



# **Red Hat Enterprise Virtualization**

## **3.2**

### **Administration Guide**

---

Administering Red Hat Enterprise Virtualization Environments.  
Edition 1

Jodi Biddle  
Zac Dover

Jodi Biddle  
Steve Gordon

Andrew Burden  
Tim Hildred

Dan Macpherson  
Cheryn Tan

Dayle Parker

Anjana Sriram

# Red Hat Enterprise Virtualization 3.2 Administration Guide

---

## Administering Red Hat Enterprise Virtualization Environments. Edition 1

Jodi Biddle  
jbiddle@redhat.com

Andrew Burden  
aburden@redhat.com

Zac Dover  
zdover@redhat.com

Steve Gordon  
sgordon@redhat.com

Tim Hildred  
thildred@redhat.com

Dan Macpherson  
dmacpher@redhat.com

Dayle Parker  
dayparke@redhat.com

Anjana Sriram  
Red Hat Engineering Content Services  
asriram@redhat.com

Cheryn Tan  
chetan@redhat.com

## **Legal Notice**

Copyright 2013 Red Hat, Inc. This document is licensed by Red Hat under the Creative Commons Attribution-ShareAlike 3.0 Unported License. If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed. Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, 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. Node.js is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project. The OpenStack Word Mark and OpenStack Logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community. All other trademarks are the property of their respective owners.

## **Keywords**

## **Abstract**

This book contains information and procedures relevant to Red Hat Enterprise Virtualization administrators.

# Table of Contents

<b>Preface</b>	<b>13</b>
1. Document Conventions	13
1.1. Typographic Conventions	13
1.2. Pull-quote Conventions	13
1.3. Notes and Warnings	14
2. Getting Help and Giving Feedback	14
2.1. Do You Need Help?	14
2.2. We Need Feedback!	14
<b>Chapter 1. Using this Guide</b>	<b>15</b>
1.1. Administration Guide Prerequisites	15
1.2. Administration Guide Layout	15
1.3. Example Workflows	15
1.3.1. Administration Guide Example Workflows Overview	15
1.3.2. Administration Guide Example Workflow: New iSCSI Data Center	15
1.3.3. Administration Guide Example Workflow: Newly Virtualized Workload	15
1.3.4. Administration Guide Example Workflow: Template for Group Use	15
<b>Chapter 2. Basics</b>	<b>17</b>
2.1. Introduction	17
2.1.1. Red Hat Enterprise Virtualization Architecture	17
2.1.2. Red Hat Enterprise Virtualization System Components	17
2.1.3. Red Hat Enterprise Virtualization Resources	17
2.1.4. Red Hat Enterprise Virtualization API Support Statement	18
2.1.5. SPICE	18
2.1.6. Administering and Maintaining the Red Hat Enterprise Virtualization Environment	18
2.2. Using the Administration Portal Graphical Interface	19
2.2.1. Graphical User Interface Elements	19
2.2.2. Tree Mode and Flat Mode	19
2.2.3. Using the Guide Me Facility	20
2.2.4. Performing Searches in Red Hat Enterprise Virtualization	20
2.2.5. Saving a Query String as a Bookmark	20
<b>Part I. Administering the Resources</b>	<b>22</b>
<b>Chapter 3. Data Centers</b>	<b>23</b>
3.1. Introduction to Data Centers	23
3.2. The Storage Pool Manager (SPM)	23
3.3. SPM Priority	23
3.4. Using the Events Tab to Identify Problem Objects in Data Centers	24
3.5. Data Center Tasks	24
3.5.1. Creating a New Data Center	24
3.5.2. Explanation of Settings in the New Data Center and Edit Data Center Windows	24
3.5.3. Editing a Resource	25
3.5.4. Creating a New Logical Network in a Data Center or Cluster	25
3.5.5. Re-Initializing a Data Center: Recovery Procedure	26
3.5.6. Removing a Data Center	26
3.5.7. Force Removing a Data Center	26
3.6. Data Centers and Storage Domains	27
3.6.1. Attaching an Existing Data Domain to a Data Center	27
3.6.2. Attaching an Existing ISO domain to a Data Center	27
3.6.3. Attaching an Existing Export Domain to a Data Center	27

3.6.4. Detaching a Storage Domain from a Data Center	27
3.6.5. Activating a Storage Domain from Maintenance Mode	28
3.7. Data Centers and Permissions	28
3.7.1. Managing System Permissions for a Data Center	28
3.7.2. Data Center Administrator Roles Explained	28
3.7.3. Assigning an Administrator or User Role to a Resource	28
3.7.4. Removing an Administrator or User Role from a Resource	29
<b>Chapter 4..Clusters</b>	<b>30</b>
4.1. Introduction to Clusters	30
4.2. Cluster Tasks	30
4.2.1. Creating a New Cluster	30
4.2.2. Explanation of Settings and Controls in the New Cluster and Edit Cluster Windows	
4.2.2.1. General Cluster Settings Explained	31
4.2.2.2. Optimization Settings Explained	32
4.2.2.3. Resilience Policy Settings Explained	33
4.2.2.4. Cluster Policy Settings Explained	33
4.2.3. Editing a Resource	34
4.2.4. Importing an Existing Red Hat Storage Cluster	34
4.2.5. Explanation of Settings in the Add Hosts Window	35
4.2.6. Setting Load and Power Management Policies for Hosts in a Cluster	36
4.2.7. Creating a New Logical Network in a Data Center or Cluster	37
4.2.8. Removing a Cluster	37
4.2.9. Designate a Specific Traffic Type for a Logical Network with the Assign/UnAssign Window	
4.2.10. Explanation of Settings in the Assign/UnAssign Window	38
4.3. Clusters and Permissions	38
4.3.1. Managing System Permissions for a Cluster	38
4.3.2. Cluster Administrator Roles Explained	38
4.3.3. Assigning an Administrator or User Role to a Resource	39
4.3.4. Removing an Administrator or User Role from a Resource	39
<b>Chapter 5..Logical Networks</b>	<b>40</b>
5.1. Introduction to Logical Networks	40
5.2. Port Mirroring	40
5.3. Required Networks, Optional Networks, and Virtual Machine Networks	40
5.4. Logical Network Tasks	41
5.4.1. Creating a New Logical Network in a Data Center or Cluster	41
5.4.2. Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows	41
5.4.3. Editing a Logical Network	42
5.4.4. Designate a Specific Traffic Type for a Logical Network with the Assign/UnAssign Window	
5.4.5. Explanation of Settings in the Assign/UnAssign Window	43
5.4.6. Adding Multiple VLANs to a Single Network Interface using Logical Networks	43
5.4.7. Using the Networks Tab	44
5.5. Logical Networks and Permissions	45
5.5.1. Managing System Permissions for a Network	45
5.5.2. Network Administrator and User Roles Explained	45
5.5.3. Assigning an Administrator or User Role to a Resource	45
5.5.4. Removing an Administrator or User Role from a Resource	45
<b>Chapter 6..Hosts</b>	<b>47</b>
6.1. Introduction to Red Hat Enterprise Virtualization Hosts	47
6.2. Red Hat Enterprise Virtualization Hypervisor Hosts	47
6.3. Red Hat Enterprise Linux Hosts	47
6.4. Host Tasks	47
6.4.1. Adding a Red Hat Enterprise Linux Host	47

6.4.2. Explanation of Settings and Controls in the New Host and Edit Host Windows	48
6.4.2.1. Host General Settings Explained	48
6.4.2.2. Host Power Management Settings Explained	48
6.4.2.3. SPM Priority Settings Explained	50
6.4.2.4. Host Console Settings Explained	50
6.4.3. Configuring Host Power Management Settings	51
6.4.4. Configuring Host Storage Pool Manager (SPM) Settings	51
6.4.5. Editing a Resource	51
6.4.6. Approving Newly Added Red Hat Enterprise Virtualization Hypervisor Hosts	52
6.4.7. Moving a Host to Maintenance Mode	52
6.4.8. Activating a host from maintenance mode	52
6.4.9. Removing a Host	52
6.4.10. Customizing Hosts with Tags	53
6.5. Hosts and Networking	53
6.5.1. Editing Host Network Interfaces and Adding Logical Networks to Hosts	53
6.5.2. Creating a Bond Device using the Administration Portal	54
6.5.3. Example Uses of Custom Bonding Options with Host Interfaces	54
6.5.4. Saving a host network configuration	55
6.6. Host Resilience	55
6.6.1. Host High Availability	55
6.6.2. Power Management by Proxy in Red Hat Enterprise Virtualization	55
6.6.3. Setting Fencing Parameters on a Host	55
6.6.4. Using Host Power Management Functions	56
6.6.5. Manually Fencing or Isolating a Non Responsive Host	57
6.7. Hosts and Permissions	57
6.7.1. Managing System Permissions for a Host	57
6.7.2. Host Administrator Roles Explained	57
6.7.3. Assigning an Administrator or User Role to a Resource	58
6.7.4. Removing an Administrator or User Role from a Resource	58
<b>Chapter 7. Storage</b>	<b>59</b>
7.1. Introduction to Storage in Red Hat Enterprise Virtualization	59
7.2. Understanding Storage Domains	59
7.3. Storage Metadata Versions in Red Hat Enterprise Virtualization	59
7.4. Preparing and Adding File-based Storage	60
7.4.1. Preparing NFS Storage	60
7.4.2. Attaching NFS Storage	60
7.4.3. Preparing Local Storage	61
7.4.4. Adding Local Storage	62
7.5. Adding POSIX Compliant File System Storage	62
7.5.1. POSIX Compliant File System Storage in Red Hat Enterprise Virtualization	62
7.5.2. Attaching POSIX Compliant File System Storage	62
7.6. Adding Block-based Storage	63
7.6.1. Preparing iSCSI Storage	63
7.6.2. Adding iSCSI Storage	63
7.6.3. Adding FCP Storage	65
7.6.4. Un-useable LUNs in Red Hat Enterprise Virtualization	65
7.7. Storage Tasks	66
7.7.1. Importing existing ISO or export storage domains	66
7.7.2. Populating the ISO Storage Domain	67
7.7.3. Moving storage domains to maintenance mode	67
7.7.4. Editing a Resource	67
7.7.5. Activating storage domains	67
7.7.6. Removing a storage domain	68
7.7.7. Destroying a storage domain	68
7.7.8. Detaching the Export Domain	68

7.7.9. Attaching an Export Domain to a Data Center	68
7.8. Red Hat Storage Volumes	69
7.8.1. Introduction to Red Hat Storage Volumes	69
7.8.2. Gluster Storage Terminology	69
7.8.3. Creating A Storage Volume	71
7.8.4. Adding Bricks to a Volume	71
7.8.5. Explanation of Settings in the Add Bricks window	72
7.8.6. Enabling Red Hat Storage Volumes for Virtualization	72
7.8.7. Starting Volumes	73
7.8.8. Tuning Volumes	73
7.8.9. Editing Volume Options	73
7.8.10. Reset Volume Options	73
7.8.11. Removing Bricks from a Volume	74
7.8.12. Stopping Red Hat Storage Volumes	74
7.8.13. Deleting Red Hat Storage Volumes	74
7.9. Storage and Permissions	74
7.9.1. Managing system permissions for a storage domain	74
7.9.2. Storage Administrator Roles Explained	75
7.9.3. Assigning an Administrator or User Role to a Resource	75
7.9.4. Removing an Administrator or User Role from a Resource	75
<b>Chapter 8. Virtual Machines</b>	<b>76</b>
8.1. Introduction to Virtual Machines	76
8.2. Supported Virtual Machine Operating Systems	76
8.3. Virtual Machine Performance Parameters	76
8.4. Creating Virtual Machines	76
8.4.1. Creating a New Virtual Machine from an Existing Template	76
8.4.2. Creating a New Virtual Machine from a Blank Template	77
8.4.3. Explanation of Settings and Controls in the New Virtual Machine and Edit Virtual Machine Windows	77
8.4.3.1. Virtual Machine General Settings Explained	77
8.4.3.2. Virtual Machine Initial Run Settings Explained	78
8.4.3.3. Virtual Machine Console Settings Explained	78
8.4.3.4. Virtual Machine Host Settings Explained	79
8.4.3.5. Virtual Machine High Availability Settings Explained	80
8.4.3.6. Virtual Machine Resource Allocation Settings Explained	80
8.4.3.7. Virtual Machine Boot Options Settings Explained	81
8.4.3.8. Virtual Machine Custom Properties Settings Explained	82
8.4.4. Creating a Cloned Virtual Machine from an Existing Template	83
8.4.5. Completing the Configuration of a Virtual Machine by Defining Network Interfaces and Hard Disks	84
8.4.6. Installing a Guest Operating System onto a Virtual Machine	85
8.4.7. Installing Windows on VirtIO Optimized Hardware	85
8.4.8. Virtual Machine Run Once Settings Explained	85
8.5. Using Virtual Machines	86
8.5.1. SPICE	86
8.5.2. Powering on a Virtual Machine	87
8.5.3. Installing SPICE Plugins in Windows and Linux	87
8.5.4. Logging in to a Virtual Machine	87
8.6. Shutting Down or Pausing Virtual Machines	87
8.6.1. Shutting Down or Pausing Virtual Machines	87
8.6.2. Shutting Down a Virtual Machine	88
8.6.3. Pausing a Virtual Machine	88
8.7. Managing Virtual Machines	88
8.7.1. Editing a Resource	88
8.7.2. Removing a Virtual Machine	88



8.7.3. Adding and Editing Virtual Machine Disks	89
8.7.4. Adding and Editing Virtual Machine Network Interfaces	89
8.7.5. Explanation of Settings in the Virtual Machine Network Interface Window	90
8.7.6. Hot Plugging Virtual Machine Disks	91
8.7.7. Hot Plugging Network Interfaces	91
8.7.8. Removing Disks and Network Interfaces from Virtual Machines	92
8.8. Virtual Machines and Permissions	92
8.8.1. Managing System Permissions for a Virtual Machine	92
8.8.2. Virtual Machines Administrator Roles Explained	92
8.8.3. Virtual Machine User Roles Explained	93
8.8.4. Assigning an Administrator or User Role to a Resource	93
8.8.5. Removing an Administrator or User Role from a Resource	93
8.9. Backing Up and Restoring Virtual Machines with Snapshots	94
8.9.1. Creating a Snapshot of a Virtual Machine	94
8.9.2. Using a Snapshot to Restore a Virtual Machine	94
8.9.3. Creating a Virtual Machine from a Snapshot	95
8.9.4. Deleting a Snapshot	95
8.10. Importing and Exporting Virtual Machines	96
8.10.1. Exporting and Importing Virtual Machines	96
8.10.2. Overview of the Export-Import Process	96
8.10.3. Performing an Export-Import of Virtual Resources	96
8.10.4. Exporting a Virtual Machine to the Export Domain	98
8.10.5. Importing a Virtual Machine into the Destination Data Center	98
8.11. Migrating Virtual Machines Between Hosts	99
8.11.1. What is Live Migration?	99
8.11.2. Live Migration Prerequisites	99
8.11.3. Automatic Virtual Machine Migration	99
8.11.4. Preventing Automatic Migration of a Virtual Machine	99
8.11.5. Manually Migrating Virtual Machines	100
8.11.6. Setting Migration Priority	100
8.11.7. Cancelling ongoing virtual machine migrations	101
8.11.8. Event and Log Notification upon Automatic Migration of Highly Available Virtual Servers	101
8.12. Improving Uptime with Virtual Machine High Availability	101
8.12.1. Why Use High Availability?	101
8.12.2. What is High Availability?	101
8.12.3. High Availability Considerations	101
8.12.4. Configuring a Highly Available Virtual Machine	102
8.13. Other Virtual Machine Tasks	102
8.13.1. Enabling SAP monitoring for a virtual machine from the Administration Portal	102
8.13.2. Configuring Red Hat Enterprise Linux 5.4 or Higher Virtual Machines to use SPICE	103
8.13.2.1. Using SPICE on virtual machines running versions of Red Hat Enterprise Linux released prior to 5.4	103
8.13.2.2. Installing qxl drivers on virtual machines	103
8.13.2.3. Configuring qxl drivers on virtual machines	103
8.13.2.4. Configuring a virtual machine's tablet and mouse to use SPICE	104
8.13.3. KVM virtual machine timing management	104
<b>Chapter 9. Templates</b>	<b>106</b>
9.1. Introduction to Templates	106
9.2. Template Tasks	106
9.2.1. Creating a Template from an Existing Virtual Machine	106
9.2.2. Explanation of Settings and Controls in the New Template Window	106
9.2.3. Editing a Resource	107
9.2.4. Explanation of Settings and Controls in the Edit Template Window	107

9.2.5. Deleting a Template	107
9.2.6. Exporting Templates	108
9.2.6.1. Migrating Templates to the Export Domain	108
9.2.7. Importing Templates	108
9.2.7.1. Importing a Template into a Data Center	108
9.3. Sealing Templates in Preparation for Deployment	109
9.3.1. Sealing a Linux Virtual Machine for Deployment as a Template	109
9.3.2. Sealing a Windows Template	109
9.3.2.1. Considerations when Sealing a Windows Template with Sysprep	109
9.3.2.2. Sealing a Windows XP Template	109
9.3.2.3. Sealing a Windows 7 or Windows 2008 Template	110
9.4. Templates and Permissions	110
9.4.1. Managing System Permissions for a Template	110
9.4.2. Template Administrator Roles Explained	110
9.4.3. Template User Roles Explained	110
9.4.4. Assigning an Administrator or User Role to a Resource	111
9.4.5. Removing an Administrator or User Role from a Resource	111
<b>Chapter 10.. Pools</b>	<b>112</b>
10.1. Introduction to Virtual Machine Pools	112
10.2. Virtual Machine Pool Tasks	112
10.2.1. Creating a Virtual Machine Pool	112
10.2.2. Editing a Resource	112
10.2.3. Prestarting Virtual Machines in a Pool	113
10.2.4. Adding Virtual Machines to a Virtual Machine Pool	113
10.2.5. Detaching Virtual Machines from a Virtual Machine Pool	113
10.2.6. Removing a Virtual Machine Pool	113
10.3. Pools and Permissions	114
10.3.1. Managing System Permissions for a Virtual Machine Pool	114
10.3.2. Virtual Machine Pool Administrator Roles Explained	114
10.3.3. Assigning an Administrator or User Role to a Resource	114
10.3.4. Removing an Administrator or User Role from a Resource	114
<b>Chapter 11.. Virtual Machine Disks</b>	<b>116</b>
11.1. Understanding Virtual Machine Storage	116
11.2. Understanding Virtual Disks	116
11.3. Shareable Disks in Red Hat Enterprise Virtualization	117
11.4. Creating Unassociated Virtual Machine Hard Disks	117
11.5. Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows	117
11.6. Moving a Virtual Machine Hard Disk Between Data Domains	118
11.7. Virtual Disks and Permissions	119
11.7.1. Managing System Permissions for a Virtual Disk	119
11.7.2. Virtual Disk User Roles Explained	119
11.7.3. Assigning an Administrator or User Role to a Resource	119
11.7.4. Removing an Administrator or User Role from a Resource	119
<b>Part II. Administering the Environment</b>	<b>121</b>
<b>Chapter 12.. Users and Roles</b>	<b>122</b>
12.1. Introduction to Users	122
12.2. Directory Users	122
12.2.1. Directory Services Support in Red Hat Enterprise Virtualization	122
12.3. User Authorization	123
12.3.1. User Authorization Model	123
12.3.2. User Actions	123

12.3.3. User Permissions	123
12.4. Red Hat Enterprise Virtualization Manager User Properties and Roles	124
12.4.1. User Properties	124
12.4.2. User and Administrator Roles	124
12.4.3. User Roles Explained	124
12.4.4. Administrator Roles Explained	125
12.5. Red Hat Enterprise Virtualization Manager User Tasks	126
12.5.1. Adding Users	126
12.5.2. Viewing User Information	126
12.5.3. Viewing User Permissions on Resources	127
12.5.4. Removing Users	127
12.5.5. Configuring Roles	127
12.5.6. Creating a New Role	127
12.5.7. Editing or Copying a Role	128
12.6. User Role and Authorization Examples	128
<b>Chapter 13.. Quotas</b>	<b>130</b>
13.1. Introduction to Quota	130
13.2. Shared Quota and Individually-defined Quota	130
13.3. Quota Accounting	130
13.4. Enabling and Changing a Quota Mode in a Data Center	131
13.5. Creating a New Quota Policy	131
13.6. Explanation of Quota Threshold Settings	131
13.7. Assigning a Quota to an Object	132
13.8. Using Quota to Limit Resources by User	132
13.9. Editing Quotas	133
13.10. Removing Quotas	134
<b>Chapter 14. Event Notifications</b>	<b>135</b>
14.1. Configuring Event Notifications	135
14.2. Parameters for event notifications in notifier.conf	135
14.3. Canceling Event Notifications	136
<b>Chapter 15.. Utilities</b>	<b>137</b>
15.1. Managing Domains with the Domain Management Tool	137
15.1.1. What is the Domain Management Tool?	137
15.1.2. Syntax for the Domain Management Tool	137
15.1.3. Using the Domain Management Tool	138
15.1.4. Adding Domains to Configuration	138
15.1.5. Editing a Domain in the Configuration	138
15.1.6. Deleting a Domain from the Configuration	138
15.1.7. Validating Domain Configuration	138
15.1.8. Listing Domains in Configuration	138
15.1.9. Getting a Configuration Value	139
15.1.10. Setting a Configuration Value	139
15.2. Editing the Configuration of the Red Hat Virtualization Manager with the Configuration Tool	
15.2.1. Configuration Tool	139
15.2.2. Syntax for rhevm-config Command	139
15.2.3. The admin@internal user	139
15.2.4. Changing the password for admin@internal	140
15.2.5. Configuration tool examples	140
15.2.6. Red Hat Enterprise Virtualization Manager Configuration Options Explanations, Limitations, and Best Practices	140
15.3. Uploading Virtual Machine Images with the Image Uploader Tool	149
15.3.1. Virtual Machine Image Uploader	149
15.3.2. Syntax for the rhevm-image-uploader Command	149

15.3.3. Creating an OVF Archive That Is Compatible with the Image Uploader	150
15.3.4. Basic rhvm-image-uploader Usage Examples	151
15.4. Editing USB Filters with the USB Filter Editor	151
15.4.1. Installing the USB Filter Editor	151
15.4.2. The USB Filter Editor Interface	151
15.4.3. Adding a USB policy	152
15.4.4. Removing a USB Policy	152
15.4.5. Searching for USB Device Policies	153
15.4.6. Exporting a USB Policy	153
15.4.7. Importing a USB Policy	154
15.5. Collecting Logs with the Log Collector Tool	154
15.5.1. Log Collector	154
15.5.2. Syntax for rhvm-log-collector Command	154
15.5.3. Basic Log Collector Usage	156
15.6. Uploading ISO Files with the ISO Uploader Tool	156
15.6.1. ISO Uploader	156
15.6.2. Syntax for rhvm-iso-uploader Command	156
15.6.3. Specifying an NFS Server	157
15.6.4. Basic ISO Uploader Usage	157
15.6.5. Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain	158
15.6.6. VirtIO and Guest Tool Image Files	158
15.7. Guest Drivers and Agents	158
15.7.1. Installing Guest Agents and Drivers	158
15.7.2. Automating Guest Additions on Windows Guests with Red Hat Enterprise Virtualization Application Provisioning Tool(APT)	158
15.7.3. Red Hat Enterprise Virtualization Guest Drivers and Guest Agents	159
15.7.4. Subscribing to channels	160
15.7.4.1. Subscribing to Channels Using Subscription Manager	160
15.7.4.2. Subscribing to channels using RHN Classic	161
<b>Chapter 16.. Log Files .....</b>	<b>162</b>
16.1. Red Hat Enterprise Virtualization Manager Installation Log Files	162
16.2. Red Hat Enterprise Virtualization Manager Log Files	162
16.3. Red Hat Enterprise Virtualization Host Log Files	162
16.4. Remotely Logging Host Activities	162
16.4.1. Setting up a Virtualization Host Logging Server	162
16.4.2. Configuring Red Hat Enterprise Virtualization Hypervisor Hosts to Use a Logging Server	163
16.4.3. Configuring Red Hat Enterprise Linux Hosts to Use a Logging Server	163
<b>Chapter 17.. Updating the Red Hat Enterprise Virtualization Environment .....</b>	<b>164</b>
17.1. Upgrades between Minor Releases	164
17.1.1. Checking for Red Hat Enterprise Virtualization Manager Updates	164
17.1.2. Updating Red Hat Enterprise Virtualization Manager	164
17.1.3. Updating Red Hat Enterprise Virtualization Manager Reports	164
17.1.4. Updating Red Hat Enterprise Virtualization Hypervisors	165
17.1.5. Updating Red Hat Enterprise Linux Virtualization Hosts	165
17.1.6. Updating the Red Hat Enterprise Virtualization Guest Tools	166
17.2. Upgrading to Red Hat Enterprise Virtualization 3.2	166
17.2.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.2	166
17.2.2. Upgrading to Red Hat Enterprise Virtualization Manager 3.1	167
17.2.3. Changing the Cluster Compatibility Version	168
17.2.4. Changing the Data Center Compatibility Version	169
<b>Chapter 18.. Backups .....</b>	<b>170</b>
18.1. Backing Up the Engine Database Using the backup.sh Script	170

18.2. Restoring the Engine Database Using the restore.sh Script	170
18.3. Backing Up Manager Configuration Files	170
18.4. Restoring Red Hat Enterprise Virtualization Manager Configuration Files	171
<b>Part III. Gathering Information About the Environment</b>	<b>173</b>
<b>Chapter 19.. Reports, History Database Reports, and Dashboards</b>	<b>174</b>
19.1. Reports	174
19.1.1. Reports	174
19.1.2. JasperReports and JasperServer in Red Hat Enterprise Virtualization	174
19.1.3. Online Help for JasperReports	174
19.1.4. Jasper Reports System Requirements	174
19.1.5. Users in the Red Hat Enterprise Virtualization Reports Portal	174
19.1.6. Logging in to Access the Reports Portal	174
19.1.7. Accessing the Red Hat Enterprise Virtualization Reports User Management Menu	175
19.1.8. Reports Portal User Roles	175
19.1.9. Navigating Reports and Dashboards	176
19.1.10. Report Parameters	176
19.1.11. Right-click Reporting Integration with the Red Hat Enterprise Virtualization Administration Portal	177
19.1.12. Executive Reports	178
19.1.12.1. Executive reports: Active Virtual Machines by OS	178
19.1.12.2. Executive Reports: Cluster Capacity Vs Usage	178
19.1.12.3. Executive Reports: Host Operating System Break Down	178
19.1.12.4. Executive Reports: Summary of Host Usage Resources	179
19.1.13. Inventory Reports	179
19.1.13.1. Inventory Reports: Hosts Inventory	179
19.1.13.2. Inventory Reports: Storage Domain Over Time	179
19.1.13.3. Inventory Reports: Virtual Machines Inventory	180
19.1.13.4. Inventory Reports: Cloud Provider Virtual Machine Inventory	180
19.1.14. Service Level Reports	181
19.1.14.1. Service Level Reports: Cluster Host Uptime	181
19.1.14.2. Service Level Reports: Cluster Quality of Service - Hosts	181
19.1.14.3. Service Level Reports: Cluster Quality of Service - Virtual Machines	182
19.1.14.4. Service Level Reports: Single Host Uptime	182
19.1.14.5. Service Level Reports: Top 10 Downtime Hosts	183
19.1.14.6. Service Level Reports: High Availability Virtual Servers Uptime	183
19.1.15. Trend Reports	184
19.1.15.1. Trend Reports: Five Least Utilized Hosts (Over Time)	184
19.1.15.2. Trend Reports: Five Least Utilized Virtual Machines (Over Time)	184
19.1.15.3. Trend Reports: Five Most Utilized Hosts (Over Time)	185
19.1.15.4. Trend Reports: Five Most Utilized Virtual Machines (Over Time)	185
19.1.15.5. Trend Reports: Multiple Hosts Resource Usage (Over Time)	186
19.1.15.6. Trend Reports: Multiple Virtual Machines Resource Usage (Over Time)	186
19.1.15.7. Trend Reports: Single Host Resource Usage (Days of Week)	187
19.1.15.8. Trend Reports: Single Host Resource Usage (Hour of Day)	187
19.1.15.9. Trend Reports: Single Virtual Machine Resources (Days of Week)	188
19.1.15.10. Trend Reports: Single Virtual Machine Resources (Hour of Day)	188
19.1.15.11. Trend Reports: Single Virtual Machine Resources (Over Time)	189
19.1.16. Ad Hoc Reports	189
19.1.17. Reports Schema: Tag History and ENUM Views	190
19.2. History Database Reports	190
19.2.1. Red Hat Enterprise Virtualization History Database	190
19.2.2. Tracking Configuration History	190
19.2.3. Recording Statistical History	191
19.2.4. Tracking Tag History	191

19.2.5. Connecting to the History Database	191
19.2.6. Allowing Read-only Access to the History Database	192
19.2.7. History Database Report Examples	192
19.2.8. Statistics History Views	194
19.2.8.1. Statistics History Views	194
19.2.8.2. Datacenter Statistics Views	194
19.2.8.3. Storage Domain Statistics Views	195
19.2.8.4. Host Statistics Views	195
19.2.8.5. Host Interface Statistics Views	197
19.2.8.6. Virtual Machine Statistics Views	197
19.2.8.7. Virtual Machine Interface Statistics Views	198
19.2.8.8. Virtual Machine Disk Statistics Views	199
19.2.9. Configuration History Views	200
19.2.9.1. Configuration History Views	200
19.2.9.2. Data Center Configuration	200
19.2.9.3. Datacenter Storage Domain Map	200
19.2.9.4. Storage Domain Configuration	200
19.2.9.5. Cluster Configuration	200
19.2.9.6. Host Configuration	201
19.2.9.7. Host Interface Configuration	201
19.2.9.8. Virtual Machine Configuration	202
19.2.9.9. Virtual Machine Interface Configuration History	203
19.2.9.10. Disks to Virtual Machine Map	204
19.2.9.11. Virtual Machine Disk Configuration	204
19.3. Dashboards	204
19.3.1. Dashboards	204
19.3.2. Header Inventory Dashboard	205
19.3.3. Header Trends Dashboard	205
19.3.4. Header Uptime Dashboard	205
19.3.5. Integrated Reporting Dashboard in the Red Hat Enterprise Virtualization Administration Portal	205
<b>Firewalls</b>	<b>207</b>
A.1. Red Hat Enterprise Virtualization Manager Firewall Requirements	207
A.2. Virtualization Host Firewall Requirements	207
A.3. Directory Server Firewall Requirements	208
A.4. Database Server Firewall Requirements	209
<b>VDSM and Hooks</b>	<b>210</b>
B.1. VDSM	210
B.2. VDSM Hooks	210
B.3. Extending VDSM with Hooks	210
B.4. Supported VDSM Events	210
B.5. The VDSM Hook Environment	210
B.6. The VDSM Hook Domain XML Object	211
B.7. Defining Virtual Machine Custom Properties	211
B.8. Setting Virtual Machine Custom Properties	211
B.9. Evaluating Virtual Machine Custom Properties in a VDSM Hook	211
B.10. Using the VDSM Hooking Module	212
B.11. VDSM Hook Execution	212
B.12. VDSM Hook Return Codes	212
B.13. VDSM Hook Examples	213
<b>Red Hat Enterprise Virtualization User Interface Plugins</b>	<b>214</b>
C.1. Red Hat Enterprise Virtualization User Interface Plugins	214
C.2. Red Hat Enterprise Virtualization User Interface Plugin Lifecycle	214

C.2.1. Red Hat Enterprise Virtualization User Interface Plugin Lifecycle	214
C.2.2. Red Hat Enterprise Virtualization User Interface Plugin Discovery	214
C.2.3. Red Hat Enterprise Virtualization User Interface Plugin Loading	214
C.2.4. Red Hat Enterprise Virtualization User Interface Plugin Bootstrapping	214
C.3. User Interface Plugin-related Files and their Locations	215
C.4. Example User Interface Plugin Deployment	215
C.5. Installing Red Hat Support Plugin	215
C.6. Using Red Hat Support Plugin	216
<b>Red Hat Enterprise Virtualization and SSL</b>	<b>218</b>
D.1. Replacing the Red Hat Enterprise Virtualization Manager SSL Certificate	218
<b>Using Search, Bookmarks, and Tags to Find Your Way Around</b>	<b>219</b>
E.1. Search	219
E.1.1. Performing Searches in Red Hat Enterprise Virtualization	219
E.1.2. Search Syntax and Examples	219
E.1.3. Search Auto-Completion	219
E.1.4. Search Result Type Options	219
E.1.5. Search Criteria	219
E.1.6. Search: Multiple Criteria and Wildcards	220
E.1.7. Search: Determining Search Order	220
E.1.8. Searching for Data Centers	220
E.1.9. Searching for Clusters	221
E.1.10. Searching for Hosts	221
E.1.11. Searching for Networks	222
E.1.12. Searching for Storage	222
E.1.13. Searching for Disks	222
E.1.14. Searching for Volumes	223
E.1.15. Searching for Virtual Machines	223
E.1.16. Searching for Pools	224
E.1.17. Searching for Templates	225
E.1.18. Searching for Users	225
E.1.19. Searching for Events	226
E.2. Bookmarks	226
E.2.1. Saving a Query String as a Bookmark	226
E.2.2. Editing a Bookmark	226
E.2.3. Deleting a Bookmark	227
E.3. Tags	227
E.3.1. Using Tags to Customize Interactions with Red Hat Enterprise Virtualization	227
E.3.2. Creating a Tag	227
E.3.3. Modifying a Tag	227
E.3.4. Deleting a Tag	227
E.3.5. Adding and Removing Tags to and from Objects	228
E.3.6. Searching for Objects using Tags	228
<b>Revision History</b>	<b>229</b>





## Preface

### 1. 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](#) 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 include the Liberation Fonts set by default.

#### 1.1. 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 keys 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 key, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from an individual key by the plus sign that connects each part of a key combination. For example:

Press **Enter** to execute the command.

Press **Ctrl+Alt+F2** to switch to a virtual terminal.

The first example highlights a particular key to press. The second example highlights a key combination: a set of three keys pressed simultaneously.

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, select 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:

Publi<sup>can</sup> is a DocBook publishing system.

#### 1.2. 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:

```
static int kvm_vm_ioctl_deassign_device(struct kvm *kvm,
                                       struct kvm_assigned_pci_dev *assigned_dev)
{
    int r = 0;
    struct kvm_assigned_dev_kernel *match;

    mutex_lock(&kvm->lock);

    match = kvm_find_assigned_dev(&kvm->arch.assigned_dev_head,
                                assigned_dev->assigned_dev_id);
    if (!match) {
        printk(KERN_INFO "%s: device hasn't been assigned before, "
                   "so cannot be deassigned\n", __func__);
        r = -EINVAL;
        goto out;
    }

    kvm_deassign_device(kvm, match);

    kvm_free_assigned_device(kvm, match);

out:
    mutex_unlock(&kvm->lock);
    return r;
}
```

### 1.3. 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.

## 2. Getting Help and Giving Feedback

### 2.1. Do You Need Help?

If you experience difficulty with a procedure described in this documentation, visit the Red Hat Customer Portal at <http://access.redhat.com>. Through the customer portal, you can:

- » search or browse through a knowledgebase of technical support articles about Red Hat products.
- » submit a support case to Red Hat Global Support Services (GSS).
- » access other product documentation.

Red Hat also hosts a large number of electronic mailing lists for discussion of Red Hat software and technology. You can find a list of publicly available mailing lists at <https://www.redhat.com/mailman/listinfo>. Click on the name of any mailing list to subscribe to that list or to access the list archives.

### 2.2. 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 in Bugzilla: <http://bugzilla.redhat.com/> against the product **Red Hat Enterprise Virtualization Manager**.

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, please include the section number and some of the surrounding text so we can find it easily.

## Chapter 1. Using this Guide

### 1.1. Administration Guide Prerequisites

You need a functioning Red Hat Enterprise Virtualization environment to use this guide. You can use the *Red Hat Enterprise Virtualization Installation Guide* or the *Red Hat Enterprise Virtualization Quick Start Guide* to install your environment and complete the initial configuration tasks.

A basic Red Hat Enterprise Virtualization environment has:

- » at least one data center,
- » at least one cluster,
- » at least one host,
- » at least one data storage domain,
- » at least one logical network: the **rhev** management network,
- » and at least one user: the internal **admin** user.

The Red Hat Enterprise Virtualization Administration Guide contains information about managing existing Red Hat Enterprise Virtualization environments. If your environment is missing one of the listed elements, please find the topic in this guide or in the *Installation Guide* or *Quick Start Guide* that describes how to add what your environment is missing.

[Report a bug](#)

### 1.2. Administration Guide Layout

In the Red Hat Enterprise Virtualization Administration Guide it is assumed that administrators want to perform actions on objects or with objects. For example, you want to add a new logical network to a cluster. "Add a new logical network" is an action, and "cluster" is an object.

The Red Hat Enterprise Virtualization Administration Guide uses objects to group content. The objects are ordered according to their likely order of usage by administrators. The objects are:

- » Data Centers;
- » Clusters;
- » Networks;
- » Hosts;
- » Storage;
- » Virtual Machines, Templates, and Pools;
- » Users and Roles;
- » Quotas;
- » Monitoring, Reports, and Dashboards;
- » Firewalls;
- » VDSM and Hooks;
- » Utilities; and
- » Backups.

To use this guide, find the object you are interested in affecting, then find the action or task you want to perform.

[Report a bug](#)

### 1.3. Example Workflows

#### 1.3.1. Administration Guide Example Workflows Overview

Example workflows can help you become comfortable with using the Red Hat Enterprise Virtualization Administration Guide. They are common tasks performed by administrators of Red Hat Enterprise Virtualization environments. Each workflow begins with a scenario, and then gives links to the tasks for each scenario in the order that they should be performed.

[Report a bug](#)

#### 1.3.2. Administration Guide Example Workflow: New iSCSI Data Center

Your employer has purchased some new hypervisors and storage to add to your environment. All the hardware has been configured by your IT department. The storage is deployed as iSCSI storage. The hypervisors run Red Hat Enterprise Linux. The storage traffic is carried over a storage network separate from management traffic. Control over this hardware is delegated to one of your colleagues.

- » [Section 5.4.1, "Creating a New Logical Network in a Data Center or Cluster"](#)
- » [Section 7.6.2, "Adding iSCSI Storage"](#)
- » [Section 3.7.3, "Assigning an Administrator or User Role to a Resource"](#)

[Report a bug](#)

#### 1.3.3. Administration Guide Example Workflow: Newly Virtualized Workload

You have recently virtualized an important workload. You need to maximize the uptime of the virtual machine it runs on. You clone the virtual machine to a template so that it is easy to re-provision if necessary. You hand control of the virtual machine and the cluster it runs on to another administrator.

- » [Section 8.12.4, "Configuring a Highly Available Virtual Machine"](#)
- » [Section 9.2.1, "Creating a Template from an Existing Virtual Machine"](#)
- » [Section 8.11.2, "Live Migration Prerequisites"](#)
- » [Section 8.11.3, "Automatic Virtual Machine Migration"](#)
- » [Section 8.11.6, "Setting Migration Priority"](#)
- » [Section 6.6.1, "Host High Availability"](#)
- » [Section 6.6.3, "Setting Fencing Parameters on a Host"](#)
- » [Section 3.7.3, "Assigning an Administrator or User Role to a Resource"](#)

[Report a bug](#)

#### 1.3.4. Administration Guide Example Workflow: Template for Group Use

You have a group of users who want to provision virtual machines running Red Hat Enterprise Linux 6. You have to add an ISO storage domain and upload an ISO to it. You install Red Hat Enterprise Linux 6 on a virtual machine, and make a template out of it. You make the group template users.

- » [Section 7.4.1, "Preparing NFS Storage"](#)
- » [Section 7.4.2, "Attaching NFS Storage"](#)
- » [Section 7.7.2, "Populating the ISO Storage Domain"](#)
- » [Section 15.6.1, "ISO Uploader"](#)
- » [Section 8.4.2, "Creating a New Virtual Machine from a Blank Template"](#)
- » [Section 8.4.6, "Installing a Guest Operating System onto a Virtual Machine"](#)

- » [Section 9.2.1, "Creating a Template from an Existing Virtual Machine"](#)
- » [Section 12.5.1, "Adding Users"](#)
- » [Section 3.7.3, "Assigning an Administrator or User Role to a Resource"](#)

[Report a bug](#)

## Chapter 2. Basics

### 2.1. Introduction

#### 2.1.1. Red Hat Enterprise Virtualization Architecture

A Red Hat Enterprise Virtualization environment consists of:

- ▀ Virtual machine **hosts** using the Kernel-based Virtual Machine (KVM).
- ▀ **Agents and tools** running on hosts including VDSM, QEMU, and libvirt. These tools provide local management for virtual machines, networks and storage.
- ▀ The **Red Hat Enterprise Virtualization Manager**; a centralized management platform for the Red Hat Enterprise Virtualization environment. It provides a graphical interface where you can view, provision and manage resources.
- ▀ **Storage domains** to hold virtual resources like virtual machines, templates, ISOs.
- ▀ A **database** to track the state of and changes to the environment.
- ▀ Access to an external **Directory Server** to provide users and authentication.
- ▀ **Networking** to link the environment together. This includes physical network links, and logical networks.

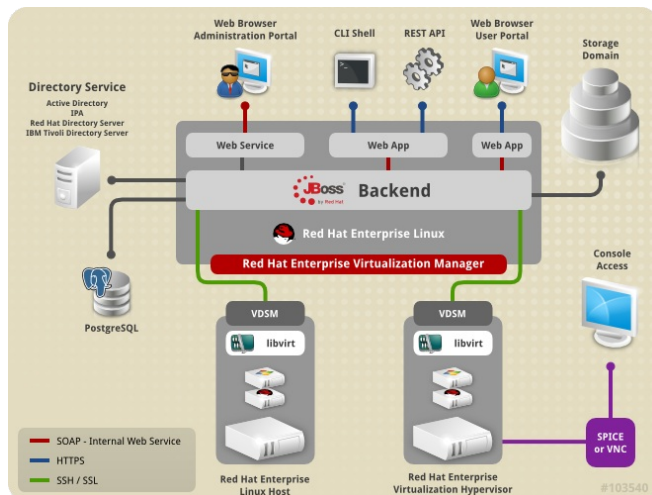


Figure 2.1. Red Hat Enterprise Virtualization Platform Overview

[Report a bug](#)

#### 2.1.2. Red Hat Enterprise Virtualization System Components

The Red Hat Enterprise Virtualization environment consists of one or more hosts (either Red Hat Enterprise Linux 6.3 or later hosts or Red Hat Enterprise Virtualization Hypervisor hosts) and at least one Red Hat Enterprise Virtualization Manager.

Hosts run virtual machines using KVM virtualization technology (Kernel-based Virtual Machine).

The Red Hat Enterprise Virtualization Manager runs on a Red Hat Enterprise Linux 6 server and provides interfaces for controlling the Red Hat Enterprise Virtualization environment. It manages virtual machine and storage provisioning, connection protocols, user sessions, virtual machine images, and high availability virtual machines.

The Red Hat Enterprise Virtualization Manager is accessed through the Administration Portal using a web browser.

[Report a bug](#)

#### 2.1.3. Red Hat Enterprise Virtualization Resources

The components of the Red Hat Enterprise Virtualization environment fall into two categories: physical resources, and logical resources. Physical resources are physical objects, such as host and storage servers. Logical resources are non-physical groupings and processes, such as logical networks and virtual machine templates.

- ▀ **Data Center** - A data center is the highest level container for all physical and logical resources within a managed virtual environment. It is a collection of clusters of virtual machines, storage, and networks.
- ▀ **Clusters** - A cluster is a set of physical hosts that are treated as a resource pool for virtual machines. Hosts in a cluster share the same network infrastructure and storage. They form a migration domain within which virtual machines can be moved from host to host.
- ▀ **Logical Networks** - A logical network is a logical representation of a physical network. Logical networks group network traffic and communication between the Manager, hosts, storage, and virtual machines.
- ▀ **Hosts** - A host is a physical server that runs one or more virtual machines. Hosts are grouped into clusters. Virtual machines can be migrated from one host to another within a cluster.
- ▀ **Storage Pool** - The storage pool is a logical entity that contains a standalone image repository of a certain type, either iSCSI, Fibre Channel, NFS, or POSIX. Each storage pool can contain several domains, for storing virtual machine disk images, 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 virtual machines can be created in a **Pool**. Virtual machines are created, managed, or deleted by power users and accessed by users.
- ▀ **Template** - A template is a model virtual machine with pre-defined settings. A virtual machine that is based on a particular template acquires the settings of the template. Using templates is the quickest way of creating a large number of virtual machines in a single step.
- ▀ **Virtual Machine Pool** - A virtual machine pool is a group of identical virtual machines that are available on demand by each group member. Virtual machine pools can be set up for different purposes. For example, one pool can be for the Marketing department, another for Research and Development, and so on.
- ▀ **Snapshot** - A snapshot is a view of a virtual machine's operating system and all its applications at a point in time. It can be used to save the settings of a virtual machine before an upgrade or installing new applications. In case of problems, a snapshot can be used to restore the virtual machine to its original state.

