for the search.

The objects tagged with the tag criteria that you specified are listed in the results list.

Part II. Managing System Components

Managing Data Centers

The data center is the highest level virtual entity for all physical and logical resources within a managed virtual environment. The data center contains the following physical resources; CPU 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. These resources are available to the virtual machines running in the data center. This section describes how to create and manage data centers, clusters and networking so that you can effectively set up the physical infrastructure for your virtual environment. Your Red Hat Enterprise Virtualization environment may contain several data centers with a complex topology of clusters, hosts, and networks; or it can consist of a single data center that uses the default data center provided by Red Hat Enterprise Virtualization. When you are setting up the physical infrastructure, as the administrator you should understand the set up and maintenance needs of the virtual machines that you intend to run in the data center. For example, consider whether some virtual machines will require dedicated resources, or can be migrated freely across hosts; whether you will need to set up desktop pools or clusters of servers; whether some machines can be expected to have consistent or varying workloads.

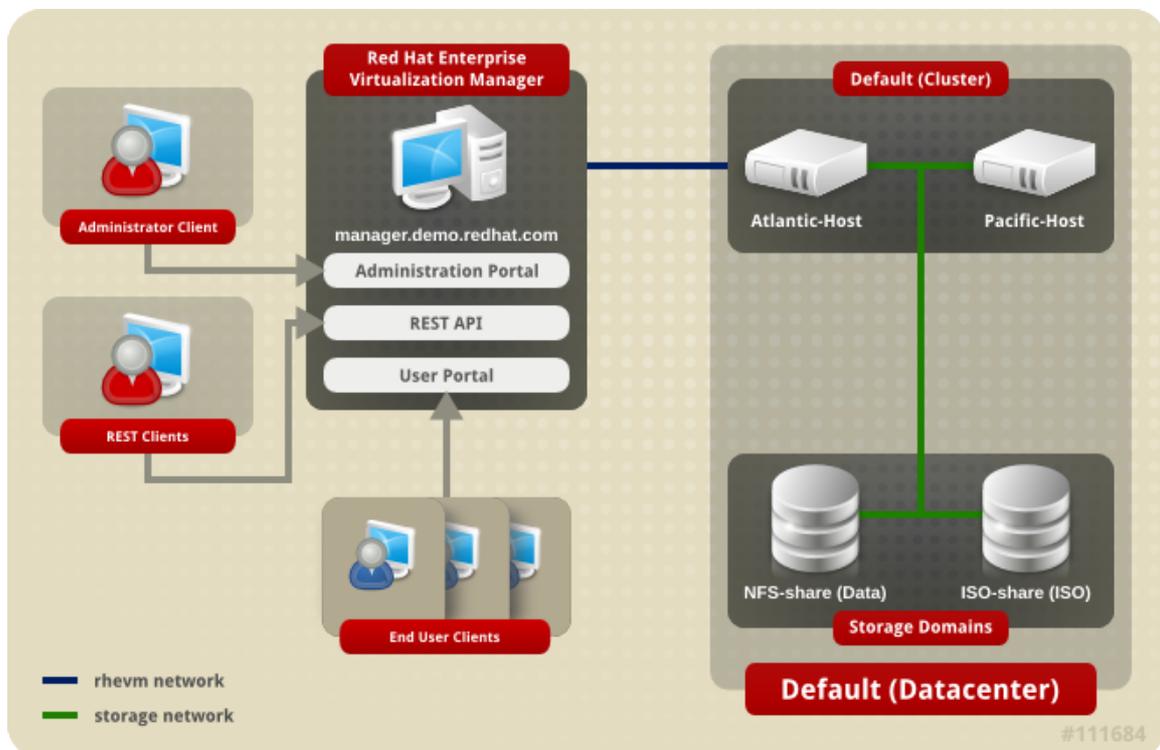


Figure 2.1. Data Centers

In all instances, it is assumed that you have successfully installed the Red Hat Enterprise Virtualization Manager, and have fulfilled the prerequisites as detailed in the *Installation Guide* and *Quick Start Guide*.

All resources from a data center, cluster, networks, hosts, through to storage and permissions provide default values when working in flat mode; while tree mode ensures that the resource is located in the correct hierarchy. This guide assumes the use of tree mode, unless the flat mode is specified. For a brief overview of terminology, see [Section 2, “Red Hat Enterprise Virtualization Resources”](#).

2.1. Data Centers

A data center is a logical entity that defines the set of physical and logical resources used in a managed virtual environment, consider it a container with clusters of hosts, virtual machines, storage

and networks. Though Red Hat Enterprise Virtualization contains a default data center at installation, you can set up additional data centers. All data centers are managed from the single Administration Portal. For example, an organization may have different data centers for different physical locations, business units, or for reasons of security. It is recommended that you do not remove the default data center, instead set up new appropriately named data centers.

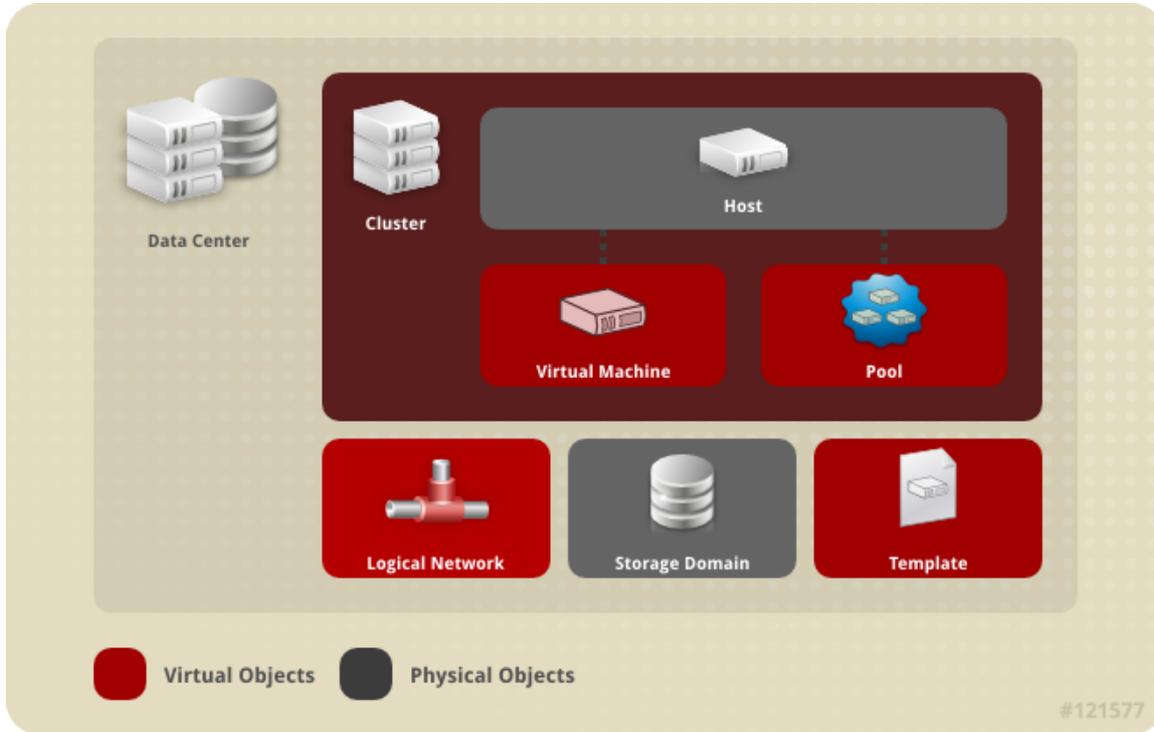


Figure 2.2. Data Center Objects

2.1.1. Data Center Properties

Data Centers can be defined, managed and viewed using the **Data Centers** tab on the Red Hat Enterprise Virtualization Manager Administration Portal.

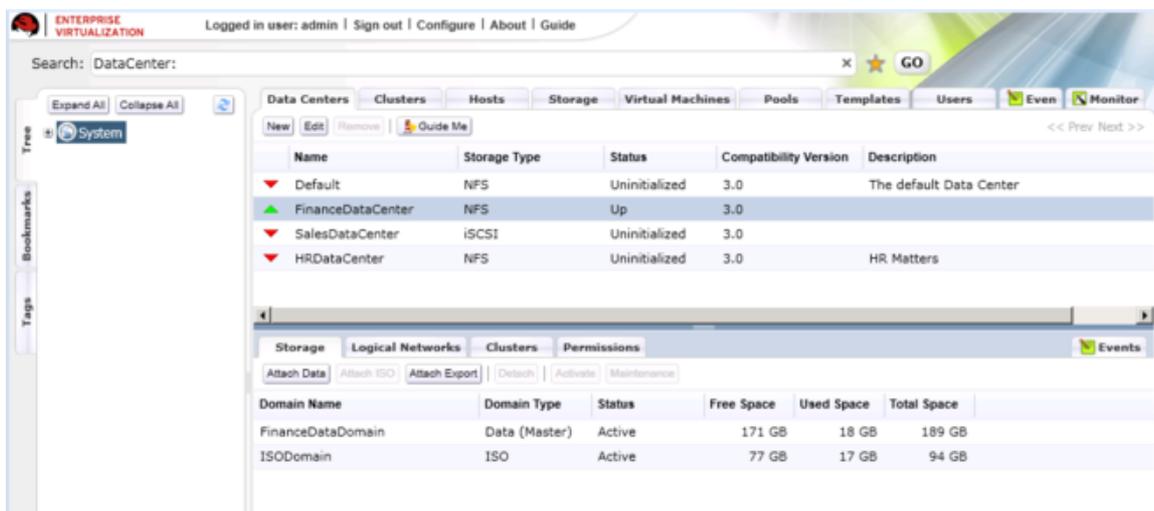


Figure 2.3. Data Center Tab

The table below describes the properties of a data center as displayed in the **New Data Center** and **Edit Data Center** dialog boxes. Missing mandatory fields and invalid entries are outlined in red when

you click **OK** to close the dialog box. In addition, field prompts indicate the expected values or range of values.

Table 2.1. Data Center Properties

Field/Tab	Description/Action
Name	The name of the data center. This must be a unique name with any combination of uppercase or lowercase letters, numbers, hyphens and underscores. Maximum length is 40.
Description	The description of the data center. While this is not a mandatory field, it is recommended.
Type	Any one of, NFS, iSCSI, FCP or Local Storage. This cannot be changed after creation without significant disruption and high possibility of data loss. The type of data domain dictates the type of the data center. All storage in a data center must be of one type only. For example, if iSCSI is selected as the type, only iSCSI data domains can be attached to the data center.
Compatibility Level	The version of the Red Hat Enterprise Virtualization, either 2.2 or 3.0. When you upgrade to a new version, it is necessary to upgrade the hosts, and then the clusters, and finally the data center. After upgrading the 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.

2.1.2. Data Centers Operations

You can create, edit and configure a data center from the Administration Portal from the **Data Centers** tab.

To create a data center from the **Data Centers** tab :

1. Click **New > New Data Center**.

Enter the properties in the **New Data Center** dialog box and click **OK**.

2. The **Guide Me** dialog box displays the configuration tasks that must be completed before the data center can be activated. See [Section 1.1.2, "Using the Guide Me facility"](#). Click **Configure Later** to close the dialog box.
3. The new data center displays in appropriate search results or lists of data centers, with a status of **Uninitialized**. An uninitialized data center requires further configuration, for example, storage domains must be attached to it. Either click the **Configure Storage** button on the **Guide Me** dialog box or select the new data center in the list, and click the **Storage** tab in the Details pane. You can define existing storage for the data center, or attach existing storage domains to the data center.

To configure a new data center using the **Guide Me** dialog box:

1. Select the required data center from the Tree panel, and click the **Guide Me** button on the **Data Centers** tab.

- The **Guide Me** dialog box displays the configuration tasks that must be completed before the data center can be activated. To set up a data center, a number of tasks must be completed in sequence. You can perform this manually, by using the various tabs and dialog boxes, or you can use the **Guide Me** dialog boxes, which provide sequential, context-sensitive steps that lead to a complete and accurate configuration. The **Guide Me** dialog box prompts you to complete the required tasks; because it is context sensitive it only displays the actions that are appropriate to the resource that is being configured. If you are configuring a new data center, the dialog box prompts you to configure clusters, hosts, and storage. In addition, you can attach an ISO library as an option.

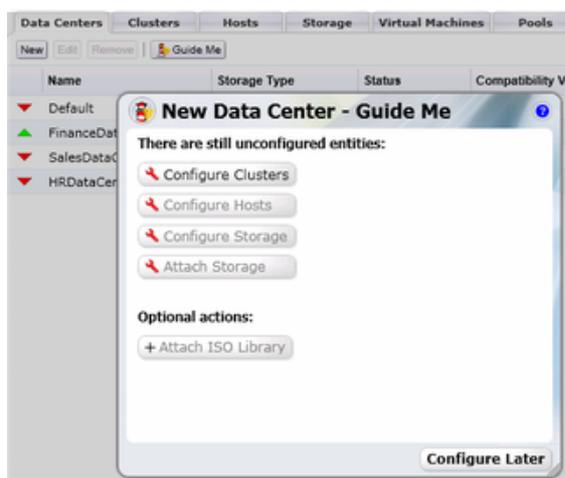


Figure 2.4. New Data Center Guide Me Dialog Box

- Click each of the buttons in order, and follow the prompts and dialog boxes to set up all the required objects for the data center. Configuring clusters, hosts and storage are described in detail later in this document.
- Click **Configure Later** to close the dialog box.

The **Guide Me** dialog box can also be accessed from the **Guide Me** button on the resource toolbar.

2.1.3. Data Centers Related Entities

To configure the data center, the related objects that belong to the data center also need to be set up. The Storage, Logical Networks, Clusters and Permissions tabs on the Details pane of a selected data center enable you to correctly define the objects in the data center, where they are to be stored, and how they are to work together.

Data Centers Storage Entities

A data center requires storage domains to be attached to it. Use data domains to store images and data of the virtual machines, an ISO domain to store the images used to install and run the virtual machines, and optionally an export domain to export virtual machines and templates from one data center to another.

Table 2.2. Data Centers Storage Tab Properties

Field/Tab	Description/Action
Domain Name	The name of the storage domain.
Domain Type	Either ISO, data or export domain. While ISO domains can be shared across data centers, data domains cannot be shared. An export

Field/Tab	Description/Action
	domain can only be attached to only one data center at a time.
Status	Either Active or Inactive.
Free Space, Used Space and Total Space	The amount of hard drive space on the storage domain that is available or used, or the total amount of hard drive space.

Data Center Logical Network Entities

The logical networks that are required for the virtual machines to communicate to hosts and storage, and for the Manager to communicate with all entities need to be defined for a data center.

Table 2.3. Data Center Logical Network Tab Properties

Field/Tab	Description/Action
Name	The name of the network belonging to the data center.
Description	The description of the network belonging to the data center.

Data Center Clusters Entities

The clusters in a data center are the migration domains for virtual machines. The ability to migrate virtual machines from one host to another in a cluster is a key feature of virtualization. In an upgrade, the clusters need to be upgraded before the data center can work in the upgraded version.

Table 2.4. Data Center Clusters Tab Properties

Field/Tab	Description/Action
Name	The name of the clusters belonging to the data center.
Compatibility	Compatibility Version- Either 2.2 or 3.0. When you upgrade to a new version, it is necessary to upgrade the hosts, and then the clusters. After upgrading the Manager, the hosts, clusters and data centers may still be in the earlier version. Ensure that you have upgraded all the hosts, before you upgrade the Compatibility Level of the cluster.
Description	The description of the cluster.

Data Center Permissions Entities

The users who have permissions to work in the data center, and their roles determine who can work in a data center, and what operations they can carry out.

Table 2.5. Data Center Permissions Tab Properties

Field/Tab	Description/Action
User	The user name of an existing user in Directory Services.

Field/Tab	Description/Action
Role	The role of the user. The role is a combination of user, permission level, and object. Roles can be default or customized roles.
Inherited Permissions	The hierarchical permissions available to the user.

2.1.4. Data Centers Troubleshooting

The **Data Centers** tab enables you to track problems or change the properties or entities attached to a data center. You can use the **Events** tab on the Details pane to view a complete list of events to track down possible errors.

If the data center has an Active status, you can only edit the **Name**, **Description** and **Compatibility Level** of the data center. If it is Inactive, you can also change the storage **Type**.

To edit a data center from the **Tree** panel:

1. Select the required data center and click the **Edit** button on the **Data Centers** tab.
2. Change the properties in the **Edit Data Center** dialog box and click **Save**.



Important

If you have changed the Storage type (of an inactive data center only), ensure that you define the storage domains for the data center.

Using the Event Tab to Troubleshoot

The **Events** tab for the data center, displays all events that occur in the system, for that selected data center. The types of events that appear in the **Events** tab are audits, warnings, and errors. In addition, the names of the user, host, virtual machine, and/or template involved in the event are listed. This makes it possible to determine the cause of the event. Select the data center in the Tree pane, and then click the **Events** tab. This will enable you to identify the problem objects, and take corrective action.

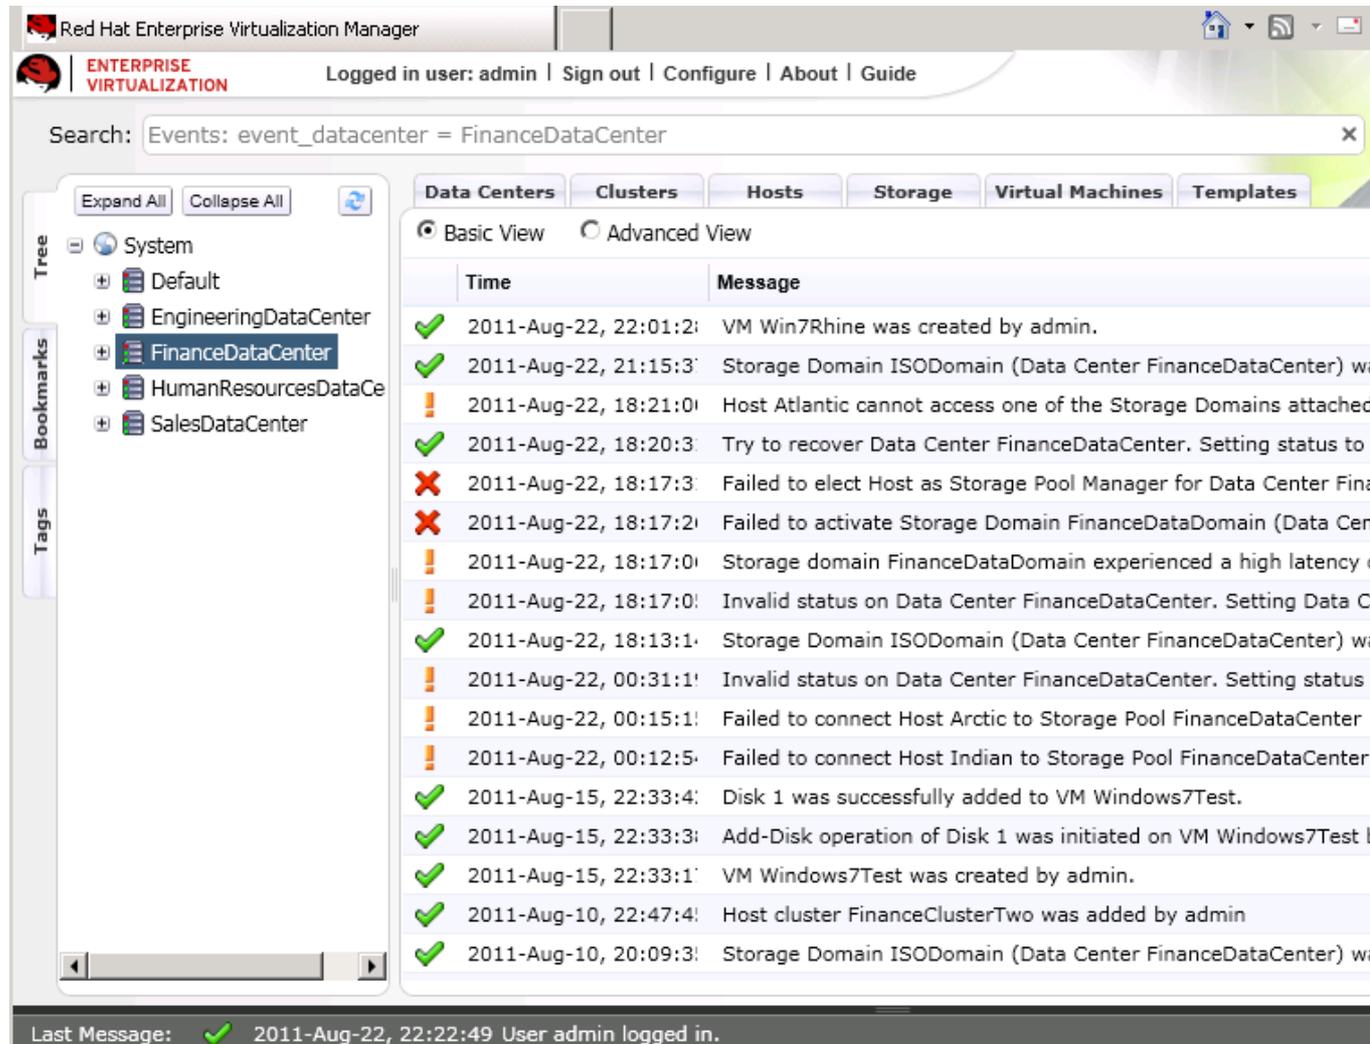


Figure 2.5. The **Events** Tab

Removing or Re-initializing a Data Center

Data centers can be removed entirely from the platform, or they can be reinitialized. Both actions are irreversible and should not be undertaken without careful consideration. Data Centers that are not in use can be permanently removed. Deleting unused data centers saves system resources, as existing hosts are checked (or pinged) at fixed intervals. Data centers can only be removed if there are no running hosts within any cluster belonging to the data center.



Note

The clusters, hosts and storage domains are not removed, and can be allocated to a different data center.

To remove a data center:

1. Click the **Data Centers** tab.
2. If the required data center is not visible, perform a search (see [Section 1.2, “Search”](#)).
3. Select the data center to be removed. Ensure that there are no running hosts in any cluster. You can use the **Remove** button or right-click on the data center and select **Remove**.
4. Click the **Remove** button.

A message prompts you to confirm removal.

5. Click **OK**. The data center is deleted and is no longer displayed on the **Data Centers** tab.

Data Centers that are not in use can be reinitialized. All data, connections and virtual machines are removed, and the data center is available to be set up again. Data centers can only be reinitialized if there are no running hosts within any cluster belonging to the data center.



Warning

The storage domains are initialized, and all existing data is removed.

To reinitialize a data center:

1. Click the **Data Centers** tab.
2. If the required data center is not visible, perform a search (see [Section 1.2, “Search”](#)).
3. Select the data center to be reinitialized. Ensure that there are no running hosts in any cluster. Right-click on the data center and select **Reinitialize Data Center**.
4. A warning message displays and prompts you to confirm initialization. You are also prompted to select a new storage domain for the data center.

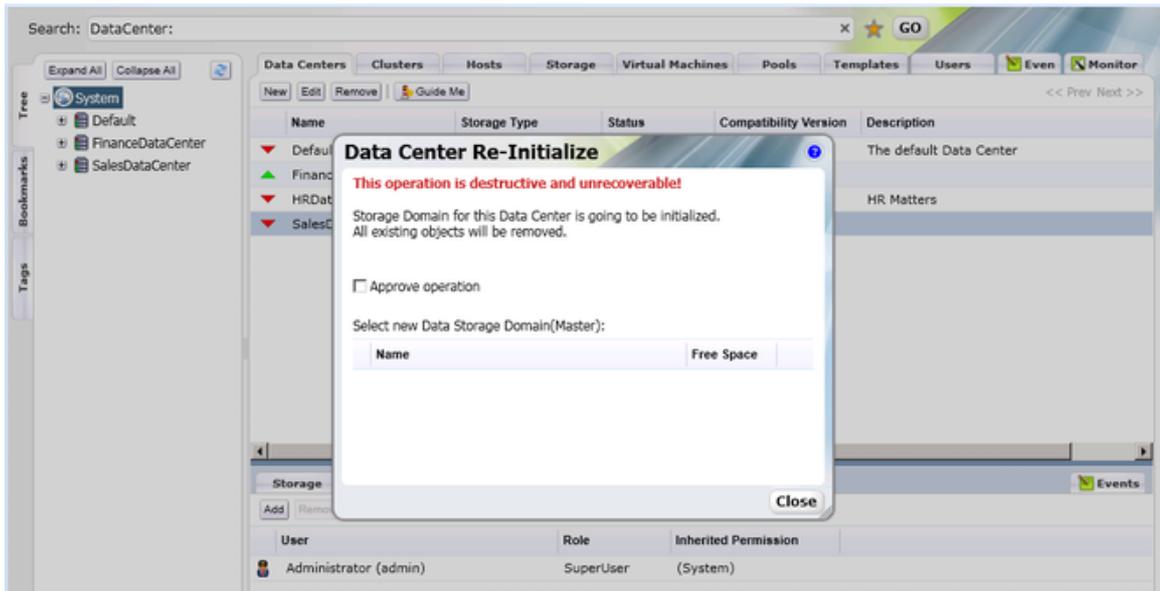


Figure 2.6. Data Center Re-Initialize Dialog Box

5. Select the **Approve Operation** checkbox.

The storage domain is initialized and all objects removed from it. This includes all virtual machine images, templates and snapshots. You can now set up the data center again, attaching the initialized, or a new storage domain.



Warning

Removal and initialization of data centers are destructive and irreversible operations.

2.1.5. System Permissions

The system administrator, as the superuser, can manage all aspects of the platform, that is, data centers, storage pools, users, roles and permissions by default; however more specific administrative roles and permissions can be assigned to other users. For example, the enterprise may need a data center administrator for a specific data center, or a particular cluster may need an administrator. All system administration roles for physical resources have a hierarchical permission system. For example, a data center administrator will automatically have permission to manage all the objects in that data center, storage, cluster and hosts; while a cluster administrator can manage all objects in the particular cluster.

While the system administrator of the data center has the full range of permissions, a data center administrator is a system administration role for a specific data center only. This is a hierarchical model, and means that if a user is assigned the data center administrator role for a data center, all objects in the data center can be managed by the user. The data center administrator role permits the following actions:

- Creation and removal of specific clusters.
- Addition and removal of hosts, virtual machines, pools.
- Permission to attach users to virtual machines within a single data center.

Chapter 2. Managing Data Centers

This is useful in a data center where there are multiple data center, each of which require their own system administrators. A data center administrator has permissions for the assigned data center only, not for all data centers in the system.

To assign a system administrator role to a data center:

1. Click the **Data Center** tab. A list of data centers displays. If the required data center is not visible, perform a search (see [Section 1.2, “Search”](#)).

Select the required data center, and click the **Permissions** tab on the Details pane. The **Permissions** tab displays a list of users and their current roles and permissions, if any.

2. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
3. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **DataCenterAdmin**.

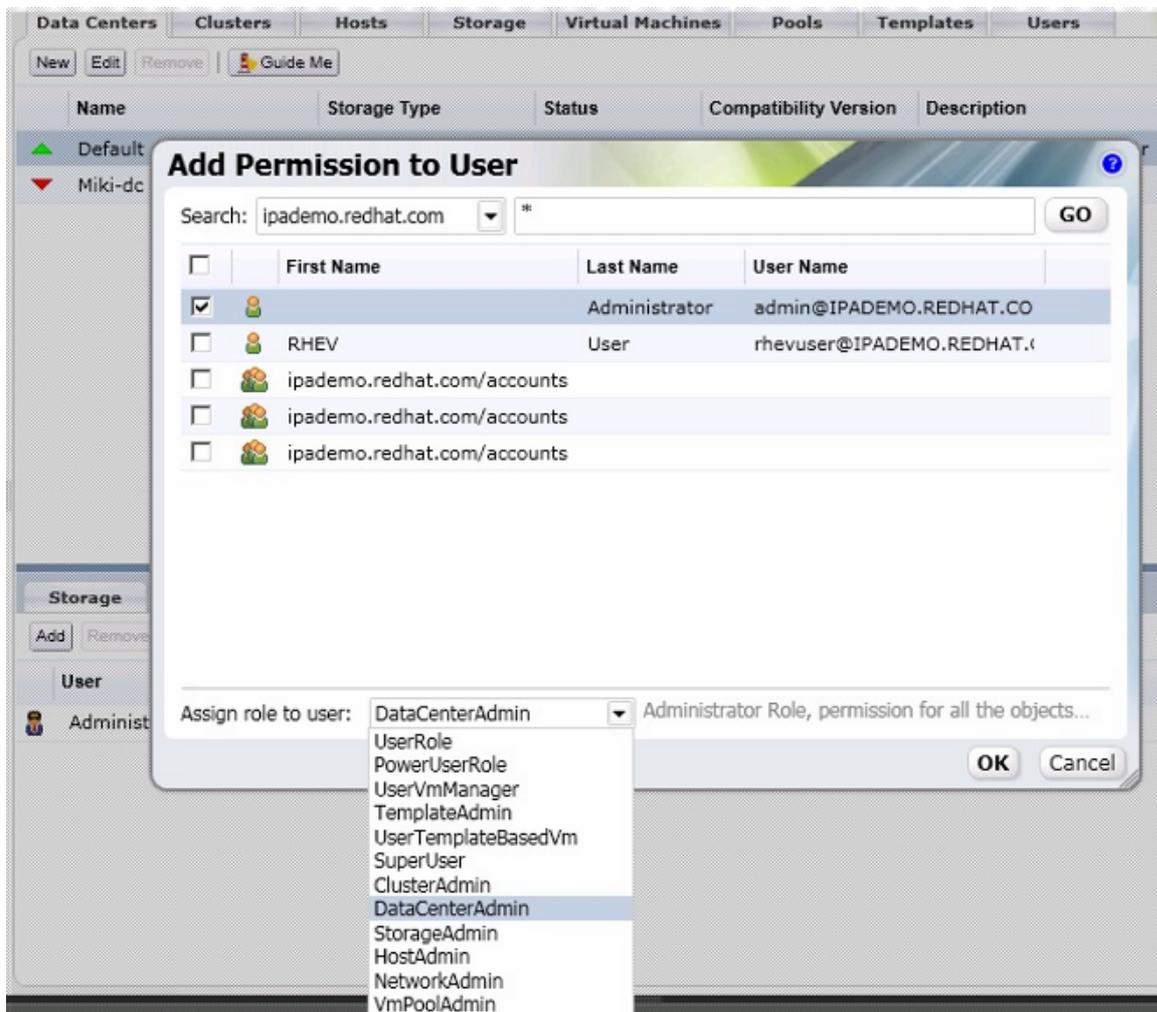


Figure 2.7. Assign DataCenterAdmin Permission

4. Click **OK**. The name of the user displays in the **Permissions** tab, with an icon and the assigned Role.



Note

You can only assign roles and permissions to existing users. See (see [Chapter 5, Users](#)).

You can also change the system administrator of a data center, by removing the existing system administrator, and adding the new system administrator, as described in the previous procedure.

To remove a system administrator role:

1. Click the **Data Center** tab. A list of data centers displays. If the required data center is not visible, perform a search (see [Section 1.2, "Search"](#)).

2. Select the required data center and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.

3. Select the check box of the appropriate user.

4. Click **Remove**. The user is removed from the Permissions tab. As this is hierarchical, the user will also be removed from the clusters, hosts and other objects.

2.2. Clusters

This section describes how to create, activate and manage host clusters in a data center.

A cluster is a collection of physical hosts that share the same storage domains and have the same type of CPU. Because virtual machines can be migrated across hosts in the same cluster, the cluster is the highest level at which power and load-sharing policies can be defined. The Red Hat Enterprise Virtualization platform contains a default cluster in the default data center at installation.

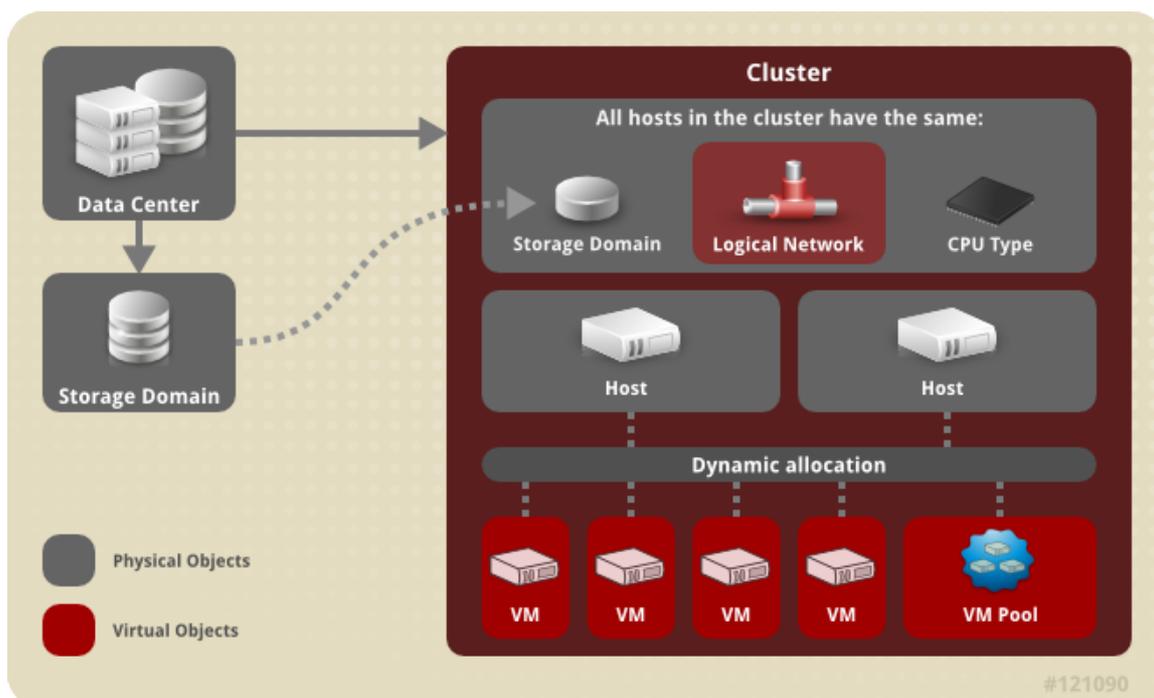


Figure 2.8. Cluster

Each cluster is a logical group that consists of multiple hosts. Every cluster in the system must belong to a data center, and every host in the system must belong to a cluster. This enables the system to dynamically allocate a virtual machine to any host in the applicable cluster, according to policies defined on the **Clusters** tab and in the Configuration tool during runtime, thus maximizing memory and disk space, as well as virtual machine uptime.

At any given time, after a virtual machine runs on a specific host in the cluster, the virtual machine can be migrated to another host in the cluster using **Migrate**. This can be very useful when a host is shut down for maintenance. The migration to another host in the cluster is transparent to the user, and the user continues working as usual. Note that a virtual machine cannot be migrated to a host outside the cluster.

2.2.1. Cluster Properties

Clusters can be defined, managed and viewed using the **Clusters** tab on the Red Hat Enterprise Virtualization Manager Administration Portal.

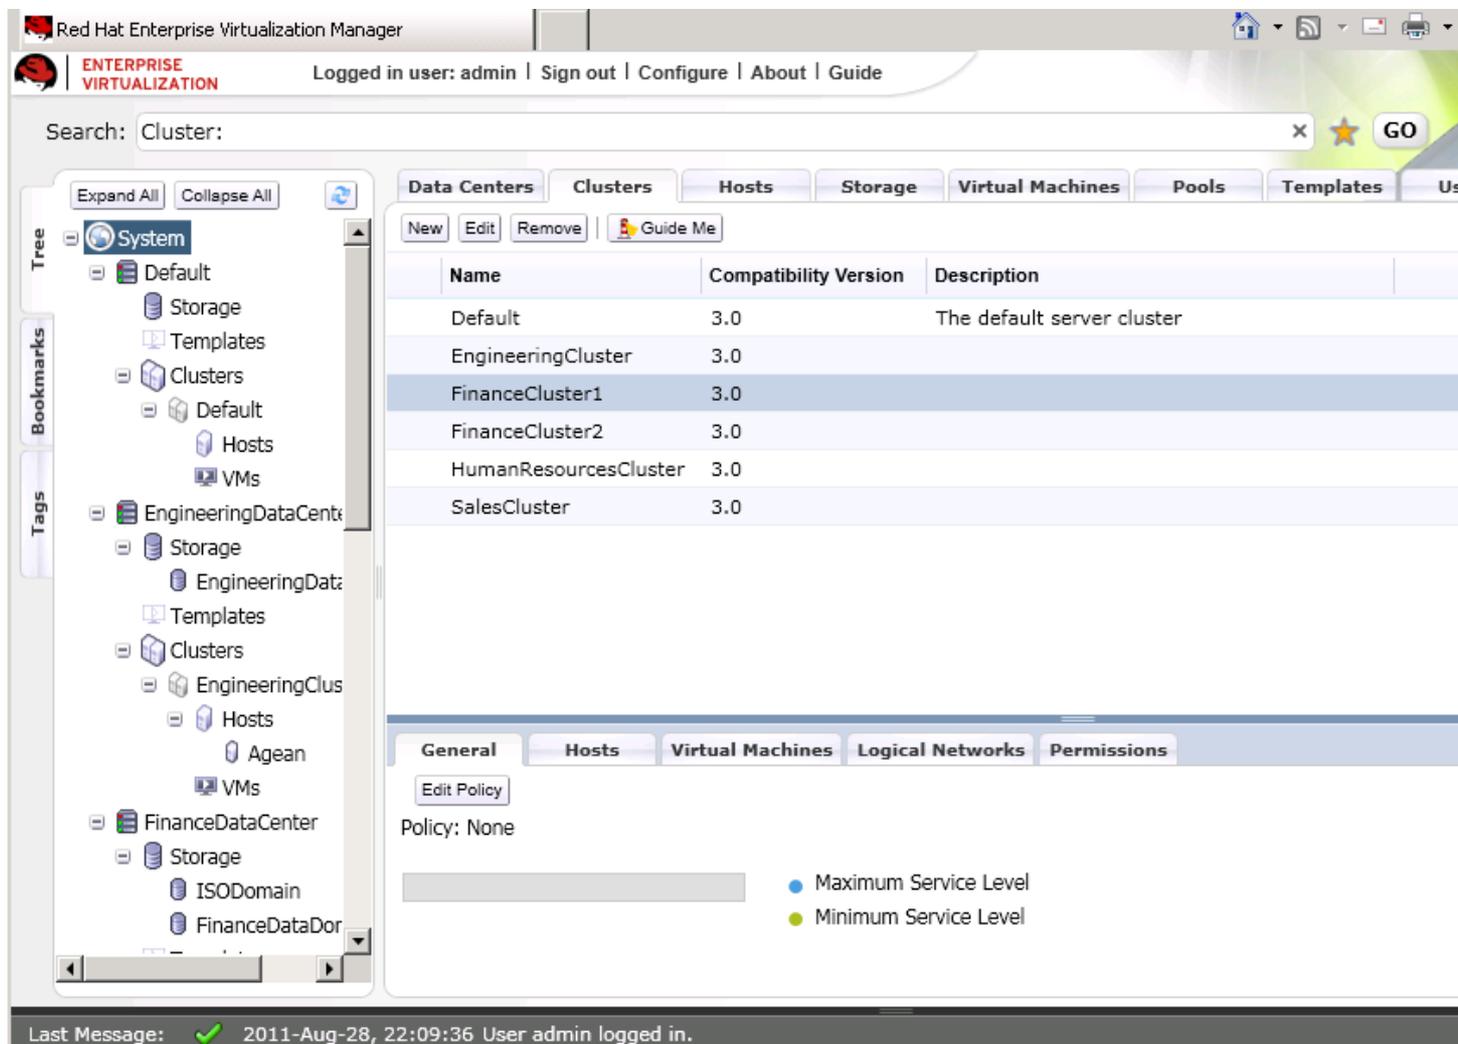


Figure 2.9. Clusters Tab

The table below describes the properties of a cluster as displayed in the **New Cluster** and **Edit Cluster** dialog boxes under **General**. Missing mandatory fields and invalid entries are outlined in red when you click **OK** to close the dialog box. In addition, field prompts indicate the expected values or range of values.

Table 2.6. Cluster Properties

Field/Tab	Description/Action
Data Center	The name of the data center. This must be a unique name with any combination of uppercase or lowercase letters, numbers, hyphens and underscores. Maximum length is 40.
Name	The name of the cluster. This must be a unique name with any combination of uppercase or lowercase letters, numbers, hyphens and underscores. Maximum length is 40.
Description	The description of the cluster. While this is not a mandatory field, it is recommended.
Type	Any one of, Intel or AMD. This cannot be changed after creation without significant disruption and high possibility of data loss. The type of data domain dictates the type of the data center.
Compatibility Level	The version of the Red Hat Enterprise Virtualization, either 2.2 or 3.0.

2.2.2. Cluster Operations

Creating a New Host Cluster

Before creating a new cluster, ensure that there is at least one host available to be assigned to it. The hosts in a cluster all run the same type of CPU.



Important

The default **rhev** network cannot be modified once a cluster has been attached to a data center. Any configuration required for the **rhev** network, such as enabling VLAN tagging, must be performed before a cluster is attached, and the data center is still in the **Uninitialized** state.

Perform the following steps to create a new cluster:

1. On the Tree pane, click the **+** sign next to the **System** tab. A list of data centers displays. Select the required **Data Center** from this list.
2. Open clusters by either expanding the required **Data Center** tab in the Tree pane and clicking on **Clusters** or select **Clusters** from the Details pane.
3. Click the **New** button. The **New Cluster** dialog box displays.
4. On the **General** tab, perform the following:
5. The second option in **New Cluster** window is **Memory Optimization**.
6. The third option in **New Cluster** window is **Resilience Policy**.
7. Click **OK** to create the cluster. The new host cluster is added to the data center and displays on the **Clusters** tab.

The **New Cluster - Guide Me** dialog box displays. For more information on this feature, see [Section 1.1.2, "Using the Guide Me facility"](#).

8. The **Guide Me** dialog box prompts you to add hosts to the new cluster. Click the **Configure Hosts** button, the **New Host** dialog box displays.

Enter the details of the host to assign to the cluster. Click **OK** to close the **New Host** dialog box and return to the **Clusters** tab. The **Hosts** tab on the Details pane displays the newly added hosts. Adding hosts is described in [Chapter 4, Red Hat Enterprise Virtualization Hosts](#).

2.2.3. Cluster Related Entities

To configure the cluster, the related objects that belong to the cluster also need to be set up. Clusters allow you to better control the scheduling of virtual machines and to perform load balancing.

Cluster General Entities

You can define load-balancing policies for each of your clusters.

Table 2.7. Cluster General Tab Properties

Field/Tab	Description/Action
Policy	The Edit Policy dialog box allows an administrator to set load balancing policy for a cluster to define the number of virtual machines that will be started on each host in the cluster. Set the Policy value to None to have no load or power sharing between hosts. This is the default mode.
Maximum Service Level	Use Even Distribution to evenly distribute the processing load across all hosts in the cluster. The host's CPU load is measured and used to apply the policy. Use the blue slider to specify the Maximum Service Level a host is permitted to have. A host that has reached the maximum service level defined will not have further virtual machines started on it. You can also specify the time interval in minutes that a host is permitted to run at the maximum service level before virtual machines are migrated off it. Selecting an even distribution load balancing policy causes virtual machines to be started on alternate hosts in the cluster to ensure that the CPU load on each host in the cluster is even.
Minimum Service Level	Use Power Saving to automatically reduce power consumption on under-utilized hosts. The green slider to specifies the Minimum Service Level a host is permitted to have. When a host has reached the minimum service level defined, virtual machines are migrated to other hosts. An administrator can then power down the hosts with low usage levels to conserve power. You can also specify the time interval (in minutes) that a host is permitted to run at the minimum service level before the migrations of any remaining virtual machines are triggered.

Cluster Hosts Entities

A cluster is a collection of physical hosts that share the same storage domains and have the same type of CPU. The host tab displays all the information related to all the hosts in a cluster.

Table 2.8. Cluster Hosts Tab Properties

Field/Tab	Description/Action
Name	The name of the host.
Host/IP	The Host/IP address.
Status	The status of the host.
Load	The number of virtual machines in each host.

Cluster Virtual Machines Entities

The clusters in a data center are the migration domains for virtual machines. The ability to migrate virtual machines from one host to another in a host is a key feature of virtualization. Virtual machines can be migrated across hosts in the same cluster hence, the cluster is the highest level at which power and load-sharing policies can be defined. When upgrading the Manager, the clusters need to be upgraded before the data center can work in the upgraded Manager. The Virtual Machines tab displays all the information related to all the virtual machines in a cluster.

Table 2.9. Cluster Virtual Machines Tab Properties

Field/Tab	Description/Action
Name	The name of the virtual machine belonging to a host.
Status	The status of all the virtual machine running on different hosts.
Host	The name of the host that a virtual machine is running on.
Uptime	The amount of time that a virtual machine has been up.

Cluster Logical Networks Entities

Logical networks are required both for the virtual machines to communicate to hosts and storage, and for the Manager to communicate with all entities. Logical networks need to be defined for each cluster.

Table 2.10. Cluster Logical Networks Tab Properties

Field/Tab	Description/Action
Name	The name of the logical networks in a cluster.
Status	The status of the logical networks.
Role	The hierarchical permissions available to the logical network.
Description	The description of the logical networks.

Cluster Permissions Entities

Cluster permissions define which users and roles can work in a cluster, and what operations the users and roles can perform.

Table 2.11. Cluster Permissions Tab Properties

Field/Tab	Description/Action
User	The user name of an existing user in the directory services.

Field/Tab	Description/Action
Role	The role of the user. The role comprises of user, permission level and object. Roles can be default or customized roles.
Inherited Permissions	The hierarchical permissions available to the user.

2.2.4. Cluster Troubleshooting

Configuring Cluster Policies

Defining the load-balancing or power sharing modes for hosts in the cluster is highly recommended. You can choose to set the policy on either load balancing or power saving, but not both.

To set load and power management policies for hosts in a cluster:

1. On the Tree pane click on the + sign next to the **System** tab. A list of Data Centers displays. Select the required **Data Center** from this list.
2. Open clusters by either expanding the required **Data Center** tab in the Tree pane and clicking on **Clusters** or select **Clusters** from the Results pane.
3. A list of clusters displays. Select the required cluster. The Details pane for the cluster displays.

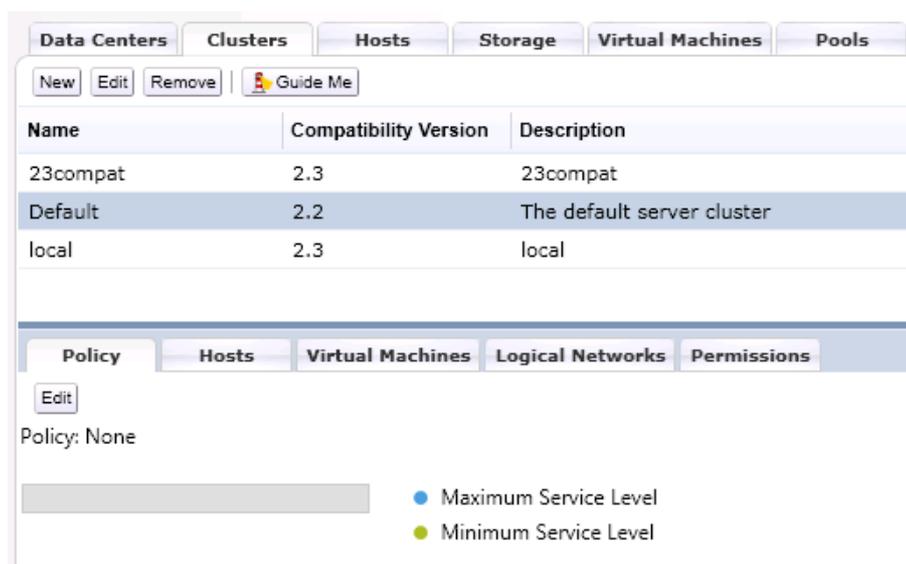


Figure 2.12. Cluster Policy Tab

4. On the Details pane, click the **Policy** tab. Click the **Edit** button. The **Edit Policy** dialog box displays, typically with the **None** option selected.

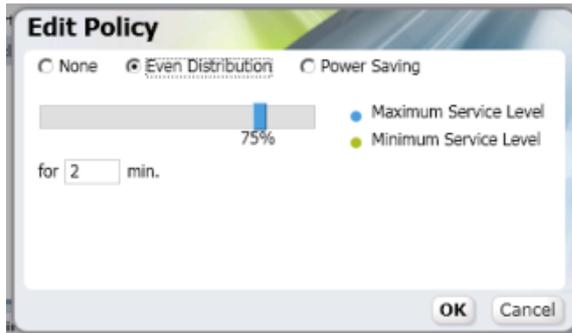


Figure 2.13. Edit Policy Dialog Box

5. Define the Load and Power Mode for hosts in the cluster. Select one of the following:
 - **None**; to have no load or power sharing between hosts. This is the default mode.
 - **Even Distribution**; to evenly distribute the processing load across all hosts in the cluster. The host's CPU load is measured and used to apply the policy. Use the blue slider to specify the **Maximum Service Level** a host is permitted to have. For example, a host that has reached the maximum service level defined will not have further virtual machines started on it. You can also specify the time interval in minutes that a host is permitted to run at the maximum service level before virtual machines are migrated off it.
 - **Power Sharing**; to distribute the amount of power consumed across all running hosts. Use the green slider to specify the **Minimum Service Level** a host is permitted to have. For example, a host that has reached the minimum service level defined virtual machines will be migrated to other hosts, enabling the hosts with low usage levels to be switched off to conserve power. You can also specify the time interval (in minutes) that a host is permitted to run at the minimum service level before a power down is triggered.
6. Click **OK** to define the policy for the cluster.

Maintaining a Cluster

You can edit cluster details, view hosts, virtual machines and logical networks, and add logical networks to a cluster. Logical Networking is described in [Section 2.3, "Logical Networks"](#).

To edit a cluster:

1. Click the **Clusters** tab.

A list of server clusters displays. Select the cluster that you want to edit.

2. Click the **Edit** button.

The **Edit Cluster** dialog box displays. The Edit Cluster dialog box is identical to the New Cluster dialog box. Modify the fields as described in [Creating a New Host Cluster](#)

3. Click **OK**.

The changes to the server cluster details are displayed in the list.

To view hosts in a cluster:

1. Click the **Clusters** tab. A list of server clusters displays. Select the appropriate cluster. The Details pane displays.
2. Click the **Hosts** tab. A list of hosts displays.

Name	Host/IP	Status	Load
▼ deathstar.englab.b		NonOperational	0 VMs
▲ fatima.englab.b		Up	4 VMs
▲ velankani.englab.b		Up	10 VMs

15, 10:56 Host cluster Shore was added by rhevadmin

Figure 2.14. The Hosts tab on the Cluster Details pane

To view virtual machines in a cluster:

1. Click the **Clusters** tab.

A list of clusters displays. Select the appropriate cluster. The Details pane displays.

2. Click the **Virtual Machines** tab.

A list of virtual machines displays. This includes both virtual servers and virtual desktops.

Removing a Cluster

Clusters that are not in use can be permanently removed. Deleting unused clusters saves system resources, as existing hosts are checked(or pinged) at fixed intervals.



Warning

It is recommended that the default cluster should not be removed.

To remove a cluster:

1. Click the **Cluster** tab.
2. If the required cluster is not visible, perform a search (see [Section 1.2, "Search"](#)).
3. Select the cluster to be removed. Ensure that there are no running hosts.
4. Click the **Remove** button.

A message prompts you to confirm removal. The dialog box lists all the clusters that are selected for removal.

5. Click **OK**.

The cluster is deleted and disappears from the **Clusters** tab.



Note

The hosts and storage domains can still be used, and allocated to a different cluster.

2.2.5. Cluster Permissions

Managing System Permissions for a Cluster

While the system administrator of the data center has the full range of permissions, a cluster administrator is a system administration role for a specific cluster only. This is a hierarchical model, and means that if a user is assigned the cluster administrator role for a cluster, all objects in the cluster can be managed by the user. The cluster administrator role permits the following actions:

- Creation and removal of specific clusters.
- Addition and removal of hosts, virtual machines, pools.
- Permission to attach users to virtual machines within a single cluster.

This is useful in a data center where there are multiple clusters, each of which require their own system administrators. A cluster administrator has permissions for the assigned cluster only, not for all clusters in the data center.

To assign a system administrator role to a cluster:

1. Click the **Clusters** tab.

A list of clusters displays. If the required cluster is not visible, perform a search (see [Section 1.2, “Search”](#)).

2. Select the cluster that you want to edit, and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and Inherited permissions, if any. This can include the Data Center Administrator.

3. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** textbox, and click **Go**. A list of possible matches display in the results list.
4. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **ClusterAdmin**.
5. Click **OK**.

The name of the user displays in the **Permissions** tab, with an icon and the assigned Role.



Note

You can only assign roles and permissions to existing users. See [Chapter 5, Users](#).

You can also change the system administrator of a cluster, by removing the existing system administrator, and adding the new system administrator, as described in the previous procedure.

To remove a system administrator role:

1. Click the **Clusters** tab. A list of clusters displays. If the required cluster is not visible, perform a search (see [Section 1.2, “Search”](#)).

2. Select the required cluster and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.

3. Select the check box of the appropriate user.
4. Click **Remove**. The user is removed from the Permissions tab. As this is hierarchical, the user will also be removed from the hosts and other objects.

2.3. Logical Networks

A Logical Network is a way of representing networks in your data center that have the same connectivity properties. The logical network in a data center ensures that the hosts, manager and storage remain connected. When the data center is created, the management network is defined by default. However new logical networks, for example for data, storage, or display, can be defined by the administrator. In general, logical networks are assigned by functionality and physical topology. New networks, for example for data, storage, or display, can be added to enhance network speed and performance. In addition, other networks can be used to segregate virtual machine traffic from the management networks, or isolate traffic between groups of virtual machines in the same cluster.

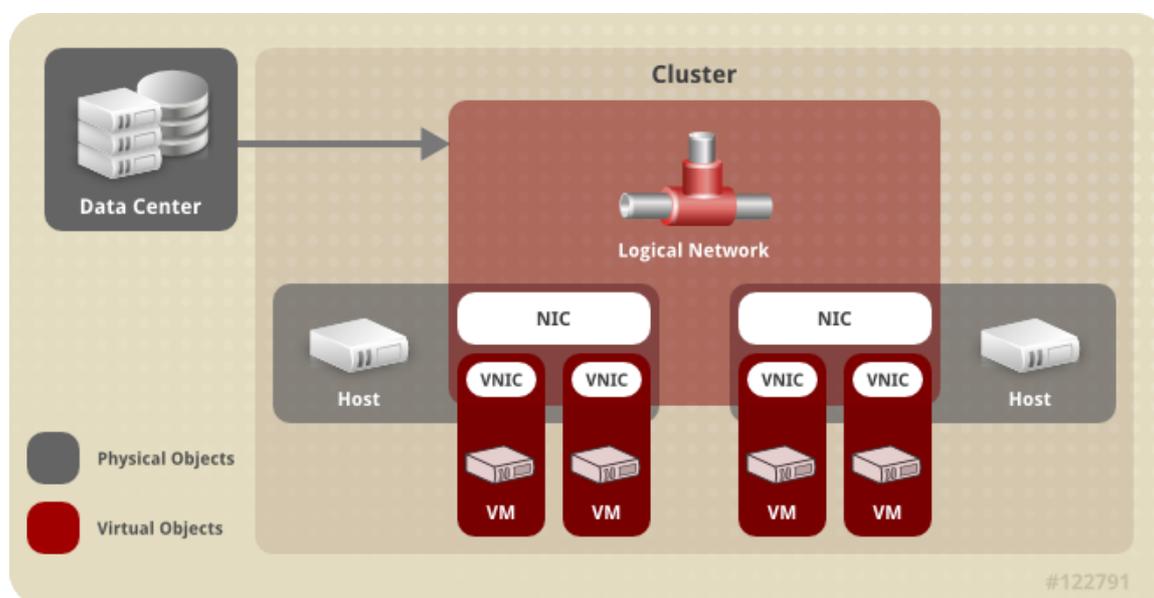


Figure 2.15. Data Center Objects

For example, a data center may have the following networks,

- Guest data network
- Storage network access
- Management network
- Display network (for SPICE or VNC)



Warning

Do not change networking in a data center or a cluster if any hosts are running. This may make the host unreachable.

2.3.1. Logical Networks Properties

A logical network is assigned as a required resource of a cluster in a data center, and by extension all hosts in a cluster must have the same set of logical networks. The implementation itself may vary from host to host (IP and bonding properties). Therefore, to configure a network, define the logical network and then apply this network to the NIC on each host. By default the management network (rhevmm) is defined for a data center. Red Hat Enterprise Virtualization platform allows you to use VLAN ID tagging and supports STP in logical networks

Table 2.12. Logical Network Properties

Field/Tab	Description/Action
Name	The name of the network belonging to the data center.
Description	The description of the network.
STP Support, VLAN Tagging	STP support and VLAN tagging are optional attributes.

2.3.2. Logical Network Operations

A logical network is assigned as a required resource of a cluster in a data center, and by extension all hosts in a cluster must have the same set of logical networks implemented. A Logical Network is an abstract entity that has to be assigned to a specific NIC on the HOST. All hosts in a cluster must be assigned an operational NIC in the logical network to be operational. The implementation itself may vary from host to host (for example, the IP and bonding properties); therefore, to configure a network, you need to first define the network and then apply this network to each hosts. You can create and add new logical network either to the data center, or to a specific cluster, assign logical networks to clusters and edit and maintain networks.

To create a new logical network in a data center or cluster

1. Navigate to the **Tree** pane and click the required data center or cluster. On the Details pane, select the **Logical Networks** tab. This displays the existing logical networks. At the minimum, the default **rhevmm** network is listed.
2. Click **New** (or **Add Network** on the **Clusters** Detail pane). The **New Logical Network** dialog box displays. Enter the **Name** and **Description**, and select the **Assign Networks to Cluster(s)** check box to add the storage network to a selected cluster in the data center.
3. Click **OK** to create the new logical network. The new logical network is created and displays on the **Logical Networks** tab with a Non Operational status. The next step is to add the logical network to each host in cluster, to match the NIC on the host to the network. After this is done, the network will become operational.

To assign or detach a logical networks in a cluster

1. Navigate to the **Tree** pane and click the required cluster. On the Details pane, select the **Logical Networks** tab. This displays the existing logical networks. At the minimum, the default **rhev** network is listed.
2. Click the **Assign/Detach** button. The **Assign/Detach Networks** dialog box displays a list of the available networks. Select the logical network or networks and click **OK**. The assigned logical network displays on the **Logical Networks** tab.

To add or edit a logical network to hosts in a cluster

1. From the **Tree** pane, select the required cluster. The **Hosts** tab displays a list of available hosts.
2. For each of your installed hosts, place the host into Maintenance mode and on the Details pane, select the **Network Interfaces** tab.
3. A list of network interfaces available for this host displays. One of them will already have the management network (rhev) configured.
4. Select the network interface to attach the newly added network and click the **Edit/ Add** button. The **Edit Network Interface** dialog box displays.

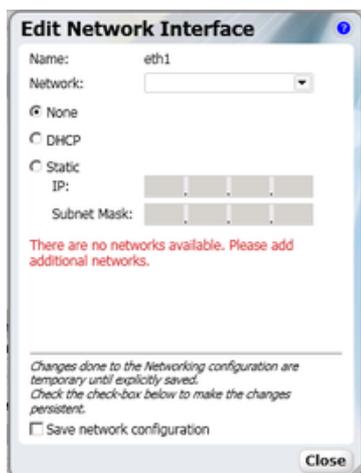


Figure 2.16. Edit network interface

Enter the new (or additional) **Network Name**. Select one of **None**, **DHCP**, **Static** radio buttons.

For **Static** enter the **IP** and **Subnet Mask** and **Default Gateway**.

5. Select the **Check Connectivity** check box to run a network check, provided the host is in Maintenance mode.
6. Select the **Save network configuration** check box and click **Close**.

To use a logical network as a display network for SPICE:

1. Click the **Clusters** tab. The list of clusters displays. Select the appropriate cluster. Click the **Logical Networks** tab in the Details pane.
2. Select the network to be used as the display network for SPICE. For more information on the SPICE protocol, see [Section 1.1.3, "About SPICE"](#).
3. Click the **Set as Display** button, and click **OK**. The role of the network appears as **Display** in the pane. The selected network will be used for SPICE/vnc traffic.

Use Logical Networks to Add Multiple VLANs to a Single Network Interface

Multiple VLANs can be added to a single interface if that interface *does not* carry the **rhev** logical network.

1. Navigate to the **Tree** pane and click the required data center or cluster. On the Details pane of the selected data center or cluster, select the **Logical Networks** tab.
2. Click **New** (or **Add Network** on the **Clusters** Detail pane). The **New Logical Network** dialog box displays. Enter the **Name** and **Description**. Select the **Enable VLAN tagging** check box, and enter the desired VLAN tag to add to traffic on this logical network in the provided text field. Select the **Assign Networks to Cluster(s)** check box to add the storage network to a selected cluster in the data center.
3. Click **OK** to create the new logical network. The new logical network is created and displays on the **Logical Networks** tab with a Non Operational status.
4. Select the **Hosts** tab, and one of the hosts in the cluster to which your VLAN tagged logical network was assigned. Click the **Network Interfaces** tab in the details pane. Select a network interface that does not carry the **rhev** logical network, and click the **Add/Edit** button. The **Edit Network Interface** displays.
5. Select the newly created, VLAN tagged **Network Name** from the drop down menu. Select one of **None**, **DHCP**, **Static** radio buttons.

For **Static** enter the **IP** and **Subnet Mask** and **Default Gateway**.

6. Select the **Check Connectivity** check box to run a network check, provided the host is in Maintenance mode.
7. Select the **Save network configuration** check box and click **OK**.
8. The **Network Interfaces** tab of the details pane now shows VLAN information for the edited network interface. In the **Vlan** column newly created VLAN devices are shown, with names based on the network interface name and VLAN tag. For example, if you have edited **eth1** to carry traffic for VLAN 43, there will be a device called **eth1.43** in the same row as the **eth1** device.

	Name	Address	MAC	Speed (Mbps)	Rx (Mbps)	Tx (Mbps)	Drops (Pkts)	Bond	Vlan
<input checked="" type="checkbox"/>	eth1		B8:AC:6F:BC:E1:13	1000	< 1	< 1	0		eth1.43
<input type="checkbox"/>	eth0	10.64.15.241	00:1B:21:6D:70:D0	1000	< 1	< 1	0		

Figure 2.17. Multiple VLANs on One Network Interface

9. Next, 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.

2.3.3. Logical Networks Troubleshooting

When checking for errors, ensure that every host in the cluster is connected to the logical networks, all hosts in the cluster must have the same network configuration. Each cluster may have a different set of logical networks but all the logical networks must exist in the data center definition. Ensure that all hosts are in maintenance before adding or editing networks to the cluster.



Warning

Do not change networking in a data center or a cluster if any hosts are running. This may make the host unreachable.

To manage logical networks in a cluster

1. Click the **Clusters** tab. The list of clusters displays.
2. Select the appropriate cluster. Click the **Logical Networks** tab in the Details pane. A list of available networks displays. You can now add, assign, detach networks, or set one network as the display network. These operations are described earlier in this section.

To edit a logical network

1. Ensure that all hosts in the data center are in Maintenance mode and navigate to the **Tree** pane and click the required data center. On the Details pane, select the **Logical Networks** tab.
2. Click **Edit**. The **Edit Logical Network** dialog box displays. Edit the **Name** and **Description**, and select the **Assign Networks to Cluster(s)** check box to add the Storage network to a selected cluster in the data center.
3. Click **OK** to save the changes.

2.3.4. Logical Networks System Permissions

While the superuser or system administrator of the platform has the full range of permissions, a Network Administrator is a system administration role for the network of a cluster only. This is a hierarchical model, and means that a user with the Data Center Administrator role, or a Cluster Administrator role for a data center, also has network permissions. The Network Administrator role permits the configuration of the network in a cluster, and in all hosts in the cluster. This is useful in an enterprise where there is a lot of network traffic that needs to be managed in an optimal fashion.

To assign a network administrator role in a cluster

1. Click the **Cluster** tab.

A list of clusters displays. If the required cluster is not visible, perform a search (see [Section 1.2, "Search"](#)).
2. Select the cluster that you want to edit, and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.

Policy Hosts Virtual Machines Logical Networks Permissions			
Add Remove			
User	Role	Inherited Permission	
 Susan Burgess	SuperUser	(System)	
 David Jorm	RHEVMUser	(System)	
 rhevadmin	SuperUser	(System)	

Figure 2.18. Cluster Permissions

- Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **NetworkAdmin**.

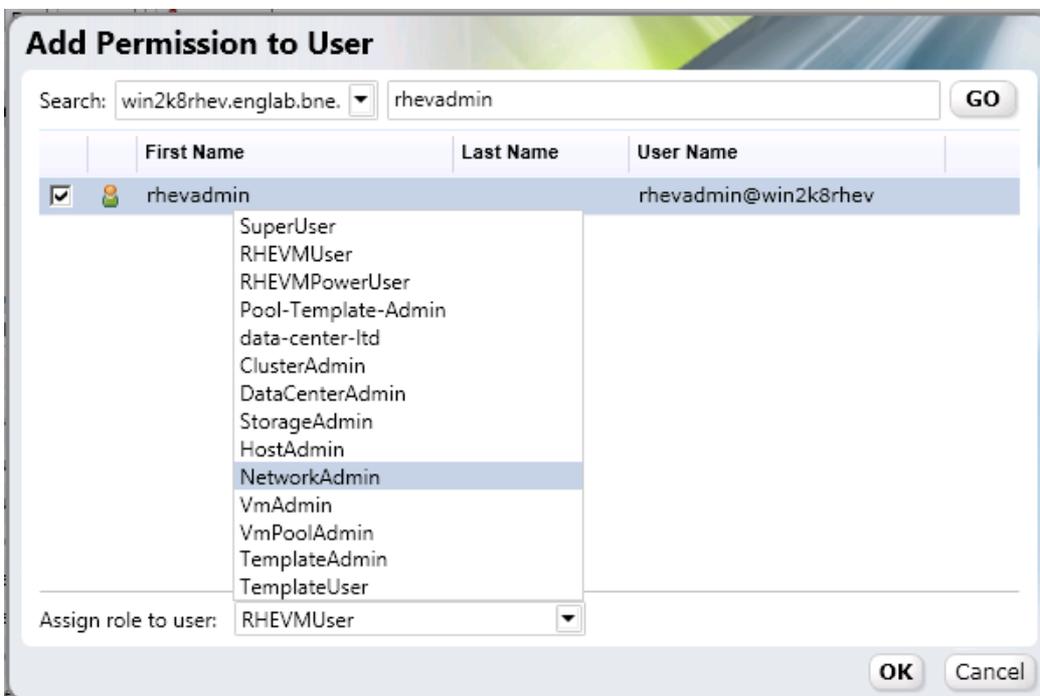


Figure 2.19. Assign Network Admin Permissions

- Click **OK**.

The name of the user displays in the **Permissions** tab, with an icon and the assigned Role.

Note

You can only assign roles and permissions to existing users. See (see [Chapter 5, Users](#)).

You can also change the network administrator of a cluster, by removing the existing administrator, and adding the new administrator, as described in the previous procedure.

To remove a network administrator role:

1. Click the **Clusters** tab. A list of clusters displays. If the required cluster is not visible, perform a search (see [Section 1.2, "Search"](#)).

2. Select the required cluster and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.

3. Select the check box of the appropriate user.
4. Click **Remove**. The user is removed from the Permissions tab. As this is hierarchical, the user will also be removed from the clusters, hosts and other objects.

Storage

Red Hat Enterprise Virtualization uses a centralized storage system for virtual machine disk images, ISO files and snapshots. Storage networking can be implemented using Network File System (NFS), Internet Small Computer System Interface (iSCSI) or Fiber Channel Protocol (FCP). You can also set up local storage on a host as a data domain if you are setting up a very small and limited environment, for example, as an evaluation. This section describes how to set up and manage the variety of storage types that can be used in Red Hat Enterprise Virtualization. Setting up storage is a vital prerequisite for a new data center, which cannot be initialized until storage domains are attached and activated. In addition to the administrator of the platform, and the data center, a storage administrator role can be assigned to a user who will have system permissions to a particular storage domain.

3.1. Managing Storage

A Red Hat Enterprise Virtualization system administrator needs to create, configure, attach and maintain storage for the virtualized enterprise. A familiarity with the storage types and their use is highly recommended. This document does not describe the concepts, protocols, requirements or general usage of NFS, iSCSI or FCP.

The Administration Portal enables administrators to assign and manage storage effectively and efficiently. The **Storage** tab on the Administration Portal provides an efficient graphical way to view and manage networked storage. The Storage Results list displays all the storage domains, and the Details pane enables access to general information about the domain.

The Administration Portal has three types of domains:

- Data domains hold the disk images of all the virtual machines running in the system, operating system images and data disks. In addition, snapshots of the virtual machines are also stored in the data domain. The data cannot be shared across data centers, and the data domain must be of the same type as the data center. For example, a data center of a iSCSI type, must have an iSCSI data domain. A data domain cannot be shared between data centers. A local storage domain can only be a data domain.
- ISO domains store ISO files (or logical CDs) used to install and boot operating systems and applications for the virtual machines. Because an ISO domain is a logical entity replacing a library of physical CDs or DVDs, an ISO domain removes the data center's need for physical media. An ISO domain can be shared across different data centers.
- An export domain is a temporary storage repository that is used to copy/move images between data centers and Red Hat Enterprise Virtualization Manager installations. In addition, the export domain 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.