- **User Types** - Red Hat Enterprise Virtualization supports multiple levels of administrators and users with distinct levels of permissions. System administrators can manage objects of the physical infrastructure, such as data centers, hosts, and storage. Users access virtual machines available from a virtual machine pool or standalone virtual machines made accessible by an administrator.
- **Events and Monitors** - Alerts, warnings, and other notices about activities help the administrator to monitor the performance and status of resources.
- **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.

[Report a bug](#)

### 2.1.4. Red Hat Enterprise Virtualization API Support Statement

Red Hat Enterprise Virtualization exposes a number of interfaces for interacting with the components of the virtualization environment. These interfaces are in addition to the user interfaces provided by the Red Hat Enterprise Virtualization Manager Administration, User, and Reports Portals. Many of these interfaces are fully supported. Some however are supported only for read access or only when your use of them has been explicitly requested by Red Hat Support.

#### Supported Interfaces for Read and Write Access

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

##### Representational State Transfer (REST) API

The REST API exposed by the Red Hat Enterprise Virtualization Manager is a fully supported interface for interacting with Red Hat Enterprise Virtualization Manager.

##### Software Development Kit (SDK)

The SDK provided by the *rhevmsdk* package is a fully supported interface for interacting with Red Hat Enterprise Virtualization Manager.

##### Command Line Shell

The command line shell provided by the *rhevmscli* package is a fully supported interface for interacting with the Red Hat Enterprise Virtualization Manager.

##### VDSM Hooks

The creation and use of VDSM hooks to trigger modification of virtual machines based on custom properties specified in the Administration Portal is supported on Red Hat Enterprise Linux virtualization hosts. The use of VDSM Hooks on virtualization hosts running Red Hat Enterprise Virtualization Hypervisor is not currently supported.

#### Supported Interfaces for Read Access

Direct interaction with these interfaces is supported and encouraged only for read access. Use of these interfaces for write access is not supported unless explicitly requested by Red Hat Support:

##### Red Hat Enterprise Virtualization Manager History Database

Read access to the Red Hat Enterprise Virtualization Manager history database using the database views specified in the Administration Guide is supported. Write access is *not* supported.

##### Libvirt on Virtualization Hosts

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

#### Unsupported Interfaces

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

##### The vdsClient Command

Use of the **vdsClient** command to interact with virtualization hosts is *not* supported unless explicitly requested by Red Hat Support.

##### Red Hat Enterprise Virtualization Hypervisor Console

Console access to Red Hat Enterprise Virtualization Hypervisor outside of the provided text user interface for configuration is *not* supported unless explicitly requested by Red Hat Support.

##### Red Hat Enterprise Virtualization Manager Database

Direct access to and manipulation of the Red Hat Enterprise Virtualization Manager database is *not* supported unless explicitly requested by Red Hat Support.



#### Important

Red Hat Support will not debug user created scripts or hooks except where it can be demonstrated that there is an issue with the interface being used rather than the user created script itself. For more general information about Red Hat support policies see <https://access.redhat.com/support/offerings/production/soc.html>.

[Report a bug](#)

### 2.1.5. SPICE

The SPICE protocol facilitates graphical connections to virtual machines. The SPICE protocol allows:

- video at more than 30 frames per second
- bi-directional audio (for soft-phones/IP phones)
- bi-directional video (for video telephony/video conferencing)
- connection to multiple monitors with a single virtual machine
- USB redirection from the client's USB port into the virtual machine
- connection to a proxy from outside of the network the hypervisor is attached to

[Report a bug](#)

### 2.1.6. Administering and Maintaining the Red Hat Enterprise Virtualization Environment

The Red Hat Enterprise Virtualization environment requires an administrator to keep it running. As an

administrator, your tasks include:

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

[Report a bug](#)

## 2.2. Using the Administration Portal Graphical Interface

### 2.2.1. Graphical User Interface Elements

The Red Hat Enterprise Virtualization Administration Portal consists 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.

#### User Interface Elements

##### 1 Header

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

##### 2 Search Bar

The **Search** bar allows you to 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 are 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** tabs allow you to view events for each resource.

The Administration Portal provides the following tabs: Data Centers, Clusters, Hosts, Storage, Disks, Virtual Machines, Pools, Templates, Users, and Events, and a Dashboard tab if you have installed the Data Warehouse and Reporting services.

##### 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. Information on a selected item is displayed in the details pane.

##### 5 Details Pane

The **Details** pane shows detailed information about a selected item in the results list. If multiple items are selected, the details pane displays information on the first selected item only.

##### 6 Tree/Bookmarks/Tags Pane

The **Tree** pane displays a navigable hierarchy of the resources in the virtualized environment.

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

**Tags** are applied to groups of resources and are used to search for all resources associated with that tag.

##### 7 Alerts/Events Pane

The **Alerts** tab lists all high severity events such as errors or warnings. The **Events** tab shows an audit of events for all resources. The **Tasks** tab lists the current running tasks. You can view this panel by clicking the maximize/ minimize button.

**Important**

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

[Report a bug](#)

### 2.2.2. 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. Working in tree mode is highly recommended for most operations.

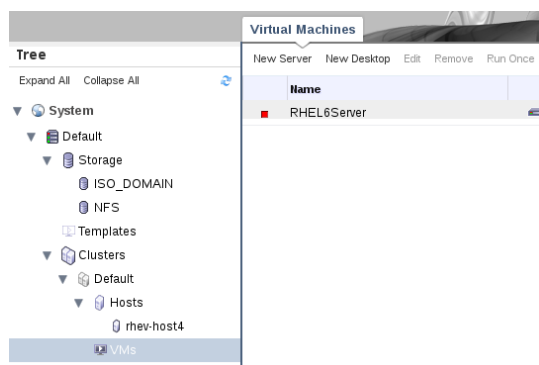


Figure 2.2. Tree Mode

Flat mode allows you to search across data centers, or storage domains. It does not limit you to viewing the resources of a single hierarchy. 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, such as **Pools** and **Users** are not in the data center hierarchy and can be accessed only in flat mode.

To access flat mode, click on the **System** item in the **Tree** pane on the left side of the screen. You are in flat mode if the **Pools** and **Users** resource tabs appear.




Figure 2.3. Flat Mode

[Report a bug](#)

### 2.2.3. Using the Guide Me Facility

When setting up resources such as data centers and clusters, a number of tasks must be completed in sequence. The context-sensitive **Guide Me** window prompts for actions that are appropriate to the resource being configured. The **Guide Me** window can be accessed at any time by clicking the **Guide Me** button on the resource toolbar.

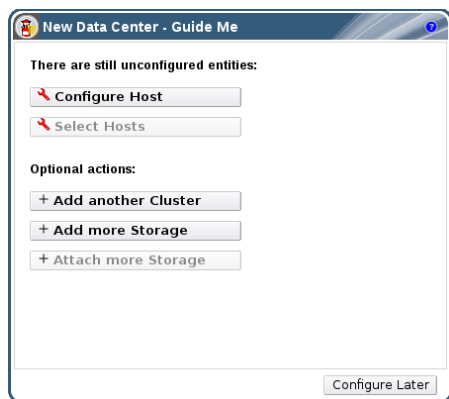


Figure 2.4. New Data Center Guide Me Window

[Report a bug](#)

### 2.2.4. Performing Searches in Red Hat Enterprise Virtualization

The Administration Portal enables the management of thousands of resources, such as virtual machines, hosts, users, and more. To perform a search, enter the search query (free-text or syntax-based) in the search bar. Search queries can be saved as bookmarks for future reuse, so you do not have to reenter a search query each time the specific search results are needed.

See Also:

- ▮ [Section E.1, "Search"](#)

[Report a bug](#)

### 2.2.5. Saving a Query String as a Bookmark

#### Summary

A bookmark can be used to remember a search query, and shared with other users.

#### Procedure 2.1. Saving a Query String as a Bookmark

1. Enter the desired search query in the search bar and perform the search.
2. Click the star-shaped **Bookmark** button to the right of the search bar to open the **New Bookmark** window.




Figure 2.5. Bookmark Icon

3. Enter the **Name** of the bookmark.
4. Edit the **Search string** field (if applicable).
5. Click **OK** to save the query as a bookmark and close the window.
6. The search query is saved and displays in the **Bookmarks** pane.

#### Result

You have saved a search query as a bookmark for future reuse. Use the **Bookmark** pane to find and select the bookmark.

See Also:

- ▮ [Section E.1, "Search"](#)



» [Section E.2, "Bookmarks"](#)

[Report a bug](#)

## **Part I. Administering the Resources**

## Chapter 3. Data Centers

### 3.1. Introduction to Data Centers

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

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

All data centers are managed from the single Administration Portal.

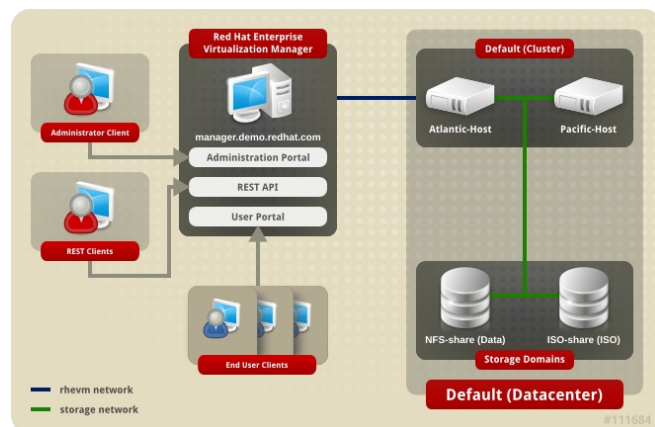


Figure 3.1. Data Centers

Red Hat Enterprise Virtualization creates a default data center during installation. It is recommended that you do not remove the default data center; instead, set up new appropriately named data centers.

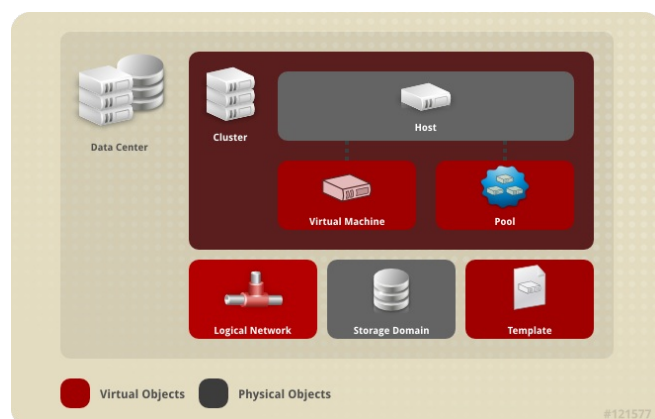


Figure 3.2. Data Center Objects

[Report a bug](#)

### 3.2. The Storage Pool Manager (SPM)

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

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

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

[Report a bug](#)

### 3.3. SPM Priority

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

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

See Also:

- ▮ [Section 6.4.4, "Configuring Host Storage Pool Manager \(SPM\) Settings"](#)

[Report a bug](#)

### 3.4. Using the Events Tab to Identify Problem Objects in Data Centers

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

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

[Report a bug](#)

## 3.5. Data Center Tasks

### 3.5.1. Creating a New Data Center

#### Summary

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



#### Note

The storage **Type** can be edited until the first storage domain is added to the data center. Once a storage domain has been added, the storage **Type** cannot be changed. If you set the **Compatibility Version** as **3.1**, it cannot be changed to **3.0** at a later time; version regression is not allowed.

#### Procedure 3.1. Creating a New Data Center

1. Select the **Data Centers** resource tab to list all data centers in the results list.
2. Click **New** to open the **New Data Center** window.
3. Enter the **Name** and **Description** of the data center.
4. Select the storage **Type**, **Compatibility Version**, and **Quota Mode** of the data center from the drop-down menus.
5. Click **OK** to create the data center and open the **New Data Center - Guide Me** window.
6. The **Guide Me** window lists the entities that need to be configured for the data center. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the data center and clicking the **Guide Me** button.

#### Result

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

#### See Also:

- » [Section 2.2.3, "Using the Guide Me Facility"](#)
- » [Section 3.5.2, "Explanation of Settings in the New Data Center and Edit Data Center Windows"](#)

[Report a bug](#)

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

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

Table 3.1. Data Center Properties

Field	Description/Action
<b>Name</b>	The name of the data center. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
<b>Description</b>	The description of the data center. This field is recommended but not mandatory.
<b>Type</b>	<p>The storage type. Choose one of</p> <ul style="list-style-type: none"> <li>» NFS</li> <li>» iSCSI</li> <li>» Fibre Channel</li> <li>» Local on Host</li> <li>» POSIX compliant FS</li> </ul> <p>The type of data domain dictates the type of the data center and cannot be changed after creation without significant disruption. 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.</p>
<b>Compatibility Version</b>	<p>The version of Red Hat Enterprise Virtualization. Choose one of:</p> <ul style="list-style-type: none"> <li>» 3.0</li> <li>» 3.1</li> <li>» 3.2</li> </ul> <p>After upgrading the Red Hat Enterprise Virtualization Manager, the hosts, clusters and data centers may still be in the earlier version. Ensure that you have upgraded all the hosts, then the clusters, before you upgrade the Compatibility Level of the data center.</p>
<b>Quota Mode</b>	<p>Quota is a resource limitation tool provided with Red Hat Enterprise Virtualization. Choose one of:</p> <ul style="list-style-type: none"> <li>» Disabled - Select if you do not want to implement Quota</li> <li>» Audit - Select if you want to edit the Quota settings</li> <li>» Enforced - Select to implement Quota</li> </ul>

[Report a bug](#)

### 3.5.3. Editing a Resource

#### Summary

Edit the properties of a resource.

#### Procedure 3.2. Editing a Resource

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

#### Result

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

#### See Also:

- » [Section 3.5.2, "Explanation of Settings in the New Data Center and Edit Data Center Windows"](#)

[Report a bug](#)

### 3.5.4. Creating a New Logical Network in a Data Center or Cluster

#### Summary

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

#### Procedure 3.3. Defining Logical Networks in a Cluster

1. Use the **Data Centers** or **Clusters** resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Logical Networks** tab of the details pane to list the existing logical networks.
3. Click **New** in the **Data Centers** details pane to open the **New Logical Network** window.  
Click **Add Network** in the **Clusters** details pane to open the **New Logical Network** window.

Figure 3.3. New Logical Network

4. Enter the **Name** and **Description** of the logical network.
5. Select the check boxes to enable a **VM Network**, **Enable VLAN tagging**, and **Override MTU**.
6. Select the cluster(s) you want to assign the network to. Note that the network is added as required Network to the selected clusters.
7. Tick the **Allow all users to use this Network** checkbox to confer permissions for the new network to all users.
8. Click **OK** to create the logical network.

**Result**

You have defined this logical network as a resource required by a cluster or clusters in the data center. You can now add this resource to the hosts in the cluster.

**See Also:**

- » [Section 5.4.2, "Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows"](#)

[Report a bug](#)

**3.5.5. Re-Initializing a Data Center: Recovery Procedure****Summary**

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

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

**Procedure 3.4. Re-Initializing a Data Center**

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

**Result**

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

[Report a bug](#)

**3.5.6. Removing a Data Center****Summary**

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

**Procedure 3.5. Removing a Data Center**

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

**Result**

The data center has been removed.

[Report a bug](#)

**3.5.7. Force Removing a Data Center****Summary**

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

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

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

#### Procedure 3.6. Force Removing a Data Center

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

#### Result

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

[Report a bug](#)

## 3.6. Data Centers and Storage Domains

### 3.6.1. Attaching an Existing Data Domain to a Data Center

#### Prerequisites:

- » [Section 7.4, "Preparing and Adding File-based Storage"](#)
- » [Section 7.6, "Adding Block-based Storage"](#)

#### Summary

Data domains that are **Unattached** can be attached to a data center. The data domain must be of the same **Storage Type** as the data center.

#### Procedure 3.7. Attaching an Existing Data Domain to a Data Center

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

#### Result

The data domain is attached to the data center and can be activated.

[Report a bug](#)

### 3.6.2. Attaching an Existing ISO domain to a Data Center

#### Summary

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

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

#### Procedure 3.8. Attaching an Existing ISO Domain to a Data Center

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

#### Result

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

[Report a bug](#)

### 3.6.3. Attaching an Existing Export Domain to a Data Center

#### Summary

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

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

#### Procedure 3.9. Attaching an Existing Export Domain to a Data Center

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

#### Result

The Export domain is attached to the data center and can be activated.

[Report a bug](#)

### 3.6.4. Detaching a Storage Domain from a Data Center

#### Summary

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

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



#### Note

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

**Procedure 3.10. Detaching a Storage Domain from a Data Center**

1. Click the **Data Centers** resource tab and select the appropriate data center.
2. Select the **Storage** tab in the details pane to list the storage domains attached to the data center.
3. Select the storage domain to detach. If the storage domain is **Active**, click **Maintenance** to move the domain into maintenance mode.
4. Click **Detach** to open the **Detach Storage** confirmation window.
5. Click **OK**.

**Result**

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

[Report a bug](#)

**3.6.5. Activating a Storage Domain from Maintenance Mode****Summary**

Storage domains in maintenance mode must be activated to be used.

**Procedure 3.11. Activating a Data Domain from Maintenance Mode**

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

**Result**

The storage domain is activated and can be used in the data center.

[Report a bug](#)

**3.7. Data Centers and Permissions****3.7.1. Managing System Permissions for a Data Center**

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **ClusterAdmin** has administrator privileges only for the assigned cluster, a **StorageAdmin** has administrator privileges only for the assigned storage domain, and so forth.

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

The data center administrator role permits the following actions:

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

**Note**

You can only assign roles and permissions to existing users.

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

**See Also:**

- » [Section 3.7.2, "Data Center Administrator Roles Explained"](#)

[Report a bug](#)

**3.7.2. Data Center Administrator Roles Explained****Data Center Permission Roles**

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

**Table 3.2. Red Hat Enterprise Virtualization System Administrator Roles**

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

**See Also:**

- » [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**3.7.3. Assigning an Administrator or User Role to a Resource****Summary**

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

**Procedure 3.12. Assigning a Role to a Resource**



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

**Result**

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

**See Also:**

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**3.7.4. Removing an Administrator or User Role from a Resource****Summary**

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

**Procedure 3.13. Removing a Role from a Resource**

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

**Result**

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

[Report a bug](#)

## Chapter 4. Clusters

### 4.1. Introduction to Clusters

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

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

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

The Red Hat Enterprise Virtualization platform contains a default cluster in the default data center at installation.

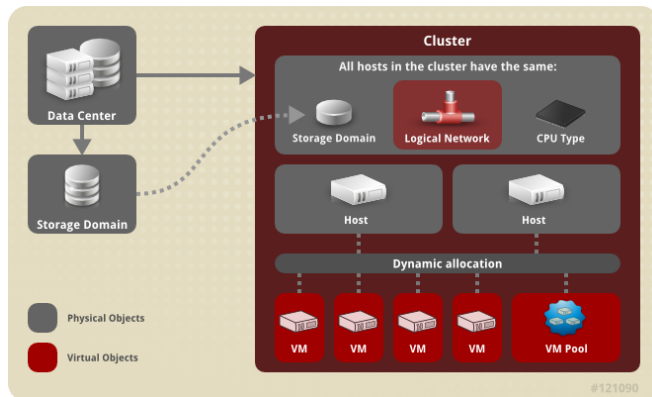


Figure 4.1. Cluster

#### See Also:

- » [Chapter 3, Data Centers](#)
- » [Chapter 5, Logical Networks](#)
- » [Chapter 6, Hosts](#)
- » [Chapter 7, Storage](#)

[Report a bug](#)

### 4.2. Cluster Tasks

#### 4.2.1. Creating a New Cluster

##### Prerequisites:

- » [Section 3.5.1, "Creating a New Data Center"](#)

##### Summary

A data center can contain multiple clusters and a cluster can hold multiple hosts. All hosts in a cluster must run the same CPU type (Intel or AMD). It is recommended that you create your hosts before you create your cluster to ensure CPU type optimization; however, you can configure the hosts at a later time using the **Guide Me** button.

**Important**

The default **rhev** network cannot be modified once a Cluster has been attached to a data center. Perform all configuration for the **rhev** network, such as enabling VLAN tagging, before a Cluster is attached and while the data center remains in the **Uninitialized** state.

##### Procedure 4.1. Creating a New Cluster

1. Select the **Clusters** resource tab to list all clusters in the results list.
2. Click **New** to open the **New Cluster** window.
3. Select the **Data Center** the cluster will belong to from the drop-down menu.
4. Enter the **Name** and **Description** of the cluster.
5. Select the **CPU Name** and **Compatibility Version** from the drop-down menus. It is important to match the CPU processor family with the minimum CPU processor type of the hosts you intend to attach to the cluster, otherwise the host will be non-operational.
6. Tick either the **Enable Virt Service** or **Enable Gluster Service** check boxes, depending on whether the cluster should be populated with virtual machine hosts or Gluster-enabled nodes. Note that you cannot add Red Hat Enterprise Virtualization Hypervisor hosts to a Gluster-enabled cluster.
7. Click the **Optimization** tab to select the memory page sharing threshold for the cluster, and optionally enable CPU thread handling on the cluster's hosts.
8. Click the **Resilience Policy** tab to select the virtual machine migration policy.
9. Click **OK** to create the cluster and open the **New Cluster - Guide Me** window.
10. The **Guide Me** window lists the entities that need to be configured for the cluster. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the cluster and clicking the **Guide Me** button.

##### Result

The new cluster is added to the virtualization environment.

##### See Also:

- » [Section 4.2.2, "Explanation of Settings and Controls in the New Cluster and Edit Cluster Windows"](#)
- » [Section 2.2.3, "Using the Guide Me Facility"](#)

[Report a bug](#)

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

4.2.2.1. General Cluster Settings Explained

The screenshot shows the 'New Cluster' window with the 'General' tab selected. The window has a sidebar with tabs: 'General', 'Optimization', 'Resilience Policy', and 'Cluster Policy'. The main area contains the following settings:

- Data Center:** A dropdown menu set to 'Default'.
- Name:** A text input field.
- Description:** A text input field.
- CPU Name:** A dropdown menu.
- Compatibility Version:** A dropdown menu set to '3.2'.
- Enable Virt Service:** An unselected radio button.
- Enable Gluster Service:** A selected radio button.
- Import existing gluster configuration:** A checked checkbox.
- Enter the details of any server in the cluster:** A section header.
- Address:** A text input field.
- Fingerprint:** A text input field with a small icon on the right.
- Root Password:** A text input field.

Figure 4.2. New Cluster window

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

Table 4.1. General Cluster Settings

Field	Description/Action
<b>Data Center</b>	The data center that will contain the cluster.
<b>Name</b>	The name of the cluster. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
<b>Description</b>	The description of the cluster. This field is recommended but not mandatory.
<b>CPU Name</b>	<p>The CPU type of the cluster. Choose one of:</p> <ul style="list-style-type: none"> <li>» Intel Conroe Family</li> <li>» Intel Penryn Family</li> <li>» Intel Nehalem Family</li> <li>» Intel Westmere Family</li> <li>» Intel SandyBridge Family</li> <li>» Intel Haswell</li> <li>» AMD Opteron G1</li> <li>» AMD Opteron G2</li> <li>» AMD Opteron G3</li> <li>» AMD Opteron G4</li> <li>» AMD Opteron G5</li> </ul> <p>All hosts in a cluster must run the same CPU type (Intel or AMD); this cannot be changed after creation without significant disruption. The CPU type should be set for the least powerful host. For example: an <b>Intel SandyBridge</b> host can attach to an <b>Intel Penryn</b> cluster; an <b>Intel Conroe</b> host cannot. Hosts with different CPU models will only use features present in all models.</p>
<b>Compatibility Version</b>	<p>The version of Red Hat Enterprise Virtualization. Choose one of:</p> <ul style="list-style-type: none"> <li>» 3.0</li> <li>» 3.1</li> <li>» 3.2</li> </ul> <p>You will not be able to select a version older than the version specified for the data center.</p>
<b>Enable Virt Service</b>	If this radio button is selected, hosts in this cluster will be used to run virtual machines.
<b>Enable Gluster Service</b>	If this radio button is selected, hosts in this cluster will be used as Red Hat Storage Server nodes, and not for running virtual machines. You cannot add a Red Hat Enterprise Virtualization Hypervisor host to a cluster with this option enabled.
<b>Import existing gluster configuration</b>	<p>This check box is only available if the <b>Enable Gluster Service</b> radio button is selected. This option allows you to import an existing Gluster-enabled cluster and all its attached hosts to Red Hat Enterprise Virtualization Manager.</p> <p>The following options are required for each host in the cluster that is being imported:</p> <ul style="list-style-type: none"> <li>» <b>Address:</b> Enter the IP or fully qualified domain name of the Gluster host server.</li> <li>» <b>Fingerprint:</b> Red Hat Enterprise Virtualization Manager fetches the host's fingerprint, to ensure you are connecting with the correct host.</li> <li>» <b>Root Password:</b> Enter the root password required for communicating with the host.</li> </ul>

**See Also:**

- » [Section 4.2.5, "Explanation of Settings in the Add Hosts Window"](#)

[Report a bug](#)

**4.2.2.2. Optimization Settings Explained**

Memory page sharing allows virtual machines to use up to 200% of their allocated memory by utilizing unused memory in other virtual machines. This process is based on the assumption that the virtual machines in your Red Hat Enterprise Virtualization environment will not all be running at full capacity at the same time, allowing unused memory to be temporarily allocated to a particular virtual machine.

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

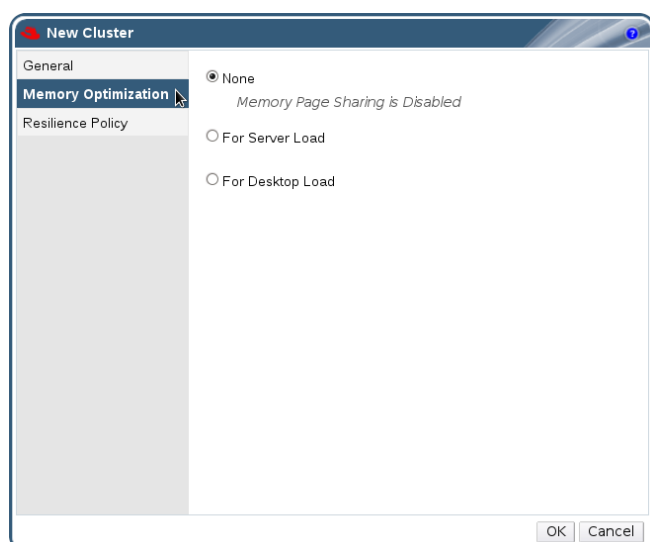


Figure 4.3. Optimization Settings

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

Table 4.2. Optimization Settings

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

[Report a bug](#)

#### 4.2.2.3. Resilience Policy Settings Explained

The resilience policy sets the virtual machine migration policy in the event of host failure. Virtual machines running on a host that unexpectedly shuts down or is put into maintenance mode are migrated to other hosts in the cluster; this migration is dependent upon your cluster policy.

#### 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.

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

Table 4.3. Resilience Policy Settings

Field	Description/Action
<b>Migrate Virtual Machines</b>	Migrates all virtual machines in order of their defined priority.
<b>Migrate only Highly Available Virtual Machines</b>	Migrates only highly available virtual machines to prevent overloading other hosts.
<b>Do Not Migrate Virtual Machines</b>	Prevents virtual machines from being migrated.

[Report a bug](#)

#### 4.2.2.4. Cluster Policy Settings Explained

Cluster policies allow you to specify the usage and distribution of virtual machines between available hosts. Define the cluster policy to enable automatic load balancing across the hosts in a cluster.

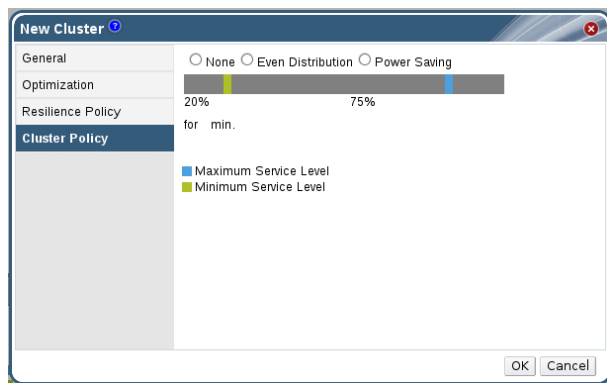


Figure 4.4. Power Management Policy Settings

The table below describes the settings for the **Edit Policy** window.

Table 4.4. Cluster General Tab Properties

Field/Tab	Description/Action
<b>None</b>	Set the policy value to <b>None</b> to have no load or power sharing between hosts. This is the default mode.
<b>Even Distribution</b>	Distributes the CPU processing load evenly across all hosts in the cluster. Additional virtual machines attached to a host will not start if that host has reached the defined Maximum Service Level.
<b>Power Saving</b>	Distributes the CPU processing load across a subset of available hosts to reduce power consumption on under-utilized hosts. Hosts with a CPU load below the minimum service level past the defined time interval will migrate all virtual machines to other hosts so that it can be powered down. Additional virtual machines attached to a host will not start if that host has reached the defined maximum service level.
<b>Maximum Service Level</b>	Set by the blue slider. If the host runs at maximum service level for the defined time interval, the Red Hat Enterprise Virtualization Manager will migrate virtual machines to other hosts in the cluster until the host's CPU load is below the maximum service threshold.
<b>Minimum Service Level</b>	Set by the green slider. If the host runs below minimum service level for the defined time interval, the Red Hat Enterprise Virtualization Manager will migrate virtual machines to other hosts in the cluster and power down the host.
<b>Time Interval Entry Field</b>	Sets the time (in minutes) that a host can run a CPU load outside of the defined service levels before the cluster policy takes action. The defined time interval protects against temporary spikes in CPU load activating cluster policies and instigating unnecessary virtual machine migration. Maximum two characters.

[Report a bug](#)

### 4.2.3. Editing a Resource

#### Summary

Edit the properties of a resource.

#### Procedure 4.2. Editing a Resource

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

#### Result

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

[Report a bug](#)

### 4.2.4. Importing an Existing Red Hat Storage Cluster

#### Summary

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

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



Currently, a Red Hat Storage node can only be added to a cluster which has its compatibility level set to **3.1**.

### Procedure 4.3. Importing an Existing Red Hat Storage Cluster to Red Hat Enterprise Virtualization Manager

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

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

Figure 4.5. Import Existing Cluster Configuration

6. In the **Address** field, enter the hostname or IP address of any server in the cluster.  
The host **Fingerprint** displays to ensure you are connecting with the correct host. If a host is unreachable or if there is a network error, an error **Error in fetching fingerprint** displays in the **Fingerprint** field.
7. Enter the **Root Password** for the server, and click **OK**.
8. The **Add Hosts** window opens, and a list of hosts that are a part of the cluster displays.

Name	Hostname/IP	Root Password	Fingerprint
10.70.35.154	10.70.35.154	*****	c2:72:9c:e5:8b:49:67:ab:6e:13:7b:2a:b0:af:d7:14
10.70.35.121	10.70.35.121	*****	08:12:48:5e:80:1f:84:43:e5:ae:f8:77:43:e3:d9:63

Figure 4.6. Add Hosts Window

9. For each host, enter the **Name** and the **Root Password**.
10. If you wish to use the same password for all hosts, select the **Use a Common Password** checkbox to enter the password in the provided text field.  
Click **Apply** to set the entered password all hosts.  
Make sure the fingerprints are valid and submit your changes by clicking **OK**.

#### Result

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

#### See Also:

- [Section 4.2.5, "Explanation of Settings in the Add Hosts Window"](#)
- [Section 4.2.2, "Explanation of Settings and Controls in the New Cluster and Edit Cluster Windows"](#)

[Report a bug](#)

### 4.2.5. Explanation of Settings in the Add Hosts Window

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

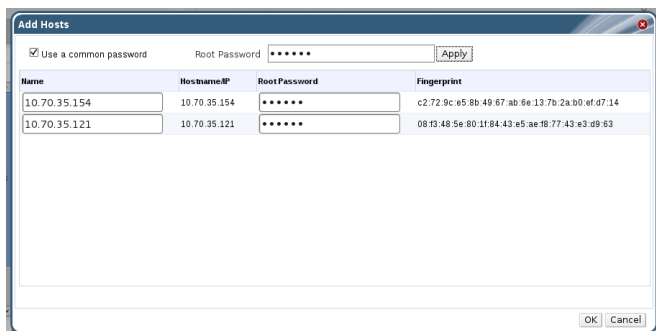


Figure 4.7. Add Hosts Window

Table 4.5. Add Gluster Hosts Settings

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

**See Also:**

- » [Section 4.2.4, "Importing an Existing Red Hat Storage Cluster"](#)

[Report a bug](#)

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

##### Summary

Cluster policies allow you to specify the usage and distribution of virtual machines between available hosts. Define the cluster policy to enable automatic load balancing across the hosts in a cluster.

A host that exceeds the **Maximum Service Level** will share its CPU processor load, additional virtual machines, to other hosts.

A host that does not exceed the **Minimum Service Level** will migrate all of its virtual machines to other hosts and power down until such time as it is required again.

##### Procedure 4.4. Setting Load and Power Management Policies for Hosts

1. Use the resource tabs, tree mode, or the search function to find and select the cluster in the results list.
2. Click the **Edit Policy** button, found in the **General** tab of the details pane, to open the **Edit Policy** window.

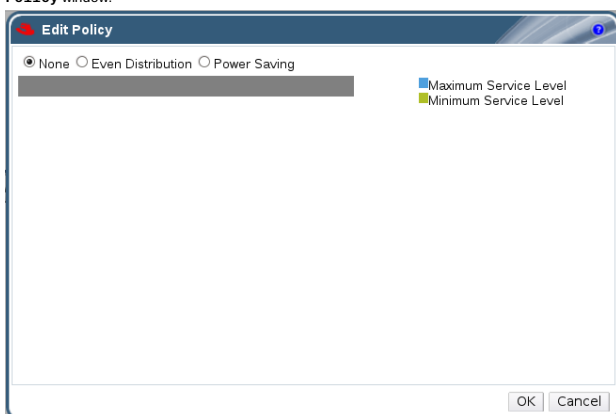


Figure 4.8. Edit Cluster Policy

3. Select one of the following policies:
  - » **None**
  - » **Even Distribution** - use the blue slider to specify the **Maximum Service Level** for a host.
  - » **Power Saving** - use the green slider to specify the **Minimum Service Level** for a host. Use the blue slider to specify the **Maximum Service Level** for a host.
4. Specify the defined time interval (in minutes) in the text field.
5. Click **OK**.
6. **Result**  
You have updated the cluster policy for the cluster.



**See Also:**

» [Section 4.2.4, "Cluster Policy Settings Explained"](#)

[Report a bug](#)

**4.2.7. Creating a New Logical Network in a Data Center or Cluster****Summary**

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

**Procedure 4.5. Defining Logical Networks in a Cluster**

1. Use the **Data Centers** or **Clusters** resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Logical Networks** tab of the details pane to list the existing logical networks.
3. Click **New** in the **Data Centers** details pane to open the **New Logical Network** window.

Click **Add Network** in the **Clusters** details pane to open the **New Logical Network** window.

Figure 4.9. New Logical Network

4. Enter the **Name** and **Description** of the logical network.
5. Select the check boxes to enable a **VM Network**, **Enable VLAN tagging**, and **Override MTU**.
6. Select the cluster(s) you want to assign the network to. Note that the network is added as required Network to the selected clusters.
7. Tick the **Allow all users to use this Network** checkbox to confer permissions for the new network to all users.
8. Click **OK** to create the logical network.

**Result**

You have defined this logical network as a resource required by a cluster or clusters in the data center. You can now add this resource to the hosts in the cluster.

**See Also:**

» [Section 5.4.2, "Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows"](#)

» [Section 5.3, "Required Networks, Optional Networks, and Virtual Machine Networks"](#)

[Report a bug](#)

**4.2.8. Removing a Cluster****Summary**

Move all hosts out of a cluster before removing it.

**Note**

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

**Procedure 4.6. Removing a Cluster**

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

**Result**

The cluster is removed.

[Report a bug](#)

**4.2.9. Designate a Specific Traffic Type for a Logical Network with the Assign/UnAssign Window****Prerequisites:**

» [Section 5.4.1, "Creating a New Logical Network in a Data Center or Cluster"](#)

**Summary**

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

**Procedure 4.7. Assigning or UnAssigning a Logical Network to a Cluster**

1. Use the **Clusters** resource tab, tree mode, or the search function to find and select the cluster in the results list.
2. Select the **Logical Networks** tab in the details pane to list the logical networks assigned to the cluster.
3. Click **Assign/UnAssign Networks** to open the **Assign/UnAssign Networks** window.

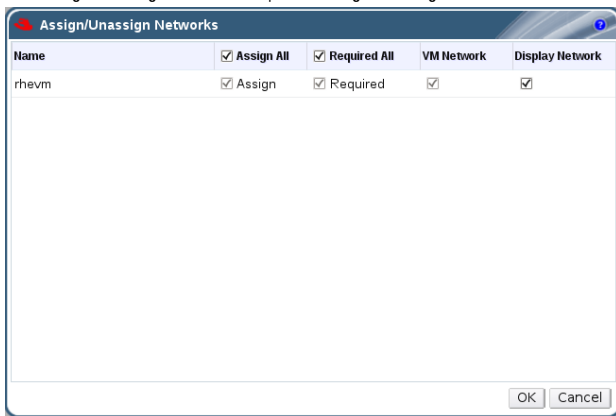


Figure 4.10. Assign/UnAssign Networks

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

**Result**

You have optimized the network traffic flow for the logical network.

**See Also:**

- » [Section 6.5.1, "Editing Host Network Interfaces and Adding Logical Networks to Hosts"](#)
- » [Section 4.2.10, "Explanation of Settings in the Assign/UnAssign Window"](#)

[Report a bug](#)

**4.2.10. Explanation of Settings in the Assign/UnAssign Window**

The table below describes the settings for the **Assigning/UnAssigning** window.

Table 4.6. Assign/UnAssign Networks Settings

Field	Description/Action
<b>Assign</b>	Assigns the logical network to all hosts in the cluster.
<b>Required</b>	A logical network becomes operational when it is attached to an active NIC on all hosts in the cluster.
<b>VM Network</b>	The logical network carries the virtual machine network traffic.
<b>Display Network</b>	The logical network carries the virtual machine SPICE and virtual network controller traffic.

[Report a bug](#)

**4.3. Clusters and Permissions****4.3.1. Managing System Permissions for a Cluster**

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **ClusterAdmin** has administrator privileges only for the assigned cluster, and so forth.

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

The cluster administrator role permits the following actions:

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

**Note**

You can only assign roles and permissions to existing users.

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

**See Also:**

- » [Section 4.3.2, "Cluster Administrator Roles Explained"](#)

[Report a bug](#)

**4.3.2. Cluster Administrator Roles Explained****Cluster Permission Roles**

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

Table 4.7. Red Hat Enterprise Virtualization System Administrator Roles

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

**See Also:**

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**4.3.3. Assigning an Administrator or User Role to a Resource****Summary**

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

**Procedure 4.8. Assigning a Role to a Resource**

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

**Result**

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

[Report a bug](#)

**4.3.4. Removing an Administrator or User Role from a Resource****Summary**

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

**Procedure 4.9. Removing a Role from a Resource**

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

**Result**

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

[Report a bug](#)

## Chapter 5. Logical Networks

### 5.1. Introduction to Logical Networks

A logical network is a named set of global network connectivity properties in your data center. When a logical network added to a host, it may be further configured with host-specific network parameters. Logical networks optimize network flow by grouping network traffic by usage, type, and requirements.

Logical networks allow both connectivity and segregation. You can create a logical network for storage communication to optimize network traffic between hosts and storage domains, a logical network specifically for all virtual machine traffic, or multiple logical networks to carry the traffic of groups of virtual machines.

The default logical network in all data centers is the management network called **rhev**. The **rhev** network carries all traffic, until another logical network is created. It is meant especially for management communication between the Red Hat Enterprise Virtualization Manager and hosts.

A logical network is a data center level resource; creating one in a data center makes it available to the clusters in a data center. A logical network that has been designated a **Required** must be configured in all of a cluster's hosts before it is operational. **Optional** networks can be used by any host they have been added to.

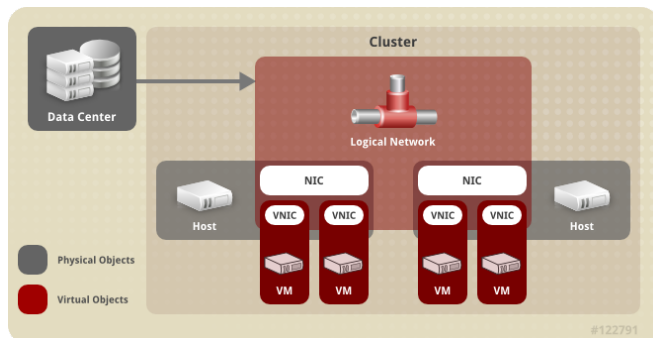


Figure 5.1. Data Center Objects



#### Warning

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



#### Important

If you plan to use Red Hat Enterprise Virtualization nodes to provide any services, remember that the services will stop if the Red Hat Enterprise Virtualization environment stops operating. This applies to all services, but you should be especially aware of the hazards of running the following on Red Hat Enterprise Virtualization:

- » Directory Services
- » DNS
- » Storage

#### See Also:

- » [Chapter 3, Data Centers](#)
- » [Chapter 4, Clusters](#)
- » [Chapter 6, Hosts](#)
- » [Chapter 7, Storage](#)
- » [Chapter 8, Virtual Machines](#)

[Report a bug](#)

### 5.2. Port Mirroring

Port mirroring copies layer 3 network traffic on given a logical network and host to a virtual interface on a virtual machine. This virtual machine can be used for network debugging and tuning, intrusion detection, and monitoring the behavior of other virtual machines on the same host and logical network.

The only traffic copied is internal to one logical network on one host. There is no increase on traffic on the network external to the host; however a virtual machine with port mirroring enabled uses more host CPU and RAM than other virtual machines.

Enable and disable port mirroring by editing network interfaces on virtual machines.

Port mirroring requires an IPv4 IP address.



#### Important

You should be aware that enabling port mirroring reduces the privacy of any other network users.

#### See Also:

- » [Section 8.7.4, "Adding and Editing Virtual Machine Network Interfaces"](#)

[Report a bug](#)

### 5.3. Required Networks, Optional Networks, and Virtual Machine Networks

In versions of Red Hat Enterprise Virtualization prior to 3.1, all of the hosts in a cluster had to use all of the cluster's logical network to be operational. Hosts that were missing logical networks were marked **Non-operational**.

Red Hat Enterprise Virtualization 3.1 and higher distinguishes between required networks and optional

networks.

**Required** networks must be applied to all hosts in a cluster for the cluster and network to be **Operational**. Logical networks are added to clusters as **Required** networks by default.

Optional networks are those logical networks that have not been explicitly declared **Required** networks. Optional networks can be implemented on only the hosts that use them. The presence or absence of these networks does not affect the **Operational** status of a host.

Use the **Assign/UnAssign Network** button to change a network's **Required** designation.

Virtual machine networks (called a **VM network** in the user interface) are logical networks designated to carry only virtual machine network traffic. Virtual machine networks can be required or optional.

#### Note

A virtual machine with a network interface on an optional virtual machine network will not start on a host without the network.

#### See Also:

- » [Section 5.4.2, "Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows"](#)

[Report a bug](#)

## 5.4. Logical Network Tasks

### 5.4.1. Creating a New Logical Network in a Data Center or Cluster

#### Summary

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

#### Procedure 5.1. Defining Logical Networks in a Cluster

1. Use the **Data Centers** or **Clusters** resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the **Logical Networks** tab of the details pane to list the existing logical networks.
3. Click **New** in the **Data Centers** details pane to open the **New Logical Network** window.

Click **Add Network** in the **Clusters** details pane to open the **New Logical Network** window.

Figure 5.2. New Logical Network

4. Enter the **Name** and **Description** of the logical network.
5. Select the check boxes to enable a **VM Network**, **Enable VLAN tagging**, and **Override MTU**.
6. Select the cluster(s) you want to assign the network to. Note that the network is added as required Network to the selected clusters.
7. Tick the **Allow all users to use this Network** checkbox to confer permissions for the new network to all users.
8. Click **OK** to create the logical network.

#### Result

You have defined this logical network as a resource required by a cluster or clusters in the data center. You can now add this resource to the hosts in the cluster.

#### See Also:

- » [Section 5.4.2, "Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows"](#)
- » [Section 5.3, "Required Networks, Optional Networks, and Virtual Machine Networks"](#)

[Report a bug](#)

Figure 5.3. New Logical Network Settings

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

Table 5.1. New Logical Network Settings

Field Name	Description
<b>Name</b>	The name of the logical network. This text field has a 15-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
<b>Description</b>	The description of the logical network. This field is recommended but not mandatory.
<b>VM Network</b>	Select this option if only virtual machines use this network. If the network is used for traffic that does not involve virtual machines, like storage communications for example, do not select <b>VM Network</b> .
<b>Enable VLAN tagging</b>	VLAN tagging is a security feature that gives all network traffic carried on the logical network a special characteristic. VLAN-tagged traffic cannot be read by interfaces that do not also have that characteristic. Use of VLANs on logical networks also allows a single network interface to be associated with multiple, differently VLAN-tagged logical networks. Enter a numeric value in the text entry field if VLAN tagging is enabled.
<b>Override MTU</b>	Set a custom maximum transmission unit for the logical network. You can use this to match the MTU supported by your new logical network to the MTU supported by the hardware it interfaces with. Enter a numeric value in the text entry field if MTU override is enabled.
<b>Attach/Detach Network to/from Cluster(s)</b>	Select the clusters that use the logical network. Logical networks are added to clusters as <b>Required</b> networks by default. You can change this later using the <b>Assign/UnAssign Network</b> button.
<b>Allow all users to use this Network</b>	Confers <b>NetworkUser</b> permissions on the new network to <b>Everyone</b> . This allows any user in the system with permissions on the relevant virtual machine or template to connect the network to the virtual machine or template.

[Report a bug](#)

### 5.4.3. Editing a Logical Network

#### Summary

Edit the settings of a logical network.

#### Procedure 5.2. Editing a Logical Network

1. Use the **Data Centers** resource tab, tree mode, or the search function to find and select the data center of the logical network in the results list.
2. Click the **Logical Networks** tab in the details pane to list the logical networks in the data center.
3. Select a logical network and click **Edit** to open the **Edit Logical Network** window.
4. You are required to detach all clusters before editing the logical network settings. Clear the cluster check boxes in the **Attach/Detach Network to/from Cluster(s)** area and click **Apply** to open the **Attach/Detach Network to/from Clusters** confirmation window.  
Select the **Approve operation** check box and click **OK** to detach the clusters and enable editing of the logical network.
5. Edit the necessary settings.
6. Attach the necessary clusters.
7. Click **OK** to save the changes.

#### Result

You have updated the settings of your logical network.

See Also:

- » [Section 5.4.2, "Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows"](#)

[Report a bug](#)

5.4.4. Designate a Specific Traffic Type for a Logical Network with the Assign/UnAssign Window

Prerequisites:

- » [Section 5.4.1, "Creating a New Logical Network in a Data Center or Cluster"](#)

Summary

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

Procedure 5.3. Assigning or UnAssigning a Logical Network to a Cluster

1. Use the **Clusters** resource tab, tree mode, or the search function to find and select the cluster in the results list.
2. Select the **Logical Networks** tab in the details pane to list the logical networks assigned to the cluster.
3. Click **Assign/UnAssign Networks** to open the **Assign/UnAssign Networks** window.

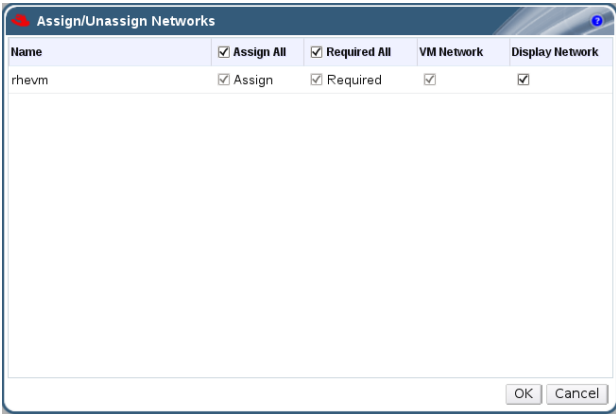


Figure 5.4. Assign/UnAssign Networks

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

Result

You have optimized the network traffic flow for the logical network.

See Also:

- » [Section 6.5.1, "Editing Host Network Interfaces and Adding Logical Networks to Hosts"](#)
- » [Section 5.4.5, "Explanation of Settings in the Assign/UnAssign Window"](#)

[Report a bug](#)

5.4.5. Explanation of Settings in the Assign/UnAssign Window

The table below describes the settings for the **Assigning/UnAssigning** window.

Table 5.2. Assign/UnAssign Networks Settings

Field	Description/Action
<b>Assign</b>	Assigns the logical network to all hosts in the cluster.
<b>Required</b>	A logical network becomes operational when it is attached to an active NIC on all hosts in the cluster.
<b>VM Network</b>	The logical network carries the virtual machine network traffic.
<b>Display Network</b>	The logical network carries the virtual machine SPICE and virtual network controller traffic.

[Report a bug](#)


5.4.6. Adding Multiple VLANs to a Single Network Interface using Logical Networks

Prerequisites:

- » [Section 5.4.1, "Creating a New Logical Network in a Data Center or Cluster"](#)

Summary

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

 **Important**

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

Procedure 5.4. Adding Multiple VLANs to a Network Interface using Logical Networks

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select in the results list a host associated with the cluster to which your VLAN-tagged logical networks are assigned.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the data center.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Drag your VLAN-tagged logical networks into the **Assigned Logical Networks** area next to

the physical network interface. The physical network interface can have multiple logical networks assigned due to the VLAN tagging.

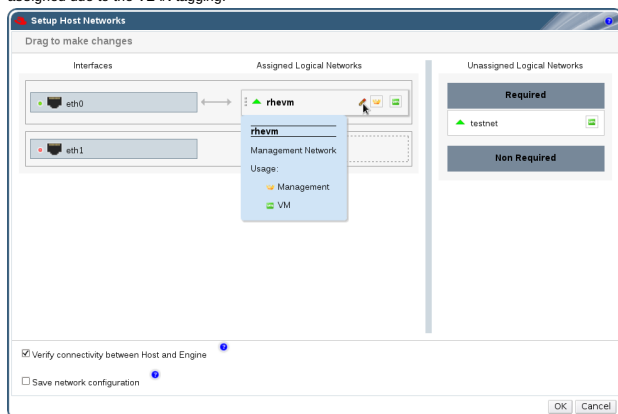


Figure 5.5. Setup Host Networks

5. Edit the logical networks by hovering your cursor over an assigned logical network and clicking the pencil icon to open the **Edit Network** window.

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

Select a **Boot Protocol** from:

- **None**,
  - **DHCP**, or
  - **Static**,
- Provide the **IP** and **Subnet Mask**.

Click **OK**.

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

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

#### Result

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

#### See Also:

- [Section 6.5, “Hosts and Networking”](#)

[Report a bug](#)

### 5.4.7. Using the Networks Tab

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

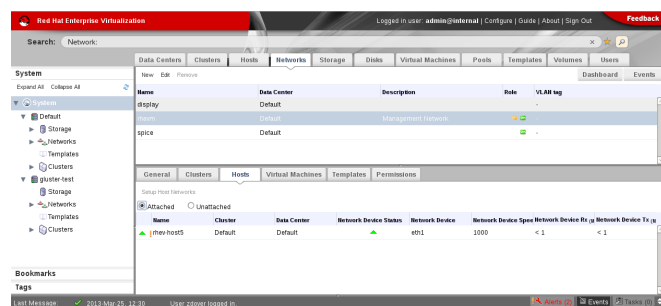


Figure 5.6. Networks Tab

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

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

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

These functions are also accessible through each individual resource tab. See the links below for how to set up networks on resources.

- [Section 5.4.1, “Creating a New Logical Network in a Data Center or Cluster”](#)
- [Section 4.2.9, “Designate a Specific Traffic Type for a Logical Network with the Assign/UnAssign Window”](#)
- [Section 5.4.6, “Adding Multiple VLANs to a Single Network Interface using Logical Networks”](#)
- [Section 8.7.4, “Adding and Editing Virtual Machine Network Interfaces”](#)
- [Section 3.7.3, “Assigning an Administrator or User Role to a Resource”](#)

[Report a bug](#)



## 5.5. Logical Networks and Permissions

### 5.5.1. Managing System Permissions for a Network

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **HostAdmin** has administrator privileges only for the assigned host, and so forth.

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

The network administrator role permits the following actions:

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

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

See Also:

- ▀ [Section 5.5.2, "Network Administrator and User Roles Explained"](#)

[Report a bug](#)

### 5.5.2. Network Administrator and User Roles Explained

#### Network Permission Roles

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

Table 5.3. Red Hat Enterprise Virtualization Network Administrator and User Roles

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

See Also:

- ▀ [Section 5.5.1, "Managing System Permissions for a Network"](#)

[Report a bug](#)

### 5.5.3. Assigning an Administrator or User Role to a Resource

#### Summary

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

#### Procedure 5.5. Assigning a Role to a Resource

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

#### Result

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

See Also:

- ▀ [Section 5.5.1, "Managing System Permissions for a Network"](#)

[Report a bug](#)

### 5.5.4. Removing an Administrator or User Role from a Resource

#### Summary

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

#### Procedure 5.6. Removing a Role from a Resource

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

5. Click **OK** to remove the user role.

**Result**

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

[Report a bug](#)

## Chapter 6. Hosts

### 6.1. Introduction to Red Hat Enterprise Virtualization Hosts

Hosts, also known as hypervisors, are the physical servers on which virtual machines run. Full virtualization is provided by using a loadable Linux kernel module called Kernel-based Virtual Machine (KVM).

KVM can concurrently host multiple virtual machines running either Windows or Linux operating systems. Virtual machines run as individual Linux processes and threads on the host machine and are managed remotely by the Red Hat Enterprise Virtualization Manager. A Red Hat Enterprise Virtualization environment has one or more hosts attached to it.

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

Red Hat Enterprise Virtualization hosts take advantage of **tuned** profiles, which provide virtualization optimizations. For more information on **tuned**, please refer to the *Tuned and ktune* section of [Red Hat Enterprise Linux 6.0 Performance Tuning Guide](#).

The Red Hat Enterprise Virtualization Hypervisor has security features enabled. Security Enhanced Linux (SELinux) and the iptables firewall are fully configured and on by default. The Manager can open required ports on Red Hat Enterprise Linux hosts when it adds them to the environment. For a full list of ports, see *Virtualization Host Firewall Requirements*.

A host is a physical 64-bit server with the Intel VT or AMD-V extensions running Red Hat Enterprise Linux 6.1 or later AMD64/Intel 64 version.



#### Important

Support is still ongoing for Red Hat Enterprise Linux 5.4 and 5.5 that already belong to 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.
- » Must have CPUs that support the AMD-V or Intel VT hardware virtualization extensions.
- » Must have CPUs that support all functionality exposed by the virtual CPU type selected upon cluster creation.
- » Has a minimum of 2 GB RAM.
- » Can have an assigned system administrator with system permissions.

Administrators can receive the latest security advisories from the Red Hat 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/>

See Also:

- » [Chapter 3, Data Centers](#)
- » [Chapter 4, Clusters](#)
- » [Chapter 5, Logical Networks](#)
- » [Chapter 8, Virtual Machines](#)

[Report a bug](#)

### 6.2. Red Hat Enterprise Virtualization Hypervisor Hosts

Red Hat Enterprise Virtualization Hypervisor hosts are installed using a special build of Red Hat Enterprise Linux with only the packages required to host virtual machines. They run stateless, not writing any changes to disk unless explicitly required to.

Red Hat Enterprise Virtualization Hypervisor hosts can be added directly to, and configured by, the Red Hat Enterprise Virtualization Manager. Alternatively a host can be configured locally to connect to the Manager; the Manager then is only used to approve the host to be used in the environment.

Unlike Red Hat Enterprise Linux hosts, Red Hat Enterprise Virtualization Hypervisor hosts cannot be added to clusters that have been enabled for Gluster service for use as Red Hat Storage nodes.



#### Important

The Red Hat Enterprise Virtualization Hypervisor is a closed system. Use a Red Hat Enterprise Linux host if additional rpm packages are required for your environment.

[Report a bug](#)

### 6.3. Red Hat Enterprise Linux Hosts

You can use a standard Red Hat Enterprise Linux 6 installation on capable hardware as a host. Red Hat Enterprise Virtualization supports hosts running Red Hat Enterprise Linux 6 Server AMD64/Intel 64 version.

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

[Report a bug](#)

### 6.4. Host Tasks

#### 6.4.1. Adding a Red Hat Enterprise Linux Host

##### Summary

A Red Hat Enterprise Linux host is based on a standard installation of Red Hat Enterprise Linux. The physical host must be set up before you can add it the Red Hat Enterprise Virtualization environment.

The Red Hat Enterprise Virtualization Manager logs into the host to perform virtualization checks, install packages, create a network bridge, and reboot the host. The process of adding a new host can take up to 10 minutes.

**Procedure 6.1. Adding a Red Hat Enterprise Linux Host**

1. Click the **Hosts** resource tab to list the hosts in the results list.
2. Click **New** to open the **New Host** window.
3. Use the drop-down menus to select the **Data Center** and **Host Cluster** for the new host.
4. Enter the **Name**, **Address**, and **Root Password** of the new host.
5. If applicable, clear the **Automatically configure host firewall** check box.
6. You can configure the **Power Management** and **SPM** using the applicable tabs now; however, as these are not fundamental to adding a Red Hat Enterprise Linux host, they are not covered in this procedure.
7. Click **OK** to add the host and close the window.

**Result**

The new host displays in the list of hosts with a status of **Installing**. Once installation is complete, the status will update to **Reboot**. The host must be activated for the status to change to **Up**.

**Note**

You can view the progress of the installation in the details pane.

**See Also:**

= [Section 6.4.2, "Explanation of Settings and Controls in the New Host and Edit Host Windows"](#)

[Report a bug](#)

**6.4.2. Explanation of Settings and Controls in the New Host and Edit Host Windows****6.4.2.1. Host General Settings Explained**

These settings apply when editing host details or adding new Red Hat Enterprise Linux hosts.

Figure 6.1. New Host Settings

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

Table 6.1. General settings

Field Name	Description
Name	The name of the cluster. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Address	The IP address, or resolvable hostname of the host.
Data Center	The data center to which the host belongs. Red Hat Enterprise Virtualization Hypervisor hosts can not be added to Gluster-enabled clusters.
Host Cluster	The cluster to which the host belongs.
Root password	The password of the host's root user. This can only be given when you add the host, it cannot be edited afterwards.

[Report a bug](#)

The screenshot shows the 'New Host' dialog box with the 'Power Management' tab selected. The dialog has a sidebar with tabs: General, Power Management (selected), SPM, and Console. The main area contains the following settings:

- ☐ Enable Power Management
- Primary (dropdown menu)
- ☐ Concurrent
- Address (text field)
- User Name (text field)
- Password (text field)
- Type (dropdown menu, currently set to 'apc')
- Port (text field)
- Slot (text field)
- Options (text field)
- Please use a comma-separated list of 'key=value' or 'key'*
- ☐ Secure
- Source (text field containing 'cluster' and 'dc' on separate lines, with an 'Up' button next to it)

At the bottom right are 'OK' and 'Cancel' buttons.

Figure 6.2. Power Management Settings

The **Power Management** settings table contains the information required on the **Power Management** tab of the **New Host** or **Edit Host** windows.

Table 6.2. Power Management Settings

Field Name	Description
<b>Primary/ Secondary</b>	Prior to Red Hat Enterprise Virtualization 3.2, a host with power management configured only recognised one fencing agent. Fencing agents configured on version 3.1 and earlier, and single agents, are treated as primary agents. The secondary option is valid when a second agent is defined.
<b>Concurrent</b>	Valid when there are two fencing agents, for example for dual-power hosts in which each power switch has two agents connected to the same power switch. <ul style="list-style-type: none"> <li>» If this checkbox is selected, both fencing agents are used concurrently when a host is fenced. This means that both fencing agents have to respond to the Stop command for the host to be stopped; if one agent responds to the Start command, the host will go up.</li> <li>» If this checkbox is not selected, the fencing agents are used sequentially. This means that to stop or start a host, the primary agent is used first, and if it fails, the secondary agent is used.</li> </ul>
<b>Address</b>	The address to access your host's power management device. Either a resolvable hostname or an IP address.
<b>User Name</b>	User account to access the power management device with. You may have to set up a user on the device, or use the default user.
<b>Password</b>	Password for the user accessing the power management device.
<b>Type</b>	The type of power management device in your host. Choose one of the following: <ul style="list-style-type: none"> <li>» <b>apc</b> - APC MasterSwitch network power switch. Not for use with APC 5.x power switch devices.</li> <li>» <b>apc_snmp</b> - Use with APC 5.x power switch devices.</li> <li>» <b>bladecenter</b> - IBM Bladecentre Remote Supervisor Adapter</li> <li>» <b>cisco_ucs</b> - Cisco Unified Computing System</li> <li>» <b>drac5</b> - Dell Remote Access Controller for Dell computers</li> <li>» <b>eps</b> - ePowerSwitch 8M+ network power switch</li> <li>» <b>ilo, ilo2, ilo3, ilo4</b> - HP Integrated Lights-Out</li> <li>» <b>ipmilan</b> - Intelligent Platform Management Interface and Sun Integrated Lights Out Management devices.</li> <li>» <b>rsa</b> - IBM Remote Supervisor Adaptor</li> <li>» <b>rsb</b> - Fujitsu-Siemens RSB management interface</li> <li>» <b>wti</b> - WTI Network PowerSwitch</li> </ul>
<b>Port</b>	The port number used by the power management device to communicate with the host.
<b>Options</b>	Power management device specific options. Give these as 'key=value' or 'key', refer to the documentation of your host's power management device for the options available.
<b>Secure</b>	Tick this checkbox to allow the power management device to connect securely to the host. This can be done via ssh, ssl, or other authentication protocols depending on and supported by the power management agent.
<b>Source</b>	Specifies whether the host will search within its <b>cluster</b> or <b>data center</b> for a fencing proxy. Use the <b>Up</b> and <b>Down</b> buttons to change the sequence in which the resources are used.

[Report a bug](#)**6.4.2.3. SPM Priority Settings Explained**

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

Table 6.3. SPM settings

Field Name	Description
<b>SPM Priority</b>	Likelihood that the host will be given the role of Storage Pool Manager (SPM). Choose between Low, Normal, and High priority.

[Report a bug](#)**6.4.2.4. Host Console Settings Explained**

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

Table 6.4. Console settings

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


[Report a bug](#)

### 6.4.3. Configuring Host Power Management Settings

#### Summary

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

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


**Important**

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

#### Procedure 6.2. Configuring Power Management Settings

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Edit** to open the **Edit Host** window.
3. Click the **Power Management** tab to display the Power Management settings.
4. Select the **Enable Power Management** check box to enable the fields.
5. The **Primary** option is selected by default if you are configuring a new power management device. If you are adding a new device, set it to **Secondary**.
6. Select the **Concurrent** check box to enable multiple fence agents to be used concurrently.
7. Enter the **Address**, **User Name**, and **Password** of the power management device into the appropriate fields.
8. Use the drop-down menu to select the **Type** of power management device.
9. Enter the **Port** number used by the power management device to communicate with the host.
10. Enter the **Options** for the power management device. Use a comma-separated list of **'key=value'** or **'key'**.
11. Select the **Secure** checkbox to enable the power management device to connect securely to the host.
12. Click **Test** to ensure the settings are correct.
13. Click **OK** to save your settings and close the window.

#### Result

You have configured the power management settings for the host. The **Power Management** drop-down menu is now enabled in the Administration Portal.

#### See Also:

- » [Section 6.4.2.2, "Host Power Management Settings Explained"](#)
- » [Section 6.6.3, "Setting Fencing Parameters on a Host"](#)

[Report a bug](#)

### 6.4.4. Configuring Host Storage Pool Manager (SPM) Settings

#### Summary

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

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

#### Procedure 6.3. Configuring SPM settings

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

#### Result

You have configured the SPM priority of the host.

#### See Also:

- » [Section 6.4.2.3, "SPM Priority Settings Explained"](#)

[Report a bug](#)

### 6.4.5. Editing a Resource

### Summary

Edit the properties of a resource.

#### Procedure 6.4. Editing a Resource

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

### Result

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

#### See Also:

- » [Section 6.4.2, "Explanation of Settings and Controls in the New Host and Edit Host Windows"](#)

[Report a bug](#)

### 6.4.6. Approving Newly Added Red Hat Enterprise Virtualization Hypervisor Hosts

#### Prerequisites:

- » [Section 6.4.1, "Adding a Red Hat Enterprise Linux Host"](#)

### Summary

You have to install your Red Hat Enterprise Virtualization Hypervisor hosts before you can approve them in the Red Hat Enterprise Virtualization Manager. Read about installing Red Hat Enterprise Virtualization Hypervisors in the *Red Hat Enterprise Virtualization Installation Guide*.

Once installed, the Red Hat Enterprise Virtualization Hypervisor host is visible in the Administration Portal but not active. Approve it so that it can host virtual machines.

#### Procedure 6.5. Approving newly added Red Hat Enterprise Virtualization Hypervisor hosts

1. In the **Hosts** tab, select the host you recently installed using the Red Hat Enterprise Virtualization Hypervisor host installation media. This host shows a status of **Pending Approval**.
2. Click the **Approve** button.

### Result

The host's status changes to **Up** and it can be used to run virtual machines.



#### Note

You can also add this host using the *Adding a Red Hat Enterprise Linux Host* procedure, which utilizes the Red Hat Enterprise Virtualization Hypervisor host's IP address and the password that was set on the **RHEV-M** screen.

[Report a bug](#)

### 6.4.7. Moving a Host to Maintenance Mode

#### Summary

Many common maintenance tasks, including network configuration and deployment of software updates, require that 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 migrated virtual machines.

#### Procedure 6.6. Moving a Host to Maintenance Mode

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Maintenance** to open the **Maintenance Host(s)** confirmation window.
3. Click **OK** to initiate maintenance mode.

#### Result:

All running virtual machines are migrated to alternative hosts. The **Status** field of the host changes to **Preparing for Maintenance**, and finally **Maintenance** when the operation completes successfully.

[Report a bug](#)

### 6.4.8. Activating a host from maintenance mode

#### Summary

A host that has been placed into maintenance mode, or recently added to the environment, must be activated before it can be used.

#### Procedure 6.7. Activating a Host from Maintenance Mode

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

### Result

The host status changes to **Unassigned**, and finally **Up** when the operation is complete. Virtual machines can now run on the host.

[Report a bug](#)

### 6.4.9. Removing a Host

#### Summary

Remove a host from your virtualized environment.

#### Procedure 6.8. Removing a host

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the



results list.

2. Place the host into maintenance mode.
3. Click **Remove** to open the **Remove Host(s)** confirmation window.
4. Select the **Force Remove** check box if the host is part of a Red Hat Storage cluster and has volume bricks on it, or if the host is non-responsive.
5. Click **OK**.

#### Result

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

[Report a bug](#)

### 6.4.10. Customizing Hosts with Tags

#### Prerequisites:

- ▮ [Section E.3.2, "Creating a Tag"](#)

#### Summary

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

#### Procedure 6.9. Customizing hosts with tags

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




Figure 6.3. Assign Tags Window

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

#### Result

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

[Report a bug](#)

## 6.5. Hosts and Networking

### 6.5.1. Editing Host Network Interfaces and Adding Logical Networks to Hosts

#### Summary

You can change the settings of host network interfaces. Moving the **rhevm** management logical network between interfaces, and adding a newly created logical network to a network interface are common reasons to edit host networking.

#### Procedure 6.10. Editing Host Network Interfaces and Adding Logical Networks to Hosts

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results.
2. Click the **Network Interfaces** tab in the details pane to list the network interfaces attached to the host and their configurations.
3. Click the **Setup Host Networks** button to open the **Setup Host Networks** window.

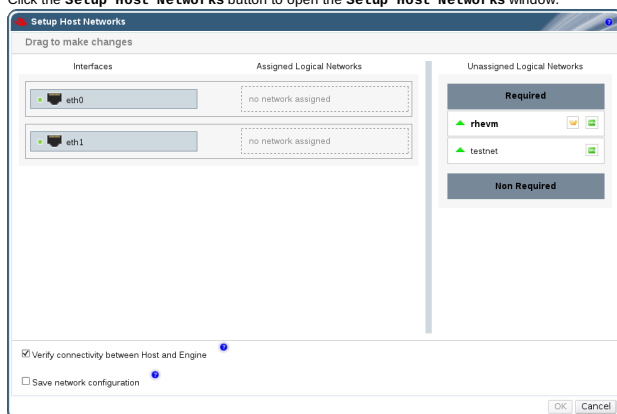


Figure 6.4. The Setup Host Networks window

4. Attach a logical network to a network interface by selecting and dragging a logical network into the **Assigned Logical Networks** area next to the network interface.

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

5. Edit the logical networks by hovering your cursor over an assigned logical network and clicking the pencil icon to open the **Edit Network** window.

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

Select a **Boot Protocol** from:

- » **None**,
- » **DHCP**, or
- » **Static**.

If you have chosen **Static**, provide the **IP** and **Subnet Mask**.

Click **OK**.

6. Select the **Verify connectivity between Host and Engine** check box to run a network check; this will only work if the host is in maintenance mode.
7. Select the **Save network configuration** check box if you want these network changes to be made persistent when the environment is rebooted.
8. Click **OK** to implement the changes and close the window.

#### Result

You have assigned logical networks to network interfaces and configured the host network.

[Report a bug](#)

### 6.5.2. Creating a Bond Device using the Administration Portal

#### Summary

You can bond two network interfaces of the same make and model together using one of the 4 supported bonding modes. This type of configuration can increase available bandwidth and reliability.

#### Note

If you would like to use the bond device to carry the **rhev** management network with VLAN tagging enabled, please see [Red Hat KB43070](#).

#### Procedure 6.11. Creating a Bond Device using the Administration Portal

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the host.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Select and drag one of the interfaces over the top of another interface and drop it to open the **Create New Bond** window.

Alternatively, right-click the interface and select another interface from the drop-down menu.

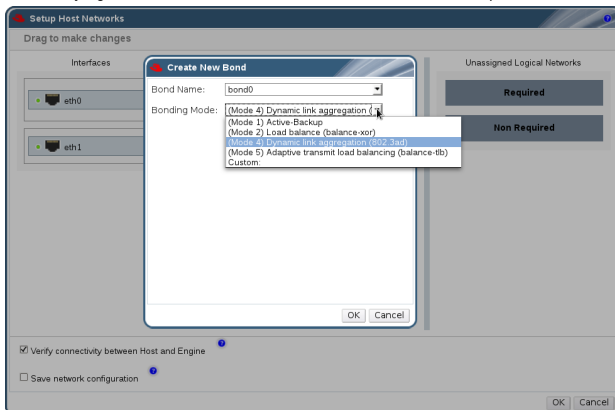


Figure 6.5. Bond Devices Window

5. Select the **Bond Name** and **Bonding Mode** from the drop-down menus.  
Bonding modes 1, 2, 4, and 5 can be selected. Any other mode can be configured using the **Custom** option.
6. Click **OK** to create the bond and close the **Create New Bond** window.
7. Assign a logical network to the newly created bonded interface.
8. Optionally choose to **Verify connectivity between Host and Engine** and **Save network configuration**.
9. Click **OK** accept the changes and close the **Setup Host Networks** window.

#### Result:

Your two interfaces are a linked bond device and can be edited as a single interface. The bond device lists in the **Network Interfaces** tab of the details pane for the selected host.

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

#### See Also:

- » [Section 6.5.1, "Editing Host Network Interfaces and Adding Logical Networks to Hosts"](#)

[Report a bug](#)

### 6.5.3. Example Uses of Custom Bonding Options with Host Interfaces

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

**Example 6.1. xmit\_hash\_policy**

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

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

**Example 6.2. ARP Monitoring**

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

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

**Example 6.3. Primary**

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

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

[Report a bug](#)

**6.5.4. Saving a host network configuration****Summary**

One of the options when configuring a host network is to save the configuration as you apply it, making the changes persistent.

Any changes made to the host network configuration will be temporary if you did not select the **Save network configuration** check box in the **Setup Host Networks** window.

Save the host network configuration to make it persistent.

**Procedure 6.12. Saving a host network configuration**

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Network Interfaces** tab on the Details pane to list the NICs on the host, their address, and other specifications.
3. Click the **Save Network Configuration** button.
4. The host network configuration is saved and the following message is displayed on the task bar: "Network Changes were saved on host <Hostname>."

**Result**

The host's network configuration is saved persistently and will survive reboots.

[Report a bug](#)

**6.6. Host Resilience****6.6.1. Host High Availability**

The Red Hat Enterprise Virtualization Manager uses fencing to keep the hosts in a cluster responsive. A **Non Responsive** host is different from a **Non Operational** host. **Non Operational** hosts can be communicated with by the Manager, but have an incorrect configuration, for example a missing logical network. **Non Responsive** hosts cannot be communicated with by the Manager.

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

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

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

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

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

[Report a bug](#)

**6.6.2. Power Management by Proxy in Red Hat Enterprise Virtualization**

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

With Red Hat Enterprise Virtualization 3.2, the ability to prioritize hosts for use as fencing proxies was introduced. You can select between:

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

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

[Report a bug](#)

**6.6.3. Setting Fencing Parameters on a Host**

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

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

#### Procedure 6.13. Setting fencing parameters on a host

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

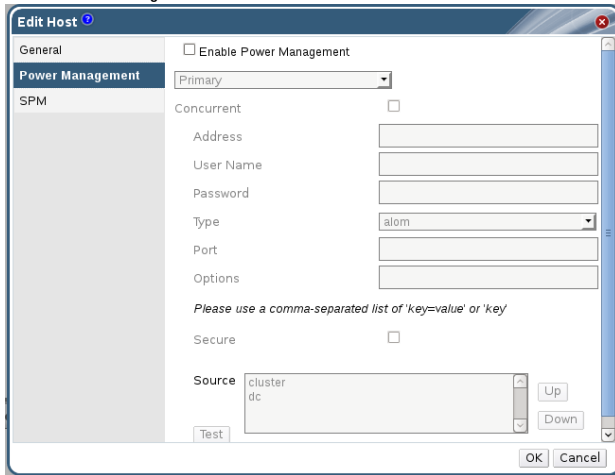


Figure 6.6. Power Management Settings

4. Select the **Enable Power Management** check box to enable the fields.
5. The **Primary** option is selected by default if you are configuring a new power management device. If you are adding a new device, set it to **Secondary**.
6. Select the **Concurrent** check box to enable multiple fence agents to be used concurrently.
7. Enter the **Address**, **User Name**, and **Password** of the power management device.
8. Select the power management device **Type** from the drop-down menu.
9. Enter the **Port** number used by the power management device to communicate with the host.
10. Enter the specific **Options** of the power management device. Use a comma-separated list of 'key=value' or 'key' entries.
11. Click the **Test** button to test the power management device. *Test Succeeded, Host Status is: on* will display upon successful verification.

**Warning**

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

12. Click **OK** to save the changes and close the window.

#### Result

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.

#### See Also:

- » [Section 6.4.2.2, "Host Power Management Settings Explained"](#)

[Report a bug](#)

### 6.6.4. Using Host Power Management Functions

#### Prerequisites:

- » [Section 6.6.3, "Setting Fencing Parameters on a Host"](#)

#### Summary

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

#### Procedure 6.14. Using Host Power Management Functions

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Power Management** drop-down menu.




Figure 6.7. Restart

3. Select one of the following options:
  - » **Restart**: This option stops the host and waits until the host's status changes to **Down**. When the agent has verified that the host is down, the highly available virtual machines are restarted on another host in the cluster. The agent then restarts this host. When the host is ready for use its status displays as **Up**.
  - » **Start**: This option starts the host and lets it join a cluster. When it is ready for use its status displays as **Up**.
  - » **Stop**: This option powers off the host. Before using this option, ensure that the virtual

machines running on the host have been migrated to other hosts in the cluster. Otherwise the virtual machines will crash and only the highly available virtual machines will be restarted on another host. When the host has been stopped its status displays as **Non-Operational**.

### Important

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

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

### Result

The selected action is performed.

[Report a bug](#)

## 6.6.5. Manually Fencing or Isolating a Non Responsive Host

### Summary

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



### Warning

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

### Procedure 6.15. Manually fencing or isolating a non-responsive host

1. On the **Hosts** tab, select the host. The status must display as **non-responsive**.
2. Manually reboot the host. This could mean physically entering the lab and rebooting the host.
3. On the Administration Portal, right-click the host entry and select the **Confirm Host has been rebooted** button.




Figure 6.8. 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**.

### Result

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

[Report a bug](#)

## 6.7. Hosts and Permissions

### 6.7.1. Managing System Permissions for a Host

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **HostAdmin** has administrator privileges only for the assigned host, and so forth.

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

The host administrator role permits the following actions:

- » Edit the configuration of the host;
- » Set up the logical networks; and
- » Remove the host.

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

### See Also:

- » [Section 6.7.2, "Host Administrator Roles Explained"](#)

[Report a bug](#)

### 6.7.2. Host Administrator Roles Explained

#### Host Permission Roles

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

Table 6.5. Red Hat Enterprise Virtualization System Administrator Roles

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

**See Also:**

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

### 6.7.3. Assigning an Administrator or User Role to a Resource

**Summary**

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

**Procedure 6.16. Assigning a Role to a Resource**

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

**Result**

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

**See Also:**

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

### 6.7.4. Removing an Administrator or User Role from a Resource

**Summary**

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

**Procedure 6.17. Removing a Role from a Resource**

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

**Result**

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

[Report a bug](#)

## Chapter 7. Storage

### 7.1. Introduction to Storage in Red Hat Enterprise Virtualization

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).
- » Other POSIX compliant file systems.
- » Internet Small Computer System Interface (iSCSI).
- » Local storage attached directly to the virtualization hosts.
- » Fibre Channel Protocol (FCP).

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

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

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

Red Hat Enterprise Virtualization platform has three types of storage domains:

- » **Data Domain:** A data domain holds the virtual hard disks and OVF files of all the virtual machines and templates in a data center. In addition, snapshots of the virtual machines are also stored in the data domain.  
The data domain cannot be shared across data centers, 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. You must attach a data domain to a data center before you can attach domains of other types to it.
- » **ISO Domain:** ISO domains store ISO files (or logical CDs) used to install and boot operating systems and applications for the virtual machines. An ISO domain removes the data center's need for physical media. An ISO domain can be shared across different data centers. ISO storage domains use NFS storage.
- » **Export Domain:** Export domains are temporary storage repositories that are used to copy and move images between data centers and Red Hat Enterprise Virtualization environments. Export domains can be used to backup virtual machines. An export domain can be moved between data centers, however, it can only be active in one data center at a time.



#### Important

Support for export storage domains backed by storage on anything other than NFS is being deprecated. While 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.

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



#### Important

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

[Report a bug](#)

### 7.2. Understanding Storage Domains

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

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

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

Virtual disks can have one of two formats, either Ccow2 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.

[Report a bug](#)

### 7.3. Storage Metadata Versions in Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization stores information about storage domains as metadata on the storage domains themselves. Each major release of Red Hat Enterprise Virtualization has seen improved implementations of storage metadata.

- » *V1 metadata (Red Hat Enterprise Virtualization 2.x series)*

Each storage domain contains metadata describing its own structure, and all of the names of physical volumes that are used to back virtual machine disk images.

Master domains additionally contain metadata for all the domains and physical volume names in the storage pool. The total size of this metadata is limited to 2 kb, limiting the number of storage domains that can be in a pool.

Template and virtual machine base images are read only.

V1 metadata is applicable to NFS, iSCSI, and FC storage domains.

- » *V2 metadata (Red Hat Enterprise Virtualization 3.0)*

All storage domain and pool metadata is stored as logical volume tags rather than written to a logical volume. Metadata about virtual machine disk volumes is still stored in a logical volume on the domains.

Physical volume names are no longer included in the metadata.

Template and virtual machine base images are read only.

V2 metadata is applicable to iSCSI, and FC storage domains.

■ **V3 metadata (Red Hat Enterprise Virtualization 3.1)**

All storage domain and pool metadata is stored as logical volume tags rather than written to a logical volume. Metadata about virtual machine disk volumes is still stored in a logical volume on the domains.

Virtual machine and template base images are no longer read only. This change enables live snapshots, live storage migration, and clone from snapshot.

Support for unicode metadata is added, for non-English volume names.

V3 metadata is applicable to NFS, POSIX, iSCSI, and FC storage domains.

#### Note

Upgrades between metadata versions are automatic. If you've upgraded from Red Hat Enterprise Virtualization 3.0 to 3.1, your existing data centers are initially in 3.0 compatibility mode. When you upgrade your hosts and change your data center compatibility from 3.0 to 3.1, the storage metadata for your storage domains is automatically upgraded to version 3.

[Report a bug](#)

## 7.4. Preparing and Adding File-based Storage

### 7.4.1. Preparing NFS Storage

#### Summary

These steps must be taken to prepare an NFS file share on a server running Red Hat Enterprise Linux 6 for use with Red Hat Enterprise Virtualization.

#### Procedure 7.1. Preparing NFS Storage

##### 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.

[Report a bug](#)

### 7.4.2. Attaching NFS Storage



**Summary**

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.

NFS data domains can be added to NFS data centers. You can add NFS, ISO, and export storage domains to data centers of any type.

**Procedure 7.2. Attaching NFS Storage**

1. Click the **Storage** resource tab to list the existing storage domains.
2. Click **New Domain** to open the **New Domain** window.

Figure 7.1. NFS Storage

3. Enter the **Name** of the storage domain.
4. Select the **Data Center**, **Domain Function / Storage Type**, and **Use Host** from the drop-down menus.  
If applicable, select the **Format** from the drop-down menu.
5. Enter the **Export Path** to be used for the storage domain.  
The export path should be in the format of **192.168.0.10:/data** or **domain.example.com:/data**
6. Click **Advanced Parameters** to enable further configurable settings. It is recommended that the values of these parameters not be modified.

### Important

All communication to the storage domain is from the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must be attached to the chosen Data Center before the storage is configured.

7. Click **OK** to create the storage domain and close the window.

**Result**

The new NFS data domain is displayed on the **Storage** tab with a status of **Locked** while the disk prepares. It is automatically attached to the data center upon completion.

[Report a bug](#)

**7.4.3. Preparing Local Storage****Summary**

A local storage domain can be set up on a host. When you set up host to use local storage, the host automatically gets added to a new data center and cluster that no other hosts can be added to. Multiple host clusters require that all hosts have access to all storage domains, which is not possible with local storage. Virtual machines created in a single host cluster cannot be migrated, fenced or scheduled.

### Important

On Red Hat Enterprise Virtualization Hypervisors the only path permitted for use as local storage is **/data/images**. This directory already exists with the correct permissions on Hypervisor installations. The steps in this procedure are only required when preparing local storage on Red Hat Enterprise Linux virtualization hosts.

**Procedure 7.3. Preparing Local Storage**

1. On the virtualization host, create the directory to be used for the local storage.

```
# mkdir -p /data/images
```

2. Ensure that the directory has permissions allowing read/write access to the **vdsm** user (UID 36) and **kvm** group (GID 36).

```
# chown 36:36 /data /data/images
```

```
# chmod 0755 /data /data/images
```

**Result**

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

[Report a bug](#)

### 7.4.4. Adding Local Storage

#### Summary

You have prepared the storage that is local to your host. Add it to the host to begin using it.

Adding local storage to a host in this manner causes the host to be put in a new data center and cluster. The local storage configuration window combines the creation of a data center, a cluster, and storage into a single process.

#### Procedure 7.4. Adding Local Storage

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Maintenance** to place the host into maintenance mode.
3. Click **Configure Local Storage** to open the **Configure Local Storage** window.

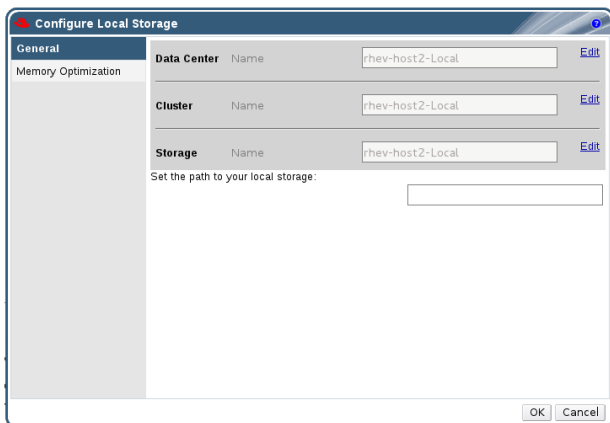


Figure 7.2. Configure Local Storage Window

4. Click the **Edit** buttons next to the **Data Center**, **Cluster**, and **Storage** fields to configure and name the local storage domain.
5. Set the path to your local storage in the text entry field.
6. If applicable, select the **Memory Optimization** tab to configure the memory optimization policy for the new local storage cluster.
7. Click **OK** to save the settings and close the window.

#### Result

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

[Report a bug](#)

## 7.5. Adding POSIX Compliant File System Storage

### 7.5.1. POSIX Compliant File System Storage in Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization 3.1 and higher supports the use of POSIX (native) file systems for storage. POSIX file system support allows you to mount file systems using the same mount options that you would normally use when mounting them manually from the command line. This functionality is intended to allow access to storage not exposed using NFS, iSCSI, or FCP.

Any POSIX compliant filesystem used as a storage domain in Red Hat Enterprise Virtualization **MUST** support sparse files and direct I/O. The Common Internet File System (CIFS), for example, does not support direct I/O, making it incompatible with Red Hat Enterprise Virtualization.

**Important**

Do not mount NFS storage by creating a POSIX compliant file system Storage Domain. Always create an NFS Storage Domain instead.

[Report a bug](#)

### 7.5.2. Attaching POSIX Compliant File System Storage

#### Summary

You want to use a POSIX compliant file system that is not exposed using NFS, iSCSI, or FCP as a storage domain.

#### Procedure 7.5. Attaching POSIX Compliant File System Storage

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click **New Domain** to open the **New Domain** window.

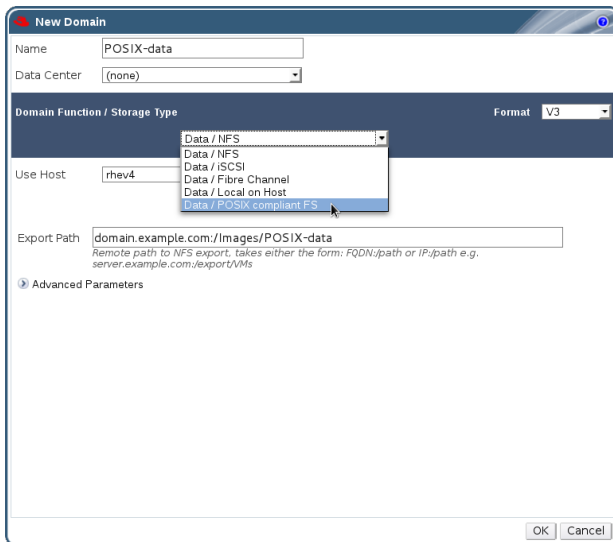


Figure 7.3. POSIX Storage

3. Enter the **Name** for the storage domain.
4. Select the **Data Center** to be associated with the storage domain. The Data Center selected must be of type **POSIX (POSIX compliant FS)**. Alternatively, select **(none)**.
5. Select **Data / POSIX compliant FS** from the **Domain Function / Storage Type** drop-down menu.  
If applicable, select the **Format** from the drop-down menu.
6. Select a host from the **Use Host** drop-down menu. Only hosts within the selected data center will be listed. The host that you select will be used to connect the storage domain.
7. Enter the **Path** to the POSIX file system, as you would normally provide it to the **mount** command.
8. Enter the **VFS Type**, as you would normally provide it to the **mount** command using the **-t** argument. See **man mount** for a list of valid VFS types.
9. Enter additional **Mount Options**, as you would normally provide them to the **mount** command using the **-o** argument. The mount options should be provided in a comma-separated list. See **man mount** for a list of valid mount options.
10. Click **OK** to attach the new Storage Domain and close the window.