Important — Export Domain Storage Type

Support for export storage domains backed by storage on anything other than NFS is being deprecated. Existing export storage domains imported from Red Hat Enterprise Virtualization 2.2 environments remain supported. New export storage domains must be created on NFS storage.

3.1.1. Storage Properties

Understanding Storage Domains

Setting up, managing and monitoring storage is essential for a data center to function efficiently at all times. A storage domain is a collection of images that have a common storage interface. A storage domain contains complete images of the virtual machines including templates and snapshots. A storage domain can be either for block devices (SAN - iSCSI or FCP) or files (NAS - NFS). On NFS, all virtual disks, templates and snapshots are simple files. On SAN (iSCSI/FCP), the LUNs are aggregated into a logical entity called a Volume Group (VG). This is done via LVM (Logical Volume Manager) See *Red Hat Enterprise Linux Logical Volume Manager Administration Guide* for more information on LVM. Each virtual disk, template or snapshot is a Logical Volume (LV) on the VG.

Virtual disks can have one of two formats, either Qcow2 or Raw. The type of storage can be either Sparse or Preallocated. Snapshots are always sparse but can be taken for disks created either as raw or sparse.

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

3.1.2. Storage Operations

3.1.2.1. Adding Storage Domains to a Data Center

Use the **Storage** tab to complete the following tasks:

- Add or edit storage domains.
- Activate, deactivate, or detach a storage domain from a data center.
- Maintain and delete storage domains

This section describes how to add a storage domain to the system. The next section describes how to configure the storage for Red Hat Enterprise Virtualization.

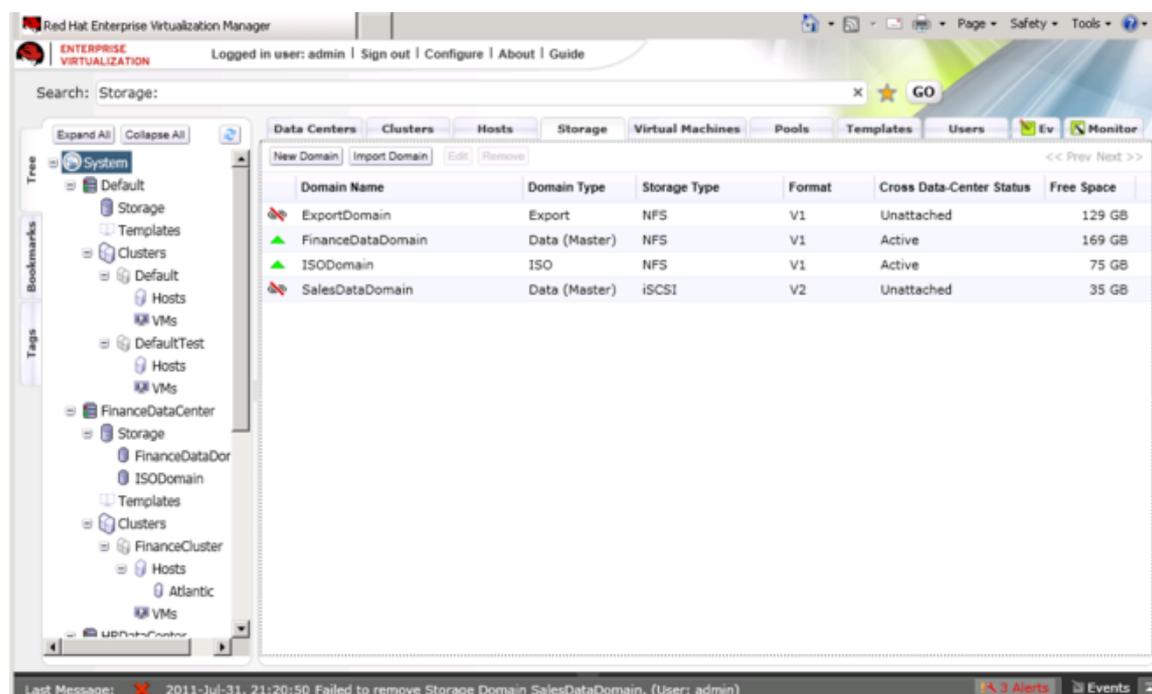


Figure 3.1. The Storage Tab

There are two ways of adding storage domains to the Administration Portal. You can set up and add a new storage domain or you can import an existing ISO or export domain from another installation of Red Hat Enterprise Virtualization Manager.

While any available host in the data center can be used to add or configure a storage domain (using the **Use Host** field), all storage domains defined in the data center must be reachable by all the hosts in the data center. If a host is unable to access a storage domain that host is likely to become nonoperational. Therefore, when adding new storage domains to an active cluster, ensure that the storage is reachable from all hosts.

!
Note

If an ISO storage domain is required, it must be added after at least one data storage domain has been added.

1. Click the **Storage** tab. The Storage list and toolbar displays.
2. Click **New Domain**. The New Domain dialog box displays.

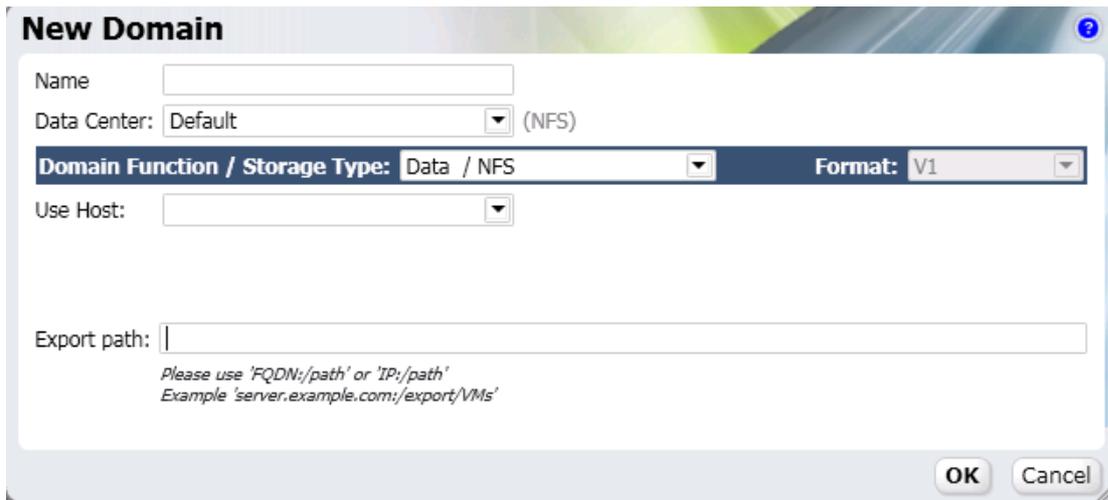


Figure 3.2. Adding New Storage

3. Enter the **Name** of the storage domain, for example, accounting-server-images. A descriptive name is recommended.
4. Select the appropriate **Data Center**. Select a **Data Center** from the drop-down list that displays all the **Data Centers** available and their storage types.
5. Select the appropriate **Domain Function / Storage Type**. Select one:
 - Data / NFS
 - ISO
 - Export

The **Domain Function / Storage Type** determines the availability of the **Format** field. Refer [Table 6.3, “Permitted Storage Combinations”](#) for more information on **Format**.

6. Select a host in **Use host**. To attach a domain, an active host must be selected.



Important

All communication to the storage domain is via the selected host and not from the Red Hat Enterprise Virtualization Manager. At least one host must be active and have access to the storage before the storage can be configured.

7. Enter the **Export path** of the storage. The export path can be either a static IP address or a resolvable hostname. For example, 192.168.0.10:/images/ISO or storage.demo.redhat.com:/exports/iso.
8. Click **OK**.
9. The storage domain displays on the **Storage** tab.

To import an existing ISO or Export storage domain:

1. Click the **Storage** tab. The Storage list and toolbar display. Refer [Figure 3.1, “The Storage Tab”](#).

- Click **Import Domain**. The New Domain dialog box displays.

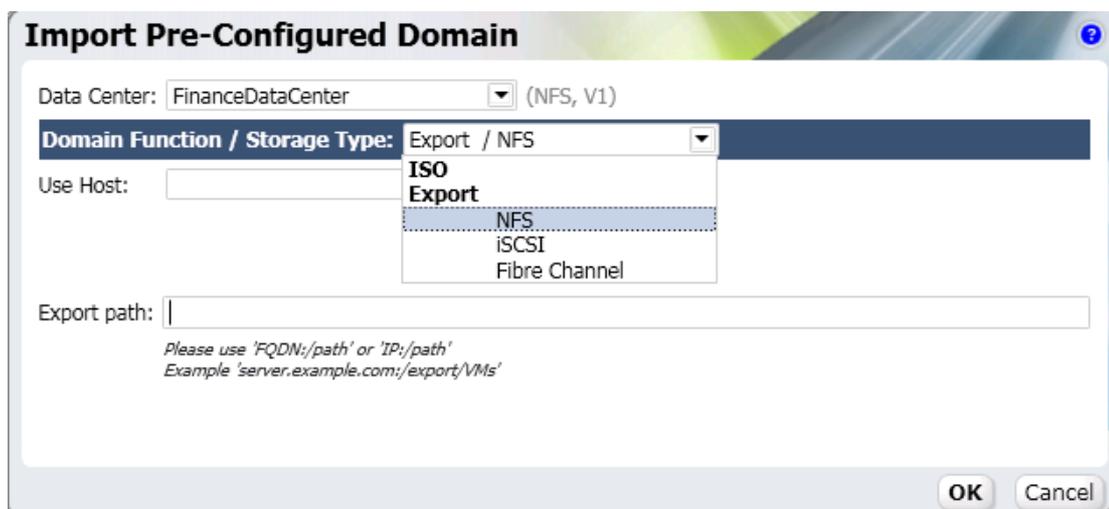


Figure 3.3. Import Domain

- Select the appropriate **Domain Function / Storage Type**. Select one:
 - Data / NFS
 - ISO
 - Export

The **Domain Function / Storage Type** determines the availability of the **Format** field. Refer [Table 6.3, “Permitted Storage Combinations”](#) for more information on **Format**.

- Select a host in **Use host**. To attach a domain, an active host must be selected.



Important

All communication to the storage domain is via the selected host and not from the Red Hat Enterprise Virtualization Manager. At least one host must be active and have access to the storage before the storage can be configured.

- Enter the **Export path** of the storage. The export path can be either a static **IP address** or a resolvable hostname. For example,

192.168.0.10:/Images/ISO

or

storage.demo.redhat.com:/exports/iso

- Click **OK**.

7. The storage domain is imported and displays on the **Storage** tab. The next step is to attach it to a data center. This is described later in this chapter, [Section 3.1.2.2, “Attaching Storage Domains to a Data Center”](#).

3.1.2.1.1. Adding NFS Storage

An NFS storage domain is an NFS file share that is attached to a data center. Once you attach an NFS file share to the data center as a storage domain it is used to provide storage to the Red Hat Enterprise Virtualization environment. How the storage domain is used depends on the function you select when attaching it.

This section details how to prepare NFS file shares on your storage infrastructure and attach them using the Red Hat Enterprise Virtualization Manager. For further information on NFS itself, see the *Red Hat Enterprise Linux — Storage Administration Guide*

New Domain

Name:

Data Center: (NFS, V1)

Domain Function / Storage Type: Format:

Use Host:

Export path:

*Please use 'FQDN:/path' or 'IP:/path'
Example 'server.example.com:/export/VMs'*

Figure 3.4. Add NFS

 **Important — NFSv4 is not Natively Supported**

NFSv4 is not natively supported by Red Hat Enterprise Virtualization. Red Hat Enterprise Virtualization will always attempt to mount NFS storage using NFSv3.

Your NFS storage server must support NFSv3 to be used with Red Hat Enterprise Virtualization. Attempts to attach NFS storage which has been exported from servers that do not support NFSv3 to the Red Hat Enterprise Virtualization environment will fail.

Preparing NFS Storage

This section outlines how to prepare an NFS file share on a server running Red Hat Enterprise Linux 6. Once created the NFS share can be attached by the Red Hat Enterprise Virtualization Manager.

1. **Install *nfs-utils***

NFS functionality is provided by the *nfs-utils* package. Before file shares can be created, check that the package is installed by querying the RPM database for the system:

```
$ rpm -qi nfs-utils
```

If the *nfs-utils* package is installed then the package information will be displayed. If no output is displayed then the package is not currently installed. Install it using **yum** while logged in as the root user:

```
# yum install nfs-utils
```

2. Configure Boot Scripts

To ensure that NFS shares are always available when the system is operational both the *nfs* and *rpcbind* services must start at boot time. Use the **chkconfig** command while logged in as root to modify the boot scripts.

```
# chkconfig --add rpcbind
# chkconfig --add nfs
# chkconfig rpcbind on
# chkconfig nfs on
```

Once the boot script configuration has been done, start the services for the first time.

```
# service rpcbind start
# service nfs start
```

3. Create Directory

Create the directory you wish to share using NFS.

```
# mkdir /exports/iso
```

Replace */exports/iso* with the name, and path of the directory you wish to use.

4. Export Directory

To be accessible over the network using NFS the directory must be exported. NFS exports are controlled using the */etc/exports* configuration file. Each export path appears on a separate line followed by a tab character and any additional NFS options. Exports to be attached to the Red Hat Enterprise Virtualization Manager must have the read, and write, options set.

To grant read, and write access to */exports/iso* using NFS for example you add the following line to the */etc/exports* file.

```
/exports/iso    *(rw)
```

Again, replace */exports/iso* with the name, and path of the directory you wish to use.

5. Reload NFS Configuration

For the changes to the */etc/exports* file to take effect the service must be told to reload the configuration. To force the service to reload the configuration run the following command as root:

```
# service nfs reload
```

6. Set Permissions

The NFS export directory must be configured for read write access and must be owned by vdsm:kvm. If these users do not exist on your external NFS server use the following command, assuming that `/exports/iso` is the directory to be used as an NFS share.

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

The permissions on the directory must be set to allow read and write access to both the owner and the group. The owner should also have execute access to the directory. The permissions are set using the `chmod` command. The following command arguments set the required permissions on the `/exports/iso` directory.

```
# chmod 0755 /exports/iso
```

Result:

The NFS file share has been created, and is ready to be attached by the Red Hat Enterprise Virtualization Manager.

Attaching NFS Storage

An NFS type **Storage Domain** is a mounted NFS share that is attached to a data center. It is used to provide storage for virtualized guest images and ISO boot media. Once NFS storage has been exported it must be attached to the Red Hat Enterprise Virtualization Manager, using the Administration Portal.

To add an NFS data, or export, storage domain you must select an NFS data center. NFS storage domains for ISO storage are able to be added to data centers of any type.

1. Click the **Storage** tab. The Storage list and toolbar display.
2. Click **New Domain**.
3. The **New Storage** dialog box displays.

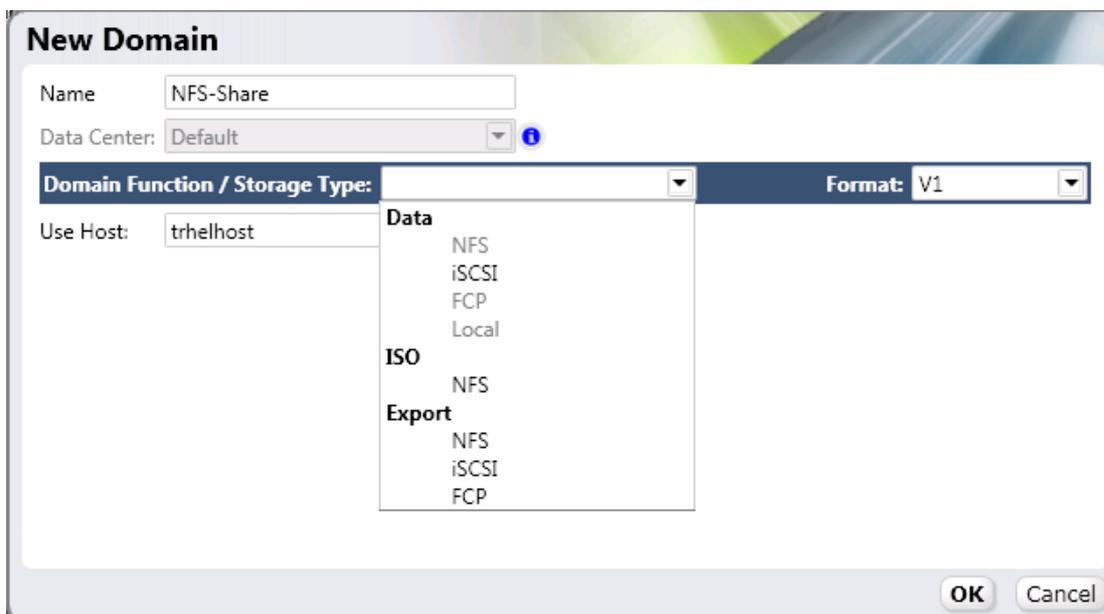


Figure 3.5. NFS Storage

- Configure the following options:

Name: Enter a suitably descriptive name.

Data Center: Select the required Data Center from the drop-down list.

Domain Function/ Storage Type: In the drop down menu, select Data → NFS. The storage domain types which are not compatible with the Default data center are grayed out. After you select your domain type, the Export Path field appears.

Export path: Enter the IP address or a resolvable hostname of the chosen host. The export path should be in the format of **192.168.0.10:/Images/ISO** or **domain.example.com:/Images/ISO**

Use Host: Select any of the hosts from the drop down menu. Only hosts which belong in the pre-selected data center will display in this list.



Active Host Required

All communication to the storage domain is via the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must exist in the system, and be attached to the chosen data center, before the storage is configured.

- Click **OK**.

Result:

The new NFS data domain displays on the Storage tab. It will remain with a *Locked status* while it is being prepared for use. When ready, it is automatically attached to the data center.

3.1.2.1.2. Adding iSCSI Storage

Red Hat Enterprise Virtualization platform supports iSCSI storage via the creation of a Storage Domain for a Volume Group. A Volume Group is a set of pre-defined Logical Unit Numbers (LUNs). Red Hat Enterprise Virtualization supports creation of a Storage Domain from a pre-existent Volume Group or a set of LUNs. Neither Volume Groups nor LUNs are able to be attached to more than one Storage Domain at a time.

For information regarding the setup and configuration of iSCSI on Red Hat Enterprise Linux, see the *Red Hat Enterprise Linux — Storage Administration Guide*.

- On the tree pane, select the **Tree** tab. On **System**, click the + icon to display the available data centers.
- Select the **Data Center** to which the domain is to be added. The storage type of the data center selected determines the type of storage domains that can be added to it. To add an iSCSI data, or export, storage domain you must select an iSCSI data center. iSCSI storage domains can not be used for ISO storage domains.
- Click the **New Domain** button.
- Click **New Storage**. The **New Storage** dialog box displays.
- From the **Domain Function / Storage Type** drop-down menu, select the appropriate storage type for the storage domain. 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. To attach a domain, the name of an active host must be selected from the list of existing hosts. Only hosts that are attached to the selected Data Center are listed.



Active Host Required

All communication to the storage domain is via the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must exist in the system, and be attached to the chosen data center, before the storage is configured.

- 7. The Red Hat Enterprise Virtualization Manager is able to map either iSCSI targets to LUNs, or LUNs to iSCSI targets. The **New Domain** dialog 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.

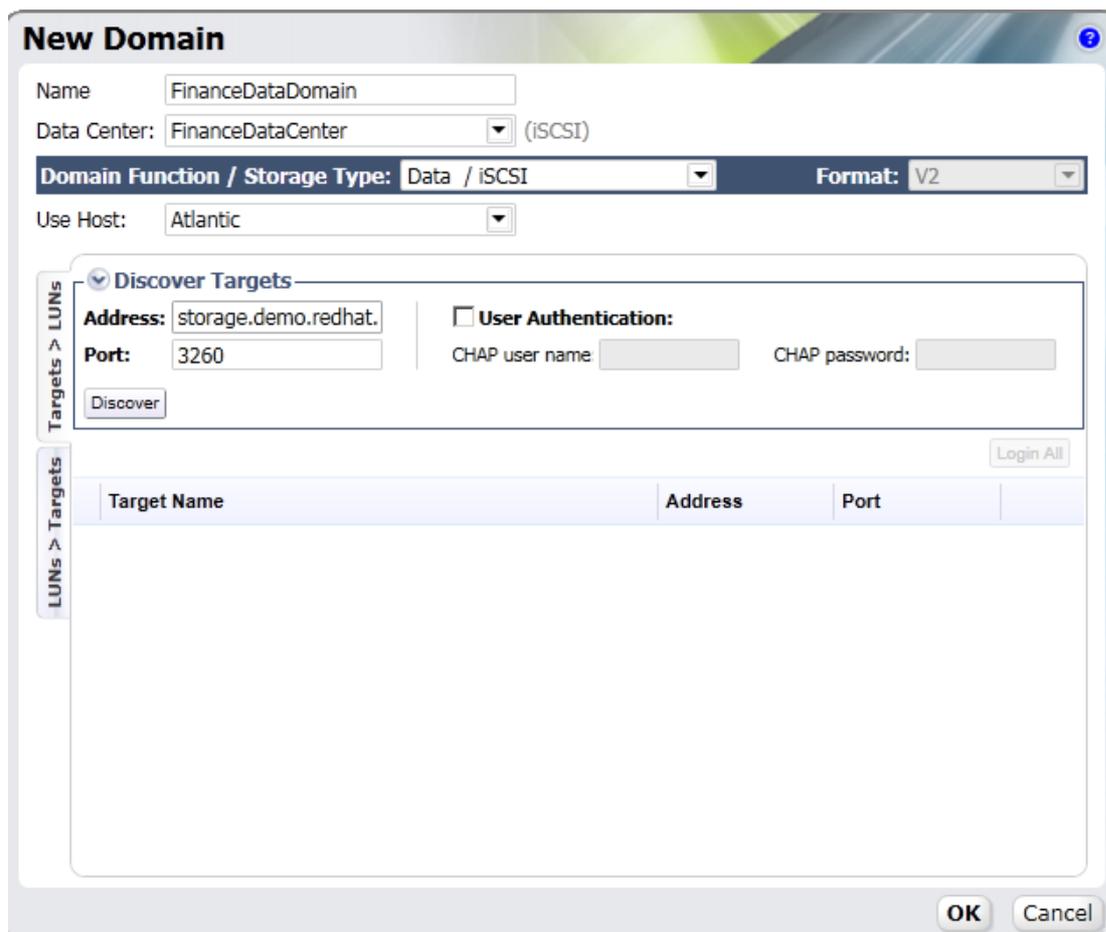


Figure 3.6. New Domain Dialog

iSCSI Target Discovery

- Click **Discover Targets** to enable target discovery options. The **New Domain** dialog automatically displays targets with unused LUNs when iSCSI is selected as the storage type. If the target that you are adding is not listed, click **Discover Targets** to enable target discovery options.

- Enter the fully qualified domain name or IP address of the iSCSI host in the **Address** field.
 - Enter the port to connect to the host on when browsing for targets in the **Port** field. The default is **3260**.
 - 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 password.
 - Click the **Discover** button.
8. Click the **+** button next to the desired target. This will expand the entry and display all unused LUNs attached to the target.

New Domain

Name: FinanceDataDomain
 Data Center: FinanceDataCenter (iSCSI)
 Domain Function / Storage Type: Data / iSCSI Format: V2
 Use Host: Atlantic

Discover Targets

Address: storage.demo.redhat. User Authentication:
 Port: 3260 CHAP user name: CHAP password:

Target Name	Address	Port	
[-] iqn.2011.com.redhat.demo.storage:finance	10.64.14.30	3260	<input type="button" value="Login"/>

LUN ID	Dev.	Size	#path	Vendor ID	Product ID	Serial
<input checked="" type="checkbox"/> 11ET_00010001		40GB	1	IET	VIRTUAL-D:	SIET_VIRT
<input checked="" type="checkbox"/> 11ET_00010002		40GB	1	IET	VIRTUAL-D:	SIET_VIRT

Figure 3.7. iSCSI LUN Selection

9. Select the check box for each LUN that you are using to create the storage domain.
10. Click **OK** to create the storage domain.

Result:

The new iSCSI storage domain displays on the storage tab. This will take some time.

3.1.2.1.2.1. Mapping iSCSI Targets to LUNs

Follow the below mentioned procedure:

1. Click the **+** button next to the desired target.

2. Select the check box for each LUN that you are using to create the storage domain.
3. Click **OK**.

Result:

The new storage domain is created.

3.1.2.1.3. Adding FCP Storage

Red Hat Enterprise Virtualization platform supports SAN storage via the creation of a Storage Domain for a Volume Group. A Volume Group is a set of pre-defined Logical Unit Numbers (LUNs). Red Hat Enterprise Virtualization supports creation of a Storage Domain from a pre-defined Volume Group or a set of LUNs. Neither Volume Groups nor LUNs are able to be attached to more than one Storage Domain at a time.

Red Hat Enterprise 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, please refer to the *Storage Administration Guide* and *DM Multipath Guide*.

Procedure 3.1. To Add FCP Storage:

1. Click the **Storage** tab. The Storage list and toolbar display.
2. Click **New Domain**.
3. The **New Domain** dialog box displays.

LUN ID	Dev. Size	#path	Vendor ID	Product ID	Serial
--------	-----------	-------	-----------	------------	--------

Figure 3.8. Adding FCP Storage

4. Configure the following options:
 1. **Name:** Enter a suitably descriptive name.
 2. **Data Center:** Select the required Data Center from the drop-down list.
 3. **Domain Function/ Storage Type:** Select **FCP**.

4. **Use Host:** Select the IP address of either the hypervisor or Red Hat Enterprise Linux host.



Active Host Required

All communication to the storage domain is via the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must exist in the system, and be attached to the chosen data center, before the storage is configured.

5. The list of existing LUNs display. On the selected LUN, select the Add LUN check box to use it as the FCP data domain.
5. Click **OK**.

Result:

The new FCP data domain displays on the **Storage** tab. It will remain with a Locked status while it is being prepared for use. When ready, it is automatically attached to the data center. Select either **Build New Domain** or **Use Existing Volume Group**.

3.1.2.1.4. Adding Local Storage

A local storage domain can be set up on a host, to be used as a data domain for a data center and cluster that contains only a single host. Virtual machines created in a single host cluster cannot be migrated, fenced or scheduled.

Preparing local storage

This section outlines how to set up a local directory with recommended settings.

- On a Red Hat Enterprise Virtualization Hypervisor host, set up the path for the local storage as /data/images. This is the only path permitted for a Red Hat Enterprise Virtualization Hypervisor.
- On a Red Hat Enterprise Linux host, set up the path for local storage in the /data directory. Any path is permitted on a Red Hat Enterprise Linux host. Follow these instructions to add local storage:
 1. Click the **Storage** tab. The Storage list and toolbar display.
 2. Click **New Domain**. The **New Domain** dialog box displays.
 3. Enter the **Name** of the domain. A suitably descriptive name is recommended.
 4. Select the **Data** option as the **Domain Type** for the data center.
 5. Select the **Storage Type** for the domain. Select **Local** from:
 - NFS
 - iSCSI
 - FCP
 - Local

6. Select the local host in the **Use host** field. This must be the host on which the local storage is set up.
7. Enter the **Path** of the storage. For example, **data/images** or **data/localimages**.
8. Click **OK**.
9. The local storage domain displays on the **Storage** tab. This may take a few moments.

3.1.2.1.5. Example - Adding a Multipath Storage Domain

This example describes how to set up an multipath iSCSI storage domain for Red Hat Enterprise Virtualization Manager. Multipathing is inherently supported in Red Hat Enterprise Virtualization Manager. In this example, each iSCSI path must be defined manually. To do this, enter an IP for every port that the iSCSI SAN has. If only a single IP is provided, only a single path to the iSCSI target will be used.

To Add Multipathed iSCSI Storage:

1. Click the **Storage** tab. The Storage list and toolbar display.
2. Click **New Domain**. The **New Domain** dialog box displays.
3. Enter the **Name** of the storage domain.
4. Enter the **Domain function** of the storage domain, as Data, ISO or Export.
5. Select iSCSI as the Storage **Type**. The dialog box displays a set of fields appropriate to the iSCSI type.

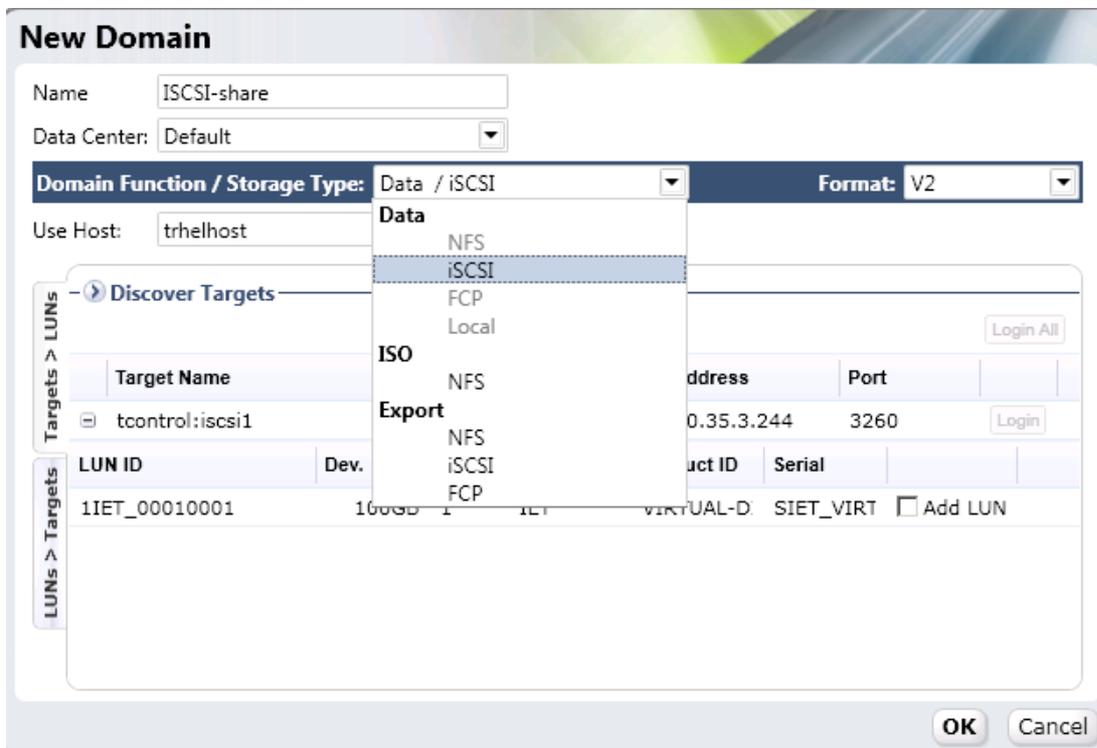


Figure 3.9. Adding iSCSI Storage

6. Select a host in the **Use host** field. To attach a domain, the name of any active host must be selected from the list.



Active Host Required

All communication to the storage domain is via the active host and not from the Red Hat Enterprise Virtualization Manager. At least one host must exist in the system before the storage can be configured.

7. Select either **Build New Domain** or **Use Preconfigured Volume Group**. In this step you can either attach a set of LUNs (create a volume group) or attach an already existing Volume Group as your storage domain. This example shows you how to build a new domain using a set of LUNs.
8. If necessary, to search for LUNs, click the **Connect to Target** button.
9. The **Connect to Targets** dialog box displays, enabling you to define a target on which to search for LUNs. Enter the requisite information in the fields.
 - a. Enter the **IP Address** of the iSCSI target.
 - b. Enter the **Port** to connect to, or leave it as the default port.
 - c. If required, enter the details for **User Authentication**.
 - d. Click the **Discover** button to find the targets.
 - e. The targets display in the list of **Discovered Targets**.
 - f. Click to either **Login to All** targets, or **Add targets manually**. If adding manually, select the LUNs (Logical Unit Numbers) from the list, and click the **Login** to login.
 - g. Click **OK**. The **Connect to Targets** dialog box closes and the LUNs display in the **New Domain** dialog box.
10. A list of LUNs display in the list of **Discovered LUNs**. Note that the Multipathing column will display a number to indicate the number of paths available to each LUN on the target. Click the check box of the LUNs to select for addition.

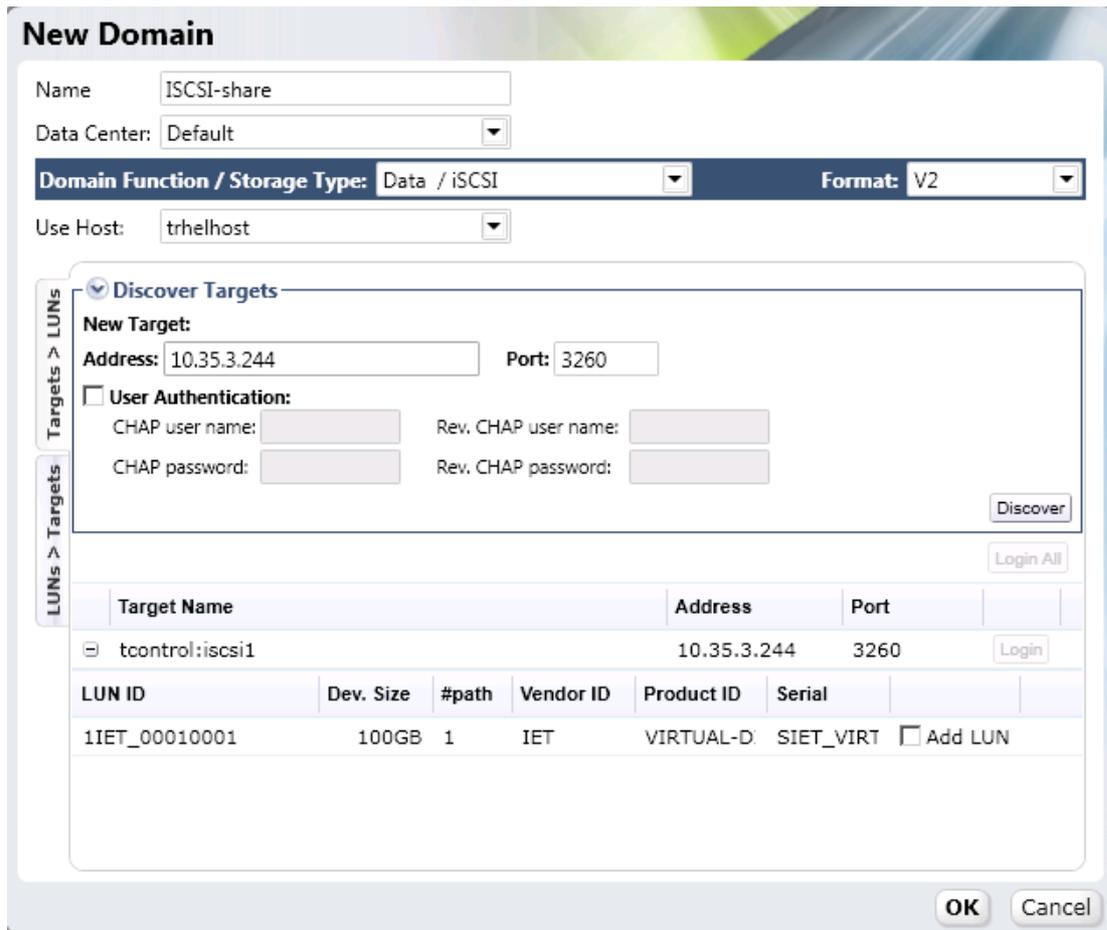


Figure 3.10. Adding Multipathed Storage

11. Click the **Add** button to use the LUNs as a storage domain.
12. The LUNs selected in the previous step display in the **Selected LUNs** grid. The set of LUNs in this list will be assigned to the new storage domain. Use the **Remove** button to remove LUNs from the Selected LUNs if necessary.
13. Click **OK** to attach the selected LUNs to the iSCSI storage domain.

3.1.2.2. Attaching Storage Domains to a Data Center

In the previous step, storage domains were created in preparation for attachment to the data center. A data center must have at least one storage domain in order to be activated. This section describes the steps to attach the data domain for virtualized disk images and subsequently the steps to attach an ISO image storage domain to a specific data center.

3.1.2.2.1. Attaching Disk Image Storage

A storage domain must be allocated to a data center to store the disk images and data of virtual machines.

To attach a data domain:

1. Click the **Data Centers** tab. Select the data center to which the storage is to be attached. If the required data center is not displayed, perform a search (see [Section 1.2, "Search"](#)).
2. The Details pane of the selected data center displays. Select the **Storage** tab.

Domain Name	Domain Type	Storage Type	Format	Cross Data-Center Status	Free Space
ExportDomain	Export	NFS	V1	Unattached	129 GB
FinanceDataDomain	Data (Master)	NFS	V1	Active	169 GB
ISODomain	ISO	NFS	V1	Active	75 GB
SalesDataDomain	Data (Master)	iSCSI	V2	Unattached	35 GB

Name	Type	Status	Avail.	Used	Total
Iso	ISO	Active	227 GB	65 GB	292 GB
DataStore	Data (Master)	Active	102 GB	88 GB	190 GB

Figure 3.11. Data Center Storage Tab

- Click the **Attach Domain** button to add the storage location where the data and disk images are stored.
- The **Attach Storage Domain** dialog box displays.
- Select the domain from the **Storage Domain** list. The names of any existing storage domains, of the type appropriate for the data center display in the list. For example, if the default data center has a storage type of NFS, only existing NFS storage domains display in the list, because only NFS storage domain types can be attached to this particular data center.
- Click **OK**. The new storage domain displays on the **Storage** tab of the Details pane.

3.1.2.2.2. ISO Uploader

The Red Hat Enterprise Virtualization Manager installation includes a tool for uploading ISO images to the ISO storage domain. This tool is referred to as the ISO uploader. It provides for the listing of storage domains and uploading of ISO files to them.

The ISO uploader command is **rhevms-iso-uploader**. You must be logged in as the root user to run it successfully. You must provide the administration credentials for the Red Hat Enterprise Virtualization environment on the command line. Full usage information, including a list of all valid options for the command, is available by running the **rhevms-iso-uploader -h** command.

3.1.2.2.2.1. Syntax

The basic syntax is of the form:

```
Usage: rhevms-iso-uploader [options] list
       rhevms-iso-uploader [options] upload [file].[file]...[file]
```

The two supported modes of operation are *list*, and *upload*.

- The *list* parameter lists the available ISO storage domains. These storage domains are the valid targets for ISO uploads. By default the list is obtained from the Red Hat Enterprise Virtualization Manager installation on the local machine.

- The *upload* parameter uploads the selected ISO file(s) to the specified ISO storage domain. By default the transfer is performed using NFS however SSH is also available.

Basic ISO uploader usage requires that, at a minimum, the either the *list* or *upload* parameter is provided. Where *upload* is selected then the name of at least one local file to upload must also be provided.

The **rhevms-iso-uploader** command has a large number of options.

General Options

--version

Displays the version number of the command in use, and exits immediately.

-h, --help

Displays command usage information, and exits immediately.

--quiet

Sets quiet mode, reducing console output to a minimum. This is off by default.

--log-file=PATH

Sets *PATH* as the log file the command should use for its own log output.

--conf-file=PATH

Sets *PATH* as the configuration file the command should use.

-v, --verbose

Sets verbose mode, providing more console output. This is off by default.

-f, --force

Where the source file being uploaded has the same file name as an existing file at the destination, force the existing file to be overwritten automatically. This is off by default.

Red Hat Enterprise Virtualization Manager Options

The options in the Red Hat Enterprise Virtualization Manager configuration group are used to specify the manager authentication details and, filter log collection from one or more virtualization hosts. If no options in this group are specified, data is not collected from any virtualization host.

-u USER, --user=USER

Sets the user as *USER*. This must be a user that exists in directory services, and is known to the Red Hat Enterprise Virtualization Manager. The user must be specified in the format *user@domain*, where *user* replaced by the username, and *domain* is replaced by the directory services domain in use.

-r FQDN, --rhevms=FQDN

Sets the Red Hat Enterprise Virtualization Manager to connect to as *FQDN*. *FQDN* must be replaced by the fully qualified domain name of the manager. By default it is assumed that the ISO uploader is being run on the same machine as the manager. Therefore the default value for this parameter is **localhost**.

ISO Storage Domain Options

The options in this configuration group are used to specify the ISO domain to which files must be uploaded

-i, --iso-domain=ISODOMAIN

Sets the storage domain named *ISODOMAIN* as the destination for uploads.

`-n, --nfs-server=NFSSERVER`
Sets the NFS path *NFSSERVER* as the destination for uploads. This option is an alternative to `--iso-domain`, the two must not be used at the same time.

Example 3.1. Specifying an NFS Server

```
# rhvm-iso-uploader --nfs-server=storage.demo.redhat.com:/iso/path upload RHEL6.0.iso
```

Connection Options

By default the ISO uploader uses NFS to upload files. Use options within this configuration group to use SSH file transfer instead.

`--ssh-user=USER`
Sets *USER* as the SSH username to use for the upload.

`--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. If no key is set the program will prompt you to enter the password of the user specified instead.

3.1.2.2.2. Examples

Example 3.2. Basic ISO Uploader Usage

In this example the ISO uploader is run to list the available ISO storage domains. The username is not provided on the command line so the tool instead prompts for it to be entered. Once the storage domains have been listed, an ISO file is uploaded to one of them over NFS.

```
# rhvm-iso-uploader list
Please provide the REST API username for RHEV-M (CTRL+D to abort): admin@directory.demo.redhat.com
Please provide the REST API password for RHEV-M (CTRL+D to abort):
ISO Storage Domain List:
  ISODomain
# rhvm-iso-uploader --iso-domain=ISODomain upload RHEL6.iso
Please provide the REST API username for RHEV-M (CTRL+D to abort): admin@directory.demo.redhat.com
Please provide the REST API password for RHEV-M (CTRL+D to abort):
```

3.1.2.2.3. Attaching an Export Storage Domain

An export domain can be attached to a data center to enable the import or export of virtual machines from one data center to another. Export domains created for a Red Hat Enterprise Virtualization 3.0 environment must be NFS based. However, export domains of other types from older environments can be imported into a 3.0 environment. An export domain can also be used to backup virtual machines and templates. To import an existing export domain, refer [Import Existing ISO or Export Storage Domain](#).



Note

At a given time, an export domain can only be attached to a single data center.

To attach an export domain:

1. Click the **Data Centers** tab.
Select the data center to which the export domain is to be attached.
2. The Details pane displays. Select the **Storage** tab.
3. Click the **Attach Export** button to add the storage location where the images are stored.
4. The **Attach Export Domain** dialog box displays, if there are export domains available.

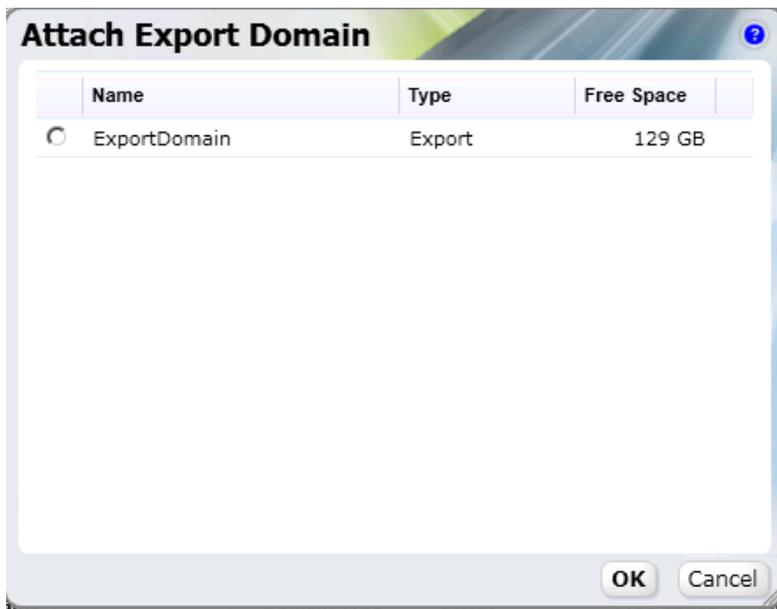


Figure 3.12. Attach Export Domain Dialog Box

5. Select the export domain from the list.
6. Click the **OK**. The new export domain displays on the **Storage** tab of the **Details** pane, with a status of *Locked*, followed by *Inactive*.
7. Select the new export domain on the **Storage** tab of the **Details** pane, and click the **Activate** button.
8. The export domain will be activated in a few moments and display an *Active* status.

3.1.3. Storage Entities

The Storage Pool Manager

The Storage Pool Manager (SPM) entity coordinates all the metadata changes across the data center. This includes creating, deleting and manipulating virtual disks (Images), snapshots, and templates,

and allocating storage for sparse block devices (on SAN). The SPM entity can be run on any host in the data center, and it is the Manager's task to grant the role to one of the hosts. All hosts in a data center must have access to all the storage domains defined in the data center. The SPM entity controls access to storage, by coordinating the metadata across the storage domains.

Red Hat Enterprise Virtualization Manager ensures that the SPM is always available. If errors occur, the Manager will try to move the SPM role to a different host. This means that if the host that is running as the SPM has problems accessing the storage, the Manager will automatically check if there is another available host that can access the storage and will move the SPM over to that host. When the SPM starts, it tries to ensure that it is the only host that was granted the role, therefore it will acquire a storage-centric lease. This process can take some time.

Multipathing

Multipathing is supported in Red Hat Enterprise Virtualization Manager by default. Setting up a multipathed storage domain is described later in this section. To configure multipathing for Red Hat Enterprise Virtualization Hypervisor hosts, see *Red Hat Enterprise Linux Hypervisor Deployment Guide*.



Warning

Do not add `user_friendly_names` and LUN aliases to a `multipath.conf` file on a Red Hat Enterprise Virtualization Hypervisor. `user_friendly_names` and LUN aliases are not supported and can lead to unpredictable system behavior.

3.1.4. Storage Permissions

3.1.4.1. Managing System Permissions for a Storage Domain

While the superuser or system administrator of the platform has the full range of permissions, a Storage Administrator is a system administration role for a storage domain only. The Storage Administrator role permits the management, creation and removal of their assigned storage domain, and configuration changes to a single storage domain. This is useful in an enterprise where there are multiple storage domains, each of which require their own system administrators. A Storage Administrator has permissions for the assigned storage domain only, not for all storage domains in the enterprise.

To assign a system administrator role to a storage domain:

1. Click the **Storage** tab.

A list of storage domains displays. If the required storage domain is not visible, perform a search (see [Section 1.2, "Search"](#)).

2. Select the storage domain that you want to edit, and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any. The System Administrator of the Red Hat Enterprise Virtualization platform, the Data Center Administrator and the Template Administrator, if existing, will display with inherited permissions.

User	Role	Inherited Permission
Susan Burgess	SuperUser	(System)
David Jorm	RHEVMUser	(System)
Cheryn	StorageAdmin	
rhevadmin	SuperUser	(System)

Figure 3.13. Storage Permissions

3. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
4. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **StorageAdmin**.

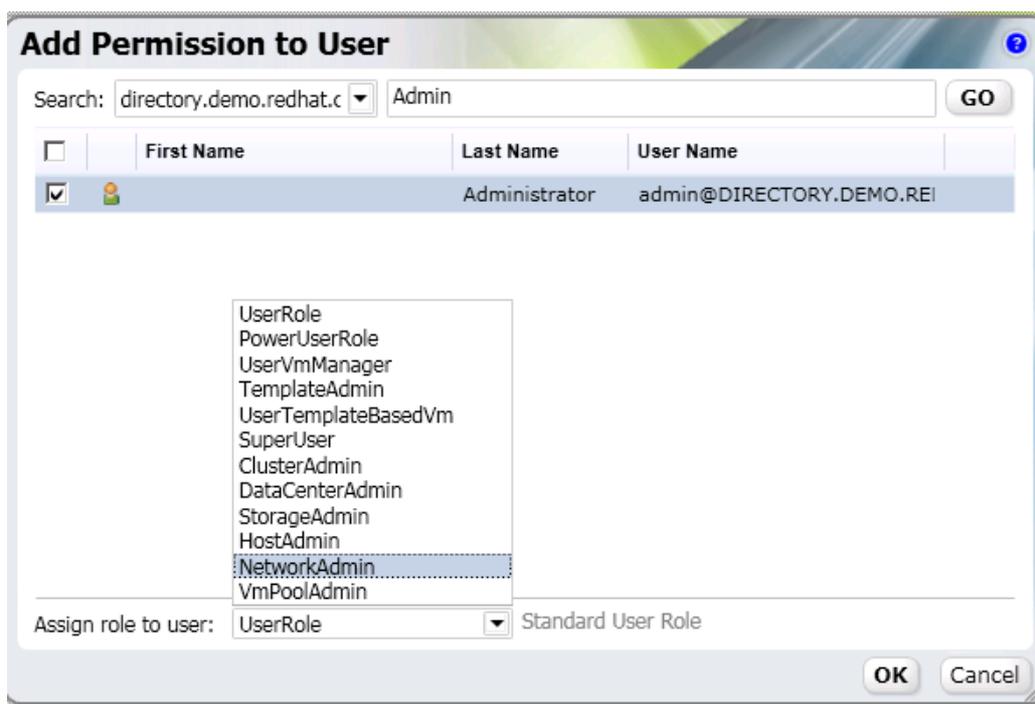


Figure 3.14. Assign StorageAdmin Permission

5. Click **OK**.

The name of the user displays in the **Permissions** tab, with an icon and the assigned Role.

**Note**

You can only assign roles and permissions to existing users. See (see [Chapter 5, Users](#)).

You can also change the system administrator of a storage domain, by removing the existing system administrator, and adding the new system administrator, as described in the previous procedure.

To remove a system administrator role:

1. Click the **Data Center** tab. A list of storage domains displays. If the required storage domain is not visible, perform a search (see [Section 1.2, "Search"](#)).
2. Select the required storage domain and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any. The System Administrator of the Red Hat Enterprise Virtualization platform, the Data Center Administrator and the Template Administrator, if existing, will display with inherited permissions. You cannot remove permissions from these users.

3. Select the check box of the appropriate user.
4. Click **Remove**. The user is removed from the Permissions tab.

3.1.5. Storage Troubleshooting

3.1.5.1. Maintaining Storage Domains

This section describes how to maintain storage domains. For example, you may need to do this to balance the load, improve performance for particular applications, or if storage domains are being replaced or retired. You can edit, reactivate and update domains. You can also deactivate domains, and detach them from the cluster and data center. Changing the storage domain properties is a sensitive task as it affects the all the Virtual Machines and Hosts in the Cluster.

**Warning**

All maintenance tasks need to be approached with extreme care. Proceed with caution before any parameters on a storage domain are changed. Failure to do so may result in the loss of all data and images. There is no guarantee that the images can be recovered.

3.1.5.1.1. Moving Storage Domains to Maintenance Mode

Storage domains in a data center need to be put into maintenance mode in a fixed order. If the data center also has an ISO domain, the ISO domain must be placed into Maintenance mode before you can place the storage domain into maintenance mode.

To move a storage domain into maintenance mode:

1. Click the **Storage** tab. The Storage page displays the list of existing storage domains, and the Storage toolbar displays.

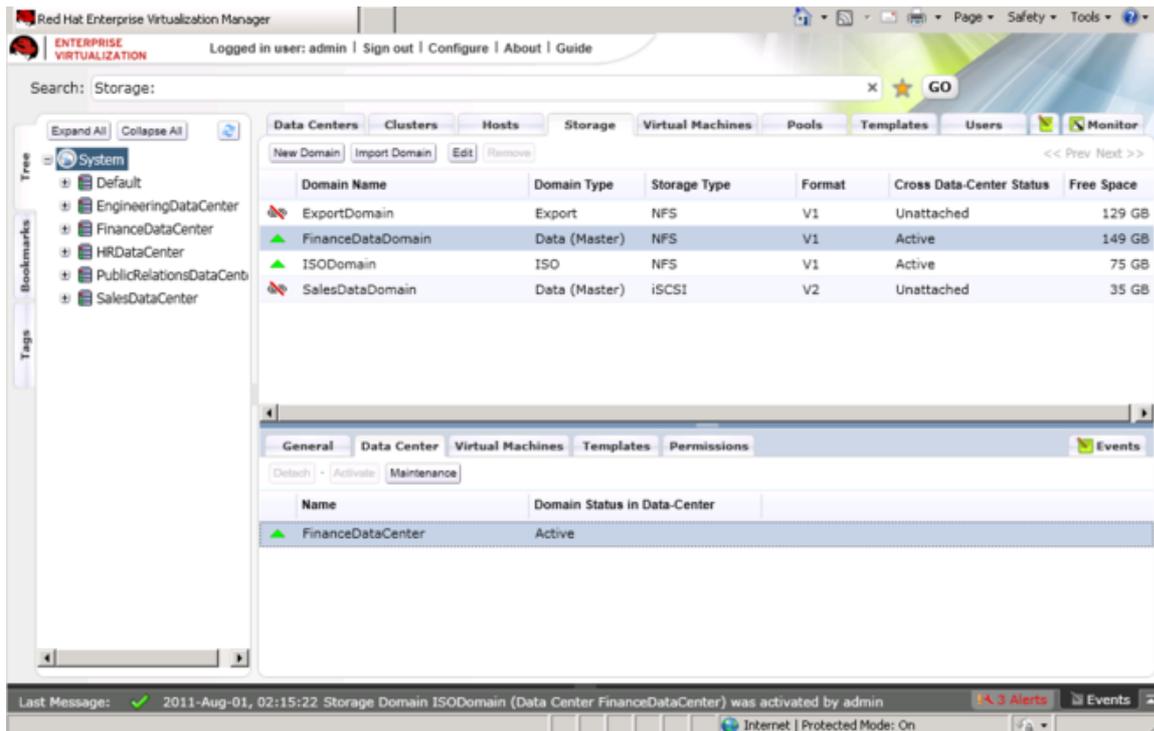


Figure 3.15. The Storage Tab

If the required storage is not displayed, perform a search (see [Section 1.2, “Search”](#)).

2. Shut down and move all the virtual machines running on the data domain. See [Section 6.7.5, “Moving Virtual Machines within a Data Center”](#).
3. Select the ISO storage domain, if any, to place in maintenance mode.
4. On the **Details** pane, click the **Data Center** tab. Click the **Maintenance** button. The ISO storage domain is deactivated, and displays as Inactive in the Storage pane.
5. Select the data domain to be moved into maintenance mode. If you attempt to move a data storage domain into maintenance mode while the ISO domain is still active, a message appears prompting you to deactivate other data domains.
6. On the Details pane, click the **Data Center** tab. Click the **Maintenance** button. The data storage domain is deactivated, and appears as Inactive in the Storage pane.

You can now edit, detach, remove or reactivate the inactive storage domains from the data-center.



Note

You can also activate, detach and place domains into maintenance mode using the Storage tab on the Details pane of the data center it is associated with.

3.1.5.1.2. Editing Storage Domains

Inactive or Active storage domains in a data center may need to be modified in a dynamically changing environment.



Warning

All maintenance tasks need to be approached with extreme caution. Proceed with caution before any parameters on a storage domain are changed. Failure to do so may result in the loss of all data and images. There is no guarantee that the images can be recovered.

To Edit Storage Domains:

1. Click the **Storage** tab. The Storage page displays the list of existing storage domains, and the Storage toolbar displays.

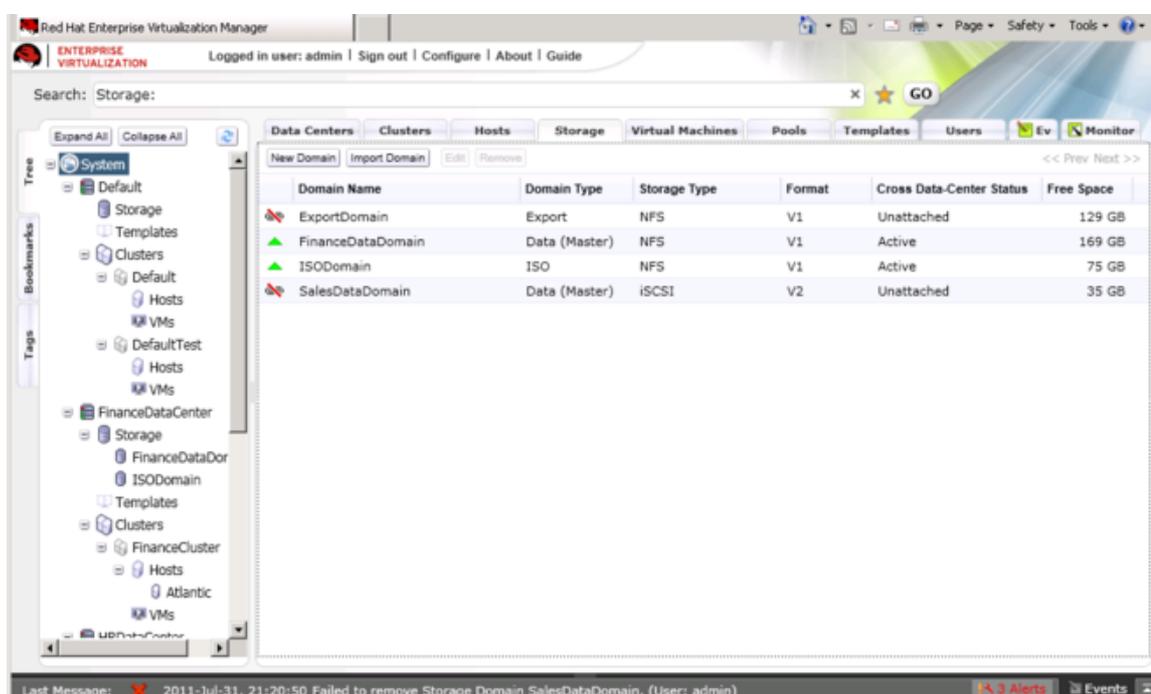


Figure 3.16. The Storage Tab

2. Select the required storage domain. Ensure that it is in Maintenance mode.
If the required storage is not displayed, perform a search (see [Section 1.2, "Search"](#)).
3. Click **Edit** on the Storage toolbar. The **Edit Storage Domain** dialog box displays. Depending on the status of the domain, some or all fields in the dialog box are enabled. The **Edit Storage Domain** dialog box contains the same fields as the **New Storage** dialog box. See [Figure 3.2, "Adding New Storage"](#).
4. Change the required fields and click **OK**.
5. You can now activate the storage and check the validity of the configuration. See [Section 3.1.2.2, "Attaching Storage Domains to a Data Center"](#)



Note

You can also activate, detach and place domains into maintenance mode using the Storage tab on the Details pane of the data center it is associated with.

3.1.5.1.3. Activating Storage Domains

Inactive storage domains in a data center need to be reactivated before they can be used. At least one data storage domain must be activated before the ISO domain can be activated, if an ISO domain exists.

To activate storage domains:

1. Click the **Storage** tab. The Storage page displays the list of existing storage domains, and the Storage toolbar displays.

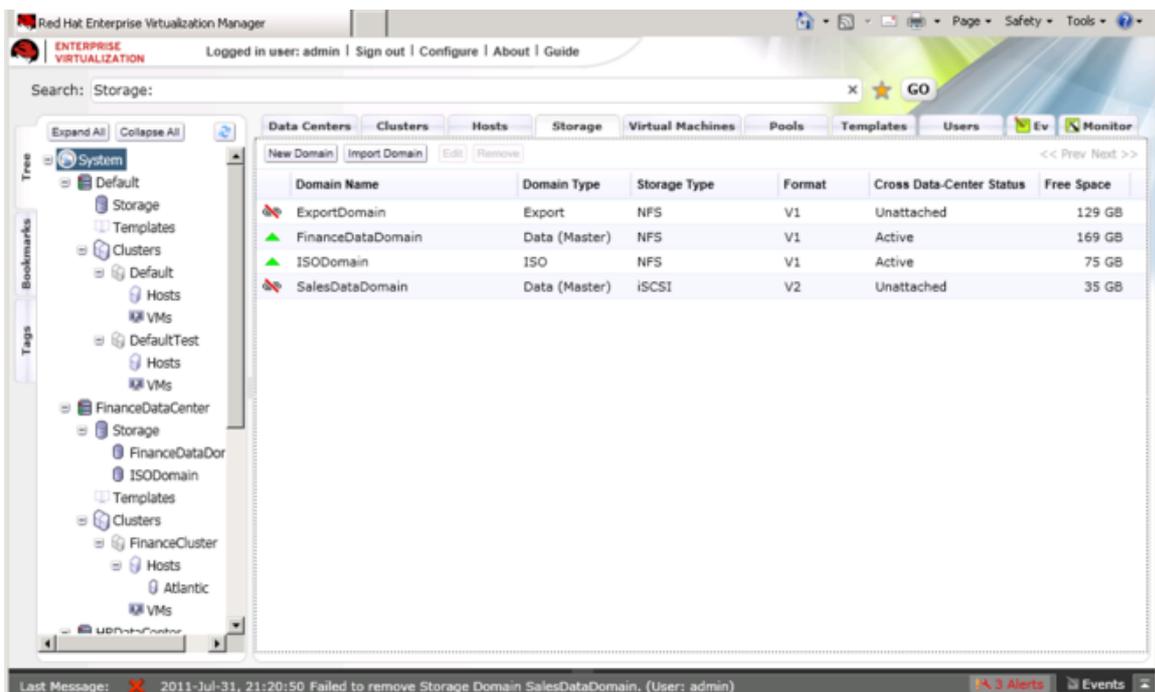


Figure 3.17. An Inactive Domain

2. Select an inactive data storage domain.

If the required storage is not displayed, perform a search (see [Section 1.2, “Search”](#)).

3. On the Details pane, click the **Data Center** tab.

Click **Activate** button on the toolbar. The domain is activated, and displays as Active in the Storage pane.



Important

If you attempt to activate the ISO domain before activating the data domain, an error message displays, and the domain is not activated.



Note

You can also activate, detach and place domains into maintenance mode using the **Storage** tab on the Details pane of the data center it is associated with.

3.1.5.2. Deleting Storage Domains

This section describes how to delete storage domains from a data center. For example, you may need to do this if storage domains are being replaced or retired. There are two ways to do this, you can choose to detach storage from a particular data center, or you may choose to remove it altogether from the system. Storage domains cannot be removed or detached if any virtual machines that reside on it are running.



Warning

Deleting storage domains is an irreversible process. Proceed with caution before any storage domains are detached or removed. All images on the storage domain are irreversibly lost on detachment and removal of a storage domain.

3.1.5.2.1. Detaching Storage Domains from a Data Center

The space available on storage domains that are merely detached from a data center remain available to be reassigned later, or assigned to other data centers. After detachment the domain will still appear in the lists of assigned or unassigned storage domains.

To detach a storage domain from a data center:

1. Click the **Storage** tab. The Storage page displays the list of existing storage, and the Details pane displays.
2. Select the storage domain to be detached. Ensure that no virtual machines are running on the domain.
3. Move the storage domain into Maintenance mode. See [Section 3.1.5.1, "Maintaining Storage Domains"](#).
4. On the Details pane, click the **Data Centers** tab.
5. Click **Detach** button on the Storage toolbar.

The Detach Storage dialog box displays a list of the domains selected for detachment.



Figure 3.18. The Detach Storage Dialog Box

6. The detached storage domain displays in the list of storage domains with a status of Detached.



Note

To check if the storage location is still available, use the **Attach Domain** or **Add ISO** button on the Storage tab in the Details pane of the data center to attach the domain again, if necessary. Refer [Section 3.1.2.2, “Attaching Storage Domains to a Data Center”](#).

3.1.5.2.2. Removing Storage Domains

Storage domains that are removed from a data center are also deleted from the system. After deletion they no longer display in the lists of storage domains, for example in the **Add Storage Domain** dialog box.

Storage domains that are removed from the system must be fully reconfigured before they can be reused.



Warning

Proceed with caution before any storage domains are detached or removed. All images on the storage domain are irreversibly lost on detachment and removal of a storage domain.

To remove a storage domain:

1. Click the **Storage** tab. The Storage page displays the list of existing storage domains, and the Storage toolbar displays.
2. Select the storage domain to be removed. Ensure that no virtual machines are running on the domain.

3. Move the domain into Maintenance mode to deactivate it. See [Section 3.1.5.1.1, “Moving Storage Domains to Maintenance Mode”](#).
4. Click **Remove** on the Storage Tool bar.
5. The **Remove Storage** dialog box displays prompting you to confirm removal, and select the host to be used to effect the removal. Select a host from the list box.

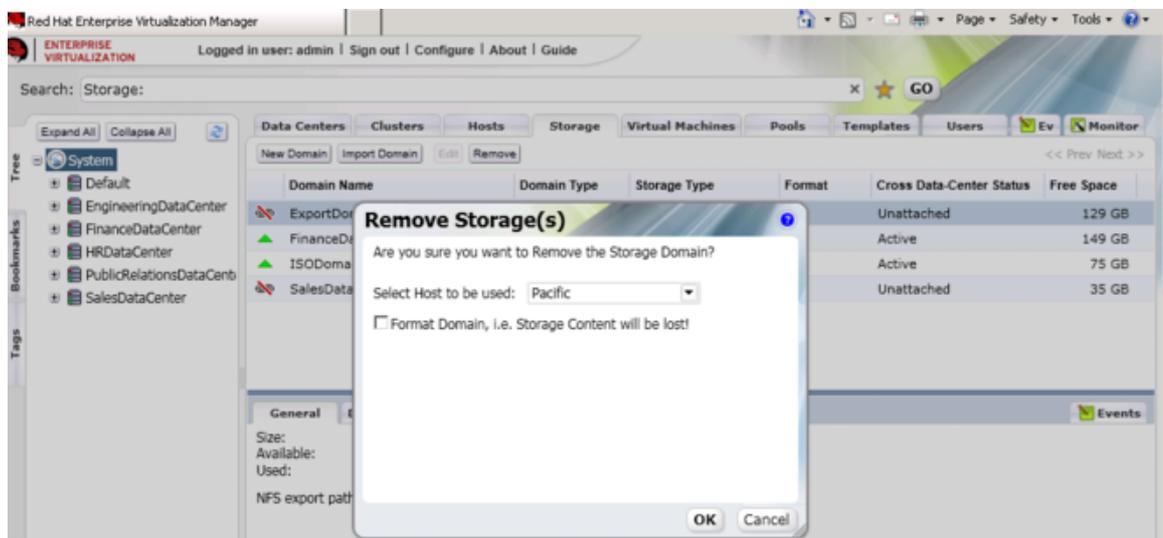


Figure 3.19. Remove Storage Dialog Box

6. Click **OK**. The storage domain is permanently removed from the system.
7. Click the **Storage** tab. The deleted storage domain no longer displays in the list of storage domains.

Note

To check that the deleted storage domain is no longer available, use the **Add Storage Domain** button on the Storage toolbar. Refer [Section 3.1.2.1, “Adding Storage Domains to a Data Center”](#). For information regarding the setup and configuration of FCP or multipathing on Red Hat Enterprise Linux, please refer to the *Storage Administration Guide* and *DM Multipath Guide*.

Red Hat Enterprise Virtualization Hosts

Hosts are the physical servers on which the virtual machines run. A Host is a compact, full-featured virtualization platform for quickly and easily deploying and managing virtualized guests. Full virtualization is provided by using a loadable Linux kernel module called Kernel-based Virtual Machine (KVM). KVM can concurrently host multiple virtualized guests running either Windows or Linux operating systems. Virtualized guests run as individual Linux processes and threads on the host machine and can be managed remotely using the Red Hat Enterprise Virtualization Manager. A Red Hat Enterprise Virtualization environment has one or more hosts attached to it.

Alternatively, host is also called hypervisor. Both Red Hat Enterprise Virtualization Hypervisors and Red Hat Enterprise Linux hosts interact with the rest of the virtualized environment in the same way. Red Hat Enterprise Virtualization Manager handles both Red Hat Enterprise Virtualization Hypervisors and Red Hat Enterprise Linux hosts with the KVM hypervisor, delivering leading performance and scalability for virtual machines on a stable and secure platform for their most mission-critical workloads. This section describes how set up and manage the host types that can be used in the Red Hat Enterprise Virtualization platform.

4.1. Managing Hosts

A host is a physical 64-bit server with the Intel VT® or AMD-V™ extensions running any of the following:

- Red Hat Enterprise Linux 6.1 or later AMD64/Intel 64 version



Note

Support is still ongoing for previous versions of Red Hat Enterprise Linux 5.4, 5.5, and already existing Clusters. However, the Red Hat Enterprise Virtualization Guest Agent is now included in the virtio serial channel, whereas before it was in a separate channel. As a result, the Guest Agent installed on Windows guests on Red Hat Enterprise Linux hosts that are upgraded from version 5 to 6 lose connection to the Manager.

A physical host on the Red Hat Enterprise Virtualization platform:

- Must belong to only one cluster in the system. For more information on clusters, refer [Section 2.2, “Clusters”](#).
- 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.
- Can have a maximum of 128 physical CPUs.
- Can have a maximum of 1 TB RAM.
- Can have an assigned system administrator with system permissions.

The Red Hat Enterprise Virtualization Hypervisor has various security features enabled. Security Enhanced Linux (SELinux) and the iptables firewall are fully configured and on by default.

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

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

Red Hat Enterprise Virtualization platform uses various network ports for management and other virtualization features. These ports must be open on a Red Hat Enterprise Linux 6.1 host or higher. For a full list of ports, see *Red Hat Enterprise Virtualization Installation Guide*.

4.1.1. Hosts Properties

The **Hosts** tab provides a graphical view of all the hosts in the system. The following properties of the host are displayed on the Hosts tab, and on the New and Edit dialog boxes. The **Network Interfaces** tab is described in [Section 4.1.2.2, “Managing Host Network Interfaces”](#).