**Result**

You have used a supported mechanism to attach an unsupported file system as a storage domain.

[Report a bug](#)

## 7.6. Adding Block-based Storage

### 7.6.1. Preparing iSCSI Storage

#### Summary

These steps must be taken to export iSCSI storage device from a server running Red Hat Enterprise Linux 6 to use as a storage domain with Red Hat Enterprise Virtualization.

#### Procedure 7.6. Preparing iSCSI Storage

1. Install the **scsi-target-utils** package using the **yum** command as root on your storage server.

```
# yum install -y scsi-target-utils
```

2. Add the devices or files you want to export to the **/etc/tgt/targets.conf** file. Here is a generic example of a basic addition to the **targets.conf** file:

```
<target iqn.YEAR-MONTH.com.EXAMPLE:SERVER.targetX>
    backing-store /PATH/TO/DEVICE1 # Becomes LUN 1
    backing-store /PATH/TO/DEVICE2 # Becomes LUN 2
    backing-store /PATH/TO/DEVICE3 # Becomes LUN 3
</target>
```

Targets are conventionally defined using the year and month they are created, the reversed fully qualified domain that the server is in, the server name, and a target number.

3. Start the **tgt** service.

```
# service tgt start
```

4. Make the **tgt** start persistently across reboots.

```
# chkconfig tgt on
```

5. Open an iptables firewall port to allow clients to access your iSCSI export. By default, iSCSI uses port 3260. This example inserts a firewall rule at position 6 in the **INPUT** table.

```
# iptables -I INPUT 6 -p tcp --dport 3260 -j ACCEPT
```

6. Save the iptables rule you just created.

```
# service iptables save
```

**Result**

You have created a basic iSCSI export. You can use it as an iSCSI data domain.

[Report a bug](#)

### 7.6.2. Adding iSCSI Storage

**Summary**

Red Hat Enterprise Virtualization platform supports iSCSI storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.

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

**Note**

You can only add an iSCSI storage domain to a data center that is set up for iSCSI storage type.

**Procedure 7.7. Adding iSCSI Storage**

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click the **New Domain** button to open the **New Domain** window.
3. Enter the **Name** of the new storage domain.

Figure 7.4. New iSCSI Domain

4. Use the **Data Center** drop-down menu to select an iSCSI data center. If you do not yet have an appropriate iSCSI data center, select **(none)**.
5. Use the drop-down menus to select the **Domain Function / Storage Type** and the **Format**. The storage domain types that are not compatible with the chosen data center are not available.
6. Select an active host in the **Use Host** field. If this is not the first data domain in a data center, you must select the data center's SPM host.

**Important**

All communication to the storage domain is 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** window automatically displays known targets with unused LUNs when iSCSI is selected as the storage type. If the target that you are adding storage from is not listed then you can use target discovery to find it, otherwise proceed to the next step.

**iSCSI Target Discovery**

- a. Click **Discover Targets** to enable target discovery options. When targets have been discovered and logged in to, the **New Domain** window automatically displays targets with LUNs unused by the environment.

**Note**

LUNs used externally to the environment are also displayed.

You can use the **Discover Targets** options to add LUNs on many targets, or multiple paths to the same LUNs.

- b. Enter the fully qualified domain name or IP address of the iSCSI host in the **Address** field.
  - c. Enter the port to connect to the host on when browsing for targets in the **Port** field. The default is **3260**.
  - d. If the Challenge Handshake Authentication Protocol (CHAP) is being used to secure the storage, select the **User Authentication** check box. Enter the **CHAP username** and **CHAP password**.
  - e. Click the **Discover** button.
  - f. Select the target you want to use from the discovery results and click the **Login** button. Alternatively, click the **Login All** to log in to all of the discovered targets.
8. Click the **+** button next to the desired target. This will expand the entry and display all unused LUNs attached to the target.
  9. Select the check box for each LUN that you are using to create the storage domain.
  10. Click **OK** to create the storage domain and close the window.

**Result**

The new iSCSI storage domain displays on the storage tab. This can take up to 5 minutes.

[Report a bug](#)

### 7.6.3. Adding FCP Storage

#### Summary

Red Hat Enterprise Virtualization platform supports SAN storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.

Red Hat 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*.

#### Note

You can only add an FCP storage domain to a data center that is set up for FCP storage type.

#### Procedure 7.8. Adding FCP Storage

1. Click the **Storage** resource tab to list all storage domains in the virtualized environment.
2. Click **New Domain** to open the **New Domain** window.
3. Enter the **Name** of the storage domain

Figure 7.5. Adding FCP Storage

4. Use the **Data Center** drop-down menu to select an FCP data center.  
If you do not yet have an appropriate FCP data center, select **(none)**.
5. Use the drop-down menus to select the **Domain Function / Storage Type** and the **Format**. The storage domain types that are not compatible with the chosen data center are not available.
6. Select an active host in the **Use Host** field. If this is not the first data domain in a data center, you must select the data center's SPM host.



#### Important

All communication to the storage domain is 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 **New Domain** window automatically displays known targets with unused LUNs when **Data / Fibre Channel** is selected as the storage type. Select the **LUN ID** check box to select all of the available LUNs.  
Alternatively, select the **Add LUN** check boxes for individual LUNs to use them with the FCP data domain.
8. Click **OK** to create the storage domain and close the window.

#### 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**.

[Report a bug](#)

### 7.6.4. Un-useable LUNs in Red Hat Enterprise Virtualization

In certain circumstances, the Red Hat Enterprise Virtualization Manager will not allow you to use a LUN to create a storage domain or virtual machine hard disk.

- LUNs that are already part of the current Red Hat Enterprise Virtualization environment are automatically prevented from being used.

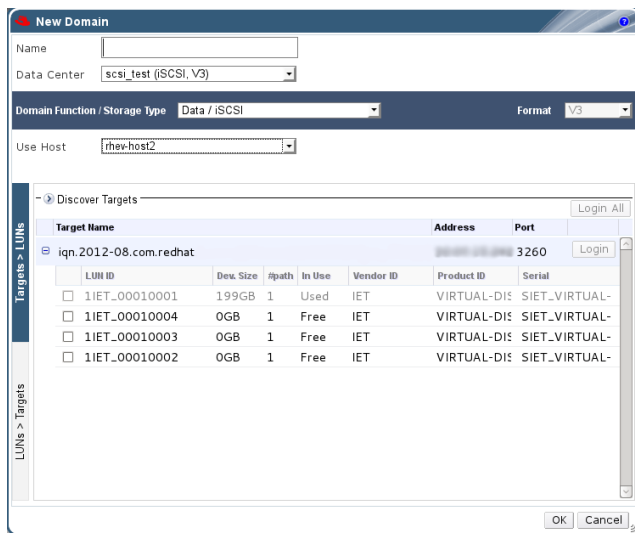


Figure 7.6. Un-useable LUNs in the Red Hat Enterprise Virtualization Administration Portal

- LUNs that are already being used by the SPM host will also display as in use. You can choose to forcefully override the contents of these LUNs, but the operation is not guaranteed to succeed.

[Report a bug](#)

## 7.7. Storage Tasks

### 7.7.1. Importing existing ISO or export storage domains

#### Summary

You have an ISO or export domain that you have been using with a different data center. You want to attach it to the data center you are using, and import virtual machines or use ISOs.

#### Procedure 7.9. Importing an Existing ISO or Export Storage Domain

1. Click the **Storage** resource tab to list all the available storage domains in the results list.
2. Click **Import Domain** to open the **Import Pre-Configured Domain** window.

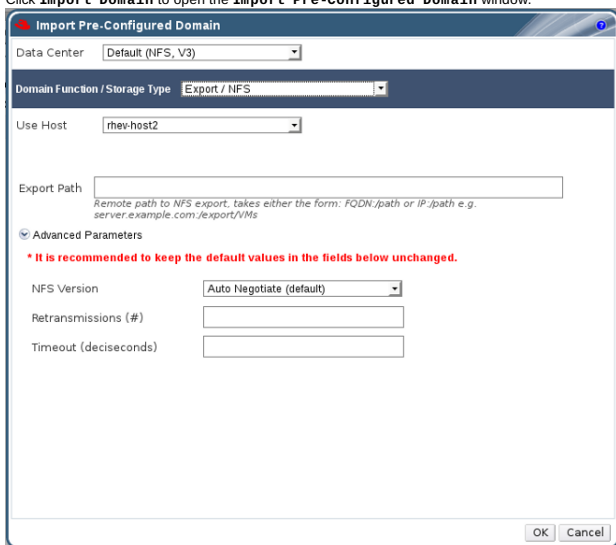


Figure 7.7. Import Domain

3. Select the appropriate **Domain Function / Storage Type** from the following:
  - » ISO
  - » Export

The **Domain Function / Storage Type** determines the availability of the **Format** field.

4. Select the SPM host from the **Use host** drop-down menu.



#### 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.

5. 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**.
6. Click **OK** to import the domain and close the window.
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.

**Result**

You have imported your export or ISO domain to your data center. Attach it to a data center to use it.

[Report a bug](#)

**7.7.2. Populating the ISO Storage Domain****Summary**

An ISO storage domain is attached to a data center, ISO images must be uploaded to it. Red Hat Enterprise Virtualization provides an ISO uploader tool that ensures that the images are uploaded into the correct directory path, with the correct user permissions.

The creation of ISO images from physical media is not described in this document. It is assumed that you have access to the images required for your environment.

**Procedure 7.10. Populating the ISO Storage Domain**

1. Copy the required ISO image to a temporary directory on the system running Red Hat Enterprise Virtualization Manager.
2. Log in to the system running Red Hat Enterprise Virtualization Manager as the **root** user.
3. Use the **rhevms-iso-uploader** command to upload the ISO image. This action will take some time, the amount of time varies depending on the size of the image being uploaded and available network bandwidth.

**Example 7.1. ISO Uploader Usage**

In this example the ISO image **RHEL6.iso** is uploaded to the ISO domain called **ISODomain** using NFS. The command will prompt for an administrative username and password. The username must be provided in the form **username@domain**.

```
# rhevms-iso-uploader --iso-domain=ISODomain upload RHEL6.iso
```

**Result**

The ISO image is uploaded and appears in the ISO storage domain specified. It is also available in the list of available boot media when creating virtual machines in the data center which the storage domain is attached to.

**See Also:**

» [Section 15.6, "Uploading ISO Files with the ISO Uploader Tool"](#)

[Report a bug](#)

**7.7.3. Moving storage domains to maintenance mode****Summary**

Detaching and removing storage domains requires that they be in maintenance mode. This is required to redesignate another data domain as the master data domain.

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

**Important**

Put any active ISO and export domains in maintenance mode using this procedure.

**Procedure 7.11. Moving storage domains to maintenance mode**

1. Use the **Storage** resource tab, tree mode, or the search function to find and select the storage domain in the results list.
2. Shut down and move all the virtual machines running on the storage domain.
3. Click the **Data Centers** tab in the details pane.
4. Click **Maintenance**. The storage domain is deactivated and has an **Inactive** status in the results list.

**Result**

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.

[Report a bug](#)

**7.7.4. Editing a Resource****Summary**

Edit the properties of a resource.

**Procedure 7.12. Editing a Resource**

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

**Result**

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

[Report a bug](#)

**7.7.5. Activating storage domains****Summary**

If you have been making changes to a data center's storage, you have to put storage domains into maintenance mode. Activate a storage domain to resume using it.

1. Use the **Storage** resource tab, tree mode, or the search function to find and select the inactive

- storage domain in the results list.
- Click the **Data Centers** tab in the details pane.
  - Select the appropriate data center and click **Activate**.

**Important**

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

**Result**

Your storage domain is active and ready for use.

[Report a bug](#)

**7.7.6. Removing a storage domain****Summary**

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

**Procedure 7.13. Removing a Storage Domain**

- Use the **Storage** resource tab, tree mode, or the search function to find and select the appropriate storage domain in the results list.
- Move the domain into maintenance mode to deactivate it.
- Detach the domain from the data center.
- Click **Remove** to open the **Remove Storage** confirmation window.
- Select a host from the list.
- Click **OK** to remove the storage domain and close the window.

**Summary**

The storage domain is permanently removed from the environment.

[Report a bug](#)

**7.7.7. Destroying a storage domain****Summary**

A storage domain encountering errors may not be able to be removed through the normal procedure. Destroying a storage domain will forcibly remove the storage domain from the virtualized environment without reference to the export directory.

When the storage domain is destroyed, you are required to manually fix the export directory of the storage domain before it can be used again.

**Procedure 7.14. Destroying a Storage Domain**

- Use the **Storage** resource tab, tree mode, or the search function to find and select the appropriate storage domain in the results list.
- Right-click the storage domain and select **Destroy** to open the **Destroy Storage Domain** confirmation window.
- Select the **Approve operation** check box and click **OK** to destroy the storage domain and close the window.

**Result**

The storage domain has been destroyed. Manually clean the export directory for the storage domain to recycle it.

[Report a bug](#)

**7.7.8. Detaching the Export Domain****Summary**

Detach the export domain from the data center to import the templates to another data center.

**Procedure 7.15. Detaching an Export Domain from the Data Center**

- Use the **Storage** resource tab, tree mode, or the search function to find and select the export domain in the results list.
- Click the **Data Centers** tab in the details pane and select the export domain.
- Click **Maintenance** to put the export domain into maintenance mode.
- Click **Detach** to open the **Detach Storage** confirmation window.
- Click **OK** to detach the export domain.

**Result**

The export domain has been detached from the data center, ready to be attached to another data center.

[Report a bug](#)

**7.7.9. Attaching an Export Domain to a Data Center****Prerequisites:**

- » [Section 9.2.7, "Importing Templates"](#)

**Summary**

Attach the export domain to a data center.

**Procedure 7.16. Attaching an Export Domain to a Data Center**

- Use the **Storage** resource tab, tree mode, or the search function to find and select the export domain in the results list.
- Click the **Data Centers** tab in the details pane.
- Click **Attach** to open the **Attach to Data Center** window.
- Select the radio button of the appropriate data center.
- Click **OK** to attach the export domain.
- Select the newly attached data center in the details pane and click **Activate**.

**Result**

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



[Report a bug](#)

## 7.8. Red Hat Storage Volumes

### 7.8.1. Introduction to Red Hat Storage Volumes

You can use the console to create and start new volumes. Volumes combine storage from more than one Red Hat Storage Server into a single global namespace. A volume is a logical collection of bricks where each brick is an export directory on a Red Hat Storage Server in the trusted storage pool. Most of the management operations of Red Hat Storage Console happen on the volume.

You can monitor volumes in your cluster from the **Volumes** tab.



#### Note

Bricks must be created externally on Red Hat Storage nodes.

[Report a bug](#)

### 7.8.2. Gluster Storage Terminology

Table 7.1. Data Center Properties

Term	Definition
<b>Brick</b>	<p>A brick is the GlusterFS basic unit of storage, represented by an export directory on a server in the trusted storage pool. A Brick is expressed by combining a server with an export directory in the following format:</p> <p><b>SERVER:EXPORT</b></p> <p>For example:</p> <p><b>myhostname:/exports/myexportdir/</b></p>
<b>Block Storage</b>	Block special files or block devices correspond to devices through which the system moves data in the form of blocks. These device nodes often represent addressable devices such as hard disks, CD-ROM drives, or memory-regions. Red Hat Storage supports XFS file system with extended attributes.
<b>Cluster</b>	A trusted pool of linked computers, working together closely thus in many respects forming a single computer. In Red Hat Storage terminology a cluster is called as trusted storage pool.
<b>Client</b>	The machine which mounts the volume (this may also be a server)
<b>Distributed File System</b>	A file system that allows multiple clients to concurrently access data spread across multiple servers/bricks in a trusted storage pool. Data sharing among multiple locations is fundamental to all distributed file system.
<b>Geo-Replication</b>	Geo-replication provides a continuous, asynchronous, and incremental replication service from site to another over Local Area Networks (LAN), Wide Area Network (WAN), and across the Internet.
<b>glusterd</b>	The Gluster management daemon that needs to run on all servers in the trusted storage pool.
<b>Metadata</b>	Metadata is data providing information about one or more other pieces of data.
<b>N-way Replication</b>	Local synchronous data replication typically deployed across campus or Amazon Web Services Availability Zones.
<b>Namespace</b>	Namespace is an abstract container or environment created to hold a logical grouping of unique identifiers or symbols. Each Red Hat Storage trusted storage pool exposes a single namespace as a POSIX mount point that contains every file in the trusted storage pool.
<b>POSIX</b>	Portable Operating System Interface (for Unix) is the name of a family of related standards specified by the IEEE to define the application programming interface (API), along with shell and utilities interfaces for software compatible with variants of the UNIX operating system. Red Hat Storage exports a fully POSIX compatible file system.
<b>RAID</b>	Redundant Array of Inexpensive Disks (RAID) is a technology that provides increased storage reliability through redundancy, combining multiple low-cost, less-reliable disk drives components into a logical unit where all drives in the array are interdependent.
<b>RRDNS</b>	Round Robin Domain Name Service (RRDNS) is a method to distribute load across application servers. RRDNS is implemented by creating multiple A records with the same name and different IP addresses in the zone file of a DNS server.
<b>Server</b>	The machine (virtual or bare metal) which hosts the actual file system in which data will be stored.
<b>Scale-Up Storage</b>	Increases the capacity of the storage device, but only in a single dimension. An example might be adding additional disk capacity to a single computer in a trusted storage pool.
<b>Scale-Out Storage</b>	Increases the capability of a storage device in multiple dimensions. For example adding a server to a trusted storage pool increases CPU, disk capacity, and throughput for the trusted storage pool.
<b>Subvolume</b>	A brick after being processed by at least one translator.
<b>Translator</b>	A translator connects to one or more subvolumes, does something with them, and offers a subvolume connection.
<b>Trusted Storage Pool</b>	A storage pool is a trusted network of storage servers. When you start the first server, the storage pool consists of that server alone.
<b>User Space</b>	Applications running in user space don't directly interact with hardware, instead using the kernel to moderate access. User Space applications are generally more portable than applications in kernel space. Gluster is a user space application.
<b>Virtual File System (VFS)</b>	VFS is a kernel software layer that handles all system calls related to the standard Linux file system. It provides a common interface to several kinds of file systems
<b>Volfile</b>	Volfile is a configuration file used by glusterfs process. Volfile will be usually located at <b>/var/lib/glusterd/vols/VOLNAME</b> .
<b>Volume</b>	A volume is a logical collection of bricks. Most of the gluster management operations happen on the volume.

[Report a bug](#)

### 7.8.3. Creating A Storage Volume

#### Summary

You can create new volumes in your storage environment. When creating a new volume, you must specify the bricks that comprise the volume and specify whether the volume is to be distributed, replicated, or striped.

**Important**

When using local storage on Red Hat Storage hosts, it is recommended to use replicated volumes, where bricks are exported from different hosts. Replicated volumes create copies of files across multiple bricks in the volume, preventing data loss when a host is fenced.

#### Procedure 7.17. Creating A Storage Volume

1. Click the **Volumes** tab. The **Volumes** tab displays a list of all volumes in the system.
2. Click the **Create Volume** button. The **Create Volume** dialog box displays.

Figure 7.8. Create Volume

3. Select the cluster from the **Volume Cluster** drop-down list.
4. In the **Name** field, enter the name of the volume.
5. Select the type of the volume from the **Type** drop-down list. You can choose the volume type as Distribute, Replicate, or Stripe.
6. Add bricks to your volume. Bricks must be created externally on the Red Hat Storage nodes. Click **Add Bricks**.
7. Select the **Access Protocol** for the new volume by selecting **GlusterFS**, or **NFS**, or both checkboxes.
8. In the **Allow Access From** field, specify the volume access control as a comma-separated list of IP addresses or hostnames.  
You can use wildcards to specify ranges of addresses such as an asterisk (\*) which specifies all IP addresses or hostnames. You need to use IP-based authentication for Gluster Filesystem and NFS exports.
9. Click **OK** to create the volume. The new volume is added and displays on the **Volume** tab.

#### Result

You've added a Red Hat Storage volume to store virtual machines on.

[Report a bug](#)

### 7.8.4. Adding Bricks to a Volume

#### Summary

You can expand your cluster by adding new bricks. You need to add at least one brick to a distributed volume, multiples of two bricks to replicated volumes, multiples of four bricks to striped volumes when expanding your storage space.

#### Procedure 7.18. Adding Bricks to a Volume

1. On the **Volumes** tab on the navigation pane, select the volume to which you want to add bricks.
2. Select the volume to add new bricks. Click the **Bricks** tab from the Details pane.
3. Click **Add Bricks** to open the **Add Bricks** window.

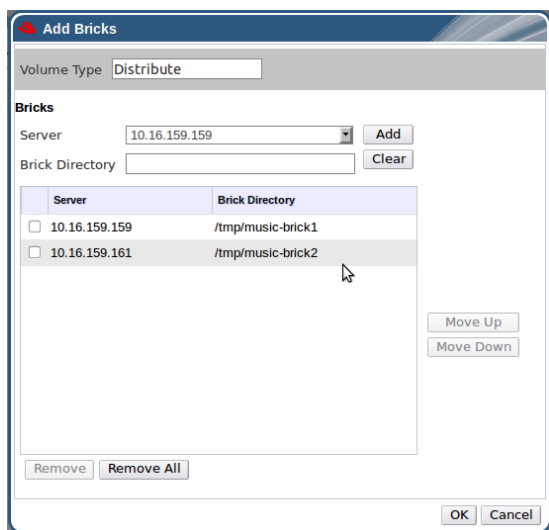


Figure 7.9. Add Bricks

4. Use the **Server** drop-down menu to select the server on which the brick will reside.
5. Enter the path for the **Brick Directory**. The directory or its parent directory must already exist.
6. Click **Add**. A list of available bricks appears, with server addresses and brick directory names.
7. Tick the checkbox of the bricks you wish to use.
8. Click **OK**.

**Result**

The new bricks are added to the volume and the bricks displays in the volume's **Bricks** tab.

[Report a bug](#)

**7.8.5. Explanation of Settings in the Add Bricks window**

Table 7.2. Add Bricks Tab Properties

Field Name	Description
<b>Volume Type</b>	Displays the type of volume. This field cannot be changed, it was set when you created the volume.
<b>Server</b>	The selected server to add new bricks.
<b>Brick Directory</b>	The directory in the server.

[Report a bug](#)

**7.8.6. Enabling Red Hat Storage Volumes for Virtualization****Summary**

You can enable a Red Hat Storage volume for virtualization using Red Hat Enterprise Virtualization Manager by creating a cluster with the **Enable Gluster Service** option selected. A volume is a logical collection of bricks, where each brick is an export directory on a server in the trusted storage pool. Most of the management operations of Red Hat Storage occur on the volume.

Red Hat recommends that you use separate data centers for Red Hat Storage nodes.

**Important**

**Enable Gluster Service** is a Technology Preview feature. Technology Preview features are not fully supported under Red Hat Subscription Service Level Agreements (SLAs), may not be functionally complete, and are not intended for production use. However, these features provide early access to upcoming product innovations, enabling customers to test functionality and provide feedback during the development process.

**Important**

Red Hat Storage currently only supports Red Hat Enterprise Virtualization 3.1. All Gluster clusters and hosts must be attached to data centers which are compatible with version 3.1.

**Procedure 7.19. Enabling Red Hat Storage Volumes for Virtualization**

1. On the **Data Centers** tab, click **New** to create a new data center. Set its **Compatibility Version** to **3.1** and its storage **Type** to **POSIX compliant FS**. See [Section 3.5.1, "Creating a New Data Center"](#) for details.
2. You can either create a new Gluster-enabled cluster and install at least one new host, or import an existing Gluster-enabled cluster and all its attached hosts.
  - A. New cluster and host: On the **Clusters** tab, click **New** to create a new cluster. Choose the **Data Center** you just created, enter all required details and select the **Enable Gluster Service** checkbox. See [Section 4.2.1, "Creating a New Cluster"](#) for details.  
On the **Hosts** tab, click **New** to add a new Red Hat Storage host. Choose the **Data Center** and **Cluster** you just created, enter all required details to install the host. See [Section 6.4.1, "Adding a Red Hat Enterprise Linux Host"](#) for details.
  - B. Existing cluster and host: On the **Clusters** tab, click **New** to import your existing Gluster enabled cluster and all its attached hosts. Choose the **Data Center** you just created, enter all required details and select the **Enable Gluster Service** checkbox. Select the **Import existing gluster configuration** checkbox and enter all required details. See [Section 4.2.4, "Importing an Existing Red Hat Storage Cluster"](#) for details.
3. On the **Volumes** tab, click **Create Volume** and enter all required information. Add bricks to the volume using the **Add Bricks** button. See [Section 7.8.3, "Creating A Storage Volume"](#) for details.
4. Select the newly created volume and click **Optimize for Virt Store**.

The volume is tuned and **group**, **storage-owner-uid**, and **storage-owner-gid** options are set.

5. Select the new volume and click **Start** to start the volume. The volume status changes to Up. You must optimize the volume before you can attach the volume to your storage domain.

#### Result

You can now create a storage domain using the optimized Red Hat Storage volume and manage it using the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

### 7.8.7. Starting Volumes

#### Summary

After a volume has been created or an existing volume has been stopped, it needs to be started before it can be used.

#### Procedure 7.20. Starting Volumes

1. In the **Volumes** tab, select the volume to be started.  
You can select multiple volumes to start by using **Shift** or **Ctrl** key.
2. Click the **Start** button.

The volume status changes to **Up**.

#### Result

You can now use your volume for virtual machine storage.

[Report a bug](#)

### 7.8.8. Tuning Volumes

#### Summary

Tuning volumes allows you to affect their performance. To tune volumes, you add options to them.

#### Procedure 7.21. Tuning Volumes

1. Click the **Volumes** tab.  
A list of volumes displays.
2. Select the volume that you want to tune, and click the **Volume Options** tab from the Details pane.  
The **Volume Options** tab displays a list of options set for the volume.
3. Click **Add** to set an option. The **Add Option** dialog box displays. Select the Option Key from the drop down list and enter the option value.

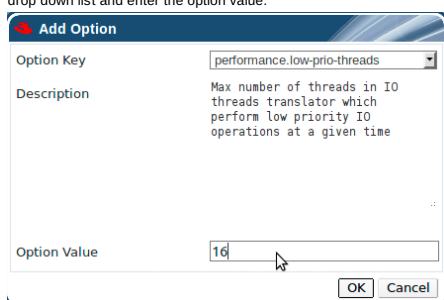


Figure 7.10. Add Option

4. Click **OK**.  
The option is set and displays in the **Volume Options** tab.

#### Result

You have tuned the options for your storage volume.

[Report a bug](#)

### 7.8.9. Editing Volume Options

#### Summary

You have tuned your volume by adding options to it. You can change the options for your storage volume.

#### Procedure 7.22. Editing Volume Options

1. Click the **Volumes** tab.  
A list of volumes displays.
2. Select the volume that you want to edit, and click the **Volume Options** tab from the Details pane.  
The **Volume Options** tab displays a list of options set for the volume.
3. Select the option you wish to edit. Click **Edit**. The **Edit Option** dialog box displays. Enter a new value for the option.
4. Click **OK**.  
The edited option displays in the **Volume Options** tab.

#### result

You've changed the options on your volume.

[Report a bug](#)

### 7.8.10. Reset Volume Options

#### Summary

You can reset options to revert them to their default values.

1. Click the **Volumes** tab.  
A list of volumes displays.

2. Select the volume and click the **Volume Options** tab from the Details pane.  
The **Volume Options** tab displays a list of options set for the volume.
3. Select the option you wish to reset. Click **Reset**. A dialog box displays, prompting to confirm the reset option.
4. Click **OK**.  
The selected option is reset.

**Note**

You can reset all volume options by clicking **Reset All** button. A dialog box displays, prompting to confirm the reset option. Click **OK**. All volume options are reset for the selected volume.

**Result**

You've reset volume options to default.

[Report a bug](#)

**7.8.11. Removing Bricks from a Volume****Summary**

You can shrink volumes, as needed, while the cluster is online and available. For example, you might need to remove a brick that has become inaccessible in a distributed volume due to hardware or network failure.

**Procedure 7.23. Removing Bricks from a Volume**

1. On the **Volumes** tab on the navigation pane, select the volume from which you wish to remove bricks.
2. Click the **Bricks** tab from the Details pane.
3. Select the bricks you wish to remove. Click **Remove Bricks**.
4. A window opens, prompting to confirm the deletion. Click **OK** to confirm.

**Result**

The bricks are removed from the volume.

[Report a bug](#)

**7.8.12. Stopping Red Hat Storage Volumes****Summary**

After a volume has been started, it can be stopped.

**Procedure 7.24. Stopping Volumes**

1. In the **Volumes** tab, select the volume to be stopped.  
You can select multiple volumes to stop by using **Shift** or **Ctrl** key.
2. Click **Stop**.

**Result**

The volume status changes is **Down**.

[Report a bug](#)

**7.8.13. Deleting Red Hat Storage Volumes****Summary**

You can delete a volume or multiple volumes from your cluster.

1. In the **Volumes** tab, select the volume to be deleted.
2. Click **Remove**. A dialog box displays, prompting to confirm the deletion. Click **OK**.

**Result**

The volume is removed from the cluster.

[Report a bug](#)

**7.9. Storage and Permissions****7.9.1. Managing system permissions for a storage domain**

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **StorageAdmin** has administrator privileges only for the assigned storage domain, and so forth.

A storage administrator is a system administration role for a specific storage domain only. This is useful in data centers with multiple storage domains, where each storage domain requires a system administrator. Use the **Configure** button in the header bar to assign a storage administrator for all storage domains in the environment.

The storage domain administrator role permits the following actions:

- » Edit the configuration of the storage domain;
- » Move the storage domain into maintenance mode; and
- » Remove the storage domain.

**Note**

You can only assign roles and permissions to existing users.

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

**See Also:**

- » [Section 7.9.2, "Storage Administrator Roles Explained"](#)

[Report a bug](#)

## 7.9.2. Storage Administrator Roles Explained

### Storage Domain Permission Roles

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

**Table 7.3. Red Hat Enterprise Virtualization System Administrator Roles**

Role	Privileges	Notes
StorageAdmin	Storage Administrator	Can create, delete, configure and manage a specific storage domain.
GlusterAdmin	Gluster Storage Administrator	Can create, delete, configure and manage Gluster storage volumes.

#### See Also:

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

## 7.9.3. Assigning an Administrator or User Role to a Resource

### Summary

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

#### Procedure 7.25. Assigning a Role to a Resource

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

#### Result

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

#### See Also:

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

## 7.9.4. Removing an Administrator or User Role from a Resource

### Summary

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

#### Procedure 7.26. Removing a Role from a Resource

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

#### Result

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

[Report a bug](#)

## Chapter 8. Virtual Machines

### 8.1. Introduction to Virtual Machines

A virtual machine is a software implementation of a computer. The Red Hat Enterprise Virtualization environment enables you to create virtual desktops and virtual servers.

Virtual machines consolidate computing tasks and workloads. In traditional computing environments, workloads usually run on individually administered and upgraded servers. Virtual machines reduce the amount of hardware and administration required to run the same computing tasks and workloads.

[Report a bug](#)

### 8.2. Supported Virtual Machine Operating Systems

Red Hat Enterprise Virtualization supports the virtualization of these guest operating systems:

**Table 8.1. Supported virtual machine operating systems**

Operating System	Architecture	SPICE support
Red Hat Enterprise Linux 3	32-bit, 64-bit	Yes
Red Hat Enterprise Linux 4	32-bit, 64-bit	Yes
Red Hat Enterprise Linux 5	32-bit, 64-bit	Yes
Red Hat Enterprise Linux 6	32-bit, 64-bit	Yes
Windows XP Service Pack 3 and newer	32 bit	Yes
Windows 7	32-bit, 64-bit	Yes
Windows 8	32-bit, 64-bit	No
Windows Server 2003 Service Pack 2 and newer	32-bit, 64-bit	Yes
Windows Server 2008	32 bit, 64 bit	Yes
Windows Server 2008 R2	64-bit	Yes
Windows Server 2012 R2	64-bit	No

Remote Desktop Protocol (RDP) is the default connection protocol for accessing Windows 8 and Windows 2012 guests from the user portal as Microsoft introduced changes to the Windows Display Driver Model which prevent SPICE from performing optimally.

[Report a bug](#)

### 8.3. Virtual Machine Performance Parameters

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

**Table 8.2. Supported virtual machine parameters**

Parameter	Number	Note
Virtualized CPUs	160	per virtual machine
Virtualized RAM	2TB	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

[Report a bug](#)

### 8.4. Creating Virtual Machines

#### 8.4.1. Creating a New Virtual Machine from an Existing Template

##### Summary

You can use a template to create a virtual machine which has already been configured with virtual disks, network interfaces, an operating system, and applications.

A virtual machine created from a template depends on the template. You cannot remove a template from the environment if there are still virtual machines that were created from it. Cloning a virtual machine from a template removes the dependency on the template.

##### Procedure 8.1. Creating a New Virtual Machine from an Existing Template

1. Click the **Virtual Machines** resource tab to list all the virtual machines in the results list. The icon to the right of the virtual machine name indicates whether it is a virtual server, a virtual machine, or a part of a virtual machine pool.
2. Click the **New Server** button to open the **New Server Virtual Machine** window, or the **New Desktop** button to open the **New Desktop Virtual Machine** window.
3. Select the **Data Center** and **Host Cluster** on which the desktop is to run. Select an existing template from the **Based on Template** drop-down menu.



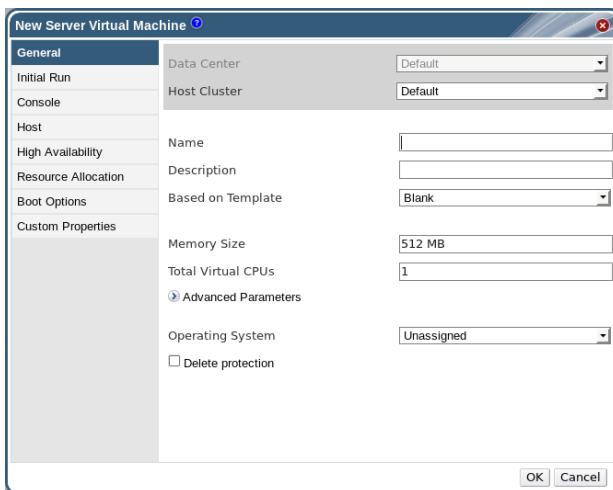


Figure 8.1. New Virtual Machine Window

4. Enter a suitable **Name** and **Description**, and accept the default values inherited from the template. You can change the rest of the fields if needed.
5. Click **OK**.

**Result**

The virtual machine is created and displayed in the **Virtual Machines** list. You can now log on to your virtual machine and begin using it, or assign users to it.

**See Also:**

- » [Section 8.4.3, "Explanation of Settings and Controls in the New Virtual Machine and Edit Virtual Machine Windows"](#)
- » [Section 8.4.2, "Creating a New Virtual Machine from a Blank Template"](#)

[Report a bug](#)

**8.4.2. Creating a New Virtual Machine from a Blank Template****Summary**

You can create a virtual machine using a blank template and configure all of its settings.

**Procedure 8.2. Creating a New Virtual Machine from a Blank Template**

1. Click the **Virtual Machines** resource tab to list all the virtual machines in the results list. The icon to the right of the virtual machine name indicates whether it is a virtual server, a virtual machine, or a part of a virtual machine pool.
2. Click the **New Server** button to open the **New Server Virtual Machine** window, or the **New Desktop** button to open the **New Desktop Virtual Machine** window.
3. On 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.
4. Alternatively, click the **Initial Run**, **Console**, **Host**, **Resource Allocation**, **Boot Options**, and **Custom Properties** tabs in turn to define options for your virtual machine.
5. Click **OK** to create the virtual machine and close the window.
6. The **New Virtual Machine - Guide Me** window opens. Use the Guide Me buttons to complete configuration or click **Configure Later** to close the window.

**Result**

The new virtual machine is created and displays in the list of virtual machines with a status of **Down**. Before you can use this virtual machine, add at least one network interface and one virtual disk, and install an operating system.

**See Also:**

- » [Section 8.4.3, "Explanation of Settings and Controls in the New Virtual Machine and Edit Virtual Machine Windows"](#)

[Report a bug](#)

**8.4.3. Explanation of Settings and Controls in the New Virtual Machine and Edit Virtual Machine Windows****8.4.3.1. Virtual Machine General Settings Explained**

These settings apply to adding or editing new virtual machines.

The **Virtual Machine: General** settings table details the information required on the **General** tab of the **New** or **Edit** windows.

Table 8.3. Virtual Machine: General Settings

Field Name	Description
<b>Data Center</b>	The data center to which the virtual machine is attached.
<b>Host Cluster</b>	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.
<b>Quota (Server only)</b>	A policy that limits the virtual machine from using more CPU, memory, or storage resources than are allocated.
<b>Name</b>	The name of virtual machine. Names must not contain any spaces, and must contain at least one character from A-Z. The maximum length of a virtual machine name is 15 characters.
<b>Description</b>	A meaningful description of the new virtual machine.
<b>Based on Template</b>	Templates can be used to create a virtual machines from existing models. This field is set to <i>Blank</i> by default, which enables creating a virtual machine from scratch.
<b>Memory Size</b>	The amount of memory assigned to the virtual machine. When allocating memory, consider the processing and storage needs of the applications that are intended to run on the virtual machine. Maximum guest memory is constrained by the selected guest architecture and the cluster compatibility level.
<b>Total Virtual CPUs</b>	The processing power allocated to the virtual machine as CPU Cores. Do not assign more cores to a virtual machine than are present on the physical host.
<b>Cores per Virtual Socket</b>	The number of cores assigned to each virtual socket.
<b>Virtual Sockets</b>	The number of CPU sockets for the virtual machine. Do not assign more sockets to a virtual machine than are present on the physical host.
<b>Operating System</b>	The operating system. Valid values include a range of Red Hat Enterprise Linux and Windows variants.
<b>Stateless (Desktop only)</b>	Select this check box if the virtual machine is to run in stateless mode. The stateless mode is used primarily for desktop virtual machines. Running a stateless desktop or server creates a new COW layer on the virtual machine hard disk image where new and changed data is stored. Shutting down the stateless virtual machine deletes the new COW layer, returning the virtual machine to its original state. This type of virtual machine is useful when creating virtual machines that need to be used for a short time, or by temporary staff.
<b>Delete protection</b>	Select this check box to make deletion of the virtual machine impossible. It is possible to delete the virtual machine only when this check box is not selected.

[Report a bug](#)

#### 8.4.3.2. Virtual Machine Initial Run Settings Explained

These settings apply to adding or editing new virtual machines.

The **Virtual Machine: Initial Run** settings table details the information required on the **Initial Run** tab of the **New** or **Edit** windows.

Table 8.4. Virtual Machine: Initial Run Settings

Field Name	Description
<b>General - Time Zone</b>	The time zone in which the virtual machine is to run. It is not necessarily the time zone for the physical host on which the virtual machine is running.
<b>Windows - Domain</b>	The domain in which the virtual machine is to run. This option is only available when Windows is selected as the operating system on the <b>Virtual Machine - General</b> tab.

[Report a bug](#)

#### 8.4.3.3. Virtual Machine Console Settings Explained

These settings apply to adding or editing new virtual machines.

The **Virtual Machine: Console** settings table details the information required on the **Console** tab of the **New** or **Edit** windows.

Table 8.5. Virtual Machine: Console Settings

Field Name	Description
<b>Protocol</b>	Defines the display protocol to be used. <b>SPICE</b> is the recommended protocol for Linux and Windows virtual machines, excepting Windows 8 and Windows Server 2012. Optionally, select <b>VNC</b> for Linux virtual machines. A VNC client is required to connect to a virtual machine using the VNC protocol.
<b>USB Support</b>	<p>Defines whether USB devices can be used on the virtual machine. This option is only available for virtual machines using the SPICE protocol. Select either:</p> <ul style="list-style-type: none"> <li>▀ <b>Disabled</b> - Does not allow USB redirection from the client machine to the virtual machine.</li> <li>▀ <b>Legacy</b> - Enables the SPICE USB redirection policy used in Red Hat Enterprise Virtualization 3.0. This option can only be used on Windows virtual machines, and will not be supported in future versions of Red Hat Enterprise Virtualization.</li> <li>▀ <b>Native</b> - Enables native KVM/ SPICE USB redirection for Linux and Windows virtual machines. Virtual machines do not require any in-guest agents or drivers for native USB. This option can only be used if the virtual machine's cluster compatibility version is set to 3.1 or higher.</li> </ul>
<b>Monitors</b>	The number of monitors for the virtual machine. This option is only available for virtual desktops using the SPICE display protocol. You can choose 1, 2 or 4. Since Windows 8 and Windows Server 2012 virtual machines do not support the SPICE protocol, they do not support multiple monitors.
<b>Enable Smartcard</b>	<p>Smartcards are an external hardware security feature, most commonly seen in credit cards, but also used by many businesses as authentication tokens. Smartcards can be used to protect Red Hat Enterprise Virtualization virtual machines. Tick or untick the checkbox to activate and deactivate Smartcard authentication for individual virtual machines.</p>
<b>Disable strict user checking</b>	<p>Click the <b>Advanced Parameters</b> arrow and select the checkbox to use this option. With this option selected, the virtual machine does not need to be rebooted when a different user connects to it.</p> <p>By default, strict checking is enabled so that only one user can connect to the console of a virtual machine. No other user is able to open a console to the same virtual machine until it has been rebooted. The exception is that a <b>SuperUser</b> can connect at any time and replace a existing connection. When a <b>SuperUser</b> has connected, no normal user can connect again until the virtual machine is re-booted.</p> <p>Disable strict checking with caution, because you can expose the previous user's session to the new user</p>

[Report a bug](#)

#### 8.4.3.4. Virtual Machine Host Settings Explained

These settings apply to adding or editing new virtual machines.

The **Virtual Machine: Host** settings table details the information required on the **Host** tab of the **New** or **Edit** windows.

Table 8.6. Virtual Machine: Host Settings

Field Name	Description
<b>Max free Memory for scheduling new VMs</b>	Reports the maximum amount of memory that remains for the hypervisor to use when creating new virtual machines. The value in this field is calculated prior to taking into account other scheduling policies, which may further reduce the amount of memory available for creating virtual machines. If an attempt to create a virtual machine fails and <b>Max free Memory for scheduling new VMs</b> only narrowly exceeds the amount of memory you need to create the virtual machine, the other scheduling policies have most likely reduced the amount of free memory to a level at which Red Hat Enterprise Virtualization cannot create a new virtual machine.
<b>Run On</b>	Defines the host on which the virtual machine is to run. Select either: <ul style="list-style-type: none"> <li>» <b>Any Host in Cluster</b> - The virtual machine can start and run on any available host in the cluster.</li> <li>» <b>Specific</b> - The virtual machine must run on a particular host in the cluster. Select the specific host from the drop-down menu of available hosts.</li> </ul>
<b>Migration Options</b>	Defines options to run and migrate the virtual machine. If the options here are not used, the virtual machine will run or migrate according to its cluster's policy. <ul style="list-style-type: none"> <li>» <b>Run VM on the selected host</b> (no migration allowed) - Starts and runs the virtual machine only on the selected host. Select this option if your virtual machine does not need to be highly available.</li> <li>» <b>Allow VM migration only upon Administrator specific request</b> (system will not trigger automatic migration of this VM) - Prevents the virtual machine from being migrated mid-operation. For example in cases of host overload or fencing, the virtual machine will remain on the host on which it was started.</li> </ul>
<b>CPU Pinning topology</b>	Enables the virtual machine's virtual CPU (vCPU) to run on a specific physical CPU (pCPU) in a specific host. This option is not supported if the virtual machine's cluster compatibility version is set to 3.0. The syntax of CPU pinning is <b>v#p[_v#p]</b> , for example: <ul style="list-style-type: none"> <li>» <b>0#0</b> - Pins vCPU 0 to pCPU 0.</li> <li>» <b>0#0_1#3</b> - Pins vCPU 0 to pCPU 0, and pins vCPU 1 to pCPU 3.</li> <li>» <b>1#1-4,^2</b> - Pins vCPU 1 to one of the pCPUs in the range of 1 to 4, excluding pCPU 2.</li> </ul> <p>In order to pin a virtual machine to a host, you must select <b>Do not allow migration</b> under <b>Migration Options</b>, and select the <b>Use Host CPU</b> check box.</p>

[Report a bug](#)**8.4.3.5. Virtual Machine High Availability Settings Explained**

These settings apply to adding or editing new server virtual machines.

The **Virtual Machine: High Availability** settings table details the information required on the **High Availability** tab of the **New** or **Edit** windows.

Table 8.7. Virtual Machine: High Availability Settings

Field Name	Description
<b>Highly Available</b>	Select this check box if the virtual machine is to be highly available. For example, in cases of host maintenance or failure, the virtual machine will be automatically re-run on another host. If the host is manually shut down by the system administrator, the virtual machine is not automatically moved to another host.
<b>Priority for Run/Migration queue</b>	Sets the priority level for the virtual machine to be migrated or restarted on another host.

[Report a bug](#)**8.4.3.6. Virtual Machine Resource Allocation Settings Explained**

These settings apply to adding or editing new virtual machines.

The **Virtual Machine: Resource Allocation** settings table details the information required on the **Resource Allocation** tab of the **New** or **Edit** windows.

Table 8.8. Virtual Machine: Resource Allocation Settings

Field Name	Sub-element	Description
Memory Allocation		The amount of physical memory guaranteed for this virtual machine.
Storage Allocation		The <b>Template Provisioning</b> option is only available when the virtual machine is created from a template.
	Thin	Provides optimized usage of storage capacity. Disk space is allocated only as it is required.
	Clone	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.
	Alias	The name of the virtual disk.
	Allocation policy	Storage allocation policy for disks created when creating a virtual machine from a template as a <b>Clone</b> . <ul style="list-style-type: none"> <li>» <b>Preallocated</b> results is 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.</li> <li>» <b>Thin Provisioned</b> results is a faster clone operation and provides optimized usage of storage capacity. Disk space is allocated only as it is required.</li> </ul>
	Target	The storage domain for the virtual disk.
	Quota	The type of quota enforced on the virtual disk.

[Report a bug](#)

#### 8.4.3.7. Virtual Machine Boot Options Settings Explained

These settings apply to adding or editing new virtual machines.

The **Virtual Machine: Boot Options** settings table details the information required on the **Boot Options** tab of the **New** or **Edit** windows.

Table 8.9. Virtual Machine: Boot Options Settings

Field Name	Description
<b>First Device</b>	<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"> <li>» <b>Hard Disk</b></li> <li>» <b>CD-ROM</b></li> <li>» <b>Network (PXE)</b></li> </ul>
<b>Second Device</b>	<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>
<b>Attach CD</b>	<p>If you have selected <b>CD-ROM</b> as a boot device, tick this check box and select a CD-ROM image from the drop-down menu. The images must be available in the ISO domain.</p>
<b>Boot Options</b>	<p>Defines the virtual machine's boot sequence, running options, and source images for installing the operating system and required drivers.</p> <ul style="list-style-type: none"> <li>» <b>Attach Floppy</b> - Attaches a floppy disk image to the virtual machine. Use this option to install Windows drivers. The floppy disk image must reside in the ISO domain.</li> <li>» <b>Attach CD</b> - Attaches an ISO image to the virtual machine. Use this option to install the virtual machine's operating system and applications. The CD image must reside in the ISO domain.</li> <li>» <b>Boot Sequence</b> - Determines the order in which the boot devices are used to boot the virtual machine. Select either <b>Hard Disk</b>, <b>CD-ROM</b> or <b>Network</b>, and use the arrow keys to move the option up or down.</li> <li>» <b>Run Stateless</b> - Deletes all changes to the virtual machine upon shutdown.</li> <li>» <b>Start in Pause Mode</b> - Starts then pauses the virtual machine to enable connection to the console, suitable for virtual machines in remote locations.</li> </ul>
<b>Linux Boot Options</b>	<p>The following options boot a Linux kernel directly instead of through the BIOS bootloader.</p> <ul style="list-style-type: none"> <li>» <b>kernel path</b> - A fully-qualified path to a kernel image to boot the virtual machine. The kernel image must be stored on either the ISO domain (path name in the format of <b>iso://path-to-image</b>) or on the host's local storage domain (path name in the format of <b>/data/images</b>).</li> <li>» <b>initrd path</b> - A fully-qualified path to a ramdisk image to be used with the previously specified kernel. The ramdisk image must be stored on the ISO domain (path name in the format of <b>iso://path-to-image</b>) or on the host's local storage domain (path name in the format of <b>/data/images</b>).</li> <li>» <b>kernel params</b> - Kernel command line parameter strings to be used with the defined kernel on boot.</li> </ul>

[Report a bug](#)

#### 8.4.3.8. Virtual Machine Custom Properties Settings Explained

These settings apply to adding or editing new virtual machines.

The **Virtual Machine: Custom Properties** settings table details the information required on the **Custom Properties** tab of the **New** or **Edit** windows.

Table 8.10. Virtual Machine: Custom Properties Settings

Field Name	Description	Recommendations and Limitations
<b>sap_agent</b>	Enables SAP monitoring on the virtual machine. Set to <b>true</b> or <b>false</b> .	-
<b>sndbuf</b>	Enter the size of the buffer for sending the virtual machine's outgoing data over the socket. Default value is 0.	-
<b>vhost</b>	<p>Disables vhost-net, which is the kernel-based virtio network driver on virtual network interface cards attached to the virtual machine. To disable vhost, the format for this property is:</p> <div> <b>LogicalNetworkName:</b> false </div> <p>This will explicitly start the virtual machine without the vhost-net setting on the virtual NIC attached to <b>LogicalNetworkName</b>.</p>	vhost-net provides better performance than virtio-net, and if it is present, it is enabled on all virtual machine NICs by default. Disabling this property makes it easier to isolate and diagnose performance issues, or to debug vhost-net errors, for example if migration fails for virtual machines on which vhost does not exist.
<b>viodiskcache</b>	Caching mode for the virtio disk. <b>writethrough</b> writes data to the cache and the disk in parallel, <b>writeback</b> does not copy modifications from the cache to the disk, and <b>none</b> disables caching.	For Red Hat Enterprise Virtualization 3.1, if viodiskcache is enabled, the virtual machine cannot be live migrated.

**Warning**

Increasing the value of the sndbuf custom property results in increased occurrences of communication failure between hosts and unresponsive virtual machines.

[Report a bug](#)

#### 8.4.4. Creating a Cloned Virtual Machine from an Existing Template

##### Summary

Cloning a virtual machine from a template is like creating a virtual machine from a template. A cloned virtual machine inherits all the settings from the original virtual machine on which its template is based. A clone does not depend on the template it was created from after it has been created.

##### Procedure 8.3. Creating a Cloned Virtual Machine from an Existing Template

1. Click the **Virtual Machines** resource tab to list all the virtual machines in the results list.
2. Click the **New Server** button to open the **New Server Virtual Machine** window, or the **New Desktop** button to open the **New Desktop Virtual Machine** window.
3. Select an existing template from the **Based on Template** drop-down menu.
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.
5. Click the **Resource Allocation** tab. The template you selected is displayed on the **Template Provisioning** field. Select **Clone**.

**New Server Virtual Machine**

General

Initial Run

Console

Host

High Availability

**Resource Allocation**

Boot Options

Custom Properties

**Memory Allocation:**

Physical Memory Guaranteed: 1024 MB

**Storage Allocation: (Available only when a template is selected)**

Template Provisioning: ☐ Thin ☒ Clone

Alias	Virtual Size	Allocation	Target
RHEL6_Thame	10 GB	Thin Provision	Storage_robot

OK Cancel

Figure 8.2. Provisioning - Clone

Select the disk provisioning mode in the **Allocation** 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. Select the **Target** storage domain for the virtual machine.
  7. Click **OK**.

**Note**

Cloning a virtual machine may take some time. A new copy of the template's disk must be created. During this time, the virtual machine's status is first **Image Locked**, then **Down**.

**Result**

The virtual machine is created and displayed in the Virtual Machines list. You can now log on to your virtual machine and begin using it, or assign users to it.

[Report a bug](#)

### 8.4.5. Completing the Configuration of a Virtual Machine by Defining Network Interfaces and Hard Disks

**Summary**

Before you can use your newly created virtual machine, the **Guide Me** window prompts you to configure at least one network interface and one virtual disk for the virtual machine.

#### Procedure 8.4. Completing the Configuration of a Virtual Machine by Defining Network Interfaces and Hard Disks

1. On the **New Virtual Machine - Guide Me** window, click the **Configure Network Interfaces** button to open the **New Network Interface** window. You can accept the default values or change them as necessary.

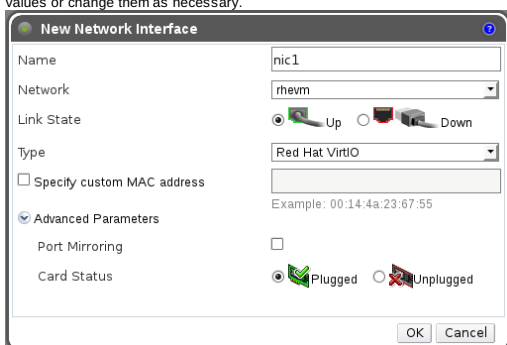


Figure 8.3. New Network Interface window

2. Enter the **Name** of the network interface.
3. Use the drop-down menus to select the **Network** and the **Type** of network interface for the new virtual machine. The **Link State** is set to **Up** by default when the NIC is defined on the virtual machine and connected to the network.

**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.

3. If applicable, select the **Specify custom MAC address** check box and enter the network interface's MAC address.
4. Click the arrow next to **Advanced Parameters** to configure the **Port Mirroring** and **Card Status** fields, if necessary.
5. Click **OK** to close the **New Network Interface** window and open the **New Virtual Machine - Guide Me** window.
6. Click the **Configure Virtual Disk** button to open the **New Virtual Disk** window.
7. Add either an **Internal** virtual disk or an **External LUN** to the virtual machine.

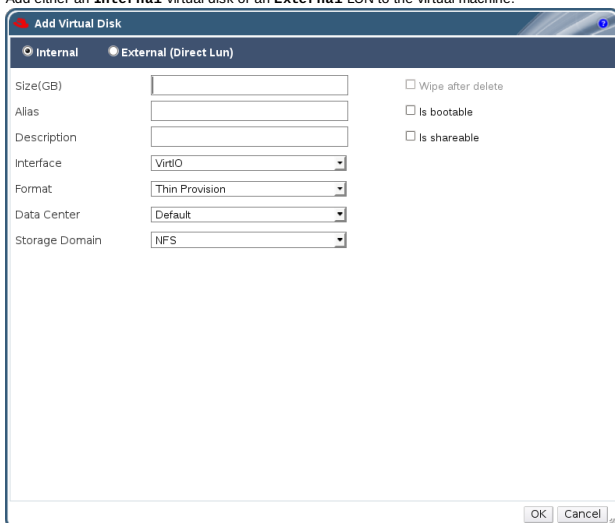


Figure 8.4. New Virtual Disk Window

8. Click **OK** to close the **New Virtual Disk** window. The **New Virtual Machine - Guide Me** window opens with changed context. There is no further mandatory configuration.
9. Click **Configure Later** to close the window.

**Result**



You have added a network interface and a virtual disk to your virtual machine.

See Also:

- » [Section 8.7.5, "Explanation of Settings in the Virtual Machine Network Interface Window"](#)
- » [Section 11.5, "Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows"](#)

[Report a bug](#)

8.4.6. Installing a Guest Operating System onto a Virtual Machine

Prerequisites:

- » [Section 8.4.2, "Creating a New Virtual Machine from a Blank Template"](#)
- » [Section 8.7.5, "Explanation of Settings in the Virtual Machine Network Interface Window"](#)
- » [Section 11.5, "Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows"](#)

Summary

An operating system has to be installed onto a virtual machine that is created from a blank template. You can install a new operating system on any virtual machine.

Procedure 8.5. Installing an operating system onto a virtual machine

1. Select the created virtual machine. It has a status of **Down**.
2. Click the **Run Once** button to open the **Run Virtual Machine** window.




Figure 8.5. Run Virtual Machine Window

3. Click the **Boot Options** tab to define the boot sequence and source images for installing the operating system.
4. Click the **Linux Boot Options** tab to define additional boot options specific to Linux virtual machines.
5. Click the **Initial Run** tab to join the virtual machine to a domain on the initial run.
6. Click the **Display Protocol** tab and select a suitable protocol to connect to the virtual machine. SPICE is the recommended protocol.
7. Click the **Custom Properties** tab to enter additional running options for virtual machines.
8. Click **OK**.

Result

You have installed an operating system onto your virtual machine. You can now log in and begin using your virtual machine, or assign users to it.

[Report a bug](#)

8.4.7. Installing Windows on VirtIO Optimized Hardware

Summary

The **virtio-win.vfd** floppy image contains Windows drivers for VirtIO optimized disk and network devices. These drivers provide a performance improvement over emulated device drivers.

The **virtio-win.vfd** is placed automatically on ISO storage domains that are hosted on the Manager server. It must be manually uploaded using the **rhev-iso-uploader** tool to other ISO storage domains.

You can install the VirtIO optimized device drivers during your Windows installation by attaching a floppy disk device to your virtual machine.

This procedure presumes that you added a **Red Hat VirtIO** network interface and a disk that uses the **VirtIO** interface to your virtual machine.

Procedure 8.6. Installing VirtIO Drivers during Windows Installation

1. Select your virtual machine from the list in the **Virtual Machines** tab.
2. Click the **Run Once** button, and the **Run Once** window displays.
3. Click **Boot Options** to expand the **Boot Options** configuration options.
4. Click the **Attach Floppy** checkbox, and select **virtio-win.vfd** from the drop down selection box.
5. Click the **Attach CD** checkbox, and select from the drop down selection box the ISO containing the version of Windows you want to install.
6. Move **CD-ROM UP** in the **Boot Sequence** field.
7. Configure the rest of your **Run Once** options as required, and click **OK** to start your virtual machine, and then click the **Console** button to open a graphical console to your virtual machine.

Result

Windows installations include an option to load additional drivers early in the the installation process. Use this option load drivers from the **virtio-win.vfd** floppy disk which was attached to your virtual machine as **A:**.

For each supported virtual machine architecture and Windows version, there is a folder on the disk containing optimized hardware device drivers.

[Report a bug](#)

8.4.8. Virtual Machine Run Once Settings Explained

The **Run Once** window defines one-off boot options for a virtual machine. For persistent boot options, use the **Boot Options** tab in the **New Virtual Machine** window. The following table details the information required for the **Run Once** window.

Table 8.11. Virtual Machine: Run Once Settings

Field Name	Description
<b>Boot Options</b>	<p>Defines the virtual machine's boot sequence, running options, and source images for installing the operating system and required drivers.</p> <ul style="list-style-type: none"> <li>» <b>Attach Floppy</b> - Attaches a floppy disk image to the virtual machine. Use this option to install Windows drivers. The floppy disk image must reside in the ISO domain.</li> <li>» <b>Attach CD</b> - Attaches an ISO image to the virtual machine. Use this option to install the virtual machine's operating system and applications. The CD image must reside in the ISO domain.</li> <li>» <b>Boot Sequence</b> - Determines the order in which the boot devices are used to boot the virtual machine. Select either <b>Hard Disk</b>, <b>CD-ROM</b> or <b>Network</b>, and use the arrow keys to move the option up or down.</li> <li>» <b>Run Stateless</b> - Deletes all changes to the virtual machine upon shutdown.</li> <li>» <b>Start in Pause Mode</b> - Starts then pauses the virtual machine to enable connection to the console, suitable for virtual machines in remote locations.</li> </ul>
<b>Linux Boot Options</b>	<p>The following options boot a Linux kernel directly instead of through the BIOS bootloader.</p> <ul style="list-style-type: none"> <li>» <b>kernel path</b> - A fully-qualified path to a kernel image to boot the virtual machine. The kernel image must be stored on either the ISO domain (path name in the format of <b>iso://path-to-image</b>) or on the host's local storage domain (path name in the format of <b>/data/images</b>).</li> <li>» <b>initrd path</b> - A fully-qualified path to a ramdisk image to be used with the previously specified kernel. The ramdisk image must be stored on the ISO domain (path name in the format of <b>iso://path-to-image</b>) or on the host's local storage domain (path name in the format of <b>/data/images</b>).</li> <li>» <b>kernel params</b> - Kernel command line parameter strings to be used with the defined kernel on boot.</li> </ul>
<b>Initial Run</b>	<p>Defines the virtual machine's domain and user log in credentials.</p> <ul style="list-style-type: none"> <li>» <b>Domain</b> - The domain in which the virtual machine runs.</li> <li>» <b>Alternate Credentials</b> - By default you will log in to the virtual machine with the same credentials you use for the Administration Portal. To log in as a different user, enter a suitable <b>User Name</b> and <b>Password</b>. This user must exist in the same domain as the virtual machine.</li> </ul>
<b>Display Protocol</b>	<p>Defines the protocol to connect to virtual machines.</p> <ul style="list-style-type: none"> <li>» <b>VNC</b> - Can be used for Linux virtual machines. Requires a VNC client to connect to a virtual machine using VNC.</li> <li>» <b>SPICE</b> - Recommended protocol for Linux and Windows virtual machines, excepting Windows 8 and Server 2012 virtual machines.</li> </ul>
<b>Custom Properties</b>	<p>Additional VDSM options for running virtual machines.</p> <ul style="list-style-type: none"> <li>» <b>sap_agent</b> - Enables SAP monitoring on the virtual machine. Set to <b>true</b> or <b>false</b>.</li> <li>» <b>sndbuf</b> - Enter the size of the buffer for sending the virtual machine's outgoing data over the socket.</li> <li>» <b>vhost</b> - Enter the name of the virtual host on which this virtual machine should run. The name can contain any combination of letters and numbers.</li> <li>» <b>viodiskcache</b> - Caching mode for the virtio disk. <b>writethrough</b> writes data to the cache and the disk in parallel, <b>writeback</b> does not copy modifications from the cache to the disk, and <b>none</b> disables caching.</li> </ul>

[Report a bug](#)

## 8.5. Using Virtual Machines

### 8.5.1. SPICE

The SPICE protocol facilitates graphical connections to virtual machines. The SPICE protocol allows:

- » video at more than 30 frames per second
- » bi-directional audio (for soft-phones/IP phones)
- » bi-directional video (for video telephony/video conferencing)
- » connection to multiple monitors with a single virtual machine
- » USB redirection from the client's USB port into the virtual machine
- » connection to a proxy from outside of the network the hypervisor is attached to


[Report a bug](#)

### 8.5.2. Powering on a Virtual Machine

#### Summary

You can start a virtual machine from the Administration Portal.

#### Procedure 8.7. Powering on a Virtual Machine

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine in the results list with a status of **Down**.
2. Click the  icon.  
Alternatively, right-click and select **Run**.

#### Result

The **Status** of the virtual machine changes to **Up**. The display protocol of the selected virtual machine is displayed. If the virtual machine has the **rhev-guest-agent** installed, its IP address is also displayed.

[Report a bug](#)

### 8.5.3. Installing SPICE Plugins in Windows and Linux

The SPICE protocol is the default graphical protocol used to connect to virtual machines. Plugins for the Internet Explorer and Firefox web browsers allow you to launch graphical virtual machine connection sessions from the Red Hat Enterprise Virtualization web portals.

#### Summary

This procedure describes the installation of the SPICE Plugin for Mozilla Firefox on Linux clients.

#### Procedure 8.8. Installing the SPICE plugin for Mozilla Firefox on Red Hat Enterprise Linux

- » Open a terminal and run the following command as root:

```
# yum install spice-xpi
```

The plugin will be installed the next time Firefox is started.

#### Result

The SPICE plugin is installed on your Red Hat Enterprise Linux client.

#### Summary


This procedure describes the installation of the SPICE ActiveX component for Internet Explorer on Windows clients.

#### Procedure 8.9. Installing the SPICE ActiveX component for Internet Explorer on Windows

1. The first time you attempt to connect to a virtual machine, an add-on notification bar displays in the browser, prompting you to install the SPICE ActiveX component. You need administrative privileges on your client machine to install the component. Contact your systems administrator if you do not have the necessary permissions.
2. When you accept the prompt to install the SPICE ActiveX component, Internet Explorer may issue a security warning. Confirm that you wish to proceed, and the component will be installed.

#### Result

The SPICE ActiveX component for Internet Explorer for Windows is installed on your client machine.



**Important**

If you installed the SPICE ActiveX component without administrative permissions, you will receive a message stating that the *usbclerk* package was not installed. This means that you will be able to connect to a virtual machine using SPICE, however you will not be able to use USB devices on your virtual machine. Contact your systems administrator to install *usbclerk* if required.

[Report a bug](#)

### 8.5.4. Logging in to a Virtual Machine

#### Prerequisites:

- » [Section 8.5.3, "Installing SPICE Plugins in Windows and Linux"](#)

#### Summary

The default protocol for graphical connections to virtual machines is SPICE. You can log in to virtual machines using the SPICE protocol from the Administration Portal. An external VNC client is required to log in to virtual machines using the VNC protocol.

#### Procedure 8.10. Logging in to a virtual machine

1. On the **Virtual Machines** resource tab, select a running virtual machine.
2. Click the **Console** button or right-click the virtual machine and select **Console** from the menu.




Figure 8.6. Connection Icon on the Virtual Machine Menu

- A. If the virtual machine's display protocol is set to SPICE, a console window to the virtual machine opens. Log in to the virtual machine's guest operating system.
- B. If the virtual machine's display protocol is set to VNC, a window containing connection details including the IP address, port number and password required opens. Use the details with your VNC client to connect to the virtual machine.

#### Result

You have connected to a virtual machine from the Administration Portal using SPICE or a VNC client.

#### See Also:

- » [Section 8.4.3.3, "Virtual Machine Console Settings Explained"](#)

[Report a bug](#)

## 8.6. Shutting Down or Pausing Virtual Machines

### 8.6.1. Shutting Down or Pausing Virtual Machines

Virtual machine should be shut down from within. 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 on the virtual machine*. Shutdown of virtual machines should be planned after due consideration, preferably at times that will least impact users.

All users should be logged off of a Windows virtual machine before shutting them down. If any users are still logged in, the virtual machine remains on with a **Powering Off** status in the Administration Portal. The virtual machine requires manual intervention to shut it down completely because the following Windows message is displayed 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?*

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.

Pausing a virtual machine puts it into *Hibernate* mode, where the virtual machine state is preserved. Applications running in RAM are written to the hard drive and CPU usage is zero.


[Report a bug](#)

### 8.6.2. Shutting Down a Virtual Machine

#### Summary

If your virtual machine has the *rhev-guest-agent* installed, or has Advanced Configuration and Power Interface (ACPI) support, you can shut it down from the Administration Portal.

#### Procedure 8.11. Shutting Down a Virtual Machine

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a running virtual machine in the results list.
2. Click the **Shut down** () button.  
Alternatively, right-click the virtual machine and select **Shut down**.

#### Result

The **Status** of the virtual machine changes to **Down**.


[Report a bug](#)

### 8.6.3. Pausing a Virtual Machine

#### Summary

If your virtual machine has the *rhev-guest-agent* installed, or has Advanced Configuration and Power Interface (ACPI) support, you can pause it from the Administration Portal. This is equal to setting it on *Hibernate* mode.

#### Procedure 8.12. Pausing a virtual machine

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a running virtual machine in the results list.
2. Click the **Suspend** () button.  
Alternatively, right-click the virtual machine and select **Suspend**.

#### Result

The **Status** of the virtual machine changes to **Paused**.

[Report a bug](#)

## 8.7. Managing Virtual Machines

### 8.7.1. Editing a Resource

#### Summary

Edit the properties of a resource.

#### Procedure 8.13. Editing a Resource

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

#### Result

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

#### See Also:

- ▮ [Section 11.5, "Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows"](#)

[Report a bug](#)

### 8.7.2. Removing a Virtual Machine

#### Summary

When you no longer require a virtual machine, remove it from the data center. Shut down the virtual machine before removing it.

#### Procedure 8.14. Removing a virtual machine

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine in the results list.
2. Shut down the virtual machine. The **Remove** button is only enabled for a virtual machine that has a status of **Down**.

3. Click **Remove**. On the **Remove Virtual Machine(s)** confirmation window, the **Remove Disk(s)** checkbox is automatically selected, which will remove the attached virtual disks together with the virtual machine. If the checkbox is de-selected, the virtual disks will remain in the environment as floating disks.
4. Click **OK** to remove the virtual machine (and associated virtual disks) and close the window.

#### Result

The virtual machine is removed from the environment and no longer displays on the **Virtual Machines** resource tab.

[Report a bug](#)

### 8.7.3. Adding and Editing Virtual Machine Disks

#### Prerequisites:

- ▮ [Section 8.4.1, "Creating a New Virtual Machine from an Existing Template"](#)

#### Summary

It is possible to add disks to virtual machines. You can add new disks, or previously created floating disks to a virtual machine. This allows you to provide additional space to and share disks between virtual machines. You can also edit disks to change some of their details.

An **Internal** disk is the default type of disk. You can also add an **External(Direct Lun)** disk. Internal disk creation is managed entirely by the Manager; external disks require externally prepared targets that already exist. Existing disks are either floating disks or shareable disks attached to virtual machines.

#### Procedure 8.15. Adding Disks to Virtual Machines

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine in the results list.
2. Click the **Disks** tab in the details pane to display a list of virtual disks currently associated with the virtual machine.
3. Click **Add** to open the **Add Virtual Disk** window.

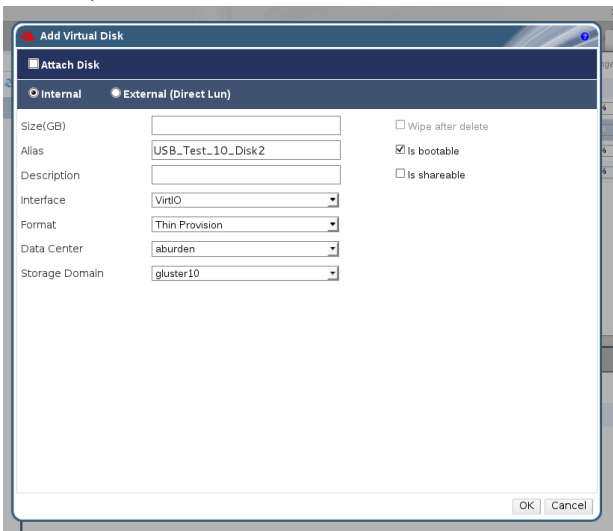


Figure 8.7. Add Virtual Disk Window

4. Use the appropriate radio buttons to switch between **Internal** and the **External (Direct Lun)** disks.
5. Select the **Attach Disk** check box to choose an existing disk from the list and select the **Activate** check box.  
Alternatively, enter the **Size**, **Alias**, and **Description** of a new disk and use the drop-down menus and check boxes to configure the disk.
6. Click **OK** to add the disk and close the window.

#### Result

Your new disk is listed in the **Virtual Disks** tab in the details pane of the virtual machine.

#### See Also:

- ▮ [Section 11.5, "Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows"](#)

[Report a bug](#)

### 8.7.4. Adding and Editing Virtual Machine Network Interfaces

#### Summary

You can add network interfaces to virtual machines. Doing so allows you to put your virtual machine on multiple logical networks. You can also edit a virtual machine's network interface to change some network interface details. This procedure can be performed on virtual machines that are running, but some actions can be performed only on virtual machines that are not running.

#### Procedure 8.16. Adding network interfaces to virtual machines

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine in the results list.
2. Select the **Network Interfaces** tab in the details pane to display a list of network interfaces that are currently associated with the virtual machine.
3. Click **New** to open the **New Network Interface** window.

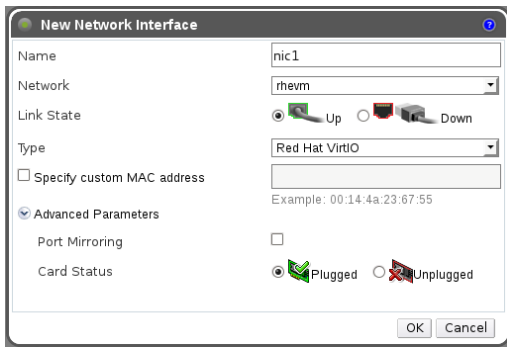


Figure 8.8. New Network Interface window

4. Enter the **Name** of the network interface.
5. Use the drop-down menus to select the **Network** and the **Type** of network interface for the new network interface. The **Link State** is set to **Up** by default when the NIC is defined on the virtual machine and connected to the network.

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.

6. If applicable, select the **Specify custom MAC address** check box and enter the network interface's MAC address.
7. Click the arrow next to **Advanced Parameters** to configure the **Port Mirroring** and **Card Status** fields, if necessary.
8. Click **OK** to close the **New Network Interface** window.

#### Result

Your new network interface is listed in the **Network Interfaces** tab in the details pane of the virtual machine.

[Report a bug](#)

#### 8.7.5. Explanation of Settings in the Virtual Machine Network Interface Window

These settings apply when you are adding or editing a virtual machine network interface. If you have more than one network interface attached to a virtual machine, you can put the virtual machine on more than one logical network.

Table 8.12. Add a network interface to a virtual machine entries

Field Name	Description
<b>Name</b>	The name of the network interface. This text field has a 21-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
<b>Network</b>	Logical network that the network interface is placed on. By default, all network interfaces are put on the <b>rhev</b> management network.
<b>Link State</b>	Whether or not the network interface is connected to the logical network. <ul style="list-style-type: none"> <li>» <b>Up:</b> The network interface is located on its slot. <ul style="list-style-type: none"> <li>» When the <b>Card Status</b> is <b>Plugged</b>, it means the network interface is connected to a network cable, and is active.</li> <li>» When the <b>Card Status</b> is <b>Unplugged</b>, the network interface will be automatically connected to the network and become active.</li> </ul> </li> <li>» <b>Down:</b> The network interface is located on its slot, but it is not connected to any network. Virtual machines will not be able to run in this state.</li> </ul>
<b>Type</b>	The virtual interface the network interface presents to virtual machines. VirtIO is faster but requires VirtIO drivers. Red Hat Enterprise Linux 5 and higher includes VirtIO drivers. Windows does not include VirtIO drivers, but they can be installed from the guest tools ISO or virtual floppy disk. rtl8139 and e1000 device drivers are included in most operating systems.
<b>Specify custom MAC address</b>	Choose this option to set a custom MAC address. The Red Hat Enterprise Virtualization Manager automatically generates a MAC address that is unique to the environment to identify the network interface. Having two devices with the same MAC address online in the same network causes networking conflicts.
<b>Port Mirroring</b>	A security feature that allows all network traffic going to or leaving from virtual machines on a given logical network and host to be copied (mirrored) to the network interface. If the host also uses the network, then traffic going to or leaving from the host is also copied.  Port mirroring only works on network interfaces with IPv4 IP addresses.
<b>Card Status</b>	Whether or not the network interface is defined on the virtual machine. <ul style="list-style-type: none"> <li>» <b>Plugged:</b> The network interface has been defined on the virtual machine. <ul style="list-style-type: none"> <li>» If its <b>Link State</b> is <b>Up</b>, it means the network interface is connected to a network cable, and is active.</li> <li>» If its <b>Link State</b> is <b>Down</b>, the network interface is not connected to a network cable.</li> </ul> </li> <li>» <b>Unplugged:</b> The network interface is only defined on the Manager, and is not associated with a virtual machine. <ul style="list-style-type: none"> <li>» If its <b>Link State</b> is <b>Up</b>, when the network interface is plugged it will automatically be connected to a network and become active.</li> <li>» If its <b>Link State</b> is <b>Down</b>, the network interface is not connected to any network until it is defined on a virtual machine.</li> </ul> </li> </ul>

[Report a bug](#)

### 8.7.6. Hot Plugging Virtual Machine Disks

#### Summary

You can hot plug virtual machine disks. Hot plugging means enabling or disabling devices while a virtual machine is running.

#### Procedure 8.17. Hot plugging virtual machine disks

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a running virtual machine in the results list.
2. Select the **Disks** tab from the details pane of the virtual machine.
3. Select the virtual machine disk you would like to hot plug.
4. Click the **Activate** or **Deactivate** button.

#### Result

You have enabled or disabled a virtual machine disk.

[Report a bug](#)

### 8.7.7. Hot Plugging Network Interfaces

#### Summary

You can hot plug network interfaces. Hot plugging means enabling and disabling network interfaces while a virtual machine is running.

#### Procedure 8.18. Hot plugging network interfaces

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a

running virtual machine in the results list.

2. Select the **Network Interfaces** tab from the details pane of the virtual machine.
3. Select the network interface you would like to hot plug and click **Edit** to open the **Edit Network Interface** window.
4. Click the **Advanced Parameters** arrow to access the **Card Status** option. Set the **Card Status** to **Plugged** if you want to enable the network interface, or set it to **Unplugged** if you want to disable the network interface.

#### Result

You have enabled or disabled a virtual network interface.

#### See Also:

- ▮ [Section 8.7.5, "Explanation of Settings in the Virtual Machine Network Interface Window"](#)

[Report a bug](#)

## 8.7.8. Removing Disks and Network Interfaces from Virtual Machines

### Summary

You can remove network interfaces and virtual hard disks from virtual machines. If you remove a disk from a virtual machine, the contents of the disk are permanently lost.

This procedure is not the same as hot plugging. You can only remove virtual hardware that is **Deactivated**.

#### Procedure 8.19. Removing disks and network interfaces from virtual machines

1. Select the virtual machine with virtual hardware you'd like to remove.
2. Select the relevant tab, either **Network Interfaces** or **Disks**, from the virtual machine details pane.
3. Select the disk or network interface you'd like to remove. To remove it, you must have first **Deactivated** it.
4. Click the **Remove** button. Click **OK** in the confirmation window. If you are removing a disk, select the **Remove Permanently** option to completely remove it from the environment. If you don't select this option, for example because the disk is a shared disk, it will remain in the **Disks** resource tab.

#### Result

The disk or network interface is no longer attached to the virtual machine.

[Report a bug](#)

## 8.8. Virtual Machines and Permissions

### 8.8.1. Managing System Permissions for a Virtual Machine

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **ClusterAdmin** has administrator privileges only for the assigned cluster, and so forth.

A **UserVmManager** is a system administration role for virtual machines in a data center. This role can be applied to specific virtual machines, to a data center, or to the whole virtualized environment; this is useful to allow different users to manage certain virtual resources.

The user virtual machine administrator role permits the following actions:

- ▮ Create, edit, and remove virtual machines; and
- ▮ Run, suspend, shutdown, and stop virtual machines.

#### Note

You can only assign roles and permissions to existing users.

Many end-users are concerned solely with the virtual machine resources of the virtualized environment. As a result, Red Hat Enterprise Virtualization provides several user roles which enable the user to manage virtual machines specifically, but not other resources in the data center.

#### See Also:

- ▮ [Section 8.8.2, "Virtual Machines Administrator Roles Explained"](#)

[Report a bug](#)

### 8.8.2. Virtual Machines Administrator Roles Explained

#### Virtual Machine Administrator Permission Roles

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

Table 8.13. Red Hat Enterprise Virtualization System Administrator Roles

Role	Privileges	Notes
DataCenterAdmin	Data Center Administrator	Can use, create, delete, manage all virtual machines within a specific data center.
ClusterAdmin	Cluster Administrator	Can use, create, delete, manage all virtual machines within a specific cluster.
NetworkAdmin	Network Administrator	Can configure and manage networks attached to virtual machines. To configure port mirroring on a virtual machine network, apply the <b>NetworkAdmin</b> role on the network and the <b>UserVmManager</b> role on the virtual machine.

#### See Also:

- ▮ [Section 12.4.2, "User and Administrator Roles"](#)



[Report a bug](#)

### 8.8.3. Virtual Machine User Roles Explained

#### Virtual Machine User Permission Roles

The table below describes the user roles and privileges applicable to virtual machine users. These roles allow access to the User Portal for managing and accessing virtual machines, but they do not confer any permissions for the Administration Portal.

Table 8.14. Red Hat Enterprise Virtualization System User Roles

Role	Privileges	Notes
UserRole	Can access and use virtual machines and pools.	Can log in to the User Portal and use virtual machines and pools.
PowerUserRole	Can create and manage virtual machines and templates.	Apply this role to a user for the whole environment with the <b>Configure</b> window, or for specific data centers or clusters. For example, if a PowerUserRole is applied on a data center level, the PowerUser can create virtual machines and templates in the data center.
UserVmManager	System administrator of a virtual machine.	Can manage virtual machines, create and use snapshots, and migrate virtual machines. A user who creates a virtual machine in the User Portal is automatically assigned the UserVmManager role on the machine.
UserTemplateBasedVm	Limited privileges to only use Templates.	Level of privilege to create a virtual machine by means of a template.
VmCreator	Can create virtual machines in the User Portal.	This role is not applied to a specific virtual machine; apply this role to a user for the whole environment with the <b>Configure</b> window. When applying this role to a cluster, you must also apply the <b>DiskCreator</b> role on an entire data center, or on specific storage domains.
NetworkUser	Logical network and network interface user for virtual machines.	If the <b>Allow all users to use this Network</b> option was selected when a logical network is created, <b>NetworkUser</b> permissions are assigned to all users for the logical network. Users can then attach or detach virtual machine network interfaces to or from the logical network.



#### Note

In Red Hat Enterprise Virtualization 3.0, the **PowerUserRole** only granted permissions for virtual machines which are directly assigned to the PowerUser, or virtual machines created by the PowerUser. Now, the **VmCreator** role provides privileges previously conferred by the **PowerUserRole**. The **PowerUserRole** can now be applied on a system-wide level, or on specific data centers or clusters, and grants permissions to all virtual machines and templates within the system or specific resource. Having a **PowerUserRole** is equivalent to having the **VmCreator**, **DiskCreator**, and **TemplateCreator** roles.

[Report a bug](#)

### 8.8.4. Assigning an Administrator or User Role to a Resource

#### Summary

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

#### Procedure 8.20. Assigning a Role to a Resource

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

#### Result

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

#### See Also:

- » [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

### 8.8.5. Removing an Administrator or User Role from a Resource

#### Summary

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

#### Procedure 8.21. Removing a Role from a Resource

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

**Result**

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

[Report a bug](#)

## 8.9. Backing Up and Restoring Virtual Machines with Snapshots

### 8.9.1. Creating a Snapshot of a Virtual Machine

**Summary**

A snapshot is a view of a virtual machine's operating system and applications at a given point in time. Take a snapshot of a virtual machine before you make a change to it that may have unintended consequences. You can use a snapshot to return a virtual machine to a previous state.

**Note**

Live snapshots can only be created for virtual machines running on 3.1-or-higher-compatible data centers. Virtual machines in 3.0-or-lower-compatible data centers must be shut down before a snapshot can be created.

**Procedure 8.22. Creating a snapshot of a virtual machine**

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine in the results list.
2. Select the **Snapshots** tab in the details pane.
3. Click **Create** in the left side-pane of the details pane to open the **Create Snapshot** window.

General   Network Interfaces   Disks   Snapshots   Applications   Permissions   Sessions		
Create   Preview   Commit   Undo   Delete   Clone		
Date	Status	Description
Current	Ok	Active VM
2013-Mar-25, 11:06	Ok	new_snapshot
2013-Mar-25, 11:03	Ok	fl7_installed
2013-Mar-25, 10:12	Ok	snapshot1

Figure 8.9. The Virtual Machines Details Pane with Snapshots tab

4. Enter a description for the snapshot.
5. Click **OK** to create the snapshot and close the window.

General   Network Interfaces   Disks   Snapshots   Applications   Permissions		
Create   Preview   Commit   Undo   Delete   Clone		
Date	Status	Description
Current	Ok	Active VM
2012-Oct-03, 15:17:04	Ok	Weekly backup

Figure 8.10. New Snapshot in the Details Pane

**Result**

The virtual machine's operating system and applications are stored in a snapshot that can be previewed or restored. The snapshot is created with a status of **Locked**, which changes to **Ok**. When you click on the snapshot, its details are shown on the **General**, **Disks**, **Network Interfaces**, and **Installed Applications** tabs in the right side-pane of the details pane.

[Report a bug](#)

### 8.9.2. Using a Snapshot to Restore a Virtual Machine

**Prerequisites:**

- » [Section 8.9.1, "Creating a Snapshot of a Virtual Machine"](#)

**Summary**

A snapshot can be used to restore a virtual machine to its previous state.

**Procedure 8.23. Using a snapshot to restore a virtual machine**

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select the virtual machine in the results list. Ensure the status is **Powered Down**.
2. Click the **Snapshots** tab in the details pane to list the available snapshots.
3. Select a snapshot to restore in the left side-pane. The snapshot details display in the right side-pane.
4. Click **Preview** to preview the snapshot. The status of the virtual machine briefly changes to **Image Locked** before returning to **Down**.

General	Network Interfaces	Disks	Snapshots	Applications	Permissions	Sessions
Create Preview Commit Undo Delete Clone						
Date	Status	Description				
Current	Ok	Active VM				
2013-Mar-25, 11:08	Ok	new_snapshot				
2013-Mar-25, 11:03	Ok	fl7_installed				
2013-Mar-25, 10:12	Ok	snapshot1				

Figure 8.11. Preview snapshot

5. Start the virtual machine and it will run with the disk image of the snapshot.
6. Click **Commit** to permanently restore the virtual machine to the condition of the snapshot. Any subsequent snapshots are erased.  
Alternatively, click the **Undo** button to deactivate the snapshot and return the virtual machine to its previous state.

**Result**

The virtual machine is restored to its state at the time of the snapshot, or returned to its state before the preview of the snapshot.

[Report a bug](#)

**8.9.3. Creating a Virtual Machine from a Snapshot****Prerequisites:**

- » [Section 8.9.1, "Creating a Snapshot of a Virtual Machine"](#)

**Summary**

You have created a snapshot from a virtual machine. Now you can use that snapshot to create another virtual machine.

**Procedure 8.24. Creating a virtual machine from a snapshot**

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select the virtual machine in the results list.
2. Click the **Snapshots** tab in the details pane to list the available snapshots for the virtual machines.
3. Select a snapshot in the list displayed and click **Clone** to open the **Clone VM from Snapshot** window.
4. Enter the **Name** and **Description** of the virtual machine to be created.

Figure 8.12. Clone a Virtual Machine from a Snapshot

5. Click **OK** to create the virtual machine and close the window.

**Result**

After a short time, the cloned virtual machine appears in the **Virtual Machines** tab in the navigation pane. It appears in the navigation pane with a status of **Image Locked**. The virtual machine will remain in this state until Red Hat Enterprise Virtualization completes the creation of the virtual machine. A virtual machine with a preallocated 20GB hard drive takes about fifteen minutes to create. Sparsely-allocated virtual disks take less time to create than do preallocated virtual disks.

When the virtual machine is ready to use, its status changes from **Image Locked** to **Down** in the **Virtual Machines** tab in the navigation pane.

[Report a bug](#)

**8.9.4. Deleting a Snapshot****Summary**

Delete a snapshot and permanently remove it from the virtualized environment.

**Procedure 8.25. Deleting a Snapshot**

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine in the results list.
2. Click the **Snapshots** tab in the details pane to list available snapshots for the virtual machine.

General	Network Interfaces	Disks	Snapshots	Applications	Permissions	Sessions
Create	Preview	Commit	Undo	Delete	Clone	
Date	Status	Description				
Current	Ok	Active VM				
2013-Mar-25, 11:08	Ok	new_snapshot				
2013-Mar-25, 11:03	Ok	fl7_installed				
2013-Mar-25, 10:12	Ok	snapshot1				

Figure 8.13. Snapshot List

3. Select the snapshot to delete.
4. In the Navigation pane, shut down the running virtual machine associated with the snapshot to be deleted.
5. Click **Delete** to open the **Delete Snapshot** confirmation window.
6. Click **OK** to delete the snapshot and close the window.

**Result**

You have removed a virtual machine snapshot. Removing a snapshot does not affect the virtual machine.

[Report a bug](#)

## 8.10. Importing and Exporting Virtual Machines

### 8.10.1. Exporting and Importing Virtual Machines

A virtual machine or a template can be moved between data centers in the same environment, or to a different Red Hat Enterprise Virtualization environment. The 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:

- Moving virtual resources between Red Hat Enterprise Virtualization environments.
- Move virtual machines and templates between data centers in a single Red Hat Enterprise Virtualization environment.
- Backing up virtual machines and templates.