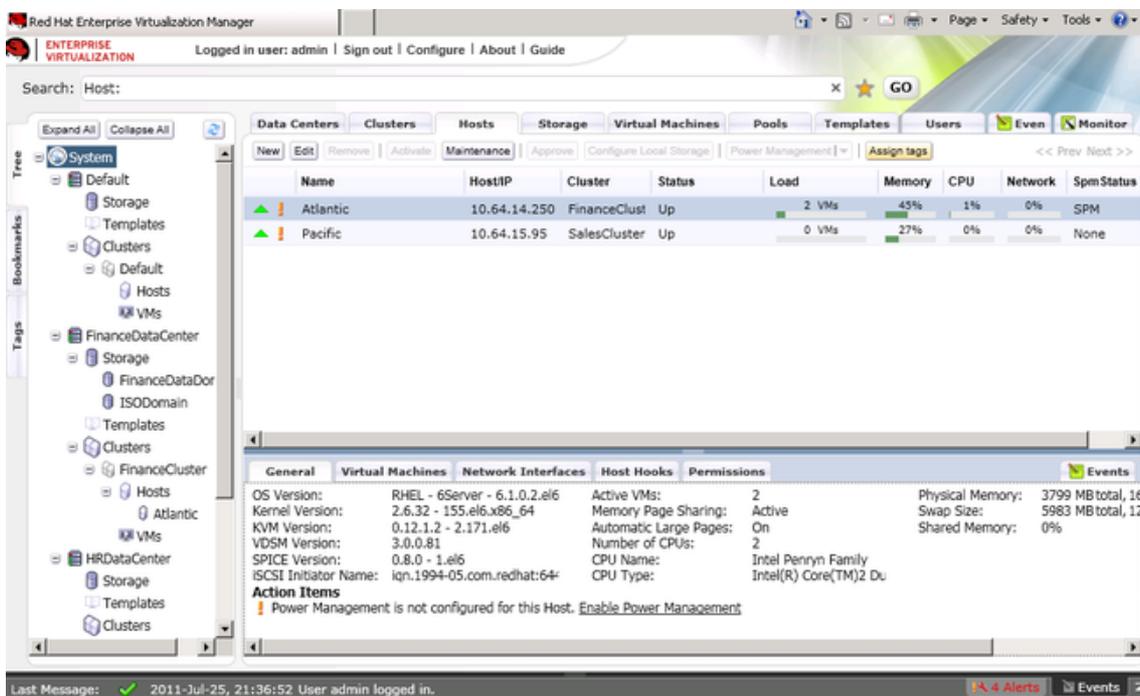


Figure 4.1. Host Details Pane

Table 4.1. Hosts Properties

Field/Tab	Description/Action
Data Center	The selected data center.
Host Cluster	The selected host cluster. All hosts in a cluster must be of the same architecture.
Name	The host name. Provide a descriptive name.
Address	The IP address, or resolvable hostname of the host (provided during installation).
Root Password	The password of the designated host; used during installation of the host.
Power Management Address	The address of the remote access card (RAC) on the host. The password of the designated host; used during installation of the host.
Power Management User Name	A valid User Name for the OOB management.

4.1.2. Host Operations

4.1.2.1. Adding Hosts

Hosts must be correctly installed before you can add them to the Red Hat Enterprise Virtualization platform. Before adding hosts ensure that they have been configured correctly with a name and IP address.



Note

There is no need to configure Network Bridge, as this is automatically created on the hypervisor and exists to allow Virtual Machines to interact with the network as if they were using physical NICs.

Once added to the Administration Portal, hosts must be either approved or activated from the **Hosts** tab on Red Hat Enterprise Virtualization Manager.



Important

If you re-install Red Hat Enterprise Virtualization Manager, you must remove the hosts to enable them to be reconnected with the correct ssh keys for the new installation of Red Hat Enterprise Virtualization Manager. In contrast, if you upgrade Red Hat Enterprise Virtualization Manager, the hosts remain connected, and no action is required from you.

4.1.2.1.1. Prerequisites

Before you can add a host to Red Hat Enterprise Virtualization platform, ensure the following criteria have been met.

- The host hardware is Red Hat Enterprise Linux certified. Refer to <https://hardware.redhat.com/> to confirm that the server has Red Hat certification.



Important

The Red Hat Enterprise Virtualization platform only supports 64-bit processors with the Intel VT or AMD-V extensions. Only the AMD64/Intel 64 version of Red Hat Enterprise Linux 6.1 and higher is compatible for use with Red Hat Enterprise Virtualization platform.

- If you are using VLAN, the network VLAN should be configured for access to the Red Hat Enterprise Virtualization Manager.
- If a host is to be highly available, and have power management, out-of-band management must be set up and configured correctly. In most instances, this requires the presence of a remote access card (RAC) in the host.

- The BIOS in the host has Intel VT or AMD-V activated.
- Red Hat Enterprise Virtualization Hypervisor or Red Hat Enterprise Linux Host has been installed with either of the supported operating systems. For detailed information on installation, including how to install multiple hosts, install from networks, or other advanced features, refer to the appropriate installation documents. Refer [Appendix H, Additional References](#).
- The host has a resolvable hostname or static IP address.
- A data partition with a minimum size of 25 GB is recommended to provide temporary storage.

4.1.2.1.2. Adding Red Hat Enterprise Virtualization Hypervisor Hosts

If you have not yet installed your Red Hat Enterprise Virtualization Hypervisor Host(s) then you must do so before you can attach them to the Red Hat Enterprise Virtualization Manager. Information on installing Red Hat Enterprise Virtualization Hypervisors is contained in the *Red Hat Enterprise Linux — Hypervisor Deployment Guide* and *Red Hat Enterprise Virtualization Installation Guide*. Once you have followed the instructions in these guides to install the Red Hat Enterprise Virtualizations you are ready to attach them to the Red Hat Enterprise Virtualization Manager.

During the installation of the Red Hat Enterprise Virtualization Hypervisor, the configuration script prompts for the IP address, or fully qualified domain name, of the Red Hat Enterprise Virtualization Manager. If the correct address is provided, the Red Hat Enterprise Virtualization Hypervisor host automatically contacts the manager and is added to the list of hosts shown in the Administration Portal. At this point while the manager knows of the hypervisor it must still be manually approved before it is ready for use. To approve a newly-configured hypervisor:

1. In the **Hosts** tab, select the newly-configured host. This host will display a status of "Pending Approval".
2. Click the **Approve** button.

The approval process is a hand shake between the Management server and the host. On successful conclusion of this process the host's status changes to *Up*. The host is now certified and is part of the Red Hat Enterprise Virtualization platform.

If you didn't provide a valid Red Hat Enterprise Virtualization Manager IP address, or fully qualified domain name, when you installed the hypervisor host then it must be added using the **Add Host** dialog. To do this you will need to know the IP address of the hypervisor and the password which was set on the **RHEV-M** screen when it was configured. See [Section 4.1.2.1.4, "To Add a Host"](#) for further information.

4.1.2.1.3. Adding Red Hat Enterprise Linux Hosts

Red Hat Enterprise Virtualization also supports hosts running Red Hat Enterprise Linux AMD64/Intel 64 version. This section describes the preparatory steps for installing the Red Hat Enterprise Linux host, as well as the steps to manually add the host to the Red Hat Enterprise Virtualization platform.

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

Red Hat Enterprise Virtualization 3.0 supports Red Hat Enterprise Linux, version 6.1 and higher as hosts.

4.1.2.1.3.1. Preparing Red Hat Enterprise Linux Hosts

To ensure a smooth and successful integration of Red Hat Enterprise Linux Hosts and Red Hat Enterprise Virtualization platform, prepare the host carefully according to the instructions in this section.

Procedure 4.1. Directions:

1. Install Red Hat Enterprise Linux

Ensure that Red Hat Enterprise Linux is correctly installed and configured on the physical host. Refer to the *Red Hat Enterprise Linux Installation Guide* for more information. Only the Base package group is required. All other packages can be removed or not selected.



Warning — DNS Configuration

The Red Hat Enterprise Linux host must have a fully resolvable network address. Valid forward and reverse lookups for the address must be available in DNS. Virtual machine migration will not work in environments where this is not the case.



Important — Authentication Files Required by useradd Must be Accessible

If you are using proprietary directory services or standard directory services with no access to authentication files for user management, the `vdsm` package will fail to create the required system user. The authentication files required by the `useradd` command must be accessible to the installer. Red Hat Directory Server (**RHDS**) recommends a security policy with a mixture of local files and **LDAP**. Following this recommendation will resolve this issue.

2. Configure VLANs

If you are using VLAN, ensure that VLANs are configured for access to the Red Hat Enterprise Virtualization Manager.

3. Check Red Hat Network Subscriptions

Ensure the host is correctly subscribed to the appropriate Red Hat Network channel **Red Hat Enterprise Virt Management Agent (v 6 x86_64)** channel, also referred to as `rhel-x86_64-rhev-mgmt-agent-6`, on Red Hat Network. If you do not have the appropriate subscription entitlements, contact Red Hat Customer Service.

- a. If the machine has not already been registered with Red Hat Network, run the `rhn_register` command as `root` to register it. To complete registration successfully you will need to supply your Red Hat Network username and password. Follow the onscreen prompts to complete registration of the system.

```
# rhn_register
```

- b. Run the `rhn-channel` command to subscribe the system to channel.

```
# rhn-channel --add --channel=rhel-x86_64-rhev-mgmt-agent-6
```

4. Edit hosts file

Add a manual host entry to the `/etc/hosts` file (on the Red Hat Enterprise Linux host) for the Red Hat Enterprise Virtualization Manager server to enable `vdsm` and other services to connect properly to the host.

Because Active Directory uses layered domain names, the Active Directory instance takes the domain name usually used by a Linux host. For example, if the server running the Red Hat Enterprise Virtualization Manager has a hostname of `server1.example.com`, Active Directory uses that address and creates a sub-address named `rhev-manager.server1.example.com`.

Edit the `/etc/hosts` file on the Red Hat Enterprise Linux Host. The following screen output sample resembles the contents of the file:

```
127.0.0.1    localhost.localdomain localhost
::1         localhost6.localdomain6 localhost6
```

Append a new line to `/etc/hosts` with the IP address and *both* variants of the Red Hat Enterprise Virtualization Manager domain names. The following screen output sample resembles the required contents of the file:

```
127.0.0.1    localhost.localdomain localhost
::1         localhost6.localdomain6 localhost6
10.0.0.1     server1.example.com rhev-manager.server1.example.com
```

5. Open firewall ports

Red Hat Enterprise Virtualization platform uses a number of network ports for management and other virtualization features.

The following steps configure `iptables` to open the required ports. These steps replace any existing firewall configuration with that required for Red Hat Enterprise Virtualization Manager. If you have existing firewall rules with which this configuration must be merged then you must manually edit the rules defined in the `iptables` configuration file, `/etc/sysconfig/iptables`.

- a. Remove any existing firewall rules.

```
# iptables --flush
```

- b. Add the ports required by Red Hat Enterprise Virtualization Manager to the `iptables` rules.

```
# iptables --append INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
# iptables --append INPUT -p icmp -j ACCEPT
# iptables --append INPUT -i lo -j ACCEPT
# iptables --append INPUT -p tcp --dport 22 -j ACCEPT
# iptables --append INPUT -p tcp --dport 16514 -j ACCEPT
# iptables --append INPUT -p tcp --dport 54321 -j ACCEPT
# iptables --append INPUT -p tcp -m multiport --dports 5634:6166 -j ACCEPT
# iptables --append INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
# iptables --append INPUT -j REJECT --reject-with icmp-host-prohibited
```

```
# iptables --append FORWARD -m physdev ! --physdev-is-bridged -j REJECT --reject-with icmp-host-prohibited
```



Note

The provided **iptables** commands add firewall rules to accept network traffic on a number of ports. These include:

- port 22 for **SSH**,
- ports 5634 to 6166 for guest console connections,
- port 16514 for **libvirt** virtual machine migration traffic,
- ports 49152 to 49216 for VDSM virtual machine migration traffic, and
- port 54321 for the Red Hat Enterprise Virtualization Manager.

- c. Save the modified rules.

```
# service iptables save
```

- d. Ensure that the **iptables** service is configured to start on boot and has been restarted, or started for the first time if it wasn't already running.

```
# chkconfig iptables on  
# service iptables restart
```

6. Configure sudo access

The Red Hat Enterprise Virtualization Manager makes use of **sudo** to perform operations as root on the host. The default configuration stored in **/etc/sudoers** contains values to allow this. If this file has been modified since Red Hat Enterprise Linux installation these values may have been removed. As root run **visudo** to ensure that the **/etc/sudoers** contains the default configuration values. Where it does not they must be added.

```
# Allow root to run any commands anywhere  
root    ALL=(ALL)  ALL
```

7. Enable SSH access for root

The Red Hat Enterprise Virtualization management daemon accesses host machines via SSH. To do this it logs in as root with an encrypted key for authentication. To ensure that SSH is configured and root is able to use it to access the system follow these additional steps.



Warning

The first time the Red Hat Enterprise Virtualization Manager is connected to the host it will install an authentication key. In the process it will overwrite any existing keys which exist in `/root/.ssh/authorized_keys`.

- a. These steps assume that the `openssh-server` package is installed on the system. Where the package is not present use **yum** to install it.

```
# yum install openssh-server
```

- b. Use **chkconfig** to verify which run-levels SSH is enabled at.

```
# chkconfig --list sshd
sshd 0:off 1:off 2:on 3:on 4:on 5:on 6:off
```

It is expected that the SSH daemon shows as **on** for run-levels **3**, **4**, and **5**. This is the default configuration.

If the configuration on the host differs use **chkconfig** to enable it for the required run-levels. The `/etc/init.d/sshd` script can then be used to ensure the service is currently started.

```
# chkconfig --level 345 sshd on
# /etc/init.d/sshd start
```

To verify this operation as successful run **chkconfig --list sshd** again and check the output. It should now show the daemon as **on** at run-level **3**, **4**, and **5**.

- c. In Red Hat Enterprise Linux the default SSH daemon configuration allows remote login by the `root` user. This is also a requirement for the Red Hat Enterprise Virtualization Manager to successfully access the machine. In some cases administrator's may have disabled this ability.

To check whether or not this is the case search the `/etc/ssh/sshd_config` for the value **PermitRootLogin**. This must be done while logged in as `root`.

```
# grep PermitRootLogin /etc/ssh/sshd_config
PermitRootLogin no
```

Where `PermitRootLogin` is set to **no** the value must be changed to **yes**. To do this edit the configuration file.

```
# vi /etc/ssh/sshd_config
```

Once the updated configuration file has been saved the SSH daemon must be told to reload it.

```
# /etc/init.d/sshd reload
Reloading sshd: [ OK ]
```

The root user should now be able to access the system via SSH.

Result:

You can now add the correctly installed and configured Red Hat Enterprise Linux host to the Red Hat Enterprise Virtualization platform.

4.1.2.1.4. To Add a Host

Before adding a host, ensure you have the correct IP and password of the host. Once you have entered the requisite details, the following steps are completed by the platform:

- Virtualization capability checks.
- Installation of requisite packages.
- Creation of bridge.
- Reboot of the host.

The process of adding a new host can take some time; the progress can be followed in the Events pane.

1. Click the **Hosts** tab. The **Hosts** tab displays a list of all hosts in the system.
2. Click the **New** button. The **New Host** dialog box displays.

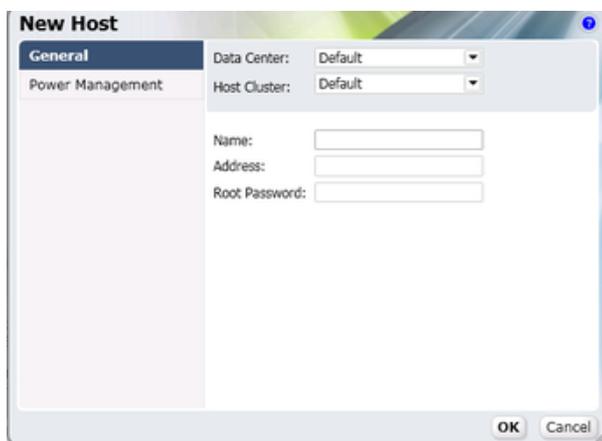


Figure 4.2. New Host Dialog Box

Enter the following details in the **General** tab:

- **Data Center**: Select the appropriate data center from the drop-down menu that you want to assign to the new host.

Note

A default data center is displayed; change the default only if you are sure that another data center can be used.

- **Host Cluster:** The cluster to which the host belongs (select from the drop-down list).
- **Name:** A descriptive name for the host.
- **Address:** The IP address, or resolvable hostname of the host (provided during installation).
- **Root password:** the password of the designated host; used during installation of the host.

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



Figure 4.3. New Host Power Management

4. **Enable Power Management:** Select this check box to turn out-of-band (OOB) power management on. If selected, the information for the following fields must also be provided:
- The **Address** of the host. This is usually the address of the remote access card (RAC) on the host.
 - A valid **User Name** for the OOB management.
 - A valid, robust **Password** for the OOB management.
 - The **Type** of the OOB management device. Select the appropriate device from the drop down list.

Table 4.2. Types of OOB management devices

alom	Sun ALOM
apc	APC
bladecenter	IBM Bladecenter Remote Supervisor Adapter
drac5	Dell Remote Access Controller for Dell computers
eps	Entry-Level Power Supply Specification

ilo	HP Integrated Lights Out standard
ipmilan	Intelligent Platform Management Interface
rsa	IBM Remote Supervisor Adaptor
rsb	Fujitsu-Siemens RSB management interface
wti	WTI PowerSwitch
Cisco_ucs	Cisco Systems

- Depending on the **Type** selected, any or all of the following fields display on the **Power Management** tab.
 - Click **Secure** to use SSH to connect to OOB management.
 - The **Port** to connect to OOB management.
 - Enter the **Slot** if a Blade server is being configured.
 - Enter any **Options** that are needed for the SSH command.
5. Click the **Test** button to test the operation of the OOB management solution. Alerts, if any, appear on the **Alerts** panel. The **Alerts** panel displays on the bottom right corner of the screen. If there are existing alerts, the alert text changes color to brighter red.

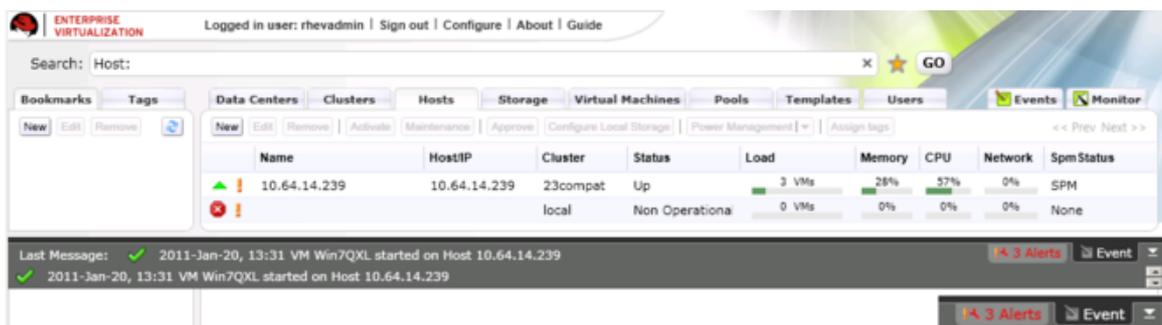


Figure 4.4. Alerts Tab

Note

The **Alerts** panel can be resized by using the **Expand/Collapse** button, or dragging the border upwards/downwards.

Important

Red Hat Enterprise Virtualization recommends the configuration of power management on the hosts. Power management enables the system to fence a troublesome host using an additional interface.



Important

If the host is required to be Highly Available, power management must be enabled and configured. Setting up power management for hosts is described in detail later in this guide. Refer [Section 13.3.1, “Setting the Parameters for Fencing”](#).

6. Click **OK**.

The new host displays in the list of hosts with a status of "Installing". Once installation is complete, the status of the newly added host is *Pending Approval*. The host must be activated for the status to change to *Up*.



Note

View the process of the host installation on the Details pane.

4.1.2.1.5. Activating a Host

After a host has been added, or an existing host has been taken down for maintenance, it needs to be activated before it can be used. Red Hat Enterprise Linux 6.1 and higher hosts need to be activated after being added or upgraded.

1. In the Hosts tab, select the host to be activated.
2. Click the **Activate** button.

The host status changes to *Up*. Virtual machines can now run on the host.

4.1.2.2. Managing Host Network Interfaces

The **Network Interfaces** tab on the **Details** pane of a host allows you to define the attachment of the logical network in the Administration Portal to the physical network interface cards (or NICs) of the host. This is a simple operation in which you attach one or more of the host's physical network interface cards (NICs) to a predefined logical network in the cluster.



Note

New logical networks cannot be defined at the host level.

The management and storage subnets are defined by default in the cluster. Typically, eth0 is allocated to the management network interface (which displays as RHEVM) and eth1 is allocated to the storage network interface (which may display as data). The Administration Portal automatically detects the attached subnets and networks, so all that is required is to match the logical network name to the correct subnet.

Each host can support up to 32 interfaces, and these are grouped by logical networks. If the default settings are not correct, or more subnets need to be added, the **Network Interfaces** tab can be used to make changes.

4.1.2.2.1. Editing Network Interfaces

You can edit the host NICs and the network using the **Network Interfaces** tab on a host's **Details** pane.

4.1.2.2.1.1. Editing Host Network Interfaces

The **Network Interfaces** tab displays the name, network name, address, MAC address, speed, and link status for each interface. The **Edit**, **Edit Management Network**, **Bond/Unbond**, **Detach** and **Save Network Configuration** buttons enable you to manage host NICs.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The **Details** pane displays.
2. Click the **Maintenance** button to migrate all virtual machines to alternative hosts, and place the host into maintenance. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.
3. Click the **Network Interfaces** tab on the **Details** pane. The **Network Interfaces** tab displays the list of NICs on the host, their address and other specifications. Select the NIC that you want to edit.
4. Click the **Edit/Add VLAN** button. The **Edit Network Interface** dialog box displays.



Figure 4.5. Host **Edit Network Interface** dialog box

- To attach the NIC to a different logical network, select a different **Network** from the list of available logical networks.
 - Select the network setting of **None**, **DHCP**, or **Static**. For **Static** setting, provide the IP, Subnet and Default Gateway information for the host.
 - Select the **Check Connectivity** check box if necessary.
 - Click **OK**.
5. Activate the host. See [Section 4.1.2.1.5, "Activating a Host"](#).

4.1.2.2.1.2. Editing Management Network

The **Network Interfaces** tab displays the name, network name, address, MAC address, speed, and link status for each interface. In the course of editing the host NICs, it may be necessary to check or edit the Management Network Interface.



Important

Communication between the Management Server and the host is via the management interface. Changing the properties of the management interface may cause the host to become unreachable.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.
2. Click the **Maintenance** button to migrate all virtual machines to alternative hosts, and place the host into maintenance. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.
3. Click the **Network Interfaces** tab on the Details pane. The **Network Interfaces** tab on the Details pane that displays the list of NICs on the host, their address and other specifications. Select the appropriate management NIC that you want to edit.
4. Click the **Edit Management Network** button. The **Edit Management Network** dialog box displays.

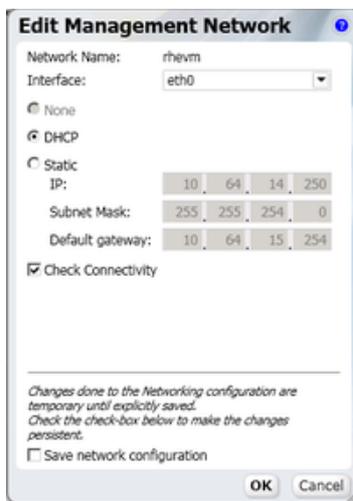


Figure 4.6. Edit Management Network Dialog Box

- To attach the RHEVM management network to a different NIC, select a different interface from the **Interface** drop-down list of available NICs.
- Select the network setting from **None**, **DHCP** or **Static**. For **Static** setting, provide the IP, Subnet and Default Gateway information for the host.
- Select the **Check Connectivity** check box if necessary.
- Select the **Save network configuration** check box to make the changes persistent, since changes done to the Networking configuration are temporary until explicitly saved.

- Click **Close**.
5. Activate the host. See [Section 4.1.2.1.5, “Activating a Host”](#).

4.1.2.2.2. Configuring Network Interfaces

After editing the NICs to ensure that the physical NICs connect to the logical networks, some further configuration may be necessary. For example, you may want to aggregate links, separate bonded links, or even detach NICs from the network. When the host is correctly configured and linked to the RHEVM network, you may want to save the network configuration.

4.1.2.2.2.1. Bonding network interfaces

Network bonding (also known as link aggregation, NIC bonding) consists of aggregating multiple network interfaces into a single logical bonded interface that correspond to a single IP address. Network bonding uses multiple network cables/ports in parallel to increase the link speed beyond the limits of any one single cable or port, and to increase the redundancy for higher availability. Red Hat Enterprise Virtualization conforms to what used to be clause 43 of IEEE 802.3-2005 Ethernet standard, usually referred to by its working group name of "IEEE 802.3ad".

Red Hat Enterprise Virtualization platform allows bonding of several NICs on a host. For example, if a host has four NICs but only two logical networks, two of the NICs can be bonded together using 802.3ad protocol to create a single channel. This channel can be mapped to a single logical network providing a higher bandwidth. To bond host NICs follow the below mentioned procedure:



Note

Ensure that the NICs have been configured correctly across the network, that is, configure your routers accordingly.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.
2. Click the **Maintenance** button to migrate all virtual machines to alternative hosts, and place the host into maintenance. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.
3. Click the **Network Interfaces** tab on the **Details** pane that displays the list of NICs on the host, their address and other specifications.
4. Select the multiple NICs that are to be bonded together.



Important

You cannot bond between NICs if the status for all is “UP” (The status arrow in the beginning is Green). To use the **Bond** feature, at least one NIC should be down (The status arrow in the beginning is Red).

5. Click the **Bond** button. The **Bond Network Interface** dialog box displays.

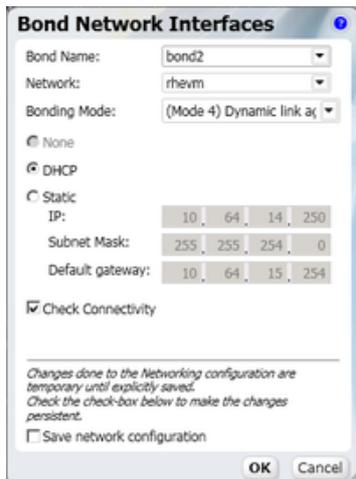


Figure 4.7. Bond Network Interface Dialog Box

- To create a bonded interface select a **Bond Name** from the list.
- Select the network setting from **None**, **DHCP** or **Static**. For **Static** setting, provide the IP, Subnet Mask and Default Gateway information for the host.
- Select the **Check Connectivity** check box if necessary.
- Select the **Save network configuration** check box to make the changes persistent, since changes done to the Networking configuration are temporary until explicitly saved.
- Click **OK**.

4.1.2.2.2. Detaching NICs

The **Network Interfaces** tab displays the name, network name, address, MAC address, speed, and link status for each interface. In the course of editing the host NICs, it may be necessary to detach a particular NIC.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.
2. Click the **Maintenance** button to migrate all virtual machines to alternative hosts, and place the host into maintenance. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.
3. Click the **Network Interfaces** tab on the **Details** pane that displays the list of NICs on the host, their address, and other specifications.
4. Select the NIC (or NICs) to be detached, and click the **Detach** button. The **Detach Network Interface** dialog box displays.



Figure 4.8. **Detach Network Interface** Dialog Box

5. The dialog box lists the NICs selected for detachment.
6. Click **OK** to confirm the detachment.
7. Activate the host. See [Section 4.1.2.1.5, “Activating a Host”](#).

4.1.2.2.3. Saving Host Network Configuration

When the host is correctly configured and linked to the network, you may want to save the network configuration.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.
2. Click the **Maintenance** button to migrate all virtual machines to alternative hosts, and place the host into maintenance. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.
3. Click the **Network Interfaces** tab on the Details pane that displays the list of NICs on the host, their address, and other specifications.
4. Click the **Save Network Configuration** button.
5. The host network configuration is saved and the following message is displayed on the task bar: “Network Changes were saved on host <Host name>.”

4.1.3. Hosts Entities

The **General** tab on the Details pane provides information on an individual host, including hardware and software versions, and whether updates are available (in the case of Hypervisor hosts).

1. Click the **Hosts** tab. If the host you want to view is not displayed, perform a search (see [Section 1.2, “Search”](#)).

A list of hosts displays. Select the appropriate host. The Details pane displays. The Details pane displays general information, network interface information and virtual machine information.

2. Click the **General** tab. The General tab displays the following information:
 - Version Information for OS, Kernel, KVM, VDSM, and SPICE.
 - Host type for example iSCSI Initiator Name.
 - Active VMs.

- Status of Memory Page Sharing (Active/Inactive) and Automatic Large Pages(On/Off).
- CPU information like number of CPUs attached, CPU name and type, total Physical Memory allocated to the selected host, Swap Size, and Shared Memory proportion.
- In addition, if an updated version of the host is available, an Alert appears.

4.1.3.1. Viewing Virtual Machines on Hosts

The **Virtual Machines** tab on the Details pane provides information on virtual machines running on the host.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.
2. Click the **Virtual Machines** tab on the Details pane. A list of virtual machines running on the host displays. This includes both virtual servers and virtual desktops. Virtual machines can be scheduled to run on the approved host. The list also displays cluster, network and display information.
3. You can **Suspend**, **Shut down** and **Migrate** a virtual machine from this tab.

4.1.3.2. Viewing Network Interfaces on Hosts

The **Network Interfaces** tab on the Details pane provides information about logical and physical networks of the host.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.
2. Click the **Network Interfaces** tab on the Details pane. A list of virtual machines running on the host displays. This allows you to define the attachment of the logical network in the Administration Portal to the physical network interface cards (or NICs) of the host. The **Network Interfaces** tab is discussed in detail in [Section 4.1.2.2, “Managing Host Network Interfaces”](#).

4.1.3.3. Viewing Network Host Hooks

The **Host Hooks** tab on the Details pane provides information about logical and physical networks of the host.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.
2. Click the **Host Hooks** tab on the Details pane. Hooks can be implemented to execute scripts when key events are triggered. When called, a hook executes the scripts in `/usr/libexec/vdsm/hooks/hook-name/` in alphabetical order. The **Host Hooks** tab displays information about Event, Script, and Property name and value. More information about VDSM Hooks can be found in Appendix B.

4.1.3.4. Viewing Permissions on Hosts

The **Permissions** tab on the Details pane provides information about user roles and their inherent permissions.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The Details pane displays.

2. Click the **Permissions** tab on the **Details** pane. The **Permissions** tab displays a list of users, their roles, and their inherited permissions.

4.1.3.5. Viewing Events on Hosts

The **Events** tab on the **Details** pane provides information about an important or unusual event.

1. Click the **Hosts** tab. A list of hosts displays. Select the appropriate host. The **Details** pane displays.
2. Click the **Events** tab on the **Details** pane. The **Events** tab displays information about any important events that a system administrator should know about, such as problems with storage or notifications that have been sent to users for unexpected events, such as when virtual machines are unexpectedly shut down.

4.1.4. Host Troubleshooting

4.1.4.1. Maintaining Hosts

You can use the Administration Portal to complete many host maintenance tasks. For example, you may have to change the network configuration details of the host, or the cluster to which it belongs. All virtual machines are migrated automatically, however the migration increases the load on the network and on other hosts. If a large number of virtual machines are running on the host (that is to be moved into maintenance mode), the migration of the virtual machines may take a considerable amount of time.



Warning

Maintaining hosts may involve the shut down, de-activation and restart of the physical host. Moving hosts into Maintenance must be planned and considered carefully.

4.1.4.1.1. Moving a Host into Maintenance Mode

Hosts must occasionally be brought down for maintenance. Red Hat Enterprise Virtualization platform attempts to migrate all the virtual machines running on the host to other hosts automatically. In some instances this may not be possible, and you may have to manually migrate or shut down a virtual machine, before the host can be placed in maintenance. Follow the below mentioned steps to move a host into maintenance mode:

1. Click the **Hosts** tab. A list of hosts displays. If the host you want to edit is not displayed, perform a search (see [Section 1.2, "Search"](#)).
2. Select the appropriate host. The **Details** pane displays information about the host.
3. Click the **Maintenance** button to migrate all virtual machines to alternative hosts, and place the host into maintenance. The **Maintenance Host(s)** dialog box appears. Click **OK**.



Figure 4.9. Host Details Pane

4. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance.
5. Perform any required tasks. When the host is ready to be reactivated, click the **Activate** button to bring the host back up.

The **Status** field of the host changes to *Up*.



Note

If Red Hat Enterprise Virtualization Manager is unable to communicate with and control the host, its status displays *Non-responsive*.

4.1.4.1.2. Editing Host Details

You can edit the details of a host, such as its name or network configuration. You can also change the cluster to which the host belongs.

Before changing the cluster that the host belongs to, you must first place it into maintenance mode (see [Section 4.1.4.1.1, “Moving a Host into Maintenance Mode”](#)). Follow this procedure to edit host details:



Warning

Maintaining hosts may involve the shut down, de-activation and restart of the physical hosts. If any virtual machines are running on the host, be aware that you may lose data and configuration details if the virtual machine have not been shut down. Moving hosts into maintenance must be carefully planned and executed with due care and consideration.

1. Click the **Hosts** tab. A list of hosts is displayed. If the host you want to edit is not displayed, perform a search (see [Section 1.2, “Search”](#)).
2. Select the host that you want to edit. Click the **Edit** button. The **Edit Host** dialog box opens.

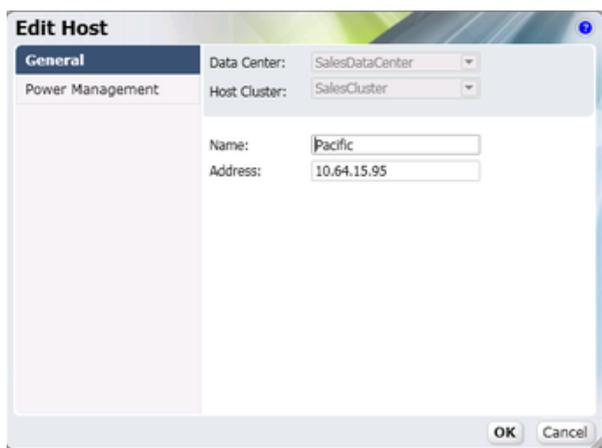


Figure 4.10. Edit Host Dialog Box

3. Edit the details as required (see [Section 4.1.2.1.3, “Adding Red Hat Enterprise Linux Hosts”](#)). Click **Save** to save the changes.

The details of the host are updated in the **Hosts** tab, and the status changes appropriately.

4.1.4.1.3. Customizing Hosts

You can assign user defined tags to objects and aggregate these objects into a group; for example you can create a group of hosts running in a department or location.

To tag a host:

1. Click the **Hosts** tab.

A list of hosts is displayed. If the host you want to edit is not displayed, perform a search (see [Section 1.2, “Search”](#)).

2. Select the appropriate host, and click the **Assign Tags** button.

The **Assign Tags** dialog box opens. It displays a list of available tags.

3. Select the required tags.
4. Click **Close**.

The tagged host displays in the result of searches for the assigned tag.

4.1.4.1.3.1. Deleting a Physical Host

Hosts that are no longer being used by the Red Hat Enterprise Virtualization platform can be permanently removed. Deleting unused hosts saves system resources, as existing hosts are checked (or *pinged*) at fixed intervals. Ensure that any virtual machines are migrated off the host, or shut down if they are no longer required. Follow this procedure to delete a host:

1. Click the **Hosts** tab. If a host that you want to delete is not displayed, perform a search (see [Section 1.2, “Search”](#)). Select the host to be deleted.
2. Place the host into *Maintenance* mode (see [Section 4.1.4.1.1, “Moving a Host into Maintenance Mode”](#)).
3. Click the **Remove** button. A confirmation message displays.

- Click **OK**. The host is removed from Red Hat Enterprise Virtualization platform and deleted from the **Hosts** tab.

4.1.5. Hosts Permissions

While the superuser or system administrator of the platform has the full range of permissions, a Host Administrator is a system administration role for a specific host only. This is a hierarchical model, it means that the Data Center Administrator and the Cluster Administrator have permissions to manage the hosts. This is useful in an enterprise where there are multiple hosts, perhaps running specific sets of virtual machines, each of which require their own system administrators. A Host Administrator has permissions for the assigned host only, not for all hosts in the cluster.

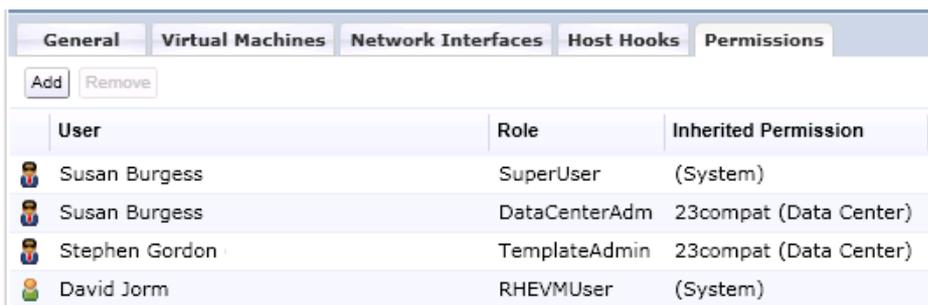
To assign a system administrator role for a host:

- Click the **Hosts** tab.

A list of hosts displays. If the required host is not visible, perform a search (see [Section 1.2, “Search”](#)).

- Select the host that you want to edit, and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.



User	Role	Inherited Permission
 Susan Burgess	SuperUser	(System)
 Susan Burgess	DataCenterAdm	23compat (Data Center)
 Stephen Gordon	TemplateAdmin	23compat (Data Center)
 David Jorm	RHEVMUser	(System)

Figure 4.11. Hosts Permissions

- Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
- Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **HostAdmin**.

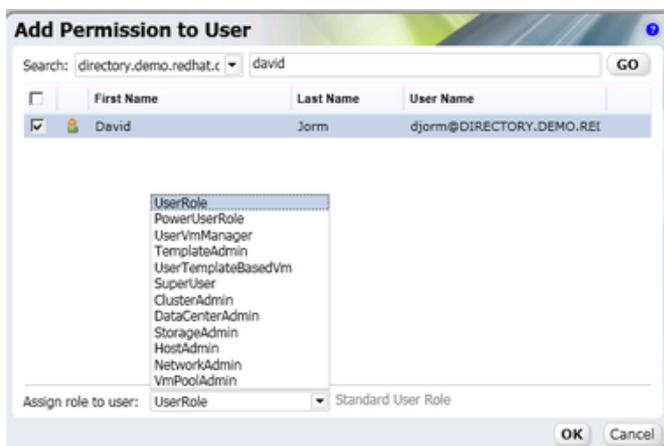


Figure 4.12. Assign HostAdmin Permission

5. Click **OK**.

The name of the user displays in the **Permissions** tab, with an icon and the assigned Role.



Note

You can only assign roles and permissions to existing users. See [Chapter 5, Users](#).

You can also change the system administrator of a host, by removing the existing system administrator, and adding the new system administrator, as described in the previous procedure.

To remove a system administrator role:

1. Click the **Hosts** tab. A list of hosts displays. If the required host is not visible, perform a search (see [Section 1.2, "Search"](#)).
2. Select the required host and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any. The Super User, Data Center Administrator and Cluster Administrator, if any, will display in the Inherited Permissions tab. However, none of these higher level roles can be removed.

3. Select the check box of the appropriate user.
4. Click **Remove**. The user is removed from the Permissions tab.

Users

This section describes the types of users in Red Hat Enterprise Virtualization Manager, how to set up user roles that control user permission levels, and how to manage users on the Red Hat Enterprise Virtualization platform. Red Hat Enterprise Virtualization relies on directory services for user authentication and information. Currently the two supported providers of directory services for use with the Red Hat Enterprise Virtualization Manager are Identity, Policy, and Audit (IPA) and Microsoft Active Directory.

There are two basic types of users in the Red Hat Enterprise Virtualization platform, end users who use and manage the virtual machines, and administrative users who are responsible for the supply of virtual machines and virtual infrastructure to the end users. Users are assigned roles that allow them to perform their tasks as required. The role with the highest level of permissions is the admin role, which allows a user to set up, manage, and optimize all aspects of the Red Hat Enterprise Virtualization platform. By setting up and configuring roles with permissions to perform actions and create objects, users can be provided with a range of permissions that allow the safe delegation of some administrative tasks to users without granting them complete administrative control.

Red Hat Enterprise Virtualization Manager provides a rich user interface that allows an administrator to manage their virtual infrastructure from a web browser allowing even the most advanced configurations such as network bonding and VLANs to be centrally managed from a graphical console.



Note

Users are not created in Red Hat Enterprise Virtualization platform, but in the Directory Services domain. Red Hat Enterprise Virtualization Manager can be configured to use multiple Directory Services domains.

5.1. Authorization Model

Red Hat Enterprise Virtualization applies authorization controls to each action performed in the system. Authorization is applied based on the combination of the three components in any action:

- The user performing the action
- The type of action being performed
- The object on which the action is being performed

Actions

For an action to be successfully performed, the user must have the appropriate permission for the object being acted upon. Each type of action corresponds to a permission. There are many different permissions in the system, so for simplicity they are grouped together in roles.

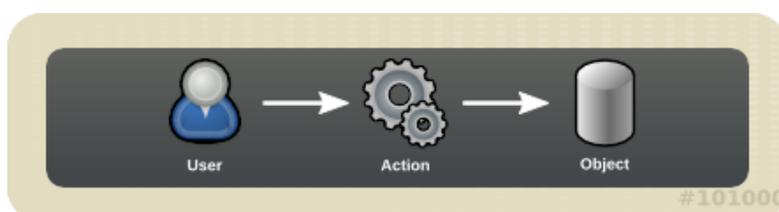


Figure 5.1. Actions

Permissions

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

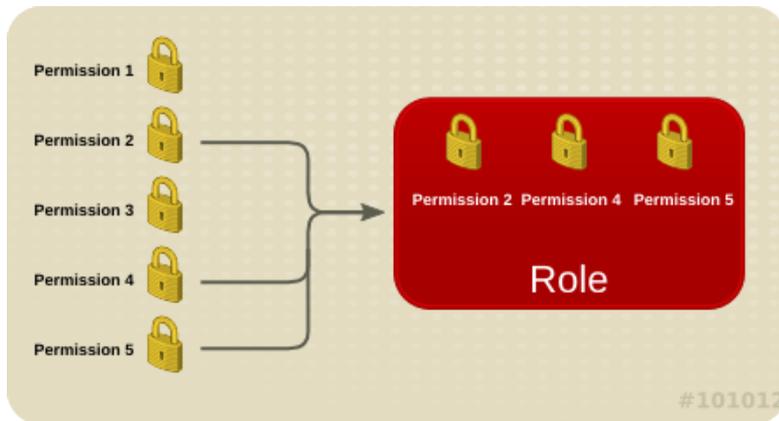


Figure 5.2. Permissions & Roles

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.

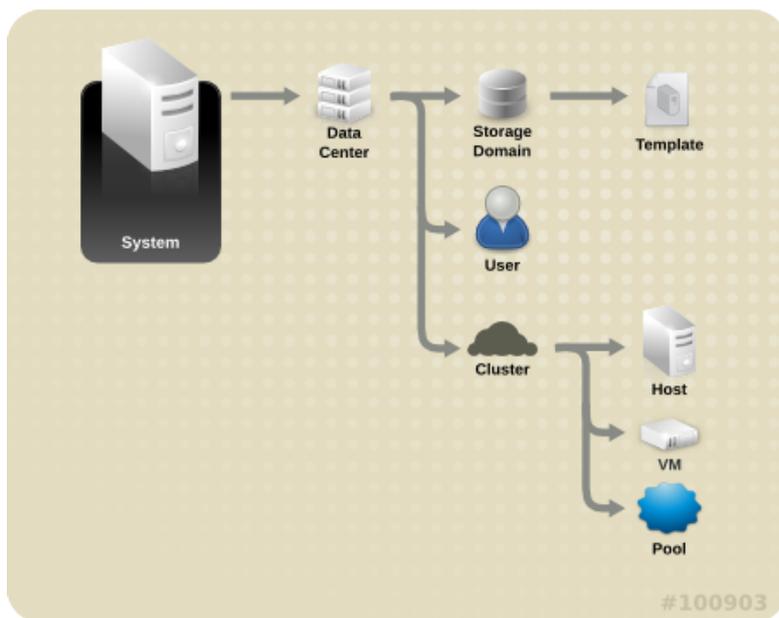


Figure 5.3. Red Hat Enterprise Virtualization Object Hierarchy



Important — Actions can impact multiple objects

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.

5.2. User Properties

Roles and Permissions can be considered as the properties of the User object. Roles are predefined sets of privileges that can be configured from Red Hat Enterprise Virtualization Manager, permitting

access and management to different levels of resources in the data center, to specific physical and virtual resources. Multi-level administration includes a hierarchy of permissions that can be configured to provide a finely grained model of permissions, or a wider level of permissions as required by your enterprise. 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 log into and 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, effectively acting as system administrator on the virtual machine.

5.2.1. Roles

Red Hat Enterprise Virtualization platform provides a range of pre-configured or default roles, from the Superuser or system administration of the platform, to an end user with permissions to access a single virtual machine only. There are two types of system administration roles, roles with system permissions to physical resources, such as hosts and storage; and roles with system permissions to virtual resources such as virtual machines and pools. While you cannot change the default roles, you can clone them, and then customize the new roles as required.

There are two types of roles in Red Hat Enterprise Virtualization, **administrator** roles and **user** roles. The privileges provided by these roles are shown in this section.



Note

The default roles cannot be removed from the platform, or privileges cannot be modified, however the name and descriptions can be changed.

Role Types

- **Administrator Role:** Allows access to the *Administration Portal* for managing virtual resources. An administrator role does not confer any permissions for the User Portal.
- **User Role:** Allows access to the *User Portal* for managing and accessing virtual machines. A user role does not confer any permissions for the Administration Portal.

For example, if a user has an **administrator** role on a cluster, they could manage all virtual machines in the cluster using the *Administration Portal*. They could not access any of these virtual machines in the *user portal*; this would require a **user** role.

Table 5.1. Red Hat Enterprise Virtualization User Roles

Role	Privileges	Notes
UserRole	User privileges	View resource state and details. View all the resource tabs. Can perform basic operations on the virtual machine and connect to the virtual machine. This role also has permissions to use a virtual machine in a pool.
PowerUserRole	Allowed to create and manage virtual machines and templates	Can change the CD and add, remove, and set access privileges for all the users and groups, for all physical and

Role	Privileges	Notes
		virtual resources in the data center.
UserVmManager	Access to Virtual Machines and Pools.	Level of privileges allow the user to administer virtual machines, including configuring network and storage, manipulating snapshots and migrating virtual machines. System administrator of a virtual machine.
UserTemplateBasedVm	Limited privileges to only use Templates	Level of privilege to create a virtual machine by means of a template.

Table 5.2. Red Hat Enterprise Virtualization System Administrator Roles

Role	Privileges	Notes
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, cluster or host. A network administrator of a data center or cluster will have inherited network permissions for virtual pools within the cluster as well.
VMPoolAdmin	System Administrator role of a virtual pool.	Can create, delete, and configure a virtual pool, and perform basic operations on a virtual machine.
TemplateAdmin	Can perform all operations on templates.	Has privileges to create, delete and configure storage domains as well as network details in addition to moving templates between domains.
SuperUser	Full permissions across all objects and levels	Can manage all objects across all data centers.
ClusterAdmin	Cluster Administrator	Can use, create, delete, manage all physical and virtual resources in a specific cluster, including hosts, templates and virtual machines.
DataCenterAdmin	Data Center Administrator	Can use, create, delete, manage all physical and virtual resources within a specific data center, including clusters, hosts, templates and virtual machines.

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.
StorageAdmin	Storage Administrator	Can create, delete, configure and manage a specific storage domain.

5.2.2. Permissions

The following table details the actions for each object in the data center, for each of which permission may be assigned. This results in a high level of control over actions at multiple levels.

Table 5.3. Permissions Actions on Objects

Object	Action
System - Configure RHEV-M	Manipulate Users, Manipulate Permissions, Manipulate Roles, Generic Configuration
Data Center - Configure Data Center	Create, Delete, Edit Data Center Properties, Edit Network
Storage - Configure Storage Domain	Create, Delete, Edit Storage Domain Properties, Manipulate Status
Cluster - Configure Cluster	Create, Delete, Edit Cluster Properties, Edit Network
Host - Configure Host	Create, Delete, Edit Host Properties, Manipulate Status, Edit Network
Template - Basic Operations	Edit Template Properties, Edit Network
Template - Provisioning Operations	Create, Delete, Import/Export, Copy Templates
VM - Basic Operations	Basic Operations, Change CD, Remote Log In
VM - Provisioning Operations	Edit Properties, Create, Delete, Import/Export Virtual Machines, Edit Network, Edit Storage, Edit Snapshots
VM - Administration Operations	Move VM, Migrate
VM Pool - Basic Operations	Basic Operations
VM Pool - Provisioning Operations	Create, Delete, Edit Properties

5.3. Users Operations

Users can be added or removed from the system, and assigned roles, and given permissions to various objects, enabling them to effectively perform their required work. The **Users** Details pane displays information on the status and privileges of users, enabling the system administrator to assign or change roles, allot virtual machines, set up event notifications and allocate Directory Service groups. Because of the level of detail that is possible, a multi-level administration system can be defined.



Note

Login to the system is verified against the Directory Service records of the organization.

5.3.1. Adding Users and Groups

Existing users must be added to the Administration Portal before being assigned roles, or allotted desktops.

Adding Users

1. Click the **Users** tab. The list of authorized users for Red Hat Enterprise Virtualization platform displays.
2. Click **Add**. The **Add Users and Groups** dialog box displays.

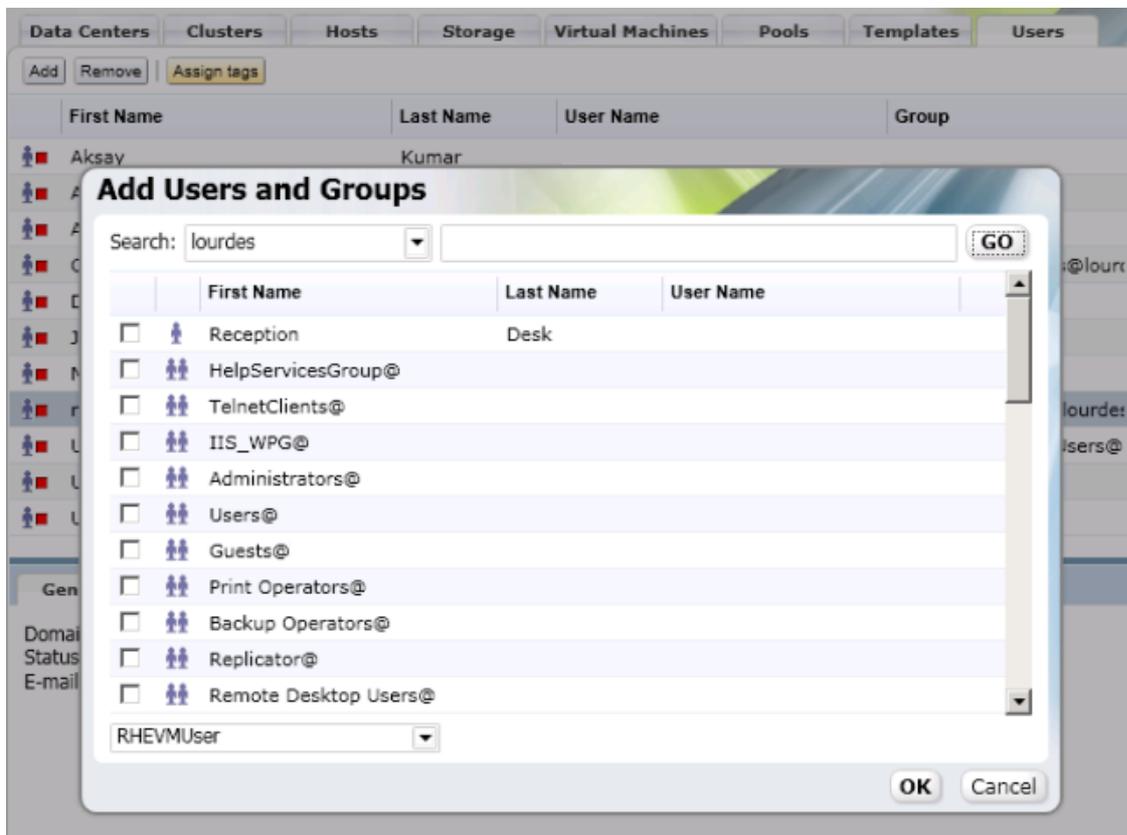


Figure 5.4. Add Users and Groups Dialog Box

3. The default **Search** domain displays. If there are more than one search domains, select the appropriate search domain. Enter a name or part of a name in the search text field, and click **GO**. Alternatively, click **GO** to view a list of all users and groups.
4. Select the group, user or users check boxes. The added user displays on the **Users** tab.

Viewing User Information

Users are not created from within the platform, Red Hat Enterprise Virtualization Manager accesses user information from the organization's Directory Service. This means that you can only assign roles

to users who already exist in your Directory Services domain. To assign permissions to users, use the **Permissions** tab on the Details pane of the relevant resource.

Example 5.1. Assigning a user permissions to use a particular virtual machine

To assign a user to a particular virtual machine, use the **Permissions** tab on the Details pane of the selected virtual machine. See [Section 6.5, “Managing Permissions to Virtual Machines”](#).

To view general user information:

1. Click the **Users** tab. The list of authorized users for Red Hat Enterprise Virtualization platform displays.
2. Select the user, or perform a search if the user is not visible on the results list.
3. The Details pane displays for the selected user, usually with the **General** tab displaying general information, such as the domain name, email and status of the user.
4. The other tabs allow you to view groups, permissions, and events for the user.

For example, to view the groups to which the user belongs, click the **Directory Groups** tab.

5.4. Users Troubleshooting

When troubleshooting problems related to users, the first thing to recall is that users must be correctly added to and authenticated at the Directory Services level, not on the Red Hat Enterprise Virtualization platform. Problems with permissions can occur when adequate levels of permissions to all required objects have not been assigned. Users, particularly those with administrator roles require to be notified when events or triggers occur.



Important — Actions can impact multiple objects

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.

5.4.1. Managing Event Notifiers

This section describes how to set up and manage event notifications for users. Events are displayed on the **Events** tab, however, users can be notified by email about selected events. For example, a system administrator might like to know when there is a problem with storage, or a team lead may want to be notified if virtual machines shut down.

To set up event notifications:

1. Click the **Users** tab. The list of authorized users for Red Hat Enterprise Virtualization platform displays.
2. Select the user who requires notification, or perform a search if the user is not visible on the results list.

Click the **Event Notifier** tab. The **Event Notifier** tab displays a list of events for which the user will be notified, if any.

- Click the **Manage Events** button. The **Add Event Notification** dialog box displays a list of events, for Hosts, Storage, Virtual Machines and General Management events. You can select all, or pick individual events from the list. Click the **Expand All** button to see complete lists of events.

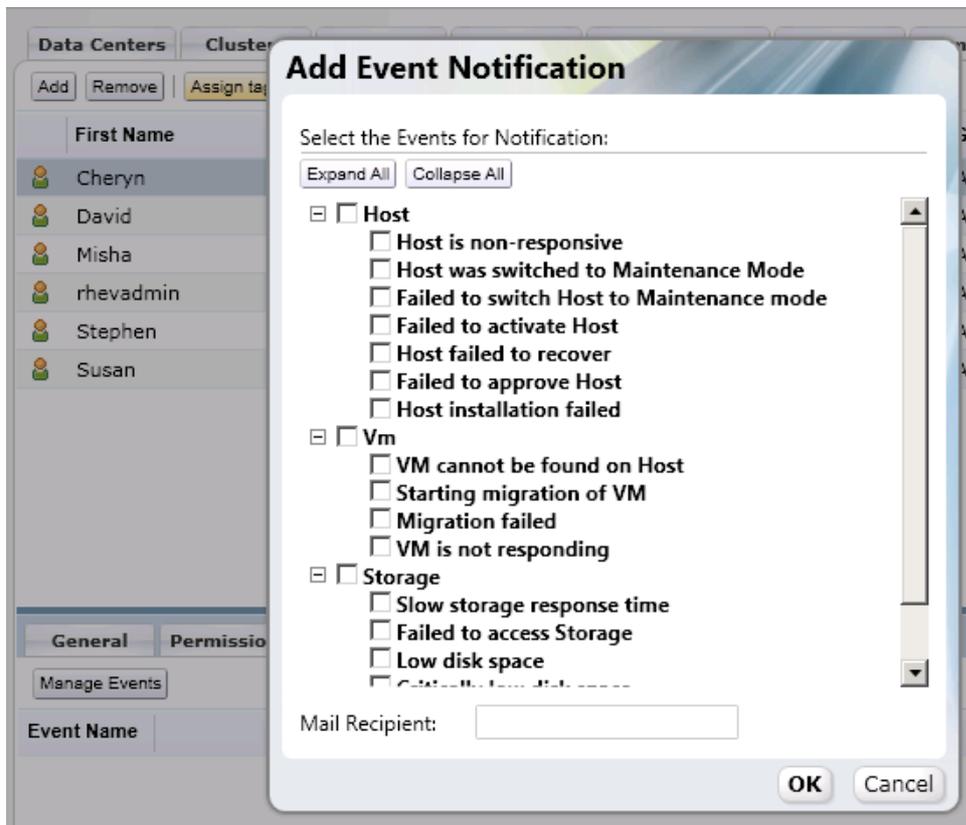


Figure 5.5. The Add Events Dialog Box

- Click **OK**. The selected events display on the **Event Notifier** tab for the user.

To cancel event notification:

- In the **Users** tab, select the user or the user group.
- Select the **Event Notifier** tab. The details pane displays the events for which the user will receive notifications.
- Click the **Manage Events** button. The **Add Event Notification** dialog box displays a list of events, for Hosts, Storage, Virtual Machines and General Management events. To remove an event notification, deselect events from the list. Click the **Expand All** button to see the complete lists of events.
- Click **OK**. The deselected events are removed from the display on the **Event Notifier** tab for the user.

5.4.2. Removing Users

A system administrator will need to remove users, for example, when they leave the company.



Note

A user can only be removed, if all virtual machines have been detached from the user.

To remove a user:

1. Click the **Users** tab. The list of authorized users for Red Hat Enterprise Virtualization platform displays.

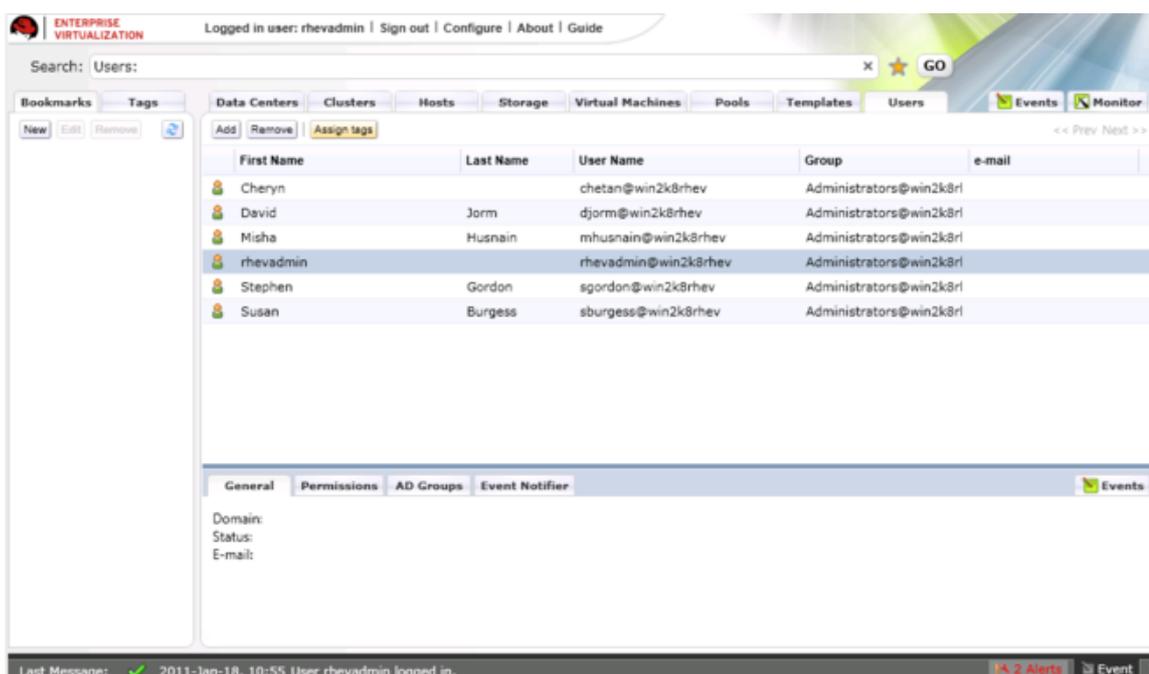


Figure 5.6. Users Tab

2. Select the user to be removed.
3. Click the **Virtual Machines** tab in the Details pane. If the user is running any virtual machines, remove the virtual machines from the user by clicking the **Remove** button on the Virtual Machines tab. See [Section 6.5, "Managing Permissions to Virtual Machines"](#).
4. Click the **Remove** button. A message displays prompting you to confirm the removal. Click **OK**.
5. The user is removed from Red Hat Enterprise Virtualization



Note

All user information is read from the directory service. Removing a user from the Red Hat Enterprise Virtualization system deletes the record in the Red Hat Enterprise Virtualization database, denying the user the ability to log on to the desktop. It removes the association in the Directory Service between the desktop and the user. All other user properties remain intact.

Part III. Managing Virtualization Infrastructure

Managing Virtual Resources

The Red Hat Enterprise Virtualization platform supports both virtual servers and virtual desktops. The **Virtual Machines** tab on the Administration Portal provides an efficient graphical way to view and manage virtual machines. For more information on virtual machines, virtual disk formats and storage of guest images, see *Red Hat Enterprise Linux Virtualization Guide*.

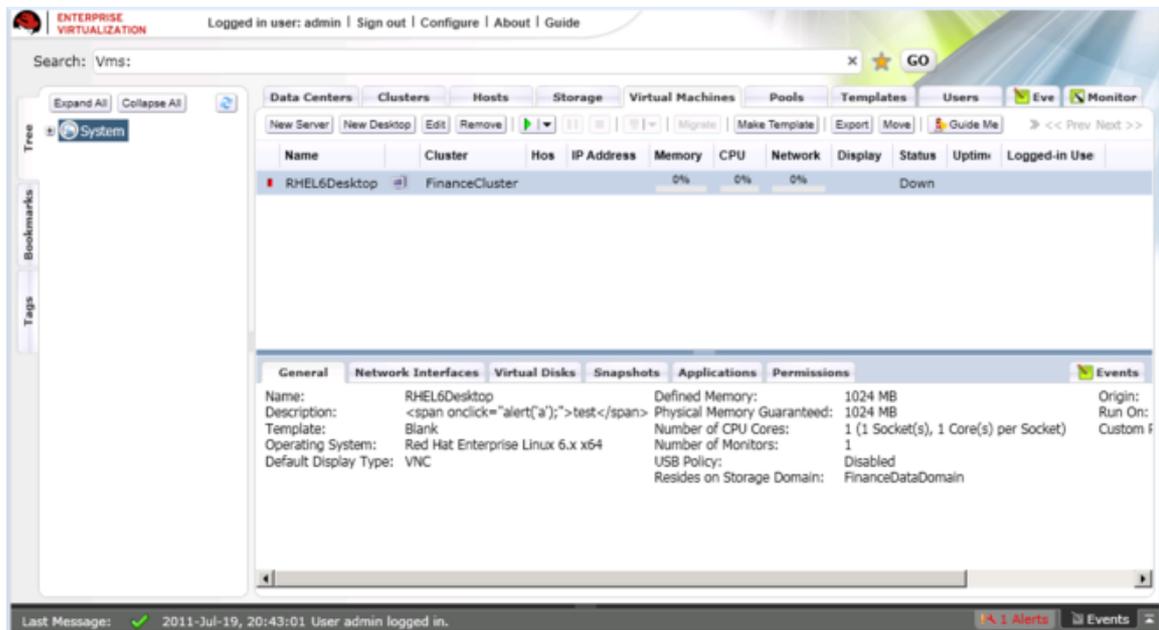


Figure 6.1. The Virtual Machines Tab

Administrative tasks for virtual machines include:

- Creating virtual machines manually or from templates.
- Starting, suspending and migrating virtual machines.
- Backing up and restoring virtual machines by taking a snapshot.
- Importing or exporting virtual machines.
- Converting virtual machines from foreign hypervisors.
- Assigning virtual machines to users.

This chapter describes how to create and maintain virtual machines. A virtual desktop fulfills the tasks of a physical desktop without the actual hardware. Virtual machines in a cluster can be migrated to other hosts within the same cluster. An understanding of how virtual machines access networked storage is helpful.

6.1. About Virtual Machines

This section briefly describes the storage, processing and network parameters of virtual machines in the Red Hat Enterprise Virtualization platform.

It is recommended that the overcommit memory feature be used in order to achieve density. However this may result in some virtual machines running with a 'Not enough memory' error message. This is a transient state, running the virtual machines again should succeed. For more information on setting memory overcommit, see [Creating a New Host Cluster](#).

6.1.1. Supported Virtual Machines

Red Hat Enterprise Virtualization presently supports the following virtual machines:

Supported Guests

Red Hat Enterprise Virtualization presently supports the following virtualized guest operating systems:

- Red Hat Enterprise Linux 3 (32 bit and 64 bit)
- Red Hat Enterprise Linux 4 (32 bit and 64 bit)
- Red Hat Enterprise Linux 5 (32 bit and 64 bit)
- Red Hat Enterprise Linux 6 (32 bit and 64 bit)
- Windows XP Service Pack 3 and newer (32 bit only)
- Windows 7 (32 bit and 64 bit)
- Windows Server 2003 Service Pack 2 and newer (32 bit and 64 bit)
- Windows Server 2008 (32 bit and 64 bit)
- Windows Server 2008 R2 (64 bit only)

Para-virtualized drivers (the virtio drivers) that increase the performance for a virtual machine's block and network devices are available for the following operating systems and versions.

Table 6.1. Para-virtualized driver support

Guest operating system	Para-virtualized drivers
Red Hat Enterprise Linux 4.8 and newer (32 bit and 64 bit)	Block and network drivers
Red Hat Enterprise Linux 5.4 and newer (32 bit and 64 bit)	Block and network drivers
Red Hat Enterprise Linux 6.0 and newer (32 bit and 64 bit)	Block and network drivers
Windows XP	Block and network drivers
Windows 7 (32 bit and 64 bit)	Block and network drivers
Windows Server 2003 R2 (32 bit and 64 bit)	Block and network drivers
Windows Server 2008 (32 bit and 64 bit)	Block and network drivers



Important

If a network interface on a Windows virtual machine is configured using the default network drivers, the network configuration settings are lost if the Red Hat Enterprise Virtualization para-virtualized network drivers are installed subsequently. To avoid this issue, you should install the Red Hat Enterprise Virtualization para-virtualized network drivers before configuring network interfaces on Windows virtual machines.

6.1.2. Virtual Machine Performance Parameters

Red Hat Enterprise Virtualization virtual machines can support the following parameters:

Table 6.2. Supported virtual machine parameters

Parameter	Number	Note
Virtualized CPUs	16	per virtual machine
Virtualized RAM	256GB	For a 64 bit virtual machine
Virtualized RAM	4GB	per 32 bit virtual machine. Note, the virtual machine may not register the entire 4GB. The amount of RAM that the virtual machine recognizes is limited by its operating system.
Virtualized storage devices	8	per virtual machine
Virtualized network interface controllers	8	per virtual machine
Virtualized PCI devices	32	per virtual machine

6.1.3. Understanding Virtual Machine Storage

Red Hat Enterprise Virtualization platform 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 images data, and meta-data and the pool's meta-data.

In an NFS data center, the SPM creates the virtual disk on top of a regular file system as a normal disk in preallocated (RAW) format. Where sparse allocation is chosen additional layers on the disk will be created in thinly provisioned Qcow2 (sparse) format. For iSCSI and SAN, the SPM creates a Volume group (VG) on top of the Logical Unit Numbers (LUNs) provided. During the virtual disk creation, either a preallocated format (RAW) or a thinly provisioned Qcow2 (sparse) format is created.

For a virtual disk with a preallocated format, a Logical Volume (LV) of the specified size in GB is created. If necessary, the virtual machine (VM) can be mounted on a Red Hat Enterprise Linux server using `kpartx`, `vgscan`, `vgchange` and `mount` to investigate the VM's processes or problems.

For a virtual disk with a thin provision format, a 512M LV is created initially. The LV is continuously monitored by the host on which the VM is running. As soon as the usage nears a threshold the host notifies the SPM, and the SPM extends the LV by 512M. The host is responsible for resuming the VM after the LV has been extended. If the VM goes into a pause state it means that the SPM could not extend the disk on time. This can occur if the SPM is too busy or there is not enough storage space.

From a performance point of view, a virtual disk with a preallocated (RAW) format is significantly faster than a virtual disk with a thin provisioning (Qcow2) format. It is recommended that the thin provision format be used for non-IO intensive virtual desktops.