There are three stages of exporting and importing virtual resources:

- First you export your virtual machines and templates to an export domain.
- Second, you detach the export domain from one data center, and attach it to another. You can attach it to a different data center in the same Red Hat Enterprise Virtualization environment, or attach it to a data center in a separate Red Hat Enterprise Virtualization environment that is managed by another installation of the Red Hat Enterprise Virtualization Manager.
- Third, you import your virtual machines and template into the data center you attached the export domain to.

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 must exist in the destination data center for the virtual machine to work, or the virtual machine must be exported with the **Collapse Snapshots** option selected.

[Report a bug](#)

### 8.10.2. Overview of the Export-Import Process

The export domain allows you to move virtual machines and templates between Red Hat Enterprise Virtualization environments.

Exporting and importing resources requires that an active export domain be attached to the data center. An export domain is a temporary storage area containing two directories per exported virtual resource. One directory consists of all the OVF (Open Virtualization Format) files pertaining to the virtual machine. The other holds the virtual resource's disk image, or images.

You can also import virtual machines from other virtualization providers, for example, Xen, VMware or Windows virtual machines, using the V2V feature. 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

An export domain can be active in only one data center. This means that the export domain can be attached to either the source data center or the destination data center.

Exporting virtual resources across data centers requires some preparation. Make sure that:

- an export domain exists, and is attached to the source data center.
- the virtual machine is shut down.
- if the virtual machine was created from a template, the template resides on the destination data center, or is exported alongside the virtual machine.

When the virtual machine, or machines, have been exported to the export domain, you can import them into the destination data center. If the destination data center is within the same Red Hat Enterprise Virtualization environment, delete the originals from the source data center after exporting them to the export domain.

[Report a bug](#)

### 8.10.3. Performing an Export-Import of Virtual Resources

#### Summary

This procedure provides a graphical overview of the steps required to import a virtual resource to its destination.

#### Procedure 8.26. Performing an export-import of virtual resources

1. Attach the export domain to the source data center.




Figure 8.14. Attach Export Domain

2. Export the virtual resource to the export domain.




Figure 8.15. Export the Virtual Resource

3. Detach the export domain from the source data center.

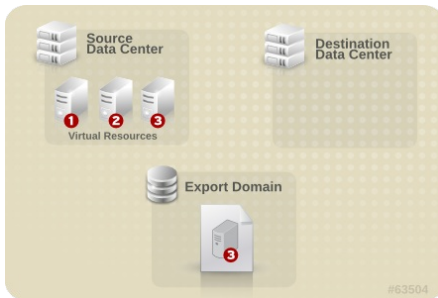


Figure 8.16. Detach Export Domain

4. Attach the export domain to the destination Data center.

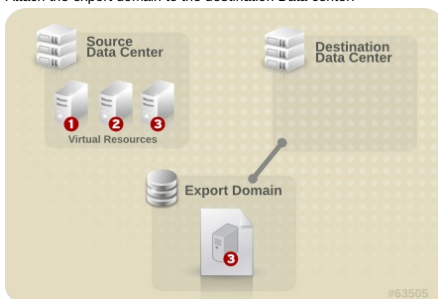


Figure 8.17. Attach the Export Domain

5. Import the virtual resource into the destination data center.

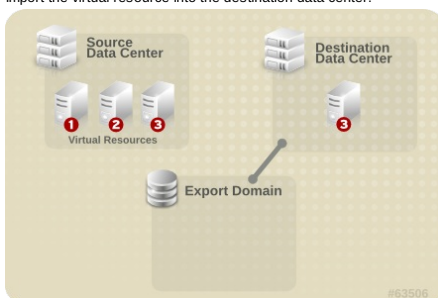


Figure 8.18. Import the virtual resource

#### Result

Your virtual resource is exported to the destination data center.

**See Also:**

- » [Section 3.6.3, "Attaching an Existing Export Domain to a Data Center"](#)
- » [Section 3.6.4, "Detaching a Storage Domain from a Data Center"](#)
- » [Section 8.10.4, "Exporting a Virtual Machine to the Export Domain"](#)
- » [Section 8.10.5, "Importing a Virtual Machine into the Destination Data Center"](#)

[Report a bug](#)

### 8.10.4. Exporting a Virtual Machine to the Export Domain

#### Summary

Export a virtual machine to the export domain so that it can be imported into a different data center. Before you begin, the export domain must be attached to the data center that contains the virtual machine to be exported.

#### Procedure 8.27. Exporting a Virtual Machine to the Export Domain

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine in the results list. Ensure the virtual machine has a status of **Down**.
2. Click **Export** to open the **Export Virtual Machine** window.
3. Select the **Force Override** check box to override existing images of the virtual machine on the export domain.  
Select the **Collapse Snapshots** check box to create a single export volume per disk. Selecting this option will remove snapshot restore points and include the template in a template-based virtual machine. This removes any dependencies a virtual machine has on a template.
4. Click **OK** to export the virtual machine and close the window.

#### Result

The export of the virtual machine begins. The virtual machine displays in the Virtual Machines list with an **Image Locked** status as it is exported. Depending on the size of your virtual machine hard disk images, and your storage hardware, this can take up to an hour. Use the **Events** tab to view the progress.

When complete, the virtual machine has been exported to the export domain and displays on the **VM Import** tab of the export domain's details pane.

[Report a bug](#)

### 8.10.5. Importing a Virtual Machine into the Destination Data Center

#### Summary

You have a virtual machine on an export domain. Before the virtual machine can be imported to a new data center, the export domain must be attached to the destination data center.

#### Procedure 8.28. Importing a Virtual Machine into the Destination Data Center

1. Use the **Storage** resource tab, tree mode, or the search function to find and select the export domain in the results list. The export domain must have a status of **Active**.
2. Select the **VM Import** tab in the details pane to list the available virtual machines to import.
3. Select one or more virtual machines to import and click **Import** to open the **Import Virtual Machine(s)** window.

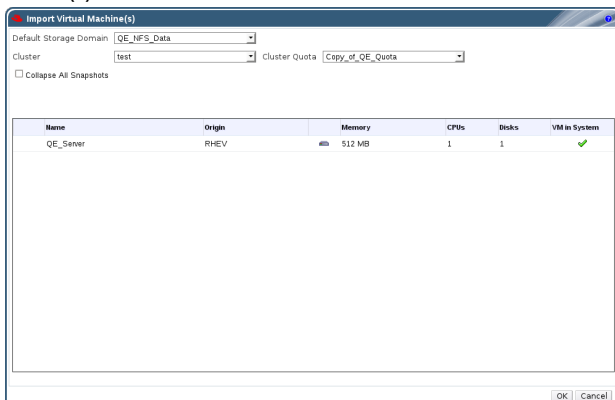


Figure 8.19. Import Virtual Machine

4. Use the drop-down menus to select the **Default Storage Domain**, **Cluster**, and **Cluster Quota** in the data center.
5. Select the **Collapse All Snapshots** check box to remove snapshot restore points and include templates in template-based virtual machines.
6. Click **OK** to import the virtual machines.

The **Import Conflict** window opens if the virtual machine exists in the virtualized environment.

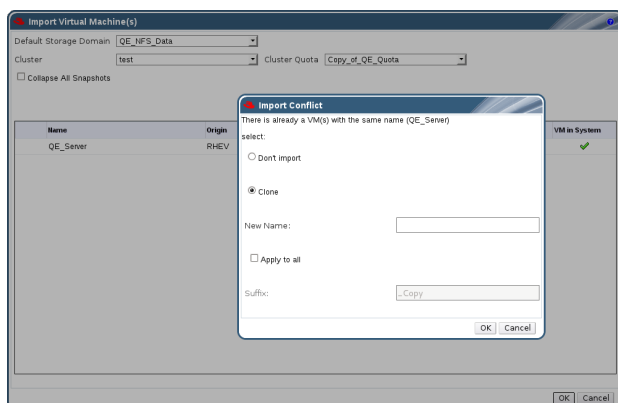


Figure 8.20. Import Conflict Window

7. Choose one of the following radio buttons:

- » **Don't import**
- » **Clone** and enter a unique name for the virtual machine in the **New Name** : field.

Or select the **Apply to all** check box to import all duplicated virtual machines with the same suffix.

8. Click **OK** to import the virtual machines and close the window.

#### Result

You have imported the virtual machine to the destination data center. This may take some time to complete.

[Report a bug](#)

## 8.11. Migrating Virtual Machines Between Hosts

### 8.11.1. 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.

[Report a bug](#)

### 8.11.2. Live Migration Prerequisites

Live migration is used to seamlessly move virtual machines to support a number of common maintenance tasks. Ensure that your Red Hat Enterprise Virtualization environment is correctly configured to support live migration well in advance of using 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.
- » The migrating virtual machine must not have the **cache!=none** custom property set.

In addition, for best performance, the storage and management networks should be split to avoid network saturation. Virtual machine migration involves transferring large amounts of data between hosts.

Live migration is performed using the management network. Each live migration event is limited to a maximum transfer speed of 30 MBps, and the number of concurrent migrations supported is also limited by default. Despite these measures, concurrent migrations have the potential to saturate the management network. It is recommended that separate logical networks are created for storage, display, and virtual machine data to minimize the risk of network saturation.

[Report a bug](#)

### 8.11.3. Automatic Virtual Machine Migration

Red Hat Enterprise Virtualization Manager automatically initiates live migration of all virtual machines running on a host when the host is moved into maintenance mode. The destination host for each virtual machine is assessed as the virtual machine is migrated, in order to spread the load across the cluster.

The Manager automatically initiates live migration of virtual machines in order to maintain load balancing or power saving levels in line with cluster policy. While no cluster policy is defined by default, it is recommended that you specify the cluster policy which best suits the needs of your environment. You can also disable automatic, or even manual, live migration of specific virtual machines where required.

[Report a bug](#)

### 8.11.4. Preventing Automatic Migration of a Virtual Machine

#### Summary

Red Hat Enterprise Virtualization Manager allows you to disable automatic migration of virtual machines. You can also disable manual migration of virtual machines by setting the virtual machine to run only on a specific host.

The ability to disable automatic migration and require a virtual machine to run on a particular host is useful when using application high availability products, such as Red Hat High Availability or Cluster Suite.

#### Procedure 8.29. Preventing automatic migration of a virtual machine

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine or virtual server in the results list.
2. Click **Edit** to open the **Edit Desktop Virtual Machine** or **Edit Server Virtual Machine** window.

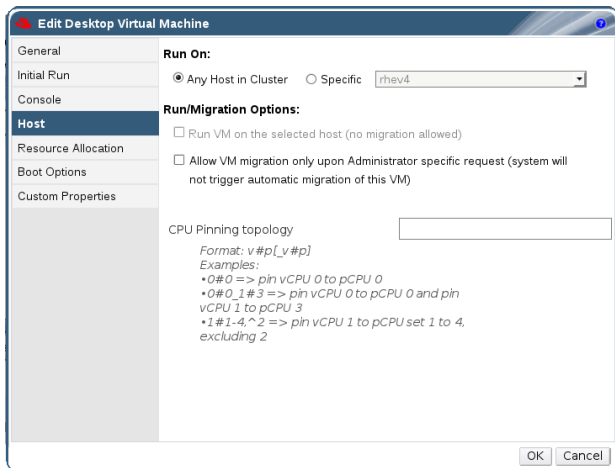


Figure 8.21. Edit Desktop Virtual Machine Window

3. Click the **Host** tab.
4. Use the **Run On** radio buttons to designate the virtual machine to run on **Any Host in Cluster** or a **Specific** host. If applicable, select a specific host from the drop-down menu.

**Warning**

Explicitly assigning a virtual machine to a specific host and disabling migration is mutually exclusive with Red Hat Enterprise Virtualization high availability. Virtual machines that are assigned to a specific host can only be made highly available using third party high availability products like Red Hat High Availability.

5. Use the drop-down menu to affect the **Migration Options**. Select **Do not allow migration** to enable the **Use Host CPU** check box.
6. If applicable, enter relevant **CPU Pinning topology** commands in the text field.
7. Click **OK** to save the changes and close the window.

**Result**

You have changed the migration settings for the virtual machine.

**See Also:**

» [Section 4.2.2.3, "Resilience Policy Settings Explained"](#)

[Report a bug](#)

**8.11.5. Manually Migrating Virtual Machines****Summary**

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. When bringing a server down for maintenance, migration is triggered automatically, so manual migration is not required. Migration of virtual machines does not cause any service interruption.

The migrating virtual machine must not have the **cache!=none** custom property set.

**Procedure 8.30. Manually migrating virtual machines**

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a running virtual machine in the results list.  
Click **Migrate** to open the **Migrate Virtual Machine(s)** window.
2. Use the radio buttons to select whether to **Select Host Automatically** or to **Select Destination Host**, specifying the host using the drop-down menu.

**Note**

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.

3. Click **OK** to commence migration and close the window.

**Result**

The virtual machine is migrated. Once migration is complete the **Host** column will update to display the host the virtual machine has been migrated to.

[Report a bug](#)

**8.11.6. Setting Migration Priority****Summary**

Red Hat Enterprise Virtualization Manager queues concurrent requests for migration of virtual machines off of 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.

#### Note

You can only set the migration priority for virtual servers. You can not set migration priority for virtual desktops.

#### Procedure 8.31. Setting Migration Priority

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual server in the results list.
2. Click **Edit** to open the **Edit Server Virtual Machine** window.
3. Select the **High Availability** tab.
4. Use the radio buttons to set the **Priority for Run/Migrate Queue** of the virtual machine to one of **Low**, **Medium**, or **High**.
5. Click **OK** to save changes and close the window.

#### Result

The virtual machine's migration priority has been modified.

[Report a bug](#)

### 8.11.7. Cancelling ongoing virtual machine migrations

#### Summary

A virtual machine migration is taking longer than you expected. You'd like to be sure where all virtual machines are running before you make any changes to your environment.

#### Procedure 8.32. Cancelling ongoing virtual machine migrations

1. Select the migrating virtual machine. It is displayed in the **Virtual Machines** resource tab with a status of **Migrating from**.
2. Click the **Cancel Migration** button at the top of the results list. Alternatively, right-click on the virtual machine and select **Cancel Migration** from the context menu.

#### Result

The virtual machine status returns from **Migrating from** status to **Up** status.

[Report a bug](#)

### 8.11.8. Event and Log Notification upon Automatic Migration of Highly Available Virtual Servers

When a virtual server is automatically migrated because of the high availability function, the details of an automatic migration are documented in the **Events** tab and in the engine log to aid in troubleshooting, as illustrated in the following examples:

#### Example 8.1. Notification in the Events Tab of the Web Admin Portal

Highly Available *Virtual\_Machine\_Name* failed. It will be restarted automatically.

*Virtual\_Machine\_Name* was restarted on Host *Host\_Name*

#### Example 8.2. Notification in the Manager engine.log

This log can be found on the Red Hat Enterprise Virtualization Manager at `/var/log/ovirt-engine/engine.log`:

Failed to start Highly Available VM. Attempting to restart. VM Name: *Virtual\_Machine\_Name*, VM Id: *Virtual\_Machine\_ID\_Number*

[Report a bug](#)

## 8.12. Improving Uptime with Virtual Machine High Availability

### 8.12.1. Why Use High Availability?

High availability is recommended for virtual machines running critical workloads.

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.

[Report a bug](#)

### 8.12.2. What is High Availability?

High availability means that a virtual machine will be automatically restarted if its process is interrupted. This happens if the virtual machine is terminated by methods other than powering off from within the guest or sending the shutdown command from the Manager. When these events occur, the highly available virtual machine is 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 hosts and storage, 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.

With high availability, interruption to service is minimal because virtual machines are restarted within seconds with no user intervention required. High availability keeps your resources balanced by restarting guests on a host with low 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.

[Report a bug](#)

### 8.12.3. 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 running the highly available virtual machine 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 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 CPUs on the destination host that are not in use 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.

[Report a bug](#)

#### 8.12.4. Configuring a Highly Available Virtual Machine

##### Summary

High availability must be configured individually for each virtual server.

##### Note

You can only set high availability for virtual servers. You can not set high availability for virtual desktops.

##### Procedure 8.33. Configuring a Highly Available Virtual Machine

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual server in the results list.
2. Click **Edit** to open the **Edit Server Virtual Machine** window.
3. Click the **High Availability** tab.

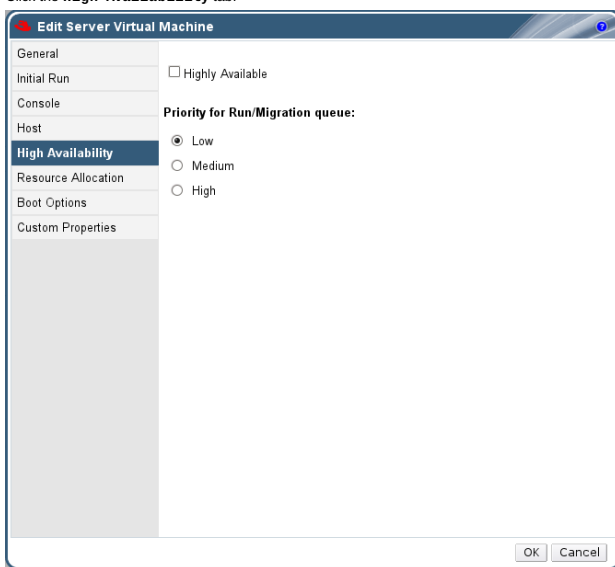


Figure 8.22. Set virtual machine high availability

4. Select the **Highly Available** check box to enable high availability for the virtual server.
5. Use the radio buttons to set the **Priority for Run/Migrate Queue** of the virtual machine to one of **Low**, **Medium**, or **High**. 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.
6. Click **OK** to save changes and close the window.

##### Result

You have configured high availability for a virtual machine. You can check if a virtual machine is highly available when you select it and click on its **General** tab in the details pane.

##### See Also:

- » [Section 8.4.3.5, "Virtual Machine High Availability Settings Explained"](#)
- » [Section 4.2.2.3, "Resilience Policy Settings Explained"](#)

[Report a bug](#)

## 8.13. Other Virtual Machine Tasks

### 8.13.1. Enabling SAP monitoring for a virtual machine from the Administration Portal

##### Summary

Enable SAP monitoring on a virtual machine to be recognized by SAP monitoring systems.

##### Procedure 8.34. Enabling SAP monitoring for a Virtual Machine from the Administration Portal

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select a virtual machine with a status of **Down** in the results list.
2. Click **Edit** button to open the **Edit Virtual Machine** window.
3. Select the **Custom Properties** tab.




Figure 8.23. Enable SAP

4. Use the drop-down menu to select **sap\_agent**. Ensure the secondary drop-down menu is set to **True**.  
If previous properties have been set, select the plus sign to add a new property rule and select **sap\_agent**.
5. Click **OK** to save changes and close the window.

**Result**

You have enabled SAP monitoring for your virtual machine.

[Report a bug](#)

### 8.13.2. Configuring Red Hat Enterprise Linux 5.4 or Higher Virtual Machines to use SPICE

#### 8.13.2.1. Using SPICE on virtual machines running versions of Red Hat Enterprise Linux released prior to 5.4

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 newer. Older versions are not supported. Installing SPICE on a virtual machine running Red Hat Enterprise Linux significantly improves the performance of the graphical user interface.

#### Note

Typically, this is most useful for virtual machines where the user requires the use of the graphical user interface. System administrators who are creating virtual servers may prefer not to configure SPICE if their use of the graphical user interface is minimal.

[Report a bug](#)

#### 8.13.2.2. Installing qxl drivers on virtual machines

**Summary**

This procedure installs qxl drivers on virtual machines running Red Hat Enterprise Linux 5.4 or higher.

**Procedure 8.35. Installing qxl drivers on a virtual machine**

1. Log in to a Red Hat Enterprise Linux virtual machine.
2. Open a terminal.
3. Run the following command as root:

```
# yum install xorg-x11-drv-qxl
```

**Result**

The qxl drivers have been installed and must now be configured.

[Report a bug](#)

#### 8.13.2.3. Configuring qxl drivers on virtual machines

**Summary**

You can configure qxl drivers using either a graphical interface or the command line. Perform only one of the following procedures.

**Procedure 8.36. Configuring qxl drivers in GNOME**

1. Click **System**.
2. Click **Administration**.
3. Click **Display**.
4. Click the **Hardware** tab.
5. Click **Video Cards Configure**.
6. Select **qxl** and click **OK**.
7. Restart X-Windows by logging out of the virtual machine and logging back in.

**Procedure 8.37. Configuring 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
```

#### Result

You have configured qxl drivers to enable your virtual machine to use SPICE.

[Report a bug](#)

#### 8.13.2.4. Configuring a virtual machine's tablet and mouse to use SPICE

##### Summary

Edit the `/etc/X11/xorg.conf` file to enable SPICE for your virtual machine's tablet devices.

##### Procedure 8.38. Configuring a virtual machine's tablet and mouse to use SPICE

1. Verify that the tablet device is available on your guest:

```
# /sbin/lshusb -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.

#### Result

You have enabled tablet devices on your virtual machine to use SPICE.

[Report a bug](#)

#### 8.13.3. 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. The NTP servers you are trying to use must be operational and accessible to your hosts and virtual machines.

#### 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](#)). 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 `rhevm-config` tool to recive alerts when hosts drift out of sync.


You can use the `rhevm-config` tool to configure alerts when your hosts drift out of sync.

There are 2 relevant parameters for time drift on hosts: `EnableHostTimeDrift` and `HostTimeDriftInSec`. `EnableHostTimeDrift`, with a default value of false, can be enabled to recieve alert notifications of host time drift. The `HostTimeDriftInSec` parameter is used to set the maximum allowable drift before alerts start being sent.

Alerts are sent once per hour per host.

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.

Red Hat Enterprise Linux	Additional virtual machine kernel parameters
5.4 AMD64/Intel 64 with the para-virtualized clock	Additional parameters are not required
5.4 AMD64/Intel 64 without the para-virtualized clock	notsc lpj=n
5.4 x86 with the para-virtualized clock	Additional parameters are not required
5.4 x86 without the para-virtualized clock	clocksource=acpi_pm lpj=n
5.3 AMD64/Intel 64	notsc
5.3 x86	clocksource=acpi_pm
4.8 AMD64/Intel 64	notsc
4.8 x86	clock=pmtmr
3.9 AMD64/Intel 64	Additional parameters are not required
3.9 x86	Additional parameters are not required

Using the Real-Time Clock with Windows virtual machines

Windows uses the 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](#).

See Also:

» [Section 15.2.1, "Configuration Tool"](#)

[Report a bug](#)

## Chapter 9. Templates

### 9.1. Introduction to Templates

A template is a copy of a preconfigured virtual machine, used to simplify the subsequent, repeated creation of similar virtual machines. Templates capture installed software and software configurations, as well as the hardware configuration, of the original virtual machine.

When you create a template from a virtual machine, a read-only copy of the virtual machine's disk is taken. The read-only disk becomes the base disk image of the new template, and of any virtual machines created from the template. As such, the template cannot be deleted whilst virtual machines created from the template exist in the environment.

Virtual machines created from a template use the same NIC type and driver as the original virtual machine, but utilize separate and unique MAC addresses.

#### Note

A virtual machine may require to be sealed before being used to create a template.

[Report a bug](#)

### 9.2. Template Tasks

#### 9.2.1. Creating a Template from an Existing Virtual Machine

##### Summary

Create a template from an existing virtual machine to use as the blueprint for additional virtual machines.

##### Procedure 9.1. Creating a Template from an Existing Virtual Machine

1. Use the **Virtual Machines** resource tab, tree mode, or the search function to find and select the virtual machine in the results list.
2. Ensure that the virtual machine is powered down and has a status of **Down**.
3. Click **Make Template** to open the **New Template** window.

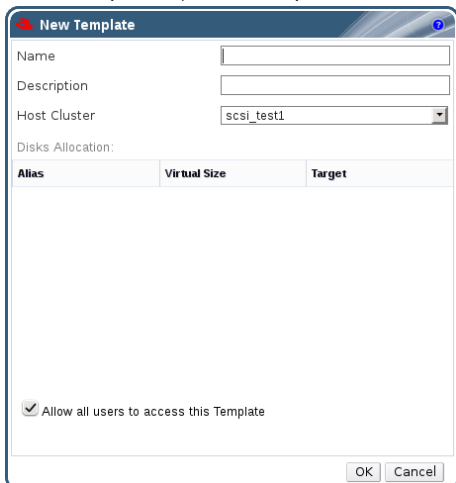


Figure 9.1. New Template Window

4. Enter the **Name** and **Description** of the template.
5. Select the **Host Cluster** and storage **Target** from the drop-down menus; by default, these are set the same as the source virtual machine.
6. The **Allow all users to access this Template** check box is selected by default, which makes it public.
7. Click **OK** to create the template. The virtual machine displays a status of **Image Locked** while the template is being created; this may take up to an hour, depending on the virtual machine disk image size and your storage hardware.

##### Result

The template is created and added to the **Templates** tab. You can now create new virtual machines from the template.

#### Note

When a template is made, the virtual machine is copied so that both the existing virtual machine and its template are usable after template creation.

##### See Also:

- » [Section 8.9.1, "Creating a Snapshot of a Virtual Machine"](#)
- » [Section 9.2.2, "Explanation of Settings and Controls in the New Template Window"](#)

[Report a bug](#)

#### 9.2.2. Explanation of Settings and Controls in the New Template Window

The table below describes the settings for the **New Template** window.

Table 9.1. New and Edit Template Properties

Field	Description/Action
<b>Name</b>	The name of the template. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
<b>Description</b>	The description of the template. This field is recommended but not mandatory.
<b>Host Cluster</b>	The cluster with which the template will be associated. This is the same as the original virtual machines by default; you can select from any cluster in the data center.
<b>Target</b>	The storage domain with which the template will be associated. This is the same as the original virtual machines by default; you can select from any storage domain in the data center.
<b>Allow all users to access the Template</b>	A public template can be accessed by all users. A private template can only be accessed by <b>TemplateAdmin</b> and the <b>SuperUser</b> .

[Report a bug](#)

### 9.2.3. Editing a Resource

#### Summary

Edit the properties of a resource.

#### Procedure 9.2. Editing a Resource

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

#### Result

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

[Report a bug](#)

### 9.2.4. Explanation of Settings and Controls in the Edit Template Window

The table below describes the settings for the **Edit Template** window.

Table 9.2. Edit Template Window

Field Name	Description
<b>Data Center</b>	The data center to which the virtual machine is attached.
<b>Host Cluster</b>	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.
<b>Name</b>	The name of virtual machine. Names must not contain any spaces, and must contain at least one character from A-Z. The maximum length of a virtual machine name is 15 characters.
<b>Description</b>	A meaningful description of the new virtual machine.
<b>Memory Size</b>	The amount of memory assigned to the virtual machine. When allocating memory, consider the processing and storage needs of the applications that are intended to run on the virtual machine. Maximum guest memory is constrained by the selected guest architecture and the cluster compatibility level.
<b>Total Virtual CPUs</b>	The processing power allocated to the virtual machine as CPU Cores. Do not assign more cores to a virtual machine than are present on the physical host.
<b>Cores per Virtual Socket</b>	The number of cores assigned to each virtual socket.
<b>Virtual Sockets</b>	The number of CPU sockets for the virtual machine. Do not assign more sockets to a virtual machine than are present on the physical host.
<b>Operating System</b>	The operating system. Valid values include a range of Red Hat Enterprise Linux and Windows variants.
<b>Stateless (Desktop only)</b>	Select this check box if the virtual machine is to run in stateless mode. The stateless mode is used primarily for desktop virtual machines. Running a stateless desktop or server creates a new COW layer on the virtual machine hard disk image where new and changed data is stored. Shutting down the stateless virtual machine deletes the new COW layer, returning the virtual machine to its original state. This type of virtual machine is useful when creating virtual machines that need to be used for a short time, or by temporary staff.

[Report a bug](#)

### 9.2.5. Deleting a Template

#### Summary

Delete a template from your Red Hat Enterprise Virtualization environment.

**Warning**

If you have used a template to create a virtual machine, make sure that you do not delete the template as the virtual machine needs it to continue running.

**Procedure 9.3. Deleting a Template**

1. Use the resource tabs, tree mode, or the search function to find and select the template in the results list.
2. Click **Remove** to open the **Remove Template(s)** window.
3. Click **OK** to remove the template.

**Result**

You have removed the template.

**See Also:**

- » [Section 11.5, "Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows"](#)

[Report a bug](#)

**9.2.6. Exporting Templates****9.2.6.1. Migrating Templates to the Export Domain****Summary**

Export templates into the export domain to move them to another data domain, either in the same Red Hat Enterprise Virtualization environment, or another one.

**Procedure 9.4. Exporting Individual Templates to the Export Domain**

1. Use the **Templates** resource tab, tree mode, or the search function to find and select the template in the results list.
2. Click **Export** to open the **Export Template** window.

**Note**

Select the **Force Override** check box to replace any earlier version of the template on the export domain.

3. Click **OK** to begin exporting the template; this may take up to an hour, depending on the virtual machine disk image size and your storage hardware.
4. Repeat these steps until the export domain contains all the templates to migrate before you start the import process.

Use the **Storage** resource tab, tree mode, or the search function to find and select the export domain in the results list and click the **Template Import** tab in the details pane to view all exported templates in the export domain.

**Result**

The templates have been exported to the export domain.

**See Also:**

- » [Section 7.7.8, "Detaching the Export Domain"](#)

[Report a bug](#)

**9.2.7. Importing Templates****9.2.7.1. Importing a Template into a Data Center****Prerequisites:**

- » [Section 9.2.7, "Importing Templates"](#)
- » [Section 7.7.9, "Attaching an Export Domain to a Data Center"](#)

**Summary**

Import templates from a newly attached export domain.

**Procedure 9.5. Importing a Template into a Data Center**

1. Use the resource tabs, tree mode, or the search function to find and select the newly attached export domain in the results list.
2. Select the **Template Import** tab of the details pane to display the templates that migrated across with the export domain.
3. Select a template and click **Import** to open the **Import Template(s)** window.

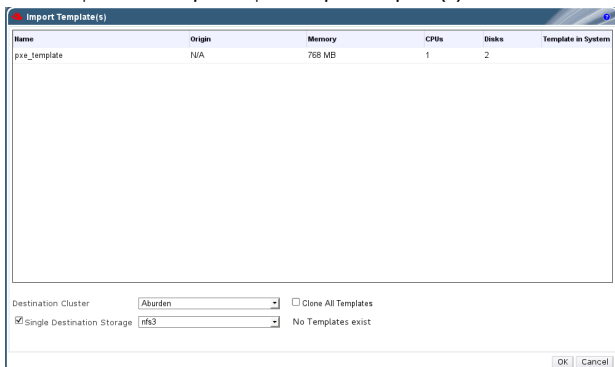


Figure 9.2. Import Templates

4. Select the templates to import.
5. Use the drop-down menus to select the **Destination Cluster** and **Storage** domain. Alter the **Suffix** if applicable.



Alternatively, clear the **Clone All Templates** check box.

- Click **OK** to import templates and open a notification window. Click **Close** to close the notification window.

#### Result

The template is imported into the destination data center. This can take up to an hour, depending on your storage hardware. You can view the import progress in the **Events** tab.

Once the importing process is complete, the templates will be visible in the **Templates** resource tab. The templates can create new virtual machines, or run existing imported virtual machines based on that template.

[Report a bug](#)

## 9.3. Sealing Templates in Preparation for Deployment

### 9.3.1. Sealing a Linux Virtual Machine for Deployment as a Template

#### Summary

Generalize (seal) a Linux virtual machine before making it into a template. This prevents conflicts between virtual machines deployed from the template.

#### Procedure 9.6. Sealing a Linux Virtual Machine

- Log in to the virtual machine. Flag the system for re-configuration by running the following command as root:

```
# touch /.unconfigured
```

- Remove ssh host keys. Run:

```
# rm -rf /etc/ssh/ssh_host_*
```

- Set **HOSTNAME=localhost.localdomain** in **/etc/sysconfig/network**

- Remove **/etc/udev/rules.d/70-\***. Run:

```
# rm -rf /etc/udev/rules.d/70-*
```

- Remove the **HWADDR=** line from **/etc/sysconfig/network-scripts/ifcfg-eth\***.

- Optionally delete all the logs from **/var/log** and build logs from **/root**.

- Shut down the virtual machine. Run:

```
# poweroff
```

#### Result

The virtual machine is sealed and can be made into a template. You can deploy Linux virtual machines from this template without experiencing configuration file conflicts.


[Report a bug](#)

### 9.3.2. Sealing a Windows Template

#### 9.3.2.1. Considerations when Sealing a Windows Template with Sysprep

A template created for Windows virtual machines must be generalized (sealed) before being used to deploy virtual machines. This ensures that machine-specific settings are not reproduced in the template.

The **Sysprep** tool is used to seal Windows templates before use.


**Important**  
 Do not reboot the virtual machine during this process.

Before starting the **Sysprep** process, verify the following settings are configured:

- » The Windows **Sysprep** parameters have been correctly defined.  
If not, click **Edit** and enter the required information in the **Operating System** and **Domain** fields.
- » The correct product key has been entered in the **rhev-config** configuration tool.  
If not, run the configuration tool on the Manager as the root user, and enter the required information. The configuration keys that you need to set are **ProductKey** and **SysPrepPath**. For example, the Windows 7 configuration value is **ProductKeyWindow7** and **SysPrepWindows7Path**. Set these values with this command:

```
# rhevm-config --set ProductKeyWindow7=<validproductkey> --cver=general
```

#### See Also:


- » [Section 15.2.2, "Syntax for rhevm-config Command"](#)

[Report a bug](#)

#### 9.3.2.2. Sealing a Windows XP Template

#### Summary

Seal a Windows XP template using the **Sysprep** tool before using the template to deploy virtual machines.


**Note**  
 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 9.7. Sealing a Windows XP Template

- 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>
- Create a new directory: **c:\sysprep**.
- Open the **deploy.cab** file and add its contents to **c:\sysprep**.
- Execute **sysprep.exe** from within the folder and click **OK** on the welcome message to display the Sysprep tool.

5. Select the following check boxes:
  - » **Don't reset grace period for activation**
  - » **Use Mini-Setup**
6. Ensure that the shutdown mode is set to **Shut down** and click **Reseal**.
7. Acknowledge the pop-up window to complete the sealing process; the virtual machine shuts down automatically upon completion.

**Result**

The Windows XP template is sealed and ready for deploying virtual machines.

[Report a bug](#)

**9.3.2.3. Sealing a Windows 7 or Windows 2008 Template****Summary**

Seal a Windows 7 or Windows 2008 template before using the template to deploy virtual machines.

**Procedure 9.8. Sealing 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 opens. 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** to open the **Edit String** window.
  5. Enter the following information in the provided fields:
    - » Value name: **UnattendFile**
    - » Value data: **a:\sysprep.inf**
  6. Launch **Sysprep** from **C:\Windows\System32\sysprep\sysprep.exe**.
  7. Enter the following information into the **Sysprep** tool:
    - » Under **System Cleanup Action**, select **Enter System Out-of-Box-Experience (OOBE)**.
    - » Select the **Generalize** check box if you need to change the computer's system identification number (SID).
    - » Under **Shutdown Options**, select **Shutdown**.
- Click **OK** to complete the sealing process; the virtual machine shuts down automatically upon completion.

**Result**

The Windows 7 or Windows 2008 template is sealed and ready for deploying virtual machines.

[Report a bug](#)


**9.4. Templates and Permissions****9.4.1. Managing System Permissions for a Template**

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **ClusterAdmin** has administrator privileges only for the assigned cluster, and so forth.

A template administrator is a system administration role for templates in a data center. This role can be applied to specific virtual machines, to a data center, or to the whole virtualized environment; this is useful to allow different users to manage certain virtual resources.

The template administrator role permits the following actions:

- » Create, edit, export, and remove associated templates; and
- » import and export templates.


**Note**

You can only assign roles and permissions to existing users.

**See Also:**

- » [Section 9.4.2, "Template Administrator Roles Explained"](#)

[Report a bug](#)

**9.4.2. Template Administrator Roles Explained****Template Administrator Permission Roles**

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

**Table 9.3. Red Hat Enterprise Virtualization System Administrator Roles**

Role	Privileges	Notes
TemplateAdmin	Can perform all operations on templates.	Has privileges to create, delete and configure templates' storage domains and network details, and to move templates between domains.
NetworkAdmin	Network Administrator	Can configure and manage networks attached to templates.

**See Also:**

- » [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**9.4.3. Template User Roles Explained****Template User Permission Roles**

The table below describes the user roles and privileges applicable to using and administrating templates in the User Portal.

Table 9.4. Red Hat Enterprise Virtualization Template User Roles

Role	Privileges	Notes
TemplateCreator	Can create, edit, manage and remove virtual machine templates within assigned resources.	The <b>TemplateCreator</b> role is not applied to a specific template; apply this role to a user for the whole environment with the <b>Configure</b> window. Alternatively apply this role for specific data centers, clusters, or storage domains.
TemplateOwner	Can edit and delete the template, assign and manage user permissions for the template.	The <b>TemplateOwner</b> role is automatically assigned to the user who creates a template. Other users who do not have <b>TemplateOwner</b> permissions on a template cannot view or use the template.
UserTemplateBasedVm	Can use the template to create virtual machines.	Cannot edit template properties.
NetworkUser	Logical network and network interface user for templates.	If the <b>Allow all users to use this Network</b> option was selected when a logical network is created, <b>NetworkUser</b> permissions are assigned to all users for the logical network. Users can then attach or detach template network interfaces to or from the logical network.

**See Also:**

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**9.4.4. Assigning an Administrator or User Role to a Resource****Summary**

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

**Procedure 9.9. Assigning a Role to a Resource**

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

**Result**

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

**See Also:**

» [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**9.4.5. Removing an Administrator or User Role from a Resource****Summary**

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

**Procedure 9.10. Removing a Role from a Resource**

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

**Result**

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

[Report a bug](#)

## Chapter 10. Pools

### 10.1. Introduction to Virtual Machine Pools

Virtual machines can be placed in pools for use on demand by any user in a particular group. A single virtual machine cannot be used concurrently by multiple users. Different virtual machine pools can be set up for different purposes.

Users do not always get the same desktop, but get an available virtual machine of the required type, from the appropriate virtual machine pool.

Virtual machines in pools are stateless, data is not persistent across reboots. Virtual machines in a pool are started when there is a user request, and shut down when the user is finished.

Virtual machine pools can contain prestarted virtual machines waiting for user requests. This provides a faster response time, but uses system resources on idle virtual machines.

#### Note

Virtual machines are not stateless when accessed from the Administration Portal. This is because administrators need to be able to write changes to the disk if necessary.

[Report a bug](#)

### 10.2. Virtual Machine Pool Tasks

#### 10.2.1. Creating a Virtual Machine Pool

##### Summary

You can create a virtual pool that contains multiple virtual machines. Virtual machine pools provide virtual machines to end users in an on-demand basis.

##### Procedure 10.1. Creating a Virtual Machine Pool

1. In flat mode, click the **Pools** resource tab to display a list of virtual machine pools.
2. Click the **New** button to open the **New Pool** window.
3. Use the drop-down menus to select the **Data Center** and **Host Cluster**.
4. Enter the **Name**, **Description**, and **Number of VMs** for the pool.
5. If applicable, use the drop-down menu to select a template.
6. Enter the **Memory Size** to be used for each virtual machine in the pool and the **Total Virtual CPUs**.
7. If applicable, click the **Advanced Parameters** expansion button and use the drop-down menus to select the **Cores per Virtual Socket** and **Virtual Sockets**.
8. Use the drop-down menu to select the **Operating System** to be used for the virtual machines in the pool.
9. In the **Pool** tab, select one of the following pool types:
  - » **Manual** - The administrator is responsible for explicitly returning the virtual machine to the pool. The virtual machine reverts to the original base image after the administrator returns it to the pool.
  - » **Automatic** - When the virtual machine is shut down, it automatically reverts to its base image and is returned to the virtual machine pool.
10. The **Initial Run**, **Console**, **Host**, **Resource Allocation**, **Boot Options**, and **Custom Properties** tabs are not mandatory but define options for your pool. These tabs feature identical settings and controls as the **New Virtual Machines** window.
11. Click **OK** to create the pool.

##### Result

You have created and configured a virtual machine pool with the specified number of identical virtual machines.

You can view these virtual machines in the **Virtual Machines** resource tab, or in the details pane of the **Pools** resource tab; a virtual machine in a pool is distinguished from independent virtual machines by its icon.

General Virtual Machines Permissions			
Detach			
Name	Host	IP Address	Status
■ test_pool-1			Down
■ test_pool-2			Down

Figure 10.1. Virtual Machines Displaying in the Details Pane of a Virtual Machine Pool

##### See Also:

- » [Section 8.4.2, "Creating a New Virtual Machine from a Blank Template"](#)
- » [Section 8.4.3.3, "Virtual Machine Console Settings Explained"](#)
- » [Section 8.4.3, "Explanation of Settings and Controls in the New Virtual Machine and Edit Virtual Machine Windows"](#)

[Report a bug](#)

#### 10.2.2. Editing a Resource

##### Summary

Edit the properties of a resource.

##### Procedure 10.2. Editing a Resource

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

##### Result

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

##### See Also:

» [Section 8.4.3, "Explanation of Settings and Controls in the New Virtual Machine and Edit Virtual Machine Windows"](#).

[Report a bug](#)

### 10.2.3. Prestarting Virtual Machines in a Pool

#### Prerequisites:

» [Section 10.2.1, "Creating a Virtual Machine Pool"](#)

In a virtual machine pool, the machines are in a powered down state by default. When a user requests a virtual machine from a pool, a machine is powered up and assigned to the user. In contrast, a prestarted virtual machine is already running and waiting to be assigned to a user, decreasing the amount of time a user has to wait before being able to access a machine. When a prestarted virtual machine is shut down it is returned to the pool and restored to its original state. The maximum number of prestarted virtual machines is the number of virtual machines in the pool.

#### Summary

Prestarted virtual machines are suitable for environments in which users require immediate access to virtual machines which are not specifically assigned to them. Only automatic pools can have prestarted virtual machines.

#### Procedure 10.3. Prestarting Virtual Machines in a Pool

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Click **Edit** to open the **Edit Pool** window.
3. Enter the number of virtual machines to be prestarted in the **Prestarted VMs** field.
4. Select the **Pool** tab. Ensure that the **Pool Type** is set to **Automatic**.
5. Click **OK**.

#### Result

You have set a number of prestarted virtual machines in a pool. The prestarted machines are running and available for use.

[Report a bug](#)

### 10.2.4. Adding Virtual Machines to a Virtual Machine Pool

#### Summary

If you require more virtual machines than was originally provisioned in a virtual machine pool, add more machines to the pool.

#### Procedure 10.4. Adding Virtual Machines to a Virtual Machine Pool

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Click **Edit** to open the **Edit Pool** window.
3. Enter the number of additional virtual machines to add in the **Increase number of VMs in pool by** field.
4. Click **OK**.  
The total number of virtual machines in the pool has increased.

#### Result

You have added more virtual machines to the virtual machine pool.

[Report a bug](#)


### 10.2.5. Detaching Virtual Machines from a Virtual Machine Pool

#### Summary

You can detach virtual machines from a virtual machine pool. Detaching a virtual machine removes it from the pool to become an independent virtual machine.

#### Procedure 10.5. Detaching Virtual Machines from a Virtual Machine Pool

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Ensure the virtual machine has a status of **Down** as you cannot detach a running virtual machine.  
Click the **Virtual Machines** tab in the details pane to list the virtual machines in the pool.
3. Select the virtual machine that you want to remove and click **Detach** to open the **Detach Virtual Machine(s)** confirmation window.
4. Click **OK** to detach the virtual machine from the pool.



General Virtual Machines Permissions			
Detach			
Name	Host	IP Address	Status
test_pool-1			Down
test_pool-2			Down

Figure 10.2. A Detached Virtual Machine

#### Result

You have detached the virtual machine from the virtual machine pool.

[Report a bug](#)

### 10.2.6. Removing a Virtual Machine Pool

#### Summary

You can remove a virtual machine pool from a data center. You must first either delete or detach all of the virtual machines in the pool. Detaching virtual machines from the pool will preserve them as independent virtual machines.

**Procedure 10.6. Removing a Virtual Machine Pool**

1. Use the **Pools** resource tab, tree mode, or the search function to find and select the virtual machine pool in the results list.
2. Click **Remove** to open the **Remove Pool(s)** confirmation window.
3. Click **OK** to remove the pool.

**Result**

You have removed the pool from the data center.

[Report a bug](#)

**10.3. Pools and Permissions****10.3.1. Managing System Permissions for a Virtual Machine Pool**

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **ClusterAdmin** has administrator privileges only for the assigned cluster, and so forth.

A virtual machine pool administrator is a system administration role for virtual machine pools in a data center. This role can be applied to specific virtual machine pools, to a data center, or to the whole virtualized environment; this is useful to allow different users to manage certain virtual machine pool resources.

The virtual machine pool administrator role permits the following actions:

- » Create, edit, and remove pools; and
- » Add and detach virtual machines from the pool.

**Note**

You can only assign roles and permissions to existing users.

**See Also:**

- » [Section 10.3.2, "Virtual Machine Pool Administrator Roles Explained"](#)

[Report a bug](#)

**10.3.2. Virtual Machine Pool Administrator Roles Explained****Pool Permission Roles**

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

**Table 10.1. Red Hat Enterprise Virtualization System Administrator Roles**

Role	Privileges	Notes
VmPoolAdmin	System Administrator role of a virtual pool.	Can create, delete, and configure a virtual pool, assign and remove virtual pool users, and perform basic operations on a virtual machine.
ClusterAdmin	Cluster Administrator	Can use, create, delete, manage all virtual machine pools in a specific cluster.

**See Also:**

- » [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**10.3.3. Assigning an Administrator or User Role to a Resource****Summary**

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

**Procedure 10.7. Assigning a Role to a Resource**

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

**Result**

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

**See Also:**

- » [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

**10.3.4. Removing an Administrator or User Role from a Resource****Summary**

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

**Procedure 10.8. Removing a Role from a Resource**

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

inherited permissions for the selected resource.

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

**Result**

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

[Report a bug](#)

## Chapter 11. Virtual Machine Disks

### 11.1. Understanding Virtual Machine Storage

Red Hat Enterprise Virtualization supports three storage types: NFS, iSCSI and FCP.

In each type, a host known as the Storage Pool Manager (SPM) manages access between hosts and storage. The SPM host is the only node that has full access within the storage pool; the SPM can modify the storage domain metadata, and the pool's metadata. All other hosts can only access virtual machine hard disk image data.

By default in an NFS, local, or POSIX compliant data center, the SPM creates the virtual disk using a thin provisioned format as a file in a file system.

In iSCSI and other block-based data centers, the SPM creates a volume group on top of the Logical Unit Numbers (LUNs) provided, and makes logical volumes to use as virtual machine disks. Virtual machine disks on block-based storage are preallocated by default.

If the virtual disk is preallocated, a logical volume of the specified size in GB is created. The virtual machine can be mounted on a Red Hat Enterprise Linux server using `kpartx`, `vgscan`, `vgchange` and `mount` to investigate the virtual machine's processes or problems.

If the virtual disk is a thin provisioned, a 1 GB logical volume is created. The logical volume is continuously monitored by the host on which the virtual machine is running. As soon as the usage nears a threshold the host notifies the SPM, and the SPM extends the logical volume by 1 GB. The host is responsible for resuming the virtual machine after the logical volume has been extended. If the virtual machine goes into a paused state it means that the SPM could not extend the disk in time. This occurs if the SPM is too busy or if there is not enough storage space.

A virtual disk with a preallocated (RAW) format has significantly faster write speeds than a virtual disk with a thin provisioning (Qcow2) format. Thin provisioning takes significantly less time to create a virtual disk. The thin provision format is suitable for non-IO intensive virtual machines.

[Report a bug](#)

### 11.2. Understanding Virtual Disks

Virtual disks are of two types, **Thin Provisioned** or **Preallocated**. Preallocated disks are RAW formatted. Thin provisioned disks are Qcow2 formatted.

#### » Preallocated

A preallocated virtual disk has reserved storage of the same size as the virtual disk itself. The backing storage device (file/block device) is presented as is to the virtual machine with no additional layering in between. 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 deduplicated in the hard disk image file. (If these assumptions are incorrect, do not select Preallocated for NFS virtual disks).

#### » Thin Provisioned

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.

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

**Table 11.1. Permitted Storage Combinations**

Storage NFS or iSCSI/FCP	Format RAW or Qcow2	Type Sparse or Preallocated	Note
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.
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).

[Report a bug](#)



### 11.3. Shareable Disks in Red Hat Enterprise Virtualization

Some applications require storage to be shared between servers. Red Hat Enterprise Virtualization allows you to mark virtual machine hard disks as **shareable** and attach those disks to virtual machines. That way a single virtual disk can be used by multiple cluster-aware guests.

Shared disks are not to be used in every situation. For applications like clustered database servers, and other highly available services, shared disks are appropriate. Attaching a shared disk to multiple guests that are not cluster-aware is likely to cause data corruption because their reads and writes to the disk are not coordinated.

You cannot take a snapshot of a shared disk. Virtual disks that have snapshots taken of them cannot later be marked shareable.

You can mark a disk shareable either when you create it, or by editing the disk later.

[Report a bug](#)

### 11.4. Creating Unassociated Virtual Machine Hard Disks

#### Summary

You can create a virtual hard disk independently from any virtual machines. You can then attach your disk to a virtual machine, or many virtual machines if it's shareable.

#### Procedure 11.1. Creating Unassociated Virtual Machine Hard Disks

1. Select the **Disks** resource tab.
2. Click **Add** to open the **Add Virtual Disk** window.

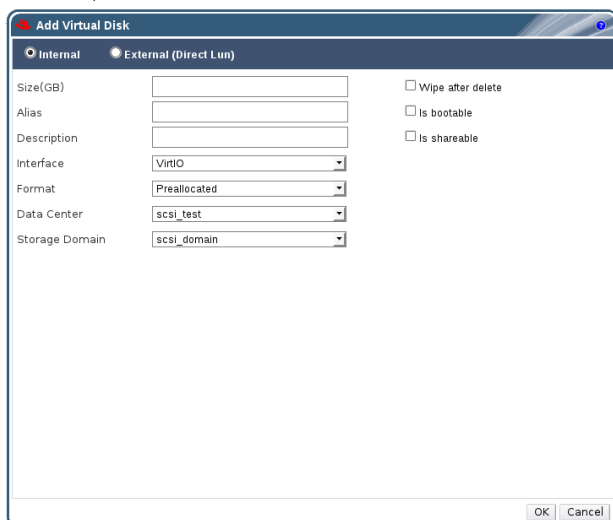


Figure 11.1. Add Virtual Disk Window

3. Use the radio buttons to select **Internal** or **External(Direct Lun)**.
4. Enter the **Size(GB)**, **Alias**, and **Description** of the virtual disk.
5. Use the drop-down menus to select the **Interface**, **Format**, **Data Center**, and **Storage Domain** of the virtual disk.
6. Select the appropriate check boxes to wipe the disk after delete, make the disk bootable, and make the disk shareable.
7. Click **OK** to create the virtual disk and close the window.

#### Result

You have created a virtual disk which can be shared between different virtual machines.

#### See Also:

- ▮ [Section 11.5, "Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows"](#)

[Report a bug](#)

### 11.5. Explanation of Settings in the New Virtual Machine Disk and Edit Virtual Machine Disk Windows

Table 11.2. Add a disk settings: Internal

Field Name	Description
<b>Size(GB)</b>	Size of new virtual disk in GB.
<b>Alias</b>	The name of the template, limited to 40 characters.
<b>Description</b>	Optionally describe the new virtual disk.
<b>Interface</b>	The virtual interface the disk presents to virtual machines. VirtIO is faster but requires VirtIO drivers. Red Hat Enterprise Linux 5 and higher includes VirtIO drivers. Windows does not include VirtIO drivers, but they can be installed from the guest tools ISO or virtual floppy disk. IDE devices do not require special drivers.
<b>Format</b>	The provisioning policy for the new virtual disk. Preallocated allots the entire size of the disk on the data domain at creation time. Thin provision starts at 1 GB and sets a maximum limit on the size the disk can grow to. Preallocated disks take more time to create than thin provisioned ones, but have better read and write performance. Preallocated is the recommended format for server disks. Thin provisioned disks are faster to create and allow for storage over-commitment. Thin provisioned is the recommended format for desktop disks.
<b>Data Center</b>	The data center in which this disk is available.
<b>Storage Domain</b>	The data domain that the disk is stored on.
<b>Wipe after delete</b>	Enhanced security for deletion of sensitive material when the disk is deleted.
<b>Is bootable</b>	Sets the bootable flag on the disk.
<b>Is shareable</b>	This disk can be attached to more than one virtual machine at a time.

The **External(Direct LUN)** settings include all of the entries from the **Internal** except for **Size(GB)** and some additional entries.

Table 11.3. Add a disk settings: External(Direct LUN)

Field Name	Description
<b>Use Host</b>	The host to use to mount the LUN.
<b>Storage Type</b>	Type of external LUN to add. Either iSCSI or Fibre Channel targets are allowed.
<b>Discover Targets</b>	Can be expanded when you are using iSCSI external LUNs and <b>Targets &gt; LUNs</b> is selected.
<b>Address</b>	The hostname or IP address of the target server. The <b>Address</b> field is visible when you are using iSCSI external LUNs and <b>Targets &gt; LUNs</b> is selected.
<b>Port</b>	The port that has been opened on the target server. The default is 3260. The <b>Port</b> field is visible when you are using iSCSI external LUNs and <b>Targets &gt; LUNs</b> is selected.
<b>User Authentication</b>	The iSCSI server requires User Authentication. The <b>User Authentication</b> field is visible when you are using iSCSI external LUNs and <b>Targets &gt; LUNs</b> is selected.
<b>CHAP username and password</b>	The username and password of a user with permission to login to LUNs. The <b>CHAP username and password</b> fields are accessible when you are using password protected iSCSI external LUNs and <b>Targets &gt; LUNs</b> is selected.

Fill in all the fields correctly and your external LUNs are displayed to be used as virtual disks.

Using LUNs directly as virtual machine hard disk images removes a layer of abstraction between your virtual machines and their data.

Here are some things that you should keep in mind when you choose to use a direct LUN as a virtual machine hard disk image:

- » Live storage migration of direct LUN hard disk images is not supported.
- » Direct LUN disks are not included in virtual machine exports.
- » Direct LUN disks are not included in virtual machine snapshots.

See Also:

- » [Section 7.6.4, "Un-useable LUNs in Red Hat Enterprise Virtualization"](#)

[Report a bug](#)

## 11.6. Moving a Virtual Machine Hard Disk Between Data Domains

### Summary

You would like to move a virtual hard disk from one data domain to another. You might want to do this to take advantage of high performance storage, or because you would like to decommission one of your storage domains that contains virtual hard disks.

If your virtual machine was created from a template, then the template must exist on any data domain that you migrate the virtual machine disks to. Use an export domain to move templates between storage domains.

#### Procedure 11.2. Moving a Virtual Machine Hard Disk Between Data Domains

1. Select the **Disks** resource tab.
2. Select the virtual disk or disks to move.
3. Click **Move** to open the **Move Disk(s)** window.
4. Use the drop-down menu or menus to select the **Target** data domain.
5. Click **OK** to move the disks and close the window.

**Result**

Disks have a **Locked** status while being moved. Upon completion, the virtual disk has been moved from the source domain to the target domain.

**See Also:**

- » [Section 9.2.6.1, "Migrating Templates to the Export Domain"](#)

[Report a bug](#)

## 11.7. Virtual Disks and Permissions

### 11.7.1. Managing System Permissions for a Virtual Disk

The system administrator, as the **SuperUser**, manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for empowering a user with certain administrative privileges that limit them to a specific resource: a **DataCenterAdmin** role has administrator privileges only for the assigned data center, a **ClusterAdmin** has administrator privileges only for the assigned cluster, a **StorageAdmin** has administrator privileges only for the assigned storage domain, and so forth.

Red Hat Enterprise Virtualization Manager provides two default virtual disk user roles, but no default virtual disk administrator roles. One of these user roles, the **DiskCreator** role, enables the administration of virtual disks from the User Portal. This role can be applied to specific virtual machines, to a data center, to a specific storage domain, or to the whole virtualized environment; this is useful to allow different users to manage different virtual resources.

The virtual disk creator role permits the following actions:

- » Create, edit, and remove virtual disks associated with a virtual machine or other resources; and
- » Edit user permissions for virtual disks.

**Note**

You can only assign roles and permissions to existing users.

**See Also:**

- » [Section 11.7.2, "Virtual Disk User Roles Explained"](#)

[Report a bug](#)

### 11.7.2. Virtual Disk User Roles Explained

#### Virtual Disk User Permission Roles

The table below describes the user roles and privileges applicable to using and administering virtual machine disks in the User Portal.

**Table 11.4. Red Hat Enterprise Virtualization System Administrator Roles**

Role	Privileges	Notes
DiskOperator	Virtual disk user.	Can use, view and edit virtual disks. Inherits permissions to use the virtual machine to which the virtual disk is attached.
DiskCreator	Can create, edit, manage and remove virtual machine disks within assigned clusters or data centers.	This role is not applied to a specific virtual disk; apply this role to a user for the whole environment with the <b>Configure</b> window. Alternatively apply this role for specific data centers, clusters, or storage domains.

**See Also:**

- » [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

### 11.7.3. Assigning an Administrator or User Role to a Resource

#### Summary

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

#### Procedure 11.3. Assigning a Role to a Resource

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

**Result**

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

**See Also:**

- » [Section 12.4.2, "User and Administrator Roles"](#)

[Report a bug](#)

### 11.7.4. Removing an Administrator or User Role from a Resource

#### Summary

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

#### Procedure 11.4. Removing a Role from a Resource

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

**Result**

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

[Report a bug](#)

## **Part II. Administering the Environment**

## Chapter 12. Users and Roles

### 12.1. Introduction to Users

Red Hat Enterprise Virtualization uses external directory services for user authentication and information. All user accounts must be created in external directory servers; these users are called *directory users*. The exception is the **admin** user which resides in the **internal** domain created during installation.

After a directory server is attached to Red Hat Enterprise Virtualization Manager, the users in the directory can be added to the Administration Portal, making them *Red Hat Enterprise Virtualization Manager users*. Red Hat Enterprise Virtualization Manager users can be assigned different roles and permissions according to the tasks they have to perform.

There are two types of Red Hat Enterprise Virtualization Manager users - end users who use and manage virtual resources from the User Portal, and administrators who maintain the system infrastructure using the Administration Portal. **User** roles and **admin** roles can be assigned to Red Hat Enterprise Virtualization Manager users for individual resources like virtual machines and hosts, or on a hierarchy of objects like clusters and data centers.

[Report a bug](#)

### 12.2. Directory Users

#### 12.2.1. Directory Services Support in Red Hat Enterprise Virtualization

During installation Red Hat Enterprise Virtualization Manager creates its own internal administration user, **admin**. This account is intended for use when initially configuring the environment, and for troubleshooting. To add other users to Red Hat Enterprise Virtualization you will need to attach a directory server to the Manager using the Domain Management Tool, **rhevmanage-domains**.

Once at least one directory server has been attached to the Manager you will be able to add users that exist in the directory server and assign roles to them using the Administration Portal. Users will be identified by their User Principle Name (UPN) of the form **user@domain**. Attachment of more than one directory server to the Manager is also supported.

The directory servers currently supported for use with Red Hat Enterprise Virtualization 3.2 are:

- » Active Directory;
- » Identity Management (IdM); and
- » Red Hat Directory Server 9 (RHDS 9).

You must ensure that the correct DNS records exist for your directory server. In particular you must ensure that the DNS records for the directory server include:

- » A valid pointer record (PTR) for the directory server's reverse look-up address.
- » A valid service record (SRV) for LDAP over TCP port **389**.
- » A valid service record (SRV) for Kerberos over TCP port **88**.
- » A valid service record (SRV) for Kerberos over UDP port **88**.

If these records do not exist in DNS then you will be unable to add the domain to the Red Hat Enterprise Virtualization Manager configuration using **rhevmanage-domains**.

For more detailed information on installing and configuring a supported directory server, refer to the vendor's documentation:

- » Active Directory - <http://technet.microsoft.com/en-us/windows/server/dd448614>.
- » Identity Management (IdM) - [http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/6/html/Identity\\_Management\\_Guide/index.html](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Identity_Management_Guide/index.html)
- » Red Hat Directory Server (RHDS) - [http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Directory\\_Server/index.html](http://docs.redhat.com/docs/en-US/Red_Hat_Directory_Server/index.html)



#### Important

A user must be created in the directory server specifically for use as the Red Hat Enterprise Virtualization administrative user. Do *not* use the administrative user for the directory server as the Red Hat Enterprise Virtualization administrative user.



#### Important

It is not possible to install Red Hat Enterprise Virtualization Manager (*rhev*) and IdM (*ipa-server*) on the same system. IdM is incompatible with the *mod\_ssl* package, which is required by Red Hat Enterprise Virtualization Manager.



#### Important

If you are using Active Directory as your directory server, and you wish to use **sysprep** in the creation of Templates and Virtual Machines, then the Red Hat Enterprise Virtualization administrative user must be delegated control over the Domain to:

- » **Join a computer to the domain**
- » **Modify the membership of a group**

For information on creation of user accounts in Active Directory refer to <http://technet.microsoft.com/en-us/library/cc732336.aspx>.

For information on delegation of control in Active Directory refer to <http://technet.microsoft.com/en-us/library/cc732524.aspx>.

## Note

Red Hat Enterprise Virtualization Manager uses Kerberos to authenticate with directory servers. RHDS does not provide native support for Kerberos. If you are using RHDS as your directory server then you must ensure that the directory server is made a service within a valid Kerberos domain. To do this you will need to perform these steps while referring to the relevant directory server documentation:

- » Configure the **memberOf** plug-in for RHDS to allow group membership. In particular ensure that the value of the **memberofgroupattr** attribute of the **memberOf** plug-in is set to **uniqueMember**. Consult the Red Hat Directory Server 9.0 *Plug-in Guide* for more information on configuring the **memberOf** plug-in.
- » Define the directory server as a service of the form **ldap/hostname@REALMNAME** in the Kerberos realm. Replace **hostname** with the fully qualified domain name associated with the directory server and **REALMNAME** with the fully qualified Kerberos realm name. The Kerberos realm name must be specified in capital letters.
- » Generate a **keytab** file for the directory server in the Kerberos realm. The **keytab** file contains pairs of Kerberos principals and their associated encrypted keys. These keys will allow the directory server to authenticate itself with the Kerberos realm. Consult the documentation for your Kerberos principle for more information on generating a **keytab** file.
- » Install the **keytab** file on the directory server. Then configure RHDS to recognize the **keytab** file and accept Kerberos authentication using GSSAPI. Consult the Red Hat Directory Server 9.0 *Administration Guide* for more information on configuring RHDS to use an external **keytab** file.
- » Test the configuration on the directory server by using the **kinit** command to authenticate as a user defined in the Kerberos realm. Once authenticated run the **ldapsearch** command against the directory server. Use the **-Y GSSAPI** parameters to ensure the use of Kerberos for authentication.

[Report a bug](#)

## 12.3. User Authorization

### 12.3.1. User Authorization Model

Red Hat Enterprise Virtualization applies authorization controls based on the combination of the three components:

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

[Report a bug](#)

### 12.3.2. User Actions

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




Figure 12.1. Actions



## Important

Some actions are performed on more than one object. For example, copying a template to another storage domain will impact both the template and the destination storage domain. The user performing an action must have appropriate permissions for all objects the action impacts.

[Report a bug](#)

### 12.3.3. User Permissions

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

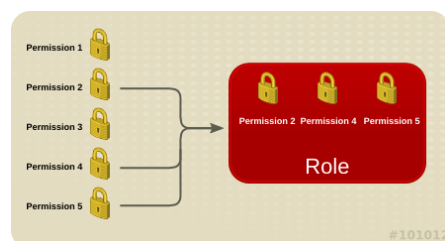


Figure 12.2. Permissions & Roles

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.

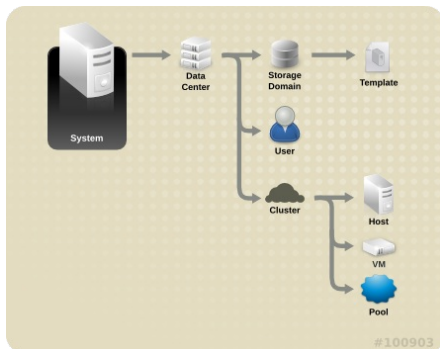


Figure 12.3. Red Hat Enterprise Virtualization Object Hierarchy

[Report a bug](#)

## 12.4. Red Hat Enterprise Virtualization Manager User Properties and Roles

### 12.4.1. User Properties

Roles and permissions are the properties of the user. Roles are predefined sets of privileges that permit access to different levels of physical and virtual resources. Multilevel administration provides a finely grained hierarchy of permissions. For example, a data center administrator has permissions to manage all objects in the data center, while a host administrator has system administrator permissions to a single physical host. A user can have permissions to use a single virtual machine but not make any changes to the virtual machine configurations, while another user can be assigned system permissions to a virtual machine.

[Report a bug](#)

### 12.4.2. User and Administrator Roles

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

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

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

See Also:

- » [Section 12.6, "User Role and Authorization Examples"](#)
- » [Section 12.4.3, "User Roles Explained"](#)
- » [Section 12.4.4, "Administrator Roles Explained"](#)

[Report a bug](#)

### 12.4.3. User Roles Explained

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

Table 12.1. Red Hat Enterprise Virtualization User Roles - Basic

Role	Privileges	Notes
UserRole	Can access and use virtual machines and pools.	Can log in to the <i>User Portal</i> , use assigned virtual machines and pools, view virtual machine state and details.
PowerUserRole	Can create and manage virtual machines and templates.	Apply this role to a user for the whole environment with the <b>Configure</b> window, or for specific data centers or clusters. For example, if a <b>PowerUserRole</b> is applied on a data center level, the <b>PowerUser</b> can create virtual machines and templates in the data center.
UserVmManager	System administrator of a virtual machine.	Can manage virtual machines, create and use snapshots, and migrate virtual machines. A user who creates a virtual machine in the <i>User Portal</i> is automatically assigned the <b>UserVmManager</b> role on the machine.

#### Note

In Red Hat Enterprise Virtualization 3.0, the **PowerUserRole** only granted permissions for virtual machines which are directly assigned to the **PowerUser**, or virtual machines created by the **PowerUser**. Now, the **VmCreator** role provides privileges previously conferred by the **PowerUserRole**. The **PowerUserRole** can now be applied on a system-wide level, or on specific data centers or clusters, and grants permissions to all virtual machines and templates within the system or specific resource. Having a **PowerUserRole** is equivalent to having the **VmCreator**, **DiskCreator**, and **TemplateCreator** roles.

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



Table 12.2. Red Hat Enterprise Virtualization User Roles - Advanced

Role	Privileges	Notes
UserTemplateBasedVm	Limited privileges to only use Templates.	Can use templates to create virtual machines.
DiskOperator	Virtual disk user.	Can use, view and edit virtual disks. Inherits permissions to use the virtual machine to which the virtual disk is attached.
VmCreator	Can create virtual machines in the User Portal.	This role is not applied to a specific virtual machine; apply this role to a user for the whole environment with the <b>Configure</b> window. Alternatively apply this role for specific data centers or clusters. When applying this role to a cluster, you must also apply the DiskCreator role on an entire data center, or on specific storage domains.
TemplateCreator	Can create, edit, manage and remove virtual machine templates within assigned resources.	This role is not applied to a specific template; apply this role to a user for the whole environment with the <b>Configure</b> window. Alternatively apply this role for specific data centers, clusters, or storage domains.
DiskCreator	Can create, edit, manage and remove virtual machine disks within assigned clusters or data centers.	This role is not applied to a specific virtual disk; apply this role to a user for the whole environment with the <b>Configure</b> window. Alternatively apply this role for specific data centers or storage domains.
TemplateOwner	Can edit and delete the template, assign and manage user permissions for the template.	This role is automatically assigned to the user who creates a template. Other users who do not have <b>TemplateOwner</b> permissions on a template cannot view or use the template.
NetworkUser	Logical network and network interface user for virtual machine and template.	Can attach or detach network interfaces from specific logical networks.

[Report a bug](#)

#### 12.4.4. Administrator Roles Explained

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

Table 12.3. Red Hat Enterprise Virtualization System Administrator Roles - Basic

Role	Privileges	Notes
SuperUser	System Administrator of the Red Hat Enterprise Virtualization environment.	Has full permissions across all objects and levels, can manage all objects across all data centers.
ClusterAdmin	Cluster Administrator.	Can use, create, delete, manage all resources in a cluster, including hosts, templates and virtual machines.
DataCenterAdmin	Data Center Administrator.	Can use, create, delete, manage all resources in a data center, including clusters, hosts, templates and virtual machines.



#### Important

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

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

Table 12.4. Red Hat Enterprise Virtualization System Administrator Roles - Advanced

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

[Report a bug](#)

## 12.5. Red Hat Enterprise Virtualization Manager User Tasks

### 12.5.1. Adding Users

#### Prerequisites:

- [Section 12.2.1, "Directory Services Support in Red Hat Enterprise Virtualization"](#)

#### Summary

Users in Red Hat Enterprise Virtualization must be added from an external directory service before they can be assigned roles and permissions.

#### Procedure 12.1. Adding Users to Red Hat Enterprise Virtualization

1. Click the **Users** tab to display the list of authorized users.
2. Click **Add**. The **Add Users and Groups** window opens.



Figure 12.4. Add Users and Groups Window

3. In the **Search** drop down menu, select the appropriate 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 check boxes for the appropriate users or groups.
5. Click **OK**.

#### Result

The added user displays on the **Users** tab.

[Report a bug](#)

### 12.5.2. Viewing User Information

#### Summary

You can view detailed information on each user in the **Users** tab.

#### Procedure 12.2. Viewing User Information

1. Click the **Users** tab to display the list of authorized users.
2. Select the user, or perform a search if the user is not visible on the results list.
3. The details pane displays for the selected user, usually with the **General** tab displaying general information, such as the domain name, email and status of the user.
4. The other tabs allow you to view groups, permissions, quotas, and events for the user.  
For example, to view the groups to which the user belongs, click the **Directory Groups** tab.

#### Result

You have viewed domain, permissions, quota, group and event information for a user.

[Report a bug](#)

### 12.5.3. Viewing User Permissions on Resources

#### Summary

Users can be assigned permissions on specific resources or a hierarchy of resources. You can view the assigned users and their permissions on each resource.

#### Procedure 12.3. Viewing User Permissions on Resources

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

#### Result

You have viewed the assigned users and their roles for a selected resource.

[Report a bug](#)

### 12.5.4. Removing Users

#### Summary

When a user account is no longer required, remove it from Red Hat Enterprise Virtualization.

#### Procedure 12.4. Removing Users

1. Click the **Users** tab to display the list of authorized users.

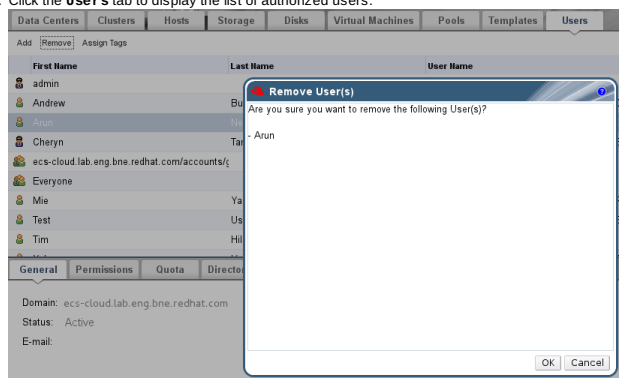


Figure 12.5. Users Tab

2. Select the user to be removed. Ensure the user is not running any virtual machines.
3. Click the **Remove** button. A message displays prompting you to confirm the removal. Click **OK**.

#### Result

The user is removed from Red Hat Enterprise Virtualization, but not from the external directory.

[Report a bug](#)

### 12.5.5. Configuring Roles

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

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

#### See Also:

- » [Section 12.3.1, "User Authorization Model"](#)

[Report a bug](#)

### 12.5.6. Creating a New Role

#### Summary

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

#### Procedure 12.5. Creating a New Role

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

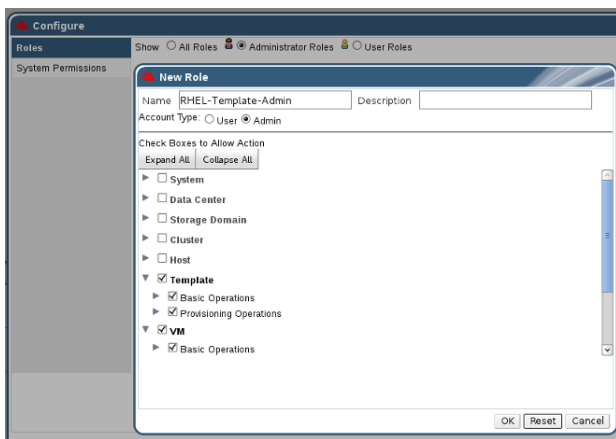


Figure 12.6. The New Role Dialog

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

**Result**

You have created a new role with permissions to specific resources. You can assign the new role to users.

[Report a bug](#)

**12.5.7. Editing or Copying a Role****Summary**

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

**Procedure 12.6. Editing or Copying a Role**

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

**Result**

You have edited the properties of a role, or cloned a role.

[Report a bug](#)

**12.6. User Role and Authorization Examples**

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

**Example 12.1. 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. 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 12.2. VM PowerUser Permissions**

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

**Example 12.3. Data Center Power User Role Permissions**

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

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

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

**Example 12.4. Network Administrator Permissions**

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

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

**Example 12.5. 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.

**Figure 12.7. 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 12.7, "UserManager Custom Role"](#). This means they apply to all other objects in the system. The role is set to have an **Account Type** of **Admin**. This means that when she is assigned this role, Rachel can only use the Administration Portal, not the User Portal.

[Report a bug](#)

## Chapter 13. Quotas

### 13.1. Introduction to Quota

Quota is a resource limitation tool provided with Red Hat Enterprise Virtualization. Quota may be thought of as a layer of limitations on top of the layer of limitations set by User Permissions.

Quota is a data-center object.

Quota allows administrators of Red Hat Enterprise Virtualization environments to limit user access to memory, CPU, and storage. Quota defines the memory resources and storage resources an administrator can assign users. As a result users may draw on only the resources assigned to them. When the quota resources are exhausted, Red Hat Enterprise Virtualization does not permit further user actions.

There are two different kinds of Quota:

**Table 13.1. The Two Different Kinds of Quota**

Quota type	Definition
Run-time Quota	This quota limits the consumption of run time resources, like CPU and memory.
Storage Quota	This quota limits the amount of storage available.

Quota, like SELinux, has three modes:

**Table 13.2. Quota Modes**

Quota Mode	Function
Enforced	This mode puts into effect the Quota that you have set in audit mode, limiting resources to the group or user affected by the quota.
Audit	This mode allows you to change Quota settings. Choose this mode to increase or decrease the amount of run-time quota and the amount of storage quota available to users affected by it.
Disabled	This mode turns off the run-time and storage limitations defined by the quota.

When a user attempts to run a virtual machine, the specifications of the virtual machine are compared to the storage allowance and the run-time allowance set in the applicable quota.

If starting a virtual machine causes the aggregated resources of all running virtual machines covered by a quota to exceed the allowance defined in the quota, then the Manager refuses to run the virtual machine.

When a user creates a new disk, the requested disk size is added to the aggregated disk usage of all the other disks covered by the applicable quota. If the new disk takes the total aggregated disk usage above the amount allowed by the quota, disk creation fails.

Quota allows for resource sharing of the same hardware. It supports hard and soft thresholds. Administrators can use a quota to set thresholds on resources. These thresholds appear, from the user's point of view, as 100% usage of that resource. To prevent failures when the customer unexpectedly exceeds this threshold, the interface supports a "grace" amount by which the threshold can be briefly exceeded. Exceeding the threshold results in a warning sent to the customer.



#### Important

Quota imposes limitations upon the running of virtual machines. Ignoring these limitations is likely to result in a situation in which you cannot use your virtual machines and virtual disks.

When quota is running in enforced mode, virtual machines and disks that do not have quotas assigned cannot be used.

To power on a virtual machine, a quota must be assigned to that virtual machine.

To create a snapshot of a virtual machine, the disk associated with the virtual machine must have a quota assigned.

When creating a template from a virtual machine, you are prompted to select the quota that you want the template to consume. This allows you to set the template (and all future machines created from the template) to consume a different quota than the virtual machine and disk from which the template is generated.

[Report a bug](#)

### 13.2. Shared Quota and Individually-defined Quota

Users with SuperUser permissions can create quotas for individual users or quotas for groups.

Group quotas can be set for Active Directory users. If a group of ten users are given a quota of 1TB of storage and one of the ten users fills the entire terabyte, then the entire group will be in excess of the quota and none of the ten users will be able to use any of the storage associated with their group.

An individual user's quota is set for only the individual. Once the individual user has used up all of his or her storage or run-time quota, the user will be in excess of the quota and the user will no longer be able to use the storage associated with his or her quota.

[Report a bug](#)

### 13.3. Quota Accounting

When a quota is assigned to a consumer or a resource, each action by that consumer or on the resource involving storage, vCPU, or memory results in quota consumption or quota release.

Since the quota acts as an upper bound that limits the user's access to resources, the quota calculations may differ from the actual current use of the user. The quota is calculated for the max growth potential and not the current usage.

#### Example 13.1. Accounting example

A user runs a virtual machine with 1 vCPU and 1024 MB memory. The action consumes 1 vCPU and 1024 MB of the quota assigned to that user. When the virtual machine is stopped 1 vCPU and 1024 MB of RAM are released back to the quota assigned to that user. Run-time quota consumption is accounted for only during the actual run-time of the consumer.

A user creates a virtual thin provision disk of 10GB. The actual disk usage may indicate only 3GB of that disk are actually in use. The quota consumption, however, would be 10GB, the max growth potential of that disk.

[Report a bug](#)

## 13.4. Enabling and Changing a Quota Mode in a Data Center

### Summary

This procedure enables or changes the quota mode in a data center. You must select a quota mode before you can define quotas. You must be logged in to the Web Administration Portal to follow the steps of this procedure.

Use **Audit** mode to test your quota to make sure it works as you expect it to. You do not need to have your quota in **Audit** mode to create or change a quota.

#### Procedure 13.1. Enabling and Changing Quota in a Data Center

1. Click the **Data Centers** tab in the Navigation Pane.
2. From the list of data centers displayed in the Navigation Pane, choose the data center whose quota policy you plan to edit.
3. Click **Edit** in the top left of the Navigation Pane.
4. An **Edit Data Center** window opens. In the **Quota Mode** drop-down, change the quota mode to **Enforced**.
5. Click the **OK** button.

### Result

You have now enabled a quota mode at the Data Center level. If, during the testing of your quota, you've set the mode to **Audit**, you must use this procedure again to change the quota mode to **Enforced** in order for your quota settings to take effect.

[Report a bug](#)

## 13.5. Creating a New Quota Policy

### Prerequisites:

- » [Section 13.4, "Enabling and Changing a Quota Mode in a Data Center"](#)

### Summary

You have enabled quota mode, either in Audit or Enforcing mode. You want to define a quota policy to manage resource usage in your data center.

#### Procedure 13.2. Creating a New Quota Policy

1. In tree mode, select the data center. The **Quota** tab appears in the Navigation Pane.
2. Click the **Quota** tab in the Navigation Pane.
3. Click **Add** in the Navigation Pane. The **New Quota** window opens.
4. Fill in the **Name** field with a meaningful name.  
Fill in the **Description** field with a meaningful name.
5. In the **Memory & CPU** section of the **New Quota** window, use the green slider to set **Cluster Threshold**.
6. In the **Memory & CPU** section of the **New Quota** window, use the blue slider to set **Cluster Grace**.
7. Click **Edit** on the bottom-right of the **Memory & CPU** field. An **Edit Quota** window opens.
8. Under the **Memory** field, select either the **Unlimited** radio button (to allow limitless use of Memory resources in the cluster), or select the **limit to** radio button to set the amount of memory set by this quota. If you select the **limit to** radio button, input a memory quota in megabytes (MB) in the **MB** field.
9. Under the **CPU** field, select either the **Unlimited** radio button or the **limit to** radio button to set the amount of CPU set by this quota. If you select the **limit to** radio button, input a number of vCPUs in the **vCpus** field.
10. Click **OK** in the **Edit Quota** window.
11. In the **Storage** section of the **New Quota** window, use the green slider to set **Storage Threshold**.
12. In the **Storage** section of the **New Quota** window, use the blue slider to set **Storage Grace**.
13. Click **Edit** in the **Storage** field. The **Edit Quota** window opens.
14. Under the **Storage Quota** field, select either the **Unlimited** radio button (to allow limitless use of Storage) or the **limit to** radio button to set the amount of storage to which quota will limit users. If you select the **limit to** radio button, input a storage quota size in gigabytes (GB) in the **GB** field.
15. Click **OK** in the **Edit Quota** window. You are returned to the **New Quota** window.
16. Click **OK** in the **New Quota** window.

### Result

You have created a new quota policy.

### See Also:

- » [Section 13.6, "Explanation of Quota Threshold Settings"](#)
- » [Section 13.7, "Assigning a Quota to an Object"](#)

[Report a bug](#)

## 13.6. Explanation of Quota Threshold Settings

Table 13.3. Quota thresholds and grace

Setting	Definition
Cluster Threshold	The amount of cluster resources available per data center.
Cluster Grace	The amount of the cluster available for the data center after exhausting the data center's Cluster Threshold.
Storage Threshold	The amount of storage resources available per data center.
Storage Grace	The amount of storage available for the data center after exhausting the data center's Storage Threshold.

If a quota is set to 100GB with 20% Grace, then consumers are blocked from using storage after they use 120GB of storage. If the same quota has a Threshold set at 70%, then consumers receive a warning when they exceed 70GB of storage consumption (but they remain able to consume storage until

they reach 120GB of storage consumption.) Both "Threshold" and "Grace" are set relative to the quota. "Threshold" may be thought of as the "soft limit", and exceeding it generates a warning. "Grace" may be thought of as the "hard limit", and exceeding it makes it impossible to consume any more storage resources.

[Report a bug](#)

### 13.7. Assigning a Quota to an Object

#### Prerequisites:

» [Section 13.5, "Creating a New Quota Policy"](#)

#### Summary

This procedure explains how to associate a virtual machine with a quota.

#### Procedure 13.3. Assigning a Quota to a Virtual Machine

1. In the navigation pane, select the Virtual Machine to which you plan to add a quota.
2. Click **Edit**. The **Edit Desktop Virtual Machine** window appears.
3. Select the quota you want the virtual machine to consume. Use the **Quota** drop-down to do this.
4. Click **OK**.

#### Result

You have designated a quota for the virtual machine you selected.

#### Summary

This procedure explains how to associate a virtual machine disk with a quota.

#### Procedure 13.4. Assigning a Quota to a Virtual Disk

1. In the navigation pane, select the Virtual Machine whose disk(s) you plan to add a quota.
2. In the details pane, select the disk you plan to associate with a quota.
3. Click **Edit**. The **Edit Virtual Disk** window appears.
4. Select the quota you want the virtual disk to consume.
5. Click **OK**.

#### Result

You have designated a quota for the virtual disk you selected.



#### Important

Quota must be selected for all objects associated with a virtual machine, in order for that virtual machine to work. If you fail to select a quota for the objects associated with a virtual machine, the virtual machine will not work. The error that the Manager throws in this situation is generic, which makes it difficult to know if the error was thrown because you did not associate a quota with all of the objects associated with the virtual machine. It is not possible to take snapshots of virtual machines that do not have an assigned quota. It is not possible to create templates of virtual machines whose virtual disks do not have assigned quotas.

[Report a bug](#)

### 13.8. Using Quota to Limit Resources by User

#### Prerequisites:

» [Section 13.5, "Creating a New Quota Policy"](#)

#### Summary

This procedure describes how to use quotas to limit the resources a user has access to.

#### Procedure 13.5. Assigning a User to a Quota

1. In the tree, click the Data Center with the quota you want to associate with a User.




Figure 13.1. Selecting the Data Center

2. Click the **Quota** tab in the navigation pane.






Figure 13.2. Selecting the Quota Tab

3. Select the target quota in the list in the navigation pane.




Figure 13.3. Selecting the Target Quota

4. Click the **Consumers** tab in the details pane.



Figure 13.4. Selecting the Consumers Tab

5. Click **Add** at the top of the details pane.




Figure 13.5. Clicking the Add Button

6. In the **Search** field, type the name of the user you want to associate with the quota.
7. Click **GO**.
8. Select the check box at the left side of the row containing the name of the target user.
9. Click **OK** in the bottom right of the **Assign Users and Groups to Quota** window.

#### Result

After a short time, the user will appear in the **Consumers** tab of the details pane.

[Report a bug](#)

## 13.9. Editing Quotas

### Summary

This procedure describes how to change existing quotas.

#### Procedure 13.6. Editing Quotas

1. On the tree pane, click on the data center whose quota you want to edit.




Figure 13.6. Selecting a Data Center to Edit Its Quota

2. Click on the **Quota** tab in the Navigation Pane.




Figure 13.7. Selecting the Quota Tab in the Navigation Pane

3. Click the name of the quota you want to edit.




Figure 13.8. Selecting a Quota to Edit

4. Click **Edit** in the Navigation pane.
5. An **Edit Quota** window opens. If required, enter a meaningful name in the **Name** field.
6. If required, you can enter a meaningful description in the **Description** field.
7. Select either the **All Clusters** radio button or the **Specific Clusters** radio button. Move the **Cluster Threshold** and **Cluster Grace** sliders to the desired positions on the **Memory & CPU** slider.
8. Select either the **All Storage Domains** radio button or the **Specific Storage Domains** radio button. Move the **Cluster Threshold** and **Cluster Grace** sliders to the desired positions on the **Memory & CPU** slider.
9. Click **OK** in the **Edit Quota** window to confirm the new quota settings.

#### Result

You have changed an existing quota.

[Report a bug](#)

## 13.10. Removing Quotas

### Summary

This procedure describes how to remove quotas.

#### Procedure 13.7. Removing Quotas

1. On the tree pane, click on the data center whose quota you want to edit.
2. Click on the **Quota** tab in the Navigation Pane.
3. Click the name of the quota you want to remove.
4. Click **Remove** at the top of the Navigation pane, under the row of tabs.
5. Click **OK** in the **Remove Quota(s)** window to confirm the removal of this quota.

#### Result

You have removed a quota.

[Report a bug](#)

## Chapter 14. Event Notifications

### 14.1. Configuring Event Notifications

#### Summary

You want to be emailed when an event happens in your Red Hat Enterprise Virtualization environment. You need to set up a mail transfer agent to deliver messages.

#### Procedure 14.1. Configuring Event Notifications

1. Ensure that you have set up the MTA agent with the appropriate variables.
2. Use the **Users** resource tab, tree mode, or the search function to find and select the user.
3. Click the **Event Notifier** tab in the details pane to list events for which the user will be notified. This list will be blank if the user does not have event notification configured.
4. Click **Manage Events** to open the **Add Event Notification** window.

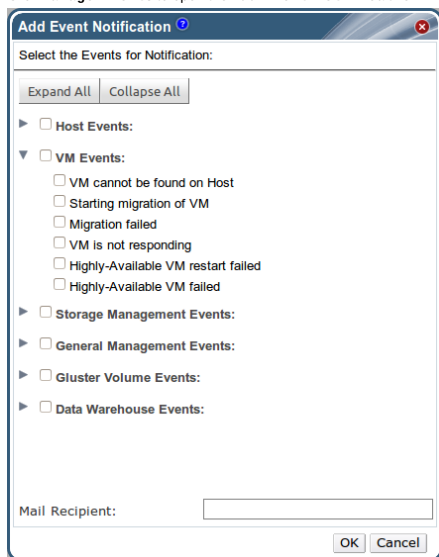


Figure 14.1. The Add Events Notification Window

5. Use the **Expand All** button, or the subject-specific expansion buttons, to view the events.
6. Select the appropriate check boxes.
7. Enter an email address in the **Mail Recipient:** field.
8. Click **OK** to save changes and close the window.
9. Add and start the **engine-notifierd** service on the Red Hat Enterprise Virtualization Manager. This activates the changes you have just made:

```
# chkconfig --add engine-notifierd
# chkconfig engine-notifierd on
# service engine-notifierd restart
```

#### Result

You now receive emails based on events in your Red Hat Enterprise Virtualization Environment. The selected events display on the **Event Notifier** tab for the user.