6.1.3.1. Understanding Virtual Disks

Virtual disks are of two types, Sparse or Preallocated, and each behaves in a different manner. The available formats are either Raw or Qcow2.

1. Sparse or Preallocated

A preallocated virtual disk has reserved storage of the same size as the virtual disk itself. This results in better performance because no storage allocation is required during runtime.

On SAN (iSCSI, FCP) this is achieved by creating a block device with the same size as the virtual disk. On NFS this is achieved by filling the backing hard disk image file with zeros. Preallocating

storage on an NFS storage domain presumes that the backing storage is not Qcow2 formatted and zeroes will not be de-duplicated in the hard disk image file. (If these assumptions are incorrect, do not select Preallocated for NFS virtual disks).

For sparse virtual disks backing storage is not reserved and is allocated as needed during runtime. This allows for storage over commitment under the assumption that most disks are not fully utilized and storage capacity can be utilized better. This requires the backing storage to monitor write requests and can cause some performance issues. On NFS backing storage is achieved simply by using files. On SAN this is achieved by creating a block device smaller than the virtual disk's defined size and communicating with the hypervisor to monitor necessary allocations. This does not require support from the underlying storage devices.

2. Raw

For raw virtual disks the backing storage device (file/block device) is presented as is to the virtual machine with no additional layering in between. This gives better performance but has several limitations.

The possible combinations of storage types and formats are described in the following table.

Table 6.3. Permitted Storage Combinations

Storage	Format	Type	Note
NFS or iSCSI/FCP	Raw or Qcow2	Sparse or Preallocated	
NFS	Raw	Preallocated	A file with an initial size which equals the amount of storage defined for the virtual disk, and has no formatting.
NFS	Raw	Sparse	A file with an initial size which is close to zero, and has no formatting.
NFS	Qcow2	Sparse	A file with an initial size which is close to zero, and has RAW formatting. Subsequent layers will be Qcow2 formatted.
SAN	Raw	Preallocated	A block device with an initial size which equals the amount of storage defined for the virtual disk, and has no formatting.
SAN	Qcow2	Preallocated	A block device with an initial size which equals the amount of storage defined for the virtual disk, and has qcow2 formatting.

Storage	Format	Type	Note
NFS or iSCSI/FCP	Raw or Qcow2	Sparse or Preallocated	
SAN	Qcow2	Sparse	A block device with an initial size which is much smaller than the size defined for the VDisk (currently 1GB), and has Qcow2 formatting for which space is allocated as needed (currently in 1GB increments).

6.2. Creating New Virtual Machines

You can create a virtual machine in several ways:

- From an existing template.
- From the blank template. This is the same as creating a virtual machine from scratch.
- As a clone from an existing template.



Note

It is best practice to create a virtual desktop from an existing template. This is the easiest way to ensure that network, storage, processing and memory is set up correctly and efficiently. It is also the best way to set up multiple identical virtual desktops. Using identical virtual desktops is also the best way to make the most efficient use of storage.

6.2.1. Creating Virtual Machines from Existing Templates

You can create a virtual machine from an existing template (either created by you, or one that came with the system). A template is a base virtual machine that is set with a unique configuration and settings. A virtual machine that is based on a particular template acquires the configurations and settings of the template. Thus, templates are used to conveniently and efficiently create a set of identical virtual machines. By default, virtual machines created from templates use *thin provisioning*. In the context of templates, thin provisioning means that all virtual machines based on a given template share the same base image as the template and must remain on the same data domain as the template.

For a complete list of all the required settings to create a virtual machine, see [Section 6.2.2, “Creating New Virtual Machines without a Template”](#).

To create a new virtual machine from an existing template:

1. Click the **Virtual Machines** tab. The Virtual Machines tab displays a list of existing virtual machines.

The icon to the right of the virtual machine name indicates whether it is a virtual server, a desktop or a part of a desktop pool.



Figure 6.2. Virtual Machine List

2. Click the **New Desktop** button.

The **New Desktop Virtual Machine** dialog box displays.

3. Select the **Data Center** and **Host Cluster** on which the desktop is to run. All templates that exist in the selected cluster display in the **Based on Template** list. Select an existing template.

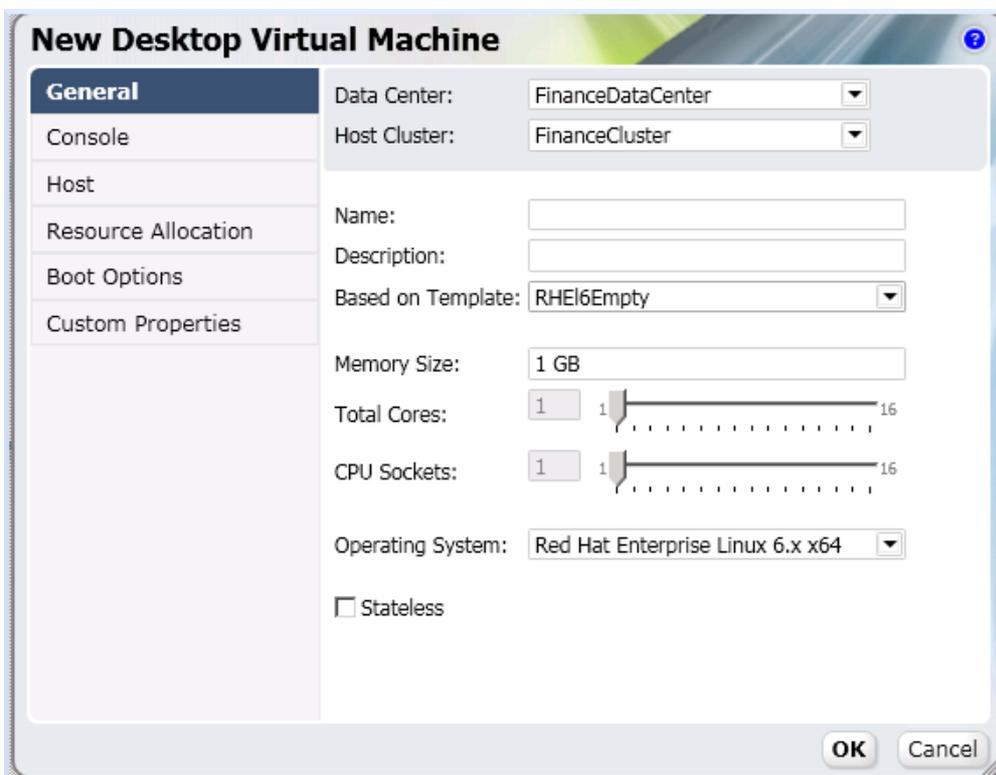


Figure 6.3. New Virtual Machine Dialog Box

4. Enter a suitable **Name** and appropriate **Description**, and accept the default values inherited from the template in the rest of the fields. You can change them if needed. See [Table 6.4, "New Virtual Machine Dialog Box Fields"](#) for field descriptions.
5. If you select a Windows Operating System, an additional **Windows Sys Prep** group displays on the **New Virtual Machine** dialog box.
6. Click **Stateless** if the virtual machine is to run in stateless mode. A stateless desktop is always created from the base template, and deleted on shutdown. This type of virtual machine is very useful when creating virtual machines that need to be used for a short time, or by temporary staff
7. Click **OK** to create the virtual machine. The virtual machine displays in the Virtual Machines list.

**Note**

It may take some time for the virtual machine to be created. During this time, the status of the virtual machine displays as **Image Locked**, followed by **Down**.

6.2.2. Creating New Virtual Machines without a Template

The Red Hat Enterprise Virtualization platform allows you to create a number of different types of virtual machines.

To create a new virtual machine from a blank template:

1. Click the **Virtual Machines** tab.

The **Virtual Machines** tab displays the existing virtual machines.

The icon to the right of the virtual machine name indicates whether it is a virtual server, a desktop or a part of a desktop pool. See [Figure 6.2, “Virtual Machine List”](#).

2. Click the **New Desktop** button.

The **New Desktop Virtual Machine** dialog box displays with the **General** tab. You only need to fill in the **Name** and **Operating System** fields. You can accept the default settings for other fields, or change them if required. If mandatory information is not entered, on clicking **OK**, the unfilled mandatory fields display with a colored border.

3. Enter information in the **General** fields of the **New Virtual Machine** dialog box:

Table 6.4. New Virtual Machine Dialog Box Fields

Field	Description	Notes
Data Center	Select an existing Data Center from the list.	The <i>Default</i> data center displays by default.
Host Cluster	The name of the host cluster to which the virtual machine is attached. It can be hosted on any physical machine in the cluster depending on the policy rules.	This is the migration domain for the virtual machine. The <i>Default</i> cluster displays by default.
Name	The name of new virtual machine. Ensure it is a unique name.	A virtual machine name must not contain any spaces, and must contain at least one character a-z. The maximum length of a virtual machine name is 15 characters. Follow the operating system's rules for virtual machine names.
Description	A meaningful description of the new virtual machine.	
Template	Select <i>Blank</i> (the default) to create a virtual machine from scratch.	Select an existing template to create a virtual machine from an existing model. See

Field	Description	Notes
		Section 6.2.1, "Creating Virtual Machines from Existing Templates"
Memory Size (MB)	The amount of memory assigned to the virtual machine.	Consider the processing and storage needs of the applications that are intended to run on the virtual machine. The maximum allowable memory for a virtual machine is 256GB, allowing even the most memory-intensive enterprise workloads to be virtualized. The total amount of memory allocated to Virtual Machines is able to exceed the amount of physical memory available to the host where memory over-commit is enabled.
Total Cores	The processing power allocated to the virtual machine, as CPU Cores, from 1 to 16 on the slider bar.	It is recommended that you do not assign too high a number to a single Virtual Machine, or more cores in total than actually exist on the physical host.
CPU Sockets	The number of CPU sockets for the virtual machine from 1 to 16 on the slider bar.	It is recommended that you do not assign too high a number to a single Virtual Machine, or more CPUs in total than actually exist on the physical host.
Operating System	The operating system. Valid values include a range of Windows and Linux variants.	This is a display only field, as no operating system is actually installed during this process.
Stateless	The virtual machine is to run in stateless mode. The stateless mode is mostly used for virtual desktops. A stateless desktop or server is always created from the base template, and deleted on shutdown. Every time the virtual machine is run, a new instance of the virtual machine is created from the base template. This type of virtual machine is very useful when creating virtual machines that	

Field	Description	Notes
	need to be used for a short time, or by temporary staff.	

- If the **Operating System** chosen is **Windows**, the **Windows Sys Prep** group displays. Enter the **Domain** and **Time Zone** in which the virtual machine is to run. This is the time zone for the virtual machine, and not necessarily the time zone for the physical host on which the virtual machine is running.
- On the **Console** tab, enter the **Protocol**, **USB Policy** and number of **Monitors**.

Table 6.5. New Virtual Machine Dialog Box Fields

Field	Description	Notes
Protocol	Define the display protocol to be used. Select either: <ul style="list-style-type: none"> • SPICE • VNC 	Select SPICE for Windows or Linux virtual machines. This is the recommended protocol. Optionally, select VNC for Linux virtual machines if desired.
USB Policy	Select Enabled or Disabled to indicate whether a USB device can be inserted into the client machine.	

- Select the **Host** tab. Define the host for the virtual machine to **Run On** by selecting either **Any Host in Cluster** if the virtual machine can start and run on any available host in the cluster, or **Specific** if the virtual machine must run on a particular host in the cluster. Select the specific host from the list of available hosts in the cluster.

Define further options for running the virtual machine using the **Run/Migration Options**. Either select **Run VM on Host (no migration allowed)**, to start and run the virtual machine only on the specified host, or select **Allow VM migration only upon Administrator specific request** (system will not trigger automatic migration of this VM) to prevent migration in mid-operation to other hosts in the cluster, for example, in case of overload or fencing of the host. However, the virtual machine may start on any host.

- Enter **Resource Allocation** for the virtual machine if required. You can define the Storage Allocation by selecting an attached **Storage Domain** from the list, as well as the Memory Allocation by entering a value in the **Physical Memory Guaranteed** field.
- Enter boot information for the virtual machine in the **Boot Options** tab of the **New Desktop Virtual Machine** dialog box:

Table 6.6. New Virtual Machine Dialog Box Fields

Field	Description	Notes
First Device	<ul style="list-style-type: none"> • Hard Disk • CD-ROM • Network (PXE) 	After installing a new virtual machine, the new virtual machine must go into Boot mode before powering up. Select the first device that the virtual machine must try to boot:

Field	Description	Notes
		<ul style="list-style-type: none"> • Hard Disk to boot from the hard disk (though if this is a blank virtual machine, it will obviously not boot from the hard disk) • CD-ROM to boot from the CD • Network (PXE) to boot from the network.
Second Device	<ul style="list-style-type: none"> • Hard Disk • CD-ROM • Network (PXE) 	Select the second device for the virtual machine to use to boot if the first device is not available. The first device selected in the previous option does not appear in the options.
Attach CD	A list of available CD-ROMs appear if Attach CD is selected.	Select the appropriate operating system ISOs available in the ISO domain.

9. Enter **Custom Properties** for the virtual machine if required.

10. Click **OK**.

If all the mandatory fields have been selected, The **New Virtual Machine - Guide Me** dialog box displays. If not, the dialog box does not close, and unfilled fields are indicated with a red border. Complete all the mandatory fields.

You can use the buttons in the **New Virtual Machine - Guide Me** dialog box immediately, or the tabs on the Details Pane to complete the configuration. Click **Configure Later**. The new virtual machine is created and displays in the list of virtual machines with the Virtual Desktop icon and Status Down icon.

6.2.3. Cloning Virtual Machines from Existing Templates

You can clone a virtual machine from an existing template (either created by you, or one that came with the system). A template is a base virtual machine that is set with a unique configuration and settings. A virtual machine that is cloned from a particular template acquires the configurations and settings of the template. A virtual machine that has been provisioned as a clone includes a complete copy of the template's hard disk image, removing the requirement that the virtual machine stay in the same data domain as the template.

To create a cloned virtual machine from an existing template:

1. Click the **Virtual Machines** tab.

The **Virtual Machines** tab displays a list of existing virtual machines.

2. Click the **New Server** button.

The **New Server Virtual Machine** dialog box displays.

3. Select an existing template from the **Based on Template** list. All templates that exist in the cluster display in the list.
4. Enter a **Name** and appropriate **Description**, and accept the default values inherited from the template in the rest of the fields. You can change them if needed. See [Table 6.4, “New Virtual Machine Dialog Box Fields”](#) for field descriptions.
5. In the **Resource Allocation** tab, on the **Provisioning** field, select *Clone* from the list.

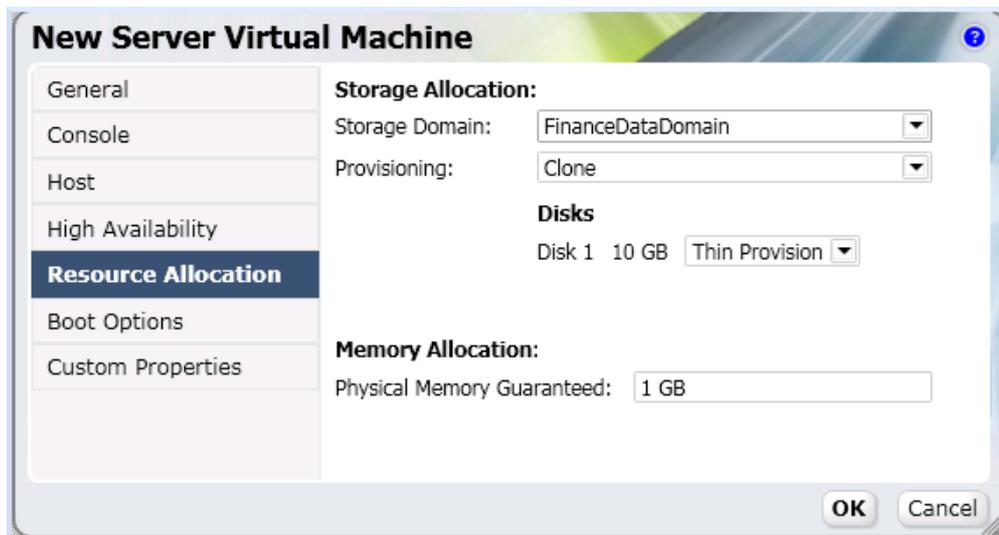


Figure 6.4. Provisioning - Clone

Select the disk provisioning mode in the **Disks** field. This selection impacts both the speed of the clone operation and the amount of disk space it requires.

- Selecting **Thin Provision** results in a faster clone operation and provides optimized usage of storage capacity. Disk space is allocated only as it is required. This is the default selection.
- Selecting **Preallocated** results in a slower clone operation and is optimized for the speed of guest read and write operations. All disk space requested in the template is allocated at the time of the clone operation.

6. Click **OK** to create the cloned virtual machine. The virtual machine displays in the Virtual Machines list.



Note

It may take some time for the virtual machine to be created. During this time, the status of the virtual machine displays as **Image Locked**, followed by **Down**.

6.3. Completing the Configuration of the Virtual Machine

Similar to a physical machine, a virtual machine requires disk storage and a network interface. Use the **New Virtual Machine - Guide Me** dialog box to define the network interfaces (NICs) and virtual disks of the new virtual machine. Click the **Guide Me** button on the main menu of the **Virtual Machines** tab to display the dialog box.

Define the NICs and Virtual Disks:

1. On the **New Virtual Machine - Guide Me** dialog box, To set up one or more network interfaces (or NICs) click the **Configure Network Interfaces** button. The **New Network Interface** dialog box displays. You can accept the default values, or change them if necessary.

Enter or select the **Name**, **Network** and **Type** of the network interface for the new virtual machine.



Note

The options on the **Network** and **Type** fields are populated by the networks available to the cluster, and the NICs available to the virtual machine.

The type of NIC that can be added to a virtual machine depends on the NIC drivers available for a guest operating system. For virtual machines running Linux operating systems, use e1000. For virtual machines running Windows operating systems, use rtl8139. The Red Hat VirtIO NIC is an enhanced, custom virtual NIC which requires a special driver and offers major speed improvements over other virtual NICs. This VirtIO driver is installed by default in Red Hat Enterprise Linux 6 guests, and must be installed manually in Windows guests. The VirtIO NIC is the optimal NIC for virtual machines in the Red Hat Enterprise Virtualization environment.

2. If required, select the **Specify Custom MAC address** checkbox, and enter the address of the NIC. Ensure that the MAC address is entered in lowercase.

Example 6.1. MAC Address Format

82:80:00:f5:9d:7c

3. Click **OK**. The dialog box closes, and the **New Virtual Machine - Guide Me** dialog box displays again, with changed context. If you have additional NICs, you can add additional network interfaces, by clicking the **Add another Network Interface** button.
4. To set up one or more virtual disks, on the **New Virtual Machine - Guide Me** dialog box, click the **Configure Virtual Disk** button.
5. The **New Virtual Disk** dialog box displays. You can accept the default values, or change them if necessary.

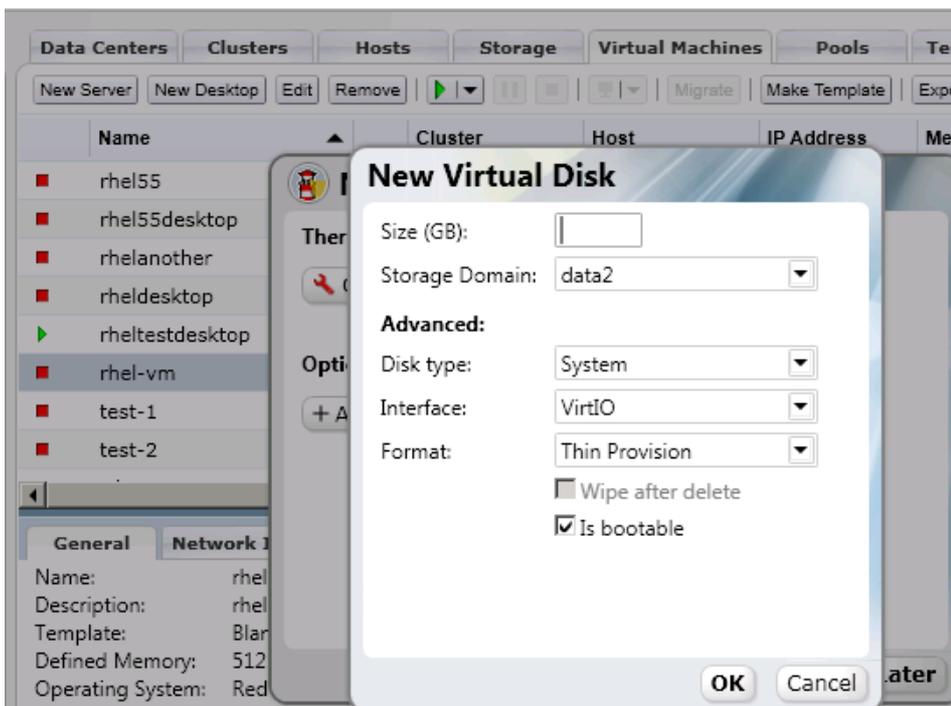


Figure 6.5. New Virtual Disk Dialog Box

- Enter the **Size** of the virtual disk in GB. Ensure that the size is appropriate to the applications that need to run on the virtual machine.
- Select the **Storage domain** where the virtual disk image is to be created.
- You can also define the **Advanced** properties of the virtual disk. These are:

Table 6.7. New Virtual Machine Dialog Box Fields

Field	Options	Notes
Disk Type	Select from <i>System</i> or <i>Data</i> options.	This refers to the disk onto which an operating system will be installed. For "Data", the default volume type is "Preallocated" and volume format is "RAW". For "System", the default volume type is "Preallocated" in case of adding disk to a virtual server and "Sparse" (thin-provisioning) in case of adding disk to a virtual desktop. Volume format for "System" disks is calculated according to the volume type and the storage domain type. Note that only one disk can be marked "Is Bootable" from the Manager.
Interface	Select the network drivers, either <i>IDE</i> or <i>PV</i> .	<i>IDE</i> is the default selection that uses an emulation of the

Field	Options	Notes
		IDE protocol. Windows 2008 virtual machines require a IDE drivers. Select <i>PV</i> to use the para-virtualized drivers.
Format	Select from <i>Preallocated</i> or <i>Thin Provision</i> .	Pre-allocated is the recommended selection for a server virtual machine, where a block of disk space is reserved for the virtual machine. Thin Provision allocates disk space as and when the virtual machine requires it. Thin Provision is the recommended selection for a desktop virtual machine.
Wipe after delete	Select if the disk is to be formatted after the virtual machine is deleted.	Selecting this option ensures that all data in the virtual machine is removed after the virtual machine is deleted.
Is bootable	Select if the disk is to be a bootable disk.	

- Click **OK**. The dialog box closes, and the **New Virtual Machine - Guide Me** dialog box displays again, with changed context. There should now be no further mandatory configuration to perform. Click **Configure Later** to close the dialog box.

Once the virtual machine is configured with virtual disk space and one or more network interfaces, the next step is to install operating systems and applications on it. The virtual machine displays in the list of virtual machines on the **Virtual Machine** tab, with a status of Down.



Note

You can also use the Details Pane on the **Virtual Machines** tab to add new virtual disks or network interfaces.

6.4. Installing Operating Systems onto Blank Virtual Machines

A virtual machine that is newly created from the "Blank" template requires an operating system and applications to be installed on it. Use the **Run Once** function to install an operating system and relevant applications onto the new virtual machine.

The **Run Once** function allows the administrator to run the virtual machine in a number of special modes, such as ACPI support, disable/enable acceleration, and others. Note that running a virtual machine in these special modes can cause performance degradation.

The **Virtual Machines** tab displays the existing virtual machines.

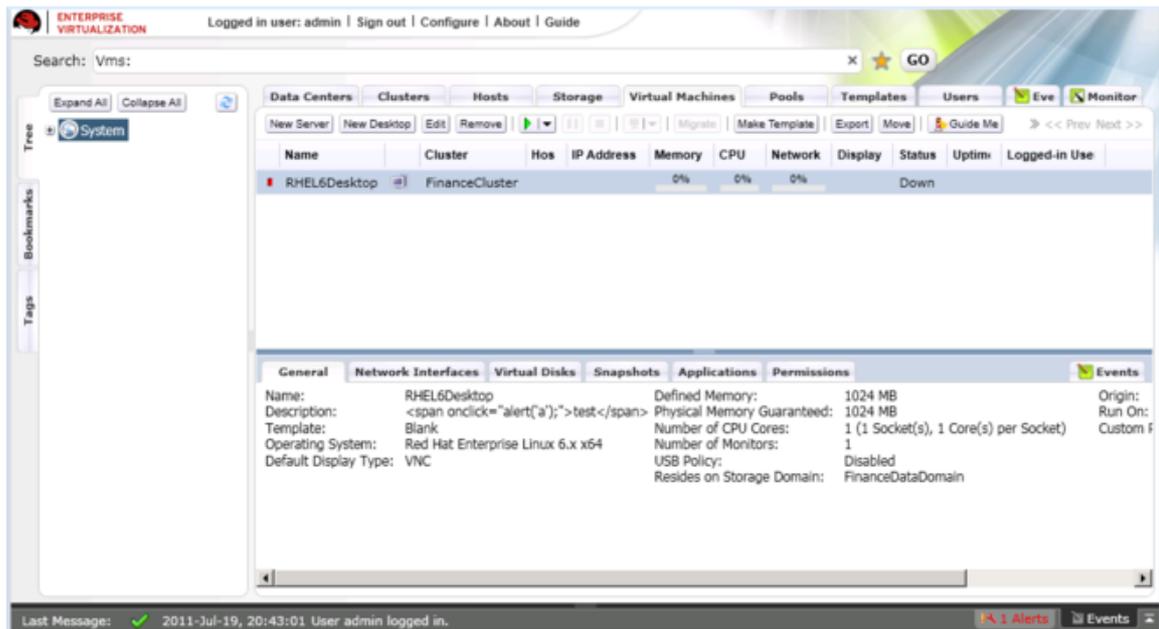


Figure 6.6. Virtual Machine List

To install an operating system onto a virtual machine:

1. Select the newly created virtual machine. It should have a status of *Down*.
2. Click the **Run Once** button on the Virtual Machines toolbar.



Figure 6.7. Run Virtual Machine Dialog Box

3. The **Run Virtual Machine** dialog box displays. The **Run Virtual Machine** dialog box consists of three sections — **Boot Options** to define how the virtual machine is to boot; **Boot Sequence** to determine which boot device is used first; and **Display Protocol** to select how the virtual machine is to connect to the system.

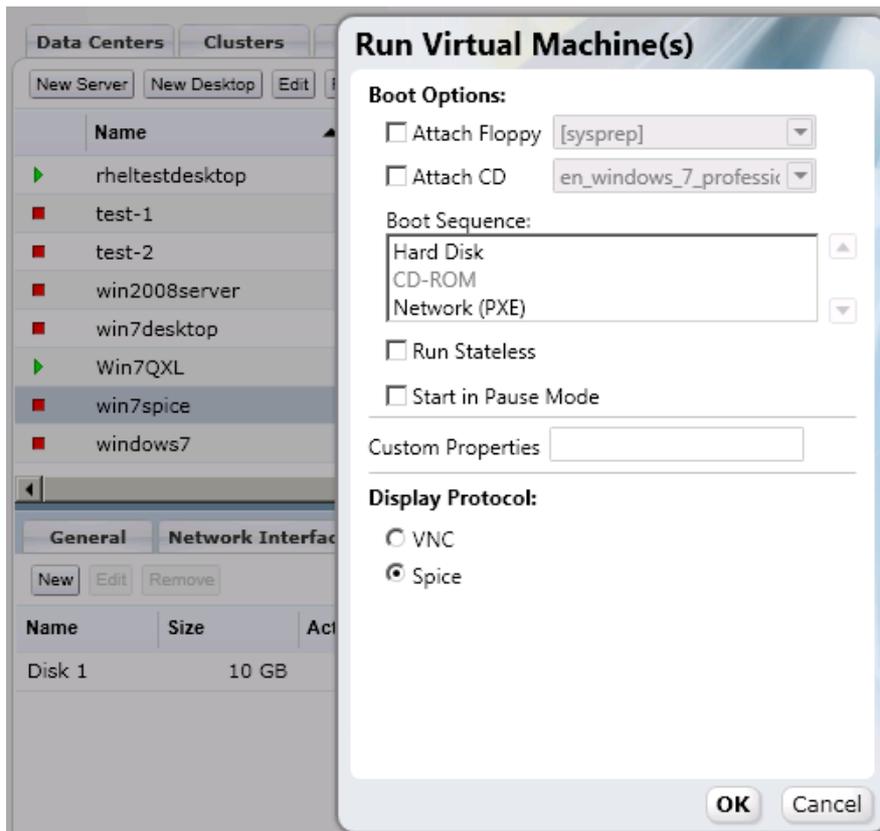


Figure 6.8. Run Virtual Machine Dialog Box

4. Define the **Boot Options**

- **Attach Floppy** – Use this option typically to install Windows drivers. It is mandatory to attach the floppy before attempting installation. The floppy must be attached, and the **Boot from CD** option selected to install drivers for the virtual machine.
- **Attach CD** – Select this option to install the operating system and applications from the CD onto the newly created virtual machine. In this case, select an ISO file from the drop-down list.
- **Boot Sequence** – After installing a new virtual machine, the new virtual machine must go into Boot mode before powering up. The Boot sequence can be altered from the previously selected one by moving the options up or down using the list buttons: **Hard Disk** to boot from the hard disk (though if this is a blank virtual machine, it will obviously not boot from the hard disk), **CD-ROM** to boot from the CD, or **Network (PXE)** to boot from the network.
- **Run Stateless** – Select this option if the virtual machine is to run in stateless mode. The stateless mode is mostly used for virtual desktops. A stateless desktop or server is always created from the base template, and deleted on shutdown. Every time the virtual machine is run, a new instance of the virtual machine is created from the base template. This type of virtual machine is very useful when creating virtual machines that need to be used for a short time, or by temporary staff.
- **Start in Pause Mode** – Select this option to run the virtual machine in *Pause* mode. In some instances, the virtual machine needs to be started and then paused to allow the administrator to connect to the display before the virtual machine goes into timeout. Connection to a virtual machine in a remote location may take longer than usual; consequently, the SPICE session may open after a timeout in an executed program has passed. To avoid such an occurrence, use the *Pause* mode. After the remote connection is made, resume the virtual machine session, either

by selecting the virtual machine in the Administration Portal and clicking the **Play** button; or by right-clicking the virtual machine's SPICE console window and selecting **Play**.

- **Linux Boot Option** – If the virtual machine is a Linux machine, you can enter customized option for the **kernel path**, **initrd path**, **kernel params** (parameters).
- **Custom Properties** – Enter any custom properties for the virtual machine to use while booting.
- **Reinitialize Sysprep** – For Windows virtual machines only. When a virtual machine runs for the first time, the system automatically attaches a virtual floppy drive containing the Sysprep configuration file to be used during Sysprep (relevant only if the virtual machine was sealed with Sysprep). The **Reinitialize Sysprep** option allows the administrator to restart the virtual machine with the attached floppy and configuration file. (For Windows virtual machines only). This option will not display for virtual machines that have never been initialized.

Define the **Display Protocol**

- Select SPICE for Windows or Linux virtual machines. This is the recommended protocol.
- Select VNC for Linux virtual machines if desired.

5. Click **OK**.

The virtual machine runs with the selected settings. The status changes to **Powering Up**, followed by **Up**.



Note

When you use the **Run Once** option, the defined parameters only apply to the current run, and do not hold for subsequent runs.

6.5. Managing Permissions to Virtual Machines

Virtual machines can be assigned to end users. A user can also be given system administrator rights to virtual machines. A virtual machine system administrator can manage a virtual machine in a similar manner to the system administrator of a physical system.

6.5.1. Assigning Virtual Machines to Users

After the virtual machines have been created and provisioned, they can be assigned to existing users. The UserRole allows end users to connect to virtual machines, perform basic operations, and use virtual machines that are part of a pool. A UserRole is an end user, not a virtual or physical system administrator role.

To assign a user to a virtual machine:

1. On the **Virtual Machines** tab, select the required virtual machine and click the **Permissions** tab from the Details pane. The **Permissions** tab displays a list of users and their current roles and permissions, if any. Note that there are no inherited permissions for virtual machines.

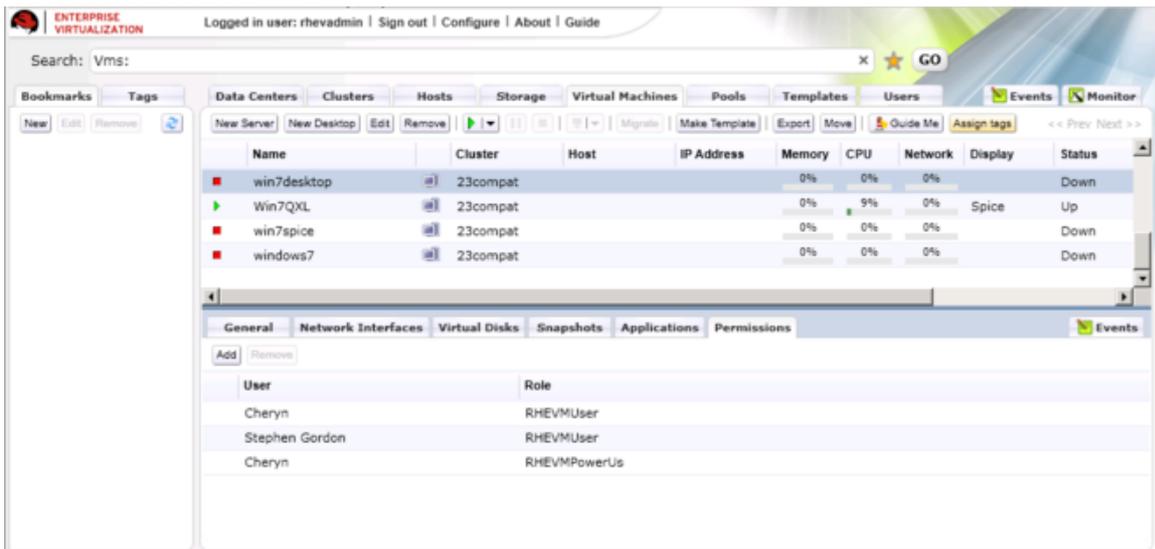


Figure 6.9. Virtual Machine Permissions

2. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
3. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select the required role, for example, **UserRole**.

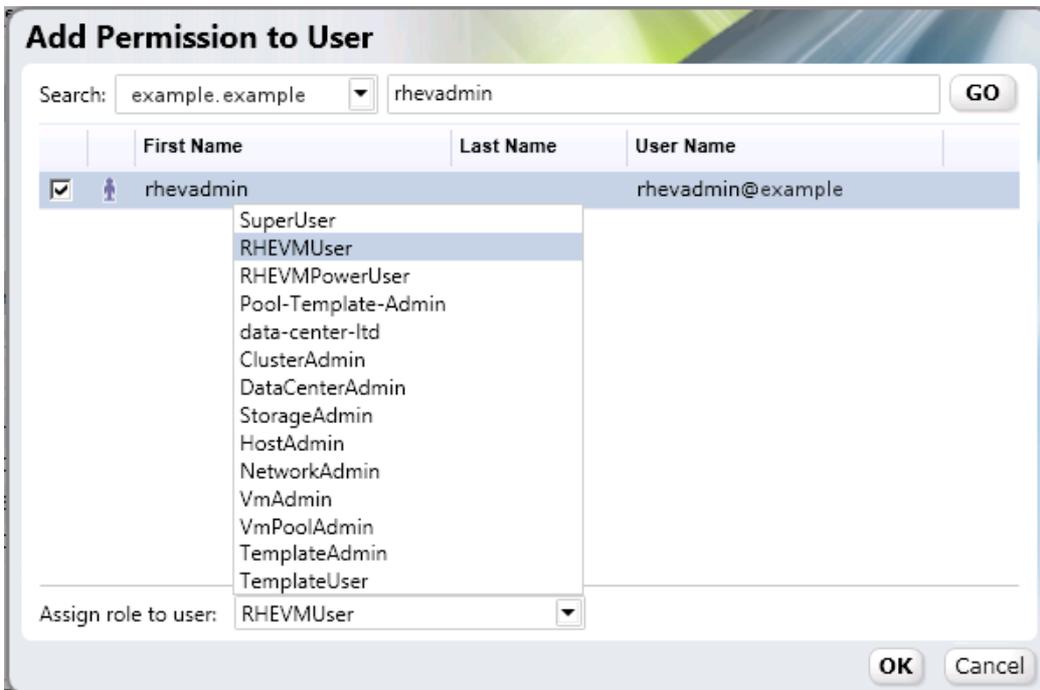


Figure 6.10. Assigning Virtual Machine User Permissions

4. Click **OK**.

The name of the user displays in the **Permissions** tab.

**Note**

You can only assign roles and permissions to existing users. See [Chapter 5, Users](#).

To remove a user role:

1. Click the **Virtual Machine** tab. A list of virtual machines displays. If the required virtual machine is not visible, perform a search (see [Section 1.2, "Search"](#)).
2. Select the required virtual machine and click the **Permissions** tab from the Details pane.
The **Permissions** tab displays a list of users and their current roles and permissions, if any.
3. Select the check box of the appropriate user.
4. Click **Remove**. A dialog prompts you to confirm the removal, click **OK** to proceed. The user is removed from the **Permissions** tab.

6.5.2. Assigning System Permissions to Virtual Machines

A **UserVmManager** role gives the user permissions to create, manage and delete the virtual machine, including its network and storage, as well as the permission to migrate the virtual machine.

To assign a system administrator role for a virtual machine:

1. Click the **Virtual Machines** tab.

A list of virtual machines displays. If the required virtual machine is not visible, perform a search (see [Section 1.2, "Search"](#)).

2. Select the virtual machine that you want to edit, and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any. Note that there are no inherited permissions for virtual machines.

3. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
4. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **UserVmManager**.
5. Click **OK**.

The name of the user displays in the **Permissions** tab.

**Note**

You can only assign roles and permissions to existing users. See [Chapter 5, Users](#).

To remove a virtual machine system administrator role:

1. Click the **Virtual Machine** tab. A list of virtual machines displays. If the required virtual machine is not visible, perform a search (see [Section 1.2, “Search”](#)).

2. Select the required virtual machine and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.

3. Select the check box of the appropriate user.
4. Click **Remove**. A dialog prompts you to confirm the removal, click **OK** to proceed. The user is removed from the **Permissions** tab.

6.6. Logging in to Virtual Machines

After creating a virtual machine from either a blank template or existing template, you can log in to the virtual machine with either the SPICE, VNC or RDP protocols to customize the virtual machine, install databases or applications, or make changes to the virtual machine.

6.6.1. Logging into Virtual Machines using SPICE

SPICE is the recommended connection protocol for Linux and Windows virtual machines. SPICE provides high quality graphics performance which is comparable to that of a physical machine, and bi-directional audio and video. It also supports multiple monitor display and USB redirection for virtual desktops. You can log in to virtual machine using SPICE if you have selected SPICE as the virtual machine's display protocol when you created it. See [Section 6.2.2, “Creating New Virtual Machines without a Template”](#) for details.

To log in to a virtual machine using SPICE

1. On the **Virtual Machines** tab select a running virtual machine.

The Details pane displays.

2. Click the **Console** button or click the **Console** option from the right-click menu.

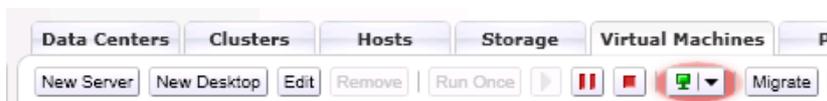


Figure 6.11. Connection Icon on the Virtual Machine Menu

3. The SPICE installation process starts if SPICE has not been previously installed. Follow the prompts to install SPICE, and proceed.
4. SPICE displays the login screen:

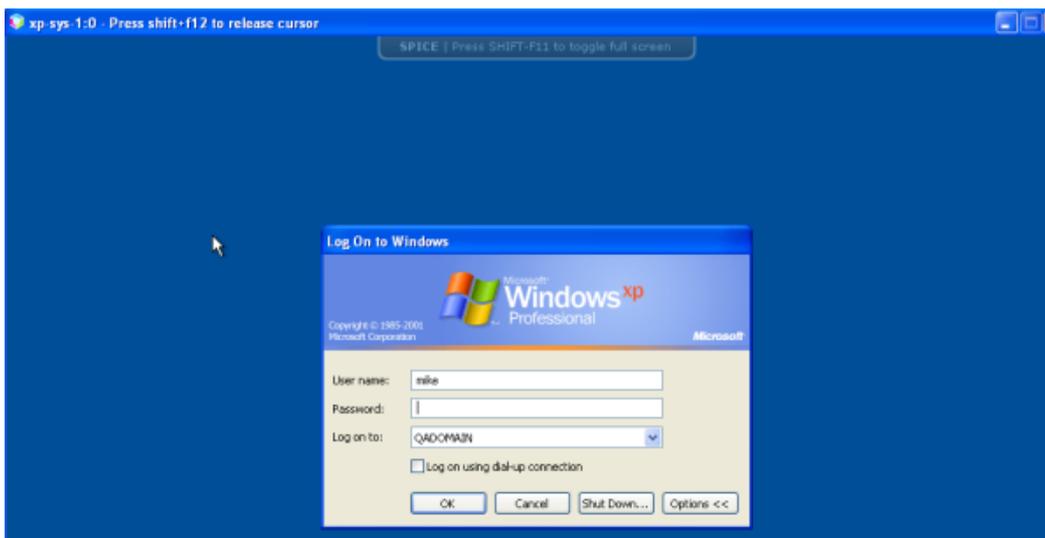


Figure 6.12. Virtual Machine Details Pane

Enter your **Username** and **Password**. Click **OK** to log onto the virtual machine.

5. When you have finished using the virtual machine, log out according to instructions specific to the operating system.

6.6.2. Logging in to Windows Virtual Machines with Remote Desktop (RDP)

You can use the Remote Desktop Protocol (RDP) to log in to Windows virtual machines, if remote access has been enabled on the virtual machine.

To log in to a virtual machine using RDP:

1. On the **Virtual Machines** tab select a running virtual machine.

The Details Pane displays.

2. Click the down arrow on the **Console** button and select the RDP option or click the **RDP** option from the right-click menu.

The RDP Windows login screen of the virtual machine displays.

3. Enter your username and password, and click **OK**. You are logged on to the virtual machine.
4. Install/uninstall applications and make the required changes to settings, if needed. If you wish, you can use the virtual machine to create a template. Refer [Chapter 7, Templates](#).
5. Shut down the virtual machine, or log out according to instructions specific to the Windows operating system.

6.6.3. Logging in to Virtual Machines with VNC

VNC is a supported display protocol for Linux and Windows virtual machines in the Administration Portal. The advantage of using VNC is that you do not have to install additional SPICE plugins. However, bear in mind that VNC is not supported in the User Portal. If your virtual machines are to be used in the User Portal, do not set VNC as the machine's default display protocol.

If you have already set your virtual machine's display protocol to VNC when you created it, you can click the **Console** button to log in to it. If your virtual machine's display protocol is set to SPICE, and you would like to change it to VNC, use the following procedure.

To log in to a virtual machine using VNC:

1. Select the virtual machine from the list on the **Virtual Machines** tab.
2. Click **Edit**. The **Edit Virtual Machine** dialog displays.

Select the **Console** tab. Set the **Protocol** field to **VNC**. Click **OK** to save your changes.

3. Back on the **Virtual Machines** tab, select the virtual machine you wish to log in to. Ensure that it is running, and click the **Console** button.

The VNC login screen of the virtual machine displays.

4. Enter your username and password, and click **OK**. You are logged on to the virtual machine.
5. Install/uninstall applications and make the required changes to settings, if needed. If you wish, you can use the virtual machine to create a template. Refer [Chapter 7, Templates](#).
6. Shut down the virtual machine, or log out from the virtual machine.

6.6.4. Console Window Menu Extension for Administrators

There are various functions available to the administrator via the SPICE menu.



Note

The functions available to the administrator differ from the ones available to the user, while the user is connected to the console.

To view the functions available to the administrator:

1. At the top left corner of the Console window, click the SPICE icon.
2. The SPICE menu displays.

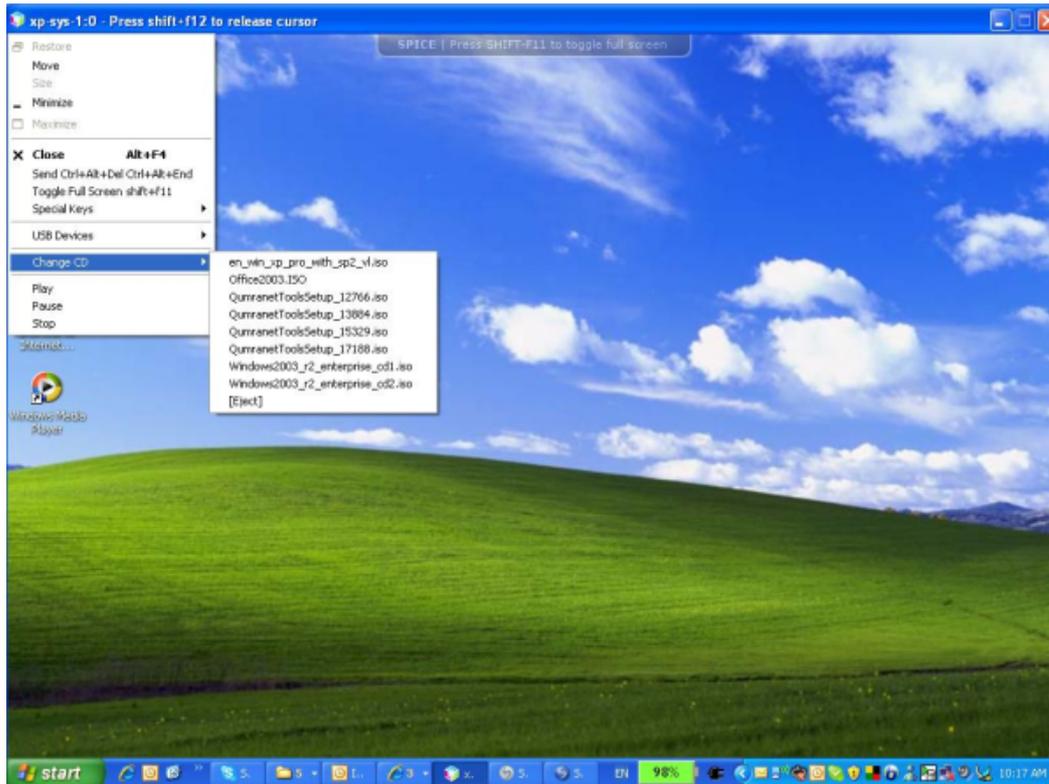


Figure 6.13. Console Window Menu for Administrators

The following features are available on the SPICE menu:

1. **Send CTRL+ALT+DEL** (or enter **Ctrl+Alt+End**): to simulate this key sequence as if entered on the virtual machine.
2. **Toggle full screen** (or enter **Shift+F11**): to switch between full-screen and window mode for the virtual machine.
3. **Special Keys** : to input special characters (selecting from the list to send a key sequence to the virtual machine).
4. **USB Devices** : allows attaching and detaching USB devices currently connected to your client.
5. **Change CD** : for the list of imported ISO image files found in the `/images` folder.
6. **Play, Pause, Stop** : to perform these basic virtual machine management operations from the Console Window menu.

6.7. Managing Virtual Machines

Some maintenance tasks are performed directly on the virtual machine (such as running, pausing, or stopping), and some maintenance tasks involve other objects (such as migrating a virtual machine to a different physical host in the same cluster).

Maintenance tasks include basic operations on a virtual machine such as editing properties, powering up or down, taking snapshots, and exporting or importing.

6.7.1. Editing Virtual Machines

You can edit the details of a virtual machine, such as its name or memory size. You cannot change the host cluster, template or Storage Domain to which the virtual machine belongs. Changes take effect after the virtual machines are shut down and restarted.



Warning

Be aware that changes to storage, operating system or networking parameters can adversely affect the virtual machine. Ensure that you have the correct details before attempting to make any changes. It is recommended that you take the precaution of backing up the virtual machine before you make changes.

To edit virtual machine details:

1. Click the **Virtual Machines** tab.
2. If the virtual machine you want to edit is not visible in the list, perform a search (see [Section 1.2, “Search”](#)).
3. Select the virtual machine and click the **Edit** button.

The **Edit Virtual Machine** dialog box displays. Disabled fields cannot be changed.

4. Edit the required details of enabled fields. Refer [Section 6.2.2, “Creating New Virtual Machines without a Template”](#) for details of the fields.



Note

Some fields cannot be changed, and are disabled by default.

5. Click **OK**.

The details of the virtual machine are updated in the **Virtual Machines** tab.

6.7.2. Powering Virtual Machines On

When you power on a virtual machine, the Red Hat Enterprise Virtualization platform automatically selects the best available host on which to run the virtual machine.

To power on a virtual machine:

1. Click the **Virtual Machines** tab.
2. If the virtual machine that you want to edit is not visible in the list, perform a search (see [Section 1.2, “Search”](#)).
3. Select the virtual machine with a status of *Down*.
4. Click  or right-click and select **Run**.

The **Status** of the virtual machine changes to *Up*. The display protocol and IP of the selected host display.

6.7.3. Shutting Down or Pausing Virtual Machines

It is recommended that a virtual machine be shut down from within its console. However, occasionally there is a need to shut down the virtual machine from the Administration Portal. The Red Hat Enterprise Virtualization platform provides for an orderly shutdown if the guest tools are installed.

It is best practice that all users are logged off from a Windows virtual machine before shutting down. If any users are still logged in, the following Windows message displays on the virtual machine, *Other people are logged on to this computer. Shutting down Windows might cause them to lose data. Do you want to continue shutting down?*, and the virtual machine remains with a *Powering Off* status in the Administration Portal.

If a virtual machine cannot be properly shut down, since, for example, the operating system is not responsive, you might need to force a shutdown, which is equivalent to pulling out the power cord of a physical machine.



Warning

Exercise extreme caution when forcing shutdown of a virtual machine, as data loss may occur. Shutdown of virtual machines should be planned after due consideration, preferably at times that will least impact users.

To pause a virtual machine:

1. Click the **Virtual Machines** tab.
2. If the virtual machine that you want to edit is not visible in the list, perform a search (see [Section 1.2, “Search”](#)).
3. Select the virtual machine.
4. Click the **Pause** () button.

The **Status** of the virtual machine changes to *Paused*.

Pausing a virtual machine puts it into *Hibernate* mode, where the virtual machine state is preserved. Applications continue running, but CPU usage is zero.

To shut down a virtual machine:

1. Click the **Virtual Machines** tab.
2. If the virtual machine that you want to edit is not visible in the list, perform a search (see [Section 1.2, “Search”](#)).
3. Select the virtual machine.
4. Click the **Power Off** () button.

The **Status** of the virtual machine changes to *Down*.

6.7.4. Migrating Virtual Machines

A running virtual machine can be migrated to any host within its designated host cluster. This is especially useful if the load on a particular host is too great, and is essential before bringing a server down for maintenance (migration is automatic in this case). Migration of virtual machines does not cause any service interruption.



Note

Virtual Machines migrate within their designated host cluster. The system determines the host to which the virtual machine is migrated, according to the load balancing and power rules set up in the Policy Engine. These rules are explained in the advanced sections later in this guide — [Section 13.4.2, “Setting a Cluster Resilience Policy”](#) and [Section 12.3.2, “Cluster Policy”](#).

To migrate a virtual machine to another host:

1. Click the **Virtual Machines** tab.
2. If the virtual machine you want to migrate is not visible in the list, perform a search (see [Section 1.2, “Search”](#)).
3. Select the virtual machine.
4. Click the **Migrate** button. The **Migrate Virtual Machine** dialog box displays.
5. Choose from **Select Host Automatically** or specify a destination from the **Select Destination Host** list. If you chose **Select Destination Host**, only active hosts within the cluster display in the list.
6. Click **Close** to close the dialog box.

The virtual machine is migrated to another host in the cluster. Shortly after, the **Host** column displays the new host to which the virtual machine has migrated.

6.7.5. Moving Virtual Machines within a Data Center

A virtual machine can be moved to a different Storage Domain within the data center. The data center requires an additional active data domain in the data center.

To move a virtual machine to another storage domain:

1. Click the **Virtual Machines** tab.
2. If the virtual machine you want to migrate is not visible in the list, perform a search (see [Section 1.2, “Search”](#)).
3. Select the virtual machine.
4. Shut down the virtual machine.
5. Click the **Move** button.
6. The **Move Virtual Machine** dialog box displays. Select from the list of available **Storage Domains**. If there are no additional Storage domains, an error message displays.
7. Click **Close**.

The virtual machine is moved to the different storage domain.

6.7.6. Removing Virtual Machines

Virtual Machines no longer in use can be removed.



Warning

Removing a virtual machine is final and cannot be reversed.

To remove a virtual machine:

1. Click the **Virtual Machine** tab.
2. If the virtual machine you want to remove is not visible in the list, perform a search (see [Section 1.2, "Search"](#)).
3. Select the virtual machine.
4. Shut down the virtual machine. The **Remove** button is only enabled for a virtual machine that has a status of Down.
5. Click the **Remove** button.
A confirmation message is displays. Click **OK**.
6. The virtual machine is removed from the platform and no longer displays on the **Virtual Machines** tab.

6.8. Using Virtual Machine Snapshots

A snapshot is a view of a virtual machine's operating system and all its applications at a given point in time. The snapshot is a very important tool in managing virtual machines. Whenever the virtual machine is powered off, you can create a snapshot of a virtual machine's hard drive. If future changes cause a problem, you can restore the virtual machine to the previous state of any of the snapshots. Restoration to a snapshot means that you return to the point in time when the snapshot was created. After you restore to that point of time, you cannot return to snapshots created after that time.

For example, given that snapshots were created on Sunday at 8 am, 10 am, 12 pm, and at 3 pm. At 6 pm, a problem arises on your virtual machine, and you decide to restore the virtual machine to the state of the snapshot created at 10 am. This restore automatically erases the snapshots created after 10 am, meaning the snapshots of 12 pm and 3 pm no longer exist. However, the snapshots taken before the restoration, in this case at 8 am, still exist.



Warning

When a restoration is performed from a snapshot, all data written to the virtual machine's hard drive after the selected snapshot creation point is lost, including subsequent snapshots.

6.8.1. Creating Snapshots of Virtual Machines

This section describes how to create a snapshot of a virtual machine. You can also **Preview**, **Commit**, **Undo** and **Delete** the snapshot.

To create a snapshot of a virtual machine:

1. Click the **Virtual Machines** tab.
2. If the virtual machine for which you want to create a **Snapshot** is not displayed, perform a search (see [Section 1.2, "Search"](#)).
3. Select the virtual machine. Ensure that the virtual machine has a status of *Down*.
4. On the Details pane, select the **Snapshots** tab. Click the **Create** button.

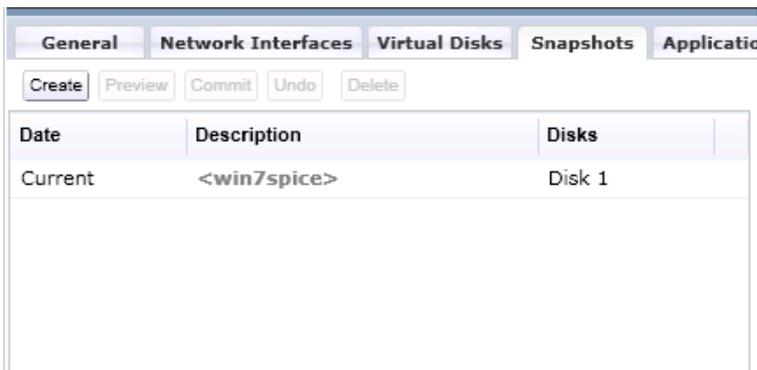


Figure 6.14. The Virtual Machines Details Pane with Snapshots tab

The **Create Snapshot** dialog box displays.

5. Enter a description for the snapshot, select all the virtual disks attached to the virtual machine and click **OK**.

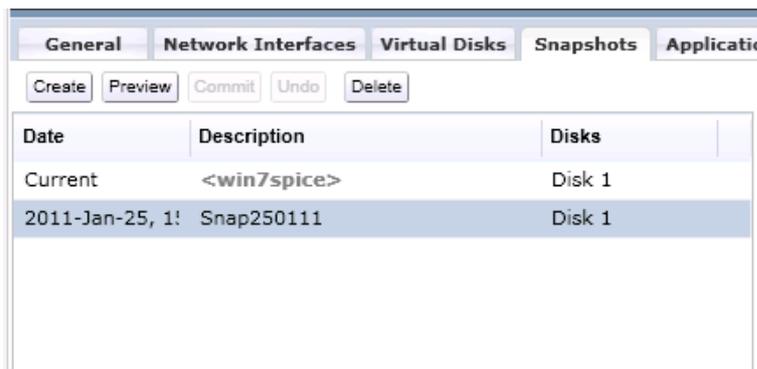


Figure 6.15. New Snapshot Dialog Box

The virtual machine's operating system and applications are stored in a snapshot that can be previewed or restored. The status of the virtual machine briefly changes to Image Locked, before returning to Down.

6.8.2. Restoring Virtual Machines from Snapshots

This section describes how to restore a virtual machine from a snapshot.

To use a snapshot to restore a virtual machine:

1. Click the **Virtual Machines** tab.

2. If the virtual machine you want to restore is not visible in the list, perform a search (see [Section 1.2, "Search"](#)). Ensure that the virtual machine is powered down and has a status of Down.
3. Click the virtual machine. On the Details Pane, click the **Snapshots** tab. A list of snapshots displays.

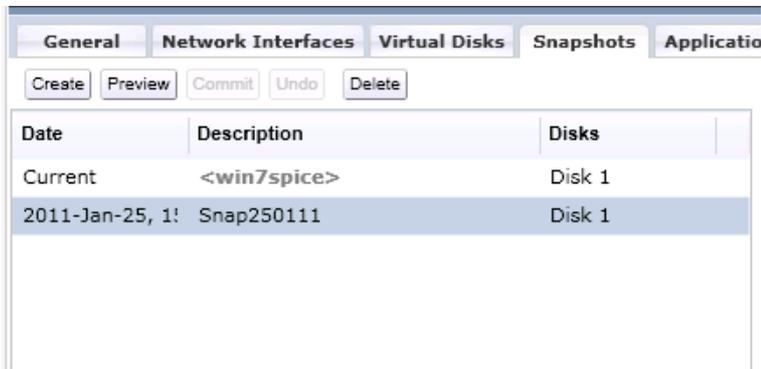


Figure 6.16. Snapshot List

4. Select the snapshot that you want to restore.
The Snapshot Details display, and the **Preview** button is enabled.
5. Click **Preview** to preview the snapshot. The Status of the virtual machine briefly changes to Image Locked, before returning to Down.
6. At this point, you can start the virtual machine and it will run with the disk image of the time the snapshot was taken. After you have checked the snapshot do one of the following:
 - a. To restore to this point:

Click **Commit**.

The virtual machine is restored to the state it was in at the time of the snapshot. Also, any subsequent snapshots are erased.

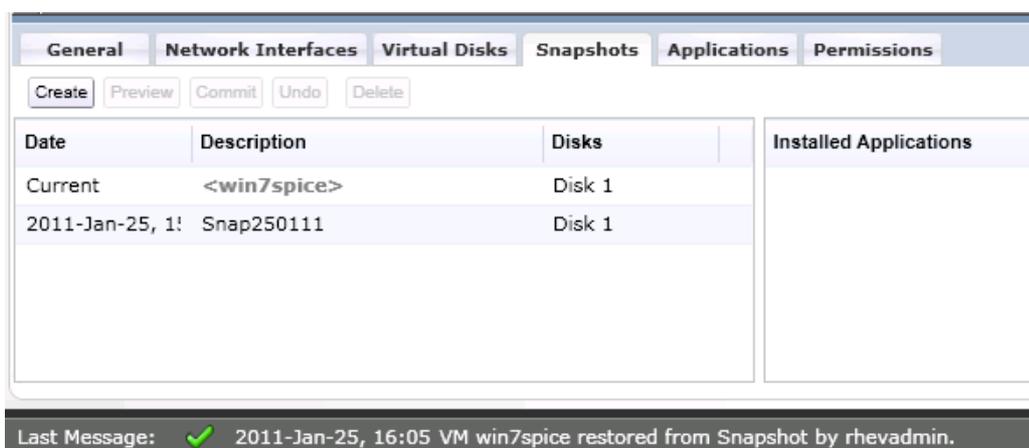


Figure 6.17. Snapshot List

- b. Alternatively, click the **Undo** button.

The virtual machine is not restored to the state of the snapshot, and returns to its current state.

6.8.3. Deleting Snapshots

This section describes how to delete a snapshot. Snapshots occupy virtual disk space, and depending on the installed applications can significantly reduce available disk space.

To delete a snapshot:

1. Click the required virtual machine. On the Details Pane, click the **Snapshots** tab. A list of snapshots displays.

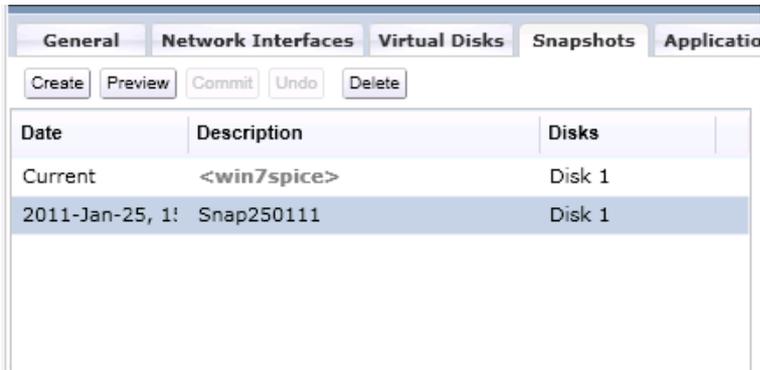


Figure 6.18. Snapshot List

2. Click **Preview** to preview the snapshot. The Status of the virtual machine briefly changes to Image Locked, before returning to Down.
3. At this point, you can start the virtual machine and it will run with the disk image of the time the snapshot was taken. After you have checked the snapshot, and are sure that you wish to delete it, click the **Delete** button.

The snapshot is deleted.

6.9. Exporting and Importing Virtual Resources

A virtual machine or a template can be imported or exported to a data center in a different Red Hat Enterprise Virtualization system. Red Hat Enterprise Virtualization Manager allows you to import and export virtual machines (and templates) stored in Open Virtual Machine Format (OVF). This feature can be used in multiple ways:

- To move virtual resources to a different installation of Red Hat Enterprise Virtualization.
- To move virtual resources to a different data center in the same installation of Red Hat Enterprise Virtualization. To do this, the original virtual resource must be deleted.
- To back up virtual resources.

There are two methods of exporting and importing virtual resources:

- Exporting or importing one or more virtual machines or templates.
- Exporting and importing a domain of virtual machines and templates. Refer to [Import Existing ISO or Export Storage Domain](#).

A virtual machine must be stopped before it can be moved across data centers. If the virtual machine was created using a template, the template is not automatically exported, however the template must exist in the destination domain for the virtual machine to work.

6.9.1. Overview of the Export-Import Process

To export or import resources, an active export domain must be attached to the data center. The export domain can be thought of as a temporary storage area that contains one directory per virtual machine. Each directory consists of all the OVF (Open Virtualization Format) files pertaining to the virtual machine. The export domain enables you to add pre-configured virtual machines or domains to a Red Hat Enterprise Virtualization Manager system. You can also import virtual machines from a different format, for example, Xen, VMware or Windows virtual machines, using the V2V feature provided. V2V converts virtual machines and places them in the export domain. For more information on V2V, see the *Red Hat Enterprise Linux V2V Guide*.



Note

Only one export domain can be active in the data center. This means that the domain can be attached to either the source data center or the destination data center.

To perform an export-import of virtual resources:

1. Attach the export domain to the source data center. See [Section 3.1.2.2.3, “Attaching an Export Storage Domain”](#)

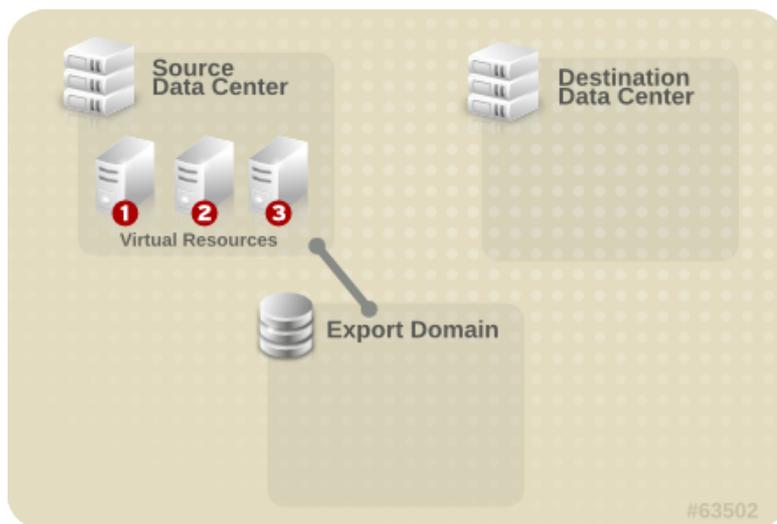


Figure 6.19. Attach Export Domain

2. Export the virtual resource to the export domain.

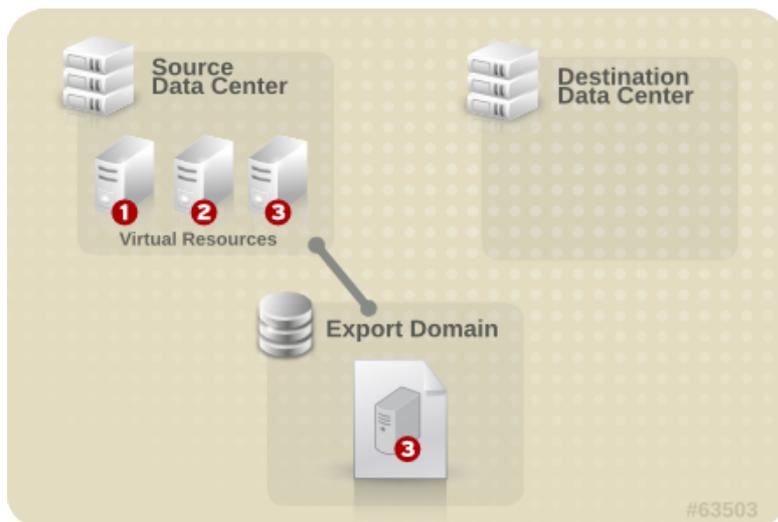


Figure 6.20. Export the Virtual Resource

3. Detach the export domain from the source data center. See [Section 3.1.5.2.1, “Detaching Storage Domains from a Data Center”](#)

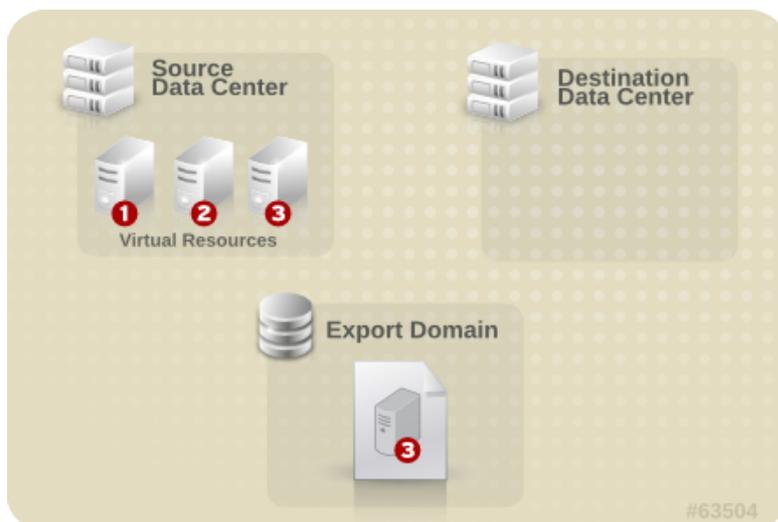


Figure 6.21. Detach Export Domain

4. Attach the export domain to the destination Data center. See [Section 3.1.2.2.3, “Attaching an Export Storage Domain”](#)

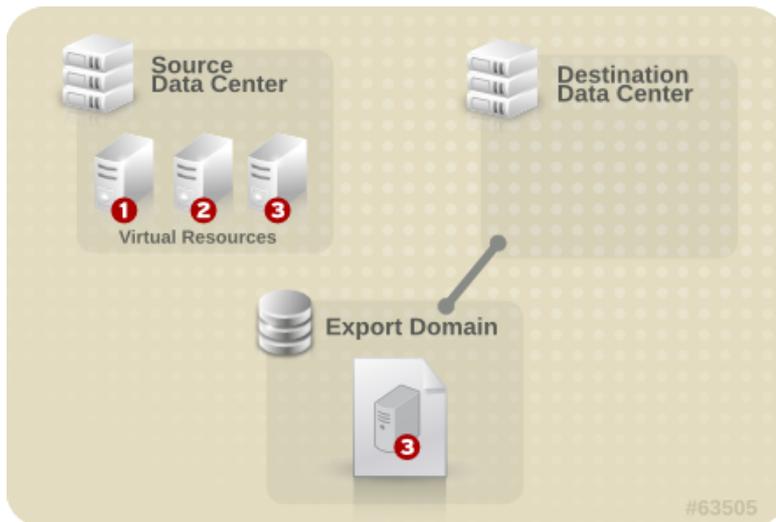


Figure 6.22. Attach the Export Domain

5. Import the virtual resource into the destination data center.

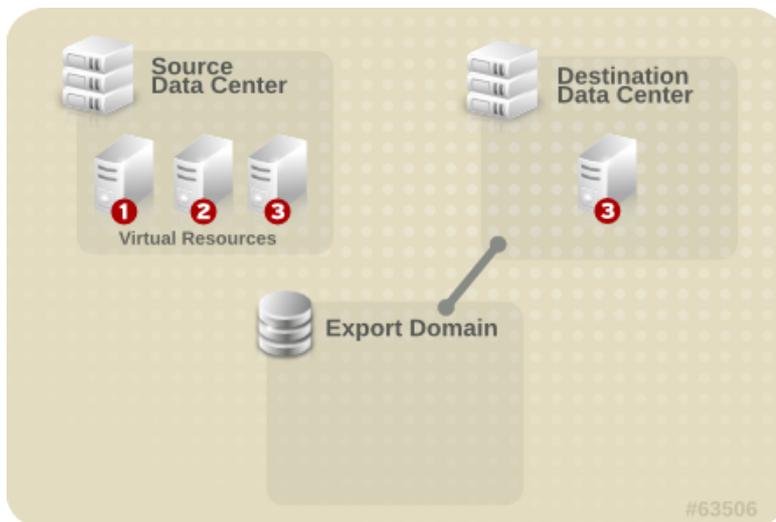


Figure 6.23. Import the virtual resource

6.9.2. Exporting Virtual Machines

Exporting virtual resources across data centers requires some preparation, for example, an export domain should exist, and be attached to the appropriate data center; the virtual machine must be shut down, and the template requirements need to be considered as well. You will also need to consider whether you want to export the virtual machine to the new data center and retain the original virtual machine, or move it to the new data center, and remove it from the source data center. You will also need to attach or detach the export domain as appropriate.

To export individual virtual machines to the export domain:

1. Click the **Virtual Machines** tab.
2. If the virtual machine you want to export is not visible in the list, perform a search (see [Section 1.2, "Search"](#)).
3. Select the virtual machine.

4. Shut down the virtual machine. Once the virtual machine is shut down, right click to display the menu.
5. Click the **Export** option.
6. The **Export Virtual Machine** dialog box displays. Select from the list of available options as appropriate, **Force Override** and **Collapse Snapshots**.

Select **Force Override** to override existing images of the virtual machine which may already exist on the export domain.

Select **Collapse Snapshots** to create a single export file per disk. Select this option if you wish to retain both the source and destination versions of the virtual machine.

7. Click **OK**.

The export of the virtual machine begins, this can take some time. The virtual machine displays in the Virtual Machines list with a Locked Status. Use the **Events** tab to view the progress.

8. The **Events** tab displays that the virtual machine has been exported.
9. The virtual machine displays on the **VM Import** tab of the export domain.

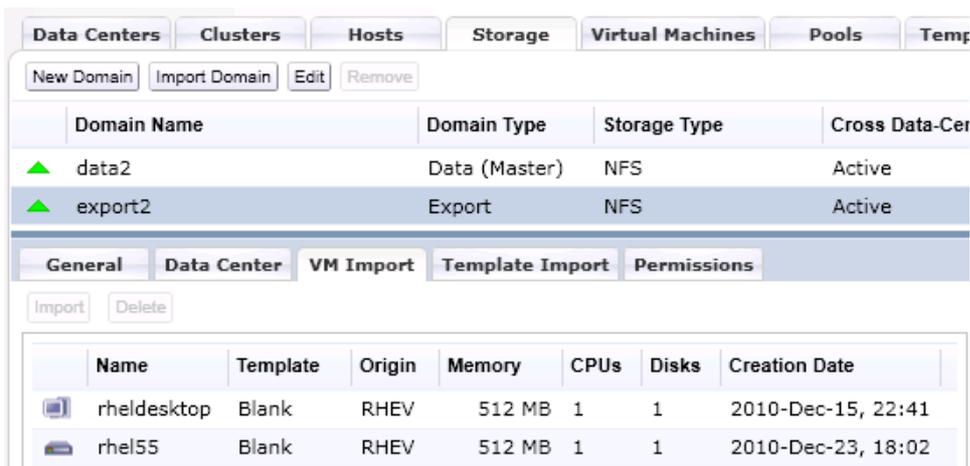


Figure 6.24. Export Virtual Machine

10. You can repeat the procedure above to export each virtual machine that you need to migrate, so that the export domain has a number of virtual machines.

6.9.3. Importing Virtual Machines into the Destination Data Center

Once the virtual machine, or machines are available in the export domain, you can import them into the destination data center. If the destination data center is within the same installation of Red Hat Enterprise Virtualization, delete the originals from the source data center after exporting them to the export domain.

To Import the Virtual Machine into the Destination Data Center

1. Detach the export domain from the source data center (see [Section 3.1.5.2.1, “Detaching Storage Domains from a Data Center”](#)), then attach it to the destination data center (see [Section 3.1.2.2.3, “Attaching an Export Storage Domain”](#)).
2. On the **Storage** tab, select the Export data domain. The Details pane of the Export storage domain displays.

- On the Details pane, select the **VM Import** tab. Select the virtual machine that is to be imported. Click **Import**.

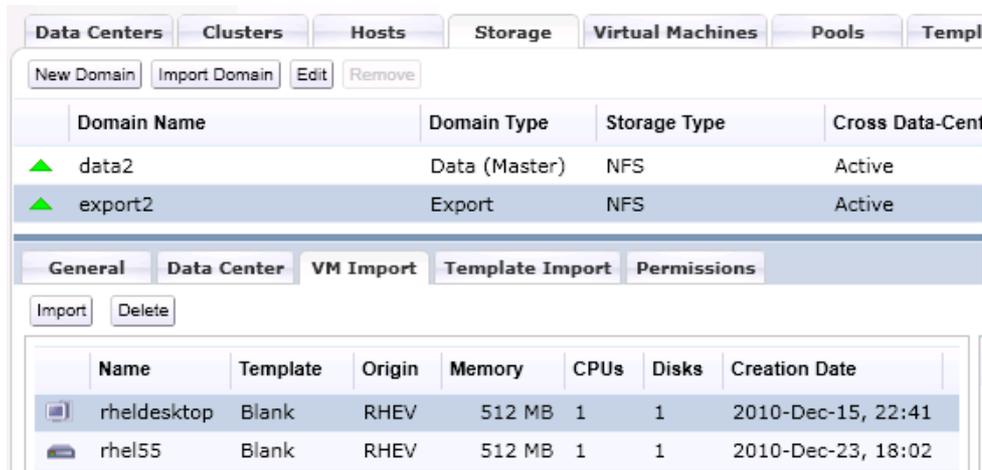


Figure 6.25. Import Virtual Machine

- The **Import Virtual Machine** dialog box displays. The names of the available virtual machines display.

Select the name of the virtual machine, and select the **Destination Cluster** and **Destination Storage** of the destination data center.

If you have not deleted the original virtual machine in the source data center, select **Collapse Snapshots**.

- Click **OK**. The virtual machine is imported into the destination data center. This can take some time. Eventually, the virtual machine displays in the **Virtual Machines** tab on the Details pane of the Storage domain belonging to the destination data center.
- You can now run the virtual machine, provided that the virtual machine was created from scratch, or from a template which still resides in the storage domain.

6.10. Backing Up Virtual Resources

Virtual Machines and Templates may need to be backed up from time to time, for example, before undertaking maintenance of hosts or storage servers. To back up virtual machines and templates, use the export domain and procedures as described in [Section 6.9, "Exporting and Importing Virtual Resources"](#).

Templates

Templates are model virtual machines that are used as a convenient and efficient way to create new virtual machines of the same type and content. Templates provide a shortcut that reduces the time required to build virtual machines. A template can contain an operating system only, or can contain all applications required by a particular department. You can also control the use of templates by defining permissions to use and administer templates.

7.1. Creating Templates from Existing Virtual Machines

A template can be created from an existing virtual machine that has been configured to meet the needs of several individuals in the organization, and has been sealed with Sysprep (Windows machines only) or a similar tool.

Before selecting an existing virtual machine as the source for a template, ensure that the virtual machine is general enough for this purpose. A virtual machine that is too specific to a particular user or group may require a lot of changes, and is therefore not practical to use as a template.

To create a template from an existing virtual machine:

1. Click the **Virtual Machines** tab.

The **Virtual Machines** tab displays a list of all virtual machines in the system.

2. Select the virtual machine that you want to use as a basis for the template definition. Ensure that the virtual machine is powered down and has a status of **Down**.



Note

Take a snapshot of the Virtual Machine at this stage if you wish to use the virtual machine (as a virtual machine) after it is used to create a template.

3. Click **Make Template**.

The **New Virtual Machine Template** displays, with the details of the selected Virtual Machine.

4. Enter, accept or change the following information. **Name** and **Description** are typically the only fields in which new information is to be entered. The rest of the fields are taken directly from the existing virtual machine.

Define the **Host Cluster** and **Storage Domain** for the virtual machines that will be created with this template. By default, these are the same as the source virtual machine, but you can select from the list, if there are multiple clusters and domains in the data center. The template is a copy of the virtual machine, hence the requirement for the **Host Cluster** and **Storage Domain**.

5. Click **OK**. The virtual machine displays a status of **Image Locked** while the template is being created. The template is created and added to the **Templates** tab. Depending on the virtual machine disk image size, this may take a long time. During this time, the action buttons for the template remain disabled. Once created, the action buttons are enabled and the template is ready for use. For example, the newly created template displays in the list of templates in the Template field on the **New Virtual Machine** dialog box.



Note

Before a Windows template is ready for application, you must first run **Sysprep** (or a similar tool) to generalize the Virtual Machine and remove "specific" personalization. Failure to do so will cause conflicts when multiple virtual machines from an un-generalized Windows templates are run. In general, templates of Linux virtual machines do not require sealing.

7.1.1. Sealing a Windows Template with Sysprep

Templates that have been created for Windows virtual machines must be generalized (sealed) before use, by means of a tool such as Sysprep. This section describes how to use **Sysprep** to seal a template before use. This ensures that machine-specific settings are not propagated through the template.



Important

Do not reboot the virtual machine during this process.

Before beginning the Sysprep process to prepare a virtual machine to be used as a template, ensure that the following settings have been configured:

- The Windows Sysprep parameters have been correctly defined. If not, click **Edit VM** and enter the required information in the **Operating System** and **Domain** fields.
- The correct Product Key has been entered in the Configuration Tool, **rhevms-config**. If not, as the root user on the Manager, run the Configuration Tool and enter the required information. The configuration keys that you need to set are **ProductKey** and **SysPrepPath**. For example, for Windows 7, the configuration value is **ProductKeyWindow7**, and **SysPrepWindows7Path**. The command to set the value is:

```
# rhevms-config --set ProductKeyWindow7=<validproductkey> --cver=general
```

For the exact syntax for getting and setting configuration values, see *Red Hat Enterprise Virtualization 3.0 Installation Guide*

Procedure 7.1. To seal a Windows XP template

1. Download **sysprep** to the virtual machine to be used as a template.

The Windows XP Sysprep tool is available at: <http://www.microsoft.com/download/en/details.aspx?id=11282>

2. Create a new folder **c:\sysprep**.
3. Open the **deploy.cab** file and put its contents in **c:\sysprep**.
4. Execute **sysprep.exe** from within the folder. Click **OK** on the welcome message.
5. The Sysprep tool displays.

Select the following check boxes:

- *Don't reset grace period for activation*
- *Use Mini-Setup*

Ensure that the Shutdown mode is set to **Shut down** before clicking **Reseal**

6. Acknowledge the pop-up window. The virtual desktop will go through the sealing process and then shut down automatically.

You can also use the procedure above to seal a Windows 2003 template. The Windows 2003 Sysprep tool is available at <http://www.microsoft.com/download/en/details.aspx?id=14830>.

Procedure 7.2. To seal a Windows 7 or Windows 2008 template

1. In the virtual machine to be used as a template, open a command line terminal and type **regedit**.
2. The **Registry Editor** window displays. On the left pane, expand **HKEY_LOCAL_MACHINE** → **SYSTEM** → **SETUP**.
3. On the main window, right click to add a new string value using **New** → **String Value**.
4. Right click on the file and select **Modify**. When the **Edit String** dialog box displays, enter the following information in the provided text boxes:
 - Value name: **UnattendFile**
 - Value data: **a:\sysprep.inf**
5. Launch Sysprep from **C:\Windows\System32\sysprep\sysprep.exe**
 - Under **System Cleanup Action**, select **Enter System Out-of-Box-Experience (OOBE)**.
 - Tick the **Generalize** check box if you need to change the computer's system identification number (SID).
 - Under **Shutdown Options**, select **Shutdown**.

Click **OK**.
6. The virtual machine will now go through the sealing process and shut down automatically.

The Windows virtual machine has now been sealed, and can be used as a template for Windows virtual machines.

7.1.2. Sealing a Linux Template with sys-unconfig

Templates that have been created for Linux virtual machines must be generalized (sealed) before use. This section describes how to use **sys-unconfig** to seal a template before use. This ensures that machine-specific settings are not propagated through the template.

1. Login to the virtual machine to be used as a template and flag the system for re-configuration by running the following command as root:

```
# touch /.unconfigured
```

2. Remove ssh host keys. Run:

```
# rm -rf /etc/ssh/ssh_host_*
```

3. Shut down the virtual machine. Run:

```
# poweroff
```

The Linux virtual machine has now been sealed, and can be used as a template for Linux virtual machines.

7.2. Managing Templates

Template details can be edited, and a template can be deleted if no virtual machines were built from it. Templates can also be exported and imported across data centers.

7.2.1. Editing Templates

You can edit the details of a template, such as its name, description, or memory size. You can also change the host cluster to which the resulting virtual machine belongs. Use the **Templates** tab to access and edit existing templates. The **Templates** tab and **Details** pane provides information on existing templates including **Creation Date**, **Cluster** and the number of virtual machines that are using a template.

To edit template details:

1. Click the **Templates** tab. The **Templates** tab displays.
2. If the template you want to edit is not visible on the list, perform a search (see [Section 1.2, “Search”](#)), and select the template.
3. Click the **Edit** button.
4. The **Edit Template** dialog box opens. This dialog box is essentially the same as the **Create New Virtual Machine** dialog box.
5. Edit details, as required. Refer [Section 6.2.2, “Creating New Virtual Machines without a Template”](#).
6. Click **OK**.

The details of the template are updated in the **Templates** tab.



Note

For a complete description of the fields, see [Table 6.4, “New Virtual Machine Dialog Box Fields”](#).

7.2.2. Copying Templates to a Different Storage Domain

You can copy a template to a new Storage domain. This will result in the identical template being available in a different Storage domain.



Important

You cannot have two copies of the same template in the same Storage domain.

To copy a template:

1. Navigate to the required template and select the template.
2. Click the **Copy** button. The **Copy Template** dialog box displays.
3. Select the **Storage Domain** that you wish to copy the template to.
4. Click **OK**.

The details of the template are copied to the new Storage Domain. The copy of the template displays the new storage domain in the **Domain** column of the Templates tab.

7.2.3. Deleting Templates

Disk space can be conserved by deleting unused templates. Templates that have been used to build virtual machines cannot be deleted unless all virtual machines created from the particular template are first removed.

To delete a template:

1. Click the **Templates** tab.
2. Select the template. If the template you want to delete is not visible on the list, perform a search (see [Section 1.2, "Search"](#)).
3. Click the **Remove** button. A confirmation message displays.
4. Click **OK**. The template is deleted, and removed from the **Templates** tab.

7.3. Exporting and Importing Templates

Like a virtual machine, a template can be imported or exported to a different Red Hat Enterprise Virtualization Manager system. Exporting templates allows you to distribute templates of virtual machines to users, including users who cannot directly access and use the templates in a specific installation of Red Hat Enterprise Virtualization Manager system.



Note

Only one export domain can be active in the data center. This means that the domain can be attached to either the source data center or the destination data center.

If a virtual machine was created using a template, the template is not automatically exported, because the template must exist in the destination domain for the virtual machine to work, the template must also be exported to the destination data center.

There are two methods of exporting and importing virtual templates:

- Exporting or Importing one or more templates.
- Exporting and Importing a domain of virtual machines and templates. Refer to [Import Existing ISO or Export Storage Domain](#) for more information.

7.3.1. Exporting Templates

Exporting templates to a different installation of Red Hat Enterprise Manager requires some preparation, for example, an export domain should exist, and be attached to the appropriate data center; any virtual machines using the templates must be shut down. You will also need to attach or detach the export domain as appropriate. See [To perform an export-import of virtual resources:](#). The **Virtual Machines** tab on the Templates Detail pane displays the virtual machines that are using the template and their status.

To export individual templates to the export domain:

1. Click the **Templates** tab. If the template you want to export is not visible in the list, perform a search to display the template on the results list (see [Section 1.2, “Search”](#)).
2. Click **Export**. The **Backup Template** dialog box displays.
3. Click **OK**. The export of the template begins, this can take some time. Use the **Events** tab to view the progress.
4. On the **Storage** tab, select the export domain. The Details pane of the Export storage domain displays. The successfully exported template displays on the **Template Import** tab of the export domain.

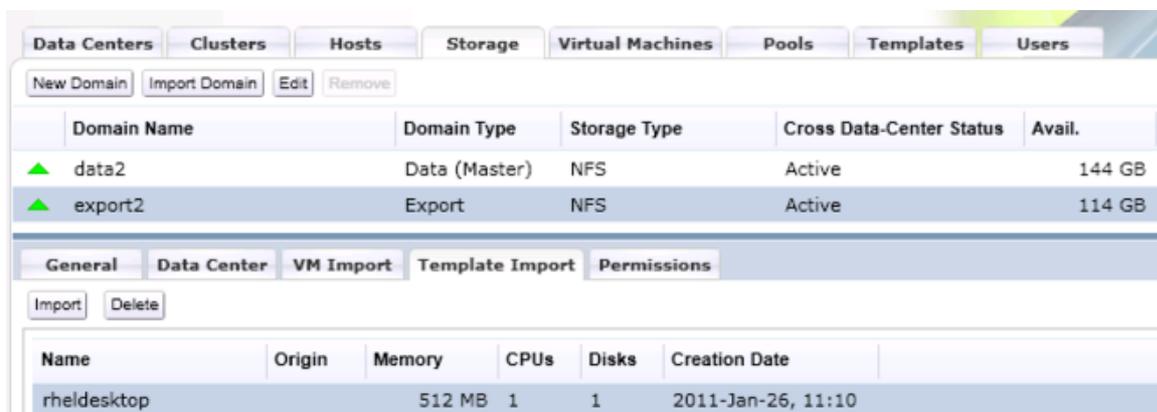


Figure 7.1. Template Import

5. You can repeat the procedure above to export each template that you need to migrate, so that the export domain has a number of templates before you start the import process.

7.3.2. Importing Templates

Once the templates are available in the export domain, they can be imported into a data center on the destination environment.

To Import the Template into the Destination Data Center

1. Detach the export domain from the source data center, and attach it to the destination data center. Refer to [To perform an export-import of virtual resources:](#) for more information.
2. On the **Storage** tab, select the export domain. The **Details** pane of the export domain displays.
3. On the **Details** pane, select the **Template Import** tab. Select the template that is to be imported.

The **Import** and **Delete** buttons are enabled on the **Template Import** tab.

4. Click **Import**. The **Import Template** dialog box displays. The names of the available templates display.
5. Select the name of the template, and select the **Destination Cluster** and **Destination Storage** of the destination data center.

Click **OK**.

6. A message displays:

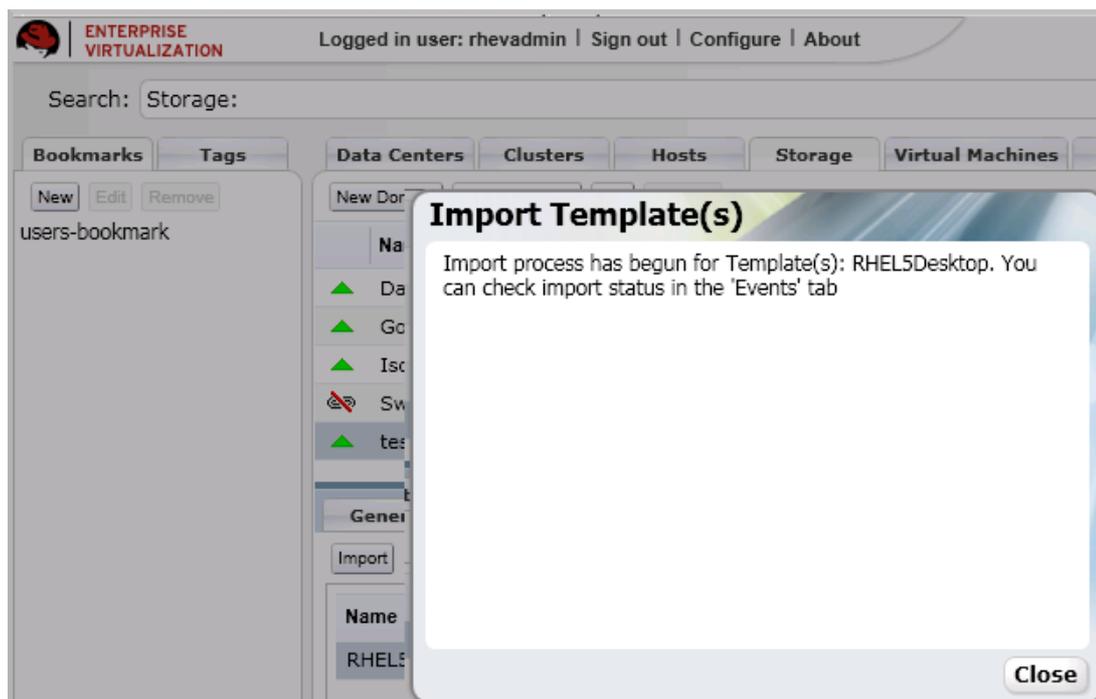


Figure 7.2. Import Templates

You can click **Close** to close the message box, and check the progress in the **Events** tab. The template is imported into the destination data center. This can take some time.

7. Eventually, the template displays in the **Template** tab on the **Details** pane of the data domain belonging to the destination data center. It also displays on the **Templates** tab with its changed cluster information indicating it's new location.



Figure 7.3. Imported Template

8. You can now use the template to create new virtual machines, or run existing imported virtual machines that are based on the template.

7.3.3. Backing Up Templates

Virtual Machines and Templates may need to be backed up from time to time, for example, before undertaking maintenance of hosts or storage servers. To back up virtual templates, use the export domain and export procedures as described in [Section 7.3, “Exporting and Importing Templates”](#).

7.4. Managing Permissions to Templates

Permissions to templates can be given to users or administrators. Templates can be assigned to end users. A template administrator can manage templates across the enterprise, configuring network and storage details.

7.4.1. Assigning Templates to Users

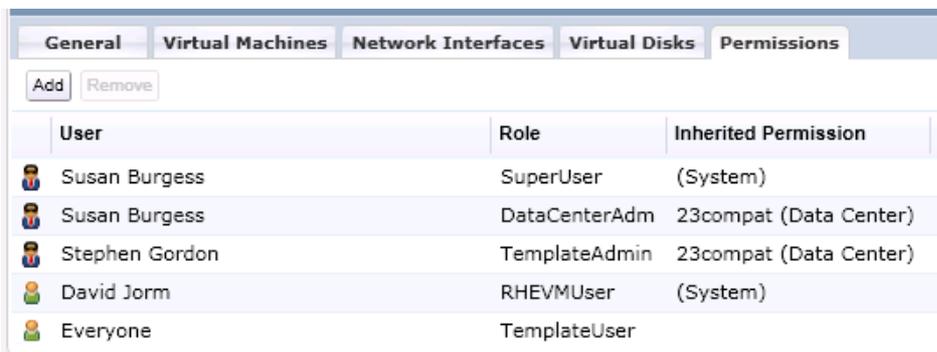
After the templates have been created and provisioned, they can be assigned to existing users. A template User role allows end users to use a template to create a virtual machine.

To assign a user to a template:

1. Click the **Templates** tab.

A list of templates displays. If the required template is not visible, perform a search (see [Section 1.2, “Search”](#)). Select the template that you want to edit, and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.



User	Role	Inherited Permission
 Susan Burgess	SuperUser	(System)
 Susan Burgess	DataCenterAdm	23compat (Data Center)
 Stephen Gordon	TemplateAdmin	23compat (Data Center)
 David Jorm	RHEVMUser	(System)
 Everyone	TemplateUser	

Figure 7.4. Template Permissions

2. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
3. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **UserTemplateBasedVM**.
4. Click **OK**.

The name of the user displays in the **Permissions** tab.



Note

You can only assign roles and permissions to existing users. See [Chapter 5, Users](#) for more information.

To remove a user role:

1. Click the **Virtual Machine** tab. A list of templates displays. If the required template is not visible, perform a search (see [Section 1.2, “Search”](#)).
2. Select the required template and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any. Note that you cannot remove users with inherited permissions.

3. Select the check box of the appropriate user.
4. Click **Remove**. The user is removed from the Permissions tab.

7.4.2. Assigning System Permissions for Templates Usage

A Template Administrator role gives the user permissions to create, manage and delete templates across the enterprise, including its network and storage details.

To assign a template administrator role:

1. Click the **Templates** tab.

A list of templates displays. If the required virtual machine is not visible, perform a search (see [Section 1.2, “Search”](#)).

User	Role	Inherited Permission
Susan Burgess	SuperUser	(System)
Susan Burgess	DataCenterAdm	23compat (Data Center)
Stephen Gordon	TemplateAdmin	23compat (Data Center)
David Jorm	RHEVMUser	(System)
Everyone	TemplateUser	

Figure 7.5. Template Permissions

2. Select the template that you want to edit, and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any.

3. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
4. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **TemplateAdmin**.

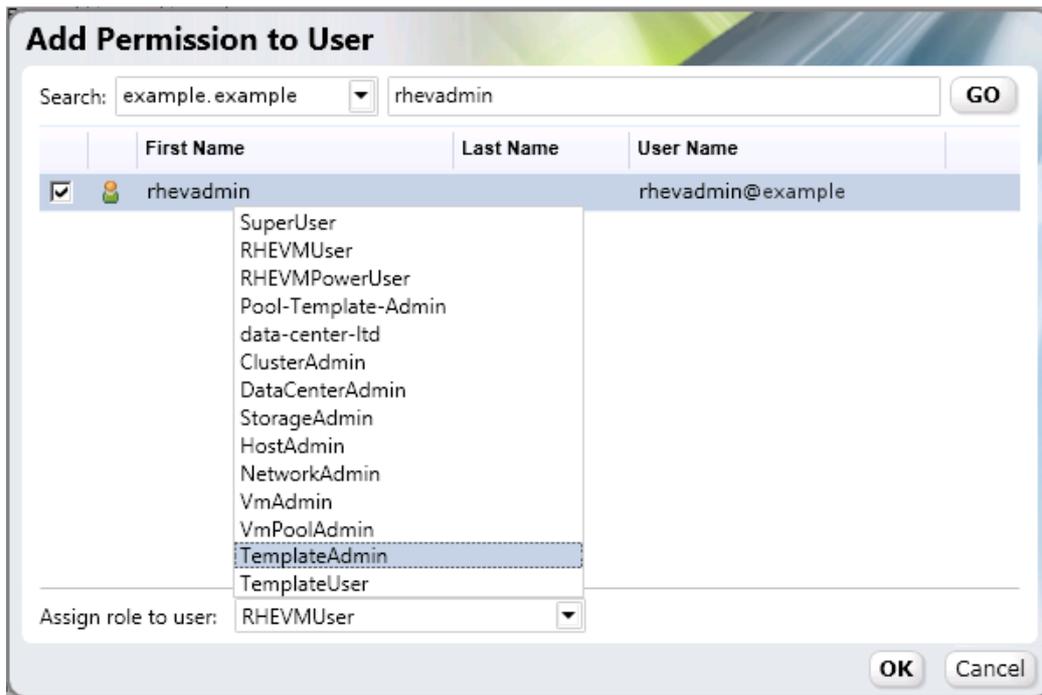


Figure 7.6. Assign TemplateAdmin Permissions

5. Click **OK**.

The name of the user displays in the **Permissions** tab.



Note

You can only assign roles and permissions to existing users (see [Chapter 5, Users](#)).

To remove a virtual machine system administrator role:

1. Click the **Templates** tab. A list of templates displays. If the required template is not visible, perform a search (see [Section 1.2, "Search"](#)).
2. Select the required template and click the **Permissions** tab from the Details pane.
The **Permissions** tab displays a list of users and their current roles and permissions, if any.
3. Select the check box of the appropriate user.
4. Click **Remove**. The user is removed from the **Permissions** tab.

Pools

Virtual desktops provide the enterprise with the benefits of high availability, scalability and interoperability. To optimize their use and maintenance, virtual machines can be placed in pools. Each desktop in a desktop pool can be used by any user in a particular group on demand, but a single desktop cannot be used concurrently by multiple users. Different desktop pools can be set up for different purposes. For example, there can be one desktop pool for the Marketing department, another for Research and Development, and so on. The user does not always get the same desktop, but gets an available desktop of the required type, from the appropriate pool.

This section describes how to create, maintain and remove pools of desktops. It includes allotting desktops to users, detaching users from desktops and assigning permissions to a virtual pool administrator.

8.1. Creating Desktop Pools

Use the **Pools** tab to create a set of identically configured desktops. A user associated with a desktop pool can use any available desktop in the pool. Virtual desktop pools must be created within a migration domain, and cannot be migrated outside a cluster.

Once a user logs on to a desktop that belongs to a pool, the desktop belongs to that user alone until it is released according to the pool type. When the user finishes working on a desktop, the desktop image reverts to base image, and becomes available to any user in the pool. Assigning users to a desktop pool is described later in this section.

To create a desktop pool:

1. Select **System** from the root of the **Tree** pane. Click the **Pools** tab. A list of pools displays.
2. Click the **New** button.
3. The **New Pool** dialog box displays the following tabs: **General**, **Pool**, **Console**, **Host**, **Resource Allocation** and **Boot Options**. If you choose Windows as the operating system, a **Windows Sys Prep** tab also displays. Ensure that you enter the mandatory information in the mandatory fields. Unfilled mandatory fields display with a colored border.
4. In the **General** group, enter the basic information about the pool, and the type of virtual machines required, as listed in the table.

Table 8.1. New Pool Dialog Box Fields

Field	Description	Notes
Data Center	Select an existing data center from the list.	The <i>Default</i> data center displays by default.
Host Cluster	The name of the host cluster to which the pool is attached. It can be hosted on any physical machine in the cluster depending on the policy rules.	This is the migration domain for the virtual machine. The <i>Default</i> cluster displays by default.
Name	The name of the new pool. Ensure it is a unique name. The name cannot be any longer than 15 characters,	A number is appended onto the name of the desktop pool to create a unique name for the virtual desktops. For example, if the desktop pool

Field	Description	Notes
	and must contain at least one alphabet, a-z.	is named HR, and has five virtual desktops, the names of the virtual desktops will be, HR-1, HR-2, HR-3, HR-4, HR-5. Ensure that the name is succinct rather than verbose. The maximum length of a virtual machine name is 15 characters. Follow the operating system's rules for virtual machine names.
Description	A meaningful description of the new pool.	
Number of VMs	Enter the Number of VMs (Virtual Machines) to create for the pool.	
Based on Template	Select an existing template to create the virtual machines from an existing model. See Section 6.2.1, "Creating Virtual Machines from Existing Templates"	The field displays the list of existing templates in the storage domain.
Memory Size (MB)	The amount of memory assigned to each virtual machine.	Consider the processing and storage needs of the applications that are intended to run on the virtual machines. The maximum allowable memory for a virtual machine is 256GB, allowing even the most memory-intensive enterprise workloads to be virtualized.
Total Cores	The processing power allocated to the virtual machine, as CPU Cores, from 1 to 16 on the slider bar.	It is recommended that you do not assign a number that is too high, or more cores than actually exist on the physical host.
CPU Sockets	The number of CPU sockets for the virtual machine from 1 to 16 on the slider bar.	It is recommended that you do not assign a number that is too high, or more CPUs than actually exist on the physical host.
Operating System	The operating system. Valid values include a range of Windows and Linux variants.	This is a display only field, as no operating system is actually installed during this process.

- In **Pool**, select one of the following pool types: *Manual* or *Automatic*.

Table 8.2. Pool Types Fields

Field	Description	Notes
Manual	The administrator is responsible for explicitly returning the desktop to the pool.	The desktop reverts to the original base image after the administrator returns it to the pool.
Automatic	When the desktop is shut down, it reverts to its base image and is then returned to the desktop pool.	

- If the **Operating System** chosen was **Windows**, the **Windows Sys Prep** group displays.

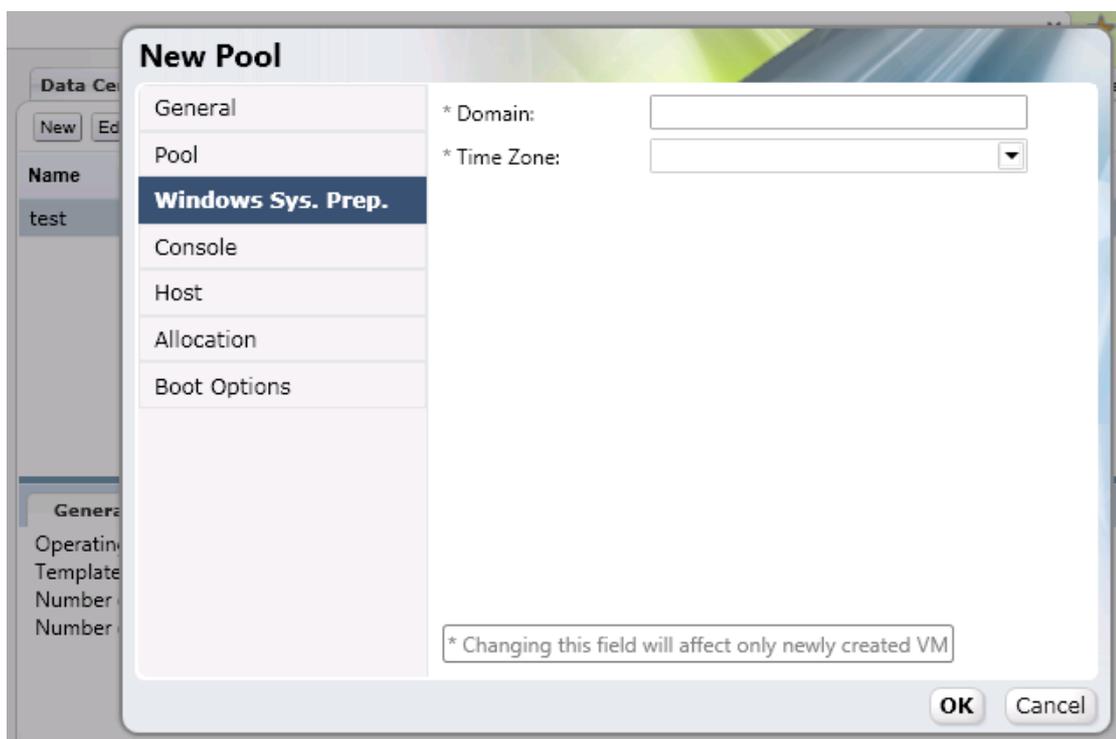


Figure 8.1. New Pool Windows Sys Prep

Enter the following information:

Table 8.3. Windows Sys Prep Fields

Field	Description	Notes
Domain	Enter the domain in which the virtual machines are to be created.	
Time Zone	Enter the time zone in which the virtual machines are to run.	This is the time zone for the virtual machines, and not necessarily the time zone for the physical host on which the virtual machines are to run.

7. Enter information in the **Console** group, to specify the communication protocol to be used, whether USB devices are allowed, and the number of monitors that the virtual machines are permitted.

Table 8.4. New Virtual Machine Dialog Box Fields

Field	Description	Notes
Protocol	Define the display protocol to be used. Select either: <ul style="list-style-type: none"> • SPICE • VNC 	Select SPICE for Windows or Linux virtual machines. This is the recommended protocol. or select VNC for Linux virtual machines if desired.
USB Policy	Select Enabled or Disabled to indicate whether a USB device can be inserted into the client machine.	The USB policy editor can be used to set up policies.
Monitors	For Windows virtual machines, the number of monitors that the virtual machine can have.	Consider the processing and storage needs of the applications that are intended to run on the virtual machine.

8. Enter information in the **Host** group if you wish to define the specific hosts on which the virtual machines are to be run and specific migration details.

Table 8.5. New Pool Host Fields

Field	Description	Notes
Run On	The host for the virtual machines to run on. <ul style="list-style-type: none"> • Select Any Host in Cluster if the virtual machines can run on any available host in the cluster. • Select Specific if the virtual machines are to run on a particular host in the cluster. Select the host from the list of available hosts in the cluster. 	
Run/Migration Options	Further options for running virtual machines <ul style="list-style-type: none"> • Select Run VM only on the selected Host to start and run the virtual machine only on the host specified in the Run On field. • Select Do not migrate VM to prevent migration in mid-operation to other hosts in the cluster, for example, in case of overload or fencing of the host. However, the 	

Field	Description	Notes
	virtual machine may start on any host.	

- Enter information in the **Resource Allocation** group. This step specifies the storage requirements for all the virtual machines in the pool.

Table 8.6. New Pool Allocation Fields

Field	Description	Notes
Storage Domain	The name of the storage domain where the images of the virtual machines of the pool will be stored.	
Provisioning	The type of storage required for the virtual machines. Select either: <ul style="list-style-type: none"> • Thin • Clone 	Refer to Section 6.1.3, "Understanding Virtual Machine Storage" for a description of these options.
Memory Allocation/Physical Memory Guarantee	The amount of physical memory that must be reserved for this pool.	

- Enter information in the **Boot Sequence** fields to specify how the virtual machines are to attempt to boot.

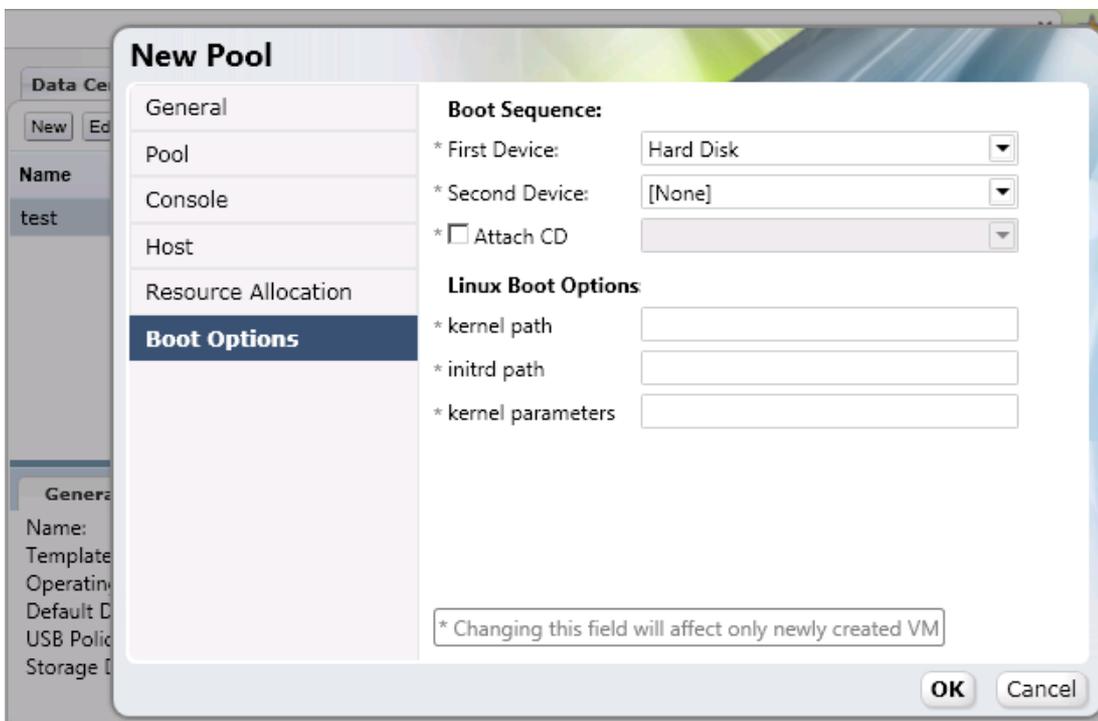


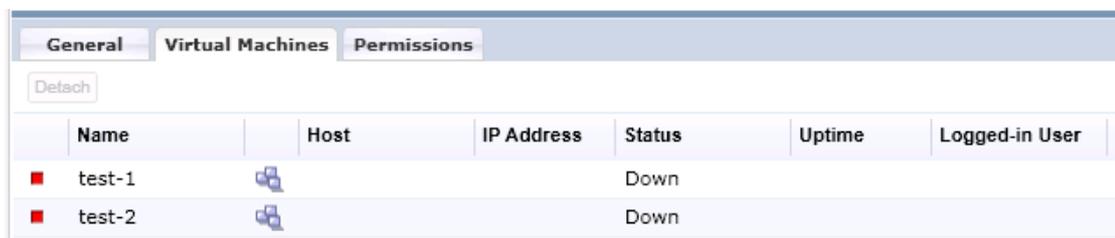
Figure 8.2. New Pool - Boot Options - Boot Sequence

Table 8.7. New Pools Boot Device Sequence Fields

Field	Description	Notes
First Device	<ul style="list-style-type: none"> • HardDisk • CD-ROM • Network (PXE) 	<p>After installing a new virtual machine, the new virtual machine must go into Boot mode before powering up. Select the first device that the virtual machine must try to boot:</p> <ul style="list-style-type: none"> • Hard Disk to boot from the hard disk (though if this is a blank virtual machine, it will obviously not boot from the hard disk) • CD-ROM to boot from the CD • Network (PXE) to boot from the network.
Second Device	<p>Select the second device that the virtual machine must try to boot:</p> <ul style="list-style-type: none"> • Hard Disk • CD-ROM • Network (PXE) 	<p>Select the second device for the virtual machine to use to boot if the first device is not available. The first device selected in the previous option does not appear in the options.</p>
Attach CD	<p>A list of available CD-ROMs appear if Attach CD is selected.</p>	<p>Select the appropriate operating system ISOs available on the system.</p>
Linux Boot Options	<p>If the virtual machine is a Linux machine, you can enter customized options for the kernel path,initrd path, kernel params (parameters).</p>	

11. Click **OK**. If any mandatory fields have been omitted, you are prompted to enter information. The required fields display with a red border. Enter the requisite information.

A pool of the specified number of identical virtual machines is created. You can view these virtual machines in the **Virtual Machines** tab (on the **Details** pane), or on the main **Virtual Machines** tab. Note the icon denoting that the virtual machines are part of a pool. During creation, the desktops have a status of **Image Locked**, followed by a status of **Down**. You need to start the desktops in the **Virtual Machines** tab. See [Chapter 6, Managing Virtual Resources](#).



	Name	Host	IP Address	Status	Uptime	Logged-in User
■	test-1			Down		
■	test-2			Down		

Figure 8.3. Desktop Pool VMs in the Details Pane

8.2. Managing Desktop Pools

Desktop pools will require maintenance from time to time, for example to add or remove users, or detach desktops.

8.2.1. Assigning Users to a Desktop Pool

The system administrator allocates users to a desktop pool. Once a user logs on to a desktop that belongs to a pool, the desktop belongs to that user until it is released according to the pool type. A virtual desktop can only be used by a single user at a time.

When the user finishes working on a desktop, the desktop image reverts to the snapshot of its state before it was taken from the pool (the original “base image”), and is then available for use by any other user in the pool.



Note

Because all desktops are identical, a user of a pooled desktop can be assigned any desktop in the pool.

The **Permissions** tab on the **Details** pane of the **Pools** tab enables you to manage users for a pool. The available users are drawn from the directory service in use.

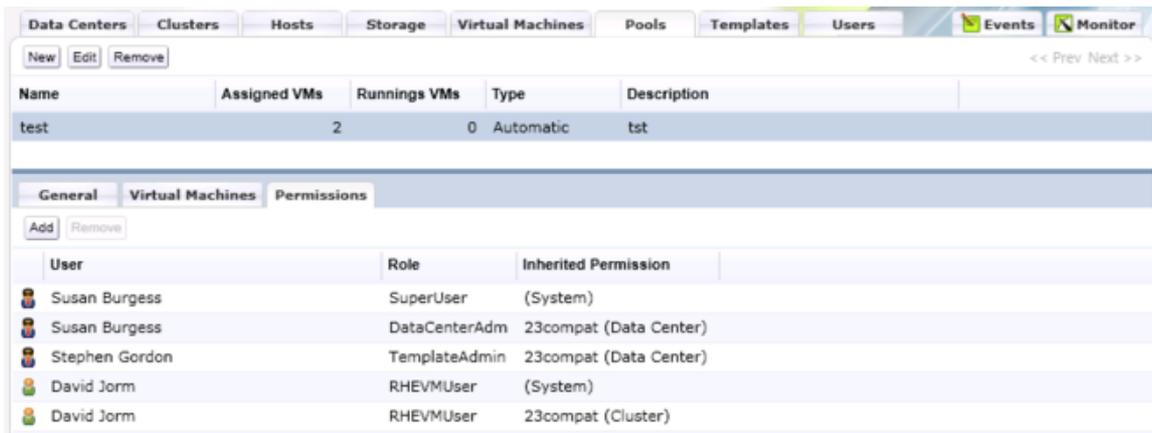


Figure 8.4. The Permissions tab in the Details Pane

1. On the **Pools** tab, select the desktop pool. The **Details** pane displays. Click the **Permissions** tab.
2. Click **Add**. The **Add Users to Pool** dialog box displays.

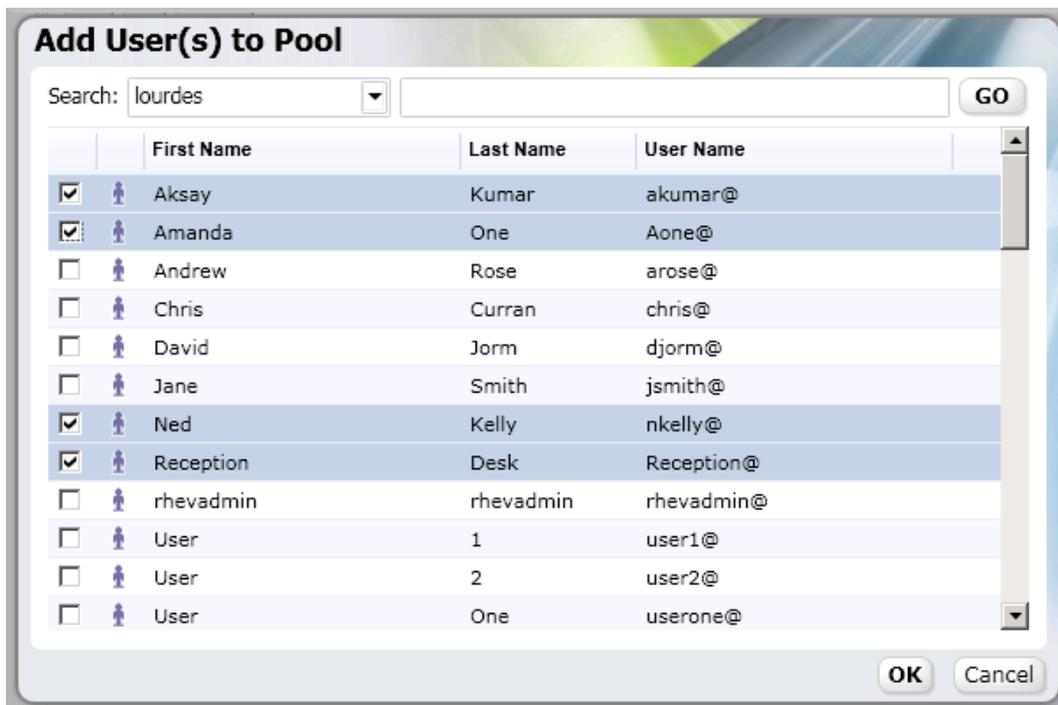


Figure 8.5. Add Users to Pool

3. Click the **Go** button to view all existing users, or enter search criteria for a set of users.
4. Select the users from the list, and click **OK**.
5. The users are assigned to the pool and display in the **Permissions** tab on the **Details** pane of the selected pool.

Removing Users from a Pool

1. On the **Pools** tab, select the desktop pool. The Details pane displays. Click the **Permissions** tab.

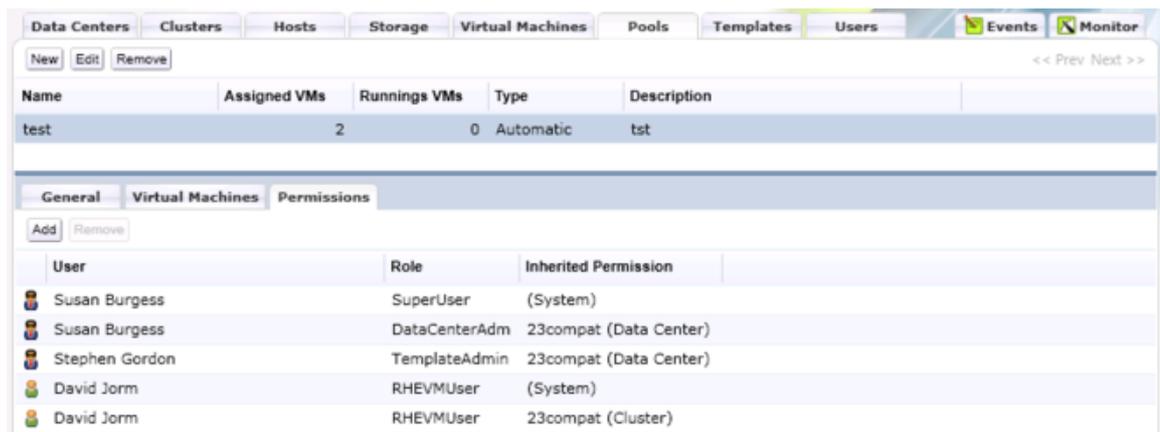


Figure 8.6. The Permissions Tab in the Details Pane

- Select the user to be removed. The **Remove** button is enabled.

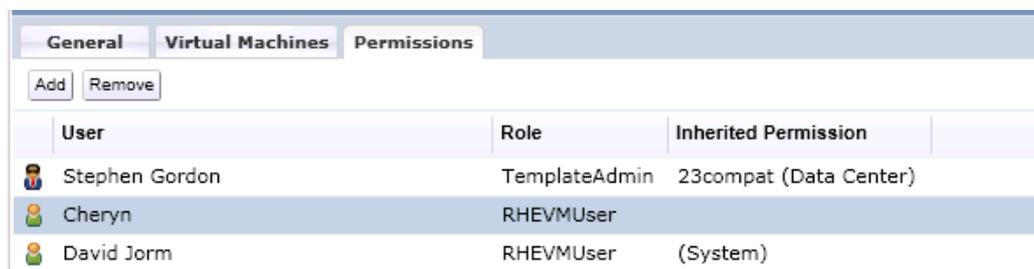


Figure 8.7. The Remove Button in the Permissions tab

- Click **Remove**. A dialog box displays, prompting to confirm the user or list of users. Click **OK**. The user is removed from the pool.

Note

The users are only detached from the desktop pool, they are not deleted.

8.2.2. Editing Desktop Pools

You can edit some details of a pool such as the **Name** and **Description**, you can change the pool type, and add virtual machines to the pool. However, you cannot change any parameters on the **Windows Sys Prep**, **Console**, **Allocation** or **Boot Sequence** groups.

Note

Changes do not apply to existing virtual machines, only to virtual machines that are created after the change is made.

To edit a desktop pool:

- Click the **Pools** tab.

2. Select the pool. If the pool you want to edit is not displayed, perform a search (see [Section 1.2, “Search”](#)).
3. Click the **Edit** button.

The **Edit pool** dialog box opens.

4. This dialog box is identical to the **New Pool** dialog box, except that fields that cannot be changed are disabled. Make other changes as appropriate. See [Section 8.1, “Creating Desktop Pools”](#).
5. Click **OK**.

The details of the desktop pools are updated.

To add virtual machines to a desktop pool:

1. Select the desktop pool.
2. Click the **Edit** button.

The **Edit pool** dialog box opens.

3. On the **General** tab, in the **Number of VMs** field, click the **Add more** button. A text box appears allowing you to indicate the number of additional virtual machines to be added to the pool.
4. Click **OK**.

The additional virtual machines display in the **Virtual Machines** tab on the **Detail** pane of the **Pools** tab. These virtual machines are identical to the existing virtual machines.

8.2.3. Detaching Desktops from a Pool

You can detach one or more desktops from a pool, or you can delete the desktop. To delete the desktop, you must first detach it from the pool and then delete it. To delete a desktop, see [Section 6.7.6, “Removing Virtual Machines”](#)

To detach desktops from a desktop pool:

1. Click the **Pools** tab. The list of desktop pools displays.
2. Select the pool containing the desktop that is to be detached.

The **Virtual Machines** tab on the **Details** pane displays a list of the virtual machine in the pool.

3. Select the desktop(s) that you want to remove, and click **Detach**. A confirmation message box displays.
4. Click **OK**.

The desktop is detached from the pool, however, it still exists in the system, and can be viewed and accessed from the **Virtual Machines** tab. Note that the icon changes to denote that the virtual machine is a standalone desktop.

Name	Cluster	Host	IP Address
rhel55	23compat		
rhel55desktop	23compat		
rhelanother	23compat		
rheltestdesktop	23compat		
test-1	23compat		
test-2	23compat		

Figure 8.8. A Detached Desktop

8.3. Managing Permissions to Pools

An administrator role can be assigned to users for a pool of virtual machines. A virtual machine pool administrator role, called VM Pool Administrator in the GUI, gives permissions to create, manage and delete the pool, including making changes to its configuration. A pool administrator can also perform basic operations on virtual machines in the pool. Inherited permissions apply to virtual machine pools, hence the superuser has permissions on all pools and the data center administrators (if any), have permissions on pools within their data center.

To assign administrator permissions to a virtual machine pool:

1. Click the **Pools** tab. A list of virtual machines displays. If the required virtual machine is not visible, perform a search (see [Section 1.2, "Search"](#)).
2. Select the virtual machine that you want to edit, and click the **Permissions** tab from the **Details** pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any. Inherited permissions apply to virtual machine pools, hence superusers and the data center administrators (if any), display in the list.

User	Role	Inherited Permission
Susan Burgess	SuperUser	(System)
Susan Burgess	DataCenterAdm	23compat (Data Center)
Stephen Gordon	TemplateAdmin	23compat (Data Center)
David Jorm	RHEVMUser	(System)
David Jorm	RHEVMUser	23compat (Cluster)

Figure 8.9. Pools Permissions

3. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter a Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches is displayed in the results list.

4. Select the check box of the user to be assigned the permissions. Scroll through the **Assign role to user** list and select **VMPoolAdmin**.

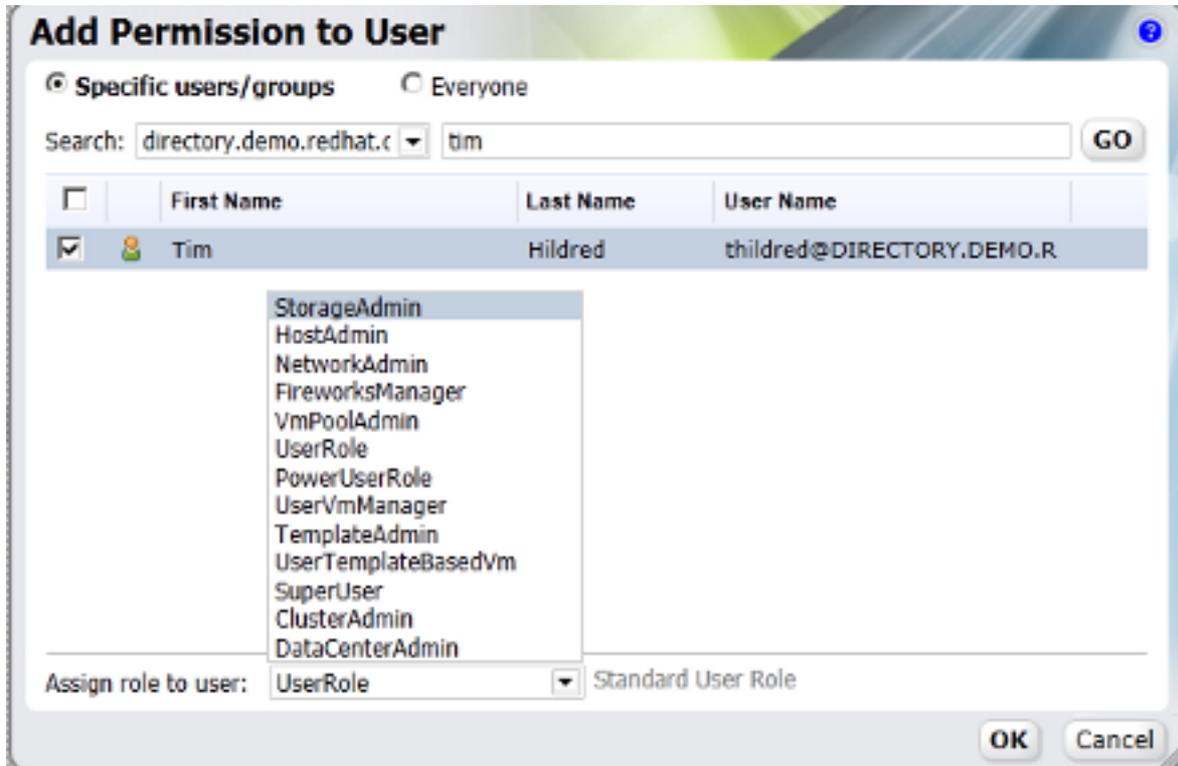


Figure 8.10. Assigning VMPoolAdmin Permission

5. Click **OK**.

The name of the user displays in the **Permissions** tab.



Note

You can only assign roles and permissions to existing users. See [Chapter 5, Users](#) for more information.

To remove the role from a user:

1. Click the **Virtual Machine** tab. A list of virtual machines displays. If the required virtual machine is not visible, perform a search (see [Section 1.2, "Search"](#)).
2. Select the required virtual machine and click the **Permissions** tab from the Details pane.

The **Permissions** tab displays a list of users and their current roles and permissions, if any. Note that you cannot remove users with inherited permissions.

3. Select the appropriate user.
4. Click **Remove**. A confirmation dialog displays, click **OK**. The user is removed from the **Permissions** tab.

8.4. Removing Desktop Pools

This section describes how to remove a desktop pool. Before removing a pool, all the virtual machines must be detached from the pool.

To remove a desktop pool:

1. Click the **Pools** tab.

The list of desktop pools displays.

2. Select the pool to be removed.

Click the **Virtual Machines** tab in the **Details** pane of the selected pool to display a list of the pool's desktops.

3. Select all the desktops in the **Virtual Machines** tab and click **Detach**. The desktops are detached from the pool. However, they will continue to display in the **Virtual Machines** tab as standalone desktops.
4. Ensure the pool is selected on the **Pools** tab, and click the **Remove** button. A message displays prompting you to confirm the removal.
5. Click **OK** to confirm the removal of the pool. The pool is removed from the data center.



Note

Since all the desktops must be detached before the pool can be removed, it means that there is no interruption of service.

Part IV. Monitoring and Reporting

Monitoring Red Hat Enterprise Virtualization

A system administrator monitors the management environment to view the virtualization infrastructure's overall performance. This helps identify key areas in the virtualization environment that require attention or optimization. A system administrator gains up-to-date information on the performance and status of virtual environment components with the **Events** and **Monitor** tabs.

The **Monitor** tab displays instant information and statistics about the entire system. It generates usage graphs for storage, memory, CPU, and network traffic of hosts and virtual machines. It also lists high severity warnings.

The **Events** tab lists all warnings, errors, and other events that occur in the system.



Note — Reporting Tools

Red Hat Enterprise Virtualization also provides administrators with reporting tools that generate textual and graphical reports with greater detail than the standard monitoring tools. This book covers these reporting tools in [Chapter 10, Red Hat Enterprise Virtualization Reports](#) and [Chapter 11, History Database Reports](#).

9.1. Using the Monitoring Tools

The **Monitor** tab provides an overall picture of the Red Hat Enterprise Virtualization platform environment to the system administrator. It lists general information, such as the number of hosts and virtual machines currently running, plus displays memory usage, CPU usage, network usage and high severity events over the last 24 hours.

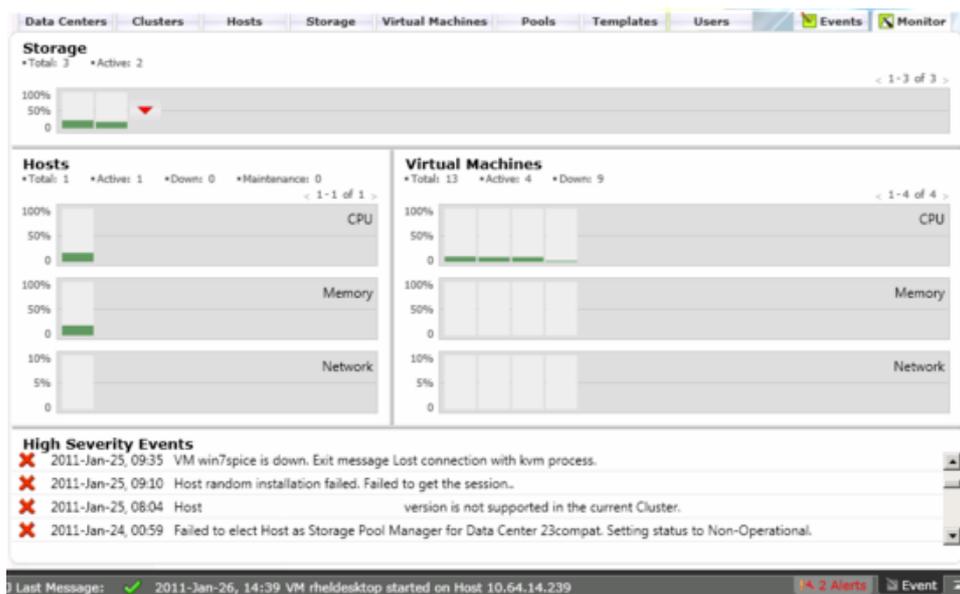


Figure 9.1. Monitor Tab

The System Monitor comprises of the following information:

- **Storage** - Graphs of the used space on each storage domain;

- **Hosts** - Graphs of CPU, memory and network traffic usage for each host;
- **Virtual Machines** - Graphs of CPU, memory and network traffic usage for each virtual machine; and
- **High Severity Events** - Viewable in the lower panel.

Drag the edges of each panel in the System Monitor to resize the panel.

9.1.1. Monitoring Storage

The Storage pane on the Monitor tab displays the storage domains usage graphs. Each domain uses a different color. Move the cursor over the graph of a storage domain to view the Available, Used, and Total capacity of the storage domain.

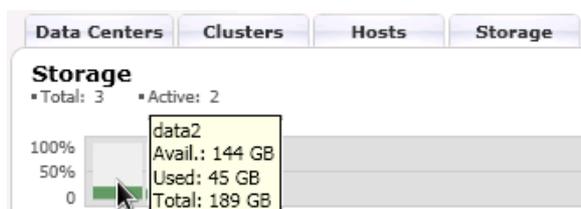


Figure 9.2. Monitor Storage

9.1.2. Monitoring Hosts

The Hosts pane displays the Memory, CPU, and Network usage graphs of the hosts. Move the cursor over the graphs of a host to view the CPU, Memory and Network usage of the hosts.

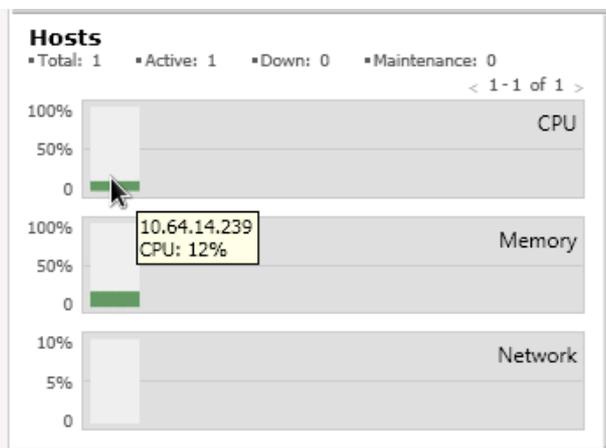


Figure 9.3. Monitor Storage

9.1.3. Monitoring Virtual Machines

The Virtual Machine pane displays the Memory, CPU, and Network usage graphs of virtual machines. Move the cursor over the graphs of a virtual machine to view the CPU, Memory and Network usage of the virtual machines.

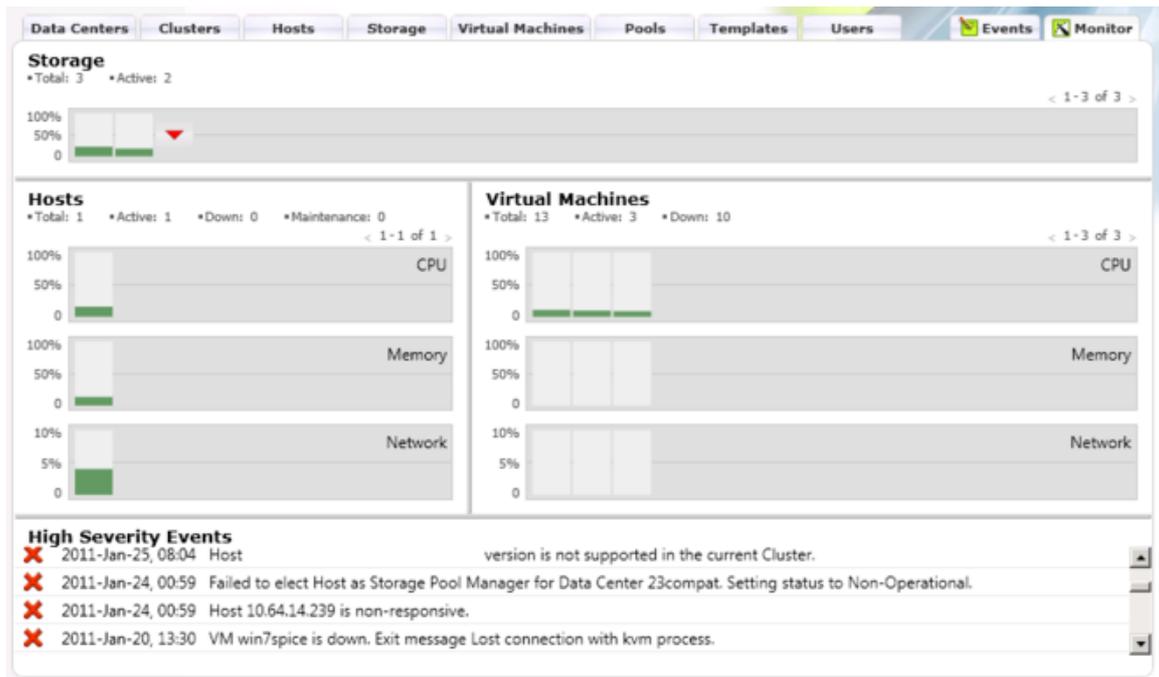


Figure 9.4. Monitor Virtual Machines

Moving the cursor over the virtual machine graphs also displays name of the virtual machine.

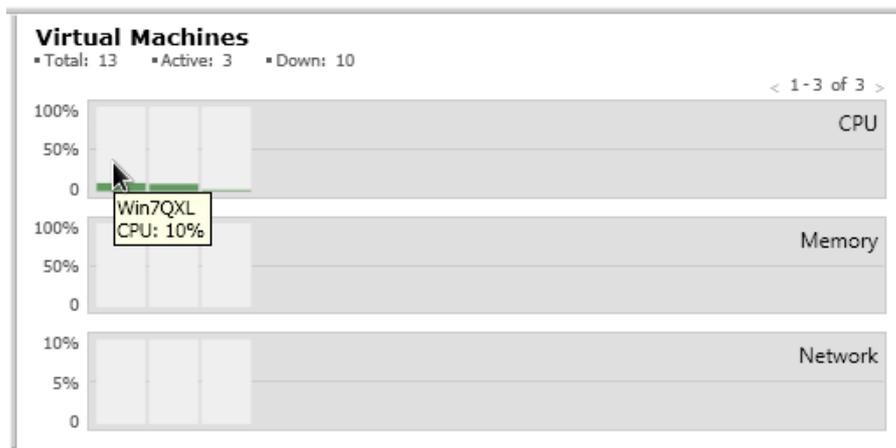


Figure 9.5. Monitor Virtual Machines

9.1.4. Monitoring High Severity Events

High Severity Events display in the lower panel of the **Monitor** tab.

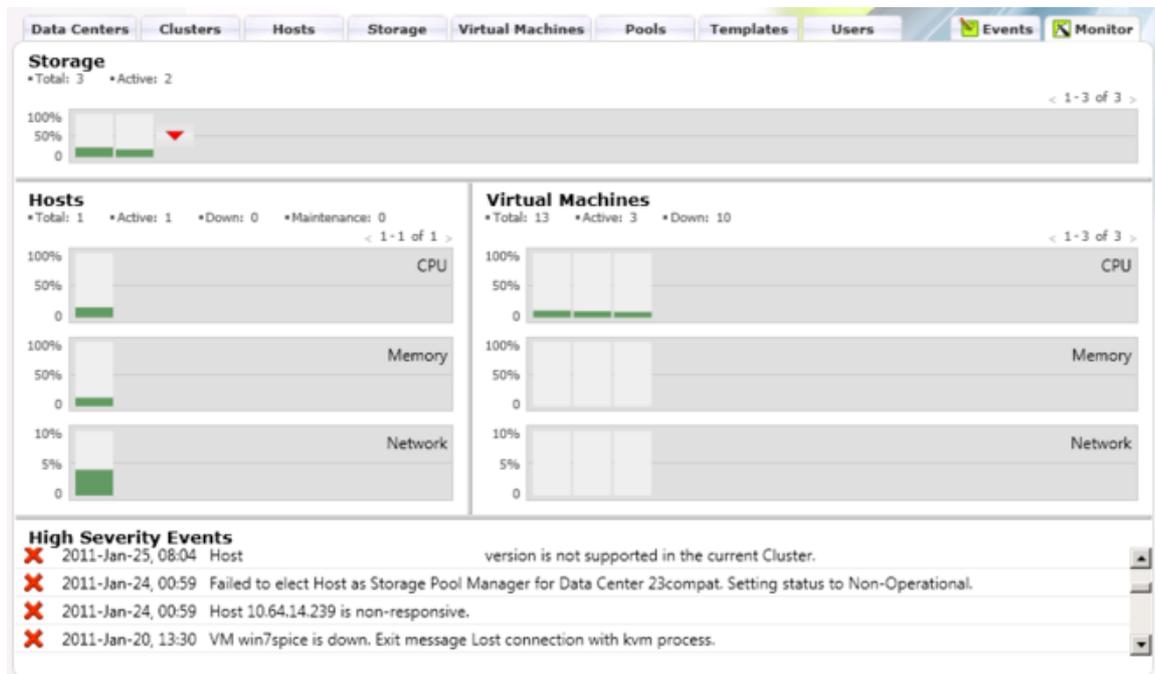


Figure 9.6. High Severity Events

9.1.5. Viewing the Event List

The **Event** list displays all system events. The types of events that appear in the Events tab are audits, warnings, and errors. The names of the user, host, virtual machine, and/or template involved in the event are also listed. This information helps determine the cause of the event. Click the column headers to sort the event list.



Figure 9.7. Event List

The following table describes the different columns of the **Event** list:

Column	Description
Event	The type of event. The possible event types are:

Column	Description
	<p>✔ Audit notification (e.g. log on).</p> <p>⚠ Warning notification.</p> <p>✘ Error notification.</p>
Time	The time that the event occurred.
Message	The message describing that an event occurred.
User	The user that received the event.
Host	The host on which the event occurred.
Virtual Machine	The virtual machine on which the event occurred.
Template	The template of the virtual machine where the event occurred.

9.1.6. Viewing Alert Information

The Alerts pane lists all important notifications about the virtualization environment.

Alerts display in the lowermost panel on top of the manager interface. Drag the top of the Alert pane to resize it or click the minimize/maximize icon in the top-right of the pane to hide or show it.

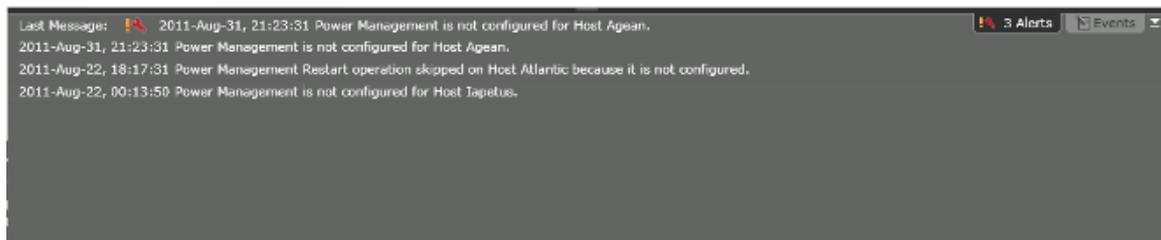


Figure 9.8. Alerts Tab

The Alerts pane also contains an Events list. Click the **Events** tab in the Alerts pane to display the Events list. See [Section 9.1.5, “Viewing the Event List”](#) for more information about the Events list.

Red Hat Enterprise Virtualization Reports

Red Hat Enterprise Virtualization provides a suite of pre-configured reports and dashboards that enable you to monitor the system. The reports module is based on open source reporting tools JasperReports and JasperServer. This chapter describes how to access Red Hat Enterprise Virtualization reports, navigate to the pre-configured reports and dashboards, and create your own ad hoc reports.

10.1. Overview

This section covers JasperReports and JasperServer, the open source reporting tools that Red Hat Enterprise Virtualization uses.

10.1.1. JasperReports & JasperServer

JasperReports

JasperReports is an open source reporting tool, capable of being embedded in Java-based applications. It produces reports which can be rendered to screen, printed, or exported to a variety of formats including PDF, Excel, CSV, Word, RTF, Flash, ODT and ODS. JasperReports integrates with JasperServer, an open source reporting server for JasperReports. Using JasperServer, reports built in JasperReports can be accessed via a web interface.

JasperServer in Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization provides a customized implementation of JasperServer, which contains a range of pre-configured reports and dashboards, plus the ability to create ad hoc reports.

10.1.2. Online Help

JasperServer provides extensive online help. Use the online help to find information on common administration tasks and the JasperServer product in general. This section provides information on the reports available for Red Hat Enterprise Virtualization and the customizations that integrate JasperServer with Red Hat Enterprise Virtualization. To navigate to the online help facility, click on **Help** in the top left-hand corner of any screen, as shown in [Figure 10.1, “Red Hat Enterprise Virtualization Reports online help”](#).



Figure 10.1. Red Hat Enterprise Virtualization Reports online help

10.1.3. Red Hat Enterprise Virtualization History Database

Red Hat Enterprise Virtualization Reports uses data from the Red Hat Enterprise Virtualization History Database. [Chapter 11, History Database Reports](#) contains instructions on how to generate reports directly from the database. The *Red Hat Enterprise Virtualization Technical Reference Guide* provides details on all the configuration and history views available in the history database. See also [Appendix D, Reports Schema](#).



Important:: Reports require sufficient data

To produce meaningful reports, sufficient data must exist in the history database. Most reports use values aggregated on a daily basis. Meaningful reports can only be produced if data for at least several days is available. In particular, because trend reports are designed to highlight long term trends in the system, a sufficient history is required to highlight meaningful trends.

10.2. System Requirements

The Red Hat Enterprise Virtualization Manager Reports tool supports the following browsers:

- In Red Hat Enterprise Linux 5.7 - Firefox 3.6
- In Red Hat Enterprise Linux 6 - Firefox 3.6
- In Windows XP - Internet Explorer 8
- In Windows 7 - Internet Explorer 8 and Internet Explorer 9
- In Windows Server 2008 - Internet Explorer 8 and Internet Explorer 9

10.3. Accessing Reports and Dashboards

This section describes how to login to the system and navigate to reports and dashboards.

10.3.1. Logging in

To access reports, navigate to the reports portal at: **http://server.example.com/rhev-
reports**. A login screen for Red Hat Enterprise Virtualization Reports displays. Enter your login credentials.

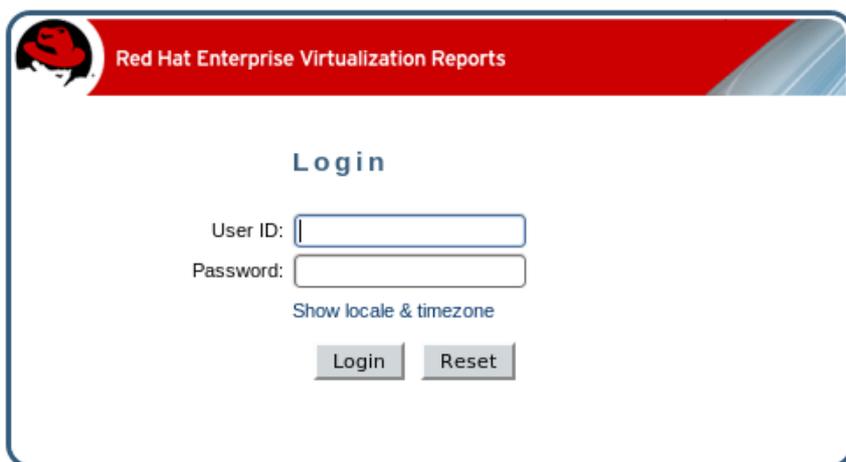


Figure 10.2. Red Hat Enterprise Virtualization Reports login screen

Red Hat Enterprise Virtualization Reports prompts you to set an administrator password during installation. Red Hat Enterprise Virtualization Reports does not provide default passwords.

The process for adding additional users is described in [Section 10.3.3, “Managing Users”](#)

Enter your credentials and click **Login**. If your credentials are valid, the main reporting screen displays.

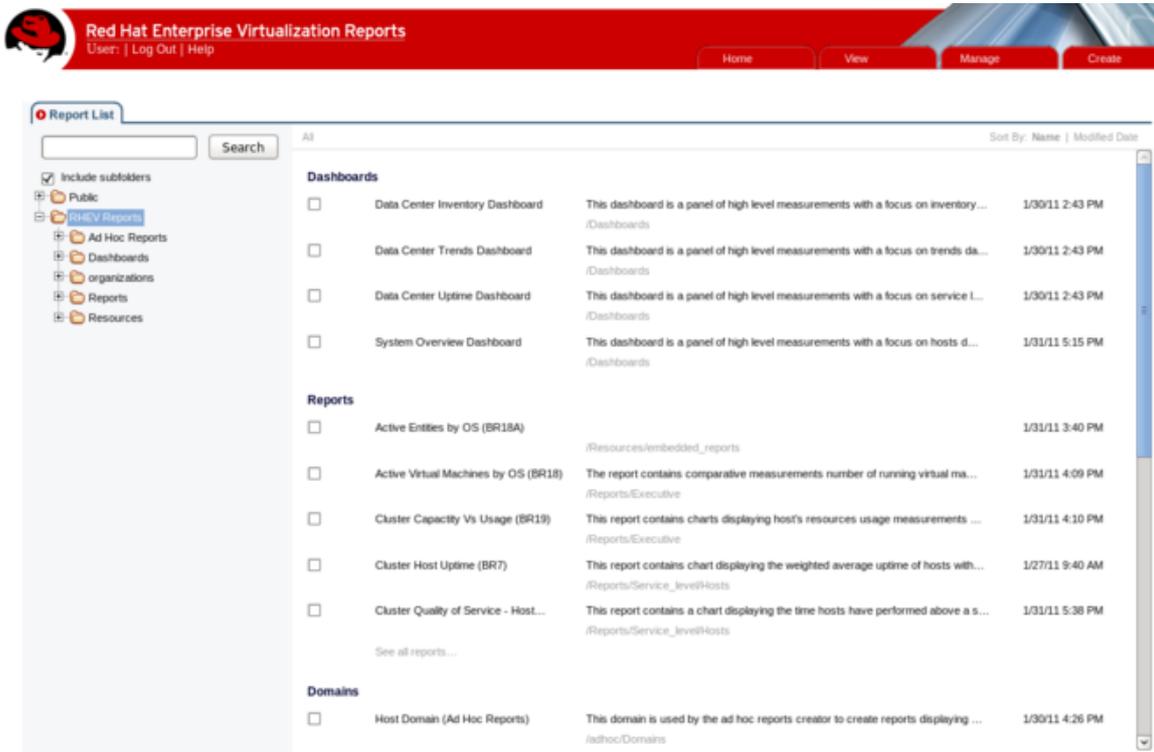


Figure 10.3. Red Hat Enterprise Virtualization Reports main screen

10.3.2. Navigating Reports

Use the navigation tree on the left of the screen to navigate to the required reports, as shown below.

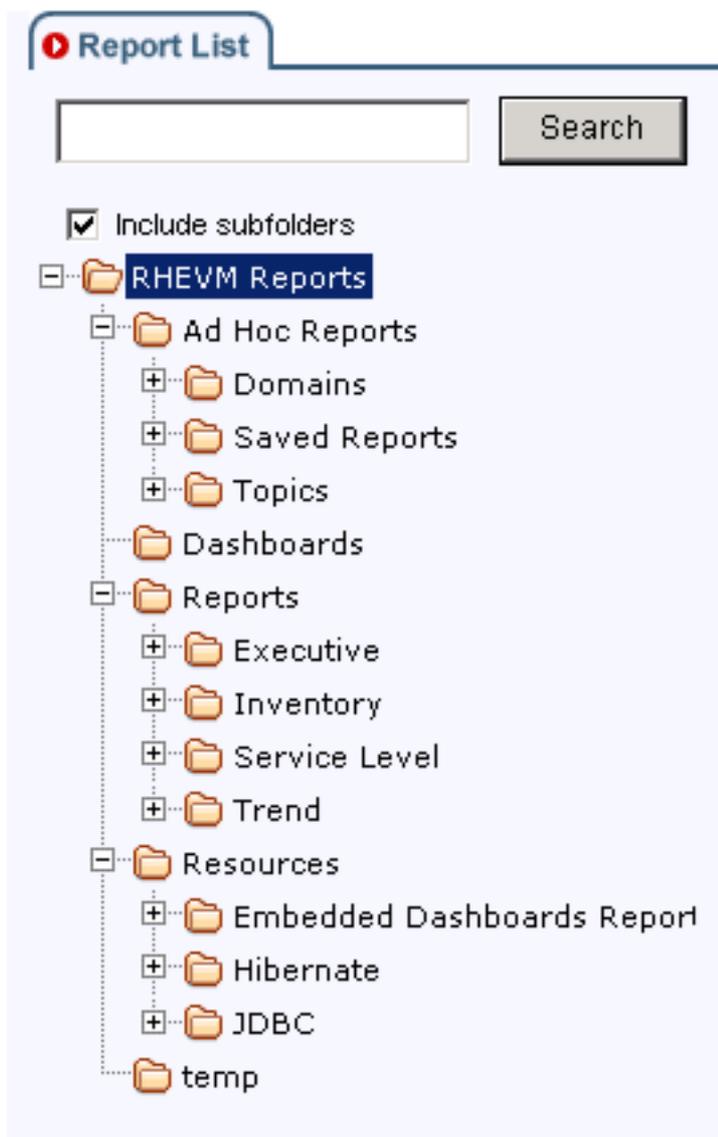


Figure 10.4. Red Hat Enterprise Virtualization Reports navigation tree - administrator's view

The folders in the navigation tree contain the following resources.

Table 10.1. Navigation Folders

Folder	Description
RHEVM Reports	Contains sub-folders for all reporting resources.
Ad Hoc Reports	Contains report types for generation of ad hoc reports, which are created with the Ad Hoc Reports Tool. See Section 10.6, “Ad Hoc Reports” for details.
Dashboards	Contains all dashboards available in the system. See Section 10.5, “Dashboards” for details.
Reports	Contains all reports available in the system. See Section 10.4, “Running Reports” for details.

Folder	Description
Resources	Contains configuration and system resources. Should not be changed by users.

10.3.3. Managing Users

To manage users, you will need to follow this procedure to get to Red Hat Enterprise Virtualization Reports' user-management menu:

Procedure 10.1. Accessing the Red Hat Enterprise Virtualization Reports user-management menu

1. While signed in to Red Hat Enterprise Virtualization Reports, hover over the **Manage** tab in the top-right part of the screen.
2. Click on **Users** in the drop-down menu that appears.
3. The **Manage Users** interface appears.

The **Manage Users** interface contains three panes:

- **Organizations**
- **Users**
- **Details**

4. Select a set of reports in the **Organizations** pane by clicking on the name of the report.
5. Select a user in the **Users** pane by clicking on the name of the user. Information about the user displays in the **Details** pane.
6. The **details** pane contains these four fields:
 - The user's name
 - Assigned Roles
 - Profile Attributes
 - Whether the user is enabled.

At the top of the **Details** pane, a menu displays three options. These options are:

- **Edit**
- **Login as User**
- **Delete User**

This menu allows you to manage users.

There are two roles, each of which provides a different level of permissions:

1. **ROLE_ADMINISTRATOR** - Can create/edit/delete reports, dashboards, ad hoc reports, and manage the server.
2. **ROLE_USER** - Can create/edit/delete ad hoc reports and view reports and dashboards.

Other roles can be created and assigned. For information on how to create and assign other roles, please refer to the JasperServer documentation.



Figure 10.5. Manage tab

Detailed documentation for user management and other system functions is provided in the online help, see [Section 10.1.2, “Online Help”](#) for details.

10.4. Running Reports

Reports provide time-sliced views of the system's status and workload. Report parameters define the scope of the report and export formats for the report, including the following: PDF, Excel, CSV, Word, RTF, Flash, ODT and ODS.

10.4.1. Report Parameters

The user defines parameters at run time to create reports. Report parameters define the scope and time scale of the report. When running a report, a prompt asks the user to provide parameters, if available.

To view the required parameters for a report, click the report in the Reports List. See [Figure 10.6, “Red Hat Enterprise Virtualization Reports - Reports List”](#). The Report Parameter Selection dialog displays. Note that the dialog is contextual and differs from report to report.

Reports		
Active Virtual Machines by OS (BR18)	The report contains comparative measurements number of running virtual machines and O...	8/11/11 12:46 PM
	/Reports/Executive	
Cluster Capacity Vs Usage (BR19)	This report contains charts displaying host's resources usage measurements (CPU co...	8/11/11 12:46 PM
	/Reports/Executive	
Cluster Host Uptime (BR7)	This report contains chart displaying the weighted average uptime of hosts within a selec...	8/11/11 12:46 PM
	/Reports/Service_level/Hosts	
Cluster Quality of Service - Hosts (BR6)	This report contains a chart displaying the time hosts have performed above a selected t...	8/11/11 12:46 PM
	/Reports/Service_level/Hosts	
Cluster Quality of Service - Virtual ...	This report contains a chart displaying the time virtual machines have performed above a ...	8/11/11 12:46 PM
	/Reports/Service_level/VMs	

Figure 10.6. Red Hat Enterprise Virtualization Reports - Reports List

Click the name of any report. The **Report Parameter Selection Dialog** appears:

Example 10.1. Report Parameter Selection

The **Report Parameter Selection Dialog** consists of a number of drop-down menus that define the report's parameters:

The parameters specified in *Example 10.1, "Report Parameter Selection"* produce the following report:

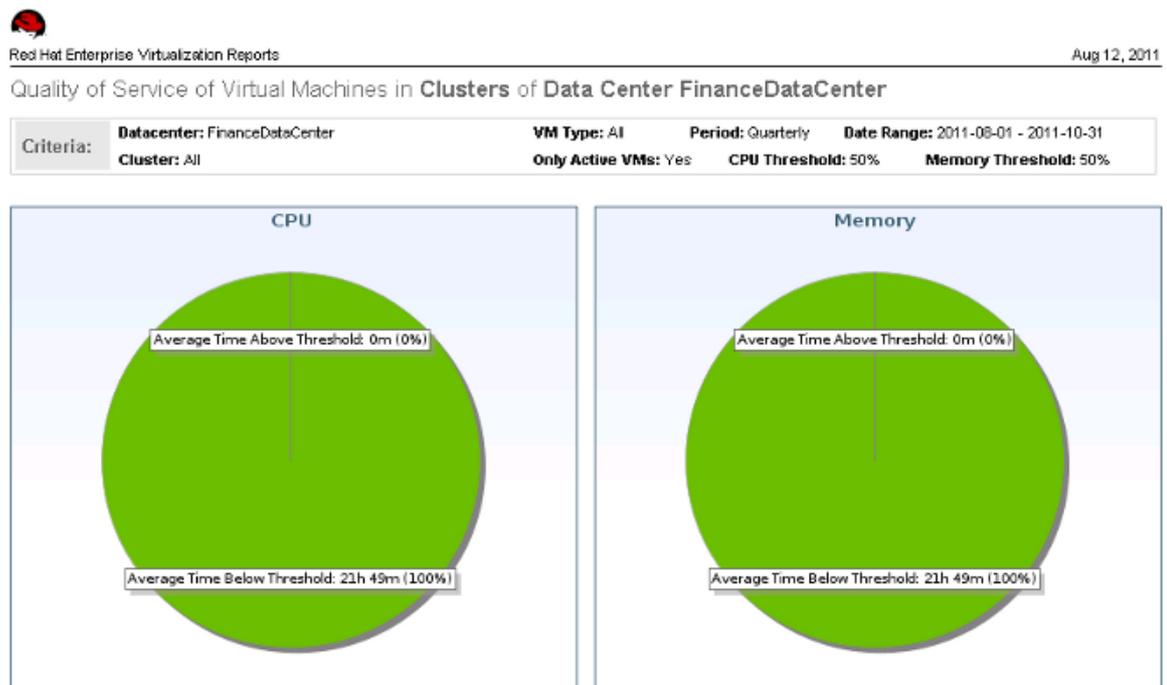


Figure 10.7. Jasper Reports Sample Report

Required parameters

Parameters marked with an asterisk (*) are required. In [Example 10.1, “Report Parameter Selection”](#), all parameters are required.

Cascading parameters

Many report parameters are cascading input fields. This means the selection made for one parameter changes the options available for another parameter. In [Example 10.1, “Report Parameter Selection”](#) the **Data Center** and **Cluster** parameters are cascading. Once a user selects a data center, only clusters within that data center are available for selection. Similarly, if a user selects a cluster, the **Host Type** field updates to show only host types that exist in the selected cluster. Cascading parameters filter out objects that do not contain child objects relevant to the report. For example, a report pertaining to virtual machines removes the selection of clusters that do not contain virtual machines. A report pertaining to both virtual machines and hosts only provides a selection from clusters containing both virtual machines and hosts.

Deleted objects

Objects deleted (removed) from the system are still recorded in the reporting history database. Select deleted objects, such as clusters, data centers and hosts, as values for report parameters if required. The bottom of the parameter options list shows deleted objects, which are suffixed with the date of removal from the system.

10.4.2. Executive Reports

Executive reports provide a high-level overview of the system status, including details appropriate for planning capacity expansion and critical decision making. More detailed reports display under **Trend**. The following executive reports are available.

10.4.2.1. Active Virtual Machines by Operating System

The **Active Virtual Machines by OS** report shows a summary of the number of active virtual machines in a given time period, broken down by operating system. The following parameters are provided to run this report:

Table 10.2. Active Virtual Machines by OS Parameters

Parameter	Description
Data Center	The report includes only virtual machines in the selected data center. The options list shows only data centers that contain virtual machines.
Cluster	The report only includes virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all virtual machines in the selected data center.
Virtual Machine Type	The report only includes virtual machines of the selected type. Possible types are Server and Desktop . The options list shows only types that exist in the selected data center and cluster. If All is selected, the report includes all virtual machine types.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly

Parameter	Description
	reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.2.2. Cluster Capacity Vs Usage

The **Cluster Capacity Vs Usage** report shows the relationship between system capacity and usage (workload) over a given time period. Capacity is expressed in terms of CPU cores and physical memory, while usage is expressed as vCPUs and virtual machine memory. The following parameters must be provided to run this report:

Table 10.3. Cluster Capacity Vs Usage Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list contains only data centers that contain clusters.
Cluster	The report only includes the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all clusters in the selected data center.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the Dates parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.2.3. Host OS Break Down

The **Host OS Break Down** report indicates the number of hosts running each operating system version over a given time period. The following parameters must be provided to run this report:

Table 10.4. Host OS Break Down Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the

Parameter	Description
	report includes all hosts in the selected data center.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.2.4. Summary of Host Usage Resources

The **Summary of Host Usage Resources** report shows a scatter plot of average host resource utilization for a given time period in terms of CPU and memory usage. The following parameters must be provided to run this report:

Table 10.5. Summary of Host Usage Resources Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all hosts in the selected data center.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.3. Inventory Reports

Inventory reports enumerate the current resources in the system. Resources include hosts, storage domains and virtual machines. The following inventory reports are available.

10.4.3.1. Hosts Inventory

The **Hosts Inventory** report shows a list of all hosts in the selected data center and cluster. The following parameters must be provided to run this report:

Table 10.6. Hosts Inventory Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all hosts in the selected data center.
Host Type	The report includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the report includes all host types.
Only Active Hosts?	If Yes is selected, the report includes only active hosts, i.e. ones not deleted from the system. If No is selected, the report includes both active and deleted hosts.

10.4.3.2. Storage Domain Size Over Time

The **Storage Domain Size Over Time** report shows a line graph contrasting the total available and total used space for a single storage domain over time for a given period. The following parameters must be provided to run this report:

Table 10.7. Storage Domain Size Over Time Parameters

Parameter	Description
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month. The list of options for the <i>Storage Domain Name</i> parameter includes only storage domains that were attached during the specified period.
Data Center	The options list for the <i>Storage Domain Name</i> parameter shows only storage domains in this selected data center.
Storage Type	The options list for the <i>Storage Domain Name</i> parameter shows only storage domains of this selected type.
Storage Domain Name	The report refers to the storage domain selected. A report is only for a single storage domain and the user must select a storage domain. The list

Parameter	Description
	of options shows only storage domains that were attached to the data center during the selected period.

10.4.3.3. Virtual Machines Inventory

The **Virtual Machines Inventory** report shows a list of all virtual machines in the selected data center and cluster. The following parameters must be provided to run this report:

Table 10.8. Virtual Machines Inventory Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all virtual machines in the selected data center.
VM Type	The report includes only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the report includes all virtual machine types.
Only Active Virtual Machines?	If Yes is selected, the report includes only active virtual machines, i.e. ones not deleted from the system. If No is selected, the report includes both active and deleted virtual machines.

10.4.4. Service Level Reports

Service level reports can be used to track system performance and availability against service level benchmarks. Service level agreement (SLA) thresholds are defined as report parameters. The following service level reports are available.

10.4.4.1. Cluster Host Uptime

The **Cluster Host Uptime** report shows the weighted average uptime of hosts within a cluster for a given period of time. This report also provides a table listing the total planned (maintenance) and unplanned down time for each host. The following parameters must be provided to run this report:

Table 10.9. Cluster Host Uptime Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in

Parameter	Description
	the selected data center. If All is selected, the report includes all hosts in the selected data center.
Host Type	The report includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the report includes all host types.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.4.2. Cluster Quality of Service - Hosts

The **Cluster Quality of Services - Hosts** report shows the amount of time hosts sustain load above a specified threshold for a given time period. Load is defined in terms of CPU usage percent and memory usage percent. The following parameters must be provided to run this report:

Table 10.10. Cluster Quality of Service - Hosts Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all hosts in the selected data center.
Host Type	The report includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the report includes all host types.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

Parameter	Description
CPU Threshold	The report measures the quality of service as the amount of time hosts sustain load above a given threshold. The <i>CPU Threshold</i> defines a load threshold as a percentage of total CPU usage on the host. The load is measured by one-minute samples, averaged over an hour. The report therefore shows sustained load, not short term peaks. A <i>CPU Threshold</i> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.
MEM Threshold	The report measures the quality of service as the amount of time hosts sustain load above a given threshold. The <i>MEM Threshold</i> defines a load threshold as a percentage of total memory usage on the host. The load is measured by one-minute samples, averaged over an hour. The report therefore shows sustained load, not short term peaks. A <i>MEM Threshold</i> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.

10.4.4.3. Cluster Quality of Service - Virtual Machines

The **Cluster Quality of Service - Virtual Machines** report shows the amount of time virtual machines sustain load above a specified threshold for a given time period. Load is defined in terms of CPU usage percent and memory usage percent. The following parameters must be provided to run this report:

Table 10.11. Cluster Quality of Service - Virtual Machines Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all virtual machines in the selected data center.
VM Type	The report includes only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the report includes all virtual machine types.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period

Parameter	Description
	range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.
CPU Threshold	The report measures quality of service as the amount of time virtual machines sustain load above a given threshold. The <i>CPU Threshold</i> defines a load threshold as a percentage of total CPU usage on the virtual machine. The load is measured by one-minute samples, averaged over an hour. The report therefore shows sustained load, not short term peaks. A <i>CPU Threshold</i> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.
MEM Threshold	The reports measures quality of service as the amount of time virtual machines sustain load above a given threshold. The <i>MEM Threshold</i> defines a load threshold as a percentage of total memory usage on the virtual machine. The load is measured by one-minute samples, averaged over an hour. The report therefore shows sustained load, not short term peaks. A <i>MEM Threshold</i> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.

10.4.4.4. Single Host Uptime

The **Single Host Uptime** report shows the total proportion of uptime, planned downtime and unplanned downtime for a single host. The following parameters must be provided to run this report:

Table 10.12. Single Host Uptime Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The list of options for the <i>Host Name</i> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>Host Name</i> parameter includes all hosts in the selected data center.
Host Type	The list of options for the <i>Host Name</i> parameter includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the list of options for the <i>Host Name</i> parameter includes all host types.

Parameter	Description
Host Name	The report refers to the host selected. A report is only for a single host and a user must select a host.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.4.5. Top 10 Downtime Hosts

The **Top 10 Downtime Hosts** report shows the total proportion of uptime, planned downtime and unplanned downtime for the 10 hosts with the greatest amount of downtime. The following parameters must be provided to run this report:

Table 10.13. Top 10 Downtime Hosts Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list contains only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all hosts in the selected data center.
Host Type	The report includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the report includes all host types.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.4.6. High Availability Virtual Servers Uptime

The **High Availability Virtual Servers Uptime** report shows the weighted average uptime of high availability virtual servers within a cluster for a given period of time. The report also

provides a table listing the total uptime and unplanned down time for each virtual server. The following parameters must be provided to run this report:

Table 10.14. High Availability Virtual Servers Uptime Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only virtual servers in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all virtual servers in the selected data center.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5. Trend Reports

Trend reports highlight trends in system utilization and availability. This is useful for planning purchase requests for capacity expansion and identifying problems and bottlenecks in the system. The following trend reports are available.

10.4.5.1. Five Least Utilized Hosts (Over Time)

The **Five Least Utilized Hosts (Over Time)** report shows the weighted average daily peak load, in terms of CPU and memory usage, for the five hosts with the lowest load factor for a given period of time. The following parameters must be provided to run this report:

Table 10.15. Five Least Utilized Hosts (Over Time) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all hosts in the selected data center.
Host Type	The report includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the report includes all host types.

Parameter	Description
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.2. Five Least Utilized Virtual Machines (Over Time)

The **Five Least Utilized Virtual Machines (Over Time)** report shows the weighted average daily peak load, in terms of CPU and memory usage, for the five virtual machines with the lowest load factor for a given period of time. The following parameters must be provided to run this report:

Table 10.16. Five Least Utilized Virtual Machines (Over Time) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all virtual machines in the selected data center.
VM Type	The report includes only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the report includes all virtual machine types.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.3. Five Most Utilized Hosts (Over Time)

The **Five Most Utilized Hosts (Over Time)** report shows the weighted average daily peak load, in terms of CPU and memory usage, for the five hosts with the highest load factor for a given period of time. The following parameters must be provided to run this report:

Table 10.17. Five Most Utilized Hosts (Over Time) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The report includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all hosts in the selected data center.
Host Type	The report includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the report includes all host types.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.4. Five Most Utilized Virtual Machines (Over Time)

The **Five Most Utilized Virtual Machines (Over Time)** report shows the weighted average daily peak load, in terms of CPU and memory usage, for the five virtual machines with the highest load factor for a given period of time. The following parameters must be provided to run this report:

Table 10.18. Five Most Utilized Virtual Machines (Over Time) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers which contain clusters.
Cluster	The report includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the report includes all virtual machines in the selected data center.
VM Type	The report includes only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the report includes all virtual machine types.

Parameter	Description
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.5. Multiple Hosts Resource Usage (Over Time)

The **Multiple Hosts Resource Usage (Over Time)** report shows the daily peak load, in terms of CPU and memory usage, for up to five selected hosts over a given period of time. The following parameters must be provided to run this report:

Table 10.19. Multiple Hosts Resource Usage (Over Time) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The list of options for the <i>Host List</i> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>Host List</i> parameter includes all hosts in the selected data center.
Host Type	The list of options for the <i>Host List</i> parameter includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the list of options for the <i>Host List</i> parameter includes all host types.
Host List	The report includes all hosts selected in the host list. Select any number of hosts up to a maximum of five.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.6. Multiple Virtual Machines Resource Usage (Over Time)

The **Multiple Virtual Machines Resource Usage (Over Time)** report shows the daily peak load, in terms of CPU and memory usage, for up to five selected virtual machines over a given period of time. The following parameters must be provided to run this report:

Table 10.20. Multiple Virtual Machines Resource Usage (Over Time) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The list of options for the <i>VM List</i> parameter include only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>VM List</i> parameter includes all virtual machines in the selected data center.
VM Type	The list of options for the <i>VM List</i> parameter includes only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the list of options for the <i>VM List</i> parameter includes all virtual machine types.
VM List	The report includes all virtual machines selected in the virtual machine list. Select any number of virtual machines up to a maximum of five.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.7. Single Host Resource Usage (Days of Week)

The **Single Host Resource Usage (Days of Week)** report shows various resource utilization metrics for a single host over a given period of time and broken down by day of the week. The metrics include CPU usage, memory usage, number of active virtual machines and network usage. The following parameters must be provided to run this report:

Table 10.21. Single Host Resource Usage (Days of Week) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.

Parameter	Description
Cluster	The list of options for the <i>Host Name</i> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>Host Name</i> parameter includes all hosts in the selected data center.
Host Type	The list of options for the <i>Host Name</i> parameter includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If All is selected, the list of options for the <i>Host Name</i> parameter includes all host types.
Host Name	The report refers to the host selected. A report is only for a single host and the user must select a host.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.8. Single Host Resource Usage (Hour of Day)

The **Single Host Resource Usage (Hour of Day)** report shows a variety of resource utilization metrics for a single host over a given period of time, broken down by hour of the day (0-23). The metrics include CPU usage, memory usage, number of active virtual machines and network usage. The following parameters must be provided to run this report:

Table 10.22. Single Host Resource Usage (Hour of Day) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The list of options for the <i>Host Name</i> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>Host Name</i> parameter includes all hosts in the selected data center.
Host Type	Only hosts of the selected type will be included in the list of options for the <i>Host Name</i> parameter. The options list shows only host types present in the selected data center and cluster. If All is selected, the list of options for the <i>Host Name</i> parameter includes all host types.

Parameter	Description
Host Name	The report refers to the host selected. A report is only for a single host and the user must select a host.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.9. Single Virtual Machine Resources (Days of Week)

The **Single Virtual Machine Resources (Days of Week)** report shows a variety of resource utilization metrics for a single virtual machine over a given period of time, broken down by day of the week. The metrics include CPU usage, memory usage, disk usage and network usage. The following parameters must be provided to run this report:

Table 10.23. Single Virtual Machine Resources (Days of Week) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The list of options for the <i>VM Name</i> parameter includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>VM Name</i> parameter includes all virtual machines in the selected data center.
VM Type	The list of options for the <i>VM Name</i> parameter includes only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the list of options for the <i>VM Name</i> parameter includes all virtual machine types.
VM Name	The report refers to the virtual machine selected. A report is only for a single virtual machine and the user must select a virtual machine.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period

Parameter	Description
	range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.10. Single Virtual Machine Resources (Hour of Day)

The **Single Virtual Machine Resources (Hour of Day)** report shows a variety of resource utilization metrics for a single virtual machine over a given period of time, broken down by hour of the day (0-23). The metrics include CPU usage, memory usage, disk usage and network usage. The following parameters must be provided to run this report:

Table 10.24. Single Virtual Machine Resources (Hour of Day) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers which contain clusters.
Cluster	The list of options for the <i>VM Name</i> parameter includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>VM Name</i> parameter includes all virtual machines in the selected data center.
VM Type	The list of options for the <i>VM Name</i> parameter includes only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the list of options for the <i>VM Name</i> parameter includes all virtual machine types.
VM Name	The report refers to the virtual machine selected. A report is only for a single virtual machine and the user must select a virtual machine.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.4.5.11. Single Virtual Machine Resources (Over Time)

The **Single Virtual Machine Resources (Over Time)** report shows a variety of resource utilization metrics for a single virtual machine over a given period of time. The metrics include CPU usage, memory usage, disk usage and network usage. The following parameters must be provided to run this report:

Table 10.25. Single Virtual Machine Resources (Over Time) Parameters

Parameter	Description
Data Center	The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
Cluster	The list of options for the <i>VM Name</i> parameter includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If All is selected, the list of options for the <i>VM Name</i> parameter includes all virtual machines in the selected data center.
VM Type	The list of options for the <i>VM Name</i> parameter lists only virtual machines of the selected type. The options list shows only virtual machine types present in the selected data center and cluster. If All is selected, the list of options for the <i>VM Name</i> parameter includes all virtual machine types.
VM Name	The report refers to the virtual machine selected. A report is only for a single virtual machine and the user must select a virtual machine.
Period Range	The report is for the period range selected. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The report covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.5. Dashboards

Dashboards are similar to reports. They include an enhanced display combining data and graphical indicators. Dashboards provide system executive summary views that cannot be edited, printed or exported like reports. The following dashboards are available.

Embedded reports are the building blocks used to construct dashboards. Embedded reports are reports that have been simplified for presentation in dashboards. The **Embedded Reports** folder, found under **Dashboards** in the navigation tree, contains the embedded reports. Users are not able to access embedded reports, nor to run them directly.

Dashboards contain embedded links to relevant reports. Click on each component on the dashboard to access them. The link takes you to a report with pre-filled parameters based on the context of the dashboard component you are viewing. Users can print the report or drill down into more detail.

10.5.1. Data Center Inventory Dashboard

The **Data Center Inventory Dashboard** provides an executive summary of the inventory of a data center over a given period of time. The dashboard includes average disk use, number of active

virtual machines and a breakdown of host operating systems. The following parameters must be provided to view this dashboard:

Table 10.26. Data Center Inventory Dashboard Parameters

Parameter	Description
Data Center	The report refers to the selected data center. The list of options shows only data centers containing either hosts, storage domains or virtual machines. The list of options for the <i>Cluster</i> parameter includes only clusters in the selected data center.
Cluster	The report refers to the cluster selected. If All is selected, the report refers to the entire data center.
Period Range	The dashboard shows data for the period range selected. Monthly dashboards cover a single month. Quarterly dashboards cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The dashboard covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.5.2. Data Center Trends Dashboard

The **Data Center Trends Dashboard** provides an executive summary of the trends in a data center over a given period of time. The dashboard includes graphs of CPU and memory usage over time for the most highly utilized hosts and virtual machines in the data center. The following parameters must be provided to view this dashboard:

Table 10.27. Data Center Trends Dashboard Parameters

Parameter	Description
Data Center	The report refers to the selected data center. The options list shows only data centers containing either hosts or virtual machines. The options list for the <i>Cluster</i> parameter includes only clusters in the selected data center.
Cluster	The report refers to the cluster selected. If All is selected, the report refers to the entire data center.
Host Type	The list of most utilized hosts in the dashboard includes only hosts of the selected type. If All is selected, the dashboard includes all host types.
VM Type	The list of most utilized virtual machines in the dashboard includes only virtual machines of the selected type. If All is selected, the dashboard includes all virtual machine types.
Period Range	The dashboard shows data for the period range selected. Monthly dashboards cover a single

Parameter	Description
	month. Quarterly dashboards cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The dashboard covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.5.3. Data Center Uptime Dashboard

The **Data Center Uptime Dashboard** provides an executive summary of the service level and uptime for a data center over a given period of time. The dashboard includes details on total uptime for each cluster in the data center for the period. The following parameters must be provided to view this dashboard:

Table 10.28. Data Center Uptime Dashboard Parameters

Parameter	Description
Data Center	The report refers to the data center selected. The options list shows only data centers containing either hosts or virtual machines.
Period Range	The dashboard shows data for the period range selected. Monthly dashboards cover a single month. Quarterly dashboards cover a three-month quarter, beginning on the month specified in the <i>Dates</i> parameter.
Dates	The dashboard covers the selected period range, beginning on this date. For a Monthly period range, the selected month is used. For a Quarterly period range, the quarter is determined as beginning on the selected month.

10.5.4. System Overview Dashboard

The **System Overview Dashboard** provides an executive summary of the hosts in a data center over a given period of time. The dashboard includes:

- A quality of service (QoS) view for each cluster, which shows the proportion of period where CPU and memory exceeded thresholds on the hosts in the cluster;
- A break down of host operating systems; and
- A summary of average host resource utilization over the period.

The following parameters must be provided to view this dashboard:

Table 10.29. System Overview Dashboard Parameters

Parameter	Description
Data Center	The report refers to the data center selected. The list of options shows only data centers containing either hosts or virtual machines. The options

Parameter	Description
	list for the <i>Cluster</i> parameter includes only clusters in the selected data center.
Cluster	The report refers to the cluster selected. If All is selected, the report pertain to the entire data center.
CPU Threshold	The report refers to the amount of Host CPU utilized.
Memory Threshold	The report refers to the amount of Memory utilized in the host by virtual machines.
Period Range	The report refers to the range of dates selected.
Dates	The report refers to the dates displayed. This parameter changes automatically when the Range parameter is set.

10.6. Ad Hoc Reports

Red Hat Enterprise Virtualization Reports provides you with a tool to create customized ad hoc reports. This tool is a component of JasperServer. To create an ad hoc report as an administrator, click **Create** in the top right-hand corner of the reports interface and select **Ad Hoc Report** as shown in [Figure 10.8, “Create Ad Hoc Report - administrator’s view”](#).



Figure 10.8. Create Ad Hoc Report - administrator's view

Clicking on **Ad Hoc Report** causes the Ad Hoc Reports dialog box to appear.

The **Working with the Ad Hoc Editor** section of the online help explains the ad hoc report interface in detail. See [Section 10.1.2, “Online Help”](#) for more information.

History Database Reports

Red Hat Enterprise Virtualization includes a comprehensive management history database, which any reporting application utilizes to generate a range of reports at data center, cluster and host levels. This chapter provides information to enable you to set up queries against the history database and generate reports.

11.1. Overview

Red Hat Enterprise Virtualization Manager uses **PostgreSQL 8.4.7** as a database platform to store information about the state of the virtualization environment, its configuration and performance. At install time, Red Hat Enterprise Virtualization Manager creates a PostgreSQL database called **rhev**. Installing the **dwh** package creates a second database called **rhev_history**, which contains historical configuration information and statistical metrics collected every minute over time from the **rhev** operational database. Tracking the changes to the database provides information on the objects in the database, enabling the user to analyze activity, enhance performance, and resolve difficulties.



Warning — Do Not Stop RHEVM History Service

The replication of data in the **rhev_history** database is performed by the **RHEVM ETL Service** (Red Hat Enterprise Virtualization Manager Extract Transform Load Service). The service is based on Talend Open Studio, a data integration tool. This service is configured to start automatically during the data warehouse package setup. It is a Java program responsible for extracting data from the **rhev** database, transforming the data to the history database standard and loading it to the **rhev_history** database.

The **RHEVM ETL** service must not be stopped.

Since the **rhev_history** database schema changes over time, the database includes a set of database views to provide a supported, versioned API with a consistent structure. A view is a virtual table composed of the result set of a database query. The database stores the definition of a view as a **SELECT** statement. The result of the **SELECT** statement populates the virtual table that the view returns. A user references the view name in **PL/PGSQL** statements the same way a table is referenced.

11.1.1. Tracking Configuration History

The ETL service tracks three types of changes:

- A new entity is added to the **rhev** database - the ETL Service replicates the change to the **rhev_history** database as a new entry.
- An existing entity is updated - the ETL Service replicates the change to the **rhev_history** database as a new entry.
- An entity is removed from the **rhev** database - A new entry in the **rhev_history** database flags the corresponding entity as removed. Removed entities are only flagged as removed. To maintain correctness of historical reports and representations, they are not physically removed.

The configuration tables in the **rhev_history** database differ from the corresponding tables in the **rhev** database in several ways. The most apparent difference is they contain fewer **configuration**

columns. This is because certain configuration items are less interesting to report than others and are not kept due to database size considerations. Also, columns from a few tables in the **rhevm** database appear in a single table in **rhevm_history** and have different column names to make viewing data more convenient and comprehensible. All configuration tables contain:

- a **history_id** to indicate the configuration version of the entity;
- a **create_date** field to indicate when the entity was added to the system;
- an **update_date** field to indicate when the entity was changed; and
- a **delete_date** field to indicate the date the entity was removed from the system.

11.1.2. Recording Statistical History

The ETL service collects data into the statistical tables every minute. Data is stored for every minute of the past 24 hours. Minute-by-minute data more than two hours old is aggregated into hourly data and stored for two months. Hourly data more than two days old is aggregated into daily data and stored for five years.

Hourly data and daily data can be found in the hourly and daily tables.

Each statistical datum is kept in its respective aggregation level table: samples, hourly, and daily history. All history tables also contain a **history_id** column to uniquely identify rows. Tables reference the configuration version of a host in order to enable reports on statistics of an entity in relation to its past configuration.

11.1.3. Tracking Tag History

The ETL Service collects tag information as displayed in the Administration Portal every minute and stores this data in the tags historical tables. The ETL Service tracks five types of changes:

- A tag is created in the Administration Portal - the ETL Service copies the tag details, position in the tag tree and relation to other objects in the tag tree.
- A entity is attached to the tag tree in the Administration Portal - the ETL Service replicates the addition to the **rhevm_history** database as a new entry.
- A tag is updated - the ETL Service replicates the change of tag details to the **rhevm_history** database as a new entry.
- An entity or tag branch is removed from the Administration Portal - the **rhevm_history** database flags the corresponding tag and relations as removed in new entries. Removed tags and relations are only flagged as removed or detached. In order to maintain correctness of historical reports and representations, they are not physically removed.
- A tag branch is moved - the corresponding tag and relations are updated as new entries. Moved tags and relations are only flagged as updated. In order to maintain correctness of historical reports and representations, they are not physically updated.

11.2. Connecting to the History Database

The **rhevm_history** database resides within the Red Hat Enterprise Virtualization Manager instance of PostgreSQL that the installer creates. To connect to the database, use a PostgreSQL compatible query or reporting tool with the credentials used in Red Hat Enterprise Virtualization Manager or any other PostgreSQL credentials available.

11.3. Example Reports

The following examples provide an introduction to reports produced from queries to the **rhev_history** database. The database gives users access to a rich data set and enables a variety of complex reporting scenarios. These examples illustrate only basic reporting requirements.

11.3.1. Resource Utilization on a Single Host

This example produces a resource utilization report for a single host. The resource utilization report provides CPU- and memory-usage percentage information from readings taken at one-minute intervals. This kind of report is useful for gaining insight into the load factor of an individual host over a short period of time. The report is defined by the following SQL query. Ensure the values provided for the **host_name** and **history_datetime** components of the **where** clause are substituted with the appropriate values for your environment and that the latest configuration is in use.

Example 11.1. Report query for resource utilization on a single host

```
select history_datetime as DateTime, cpu_usage_percent as CPU, memory_usage_percent as
Memory
  from v3_0_host_configuration_view, v3_0_host_samples_history_view
 where v3_0_host_configuration_view.host_id = v3_0_host_samples_history_view.host_id
 and host_name = 'example.labname.abc.company.com'
 and v3_0_host_configuration_view.history_id in (select max(a.history_id)
        from v3_0_host_configuration_view as a
        where v3_0_host_configuration_view.host_id = a.host_id)
 and history_datetime >= '2011-07-01 18:45'
 and history_datetime <= '2011-07-31 21:45'
```

This query returns a table of data with one row per minute:

Table 11.1. Resource Utilization for a Single Host Example Data

DateTime	CPU	Memory
2010-07-01 18:45	42	0
2010-07-01 18:46	42	0
2010-07-01 18:47	42	1
2010-07-01 18:48	33	0
2010-07-01 18:49	33	0
2010-07-01 18:50	25	1

Compose the data into a graph or chart using third party data analysis and visualization tools such as **OpenOffice.org Calc** and **Microsoft Excel**. For this example, a line graph showing the utilization for a single host over time is a useful visualization. *Figure 11.1, "Single host utilization line graph"* was produced using the **Chart Wizard** tool in **OpenOffice.org Calc**.

Resource Utilization for example.labname.abc.company.com

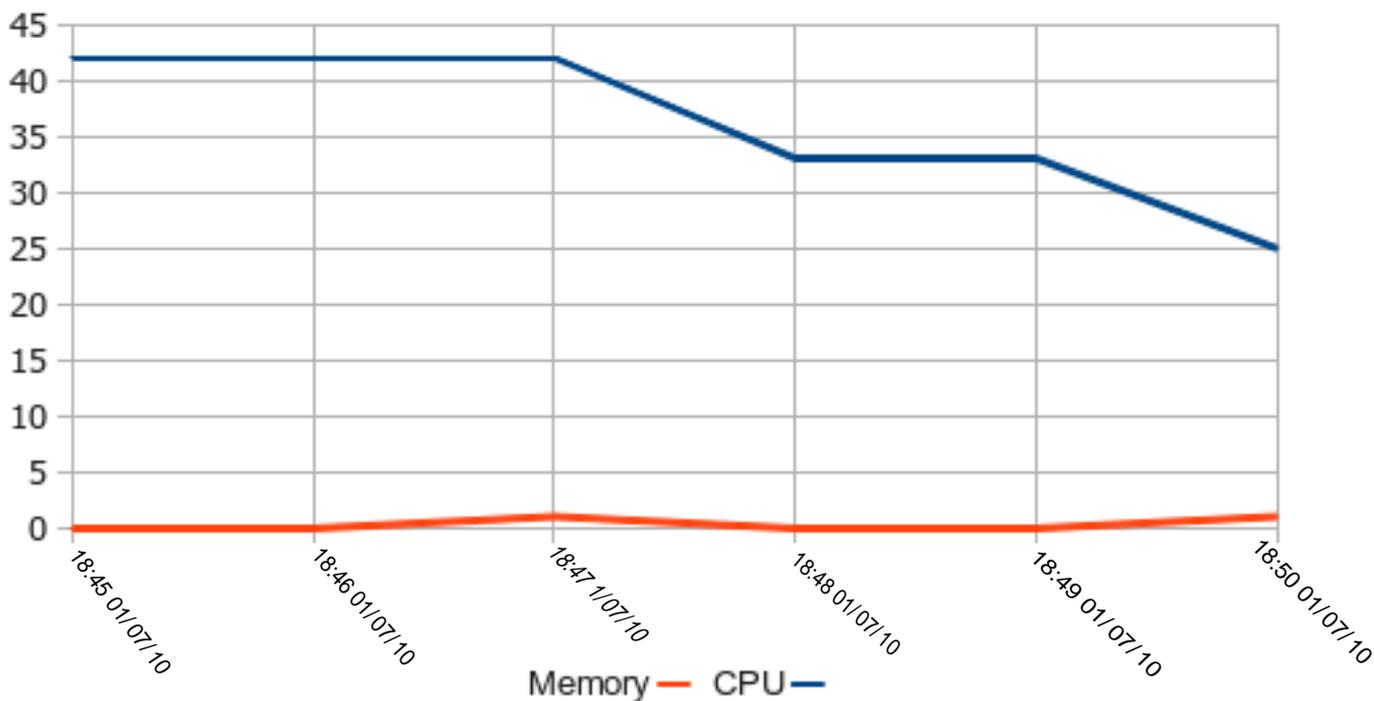


Figure 11.1. Single host utilization line graph

11.3.2. Resource Utilization Across All Hosts

This example produces an aggregated resource utilization report across all hosts in the Red Hat Enterprise Virtualization Manager environment. Aggregated usage percentages for CPU and memory are shown with an hourly temporal resolution. This kind of report reveals utilization trends for the entire environment over a long period of time and is useful for capacity planning purposes. The following SQL query defines the report. Ensure the values provided for the **history_datetime** components of the **where** clause are substituted with appropriate values for your environment.

Example 11.2. Report query for resource utilization across all hosts

```
select extract(hour from history_datetime) as Hour, avg(cpu_usage_percent) as
CPU, avg(memory_usage_percent) as Memory
from v3_0_host_hourly_history_view
where history_datetime >= '2011-07-01' and history_datetime < '2011-07-31'
group by extract(hour from history_datetime)
order by extract(hour from history_datetime)
```

This query returns a table of data with one row per hour:

Table 11.2. Resource utilization across all hosts example data

Hour	CPU	Memory
0	39	40
1	38	38

Hour	CPU	Memory
2	37	32
3	35	45
4	35	37
5	36	37

Compose the data into a graph or chart using third party data analysis and visualization tools such as **OpenOffice.org Calc** and **Microsoft Excel**. For this example, a line graph showing the total system utilization over time is a useful visualization. *Figure 11.2, "Total system utilization line graph"* was produced using the **Chart Wizard** tool in **OpenOffice.org Calc**.

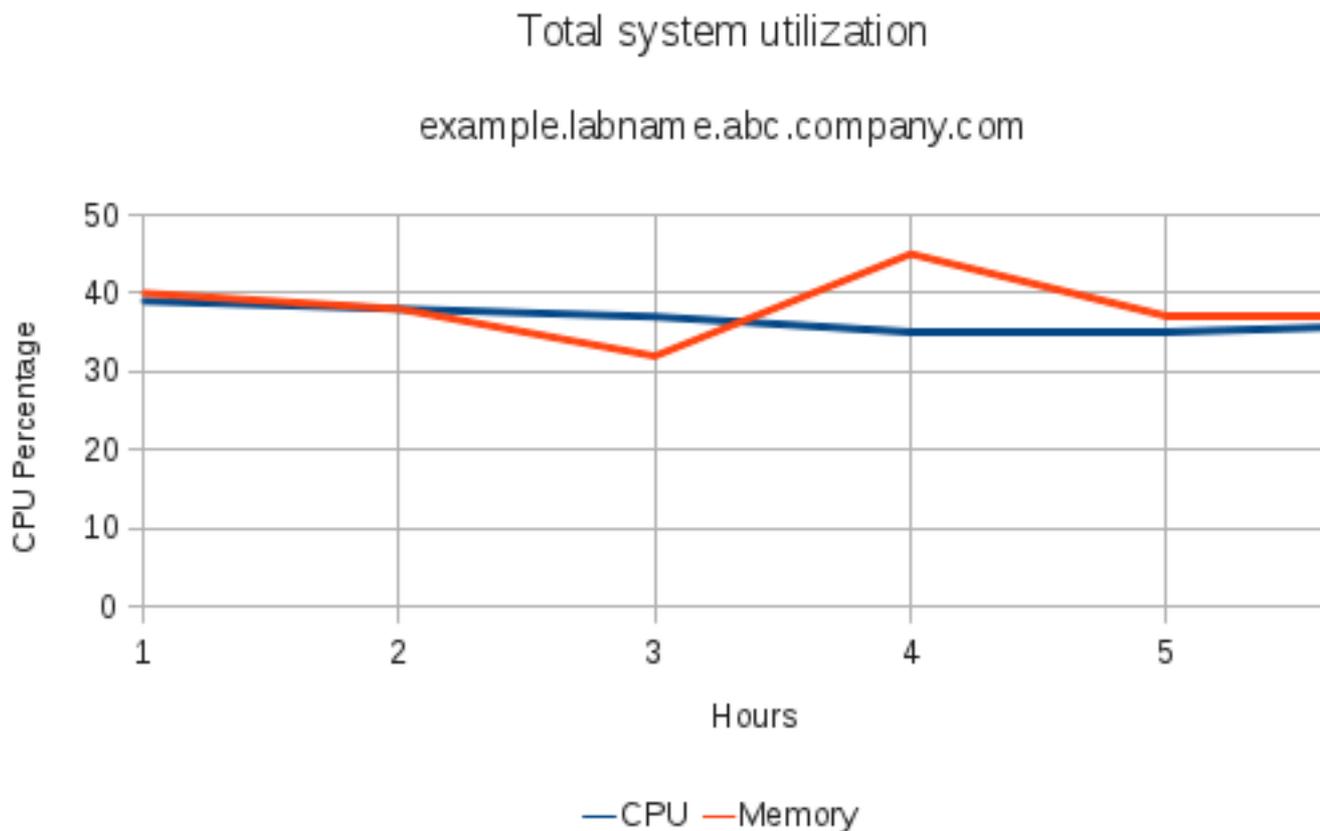


Figure 11.2. Total system utilization line graph

11.3.3. Tag Filter of Latest VM Configuration

This example filters the latest virtual machine configuration list using the history tag tables. This kind of report demonstrates usage of the tags tree built in the Red Hat Enterprise Virtualization Manager to filter lists. The following SQL query defines this report. This query uses a predefined function that receives tag history IDs and returns the tag path with latest names of the tags in the Administration Portal. Ensure the values provided for the function result components of the **where** clause are substituted with appropriate values for your environment.

Example 11.3.

```
SELECT vm_name
FROM v3_0_latest_vm_configuration_view
```

```

inner join v3_0_latest_tag_relations_history_view on
(v3_0_latest_vm_configuration_view.vm_id =
v3_0_latest_tag_relations_history_view.entity_id)
inner join v3_0_latest_tag_details_view on (v3_0_latest_tag_details_view.tag_id =
v3_0_latest_tag_relations_history_view.parent_id)
WHERE getpathinnames(v3_0_latest_tag_details_view.history_id) like '/root/tlv%'

```

This query returns a table of data with all virtual machine names that are attached to this tag:

Table 11.3. Tag Filtering of Latest VM Configuration

vm_name
RHEL6-Pool-67
RHEL6-Pool-5
RHEL6-Pool-6
RHEL6-23

11.3.4. List Current Virtual Machines' Names, Types, and Operating Systems

This example produces a list of all current virtual machines names, types and operating systems in the Red Hat Enterprise Virtualization Manager environment. This kind of report demonstrates the usage of the ENUM table. The following SQL query defines this report:

Example 11.4.

```

SELECT vm_name, vm_type_value.value as vm_type, os_value.value as operating_system
FROM v3_0_latest_vm_configuration_view
inner join v3_0_enum_translator_view as vm_type_value on (vm_type_value.enum_type
= 'VM_TYPE' and v3_0_latest_vm_configuration_view.vm_type = vm_type_value.enum_key)
inner join v3_0_enum_translator_view as os_value on (os_value.enum_type = 'OS_TYPE' and
v3_0_latest_vm_configuration_view.operating_system = os_value.enum_key)

```

This query returns a table of virtual machines with OS and VM Type data:

Table 11.4. Current Virtual Machines' Names, Types, and Operating Systems

vm_name	vm_type	operating_system
RHEL6-Pool-2	Desktop	RHEL 6 x64
RHEL6-Pool-1	Desktop	RHEL 6 x64
RHEL6-Pool-3	Desktop	RHEL 6 x64
RHEL6-Pool-4	Desktop	RHEL 6 x64
RHEL6-Pool-5	Desktop	RHEL 6 x64



Reporting Views

The *Red Hat Enterprise Virtualization Technical Reference Guide* provides a detailed reference that describes all the configuration and history views available for reporting.

Part V. Managing Advanced Functionality

Live Migration

This chapter details the live migration functionality of Red Hat Enterprise Virtualization and the configuration required to support it. Information on triggering live migration of virtual machines is also provided.

12.1. What is Live Migration?

Live migration provides the ability to move a running virtual machine between physical hosts with no interruption to service.

Live migration is transparent to the end user: the virtual machine remains powered on and user applications continue to run while the virtual machine is relocated to a new physical host.

12.2. Live Migration Prerequisites

Live migration is used to seamlessly move virtual machines to support a number of common maintenance tasks. It is important to ensure that your Red Hat Enterprise Virtualization environment is correctly configured to support live migration well in advance of needing to use it.

At a minimum for successful live migration of virtual machines to be possible:

- The source and destination host must both be members of the same cluster, ensuring CPU compatibility between them.
- The source and destination host must have a status of **Up**.
- The source and destination host must have access to the same virtual networks and VLANs.
- The source and destination host must have access to the data storage domain on which the virtual machine resides.
- There must be enough CPU capacity on the destination host to support the virtual machine's requirements.
- There must be enough RAM on the destination host that is not in use to support the virtual machine's requirements.

In addition, for best performance, it is recommended that the storage and management networks should be split to avoid network saturation. Virtual machine migration involves transferring large amounts of data between hosts.

While each live migration event is limited to a maximum transfer speed of 30 MBps where this traffic exists on the management network it has the potential to saturate it. Having a separate logical network for storage reduces the likelihood of this issue being encountered.

12.3. Automatic Virtual Machine Migration

Red Hat Enterprise Virtualization Manager is able to automatically live migrate virtual machines. The manager automatically initiates live migration of virtual machines when a host is being moved into maintenance mode or becomes non-responsive. Where this occurs live migration of all virtual machines running on the host is attempted. The destination host for each virtual machine will be assessed as that virtual machine is migrated, so as to spread the additional load across the cluster.

The manager will also automatically initiate live migration of virtual machines in order to maintain load balancing or power saving levels in the cluster to be in line with cluster policy, if it has been defined.

While by default no cluster policy is defined it is recommended that administrators choose to specify the cluster policy which best suits the needs of their environment. Administrators can also disable automatic, or even manual, live migration of specific virtual machines where required.

- To move a host to maintenance mode, see [Section 12.3.1, “Moving a Host to Maintenance Mode”](#).
- To configure cluster policy, see [Section 12.3.2, “Cluster Policy”](#).
- To prevent automatic migration of a specific virtual machine, see [Section 12.3.3, “Preventing Automatic Migration of a Virtual Machine”](#).

12.3.1. Moving a Host to Maintenance Mode

Many common maintenance tasks, including network configuration and deployment of software updates, require that virtualization hosts be placed into maintenance mode. When a host is placed into maintenance mode the Red Hat Enterprise Virtualization Manager attempts to migrate all running virtual machines to alternative hosts. The normal prerequisites for live migration apply.

In particular there must be at least one active host in the cluster with capacity to run the virtual machine(s) being migrated as a result of the host being placed into maintenance mode. To move a host to maintenance mode follow these steps:

1. Ensure the **Tree** tab is selected in the left hand pane of the Administration Portal.
2. Expand the tree view using the **+** buttons until the host to be moved into maintenance mode is visible.
3. Select the host to be moved into maintenance mode. Click the **Maintenance** button.

All running virtual machines will be migrated to alternative hosts. Once this has occurred the host is moved into maintenance mode. The **Status** field of the host will change to **Preparing for Maintenance**, and finally **Maintenance** when the operation completes successfully.

Result:

All running virtual machines have been migrated to alternatives hosts. The host has been placed into maintenance mode.

12.3.2. Cluster Policy

Load balancing and power saving are key criteria in ensuring optimal performance of a virtualized environment. Red Hat Enterprise Virtualization platform allows you to define cluster policies to specify the usage and distribution of virtual machines between the available hosts. At any given time in a virtualized environment, virtual machines are starting, stopping or resuming.

In its simplest form, it is the rules in the policy engine that determine the selection of the specific host on which a virtual machine runs. The policy engine decides which server will host the next virtual machine based on whether load balancing criteria have been defined. Additionally to maintain load balancing as defined in the cluster policy the Red Hat Enterprise Virtualization Manager will use live migration to move virtual machines around the cluster as required. Defining the load-balancing or power saving modes for hosts in the cluster is highly recommended. You can choose to set the policy as either even distribution or power saving, but not both.

1. Ensure that you have the **Tree** tab open to the left of the Administration Portal screen. Expand the objects in the tree, by clicking the **+** icon next to them, until the cluster that you are modifying is visible.

2. Select the cluster from the list and click the **Edit Policy** button. The **Edit Policy** button is located on the details pane. Note that the current cluster policy, if one is defined, is also visible in the details pane.
3. Select the policy to use, the available options are:

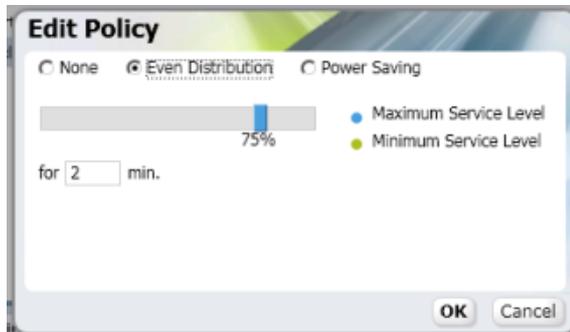


Figure 12.1. Edit Cluster Policy

- **None** — to have no load or power saving between hosts. This is the default mode. When starting a new virtual machine the host to use is selected based on CPU utilization. The active host in the cluster with the lowest CPU utilization will be chosen. While this is the same mechanism as used to select the initial host for a virtual machine under the **Even Distribution** policy no automatic live migration is performed once the virtual machine has started to maintain the distribution. Once the virtual machine has started it will continue to run on this host until it is stopped, the host goes into maintenance mode, or the host becomes non-responsive.
- **Even Distribution** — evenly distributes the processing load across all hosts in the cluster. The host's CPU load is measured and used to apply the policy. Use the blue slider to specify the **Maximum Service Level** a host is permitted to have. For example, a host that has reached the maximum service level defined will not have further virtual machines started on it. You must also specify the time interval in minutes that a host is permitted to run at the maximum service level before virtual machines are migrated off it.

Once the host has run at the maximum service level for the specified amount of time the manager will begin migrating virtual machines to other hosts in the cluster until the host's CPU load returns to a level below the maximum service threshold. When migrating the virtual machines the manager chooses the host in the cluster with the lowest CPU usage as the destination.

- **Power Saving** — distributes the load in a way that consolidates virtual machines on a subset of available hosts. This enables surplus hosts that are not in use to be powered down, saving power. Use the green slider to specify the **Minimum Service Level** a host is permitted to have. You must also specify the time interval in minutes that a host is permitted to run below the minimum service level before remaining virtual machines are migrated to other hosts in the cluster — as long as the maximum service level set also permits this.

Under-utilized hosts are then able to be turned off either manually or using the REST API to save power.

4. Click **OK** to save the policy selection for the cluster.

Result:

The cluster policy for the selected cluster has been updated.

12.3.3. Preventing Automatic Migration of a Virtual Machine

You may find that for a specific virtual machine you need to override the policies which automatically trigger migration. Red Hat Enterprise Virtualization Manager allows you to disable automatic migration of virtual machines. You are also able to disable manual migration of virtual machines by setting the virtual machine to run only on a specific host.

1. Ensure that you have the **Tree** tab open to the left of the Administration Portal Screen. Expand the objects in tree, by clicking the **+** icon next to them, until the cluster on which the virtual machine you are modifying runs is visible.
2. Expand the entry for the cluster and select the **VMs** object. The list of the cluster's virtual machines will appear.
3. Select the virtual machine from the list. Click the **Edit** button. Depending on the type of virtual machine selected the **Edit Desktop Virtual Machine** or **Edit Server Virtual Machine** dialog box will appear.
4. Click the **Host** tab on the dialog box. A number of options used to determine where the virtual machine is able to run appear.
5. Set the **Run On** option to either **Any Host in Cluster** or **Specific**.

Select **Any Host in Cluster** to allow the virtual machine to run on any host in the cluster. Subsequent options in the dialog will still allow you to ensure the virtual machine is not automatically migrated once it has started on a host. You will not however be able to disable manual migration.

Select **Specific** to set the virtual machine to only run on a specific host in the cluster. You must also select the host on which the virtual machine must run from the list. The list contains all active hosts in the cluster.



Warning — Mutually Exclusive with High Availability

Explicitly assigning a virtual machine to a specific host and disabling migration is mutually exclusive with high availability. Hosts that are assigned to a specific virtual machine and set not to migrate can not be made highly available.

6. To force the virtual machine to only run on the host selected when it starts, and to disable migration, you must select either:
 - **Run VM on the selected host (no migration allowed)** — No migration of the virtual machine is permitted. Automatic migration of the virtual machine will not be attempted and manual attempts to migrate the virtual machine will be denied. This option is only available when you have selected to run the virtual machine on a specific host.
 - **Allow VM migration only upon Administrator specific request (system will not trigger automatic migration of this VM)** — Only manual migration of the virtual machine is permitted. Automatic migration of the virtual machine will not be attempted. This option is available regardless of whether you have selected to run the virtual machine on a specific host or on any host in the cluster.

Result:

The migration settings for the virtual machine have been updated.

12.4. Manually Migrating Virtual Machines

A running virtual machine can be migrated to any host within its designated host cluster. This is especially useful if the load on a particular host is too high. Note that when bringing a server down for maintenance migration is triggered automatically, manual migration is not required. Migration of virtual machines does not cause any service interruption.

1. Click the **Virtual Machines** tab.
2. Select the entry for the virtual machine that you want to migrate.
3. With the entry for the virtual machine still selected, click the **Migrate** button; or
Right click the entry for the virtual machine and select **Migrate** in the menu that appears.
4. The Red Hat Enterprise Virtualization Manager is able to automatically select the host to migrate the virtual machine to or you are able to manually specify a host.

Choose **Select Host Automatically** to allow Red Hat Enterprise Virtualization Manager to select the destination host; or

Choose **Select Destination Host** to manually select a host. You must also select the destination host from the list provided. Only active hosts within the cluster are listed.



Note — Automatic Host Selection

Virtual Machines migrate within their designated host cluster. When the **Select Host Automatically** option is selected the system determines the host to which the virtual is migrated, according to the load balancing and power management rules set up in the cluster policy.

5. Click **OK** to commence migration and close the dialog box.

Result:

The virtual machine is migrated. Once migration has completed the **Host** column will update to display the host the virtual machine has been migrated to.

12.5. Setting Migration Priority

Red Hat Enterprise Virtualization Manager queues concurrent requests for migration of virtual machines from a given host. Every minute the load balancing process runs. Hosts already involved in a migration event are not included in the migration cycle until their migration event has completed. When there is a migration request in the queue and available hosts in the cluster to action it, a migration event is triggered in line with the load balancing policy for the cluster.

It is possible to influence the ordering of the migration queue, for example setting mission critical virtual machines to migrate before others. The Red Hat Enterprise Virtualization Manager allows you

to set the priority of each virtual machine to facilitate this. Virtual machines migrations will be ordered by priority, those virtual machines with the highest priority will be migrated first.

1. Click the **Virtual Machines** tab.
2. Select the entry for the virtual machine that you want to migrate.
3. With the entry for the virtual machine still selected, click the **Edit** button; or
Right click the entry for the virtual machine and click **Edit** in the menu that appears.
4. Click the **High Availability** tab.
5. Set the priority of the virtual machine by selecting **Low**, **Medium**, or **High** under **Priority for Run/Migrate Queue**.
6. Click **OK** to save the change to the virtual machine's priority.

Result:

The virtual machine's migration priority has been modified.

High Availability

The Red Hat Enterprise Virtualization Manager offers various high availability features which can be applied in a granular manner, from the level of a single virtual machine up to protection against multiple host failures. In addition, combining virtual machine high availability with out of band power management devices protects your virtual machines against storage failures.

This chapter describes Red Hat Enterprise Virtualization's high availability features, and how to configure your servers to run highly available virtual machines. In addition, this chapter provides information on configuring power management for your hosts.

13.1. What is High Availability?

High availability means that a virtual machine will be automatically restarted if its process is interrupted, for example if the virtual machine is terminated by methods other than powering off from within the guest or sending the shutdown command from the Red Hat Enterprise Virtualization Manager. When these events occur, the highly available virtual machine will be automatically restarted, either on its original host or another host in the cluster.

High availability is possible because the Red Hat Enterprise Virtualization Manager constantly monitors the physical servers, and automatically detects hardware failure. If host failure is detected, any virtual machine configured to be highly available is automatically restarted on another host in the cluster. In addition, all virtual machines are monitored, so if the virtual machine's operating system crashes, a signal is sent to automatically restart the virtual machine.

With high availability, interruption to service is minimal because virtual machines are restarted within seconds, and with no user intervention required. High availability keeps your resources balanced, as virtual machines are restarted on a host selected based on its current resource utilization, or based on any workload balancing or power saving policies that you configure. This ensures that there is sufficient capacity to restart virtual machines at all times.

13.1.1. Why Use High Availability?

High availability is recommended for virtual machines running critical workloads. It creates a secure and reliable environment where virtual machines are accessible during planned downtime or during unplanned downtime.

High availability can ensure that virtual machines are restarted in the following scenarios:

- When a host becomes non-operational due to hardware failure.
- When a host is put into maintenance mode for scheduled downtime.
- When a host becomes unavailable because it has lost communication with an external storage resource.
- When a virtual machine fails due to an operating system crash.

13.2. High Availability Considerations

A highly available host requires a power management device and its fencing parameters configured. In addition, for a virtual machine to be highly available when its host becomes non-operational, it needs to be started on another available host in the cluster. To enable the migration of highly available virtual machines:

- Power management must be configured for the hosts running the highly available virtual machines.
- The host on which the highly available virtual machines are running must be part of a cluster which has other available hosts.
- The destination host must be running.
- The source and destination host must have access to the data storage domain on which the virtual machine resides.
- The source and destination host must have access to the same virtual networks and VLANs.
- There must be enough virtual CPUs on the destination host that are not in use to support the virtual machine's requirements.
- There must be enough virtual RAM on the destination host that is not in use to support the virtual machine's requirements.

13.3. Host High Availability

All hosts on the Red Hat Enterprise Virtualization platform work in a cluster mode and therefore need to be fenced by the Manager either automatically or manually. Fencing allows a cluster to react to unexpected host failures as well as enforce power saving, load balancing, and virtual machine availability policies. Therefore, it is highly recommended that you configure the fencing parameters and test their correctness from time to time.

Hosts can be fenced automatically using the power management parameters, or manually by right clicking on a host and using the options on the menu. In a fencing operation, a host is re-booted, and if the host does not return to an active status within a prescribed time, Red Hat Enterprise Virtualization Manager marks it as unresponsive.

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

13.3.1. Setting the Parameters for Fencing

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

To set up fencing on a host:

1. Ensure the **Tree** tab is selected in the left hand pane of the Administration Portal. Expand the tree view using the **+** buttons until the host to be fenced is visible. The **Hosts** tab displays a list of all hosts in the system.
2. Select the host for which you wish to set up fencing and click the **Edit** button. The **Edit Host** dialog displays. Click the **Power Management** tab.

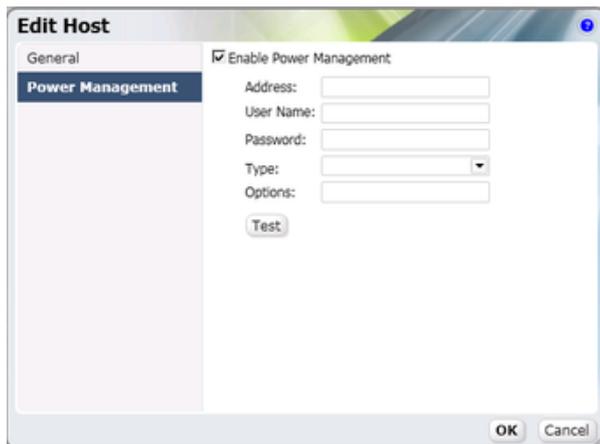


Figure 13.1. Power Management Dialog

- **Enable Power Management:** Select this checkbox to turn out-of-band (OOB) power management on. The fields for Power Management are enabled.
- The **Address** of the host. This is usually the address of the remote access card (RAC) on the host.
- A valid **User Name** for the OOB management.
- A valid, robust **Password** for the OOB management.
- The **Type** of the fencing device. Select the appropriate device from the drop down list.

alom	Sun Integrated Lights Out Manager (ILOM)
apc	APC Master MasterSwitch network power switch
bladecenter	IBM Bladecentre Remote Supervisor Adapter
drac5	Dell Remote Access Controller for Dell computers
eps	ePowerSwitch 8M+ network power switch
ilo	HP Integrated Lights Out standard
ipmilan	Intelligent Platform Management Interface
rsa	IBM Remote Supervisor Adaptor
rsb	Fujitsu-Siemens RSB management interface
wti	WTI Network PowerSwitch
cisco_ucs	Cisco UCS Integrated Management Controller



Note:

Depending on the **Type** selected, some or all of the following fields display on the **Power Management** tab.

- Click **Secure** to use SSH to connect to OOB management.

- The **Port** to connect to OOB management.
- Enter the **Slot** if a Blade server is being configured.
- Enter any **Options** that are needed for the fence-agents commands or ssh command. This is free text field that enables the administrator to enter commands that are not available via the graphical user interface. Red Hat Enterprise Virtualization Manager does not perform any checks on these options. These options should only be used by advanced users, as any errors will cause the host to become unreachable.
- Click the **Test** button to test the operation of the OOB management solution. If the power management options can be verified, the text *Test Succeeded, Host Status is: on* displays.



Warning

Power management parameters (userid, password, options, etc) are tested by Red Hat Enterprise Virtualization Manager only during setup and when required. There are no periodic tests unless manually scheduled. 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 Enterprise Virtualization Manager, fencing is likely to fail when most needed.

3. 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.

13.3.2. Using Power Management Functions on a Fenced Host

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.

To use power management options:

1. On the **Hosts** tab, select the host. Click the **Power Management** drop-down menu.

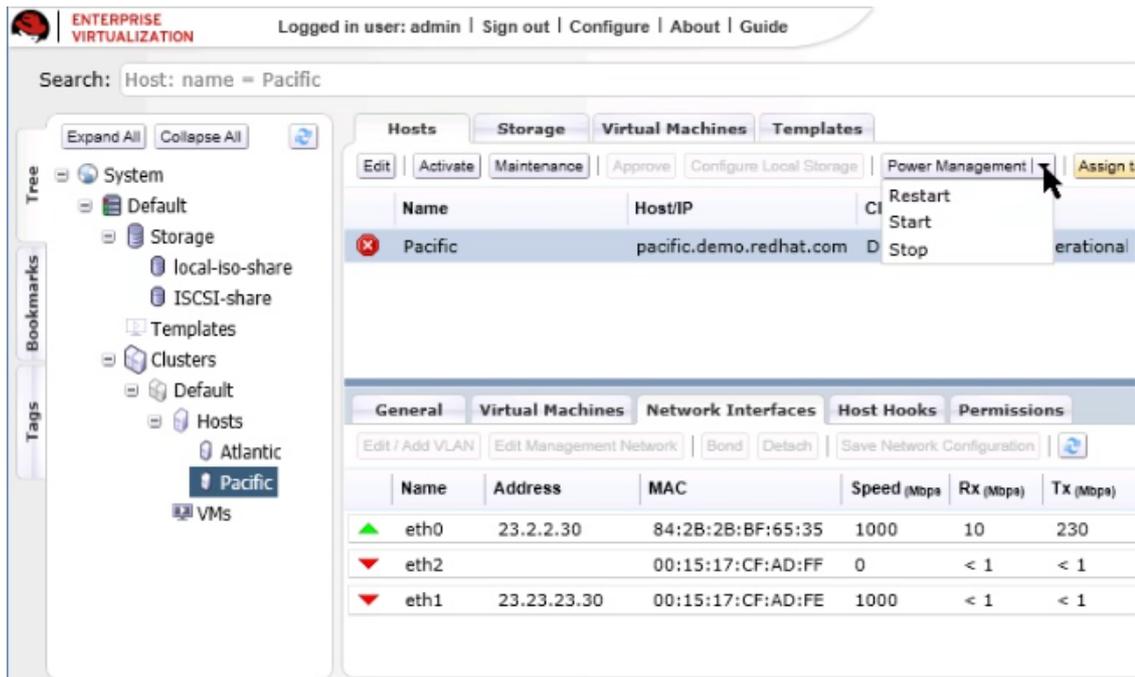


Figure 13.2. Restart

2. You can select between the following options:
 - **Restart:** This option stops the host, and waits until the host's status changes to Down. When the agent has verified that the host is down, the highly available virtual machines are restarted on another host in the cluster. The agent then restarts this host, and 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 ones will be restarted on another host. When the host has been stopped, its status displays as *Non Operational*.
3. Based on the option you chose, a relevant dialog pops up asking you to confirm your selection. Click **OK** to confirm and proceed.

13.3.3. Manually Fencing or Isolating a 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 system. You can reboot the host server manually, in order to isolate storage and networking issues.



Important

At least one host must be up and running in the data center to test fencing. Do not attempt to test the first host that is added to a data center, until at least one other host is up and running.

To manually fence a non-responsive host:

1. On the **Hosts** tab, select the host. The status must display as *Not Responding*.
2. Manually reboot the host. This could mean physically entering the lab and rebooting the host.
3. On the Administration Portal, right click the host entry and select the **Confirm Host has been rebooted** button.

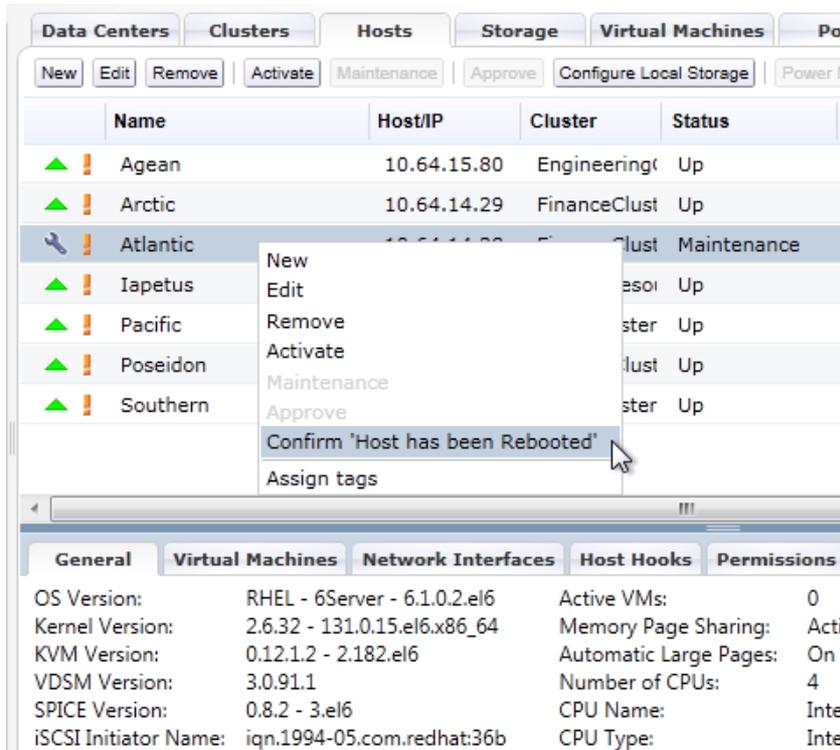


Figure 13.3. The Host Right-click menu

4. A message displays prompting you to ensure that the host has been shut down or rebooted. Select the **Approve Operation** check box and click **OK**.
5. The host to be fenced is isolated from the virtualized system, enabling any of its functions to be automatically transferred to an active host.
6. After the non-responsive host is rectified, and is reinstalled or rebooted, click the **Activate** button to restore the host status to Up.

13.4. Virtual Machine High Availability

This section provides instructions on how to configure virtual machine high availability. When a highly available virtual machine crashes due to a virtual machine error, it will be automatically restarted on the same host. In contrast, when a highly available virtual machine crashes because of a host error, it will be automatically restarted on another host in the same cluster.



Important

High availability can only be configured for virtual servers, not virtual desktops.

13.4.1. Configuring a Highly Available Virtual Machine

High availability must be configured individually for each virtual machine.

To mark a virtual machine as highly available:

1. On the **Virtual Machines** tab, select the entry for the virtual machine that you want to migrate, and click the **Edit** button.
2. On the **Edit Server Virtual Machine** dialog, click the **High Availability** tab.

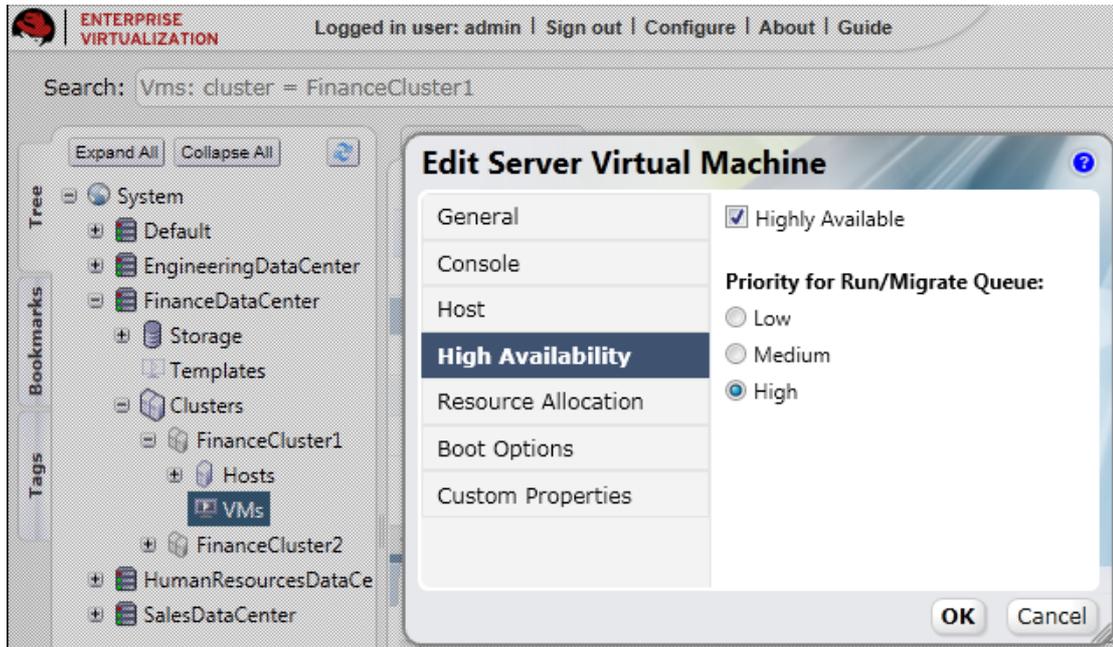


Figure 13.4. Set virtual machine high availability

3. Tick the **Highly Available** checkbox. Set the priority of the virtual machine by selecting **Low**, **Medium**, or **High** under **Priority for Run/Migrate Queue**. When migration is triggered, a queue is created in which the high priority virtual machines are migrated first. If a cluster is running low on resources, only the high priority virtual machines are migrated.
4. Click **OK** to set the virtual machine high availability.

You can check if a virtual machine is highly available when you select it and click on its **General** tab in the details pane.

13.4.2. Setting a Cluster Resilience Policy

A cluster resilience policy determines the action to be taken for the virtual machines running on a host which has shut down unexpectedly or been put into maintenance.

The next available host depends on the cluster policy you have set, as described in [Section 12.3.2, “Cluster Policy”](#). For example, if your cluster policy is set to Even Distribution, the next available host refers to the host which is running the lightest CPU workload at the time migration is triggered. The exception to this rule is the host which is elected as the Storage Pool Manager, because more of its CPU resources are in use for writing metadata from all virtual machines in the system. Therefore, the SPM tends to run less virtual machines than other hosts.

To define a cluster resilience policy:

1. On the **Clusters** tab which displays on the results list, select the cluster and click **Edit**. The **Edit Cluster** dialog displays. Select the **Resilience Policy** tab.

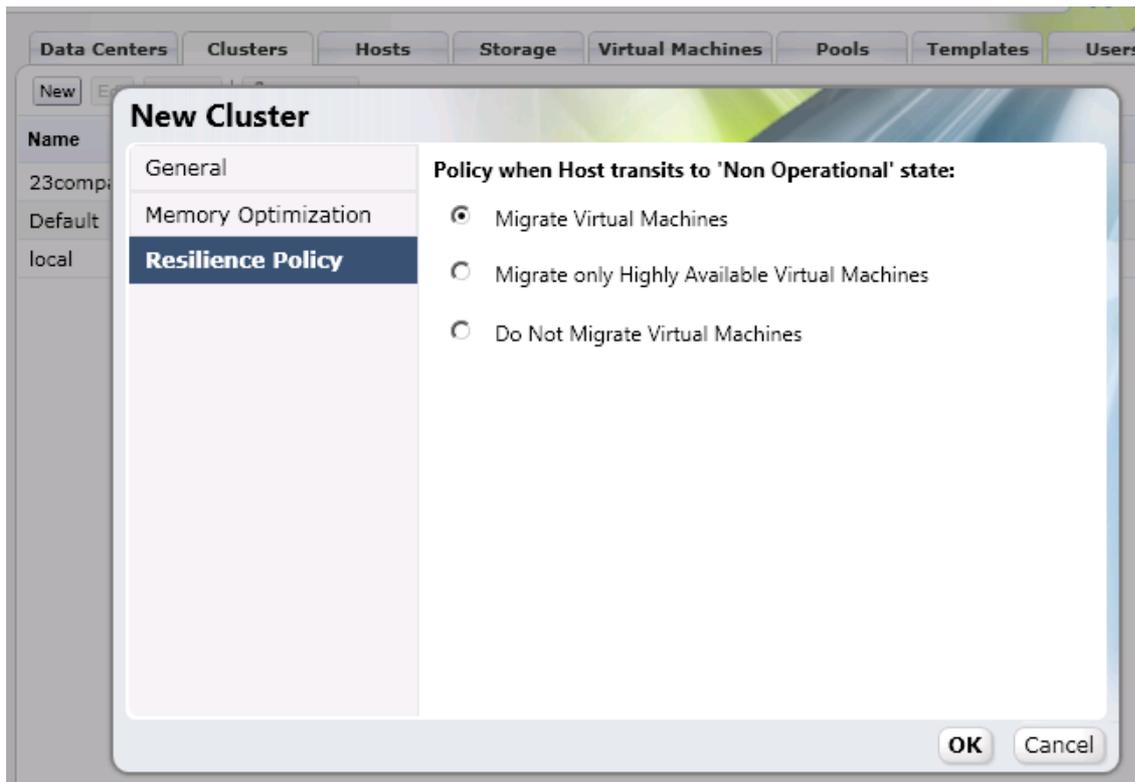


Figure 13.5. Cluster Resilience Policy

2. Select the policy to use. The available options are:
 - **Migrate Virtual Machines:** This option migrates all virtual machines when a host becomes non-operational. The virtual machines are migrated in order of their defined priority, so the highly available virtual machines are at the front of the migration queue.
 - **Migrate only Highly Available Virtual Machines:** This option migrates only the highly available virtual machines, so the other available hosts in the cluster do not become overloaded if they are to rescue virtual machines from non-operational hosts.
 - **Do Not Migrate Virtual Machines:** This option prevents your virtual machines from being migrated at all.



Note

Virtual machine migration is a network-intensive operation. For instance, on a setup where a host is running ten or more virtual machines, migrating all of them can be a long and resource-consuming process. Therefore, select the policy action to best suit your setup. If you prefer a conservative approach, disable all migration of virtual machines. Alternatively, if you have many virtual machines, but only several which are running critical workloads, select the option to migrate only highly available virtual machines.

3. Click **OK** to save the policy selection for virtual machines in the cluster.

Managing Multilevel Administration

This section describes how to set up user roles that control levels of permissions to different objects and actions in your virtualized environment. Red Hat Enterprise Virtualization supports multilevel administration. This means that users can be assigned a variety of permissions to specific objects, using a number of default roles. In addition, customized roles can be created and assigned to users.

Red Hat Enterprise Virtualization relies on directory services for user authentication. Currently the two supported providers of directory services for use with the Red Hat Enterprise Virtualization Manager are Identity, Policy, and Audit (IPA) and Active Directory.



Note

Users are not created in Red Hat Enterprise Virtualization platform, but in the Directory Services domain. Red Hat Enterprise Virtualization Manager can be configured to use multiple Directory Services domains. See the *Red Hat Enterprise Virtualization Installation Guide* for more information.

14.1. Configuring Roles

Roles are predefined sets of privileges that can be configured from Red Hat Enterprise Virtualization Manager, providing access and management permissions to different levels of resources in the data center, to specific physical and virtual resources. Permissions enable users to perform actions on objects, as explained in [Section 5.1, “Authorization Model”](#).

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.

14.1.1. Roles

There are two types of roles in Red Hat Enterprise Virtualization, **administrator** roles and **user** roles. See [Table 5.1, “Red Hat Enterprise Virtualization User Roles”](#) for details on roles.

Role Types

- Administrator - Allows access to the *Administration Portal* for managing virtual resources. An administrator role does not confer any permissions for the user portal.
- User - Allows access to the *User Portal* for managing and accessing virtual machines. A user role does not confer any permissions for the Administration Portal

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

The default roles cannot be removed from the platform, and their privileges cannot be modified. However, you can clone them, and then customize the new roles as required.

14.1.2. Creating Custom Roles

In addition to the default roles, you can set up custom roles that permit actions on objects, such as virtual machines, hosts and clusters, and assign privileges to specific entities. Use the roles to create a granular model of permissions to suit the needs of the enterprise or a group or set of users. Use the **Configure** option to work with roles. You can create a **New** role, **Edit**, **Clone** or **Remove** an existing role. In each case the appropriate dialog box displays.

Once the role is set up, you can assign the role to users as required.

To create a new role:

1. On the header bar of the Red Hat Enterprise Virtualization Manager menu, click **Configure**. The **Configure** dialog box displays. The dialog box includes a list of default User and Administrator roles, and any custom roles.
2. Click **New**. The **New Role** dialog box displays.
3. Enter the **Name** and **Description** of the new role. This name will display in the list of roles.
4. Select either **Admin** or **User** as the **Account Type**. If **Admin** is selected, this role displays with the administrator icon in the list.
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 deselect the actions you wish to permit/deny for the role you are setting up.
7. Click **OK** to apply the changes you have made. The new role displays on the list of roles.

14.1.3. Editing Roles

You may need to change the permissions, names or descriptions for the custom roles. Note that you cannot make changes to the default roles. To edit custom roles, you can use the **Edit** button on the Configure dialog box.

To edit a role:

1. On the header bar of the Red Hat Enterprise Virtualization Manager menu, click **Configure**. The **Configure** dialog box displays. The dialog box below shows the list of Administrator roles.
2. Click **Edit**. The **Edit Role** dialog box displays.

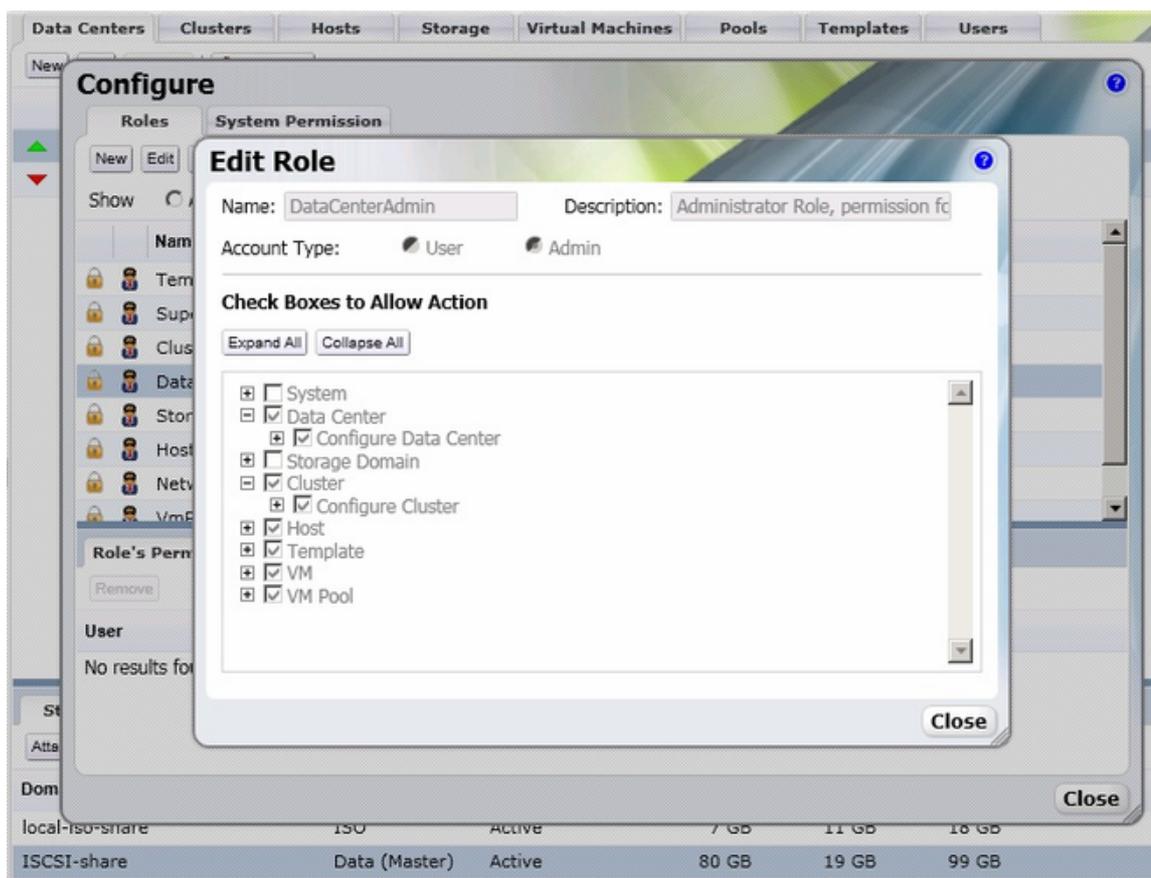


Figure 14.1. The Edit Role Dialog Box

3. If necessary, edit the **Name** and **Description** of the role. This name will display in the list of roles.
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 deselect the actions you wish to permit/deny for the role you are editing.
6. Click **OK** to apply the changes you have made.

14.1.4. Cloning Roles

You can create a new role by cloning an existing default or custom role, and changing the permissions set as required. Use the **Clone** button on the Configure dialog box.

To clone a role:

1. On the header bar of the Red Hat Enterprise Virtualization Manager menu, click **Configure**. The **Configure** dialog box displays. The dialog box includes a list of default roles, and any custom roles that exist on the platform.

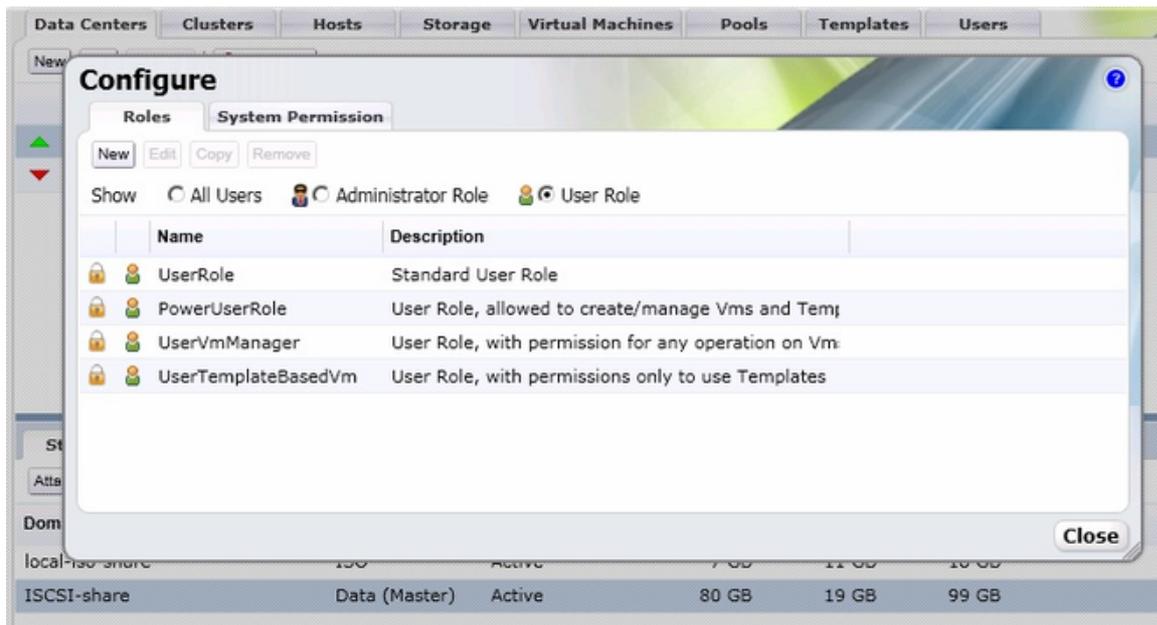


Figure 14.2. The Configure Dialog Box

2. Click **Clone**. The **Clone ROLE** dialog box displays.
3. Change the **Name** and **Description** of the new role. This name will display in the list of roles.
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 deselect the actions you wish to permit/deny for the role you are editing.
6. Click **Close** to apply the changes you have made.

14.2. User Roles Examples

This section provides a sample of how to assign a number of sample user roles, from a basic end user of a virtual desktop to a power user of the User Portal.

14.2.1. Setting Up an End User

Penelope is the receptionist at ViewGen Inc. Apart from receptionist duties, she also needs to check her mail, arrange quotes, and manage appointments for some executives. To do this, she needs a virtual machine with access to the network and a number of office applications. The administrator provides her with an account that enables her to log into and use a single virtual machine. Penelope needs a UserRole account only.

To assign a UserRole account:

1. Ensure that Penelope has a valid user account in the IPA or Active Directory domain. See the *Directory Services* chapter of the *Red Hat Enterprise Virtualization Installation Guide*.
2. In the Administration Portal, add Penelope's account to the Red Hat Enterprise Virtualization platform. See [Section 5.3.1, "Adding Users and Groups"](#).
3. Create a virtual machine (Penelope-VM) in the appropriate data center and cluster. Provision it with the applications that Penelope requires, and ensure the machine has a status of *Up*.

4. Click the **Virtual Machines** tab, and select the virtual machine (Penelope-VM). Click **Permissions** from the Details pane.
5. The **Permissions** tab displays a list of users and their current roles and permissions, if any. Note that there are no inherited permissions for virtual machines. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter Penelope's Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
6. Tick Penelope's checkbox. Scroll through the **Assign role to user** list and select **UserRole**.
7. Click **OK**. Penelope's name displays in the **Permissions** tab. Penelope can now connect to the virtual machine via the User Portal and log into the virtual machine that has been created for her. Since no one else has been assigned to this machine, she will always have access to the virtual machine, and should experience no difference to using a physical machine.

14.2.2. Setting Up a Virtual Machine Administrator

Penelope is still the receptionist at ViewGen Inc. However, in the few months she has been there she has proved to be a quick learner and is capable of maintaining her own virtual machine (Penelope-VM). The departments system administrator decides to give Penelope system administration permissions to her own virtual machine, so that she can perform tasks such as memory upgrades and backups, or reconfigure her machine independently. Penelope needs a **UserVmManager** account for her own virtual machine.

To assign a **UserVmManager** role:

1. Since Penelope is an existing user, you do not need to check her status, or set up her machine. In the Administration Portal, click the **Virtual Machines** tab, and select the virtual machine (Penelope-VM). Click **Permissions** from the Details pane.
2. The **Permissions** tab displays a list of users and their current roles and permissions, if any. Note that there are no inherited permissions for virtual machines. Click **Add** to add an existing user. The **Add Permission to User** dialog box displays. Enter Penelope's Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
3. Tick Penelope's checkbox. Scroll through the **Assign role to user** list and select **UserVmManager**.
4. Click **OK**. Penelope's name displays in the **Permissions** tab as a **UserVmManager**. Penelope can now administrate the virtual machine as well as log into the virtual machine.

14.2.3. Setting Up a Power User

Penelope is now an experienced office manager at ViewGen Inc. Penelope now has additional responsibilities, and occasionally needs to take charge of recruitment tasks, such as scheduling interviews and following up on reference checks, if the HR Manager is on vacation, or needs additional help. As per corporate policy, Penelope needs to use a particular application for this task. As this is an occasional task, Penelope would rather use a separate virtual machine to run the recruitment application.

For Penelope to create a new virtual machine through the User Portal, she needs **PowerUserRole** permissions for the data center in which the new virtual machine will reside. This is because to create a new virtual machine, she needs to be able to make changes to several components within the data center, including creating the virtual machine disk image in the storage domain. However, note that this is not the same as assigning **DataCenterAdmin** privileges to Penelope. As a **PowerUser** for a data center, Penelope can 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.

To assign a PowerUser Role:

1. Since Penelope is an existing user, you do not need to check her status. In the Administration Portal, click the **Data Centers** tab. Select the data center in which the new virtual machine will reside. Click **Permissions** from the Details pane. A list of users and their current roles and permissions displays.
2. The **Permissions** tab displays a list of users and their current roles and permissions, if any. Click **Add** to add an existing user. The **Add Permission to User** dialog displays. Enter Penelope's Name, or User Name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
3. Tick Penelope's checkbox. Scroll through the **Assign role to user** list and select **PowerUserRole**.
4. Click **OK**. In the **Permissions** tab, Penelope's name now displays with **PowerUserRole** privileges. Penelope can now create a virtual machine for herself from the User Portal.

14.3. 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 14.1. Cluster Permissions

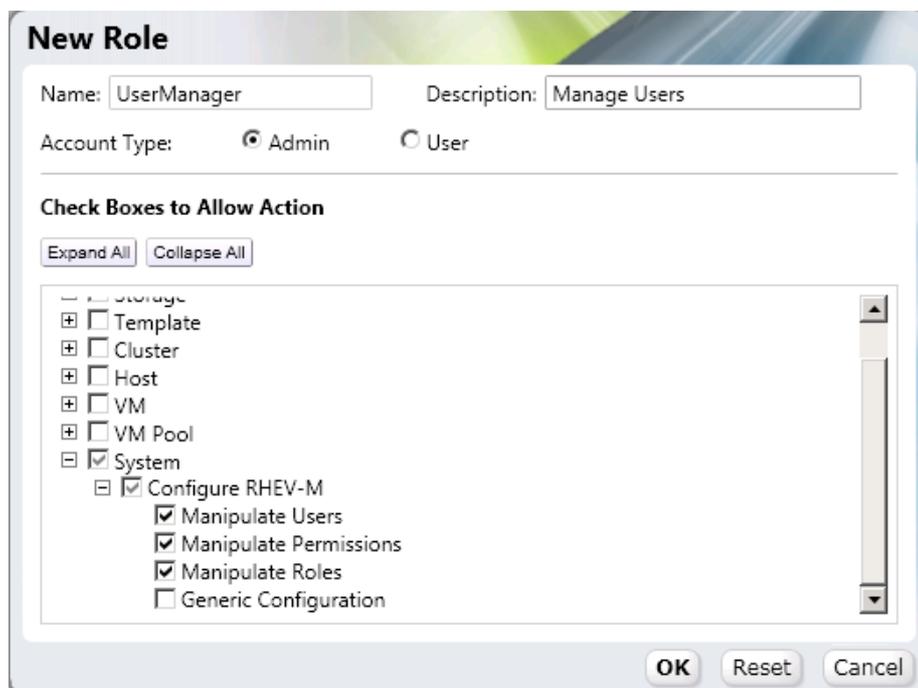
Sarah is the system administrator for the accounts department of a company. All the virtual resources for her department are organized under a Red Hat Enterprise Virtualization cluster called **accounts**. She is assigned the **ClusterAdmin** role on the accounts cluster. This enables her to manage all virtual machines in the cluster, since the virtual machines are child objects of the cluster as shown in *Figure 5.3, "Red Hat Enterprise Virtualization Object Hierarchy"*. Managing the virtual machines includes editing, adding or removing virtual resources such as disks, and taking snapshots. It does not allow her to manage any resources outside this cluster. Because **ClusterAdmin** is an administrator role, it allows her to use the Administration Portal to manage these resources, but does not give her any access via the User Portal.

Example 14.2. 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 **PowerUserRole** on the johndesktop virtual machine. This allows him to access this single virtual machine using the User Portal. Because he has **PowerUser** permissions, he can modify the virtual machine and add resources to it, such as new virtual disks. Because **PowerUserRole** is a user role, it does not allow him to use the Administration Portal.

Example 14.3. Custom Role Permissions

Rachel works in the IT department, and is responsible for managing user accounts in Red Hat Enterprise 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.



New Role

Name: Description:

Account Type: Admin User

Check Boxes to Allow Action

- Storage
- Template
- Cluster
- Host
- VM
- VM Pool
- System
 - Configure RHEV-M
 - Manipulate Users
 - Manipulate Permissions
 - Manipulate Roles
 - Generic Configuration

Figure 14.3. UserManager Custom Role

The **UserManager** custom role shown above allows manipulation of users, permissions and roles. These actions are organized under System - the top level object of the hierarchy shown in [Figure 5.3, "Red Hat Enterprise 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 only use the Administration Portal, not the User Portal.

Backing Up and Restoring the Red Hat Enterprise Virtualization Manager.

The Red Hat Enterprise Virtualization Manager is key part of a Red Hat Enterprise Virtualization environment which is responsible for maintaining important information about the environment. It is recommended that the Red Hat Enterprise Virtualization Manager be backed up so that recovering from unforeseen events is as simple and fast as possible.

15.1. Backup and Restore the rhevm Postgres Database

Information about the Red Hat Enterprise Virtualization environment is kept in a **Postgres** database, including information about virtual machines, hosts, networks, storage, users, and more. The `rhevm` database is a key component of the Red Hat Enterprise Virtualization environment so it is highly recommended that regular backups of the database be taken.

While these procedures specifically mention the `rhevm` database, they should be repeated for each database being backed up. For example, the same procedures can be used to back up and restore the `rhevm_history` database.

15.1.1. Backing up Databases in Red Hat Enterprise Virtualization

Procedure 15.1. To backup the `rhevm` database:

1. From the terminal on the Red Hat Enterprise Virtualization Manager server as root use the `pg_dump` command to dump the `rhevm` database:

```
# pg_dump -C -E UTF8 --column-inserts --disable-dollar-quoting --disable-triggers -U
postgres --format=p -f /usr/share/rhevm/db-backups/dump_RHEVDB_BACKUP_`date +%Y%m%d_%R`.sql rhevm
```

The output file name and directory are user specified, `/usr/share/rhevm/db-backups/dump_RHEVDB_BACKUP_`date +%Y%m%d_%R`.sql` should be replaced with a desired location and file name.

2. Consider mounting a remote backup location locally and using the `pg_dump` to write the database backup file directly to the mounted remote location.
3. Copy the `.sql` file to a remote backup location. This can be accomplished using `scp`, `rsync`, or some other third party back up tool.
4. Consider automating the `rhevm` database backup with ssh keys, a script, and a cron job.

Example 15.1. Example `rhevm` Postgres database backup script

This script presumes that an ssh key has been set up for the user that executes the script. The public key must then be added to the remote users `~/.ssh/authorized_hosts` using the `ssh-copy-id BACKUPUSER@BACKUPDIRECTORY` command so that the `scp` command can be done without a password.

```
#!/bin/sh
# Enter appropriate values for the BACKUPUSER, BACKUPSERVER, and BACKUPDIRECTORY
variables.
DATE = `date +%Y%m%d_%R`;
```

```

BACKUPUSER = ;
BACKUPSERVER = ;
BACKUPDIRECTORY = ;
# pg_dump -C -E UTF8 --column-inserts --disable-dollar-quoting --disable-triggers -U
postgres --format=p -f /usr/share/rhev/db-backups/dump_RHEVDB_BACKUP_`date +%Y%m%d_
`%R`.sql rhvm;
scp /usr/share/rhev/db-backups/dump_RHEVDB_BACKUP_`date +%Y%m%d_`%R`.sql
$BACKUPUSER@$BACKUPSERVER:$BACKUPDIRECTORY;
    
```

Now that you have created a backup of the rhvm database using the **pg_dump**, you can use the .sql file you created to restore the rhvm database.

15.1.2. Restoring Databases in Red Hat Enterprise Virtualization

Procedure 15.2. To restore the rhvm database:

- To restore a .sql that was created using the **pg_dump** command, use the **psql** interactive shell. From the terminal on the Red Hat Enterprise Virtualization Manager server as root:

```

# psql -U postgres -d rhvm -w /usr/share/rhev/db-backups/dump_RHEVDB_BACKUP_`date +%Y
`%m%d_`%R`.sql
    
```

Substitute the name of the .sql file being restored for *dump_RHEVDB_BACKUP_CURRENT_DATE*.

- If you are required to remove the existing rhvm database to create a new one, first stop the **jbossas** service using the **service jbossas stop** command.
- When you have created the new rhvm database and imported database backup file, start the **jbossas** service using the **service jbossas start**.

15.2. Backing up and Restoring Manager Configuration Files

The Red Hat Enterprise Virtualization Manager stores customized configurations as configuration files. These files contain specific configuration details of a given environment, and must also be backed up if a recovery is to be performed.

15.2.1. Red Hat Enterprise Virtualization Manager Configuration Files Requiring Backup

Table 15.1. Configuration files and directories requiring backup.

Location	What is it?
/etc/jbossas/jbossas.conf	Configuration file for JBoss Application Server.
/etc/rhev/	Contains Red Hat Enterprise Virtualization Manager configuration files.
/etc/yum/pluginconf.d/versionlock.list	Contains version information about currently installed Red Hat Enterprise Virtualization components.

Location	What is it?
/etc/pki/rhev/	Security certificates provided by the Red Hat Enterprise Virtualization Manager to clients.
/etc/jbossas/rhev-slimmed/	Contains the optimized enhancements for the JBoss server used by the Red Hat Enterprise Virtualization Manager.
/usr/share/rhev-reports-server/buildomatic	Contains files required to build the Red Hat Enterprise Virtualization reports server.
/usr/share/rhev/conf/iptables.example	An example of a correct iptables configuration that allows the communications required by the Red Hat Enterprise Virtualization Manager.
/usr/share/rhev/kerberos/krb5.conf	Kerberos configuration file for Red Hat Enterprise Virtualization external authorization.
/usr/share/rhev/dbscripts/create_db.sh.log	Log file from the creation of the rhev database.
/usr/share/rhev/rhev.ear/rhevmanager.war/ExternalConfig.txt	An XML configuration file for the Red Hat Enterprise Virtualization Manager Web Service.
/usr/share/rhev/rhev.ear/rhevmanager.war/ServerParameters.js	Contains port information for accessing the manager.
/usr/share/rhev-reports/reports-INSERT_VERSION_NUMBER/resources/organizations/rhevreports/Resources/JDBC/data_sources/rhev.xml	Contains JDBC connection information for the Red Hat Enterprise Virtualization databases.
/usr/share/rhev-reports/reports-INSERT_VERSION_NUMBER/users/rhevreports/rhev-002dadmin.xml	Contains plain-text, un-encrypted user and password information for the rhev-admin user.
/usr/share/rhev-reports-server/buildomatic/default_master.properties	Contains settings to handle the configuration and deployment of JasperServer.
/usr/share/rhev-reports-server/buildomatic/install.xml	Installation related macros and targets for configuration and deployment of JasperServer.
/usr/share/rhev-reports-server/buildomatic/setup.xml	Sets properties, reads and sets up configuration files, and checks the app server of JasperServer.

Location	What is it?
/root/.pgpass	Contains password information for the postgres and rhevm users.
/root/.rnd	Random seed file, used to generate secure certificates.

When all of the files in [Table 15.1, “Configuration files and directories requiring backup.”](#) have been backed up, you will be able to recover the Red Hat Enterprise Virtualization Manager to a working state after an unforeseen event.

15.2.2. Restoring Red Hat Enterprise Virtualization Manager Configuration Files

Taking regular backups of Red Hat Enterprise Virtualization Manager configuration files allows you to quickly return a broken Manager installation to a working state. With all of the files in [Table 15.1, “Configuration files and directories requiring backup.”](#) backed up, you can re-install the same version of the Manager, and restore the backed up files over top of the new installation.

All commands in the Manager restoration procedure are entered as the root user.

Procedure 15.3. To Restore the Red Hat Enterprise Virtualization Manager

1. Stop the **JBoss** service:

```
# service jbossas stop
```

2. Completely remove all previous installations of the Red Hat Enterprise Virtualization Manager:

```
# yum remove rhevm
```

3. Remove the `rhevm-slimmed` profile:

```
# rm -rf /var/lib/jbossas/server/rhevm-slimmed
```

4. Delete the old configuration for `rhevm-slimmed` profile:

```
# rm -rf /etc/jbossas/rhevm-slimmed
```

5. Remove `/etc/pki/rhevm`:

```
# rm -rf /etc/pki/rhevm
```

6. Install the Red Hat Enterprise Virtualization Manager:

```
# yum install -y rhevm
```

7. Restore the files listed in [Table 15.1, “Configuration files and directories requiring backup.”](#) to their original locations.

8. Make sure the ownership of the `.keystore` file is correct:

```
# chown jboss:jboss /etc/pki/rhev/.keystore
```

9. Make sure the permissions of the notifier.conf file is correct:

```
# chmod 644 /etc/rhev/notifier/notifier.conf
```

10. Start the JBoss service:

```
# service jbossas start
```

You have now recovered from an unforeseen event that adversely affected your installation of the Red Hat Enterprise Virtualization Manager.

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. Note that at this time hooks are not able to run on Red Hat Enterprise Virtualization Hypervisors, they must only be used on Red Hat Enterprise Linux hosts. By setting custom properties on virtual machines it is possible to pass additional parameters, specific to a given virtual machine, to the hook scripts.

VDSM

The VDSM service is used by the Red Hat Enterprise Virtualization Manager to manage Red Hat Enterprise Virtualization Hypervisors 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 hypervisor host managed by Red Hat Enterprise Virtualization Manager. It answers XML-RPC calls from clients. The Red Hat Enterprise Virtualization Manager functions as a VDSM client.

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. This is further described in [Section 16.1, "Environment"](#).



Warning VDSM hooks can cause virtual machine crashes and data loss

VDSM hooks can interfere with the operation of Red Hat Enterprise 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.

Table 16.1. Supported VDSM Events

Name	Description
before_vm_start	Before VM start.
after_vm_start	After VM start.
before_vm_cont	Before VM continue.
after_vm_cont	After VM continue.
before_vm_pause	Before VM pause.
after_vm_pause	After VM pause.

Name	Description
before_vm_hibernate	Before VM hibernate.
after_vm_hibernate	After VM hibernate.
before_vm_dehibernate	Before VM de-hibernate.
after_vm_dehibernate	After VM de-hibernate.
before_vm_migrate_source	Before VM migration, run on the source hypervisor host from which the migration is occurring.
after_vm_migrate_source	After VM migration, run on the source hypervisor host from which the migration is occurring.
before_vm_migrate_destination	Before VM migration, run on the destination hypervisor host to which the migration is occurring.
after_vm_migrate_destination	After VM migration, run on the destination hypervisor host to which the migration is occurring.
after_vm_destroy	After VM destruction.
before_vdsm_start	Before VDSM is started on the hypervisor 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 hypervisor host. after_vdsm_stop hooks are executed as the user root, and do not inherit the environment of the VDSM process.

16.1. 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.

16.1.1. Domain XML

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. 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`.



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.

16.1.2. Custom Properties

Red Hat Enterprise Virtualization Manager also allows specification of custom properties for each virtual machine. Each of these properties is provided to hooks as an environment variable. This allows users to pass virtual machine dependent parameters to the hook scripts. Additionally hook scripts are able to use the presence, or absence, of a given custom property to determine whether they should execute to completion.

Defining Custom Properties

The custom properties that are accepted by the Red Hat Enterprise Virtualization Manager — and in turn passed to custom hooks — are defined using the configuration tool, **rhev**m-**config**. Run this command as the root user on the host where Red Hat Enterprise Virtualization Manager is installed.

The configuration key **UserDefinedVMProperties** is 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 this configuration key.

Where multiple custom properties are defined they 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 it is necessary to include all of the custom properties in the command used to set the key's value.

Once the configuration key has been updated the `jbossm` service must be restarted for it to take effect.

Example 16.1. Defining *smartcard* Custom Property

1. Check the existing custom properties defined by the **UserDefinedVMProperties** configuration key using **rhev**m-**config -g UserDefinedVMProperties**.

In this case the custom property *memory* is already defined. The regular expression `^[0-9]+$` ensures that the custom property will only ever contain numeric characters.

```
# rhevm-config -g UserDefinedVMProperties
UserDefinedVMProperties : version: general
UserDefinedVMProperties : version: 2.2
UserDefinedVMProperties : memory=^[0-9]+$ version: 3.0
```

2. As 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**.

```
# rhvm-config -s UserDefinedVMProperties='memory=[0-9]+$;smartcard=(true|false)$'
--cver=3.0
```

3. Verify that the custom properties defined by the **UserDefinedVMProperties** configuration key now match your expectations.

```
# rhvm-config -g UserDefinedVMProperties
UserDefinedVMProperties : version: general
UserDefinedVMProperties : version: 2.2
UserDefinedVMProperties : memory=[0-9]+$;smartcard=(true|false)$ version: 3.0
```

4. Finally, the `jbossas` service must be restarted for the configuration change to take effect.

```
# service restart jbossas
```

Setting Custom Properties

Once custom properties are defined to Red Hat Enterprise Virtualization Manager you are able to begin setting them on virtual machines. Custom properties are set on the **Custom Properties** tab of the **New Server Virtual Machine**, **New Desktop Virtual Machine**, **Edit Server Virtual Machine**, and **Edit Desktop Virtual Machine** dialog boxes in the Administration Portal.

You are also able to set custom properties from the **Run Once** dialog box. Custom properties set from the **Run Once** dialog box will only apply to the virtual machine until it is next shutdown.

The **Custom Properties** must be provided as key to value pairs of the form **key=value**, where **key** is the name of the custom property and **value** is the value which it should be set to for this virtual machine. Where multiple custom properties are being provided they must be separated by a semi-colon, as illustrated in the example below.



Figure 16.1. Custom Properties Field

Evaluating Custom Properties in Hooks

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 16.2. 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)
```

16.1.3. 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 16.2. Hooking module functions

Name	Argument	Description
<code>tobool</code>	string	Converts a string "true" or "false" to a Boolean value
<code>read_domxml</code>	-	Reads the virtual machine's libvirt XML into a DOM object
<code>write_domxml</code>	DOM object	Writes the virtual machine's libvirt XML from a DOM object

16.2. Execution

`before_vm_start` scripts may 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 Enterprise Virtualization environment. In particular, ensure you never change the uid 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 16.3. 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", "/my_file"] )
```

The standard error stream of hook scripts is collected in VDSM's log. This information is used to debug hook scripts.

Return Codes

Hook scripts must return one of the return codes shown in [Table 16.3, "Hook return codes"](#). The return code will determine whether further hook scripts are processed by VDSM.

Table 16.3. 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

16.3. 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 16.4. CPU Pinning

Purpose:

This hook script allows the pinning of virtual machines to specific CPUs based on the *pincpu* custom property. Where the custom property is not set no action is taken.

Configuration String:

```
pincpu=^[\\^]?\\d+(-\\d+)?(,[\\^]?\\d+(-\\d+)?)*$
```

The regular expression used allows the *pincpu* custom property for a given virtual machine to specify:

- that a specific CPU be used for the virtual machine (**pincpu=0**, specifies that only CPU 0 be used), or
- that a range of CPUs be used for the virtual machine (**pincpu=1-4**, specifies that CPUs 1 through 4 be used), or
- that a specific CPU not be used for the virtual machine (**pincpu=^3**, specifies that CPU 3 not be used), or

- any comma separated combination of the above (**pincpu=1-4,6**, specifies that CPUs 1 to 4, and 6 be used).

Script:

/usr/libexec/vdsm/hooks/before_vm_start/50_pincpu

```
#!/usr/bin/python

import os
import sys
import hooking
import traceback

'''
pincpu usages
=====
pincpu=0 (use the first cpu)
pincpu=1-4 (use cpus 1-4)
pincpu=^3 (dont use cpu 3)
pincpu=1-4,6 (or all together)
'''

if os.environ.has_key('pincpu'):
    try:
        domxml = hooking.read_domxml()

        vcpu = domxml.getElementsByTagName('vcpu')[0]

        if not vcpu.hasAttribute('cpuset'):
            sys.stderr.write('pincpu: pinning cpu to: %s\n' % os.environ['pincpu'])
            vcpu.setAttribute('cpuset', os.environ['pincpu'])
            hooking.write_domxml(domxml)
        else:
            sys.stderr.write('pincpu: cpuset attribute is present in vcpu, doing nothing
\n')
    except:
        sys.stderr.write('pincpu: [unexpected error]: %s\n' % traceback.format_exc())
        sys.exit(2)
```

Example 16.5. 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 A. Utilities

A.1. Domain Management Tool

Red Hat Enterprise Virtualization Manager uses directory services to authenticate users. While during installation the manager sets up a domain named **internal** this is only used for the `admin` user. To add and remove other users from the system it is first necessary to add the directory service(s) in which they are found.

The supported directory services are Active Directory and IPA. Red Hat Enterprise Virtualization Manager includes a domain management tool, **`rhevmanage-domains`**, to add and remove domains provided by these services. In this way it is possible to grant access to the Red Hat Enterprise Virtualization environment to users stored across multiple domains. This is true even where some users are stored in a domain managed by Active Directory and others are stored in a domain managed by IPA.

You will find the **`rhevmanage-domains`** command on the machine to which Red Hat Enterprise Virtualization Manager was installed. The **`rhevmanage-domains`** command must be run as the root user.

A.1.1. Syntax

The usage syntax is:

```
Usage: rhevmanage-domains -action=ACTION [options]
```

Available actions are:

add

Add a domain to the manager's directory services configuration.

edit

Edit a domain in the manager's directory services configuration.

delete

Delete a domain from the manager's directory services configuration.

validate

Validate the manager's directory services configuration. The command attempts to authenticate to each domain in the configuration using the configured username and password.

list

List the manager's current directory services configuration.

The options available to be combined with the actions on the command line are:

-domain=DOMAIN

Specifies the domain the action must be performed on. The *-domain* parameter is mandatory for *add*, *edit*, and *delete*.

-user=USER

Specifies the domain user to use. The *-user* parameter is mandatory for *add*, and optional for *edit*.

-interactive

Specifies that the domain user's password is to be provided interactively. This option, or the *-passwordFile* option, must be used to provide the password for use with the *add* action.

-passwordFile=FILE

Specifies that the domain user's password is on the first line of the provided file. This option, or the *-interactive* option, must be used to provide the password for use with the *add* action.

-configFile=FILE

Specifies an alternative configuration file that the command must load. The *-configFile* parameter is always optional.

-report

Specifies that when performing the *validate* action all validation errors encountered will be reported in full.

Common usage examples are discussed further within this guide. For full usage information consult the **rhevmanage-domains** command's help output:

```
# rhevm-manage-domains --help
```

A.1.2. Examples

The following examples demonstrate the use of the **rhevmanage-domains** to perform basic manipulation of the Red Hat Enterprise Virtualization Manager domain configuration.

Example A.1. Adding Domains to Configuration

This example runs the **rhevmanage-domains** command to add the **directory.demo.redhat.com** domain to the Red Hat Enterprise Virtualization Manager configuration. The configuration is set to use the `admin` user when querying the domain with the password to be provided interactively.

```
# rhevm-manage-domains -action=add -domain='directory.demo.redhat.com' -user='admin' -
interactive
loaded template kr5.conf file
setting default_tkt_encotypes
setting realms
setting domain realm
success
User guid is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully added domain directory.demo.redhat.com
```

Example A.2. Edit Domain in Configuration

This example runs the **rhevmanage-domains** command to edit the **directory.demo.redhat.com** domain in the Red Hat Enterprise Virtualization Manager configuration. The configuration is updated to use the `admin` user when querying this domain with the password to be provided interactively.

```
# rhevm-manage-domains -action=edit -domain=directory.demo.redhat.com -user=admin -
interactive
loaded template kr5.conf file
setting default_tkt_encotypes
setting realms
```

```
setting domain realmo
success
User guid is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully edited domain directory.demo.redhat.com
```

Example A.3. Deleting Domains from Configuration

This example runs the **rhevmanage-domains** command to remove the **directory.demo.redhat.com** domain from the Red Hat Enterprise Virtualization Manager configuration. Users defined in the removed domain will no longer be able to authenticate with Red Hat Enterprise Virtualization Manager. The entries for the affected users will remain defined in Red Hat Enterprise Virtualization Manager until they are explicitly removed.

The domain being removed in this example is the last one listed in the Red Hat Enterprise Virtualization Manager configuration. A warning is displayed highlighting this fact and that only the admin user from the **internal** domain will be able to log in until another domain is added.

```
# rhevm-manage-domains -action=delete -domain='directory.demo.redhat.com'
WARNING: Domain directory.demo.redhat.com is the last domain in the configuration. After
deleting it you will have to either add another domain, or to use the internal admin user
in order to login.
Successfully deleted domain directory.demo.redhat.com. Please remove all users and groups
of this domain using the Administration portal or the API.
```

Example A.4. Validating Configuration

This example runs the **rhevmanage-domains** command to validate the Red Hat Enterprise Virtualization Manager configuration. The command attempts to log into each listed domain with the credentials provided in the configuration. If the attempt is successful then the domain is reported as valid.

```
# rhevm-manage-domains -action=validate
User guid is: 80b71bae-98a1-11e0-8f20-525400866c73
Domain directory.demo.redhat.com is valid.
```

Example A.5. Listing Domains in Configuration

This example runs the **rhevmanage-domains** command to list the domains defined in the Red Hat Enterprise Virtualization Manager configuration. For each configuration entry the command displays the domain, the username — in User Principle Name (UPN) format, and whether the domain is local or remote.

```
# rhevm-manage-domains -action=list
Domain: directory.demo.redhat.com
User name: admin@DIRECTORY.DEMO.REDHAT.COM
This domain is a remote domain.
```

A.2. Configuration Tool

During installation, only a subset of Red Hat Enterprise Virtualization Manager's configuration settings are modified from their defaults. You make further changes with the included configuration tool, **rhevmanage-config**.

The configuration tool does not require JBoss or Red Hat Enterprise Virtualization Manager to be running to update the configuration. Configuration key values are stored in the database and as such it must be operational for configuration changes to be saved. Changes are only applied once JBoss is restarted.

The manager's configuration is stored as a series of key to value pair mappings. The configuration tool allows you to:

- list all available configuration keys,
- list all available configuration values,
- get the value of a specific configuration key, and
- set the value of a specific configuration key.

The configuration tool also allows you to maintain multiple versions of the manager's configuration. When getting or setting the value for a configuration key the `--cver` parameter is used to specify which configuration version is to be used. The default configuration version is **general**.

A.2.1. Syntax

You will find the configuration tool on the machine to which Red Hat Enterprise Virtualization Manager was installed. Common usage examples are discussed within this guide. For full usage information consult the **rhev-config** command's help output:

```
# rhev-config --help
```

Common Tasks

List Available Configuration Keys

Use the `--list` parameter to list available configuration keys.

```
# rhev-config --list
```

The tool lists each available configuration key by name. It also returns a description of each key's purpose.

List Available Configuration Values

Use the `--all` parameter to list available configuration values.

```
# rhev-config --all
```

The tool lists each available configuration key by name as well as the current value of the key, and the configuration version.

Get Value of Configuration Key

Use the `--get` parameter to get the value of a specific key.

```
# rhev-config --get KEY_NAME
```

Replace `KEY_NAME` with the name of the key to retrieve. The tool returns the key name, value, and the configuration version. Optionally the `--cver` parameter is used to specify the configuration value from which the value should be retrieved.

Set Value of Configuration Key

Use the `--set` parameter to set the value of a specific key. You must also set the configuration version to which the change is to apply using the `--cver` parameter.

```
# rhvm-config --set KEY_NAME=KEY_VALUE --cver=VERSION
```

Replace `KEY_NAME` with the name of the key to set, and replace `KEY_VALUE` with the value to assign to it. In an environment with more than one configuration version you must also take care to replace `VERSION` with the name of the configuration version in use.

A.2.2. Examples

Example A.6. Getting a Configuration Value

```
# rhvm-config --get=SearchResultsLimit --cver=general
100
```

Example A.7. Setting a Configuration Value

```
# rhvm-config --set SearchResultsLimit=50 --cver=general
```

A.3. Log Collector

The Red Hat Enterprise Virtualization Manager installation includes a log collection tool. This allows you to easily collect relevant logs from across the Red Hat Enterprise Virtualization environment when requesting support.

The log collection command is **rhvm-log-collector**. You must be logged in as the root user to run it successfully. You must provide the administration credentials for the Red Hat Enterprise Virtualization environment on the command line. Full usage information, including a list of all valid options for the command, is available by running the **rhvm-log-collector -h** command.

A.3.1. Syntax

The basic syntax is of the form:

```
Usage: rhvm-log-collector [options] list [all, clusters, datacenters]
       rhvm-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 Enterprise Virtualization Manager. You are then able to filter 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 **rhvm-log-collector** command outputs the specific filename that it chose to use when log collection is completed.

The default action taken if no parameters are provided is to list available hosts along with the data center, and cluster, to which they belong. Where necessary the log collector will prompt you to enter usernames and passwords required to retrieve logs.

The **rhevml-log-collector** command has a large number of options. You can use these options to further refine the scope of log collection.

General Options

--version

Displays the version number of the command in use, and exits immediately.

-h, --help

Displays command usage information, and exits immediately.

--conf-file=PATH

Sets *PATH* as the configuration file the tool is to use.

--local-tmp=PATH

Sets *PATH* as the directory to which retrieved logs are to be saved. Default 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.

--quiet

Sets quiet mode, reducing console output to a minimum. This is off by default.

--log-file=PATH

Sets *PATH* as the log file the command should use for its own log output. Note that this is not to be confused with the **--local-tmp** parameter.

-v, --verbose

Sets verbose mode, providing more console output. This is off by default.

Red Hat Enterprise Virtualization Manager Options

The options in the Red Hat Enterprise Virtualization Manager configuration group are used to specify the manager authentication details and, filter log collection from one or more virtualization hosts. Note that it is possible to combine the options used to select the virtualization hosts, for example selecting all host in clusters **A** and **B** where the name of the host matches pattern **SalesHost***.

--no-hypervisors

Sets the option to skip collection of logs from the virtualization hosts.

-u USER, --user=USER

Sets the username to log in as to *USER*. This must be a username that exists in directory services, and is known to the Red Hat Enterprise Virtualization Manager. The user must be specified in the format *user@domain*, where *user* is replaced by the username, and *domain* is replaced by the directory services domain in use.

-r FQDN, --rhevml=FQDN

Sets the Red Hat Enterprise Virtualization Manager to connect to as *FQDN*. *FQDN* must be replaced by the fully qualified domain name of the manager. By default it is assumed that the log collector is

being run on the same machine as the manager. Therefore the default value for this parameter is **localhost**.

-c CLUSTER, --cluster CLUSTER

Collect all logs from the Red Hat Enterprise Virtualization Manager, as well as virtualization hosts in the cluster named *CLUSTER*. The cluster(s) for inclusion must be specified in a comma separated list of cluster names or match patterns.

-d DATACENTER, --data-center DATACENTER

Collect all logs from the Red Hat Enterprise Virtualization Manager, as well as virtualization hosts in the data center named *DATACENTER*. The data center(s) for inclusion must be specified as a comma separated list of data center names or match patterns.

-H HOSTS_LIST, --hosts=HOSTS_LIST

Collect all logs from the Red Hat Enterprise Virtualization Manager, as well as virtualization hosts included in *HOSTS_LIST*. The hosts for inclusion must be specified as a comma separated list of hostnames, fully qualified domain names, or IP addresses. Match patterns for each type of value are also valid.

SOS Report Options

The JBoss SOS plugin is always executed by log collector. To activate data collection from the JMX console the *--java-home*, *--jboss-user*, and *jboss-pass* parameters must also be provided.

--jboss-home=JBASS_HOME

JBoss installation directory path. Default is **/var/lib/jbossas**.

--java-home=JAVA_HOME

Java installation directory path. Default is **/usr/lib/jvm/java**.

--jboss-profile=JBASS_PROFILE

Quoted and space separated list of server profiles. This is used to limit log collection to the specified profiles. The default is *'rhevm-slimmed'*.

--enable-jmx

Enable the collection of run-time metrics from Red Hat Enterprise Virtualization's JBoss JMX interface.

--jboss-user=JBASS_USER

JBoss JMX invoker user to be used with twiddle. Default is *admin*.

--jboss-logsize=LOG_SIZE

Maximum size for each log file retrieved, in MB.

--jboss-stdjar=STATE

Sets collection of JAR statistics for JBoss standard JARs. Replace *STATE* with **on**, or **off**. The default is **on**.

--jboss-servjar=STATE

Sets collection of JAR statistics from any server configuration directories. Replace *STATE* with **on**, or **off**. The default is **on**.

--jboss-twiddle=STATE

Sets collection of twiddle data on, or off. Twiddle is the JBoss tool used to collect data from the JMX invoker. Replace *STATE* with **on**, or **off**. The default is **on**.

Appendix A. Utilities

`--jboss-appxml=XML_LIST`

Quoted and space separated list of applications whose XML descriptions should be retrieved. Default is `'all'`.

SSH Configuration

`--ssh-host=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 log collector connects to the Red Hat Enterprise Virtualization Manager database and dumps it for inclusion in the log report if `pg-pass` is specified. The database username, and database name also must be specified if they were changed from the default values during installation.

Where the database is not on the local machine set the `pg-dbhost`, and optionally supply a `pg-host-key`, 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. Database collection is performed by default.

`--pg-user=USER`

Sets `USER` as the username 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 `rhev`.

`--pg-dbhost=DBHOST`

Sets `DBHOST` as the hostname 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 as it is not required where the database exists on the local host.

A.3.2. Examples

Example A.8. Basic Log Collector Usage

In this example log collector is run to collect all logs from the Red Hat Enterprise Virtualization Manager and the three attached hosts. Additionally the database and JBoss logs are also collected.

```
# rhvm-log-collector
Please provide the username for rhvm (CTRL+D to abort): admin@directory.demo.redhat.com
Please provide the password for rhvm (CTRL+D to abort):
Host list (datacenter=None, cluster=None, host=None):
Data Center      | Cluster          | Hostname/IP Address
SalesDataCenter | SalesCluster    | 192.168.122.250
```

```

EngineeringDataCenter | EngineeringCluster | 192.168.122.251
FinanceDataCenter    | FinanceCluster      | 192.168.122.252
# rhvm-log-collector collect
Please provide the username for rhvm (CTRL+D to abort): admin@directory.demo.redhat.com
Please provide the password for rhvm (CTRL+D to abort):
About to collect information from 3 hypervisors. Continue? (Y/n): Y
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
Please provide the password to dump the PostgreSQL database (CTRL+D to abort):
INFO: Gathering PostgreSQL the RHEV-M database and log files from localhost...
INFO: Gathering RHEV-M information...
Please provide the password for jboss (CTRL+D to abort):
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

```

A.4. USB Filter Editor

The USB Filter Editor is a Windows tool used to configure a policy file named **usbfilter.txt**. The policy rules defined in this file allow, or deny, the passthrough of specific USB devices from client machines to virtual machines managed using Red Hat Enterprise Virtualization Manager. Once configured the policy file resides on the Red Hat Enterprise Virtualization Manager in the following location:

```

/usr/share/rhvm/rhvm.ear/userportal.war/
org.ovirt.engine.ui.userportal.UserPortal/consoles/spice/usbfilter.txt

```

Changes to the USB filter policies take effect the next time the `jbosscas` service on the Red Hat Enterprise Virtualization Manager server is restarted. To obtain the USB Filter Editor, download the **USBFilterEditor.msi** file from Red Hat Network, under the **Red Hat Enterprise Virtualization Manager (v.3 x86_64)** channel.

To install the USB Filter Editor

1. On a Windows machine, launch the **USBFilterEditor.msi** installer that was obtained from Red Hat Network.
2. You will be guided through a series of steps to install the USB Filter Editor on your machine. If you do not specify an alternative location, 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 the version of Windows in use.
3. When the installation completes, a shortcut icon is created on your desktop. This is how you launch the USB Filter Editor.



Important — Secure Copy Utility

To import and export filter policies from the Red Hat Enterprise Virtualization Manager it is necessary to use a Secure Copy (SCP) client. A Secure Copy tool for Windows machines is WinSCP (<http://winscp.net>).

A.4.1. Updating the USB Device Policy

The default USB device policy only allows virtual machines access to a limited range of USB devices. To allow the use of additional USB devices you must update the policy.

1. Click the USB Filter Editor icon on your desktop. The **Red Hat USB Filter Editor** displays the current USB policies.

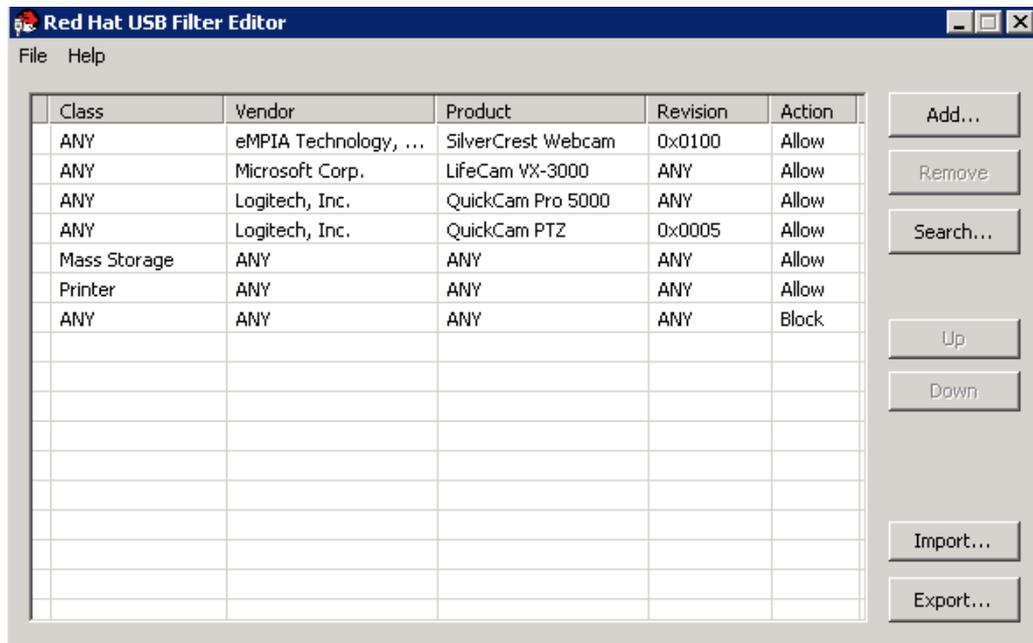


Figure A.1. Red Hat USB Filter Editor

2. For each USB device, the **Class**, **Vendor**, **Product**, **Revision** and **Action** displays. The permitted devices display with an **Allow** action, the blocked devices display with a **Block** action.
3. You can **Add** to, and **Remove** from, the list of devices that the policy allows virtual machines access to. The USB device policy rules are processed in the order in which they are listed. The **Up** and **Down** buttons enable you to move devices higher or lower in the list. The last rule in the list must always be the rule with action **Block** for any devices. This ensures that all devices that are not explicitly allowed in the policy are blocked.

The **Search** button allows you to create policy rules based on the devices attached to the system the USB Filter Editor is being run on. Additionally the **Import**, and **Export** buttons allow you to move USB device policy files to and from the Red Hat Enterprise Virtualization Manager.

Table A.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.

A.4.1.1. Adding a USB Policy

Double click the USB Filter Editor icon on your desktop. The **Red Hat USB Filter Editor** displays the current USB policies. You can specify the USB device and whether virtual machines can use them by adding a new policy.

To add a new policy, on the Red Hat USB Editor:

1. Click the **Add** button. The **Edit USB Criteria** dialog displays:

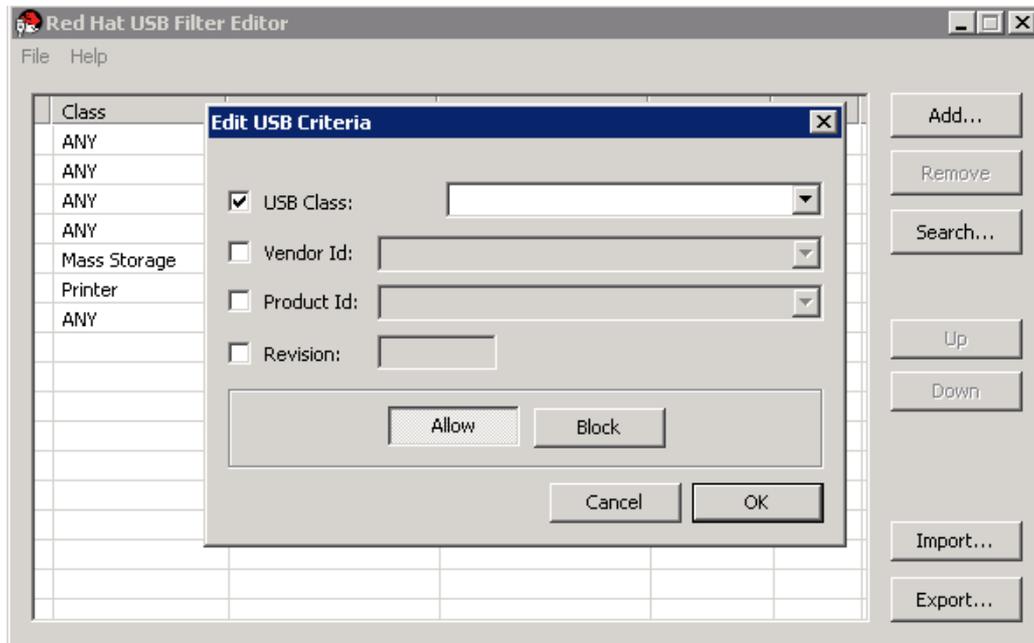


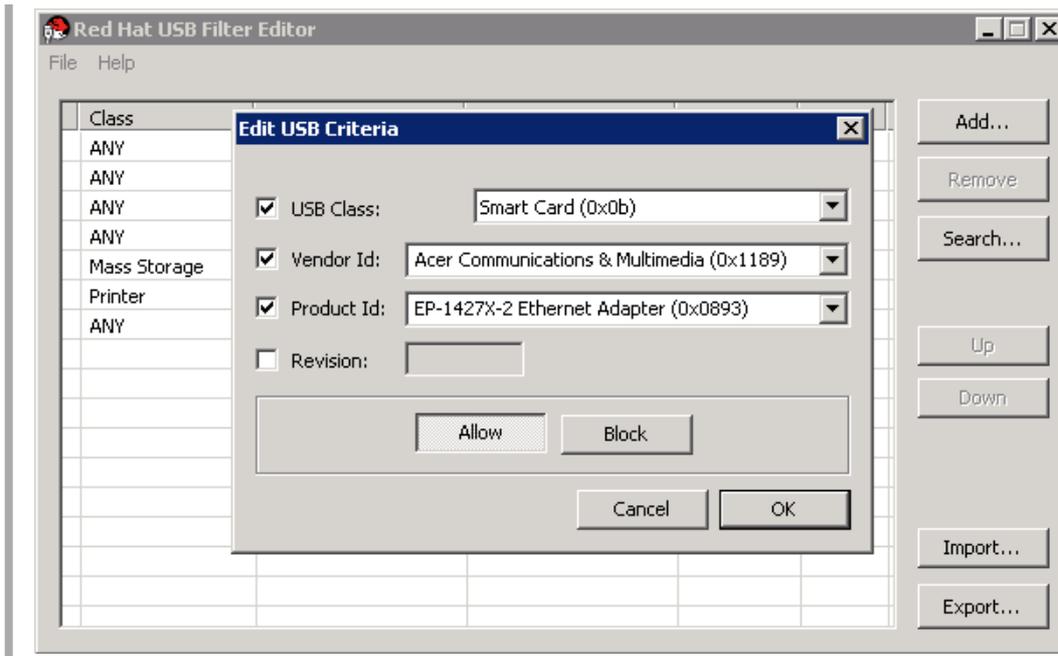
Figure A.2. Edit USB Criteria

2. Add any combination of USB devices, products, and vendors using the **USB Class**, **Vendor ID**, **Product ID**, and **Revision** check boxes and lists.

To allow virtual machines to use the specified USB device, click the **Allow** button. Or, to block virtual machines from using the specified USB device, click the **Block** button. Click **OK** to exit the dialog box. This action adds the selected filter rule to the list.

Example A.9. Adding a Device

The following is an example of how to add USB Class **Smartcard**, device **EP-1427X-2 Ethernet Adaptor**, from manufacturer **Acer Communications & Multimedia** to the list of allowed devices.



3. Click the **Save** entry in the **File** menu to save the changes.
4. For USB filter policy changes to take effect they must also be exported to the Red Hat Enterprise Virtualization Manager, see [Section A.4.2, “Export a USB Policy”](#).

A.4.1.2. Removing a USB Policy

Double click the USB Filter Editor icon on your desktop. The **Red Hat USB Filter Editor** displays the current USB policies.

To remove a policy, on the Red Hat USB Editor:

1. Select the policy to be removed.

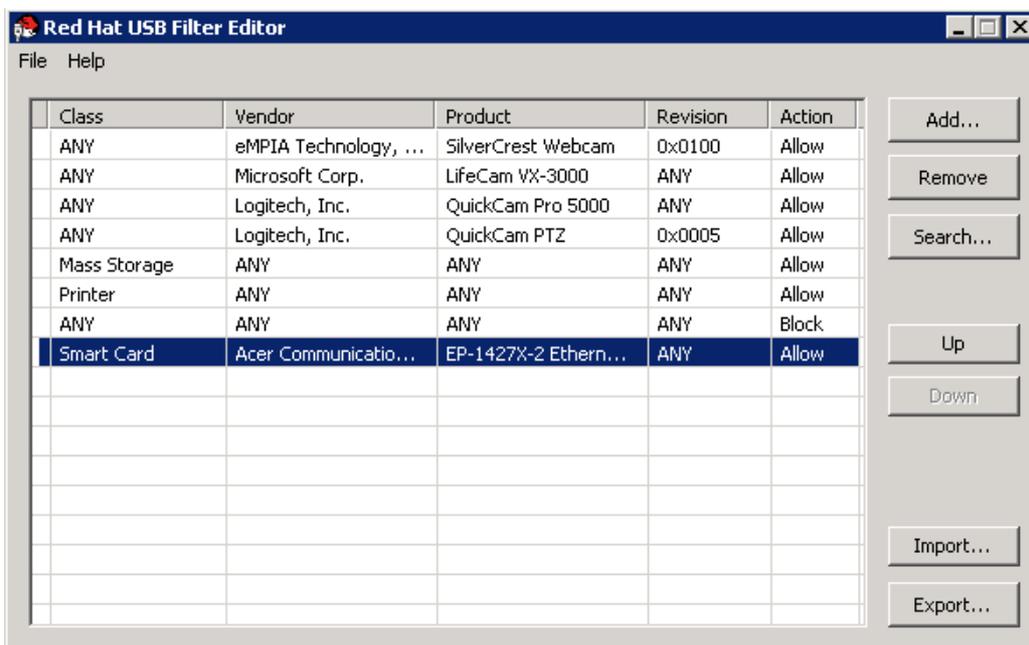


Figure A.3. Select USB Policy

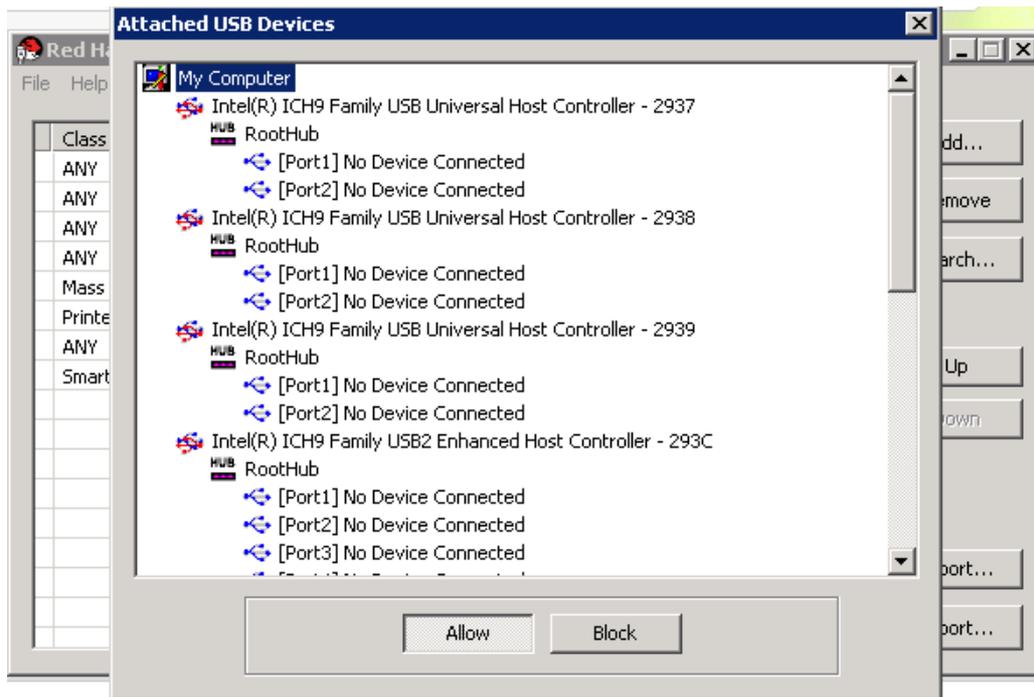


Figure A.5. Attached USB Devices

3. Select the device and click the **Allow** or **Block** button as appropriate. Double click the selected device to close the dialog box. A policy rule for the device is added to the list.
4. Use the **Up** and **Down** buttons to change the position the new policy rule in the list.
5. Click the **Save** entry in the **File** menu to save the changes.
6. For USB filter policy changes to take effect they must also be exported to the Red Hat Enterprise Virtualization Manager, see [Section A.4.2, “Export a USB Policy”](#).

A.4.2. Export a USB Policy

For USB device policy changes to take effect the updated policy must be exported and uploaded to the Red Hat Enterprise Virtualization Manager. Once the policy has been uploaded the `jbossas` service must be restarted.

1. Double click the USB Filter Editor icon on your desktop. The **Red Hat USB Filter Editor** displays the current USB policies.
2. Click the **Export** button. The **Save As** dialog displays.
3. Save the file with a filename of `usbfilter.txt`.
4. Using a Secure Copy client, such as WinSCP, upload the `usbfilter.txt` file to the server running Red Hat Enterprise Virtualization Manager. The file must be placed in the following directory on the server:

```
/usr/share/rhev/rhev.ear/userportal.war/  
org.ovirt.engine.ui.userportal.UserPortal/soles/spice/
```

5. As the root user on the server running Red Hat Enterprise Virtualization Manager restart the `jbossas` service.

```
# service jbossas restart
```

The USB Device policy will now be implemented on virtual machines running in the Red Hat Enterprise Virtualization environment.

A.4.3. Import USB Policy

If a USB device policy already exists on your Red Hat Enterprise Virtualization Manager then to make changes you must first download it and import it into the USB filter editor.

1. Using a Secure Copy client, such as WinSCP, download the **usbfilter.txt** file from the server running Red Hat Enterprise Virtualization Manager. The file exists in the following directory on the server:

```
/usr/share/rhev/rhev.ear/userportal.war/  
org.ovirt.engine.ui.userportal.UserPortal/consoles/spice/
```

2. Double click the USB Filter Editor icon on your desktop. The **Red Hat USB Filter Editor** displays.
3. Click the **Import** button. The **Open** dialog displays
4. Open the **usbfilter.txt** file that was downloaded from the server. The USB device policy will be presented in the USB Filter Editor for editing.

Appendix B. Configuring Red Hat Enterprise Linux 5.4 or Higher Virtual Machines to Use SPICE

SPICE is a remote display protocol designed for virtual environments, which enables you to view a virtualized desktop or server. SPICE delivers a high quality user experience, keeps CPU consumption low, and supports high quality video streaming.

Using SPICE on a Linux machine significantly improves the movement of the mouse cursor on the console of the virtual machine. To use SPICE, the X-Windows system requires additional qxl drivers. The qxl drivers are provided with Red Hat Enterprise Linux 5.4 and higher. Older versions are not supported. Installing SPICE on a virtual machine running Red Hat Enterprise Linux significantly improves the performance of the GUI.



Note

Typically, this is most useful for virtual desktops where the user requires the use of the GUI. System administrators who are creating virtual servers may prefer not to configure SPICE if their use of the GUI is minimal.

Procedure B.1. How to install the qxl drivers

1. Log in to the Red Hat Enterprise Linux Virtual Machine.
2. Open a terminal.
3. Run

```
# yum install xorg-x11-drv-qxl
```

This installs the qxl drivers. They must now be configured.

How to configure the qxl drivers

There are two ways to configure the qxl drivers. Perform only one of the following procedures:

- **Procedure B.2. How to configure qxl drivers in GNOME**

1. Click **System**.
2. Click **Administration**.
3. Click **Display**.
4. Click the **Hardware** tab.
 - Click **Video Cards Configure**.
5. Select **qxl** and click **OK**.
6. Restart X-Windows by logging out of the virtual machine and logging back in.

Appendix B. Configuring Red Hat Enterprise Linux 5.4 or Higher Virtual Machines to Use SPICE

- **Procedure B.3. How to configure qxl drivers on the command line:**

1. Back up `/etc/X11/xorg.conf`:

```
# cp /etc/X11/xorg.conf /etc/X11/xorg.conf.$$backup
```

2. Make the following change to the Device section of `/etc/X11/xorg.conf`:

```
Section "Device"
Identifier "Videocard0"
Driver "qxl"
Endsection
```

- **Procedure B.4. Configuring the tablet and mouse to use SPICE**

1. Verify that the tablet device is available on your guest:

```
# /sbin/lscusb -v | grep 'QEMU USB Tablet'
```

If there is no output from the command, do not continue configuring the tablet.

2. Back up `/etc/X11/xorg.conf` by running this command:

```
# cp /etc/X11/xorg.conf /etc/X11/xorg.conf.$$backup
```

3. Make the following changes to `/etc/X11/xorg.conf`:

```
Section "ServerLayout"
Identifier      "single head configuration"
Screen        0  "Screen0" 0 0
InputDevice   "Keyboard0" "CoreKeyboard"
InputDevice   "Tablet" "SendCoreEvents"
InputDevice   "Mouse" "CorePointer"
EndSection

Section "InputDevice"
Identifier     "Mouse"
Driver        "void"
#Option       "Device" "/dev/input/mice"
#Option       "Emulate3Buttons" "yes"
EndSection

Section "InputDevice"
Identifier     "Tablet"
Driver        "evdev"
Option        "Device" "/dev/input/event2"
Option "CorePointer" "true"
EndSection
```

4. Log out and log back into the virtual machine to restart X-Windows.

Appendix C. Changing Passwords on the Red Hat Enterprise Virtualization Manager

This appendix describes how to change passwords for the administrative users for Administration Portal and Red Hat Enterprise Virtualization Manager's PostgreSQL database respectively.

C.1. Changing Password for the Administrative User

The **admin@internal** user account is automatically created when you installed and configured Red Hat Enterprise Virtualization Manager. This account is stored locally in Red Hat Enterprise Virtualization Manager's PostgreSQL database, and exists separately from other directory services. Unlike the IPA or Active Directory domains, users cannot be added or deleted from the **internal** domain. The **admin@internal** user is the SuperUser of the Red Hat Enterprise Virtualization Manager, and has administrative privileges over the environment via the Administration Portal.

During installation, you were prompted for a password for the **admin@internal** user. However, if you have forgotten the password, or choose to reset the password, you can use the **rhev-config** utility.

Procedure C.1. To reset the password for the admin@internal user

1. Log in to the Red Hat Enterprise Virtualization Manager server as the **root** user.
2. Use the **rhev-config** utility to set a new password for the **admin@internal** user. Run the following command:

```
# rhev-config -s AdminPassword=<your_password>
```

You do not need to use quotes, but use escape shell characters if you include them in the password.

3. Restart the JBoss service for the changes to take effect.

```
# service jbossas restart
```

Appendix D. Reports Schema

D.1. Configuration History Views

This section describes the configuration views available to the user for querying and generating reports.



delete_date does not appear for living entities

delete_date does not appear in latest views because these views provide the latest configuration of living entities, which, by definition, have not been deleted.

D.1.1. v3_0_datacenter_configuration_view lv3_0_latest_datacenter_configuration_view

Data centers configuration history in the system.

Table D.1. v3_0_datacenter_configuration_view\lv3_0_latest_datacenter_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
datacenter_id	uuid	The unique ID of the data center in the system.
datacenter_name	varchar(40)	Name of the data center, as displayed in the edit dialog.
datacenter_description	varchar(4000)	Description of the data center, as displayed in the edit dialog.
storage_type	smallint	<ul style="list-style-type: none">• 0 - Unknown• 1 - NFS• 2 - FCP• 3 - iSCSI• 4 - Local• 6 - All
create_date	timestamp with time zone	The date this entity was added to the system.
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.1.2. v3_0_datacenter_storage_domain_map_view lv3_0_latest_datacenter_configuration_view

A historical map showing the relationships between storage domains and data centers in the system.

Table D.2. v3_0_datacenter_storage_domain_map_view\lv3_0_latest_datacenter_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
storage_domain_id	uuid	The unique ID of this storage domain in the system.
datacenter_id	uuid	The unique ID of the data center in the system.
attach_date	timestamp with time zone	The date the storage domain was attached to the data center.
detach_date	timestamp with time zone	The date the storage domain was detached from the data center.

D.1.3. v3_0_storage_domain_configuration_view lv3_0_latest_storage_domain_configuration_view

Storage domains configuration history in the system.

Table D.3. v3_0_storage_domain_configuration_view
lv3_0_latest_storage_domain_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
storage_domain_id	uuid	The unique ID of this storage domain in the system.
storage_domain_name	varchar(250)	Storage domain name.
storage_domain_type	smallint	<ul style="list-style-type: none"> • 0 - Data (Master) • 1 - Data • 2 - ISO • 3 - Export
storage_type	smallint	<ul style="list-style-type: none"> • 0 - Unknown • 1 - NFS • 2 - FCP • 3 - iSCSI • 4 - Local • 6 - All

Name	Type	Description
create_date	timestamp with time zone	The date this entity was added to the system.
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.1.4. v3_0_cluster_configuration_view v3_0_latest_cluster_configuration_view

Clusters configuration history in the system.

Table D.4. v3_0_cluster_configuration_view\v3_0_latest_cluster_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
cluster_id	uuid	The unique identifier of the datacenter this cluster resides in.
cluster_name	varchar(40)	Name of the cluster, as displayed in the edit dialog.
cluster_description	varchar(4000)	As defined in the edit dialog.
datacenter_id	uuid	The unique identifier of the datacenter this cluster resides in.
cpu_name	varchar(255)	As displayed in the edit dialog.
compatibility_version	varchar(40)	As displayed in the edit dialog.
datacenter_configuration_version	integer	The data center configuration version at the time of creation or update.
create_date	timestamp with time zone	The date this entity was added to the system.
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.1.5. v3_0_host_configuration_view v3_0_latest_host_configuration_view

Host configuration history in the system.

Table D.5. v3_0_host_configuration_view\v3_0_latest_host_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.

Appendix D. Reports Schema

Name	Type	Description
host_id	uuid	The unique ID of the host in the system.
host_unique_id	varchar(128)	This field is a combination of the host physical UUID and one of its MAC addresses, and is used to detect hosts already registered in the system.
host_name	varchar(255)	Name of the host (same as in the edit dialog).
cluster_id	uuid	The unique ID of the cluster that this host belongs to.
host_type	smallint	<ul style="list-style-type: none"> • 0 - RHEL Host • 2 - RHEV Hypervisor Node
fqn_or_ip	varchar(255)	The host's DNS name or its IP address for Red Hat Enterprise Virtualization Manager to communicate with (as displayed in the edit dialog).
memory_size_mb	integer	The host's physical memory capacity, expressed in megabytes (MB).
swap_size_mb	integer	The host swap partition size.
cpu_model	varchar(255)	The host's CPU model.
number_of_cores	smallint	Total number of CPU cores in the host.
host_os	varchar(255)	The host's operating system version.
pm_ip_address	varchar(255)	Power Management server IP address.
kernel_version	varchar(255)	The host's kernel version.
kvm_version	varchar(255)	The host's KVM version.
vdsm_version	varchar(40)	The host's VDSM version.
vdsm_port	integer	As displayed in the edit dialog.
cluster_configuration_version	integer	The cluster configuration version at the time of creation or update.
create_date	timestamp with time zone	The date this entity was added to the system.
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.1.6. v3_0_host_configuration_view v3_0_latest_host_interface_configuration_view

Host interface configuration history in the system.

Table D.6. v3_0_host_configuration_view\v3_0_latest_host_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
host_interface_id	uuid	The unique ID of this interface in the system.
host_interface_name	varchar(50)	The interface name as reported by the host.
host_id	uuid	Unique ID of the host this interface belongs to.
host_interface_type	smallint	<ul style="list-style-type: none"> • 0 - rt18139_pv • 1 - rt18139 • 2 - e1000 • 3 - pv
host_interface_speed_bps	integer	The interface speed in bits per second.
mac_address	varchar(20)	The interface MAC address.
network_name	varchar(50)	The logical network associated with the interface.
ip_address	varchar(50)	As displayed in the edit dialog.
gateway	varchar(20)	As displayed in the edit dialog.
bond	Boolean	A flag to indicate if this interface is a bonded interface.
bond_name	varchar(50)	The name of the bond this interface is part of (if it is part of a bond).
vlan_id	integer	As displayed in the edit dialog.
host_configuration_version	integer	The host configuration version at the time of creation or update.
create_date	timestamp with time zone	The date this entity was added to the system.
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.1.7. v3_0_vm_configuration_view lv3_0_latest_vm_configuration_view

A list of all virtual machines in the system.

Table D.7. v3_0_vm_configuration_view/v3_0_latest_vm_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
vm_id	uuid	The unique ID of this VM in the system.
vm_name	varchar(255)	The name of the VM.
vm_description	varchar(4000)	As displayed in the edit dialog.
vm_type	smallint	<ul style="list-style-type: none"> • 0 - Desktop • 1 - Server
cluster_id	uuid	The unique ID of the cluster this VM belongs to.
template_id	uuid	The unique ID of the template this VM is derived from. The field is for future use, as the templates are not synchronized to the history database in this version.
template_name	varchar(40)	Name of the template from which this VM is derived.
cpu_per_socket	smallint	Virtual CPUs per socket.
number_of_sockets	smallint	Total number of virtual CPU sockets.
memory_size_mb	integer	Total memory allocated to the VM, expressed in megabytes (MB).
operating_system	smallint	<ul style="list-style-type: none"> • 0 - Unknown • 1 - Windows XP • 3 - Windows 2003 • 4 - Windows 2008 • 5 - Other Linux • 6 - Other • 7 - RHEL 5 • 8 - RHEL 4 • 9 - RHEL 3 • 10 - Windows2003 x64

Name	Type	Description
		<ul style="list-style-type: none"> • 11 - Windows 7 • 12 - Windows 7 x64 • 13 - RHEL 5 x64 • 14 - RHEL 4 x64 • 15 - RHEL 3 x64 • 16 - Windows 2008 x64 • 17 - Windows 2008R2 x64 • 18 - RHEL 6 • 19 - RHEL 6 x64
ad_domain	varchar(40)	As displayed in the edit dialog.
default_host	uuid	As displayed in the edit dialog, the ID of the default host in the system.
high_availability	Boolean	As displayed in the edit dialog.
initialized	Boolean	A flag to indicate if this VM was started at least once for Sysprep initialization purposes.
stateless	Boolean	As displayed in the edit dialog.
fail_back	Boolean	As displayed in the edit dialog.
auto_suspend	Boolean	As displayed in the edit dialog.
usb_policy	smallint	As displayed in the edit dialog.
time_zone	varchar(40)	As displayed in the edit dialog.
cluster_configuration_version	integer	The cluster configuration version at the time of creation or update.
default_host_configuration_version	integer	The host configuration version at the time of creation or update.
create_date	timestamp with time zone	The date this entity was added to the system.
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.1.8. v3_0_vm_configuration_view latest_vm_interface_configuration_view

Virtual interfaces configuration history in the system

Table D.8. v3_0_vm_configuration_view\latest_vm_interface_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
vm_interface_id	uuid	The unique ID of this interface in the system.
vm_interface_name	varchar(50)	As displayed in the edit dialog.
vm_id	uuid	The ID of the virtual machine this interface belongs to.
vm_interface_type	smallint	The type of the virtual interface. <ul style="list-style-type: none"> • 0 - rt18139_pv • 1 - rt18139 • 2 - e1000 • 3 - pv
vm_interface_speed_bps	integer	The average speed of the interface during the aggregation in bits per second.
mac_address	varchar(20)	As displayed in the edit dialog.
network_name	varchar(50)	As displayed in the edit dialog.
vm_configuration_version	integer	The virtual machine configuration version at the time of creation or update.
create_date	timestamp with time zone	The date this entity was added to the system.
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.1.9. v3_0_disks_vm_map_view\v3_0_latest_disks_vm_map_view

A historical map showing the relationships between virtual disks and virtual machines in the system.

Table D.9. v3_0_disks_vm_map_view\v3_0_latest_disks_vm_map_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
vm_disk_id	uuid	The unique ID of this virtual disk in the system.
vm_id	uuid	The unique ID of the virtual machine in the system.
attach_date	timestamp with time zone	The date the virtual disk was attached to the virtual machine.

Name	Type	Description
detach_date	timestamp with time zone	The date the virtual disk was detached from the virtual machine.

D.1.10. v3_0_vm_disk_configuration_view lv3_0_latest_vm_disk_configuration_view

Virtual disks configuration history in the system.

Table D.10. v3_0_vm_disk_configuration_view\v3_0_latest_vm_disk_configuration_view

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
vm_disk_id	uuid	The unique ID of this disk in the system.
storage_domain_id	uuid	The ID of the storage domain this disk image belongs to.
vm_internal_drive_mapping	varchar	The virtual machine internal drive mapping.
vm_disk_description	varchar(4000)	As displayed in the edit dialog.
vm_disk_space_size_mb	integer	The defined size of the disk in megabytes (MB).
disk_type	integer	As displayed in the edit dialog. Only System and data are currently used. <ul style="list-style-type: none"> • 0 - Unassigned • 1 - System • 2 - Data • 3 - Shared • 4 - Swap • 5 - Temp
vm_disk_format	integer	As displayed in the edit dialog. <ul style="list-style-type: none"> • 3 - Unassigned • 4 - COW • 5 - RAW
vm_disk_interface	integer	<ul style="list-style-type: none"> • 0 - IDE • 1 - SCSI (not supported) • 2 - VirtIO
create_date	timestamp with time zone	The date this entity was added to the system.

Name	Type	Description
update_date	timestamp with time zone	The date this entity was changed in the system.
delete_date	timestamp with time zone	The date this entity was deleted from the system.

D.2. Statistics History Views

This section describes the statistics history views available to the user for querying and generating reports.

D.2.1. v3_0_datacenter_samples_history_view

lv3_0_datacenter_hourly_history_view

lv3_0_datacenter_daily_history_view

Historical statistics for each data center in the system.

Table D.11. v3_0_datacenter_samples_history_view\lv3_0_datacenter_hourly_history_view\lv3_0_datacenter_daily_history_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
history_datetime	timestamp with time zone	The timestamp of this history row (rounded to minute, hour, day as per the aggregation level).
datacenter_id	uuid	The unique ID of the data center.
datacenter_status	smallint	<ul style="list-style-type: none"> -1 - Unknown Status (used only to indicate a problem with the ETL -- PLEASE NOTIFY SUPPORT) 1 - Up 2 - Maintenance 3 - Problematic
minutes_in_status	decimal	The total number of minutes that the data center was in the status shown in the datacenter_status column for the aggregation period. For example, if a data center was up for 55 minutes and in maintenance mode for 5 minutes during an hour, two rows will show for this hour. One will have a datacenter_status of Up and minutes_in_status

Name	Type	Description
		of 55, the other will have a datacenter_status of Maintenance and a minutes_in_status of 5.
datacenter_configuration_version	integer	The data center configuration version at the time of sample.

D.2.2. v3_0_storage_domain_samples_history_view-v3_0_storage_domain_hourly_history_view\lv3_0_storage_domain_daily_history_view

Historical statistics for each storage domain in the system.

Table D.12. v3_0_storage_domain_samples_history_view\lv3_0_storage_domain_hourly_history_view-v3_0_storage_domain_daily_history_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
history_datetime	timestamp with time zone	The timestamp of this history row (rounded to minute, hour, day as per the aggregation level).
storage_domain_id	uuid	Unique ID of the storage domain in the system.
available_disk_size_gb	integer	The total available (unused) capacity on the disk, expressed in gigabytes (GB).
used_disk_size_gb	integer	The total used capacity on the disk, expressed in gigabytes (GB).
storage_configuration_version	integer	The storage domain configuration version at the time of sample.

D.2.3. v3_0_host_samples_history_view\lv3_0_host_hourly_history_view\lv3_0_host_daily_history_view

Historical statistics for each host in the system.

Table D.13. v3_0_host_samples_history_view\lv3_0_host_hourly_history_view\lv3_0_host_daily_history_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
history_datetime	timestamp with time zone	The timestamp of this history row (rounded to minute, hour, day as per the aggregation level).

Appendix D. Reports Schema

Name	Type	Description
host_id	uuid	Unique ID of the host in the system.
host_status	smallint	<ul style="list-style-type: none"> -1 - Unknown Status (used only to indicate a problem with the ETL -- PLEASE NOTIFY SUPPORT) 1 - Up 2 - Maintenance 3 - Problematic
minutes_in_status	decimal	The total number of minutes that the host was in the status shown in the status column for the aggregation period. For example, if a host was up for 55 minutes and down for 5 minutes during an hour, two rows will show for this hour. One will have a status of Up and minutes_in_status of 55, the other will have a status of down and a minutes_in_status of 5.
memory_usage_percent	smallint	Percentage of used memory on the host.
max_memory_usage	smallint	Percentage of used memory on the host.
cpu_usage_percent	smallint	Used CPU percentage on the host.
max_cpu_usage	smallint	The maximum CPU usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
ksm_cpu_percent	smallint	CPU percentage ksm on the host is using.
max_ksm_cpu_percent	smallint	The maximum KSM usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.

Name	Type	Description
active_vms	smallint	The average number of active virtual machines for this aggregation.
max_active_vms	smallint	The maximum active number of virtual machines for the aggregation period. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
total_vms	smallint	The average number of all virtual machines on the host for this aggregation.
max_total_vms	smallint	The maximum total number of virtual machines for the aggregation period. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
total_vms_vcpus	smallint	Total number of VCPUs allocated to the host.
max_total_vms_vcpus	smallint	The maximum total virtual machine VCPU number for the aggregation period. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
cpu_load	smallint	The CPU load of the host.
max_cpu_load	smallint	The maximum CPU load for the aggregation period. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
system_cpu_usage_percent	smallint	Used CPU percentage on the host.
max_cpu_usage_percent	smallint	The maximum system CPU usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it

Appendix D. Reports Schema

Name	Type	Description
		is the maximum hourly average value.
user_cpu_usage_percent	smallint	Used user CPU percentage on the host.
max_user_cpu_usage_percent	smallint	The maximum user CPU usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
swap_used_mb	integer	Used swap size usage of the host in megabytes (MB).
max_swap_used_mb	integer	The maximum user swap size usage of the host for the aggregation period in megabytes (MB), expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
host_configuration_version	integer	The host configuration version at the time of sample.

D.2.4. v3_0_host_interface_samples_history_view

lv3_0_host_interface_hourly_history_view

lv3_0_host_interface_daily_history_view

Historical statistics for each host network interface in the system.

Table D.14. v3_0_host_interface_samples_history_view\lv3_0_host_interface_hourly_history_view
lv3_0_host_interface_daily_history_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
history_datetime	timestamp with time zone	The timestamp of this history view (rounded to minute, hour, day as per the aggregation level).
host_interface_id	uuid	Unique identifier of the interface in the system.
receive_rate_percent	smallint	Used receive rate percentage on the host.
max_receive_rate_percent	smallint	The maximum receive rate for the aggregation period,

Name	Type	Description
		expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
transmit_rate_percent	smallint	Used transmit rate percentage on the host.
max_transmit_rate_percent	smallint	The maximum transmit rate for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
host_interface_configuration_version	integer	The host interface configuration version at the time of sample.

D.2.5. v3_0_vm_samples_history_view v3_0_vm_hourly_history_view\v3_0_vm_daily_history_view

Historical statistics for the virtual machines in the system.

Table D.15. v3_0_vm_samples_history_view\v3_0_vm_hourly_history_view
v3_0_vm_daily_history_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
history_datetime	timestamp with time zone	The timestamp of this history row (rounded to minute, hour, day as per the aggregation level).
vm_id	uuid	Unique ID of the virtual machine in the system.
vm_status	smallint	<ul style="list-style-type: none"> • -1 - Unknown Status (used only to indicate problems with the ETL -- PLEASE NOTIFY SUPPORT) • 0 - Down • 1 - Up • 2 - Paused • 3 - Problematic
minutes_in_status	decimal	The total number of minutes that the virtual machine was in the status shown in the status

Appendix D. Reports Schema

Name	Type	Description
		column for the aggregation period. For example, if a virtual machine was up for 55 minutes and down for 5 minutes during an hour, two rows will show for this hour. One will have a status of Up and minutes_in_status, the other will have a status of Down and a minutes_in_status of 5.
cpu_usage_percent	smallint	The percentage of the CPU in use by the virtual machine.
max_cpu_usage	smallint	The maximum CPU usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
memory_usage_percent	smallint	Percentage of used memory in the virtual machine. The guest tools must be installed on the virtual machine for memory usage to be recorded.
max_memory_usage	smallint	The maximum memory usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value. The guest tools must be installed on the virtual machine for memory usage to be recorded.
user_cpu_usage_percent	smallint	Used user CPU percentage on the host.
max_user_cpu_usage_percent	smallint	The maximum user CPU usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregation, it is the maximum hourly average value.
system_cpu_usage_percent	smallint	Used system CPU percentage on the host.

Name	Type	Description
max_system_cpu_usage_percent	smallint	The maximum system CPU usage for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
vm_ip	varchar(255)	The IP address of the first NIC. Only shown if the guest agent is installed.
current_user_name	varchar(255)	Name of user logged into the virtual machine console, if a guest agent is installed.
currently_running_on_host	uuid	The unique ID of the host the virtual machine is running on.
vm_configuration_version	integer	The virtual machine configuration version at the time of sample.
current_host_configuration_version	integer	The current host the virtual machine is running on.

D.2.6. v3_0_vm_interface_samples_history_view

v3_0_vm_interface_hourly_history_view

v3_0_vm_interface_daily_history_view

Historical statistics for the virtual machine network interfaces in the system.

Table D.16. v3_0_vm_interface_samples_history_view\v3_0_vm_interface_hourly_history_view\v3_0_vm_interface_daily_history_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
history_datetime	timestamp with time zone	The timestamp of this history row (rounded to minute, hour, day as per the aggregation level).
vm_interface_id	uuid	Unique identifier of the interface in the system.
receive_rate_percent	smallint	Used receive rate percentage on the host.
max_receive_rate_percent	smallint	The maximum receive rate for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it

Appendix D. Reports Schema

Name	Type	Description
		is the maximum hourly average value.
transmit_rate_percent	smallint	Used transmit rate percentage on the host.
max_transmit_rate_percent	smallint	The maximum transmit rate for the aggregation period, expressed as a percentage. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average rate.
vm_interface_configuration_version	integer	The virtual machine interface configuration version at the time of sample.

D.2.7. v3_0_vm_disk_daily_history_view

lv3_0_vm_disk_hourly_history_view

lv3_0_vm_disk_samples_history_view

Historical statistics for the virtual disks in the system.

Table D.17. v3_0_vm_disk_daily_history_view\lv3_0_vm_disk_hourly_history_view
lv3_0_vm_disk_samples_history_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
history_datetime	timestamp with time zone	The timestamp of this history row (rounded to minute, hour, day as per the aggregation level).
vm_disk_id	uuid	Unique ID of the disk in the system.
vm_disk_status	integer	<ul style="list-style-type: none"> • 0 - Unassigned • 1 - OK • 2 - Locked • 3 - Invalid • 4 - Illegal
minutes_in_status	decimal	The total number of minutes that the virtual machine disk was in the status shown in the status column for the aggregation period. For example, if a virtual machine disk was locked for 55 minutes and OK for 5 minutes during

Name	Type	Description
		an hour, two rows will show for this hour. One will have a status of Locked and minutes_in_status of 55, the other will have a status of OK and a minutes_in_status of 5.
vm_actual_disk_size_mb	integer	The actual size allocated to the disk.
read_rate_bytes_per_second	integer	Read rate to disk in bytes per second.
max_read_rate_bytes_per_second	integer	The maximum read rate for the aggregation period. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
write_rate_bytes_per_second	integer	Write rate to disk in bytes per second.
max_write_rate_bytes_per_second	integer	The maximum write rate for the aggregation period. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
read_latency_seconds	decimal	The virtual machine disk read latency measured in seconds.
max_read_latency_seconds	decimal	The maximum write latency for the aggregation period, measured in seconds. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
write_latency_seconds	decimal	The virtual machine disk write latency measured in seconds.
max_write_latency_seconds	decimal	The maximum write latency for the aggregation period, measured in seconds. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
flush_latency_seconds	decimal	The virtual machine disk flush latency measured in seconds.

Name	Type	Description
max_flush_latency_seconds	decimal	The maximum flush latency for the aggregation period, measured in seconds. For hourly aggregations, this is the maximum collected sample value. For daily aggregations, it is the maximum hourly average value.
vm_disk_configuration_version	integer	The virtual machine disk configuration version at the time of sample.

D.3. Tag History and ENUM Views

This section describes the tag history and ENUM views available to the user for querying and generating reports. Latest tag views show only living tags relations and the latest details version.



delete_date and detach_date do not appear for living entities

delete_date and **detach_date** do not appear in latest views because these views provide the latest configuration of living entities, which, by definition, have not been deleted.

D.3.1. v3_0_tag_relations_history_view lv3_0_latest_tag_relations_history_view

Tag relations history in the system.

Table D.18. tag_and_ENUM_VIEWS_table_v3_0_tag_relations_history_view
lv3_0_latest_tag_relations_history_view"

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
entity_id	uuid	Unique ID of the entity or tag in the system.
entity_type	smallint	<ul style="list-style-type: none"> • 2 - VM • 3 - Host • 5 - VM pool • 18 - Tag
parent_id	uuid	Unique ID of the entity or tag in the system.
attach_date	timestamp with time zone	The date the entity or tag was attached to the entity or tag.
detach_date	timestamp with time zone	The date the entity or tag was detached from the entity or tag.

D.3.2. v3_0_tag_details_view|v3_0_latest_tag_details_view

Tag details history in the system.

Table D.19. v3_0_tag_details_view|v3_0_latest_tag_details_view

Name	Type	Description
history_id	integer	The unique ID of this row in the table.
tag_id	uuid	Unique ID of the tag in the system.
tag_name	varchar(50)	Name of the tag, as displayed in the tag tree.
tag_description	varchar(4000)	Description of the tag, as displayed in the edit dialog.
tag_path	varchar(4000)	The path to the tag in the tree.
tag_level	smallint	The tag level in the tree.
create_date	timestamp with time zone	The date this tag was added to the system.
update_date	timestamp with time zone	The date this tag was changed in the system.
delete_date	timestamp with time zone	The date this tag was deleted from the system.

D.3.3. v3_0_enum_translator_view

The ENUM table is used to easily translate column numeric types to their meanings and lists ENUM values for columns in the history database.

Table D.20. v3_0_enum_translator_view

Name	Type	Description
enum_type	varchar(40)	The type of ENUM.
enum_key	smallint	The key of the ENUM.
value	varchar(40)	The value of the ENUM.

Appendix E. Search Parameters

This appendix describes in detail how Search works in Red Hat Enterprise Virtualization.

E.1. Searching for Resources

This section specifies, for each resource, the resource's unique set of properties as well as the set of associated resource types.

E.1.1. Searching for Data Centers

The following table describes all search options for Data Centers.

Table E.1. Searching for Data Centers

Property (of resource or resource-type)	Type	Description (Reference)
Clusters.clusters-prop	See property types in Section E.1.2, "Searching for Clusters"	The property of the clusters associated with the data center.
name	String	The name of the data center.
description	String	A description of the data center.
type	String	The type of data center.
status	List	The availability of the data center.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Datcenter: type = nfs and status != up

returns a list of data centers with:

- A storage type of NFS and status other than up

E.1.2. Searching for Clusters

The following table describes all search options for clusters.

Table E.2. Searching Clusters

Property (of resource or resource-type)	Type	Description (Reference)
Datcenter.datcenter-prop	See property types in Section E.1.1, "Searching for Data Centers"	The property of the data center associated with the cluster.
Datcenter	String	The data center to which the cluster belongs.
name	String	The unique name that identifies the clusters on the network.

Property (of resource or resource-type)	Type	Description (Reference)
description	String	The description of the cluster.
initialized	String	True or False indicating the status of the cluster.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Clusters: initialized = true or name = Default

returns a list of clusters which are:

- initialized; or
- named Default

E.1.3. Searching for Hosts

The following table describes all search options for hosts.

Table E.3. Searching for Hosts

Property (of resource or resource-type)	Type	Description (Reference)
Vms.Vms-prop	See property types in Section E.1.5, "Searching for Virtual Machines"	The property of the Vms associated with the host.
Templates.templates-prop	See property types in Section E.1.7, "Searching for Templates"	The property of the templates associated with the host.
Events.events-prop	See property types in Section E.1.9, "Searching for Events"	The property of the Events associated with the host.
Users.users-prop	See property types in Section E.1.8, "Searching for Users"	The property of the Users associated with the host.
name	String	The name of the host.
status	List	The availability of the host.
cluster	String	The cluster to which the host belongs.
address	String	The unique name that identifies the host on the network.
cpu_usage	Integer	The percent of processing power used.
mem_usage	Integer	The percentage of memory used.

Property (of resource or resource-type)	Type	Description (Reference)
network_usage	Integer	The percentage of network usage.
load	Integer	Jobs waiting to be executed in the <i>run-queue</i> per processor, in a given time slice.
version	Integer	The version number of the operating system.
cpus	Integer	The number of CPUs on the host.
memory	Integer	The amount of memory available.
cpu_speed	Integer	The processing speed of the CPU.
cpu_model	String	The type of CPU.
active_vms	Integer	The number of Vms currently running.
migrating_vms	Integer	The number of Vms currently being migrated.
committed_mem	Integer	The percentage of committed memory.
tag	String	The tag assigned to the host.
type	String	The type of host.
datacenter	String	The data center to which the host belongs.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Hosts: cluster = Default and Vms.os = windows7

returns a list of hosts which:

- Are part of the Default cluster and host virtual machines running the Windows 7 operating system.

E.1.4. Searching for Storage

The following table describes all search options for storage.

Table E.4. Searching for Storage

Property (of resource or resource-type)	Type	Description (Reference)
name	String	The unique name that identifies the storage on the network.

Property (of resource or resource-type)	Type	Description (Reference)
status	String	The status of the storage domain.
datacenter	String	The data center to which the storage belongs.
type	String	The type of the storage.
size	Integer	The size of the storage.
used	Integer	The amount of the storage that is used.
committed	Integer	The amount of the storage that is committed.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Storage: size > 200 or used < 50

returns a list of storage with:

- total storage space greater than 200 GB; or
- used storage space less than 50 GB.

E.1.5. Searching for Virtual Machines

The following table describes all search options for virtual machines (Vms). Vms can be either virtual servers or virtual desktops.

Table E.5. Searching for Virtual Machines

Property (of resource or resource-type)	Type	Description (Reference)
Hosts.hosts-prop	See property types in Section E.1.3, "Searching for Hosts"	The property of the hosts associated with the virtual machine.
Templates.templates-prop	Property types in the bottom portion of this table as well as custom tags	The property of the templates associated with the virtual machine.
Events.events-prop	See property types in Section E.1.9, "Searching for Events"	The property of the events associated with the virtual machine.
Users.users-prop	See property types in Section E.1.8, "Searching for Users"	The property of the users associated with the virtual machine.
name	String	The name of the virtual machine.

Property (of resource or resource-type)	Type	Description (Reference)
status	List	The availability of the virtual machine.
ip	Integer	The IP address of the virtual machine.
uptime	Integer	The number of minutes that the virtual machine has been running.
domain	String	The domain (usually Active Directory domain) that groups these machines.
os	String	The operating system on which the virtual machine was created.
creationdate	Date	The date on which the virtual machine was created.
address	String	The unique name that identifies the virtual machine on the network.
cpu_usage	Integer	The percent of processing power used.
mem_usage	Integer	The percentage of memory used.
network_usage	Integer	The percentage of network used.
memory	Integer	The maximum memory defined.
apps	String	The applications currently installed on the virtual machine.
cluster	List	The cluster to which the virtual machine belongs.
pool	List	The virtual machine pool to which the virtual machine belongs.
loggedinuser	String	The name of the user currently logged in to the virtual machine.
tag	List	The tags to which the virtual machine belongs.
datacenter	String	The data center to which the virtual machine belongs.
type	List	The virtual machine type (server or desktop).
sortby	List	Sorts the returned results by one of the resource properties.

Appendix E. Search Parameters

Property (of resource or resource-type)	Type	Description (Reference)
page	Integer	The page number of results to display.

Example

Vms: template.name = Win* and user.name = ""

Returns a list of VMs, where:

- The template on which the virtual machine is based begins with Win and the virtual machine is assigned to any user.

Example

Vms: cluster = Default and os = windowsxp

Returns a list of VMs, where:

- The cluster to which the virtual machine belongs is named Default and the virtual machine is running the Windows XP operating system.

E.1.6. Searching for Pools

The following table describes all search options for Pools.

Table E.6. Searching for Pools

Property (of resource or resource-type)	Type	Description (Reference)
name	String	The name of the pool.
description	String	The description of the pool.
type	String	The type of pool.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Pools: type = automatic

returns a list of pools with:

- Type of automatic

E.1.7. Searching for Templates

The following table describes all search options for templates. Options are offered by Auto-Completion when typing where appropriate.

Table E.7. Searching for Templates

Property (of resource or resource-type)	Type	Description (Reference)
Vms.Vms-prop	See property types in Section E.1.5, "Searching for Virtual Machines"	The property of the VMs associated with the template.
Hosts.hosts-prop	See property types in Section E.1.3, "Searching for Hosts"	The property of the hosts associated with the template.
Events.events-prop	See property types in Section E.1.9, "Searching for Events"	The property of the events associated with the template.
Users.users-prop	See property types in Section E.1.8, "Searching for Users"	The property of the users associated with the template.
name	String	The name of the template.
domain	String	The domain of the template.
os	String	The type of operating system.
creationdate	Integer	The date on which the template was created. Date format is <i>mm/dd/yy</i> .
childcount	Integer	The number of Vms created from the template.
mem	Integer	Defined memory.
description	String	The description of the template.
status	String	The status of the template.
cluster	String	The cluster associated with the template.
datacenter	String	The data center associated with the template.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Template: `Events.severity >= normal and Vms.uptime > 0`

Returns a list of templates, where:

- Events of normal or greater severity have occurred on VMs derived from the template, and the VMs are still running.

Appendix E. Search Parameters



Figure E.1. Sample Template Search

E.1.8. Searching for Users

The following table describes all search options for users.

Table E.8. Searching for Users

Property (of resource or resource-type)	Type	Description (Reference)
Vms.Vms-prop	See property types in Section E.1.5, "Searching for Virtual Machines"	The property of the VMs associated with the user.
Hosts.hosts-prop	See property types in Section E.1.3, "Searching for Hosts"	The property of the hosts associated with the user.
Templates.templates-prop	See property types in Section E.1.7, "Searching for Templates"	The property of the templates associated with the user.
Events.events-prop	See property types in Section E.1.9, "Searching for Events"	The property of the events associated with the user.
name	String	The name of the user.
lastname	String	The last name of the user.
username	String	The unique name of the user.
department	String	The department to which the user belongs.
group	String	The group to which the user belongs.
title	String	The title of the user.
status	String	The status of the user.
role	String	The role of the user.
tag	String	The tag to which the user belongs.
pool	String	The pool to which the user belongs.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Users: Events.severity > normal and Vms.status = up or Vms.status = pause

Returns a list of users where:

- Events of greater than normal severity have occurred on their Vms AND the Vms are still running; or
- The users' VMs are paused.

E.1.9. Searching for Events

The following table describes all search options you can use to search for events. Auto-completion is offered for many options as appropriate.

Table E.9. Searching for Events

Property (of resource or resource-type)	Type	Description (Reference)
Vms.Vms-prop	See property types in Section E.1.5, "Searching for Virtual Machines"	The property of the Vms associated with the event.
Hosts.hosts-prop	See property types in Section E.1.3, "Searching for Hosts"	The property of the hosts associated with the event.
Templates.templates-prop	See property types in Section E.1.7, "Searching for Templates"	The property of the templates associated with the event.
Users.users-prop	See property types in Section E.1.8, "Searching for Users"	The property of the users associated with the event.
type	List	Type of the event.
severity	List	The severity of the Event: Warning/Error/Normal.
message	String	Description of the event type.
time	Integer	Time at which the event occurred.
username	username	The username associated with the event.
event_host	String	The host associated with the event.
event_vm	String	The virtual machine associated with the event.
event_template	String	The template associated with the event.
event_storage	String	The storage associated with the event.
event_datacenter	String	The data center associated with the event.
sortby	List	Sorts the returned results by one of the resource properties.

Appendix E. Search Parameters

Property (of resource or resource-type)	Type	Description (Reference)
page	Integer	The page number of results to display.

Example

Events: `Vms.name = testdesktop` and `Hosts.name = gonzo.example.com`

Returns a list of events, where:

- The event occurred on the virtual machine named `testdesktop` while it was running on the host `gonzo.example.com`.

Appendix F. SAP Monitoring

This appendix describes how to enable the hypervisor monitoring required when running SAP software in Red Hat Enterprise Linux guests with Red Hat Enterprise Virtualization. Hypervisor monitoring for SAP applications uses the **vhostmd** VDSM hook, which is provided to hosts by the `vdsm-hook-vhostmd` package.

Red Hat Enterprise Virtualization Hypervisor hosts include the `vdsm-hook-vhostmd` package by default.

On Red Hat Enterprise Linux hosts, install the **vhostmd** VDSM hook from the terminal as the root user using the `yum install vdsm-hook-vhostmd` command.

Procedure F.1. To enable SAP monitoring for a Virtual Machine from the Administration Portal:

1. Select the virtual machine to be monitored from the **Virtual Machines** tab. The virtual machine must be in a stopped state.
2. Click the **Edit** button to display the **Edit Virtual Machine** dialog.

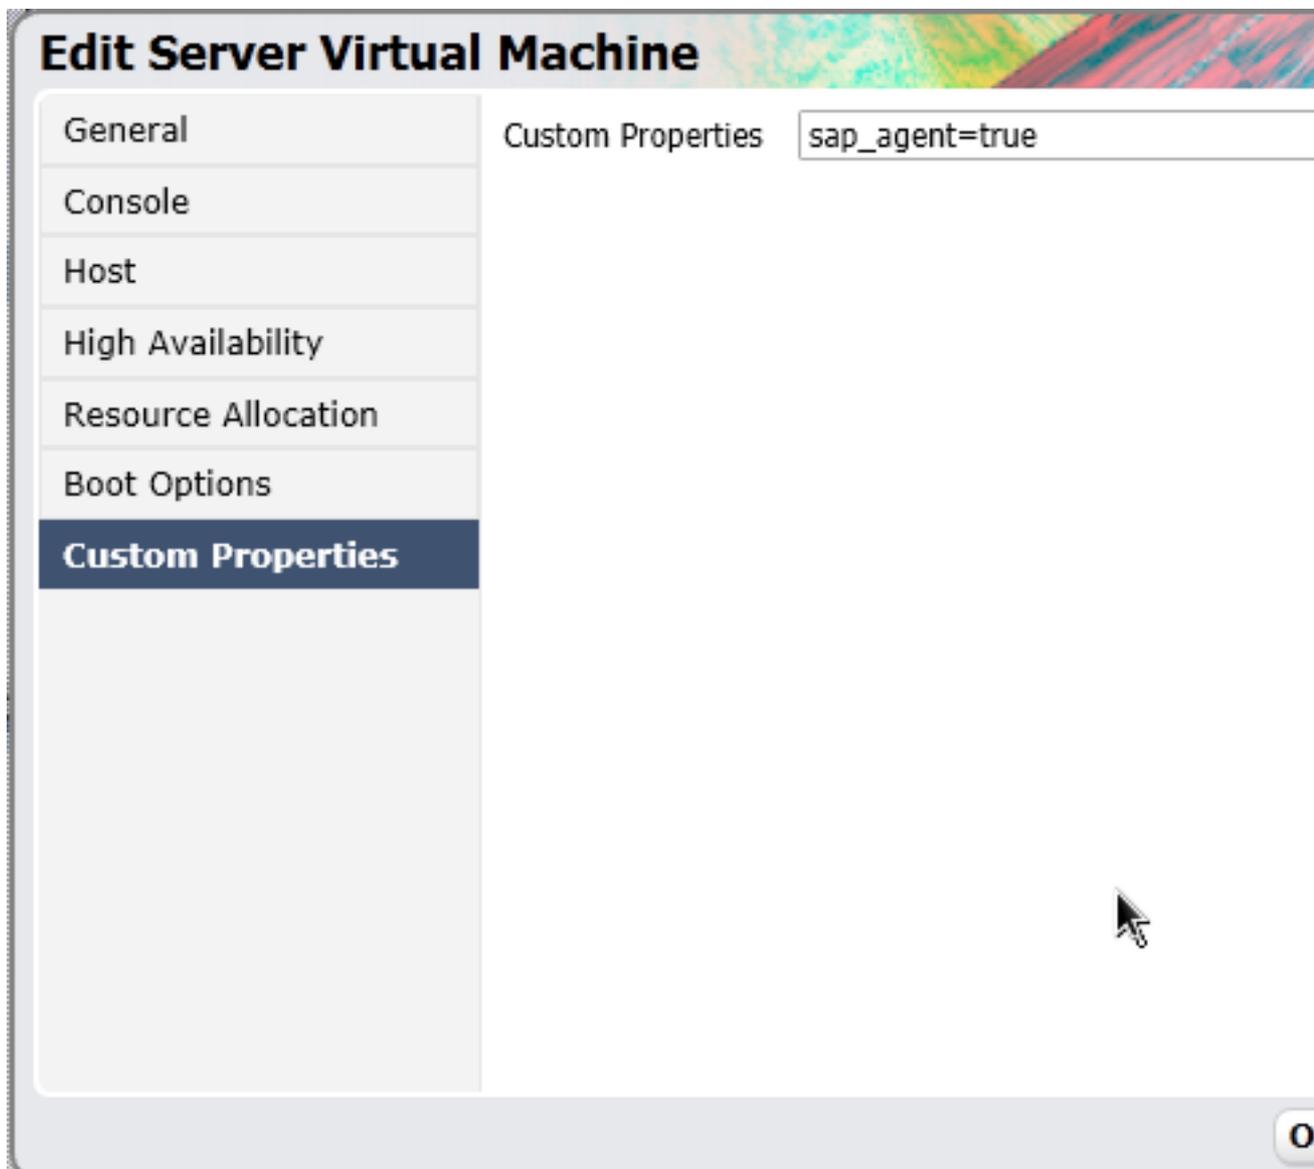


Figure F.1. Enable SAP

3. Select the **Custom Properties** tab.
4. In the **Custom Properties** text field, enter `sap_agent=true`. Click **OK**.

You have now enabled monitoring of a virtual machine for use with SAP. Repeat the procedure with each virtual machine that is part of the SAP environment.

The **Custom Properties** tab is also available from the **New Virtual Machine** dialog. SAP monitoring can be enabled for new virtual machines by following this procedure at the time of virtual machine creation.

Virtual machine monitoring for SAP can also be enabled using the REST API. See the *Red Hat Enterprise Virtualization REST API Guide* for more information.

Appendix G. KVM Virtual Machine Timing Management

Virtualization poses various challenges for virtual machine time keeping. Virtual machines which use the Time Stamp Counter (TSC) as a clock source may suffer timing issues as some CPUs do not have a constant Time Stamp Counter. Virtual machines running without accurate timekeeping can have serious affects on some networked applications as your virtual machine will run faster or slower than the actual time.

KVM works around this issue by providing virtual machines with a para-virtualized clock. The KVM pvclock provides a stable source of timing for KVM guests that support it.

Presently, only Red Hat Enterprise Linux 5.4 and higher virtual machines fully support the para-virtualized clock.

Virtual machines can have several problems caused by inaccurate clocks and counters:

- Clocks can fall out of synchronization with the actual time which invalidates sessions and affects networks.
- Virtual machines with slower clocks may have issues migrating.

These problems exist on other virtualization platforms and timing should always be tested.



Important

The Network Time Protocol (NTP) daemon should be running on the host and the virtual machines. Enable the `ntpd` service:

```
# service ntpd start
```

Add the `ntpd` service to the default startup sequence:

```
# chkconfig ntpd on
```

Using the `ntpd` service should minimize the affects of clock skew in all cases.

Determining if your CPU has the constant Time Stamp Counter

Your CPU has a constant Time Stamp Counter if the `constant_tsc` flag is present. To determine if your CPU has the `constant_tsc` flag run the following command:

```
$ cat /proc/cpuinfo | grep constant_tsc
```

If any output is given your CPU has the `constant_tsc` bit. If no output is given follow the instructions below.

Configuring hosts without a constant Time Stamp Counter

Systems without constant time stamp counters require additional configuration. Power management features interfere with accurate time keeping and must be disabled for virtual machines to accurately keep time with KVM.



Important

These instructions are for AMD revision F cpus only.

If the CPU lacks the **constant_tsc** bit, disable all power management features ([BZ#513138](#)¹). Each system has several timers it uses to keep time. The TSC is not stable on the host, which is sometimes caused by **cpufreq** changes, deep C state, or migration to a host with a faster TSC. Deep C sleep states can stop the TSC. To prevent the kernel using deep C states append "**processor.max_cstate=1**" to the kernel boot options in the **grub.conf** file on the host:

```
term Red Hat Enterprise Linux Server (2.6.18-159.el5)
  root (hd0,0)
  kernel /vmlinuz-2.6.18-159.el5 ro root=/dev/VolGroup00/LogVol00 rhgb
  quiet processor.max_cstate=1
```

Disable **cpufreq** (only necessary on hosts without the **constant_tsc**) by editing the **/etc/sysconfig/cpuspeed** configuration file and change the **MIN_SPEED** and **MAX_SPEED** variables to the highest frequency available. Valid limits can be found in the **/sys/devices/system/cpu/cpu*/cpufreq/scaling_available_frequencies** files.

Using the para-virtualized clock with Red Hat Enterprise Linux virtual machines

For certain Red Hat Enterprise Linux virtual machines, additional kernel parameters are required. These parameters can be set by appending them to the end of the **/kernel** line in the **/boot/grub/grub.conf** file of the virtual machine.



Note

The process of configuring kernel parameters can be automated using the **ktune** package

The **ktune** package provides an interactive Bourne shell script, **fix_clock_drift.sh**. When run as the superuser, this script inspects various system parameters to determine if the virtual machine on which it is run is susceptible to clock drift under load. If so, it then creates a new **grub.conf.kvm** file in the **/boot/grub/** directory. This file contains a kernel boot line with additional kernel parameters that allow the kernel to account for and prevent significant clock drift on the KVM virtual machine. After running **fix_clock_drift.sh** as the superuser, and once the script has created the **grub.conf.kvm** file, then the virtual machine's current **grub.conf** file should be backed up manually by the system administrator, the new **grub.conf.kvm** file should be manually inspected to ensure that it is identical to **grub.conf** with the exception of the additional boot line parameters, the **grub.conf.kvm** file should finally be renamed **grub.conf**, and the virtual machine should be rebooted.

The table below lists versions of Red Hat Enterprise Linux and the parameters required for virtual machines on systems without a constant Time Stamp Counter.

¹ https://bugzilla.redhat.com/show_bug.cgi?id=513138

Red Hat Enterprise Linux

5.4 AMD64/Intel 64 with the para-virtualized clock

5.4 AMD64/Intel 64 without the para-virtualized clock

5.4 x86 with the para-virtualized clock

5.4 x86 without the para-virtualized clock

5.3 AMD64/Intel 64

5.3 x86

4.8 AMD64/Intel 64

4.8 x86

3.9 AMD64/Intel 64

3.9 x86

Additional virtual machine kernel parameters

Additional parameters are not required

divider=10 notsc lpj=n

Additional parameters are not required

divider=10 clocksource=acpi_pm lpj=n

divider=10 notsc

divider=10 clocksource=acpi_pm

notsc divider=10

clock=pmtmr divider=10

Additional parameters are not required

Additional parameters are not required

Using the Real-Time Clock with Windows virtual machines

Windows uses both the Real-Time Clock (RTC) and the Time Stamp Counter (TSC). For Windows virtual machines the Real-Time Clock can be used instead of the TSC for all time sources which resolves virtual machine timing issues.

To enable the Real-Time Clock for the PMTIMER clocksource (the PMTIMER usually uses the TSC) add the following line to the Windows boot settings. Windows boot settings are stored in the boot.ini file. Add the following line to the **boot.ini** file:

```
/use pmtimer
```

For more information on Windows boot settings and the pmtimer option, refer to [Available switch options for the Windows XP and the Windows Server 2003 Boot.ini files](#)².

² <http://support.microsoft.com/kb/833721>

Appendix H. Additional References

These additional documentation resources do not form part of the Red Hat Enterprise Virtualization documentation suite. They do however contain useful information for System Administrators managing Red Hat Enterprise Virtualization environments and are available at <http://docs.redhat.com/>.

Red Hat Enterprise Linux — Deployment Guide

A guide to the deployment, configuration and administration of Red Hat Enterprise Linux.

Red Hat Enterprise Linux — DM-Multipath Guide

A guide to the use of Device-Mapper Multipathing on Red Hat Enterprise Linux.

Red Hat Enterprise Linux — Hypervisor Deployment Guide

The complete guide to installing, deploying and maintaining Red Hat Enterprise Virtualization Hypervisors.

Red Hat Enterprise Linux — Installation Guide

A guide to the installation of Red Hat Enterprise Linux.

Red Hat Enterprise Linux — Storage Administration Guide

A guide to the management of storage devices and file systems on Red Hat Enterprise Linux.

Red Hat Enterprise Linux — Virtualization Guide

A guide to the installation, configuration, administration and troubleshooting of virtualization technologies in Red Hat Enterprise Linux.

Appendix I. Revision History

Revision 2-22	Mon Dec 05 2011	Kate Grainger kgrainge@redhat.com
Bug fixes.		
Revision 2-21	Thu Nov 17 2011	Kate Grainger kgrainge@redhat.com
Bug fixes.		
Revision 2-20	Thu Nov 10 2011	Zac Dover zdoover@redhat.com
Removing references to internal resources in images.		
Revision 2-19	Wed Nov 09 2011	Zac Dover zdoover@redhat.com
Bug fixes.		
Revision 2-19	Wed Nov 09 2011	Zac Dover zdoover@redhat.com
Bug fixes.		
Revision 2-18	Thu Nov 03 2011	Kate Grainger kgrainge@redhat.com
QE and translation updates.		
Revision 2-17	Fri Oct 28 2011	Kate Grainger kgrainge@redhat.com
QE updates.		
Revision 2-16	Tues Oct 25 2011	Kate Grainger kgrainge@redhat.com
QE updates.		
Revision 2-15	Mon Oct 24 2011	Kate Grainger kgrainge@redhat.com
Translation and QE updates.		
Revision 2-14	Mon Oct 24 2011	Kate Grainger kgrainge@redhat.com
Translation and QE updates.		
Revision 2-13	Mon Oct 19 2011	Shikha Nansi snansi@redhat.com
Publishing book.		
Revision 2-12	Mon Oct 17 2011	Shikha Nansi snansi@redhat.com
QA edits.		

Appendix I. Revision History

Revision 2-11	Mon Oct 12 2011	Shikha Nansi snansi@redhat.com
QA edits.		
Revision 2-10	Mon Oct 5 2011	Shikha Nansi snansi@redhat.com
Publishing book.		
Revision 2-9	Mon Oct 3 2011	Zac Dover zdoover@redhat.com
Publishing book.		
Revision 2-8	Fri Sep 30 2011	Zac Dover zdoover@redhat.com
QA corrections in History database tables.		
Revision 2-7	Fri Sep 16 2011	Shikha Nansi snansi@redhat.com
Changes from email feedback.		
Revision 2-6	Fri Sep 16 2011	Shikha Nansi snansi@redhat.com
Changes from email feedback.		
Revision 2-5	Thu Sep 15 2011	Shikha Nansi snansi@redhat.com
Changes from email feedback.		
Revision 2-4	Wed Sep 14 2011	Zac Dover zdoover@redhat.com
Corrections to History and ENUM Views and vm_disk_history tables.		
Revision 2-3	Fri Sep 09 2011	Shikha Nansi snansi@redhat.com
Spell Check.		
Revision 2-2	Thu Sep 08 2011	Shikha Nansi snansi@redhat.com
Updated Hosts and Storage.		
Revision 1-1	Thu Sep 01 2011	Shikha Nansi snansi@redhat.com
Updated Clusters and Logical Networks.		
Revision 1-0	Tue Aug 23 2011	Susan Burgess sburgess@redhat.com
Restructured Guide for 3.0.		