[Report a bug](#)

### 14.2. Parameters for event notifications in notifier.conf

The event notifier configuration file can be found in `/etc/rhev/notifier/notifier.conf`.

Table 14.1. notifier.conf variables

Variable name	Default	Remarks
INTERVAL_IN_SECONDS	120	The interval in seconds between instances of dispatching messages to subscribers.
MAIL_SERVER	none	The SMTP mail server address. Required.
MAIL_PORT	25	The default port of a non-secured SMTP server is 25. The default port of a secured SMTP server (one with SSL enabled) is 465.
MAIL_USER	none	If SSL is enabled to authenticate the user, then this variable must be set. This variable is also used to specify the "from" user address when the MAIL_FROM variable is not set. Some mail servers do not support this functionality. The address is in RFC822 format.
MAIL_PASSWORD	none	This variable is required to authenticate the user if the mail server requires authentication or if SSL is enabled.
MAIL_ENABLE_SSL	false	This indicates whether SSL should be used to communicate with the mail server.
HTML_MESSAGE_FORMAT	false	The mail server sends messages in HTML format if this variable is set to "true".
MAIL_FROM	none	This variable specifies a "from" address in RFC822 format, if supported by the mail server.
MAIL_REPLY_TO	none	This variable specifies "reply-to" addresses in RFC822 format on sent mail, if supported by the mail server.
DAYS_TO_KEEP_HISTORY	none	This variable sets the number of days dispatched events will be preserved in the history table. If this variable is not set, events remain on the history table indefinitely.

[Report a bug](#)

### 14.3. Canceling Event Notifications

#### Summary

A user has configured some unnecessary event notifications and wants them canceled.

#### Procedure 14.2. Canceling Event Notifications

1. In the **Users** tab, select the user or the user group.
2. Select the **Event Notifier** tab in the details pane to list events for which the user receives notifications.
3. Click **Manage Events** to open the **Add Event Notification** window.
4. Use the **Expand All** button, or the subject-specific expansion buttons, to view the events.
5. Clear the appropriate check boxes to remove notification for that event.
6. Click **OK** to save changes and close the window.

#### Result

You have canceled unnecessary event notifications for the user.

[Report a bug](#)

## Chapter 15. Utilities

### 15.1. Managing Domains with the Domain Management Tool

#### 15.1.1. What is the Domain Management Tool?

Red Hat Enterprise Virtualization Manager authenticates users using directory services. To add users to Red Hat Enterprise Virtualization Manager you must first use the internal **admin** user to add the directory service that the users must be authenticated against. You add and remove directory services domains using the included domain management tool, **rhevmanage-domains**.

The **rhevmanage-domains** command is only accessible on the machine on which Red Hat Enterprise Virtualization Manager is installed. The **rhevmanage-domains** command must be run as the **root** user.

[Report a bug](#)

#### 15.1.2. Syntax for the Domain Management Tool

The usage syntax is:

```
rhevmanage-domains -action=ACTION [options]
```

Available actions are:

##### **add**

Add a domain to Red Hat Enterprise Virtualization Manager's directory services configuration.

##### **edit**

Edit a domain in Red Hat Enterprise Virtualization Manager's directory services configuration.

##### **delete**

Delete a domain from Red Hat Enterprise Virtualization Manager's directory services configuration.

##### **validate**

Validate Red Hat Enterprise Virtualization Manager's directory services configuration. This command attempts to authenticate each domain in the configuration using the configured user name and password.

##### **list**

List Red Hat Enterprise Virtualization Manager's current directory services configuration.

These options can be combined with the actions on the command line:

##### **-domain=DOMAIN**

Specifies the domain on which the action will be performed. The **-domain** parameter is mandatory for **add**, **edit**, and **delete**.

##### **-provider=PROVIDER**

Specifies the LDAP provider type of the directory server for the domain. Valid values are:

- » **ActiveDirectory** - Active Directory.
- » **IPA** - Identity Management (IdM).
- » **RHDS** - Red Hat Directory Server. Red Hat Directory Server does not come with Kerberos. Red Hat Enterprise Virtualization requires Kerberos authentication. RHDS must be made a service within a Kerberos domain to provide directory services to the Manager.



If you want to use RHDS as your directory server, you must have the **memberof** plugin installed in RHDS. To use the **memberof** plugin, your users must be **inetusers**. For more information about using the **memberof** plugin, see the *Red Hat Directory Server Plug-in Guide*.

##### **-user=USER**

Specifies the domain user to use. The **-user** parameter is mandatory for **add**, and optional for **edit**.

##### **-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.

##### **-addPermissions**

Specifies that the domain user will be given the **SuperUser** role in Red Hat Enterprise Virtualization Manager. By default, if the **-addPermissions** parameter is not specified, the **SuperUser** role is not assigned to the domain user. The **-addPermissions** option is optional. It is only valid when used in combination with the **add** and **edit** actions.

##### **-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.

##### **-configFile=FILE**

Specifies an alternate configuration file that the command must load. The **-configFile** parameter is always optional.

##### **-report**

In conjunction with the **validate** action results in the output of a report of all encountered validation errors.

For full usage information consult the **rhevmanage-domains** command's help output:

```
# rhvm-manage-domains --help
```

[Report a bug](#)

### 15.1.3. Using the Domain Management Tool

The following examples demonstrate the use of the **rhvm-manage-domains** command to perform basic manipulation of the Red Hat Enterprise Virtualization Manager domain configuration.

[Report a bug](#)

### 15.1.4. Adding Domains to Configuration

In this example, the **rhvm-manage-domains** command is used to add the **IdM** domain **directory.demo.redhat.com** to the Red Hat Enterprise Virtualization Manager configuration. The configuration is set to use the **admin** user when querying the domain; the password is provided interactively.

#### Example 15.1. rhvm-manage-domains Add Action

```
# rhvm-manage-domains -action=add -domain=directory.demo.redhat.com -
provider=IPA -user=admin -interactive
loaded template kr5.conf file
setting default_tkt_enctypes
setting realms
setting domain realm
success
User guide is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully added domain directory.demo.redhat.com. oVirt Engine restart is
required in order for the changes to take place (service ovirt-engine restart).
```

[Report a bug](#)

### 15.1.5. Editing a Domain in the Configuration

In this example, the **rhvm-manage-domains** command is used 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; the password is provided interactively.

#### Example 15.2. rhvm-manage-domains Edit Action

```
# rhvm-manage-domains -action=edit -domain=directory.demo.redhat.com -
user=admin -interactive
loaded template kr5.conf file
setting default_tkt_enctypes
setting realms
setting domain realm0
success
User guide is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully edited domain directory.demo.redhat.com. oVirt Engine restart is
required in order for the changes to take place (service ovirt-engine restart).
```

[Report a bug](#)

### 15.1.6. Deleting a Domain from the Configuration

In this example, the **rhvm-manage-domains** command is used 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 the Red Hat Enterprise Virtualization Manager. The entries for the affected users will remain defined in the 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.

#### Example 15.3. rhvm-manage-domains Delete Action

```
# rhvm-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.
```

[Report a bug](#)

### 15.1.7. Validating Domain Configuration

In this example, the **rhvm-manage-domains** command is used 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. The domain is reported as valid if the attempt is successful.

#### Example 15.4. rhvm-manage-domains Validate Action

```
# rhvm-manage-domains -action=validate
User guide is: 80b71bae-98a1-11e0-8f20-525400866c73
Domain directory.demo.redhat.com is valid.
```

[Report a bug](#)

### 15.1.8. Listing Domains in Configuration

In this example, the **rhvm-manage-domains** command is used to list the domains defined in the Red Hat Enterprise Virtualization Manager configuration. For each configuration entry the command displays the domain, the user name — in User Principle Name (UPN) format — and whether the domain is local or remote.

#### Example 15.5. rhvm-manage-domains List Action

```
# rhvm-manage-domains -action=list
Domain: directory.demo.redhat.com
User name: admin@DIRECTORY.DEMO.REDHAT.COM
This domain is a remote domain.
```

[Report a bug](#)

### 15.1.9. Getting a Configuration Value

**Example 15.6. Retrieving the Value of the SearchResultsLimit Key**

```
# rhvm-config --get=SearchResultsLimit --cver=general
100
```

[Report a bug](#)

### 15.1.10. Setting a Configuration Value

**Example 15.7. Setting the Value of the SearchResultsLimit Key**

```
# rhvm-config --set SearchResultsLimit=50 --cver=general
```

[Report a bug](#)

## 15.2. Editing the Configuration of the Red Hat Virtualization Manager with the Configuration Tool

### 15.2.1. Configuration Tool

Installing the Red Hat Enterprise Virtualization Manager modifies only a subset of configuration settings from their defaults. Further modifications are made using the configuration tool: **rhvm-config**.

The configuration tool does not require JBoss or the Red Hat Enterprise Virtualization Manager to be running to update the configuration. Configuration key values are stored in the database; configuration changes will not be saved unless the database is operational. Changes are applied when JBoss is restarted.

The Red Hat Enterprise Virtualization Manager stores configuration settings 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,
- ▀ retrieve the value of a specific configuration key, and
- ▀ set the value of a specific configuration key.

You are also able to maintain multiple versions of the Manager's configuration with the configuration tool. Use the **--cver** parameter to specify the configuration version to be used when retrieving or setting a value for a configuration key. The default configuration version is **general**.

[Report a bug](#)

### 15.2.2. Syntax for rhvm-config Command

The configuration tool is accessible on the client machine on which the Red Hat Enterprise Virtualization Manager is installed. For full usage information consult the help output of the **rhvm-config** command:

```
# rhvm-config --help
```

#### Common tasks

##### List available configuration keys

Use the **--list** parameter to list available configuration keys.

```
# rhvm-config --list
```

Each available configuration key is listed by name and description.

##### List available configuration values

Use the **--all** parameter to list available configuration values.

```
# rhvm-config --all
```

Each available configuration key is listed by name, current value of the key, and the configuration version.

##### Retrieve value of configuration key

Use the **--get** parameter to retrieve the value of a specific key.

```
# rhvm-config --get KEY_NAME
```

Replace **KEY\_NAME** with the name of the specific key to retrieve the key name, value, and the configuration version. Use the **--cver** parameter to specify the configuration version of the value to 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 specific key to set; replace **KEY\_VALUE** with the value to be set. Environments with more than one configuration version require the **VERSION** to be specified.

[Report a bug](#)

### 15.2.3. The admin@internal user

The **admin@internal** user account is automatically created upon installation of the Red Hat Enterprise Virtualization Manager. This account is stored locally in the Red Hat Enterprise Virtualization Manager's PostgreSQL database, separate from external directory services such as IdM or Active Directory. Unlike external 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.

The password for the **admin@internal** user is set during the installation of the Red Hat Enterprise Virtualization Manager. Use the **rhvm-config** utility if you need to reset the password.

[Report a bug](#)

#### 15.2.4. Changing the password for admin@internal

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.

```
# rhev-config -s AdminPassword=interactive
```

Use escape characters if your password includes any special characters.

3. Restart the ovirt-engine service for the changes to take effect.

```
# service ovirt-engine restart
```

[Report a bug](#)

#### 15.2.5. Configuration tool examples

##### Example 15.8. Getting a Configuration Value

```
# rhev-config --get=SearchResultsLimit --cver=general  
100
```

##### Example 15.9. Setting a Configuration Value

```
# rhev-config --set SearchResultsLimit=50 --cver=general
```

[Report a bug](#)

#### 15.2.6. Red Hat Enterprise Virtualization Manager Configuration Options Explanations, Limitations, and Best Practices



**Table 15.1. Red Hat Enterprise Virtualization Manager Configuration Options, Explanations, Limitations, and Best Practices**

Option Name	Description	Possible values	Recommended (default) values	Comments
AsyncTaskPollingRate	Async Task Polling Rate (in seconds)	Integer	10	How often (in seconds) RHEV-M queries the status of an async task currently running.
AsyncTaskZombieTaskLifeInMinutes	Zombie tasks lifetime in minutes	Integer	3000	How long (in minutes) a task is allowed to run before assuming it has become a zombie and should be killed. The value affects large storage manipulations especially. When using slow storage and large virtual images, or when a task is known to take longer than 3000 minutes (50 hours), the value should be increased.
AuditLogAgingThreshold	Audit Log Aging Threshold (in days)	Integer	30	How long an audit log is kept before being rotated.
AuditLogCleanupTime	Time to check for Audit Log cleanup	Time	03:35:35	At what time the Audit Log is checked for Aging and cleaned up.
AuthenticationMethod	Authentication Method used by RHEV-M	String	LDAP	The API used for querying users. Currently LDAP is the only supported value.
BlockMigrationOnSwapUsagePercentage	Host swap percentage threshold (for scheduling)	Integer	0	The maximum percentage of swap space on the host that a VM run or migration is allowed to consume on the host. If the host is swapping beyond this percentage, a VM will not migrate over and will not be started.
BootstrapMinimalVdsmVersion	Minimum VDSM version	String	4.9	The minimum version of VDSM that is acceptable when adding hosts to the Engine. Newer versions have more features.
CABaseDirectory	CA Base Directory	String	/etc/pki/ovirt-engine	Where the RHEV-M Certificate Authority is located on the RHEV-M host.
CertificateFileName	Certificate File Name	String	/etc/pki/ovirt-engine/certs/engine.cer	Points to the certificate file used by RHEV-M for SSL/TLS communication with VDSM.
CpuOverCommitDurationMinutes	The duration in minutes of CPU consumption to activate selection algorithm	Integer	2	When the cluster policy is set to Even Distribution, wait for this amount of minutes after detecting CPU overcommit before triggering virtual machine migrations to rebalance the host load. This configuration value applies only for the default.
DisableFenceAtStartupInSec	Disable Fence Operations At RHEV-M Startup in Seconds	Integer	300	Allow this amount of seconds after RHEV-M starts to detect hosts, before assuming the hosts are non-responsive and proceed to fence hosts. This value should be increased when RHEV-M is on a machine that has slow network startup (a VMWare guest, for

EmulatedMachine	Supported machine types	String	RHEL 6.2.0: v3.0 RHEL 6.3.0: v3.1 RHEL 6.4.0: v3.2	example).
WANDisableEffects	Disabled WAN Effects value to send to the SPICE console	StringMultiple	animation	The list of effects which will be disabled for SPICE. Possible values: <b>animation</b> , <b>wallpaper</b> , <b>font-smooth</b> , and <b>all</b> .
WANColorDepth	WAN Color Depth value to send to the SPICE console	Integer	16	The color depth used by the SPICE. Possible values are <b>16</b> and <b>32</b> .
EnableSpiceRootCertificateValidation	Enable Spice Root Certification Validation	String	true	If 'true', the certificate of the host on which the virtual machine is running and the RHEV-M setup CA certificate are sent to the SPICE client when attempting to connect to the virtual machine with SPICE, as an extra security mechanism.
EnableUSBAsDefault	Enable USB devices attachment to the virtual machine by default	String	true	-
EnableVdsLoadBalancing	Enables Host Load Balancing system	String	true	This config value allows the user to turn on or off (true and false, respectively) the virtual machine load balancing according to the policy configured for the cluster.
FreeSpaceCriticalLowInGB	Critical low disk space alert threshold (in GB)	Integer	5	Produces an alert when a Storage Domain has this amount of space left. This setting is also used in various preliminary tests for action sanity when users try to use storage domains, to prevent reaching this critical amount. Adding and importing disks will fail if the amount of space is less than the value specified here.
FreeSpaceLow	Limit of percentage of free disk space below which it is considered low	Integer	10	When a storage domain has this percentage of space left, it is considered low on disk space.
HighUtilizationForEvenlyDistribute	High Utilization Limit For Evenly Distribute selection algorithm	Integer	75	Maximum number of virtual machines per host in the Evenly Distribute algorithm.
HighUtilizationForPowerSave	High Utilization Limit For Power Save selection algorithm	Integer	75	A default for newly created clusters, in use with PowerSave load balancing algorithm, marks the higher limit of host utilization for populating hosts.
LDAPQueryTimeout	Read Timeout in seconds for LDAP queries	Integer	30	The amount of time an LDAP query will read before the query is stopped.
LDAPOperationTimeout	Search timeout at LDAP server side	Integer	30	The amount of time an LDAP search will operate before it is stopped.
LDAPConnectTimeout	Connect timeout in seconds for LDAP queries	Integer	30	The amount of time an LDAP query will connect before it is stopped.
LocalAdminPassw	Local	Password	Populated during	The password for

ord	Administrator Password		initial setup	<b>admin@local</b> default user.
LogPhysicalMemoryThresholdInMB	Threshold for logging low host memory in MB	Integer	1024	The minimum amount of RAM left before a host is considered low on memory. If a host's RAM is lower than this setting, it is recorded on the audit log and no action is taken.
LowUtilizationForEvenlyDistribute	Low Utilization Limit for Evenly Distribute selection algorithm	Integer	0	Minimum number of virtual machines per host in the Evenly Distribute algorithm.
LowUtilizationForPowerSave	Low Utilization Limit for Power Save selection algorithm	Integer	20	A default for newly created clusters, in use with PowerSave load balancing algorithm, marks the lower limit of host utilization for populating hosts.
MacPoolRanges	MAC Addresses Pool Ranges	String	00:1A:4A:23:18:00 - 00:1A:4A:23:18:FF	The MAC address pool range to be automatically assigned to virtual machines.
MaxDiskSize	Max size of newly created disk (in GB)	Integer	16384	The maximum size (in GB) for a new virtual disk image assigned to virtual machines.
MaxMacsCountInPool	Maximum MAC Addresses count in Pool	Integer	100000	Maximum number of MAC addresses allowed in the MAC pool.
MaxNumberOfHostsInStoragePool	Maximum number of hosts in Storage Pool	Integer	250	Limits the maximum number of hosts assigned to the clusters of a single Data Center. This can be increased after testing more hosts, if necessary.
MaxNumOfCpuPerSocket	Maximum Number of CPU per socket	Integer	16	Maximum number of virtual CPU cores that can be assigned to a single virtual CPU socket.
MaxNumOfVMCPUs	Total Numbers of Virtual Machine CPUs	Integer	160	Maximum total amount of CPU cores assigned to a virtual machine (determined by number of cores multiplied by number of sockets).
MaxNumOfVmSockets	Maximum number of sockets per virtual machine	Integer	16	Maximum number of virtual CPU sockets assigned to a virtual machine.
MaxRerunVmOnVdsCount	Maximum virtual machine rerun attempts on a host	Integer	3	Maximum number of attempts to start a virtual machine on a host before an error ("unable to start VM") is reported.
MaxStorageVdsDelayCheckSec	Max delay for check of domain in seconds	Integer	5	Maximum amount of seconds to wait for storage domain status to be returned before reporting an error.
MaxStorageVdsTimeoutCheckSec	Maximum timeout for last check of domain in seconds	Integer	30	When monitoring storage, vdsmd on the hosts reports a "lastCheck" value for each domain. This setting is used to decide whether the last check happened too long ago and domain is considered in error.
MaxVDSMemOverCommit	Max Host Memory Over-Commit (%) for virtual desktops load	Integer	200	The percentage of memory overcommit permitted to occur when using virtual

				desktop loads.
MaxVdsMemOverCommitForServers	Maximum Host Memory Over-Commit (%) for Virtual Servers load	Integer	150	The percentage of memory overcommit permitted to occur when using virtual server loads.
MaxVdsNameLength	Max VDS name length	Integer	255	Maximum name length for a Hypervisor host.
MaxVmNameLengthNonWindows	Maximum virtual machine name length for non-Windows operating system	Integer	64	Maximum name length for a non-Windows virtual machine.
MaxVmNameLengthWindows	Maximum name length in Windows	Integer	15	Maximum name length for Windows virtual machine (limitation imposed by Windows hostnames).
MaxVmsInPool	Max virtual machines in pool	Integer	1000	Maximum number of virtual machines in a single data center.
VmPoolMaxSubsequentFailures	Maximum number of subsequent VM creation failures before giving up	Integer	3	The maximum number of subsequent failed virtual machine creations that can occur in a virtual machine pool before the operation is stopped.
NumberOfFailedRunsOnVds	Number of Failed Runs on Host	Integer	3	Number of attempts to run virtual machines on hosts before setting host status to "Error".
NumberOfVmsForTopSizeVms	Number of virtual machines with highest disk size to display	Integer	10	Number of virtual machines to display in the storage domain's virtual machine tab. Will display this amount of virtual machines, sorted by the most storage space per used virtual machine.
NumberVmRefreshesBeforeSave	Number of Virtual Machine Data Refreshes Before Saving to Database	Integer	5	The number of host monitor iterations between refreshing virtual machines from VDSM (determines if virtual machines should be refreshed one upon each iteration)
oVirtISOsRepositoryPath	The RHEV-H installation files path	String	/usr/share/rhev-hypervisor	The location of RHEV-H ISO images used for upgrading Hypervisor hosts.
ProductKey2003	Product Key (for Windows 2003)	String	-	Windows serial key to be used with sysprepped virtual machines created from a template.
ProductKey2003x64	Product Key (for Windows 2003 x64)	String	-	Windows serial key to be used with sysprepped virtual machines created from a template.
ProductKey2008	Product Key (for Windows 2008)	String	-	Windows serial key to be used with sysprepped virtual machines created from a template.
ProductKey2008R2	Product Key (for Windows 2008 R2)	String	-	Windows serial key to be used with sysprepped virtual machines created from a template.
ProductKey2008x64	Product Key (for Windows 2008 x64)	String	-	Windows serial key to be used with sysprepped virtual machines created from a template.
ProductKey	Product Key (for Windows XP)	String	-	Windows serial key to be used with sysprepped virtual machines

ProductKeyWindow7	Product Key (for Windows 7)	String	-	created from a template. Windows serial key to be used with sysprepped virtual machines created from a template.
ProductKeyWindow7x64	Product Key (for Windows 7 x64)	String	-	Windows serial key to be used with sysprepped virtual machines created from a template.
ProductRPMVersion	Red Hat Enterprise Virtualization Manager RPM Version	String	Automatically populated	The PRM version of the currently used rhvm package.
RDPLoginWithFQDN	Connect to RDP console with Fully Qualified User-Name (user@domain)	String	true	-
SANWipeAfterDelete	Initializing disk image is more secure but it can be time consuming and I/O intensive depending on the size of the image	String	false	Represents the default value (checked/unchecked) of the "wipe after delete" check box in the "New Virtual Disk" dialog if the newly-created Disk belongs to a SAN (FC/SCSI) Data Center (i.e. about to be created on a SAN storage domain).
SearchResultsLimit	Max Quantity of Search Results	Integer	100	The number of results to return for search queries if no specific figure is given in the web administration portal or REST.
ServerRebootTimeout	Host Reboot Timeout (in seconds)	Integer	300	Wait this amount of seconds when a host is rebooted or fenced, before determining that the host is <b>Non Responsive</b> . Can be increased for hosts that take longer to reboot.
SpiceProxyDefault	The address of the SPICE Proxy.	String	none	When this key is set to a value, the SPICE proxy is activated (turned on). When this key is not set to a value, the SPICE proxy is not activated (turned off).
SpiceReleaseCursorKeys	Keyboard keys combination that causes the mouse cursor to be released from its grab on SPICE	String	Shift+F12	-
SpiceSecureChannels	SPICE Secure Channels	String	smain, sinputs, scursor, splayback, srecord, sdisplay, susbredit, smartcard	Which SPICE channels should be secured with SSL.
SpiceToggleFullscreenKeys	Keyboard keys combination that toggles the full-screen state of SPICE client window	String	Shift+F11	-
SpiceUsbAutoShare	Enable USB devices sharing by default in SPICE	String	true	Represents the default value (checked/unchecked) of the "Enable USB Auto-Share" check box in the SPICE Console Options dialog in the User Portal.
SpmCommandFailOverRetries	Number of retries to failover the Storage Pool Manager on failed commands	Integer	3	Number of SPM selection failover retries. In case an SPM command fails, backend performs a failover - it selects a new SPM and re-runs the command.
SPMFailOverAttempts	Number of attempts to	Integer	3	When monitoring a Storage Pool, if

	connect to the Storage Pool Manager before Failover				the current SPM fails, failover does not happen immediately (see description of <code>SpmCommandFailOverRetries</code> ). This setting defines the number of retries before deciding that the current SPM is down and a failover is required.
<code>SpmVcpuConsumption</code>	The CPI consumption of SPM embodied as number of VCPUs on the Host	Integer	1		When a host is the SPM, it is considered to be using this amount of extra virtual CPUs, to make up for the overhead that SPM operations generate.
<code>SSHInactivityTimeoutSeconds</code>	SSH Inactivity Timeout (in seconds)	Integer	600		The maximum amount of time backend allows for an SSH session to remote hosts. After this timeout the session is killed.
<code>SSHInactivityHardTimeoutSeconds</code>	SSH Inactivity Hard Timeout (in seconds)	Integer	-		
<code>NumberOfUSBSlots</code>	Number of USB slots in VMs with native USB support	Integer	4		-
<code>SSLEnabled</code>	SPICE SSL Enabled	String	true		Whether SPICE Secure channels should be SSL encrypted.
<code>StorageDomainFailureTimeoutInMinutes</code>	Storage Domain failure timeout	Integer	5		Defines the amount of time taken before deciding domain is problematic, starting at the first failure reported by VDSM (in minutes).
<code>StoragePoolRefreshTimeInSeconds</code>	Storage Pool Manager Polling Rate (in seconds)	Integer	10		Storage Pool monitoring frequency.
<code>SysPrep2K3Path</code>	Path to a Windows 2003 machine sysprep file	string		<code>/etc/ovirt-engine/sysprep/sysprep.2k3</code>	Path to the operating system specific sysprep file template.
<code>SysPrep2K8Path</code>	Path to a Windows 2008 machine sysprep file	string		<code>/etc/ovirt-engine/sysprep/sysprep.2k8x86</code>	Path to the operating system specific sysprep file template.
<code>SysPrep2K8R2Path</code>	Path to a Windows 2008 R2 machine sysprep file	string		<code>/etc/ovirt-engine/sysprep/sysprep.2k8</code>	Path to the operating system specific sysprep file template.
<code>SysPrep2K8x64Path</code>	Path to a Windows 2008 machine sysprep file	string		<code>/etc/ovirt-engine/sysprep/sysprep.2k8</code>	Path to the operating system specific sysprep file template.
<code>SysPrepWindows7Path</code>	Path to a Windows 7 machine sysprep file	string		<code>/etc/ovirt-engine/sysprep/sysprep.w7</code>	Path to the operating system specific sysprep file template.
<code>SysPrepWindows7x64Path</code>	Path to a Windows 7 x64 machine sysprep file	string		<code>/etc/ovirt-engine/sysprep/sysprep.w7x64</code>	Path to the operating system specific sysprep file template.
<code>SysPrepWindows8Path</code>	Path to a Windows 8 machine Sys-Prep file	String		<code>/etc/ovirt-engine/sysprep/sysprep.w8</code>	Path to the operating system specific sysprep file template.
<code>SysPrepWindows8x64Path</code>	Path to a Windows 8 x64 machine Sys-Prep file	String		<code>/etc/ovirt-engine/sysprep/sysprep.w8x64</code>	Path to the operating system specific sysprep file template.
<code>SysPrepWindows2012x64Path</code>	Path to a Windows 2012 x64 machine Sys-Prep file	String		<code>/etc/ovirt-engine/sysprep/sysprep.2k12x64</code>	Path to the operating system specific sysprep file template.
<code>SysPrepXPPath</code>	Path to a Windows XP machine sysprep file	string		<code>/etc/ovirt-engine/sysprep/sysprep.xp</code>	Path to the operating system specific sysprep file template.
<code>TimeoutToResetVdsInSeconds</code>	Communication timeout in seconds before attempting reset	Integer	60		The amount of time a host is nonresponsive before a fence command is issued. This is used in conjunction with <b><code>VDSAttemptsToResetCount</code></b> .

TimeToReduceFailedRunOnVdsInMinutes	Time to Reduce Failed Run on Host (in minutes)	Integer	30	The amount of time that the host will be in Error status after failing to run virtual machines.
UserDefinedVMPProperties	User-defined virtual machine properties	String	Custom properties set by user	Mostly used with VDSM hooks.
UserRefreshRate	Refresh Rate of Users' Data from Active Directory (in seconds)	Integer	3600	How often the directory server is polled for user account updates.
UtilizationThresholdInPercent	The Utilization Threshold (in percent)	Integer	80	In load balancing, this is a default value used to calculate the maximum CPU limit to determine if the host is over-utilized. This is the percent of the value that the user set in high-utilization in the cluster.
ValidNumOfMonitors	Valid Numbers of Monitors	Integer	1,2,4	Number of monitors available for SPICE-enabled virtual machines.
VdcVersion	Red Hat Enterprise Virtualization Manager Version	String	Automatically set to the current version of RHEV-M	-
VDSAttemptsToResetCount	Number of attempts to communicate with Host before trying to reset	Integer	2	The amount of times to retry communications with a host before a fence command is issued. Used in conjunction with <b>TimeoutToResetVdsInSeconds</b> .
VdsLoadBalancingIntervalInMinutes	Host Load Balancing Interval (in minutes)	Integer	1	The interval between running the virtual machines' load balancer in minutes (also defines the first invocation of the load balancer).
VdsRecoveryTimeoutInMinutes	Host Timeout when Recovering (in minutes)	Integer	3	When VDSM fails/restarts, it can sometimes be in recovering mode (VDSM reports "initializing" or "recovering from reports").
VdsRefreshRate	Time interval in seconds to poll a Host's status	Integer	2	How often a Hypervisor host's status is checked.
vdsTimeout	Host Control Communication Timeout (in seconds)	Integer	180	Timeout for a VDSM call - the time engine will wait for sync call to VDSM.
VMGracefulShutdownMessage	Message displayed in Virtual Machine when Virtual Machine is being shut down from Red Hat Enterprise Virtualization Manager	String	The Red Hat Enterprise Virtualization Manager is shutting down this Virtual Machine	-
VMMMinMemorySizeInMB	Minimal memory size of virtual machine in MB	Integer	256	-
VncKeyboardLayout	Keyboard Layout configuration for VNC	String	en-us	Possible values: ar, da, de-ch, en-us, et, fo, fr-be, fr-ch, hu, it, li, mk, nl, no, pt, ru, sv, tr, de-en-gb, es, fi, fr, fr-ca, hr, is, ja, lv, nl-be, pl, pt-br, sl, th.
WaitForVdsInitInSec	Wait for a host to complete init in SPM selection	Integer	60	This is a timeout for initializing host as in <b>VdsRecoveryTimeoutInMinutes</b> , but this timeout is shorter and is used during the SPM selection algorithm. If the selected host is initialized, wait for it to recover.
FenceQuietTimeBetweenOperations	Quiet time between Power	Integer	180	The minimum time in seconds

InSec	Management operations in seconds			between two power management operations activated manually by a user.
MaxAuditLogMessageLength	Maximum length of an Audit Log message	Integer	10000	-
SysPrepDefaultUser	Default sysprep user name	String	-	This user is used if the domain for sysprep is unknown or no domain is specified.
SysPrepDefaultPassword	Default SysPrep user password	Password	Empty	This password is used if the domain for sysprep is unknown or no domain is specified.
UserSessionTimeoutInterval	Session timeout interval in minutes	Integer	30	User session timeout. Global for all types of access - User Portal/Admin Portal/Web Admin/API.
AdminPassword	admin user password	Password	*****	Password of admin user (used if no directory service is used for authentication).
IPTablesConfig	Iptables configuration used to autoconfigure the Manager.	String		The complete set of iptables rules that are used when automatic firewall configuration is selected during rhvm-setup
OvirtIsoPrefix	oVirt ISOs files prefix	String	rhvh	
OvirtInitialSupportedIsoVersion	oVirt node initial Supported ISO Version	String	5.8	
VdsLocalDisksLowFreeSpace	Amount of free disk space on a host local storage domain that should be considered low, in MB.	Integer	1000	Setting this value lower than the default of 1000MB reduces the time available to add additional space to your data domains before virtual machine performance is affected. If you have many virtual machines, generating or receiving data, it may make sense to set this value higher.
VdsLocalDisksCriticallyLowFreeSpace	Amount of free disk space on a host local storage domain that should be considered critically low, in MB.	Integer	500	Setting this value lower than the default of 500MB reduces the time between when critical disk shortage messages begin being displayed and when virtual machine performance is affected. If you have many virtual machines, generating or receiving data quickly, you might find that the default value is too low, and does not provide enough time to add more storage.
AllowDuplicateMacAddresses	Enable duplicate MAC address for VM network interface	String	false	If enabled, allows that the same MAC address be set explicitly on several virtual NICs. Otherwise, setting a MAC address that is already in use on another virtual NIC is prohibited.
JobCleanupRateInMinutes	Frequency of jobs clean-up process	Integer	10	
SucceededJobCleanupTimeInMinutes	Time to keep successfully ended jobs	Integer	10	
FailedJobCleanupTimeInMinutes	Time to keep failed jobs	Integer	60	



VmPoolMonitorIntervalInMinutes	Interval in minutes for monitoring number of Prestarted Vms in Vm Pools	Integer	5	
VmPoolMonitorBatchSize	Maximum number of Vms that the Vm Pool Monitor will attempt to Prestart in a single cycle	Integer	5	
NetworkConnectivityCheckTimeoutInSeconds	The time to wait before rolling back network changes in case the engine losses connectivity with the host in seconds	Integer	120	
AllowClusterWithVirtGlusterEnabled	Allows to create a Cluster with both Virt and Gluster services enabled	String	false	If enabled, the user can create a cluster with both Virt and Gluster support or one of them, otherwise the user cannot create a cluster that supports both.
EnableMACAntiSpoofingFilterRules	Indicates if Network Filtering should be enabled or not	String	v3.0: false v3.1: false v3.2: true	If enabled, MAC anti-spoofing rules are set on each virtual NIC, enforcing that the ethernet frames this NIC sends have the source MAC that is assigned to it in the engine.
EnableHostTimeDrift	Indicates if host time-drift validation is enabled	String	false	If time drift validation is enabled, the Manager will require that host system time be within a given variation of the Manager system time. The allowed difference is set by HostTimeDriftInSec
EngineMode	Engine working mode	String	Active	
HostTimeDriftInSec	Allowed time drift between any Host and Engine	Integer	300	

[Report a bug](#)

## 15.3. Uploading Virtual Machine Images with the Image Uploader Tool

### 15.3.1. Virtual Machine Image Uploader

Using the `rhev-image-uploader` command, you can list export storage domains and upload virtual machines in Open Virtualization Format (OVF) to an export storage domain and have them automatically recognized in the Red Hat Enterprise Virtualization Manager. The tool only supports gzip compressed OVF files created by Red Hat Enterprise Virtualization.

The image uploader makes creating distributable virtual machine images practical.

The archive should contain images and master directories that are in the following format:

```
-- images
|   |-- [Image Group UUID]
|   |   |-- [Image UUID (this is the disk image)]
|   |   |-- [Image UUID (this is the disk image)].meta
|   |-- master
|   |   |-- vms
|   |       |-- [UUID]
|   |       |-- [UUID].ovf
```

[Report a bug](#)

### 15.3.2. Syntax for the `rhev-image-uploader` Command

The basic syntax for the image uploader command is:

```
rhev-image-uploader [options] list
rhev-image-uploader [options] upload [file].[file]...[file]
```

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

- » The **list** parameter lists the valid export storage domains available for image uploads.
- » The **upload** parameter uploads selected image file(s) to the specified image storage domain.

The image uploader command requires either the **list** or **upload** parameter be included for basic usage. The **upload** parameter requires a minimum of one local file name to upload.

There are numerous parameters to further refine the **rhev-image-uploader** command. You can set defaults for any of these in the `/etc/ovirt-engine/imageuploader.conf` file.

#### General Options

**-h, --help**

Displays command usage information and returns to prompt.

**--conf-file=PATH**

Sets **PATH** as the configuration file the tool is to use. The default is **etc/ovirt-engine/imageuploader.conf**.

**--log-file=PATH**

Sets **PATH** as the specific file name the command should use for the log output.

**--quiet**

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

**-v, --verbose**

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

**-f, --force**

Force mode is necessary when the source file being uploaded has an identical file name as an existing file at the destination; it forces the existing file to be overwritten. Force mode is off by default.

## Red Hat Enterprise Virtualization Manager Options

**-u USER, --user=USER**

Sets the user associated with the file to be uploaded. The **USER** is specified in the format **user@domain**, where **user** is the user name and **domain** is the directory services domain in use. The user must exist in directory services and be known to the Red Hat Enterprise Virtualization Manager.

**-r FQDN, --rhevm=FQDN**

Sets the fully qualified domain name of the Red Hat Enterprise Virtualization Manager server from which to upload images, where **FQDN** is replaced by the fully qualified domain name of the Manager. It is assumed that the image uploader is being run on the same client machine as the Red Hat Enterprise Virtualization Manager; the default value is **localhost:443**.

## Export Storage Domain Options

These options specify the export domain to which files are uploaded. They are alternatives; do not use these parameters together.

**-e, --export-domain=EXPORT\_DOMAIN**

Sets the storage domain **EXPORT\_DOMAIN** as the destination for uploads.

**-n, --nfs-server=NFSSERVER**

Sets the NFS path **NFSSERVER** as the destination for uploads.

**-i, --ovf-id**

Use this option if you do not want to update the UUID of the image. By default, the tool will generate a new UUID for the image. This ensures that there is no conflict between the id of the incoming image and those already in the environment.

**-d, --disk-instance-id**

Use this option if you do not want to rename the instance ID for each disk (i.e. InstanceId) in the image. By default, this tool will generate new UUIDs for disks within the image to be imported. This ensures that there are no conflicts between the disks on the imported image and those within the environment.

**-m, --mac-address**

Use this option if you do not want to remove the network components from the image that will be imported. By default, this tool will remove any network interface cards from the image to prevent conflicts with network cards on other virtual machines within the environment. Once the image has been imported, use the Administration Portal to add network interface cards back and the Manager will ensure that there are no MAC address conflicts.

**-N NEW\_IMAGE\_NAME, --name=NEW\_IMAGE\_NAME**

Supply this option if you want to rename the image.

[Report a bug](#)

## 15.3.3. Creating an OVF Archive That Is Compatible with the Image Uploader

### Summary

You can create files that can be uploaded using the **rhevm-image-uploader** tool.

#### Procedure 15.1. Creating an OVF Archive That is Compatible with the Image Uploader

1. Use the Manager to create an empty export domain. An empty export domain makes it easy to see which directory contains your virtual machine.
2. Export your virtual machine to the empty export domain you just created.
3. Log in to the storage server that serves as the export domain, find the root of the NFS share and change to the subdirectory under that mount point. You started with a new export domain, there is only one directory under the exported directory. It contains the **images/** and **master/** directories.
4. Run the **tar -zcvf my.ovf images/ master/** command to create the tar/gzip ovf archive.
5. Anyone you give the resulting ovf file to (in this example, called **my.ovf**) can import it to Red Hat Enterprise Virtualization Manager using the **rhevm-image-uploader** command.

### Result

You have created a compressed OVF image file that can be distributed. Anyone you give it to can use the **rhevm-image-uploader** command to upload your image into their Red Hat Enterprise Virtualization environment.

[Report a bug](#)

### 15.3.4. Basic rhevm-image-uploader Usage Examples

Here is an example of how to use **rhevm-image-uploader** to list storage domains:

#### Example 15.10. Uploading a file Using the rhevm-image-uploader Tool

```
# rhevm-image-uploader list
Please provide the REST API username for RHEV-M: admin@internal
Please provide the REST API password for the admin@internal RHEV-M user:
*****
Export Storage Domain Name | Datacenter | Export Domain Status
myexportdom                | Myowndc   | active
```

To upload an Open Virtualization Format (ovf) file, you need to enter an NFS server name (**-n NFSSEVER**) or export domain (**-e EXPORT\_STORAGE\_DOMAIN**) and the name of the **.ovf** file:

```
# rhevm-image-uploader -e myexportdom upload myrhe16.ovf
Please provide the REST API username for RHEV-M: admin@internal
Please provide the REST API password for the admin@internal RHEV-M user:
*****
```

[Report a bug](#)

## 15.4. Editing USB Filters with the USB Filter Editor

### 15.4.1. Installing the USB Filter Editor

#### Summary

The USB Filter Editor is a Windows tool used to configure the **usbfilter.txt** policy file. The policy rules defined in this file allow or deny the pass-through of specific USB devices from client machines to virtual machines managed using the Red Hat Enterprise Virtualization Manager. The policy file resides on the Red Hat Enterprise Virtualization Manager in the following location:

**/etc/ovirt-engine/usbfilter.txt**

Changes to USB filter policies do not take effect unless the **ovirt-engine** service on the Red Hat Enterprise Virtualization Manager server is restarted.

Download the **USBFilterEditor.msi** file from the **Red Hat Enterprise Virtualization Manager (version.number x86\_64)** channel of the Red Hat Network.

#### Note

Refer to the *Red Hat Enterprise Virtualization Manager Release Notes* for specific channel names current to your system.

#### Procedure 15.2. Installing the USB Filter Editor

1. On a Windows machine, launch the **USBFilterEditor.msi** installer obtained from Red Hat Network.
2. Follow the steps of the installation wizard. Unless otherwise specified, the USB Filter Editor will be installed by default in either **C:\Program Files\RedHat\USB Filter Editor** or **C:\Program Files(x86)\RedHat\USB Filter Editor** depending on your version of Windows.
3. A USB Filter Editor shortcut icon is created on your desktop.

#### Important

Use a Secure Copy (SCP) client to import and export filter policies from the Red Hat Enterprise Virtualization Manager. A Secure Copy tool for Windows machines is WinSCP (<http://winscp.net>).

#### Result

The default USB device policy provides virtual machines with basic access to USB devices; update the policy to allow the use of additional USB devices.

[Report a bug](#)

### 15.4.2. The USB Filter Editor Interface

- » Double-click the USB Filter Editor shortcut icon on your desktop.

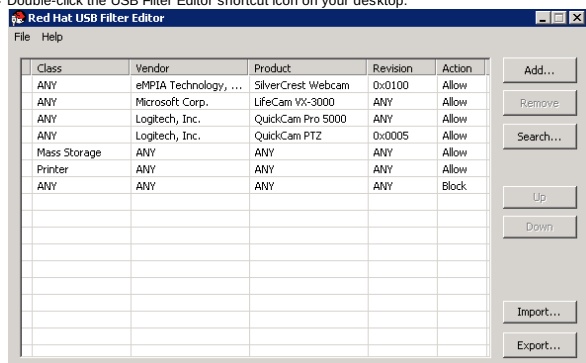


Figure 15.1. Red Hat USB Filter Editor

The **Red Hat USB Filter Editor** interface displays the **Class**, **Vendor**, **Product**, **Revision**, and **Action** for each USB device. Permitted USB devices are set to **Allow** in the **Action** column; prohibited devices are set to **Block**.

Table 15.2. USB Editor Fields

Name	Description
<b>Class</b>	Type of USB device; for example, printers, mass storage controllers.
<b>Vendor</b>	The manufacturer of the selected type of device.
<b>Product</b>	The specific USB device model.
<b>Revision</b>	The revision of the product.
<b>Action</b>	Allow or block the specified device.

The USB device policy rules are processed in their listed order. Use the **Up** and **Down** buttons to move devices higher or lower in the list. The universal **Block** rule needs to remain as the lowest entry to ensure all USB devices are denied unless explicitly allowed in the USB Filter Editor.

[Report a bug](#)

### 15.4.3. Adding a USB policy

#### Summary

Add a USB policy to the USB Filter Editor.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

#### Procedure 15.3. Adding a USB Policy

1. Click the **Add** button. The **Edit USB Criteria** window opens:

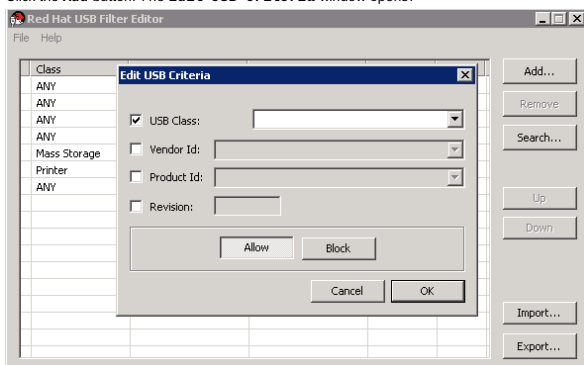


Figure 15.2. Edit USB Criteria

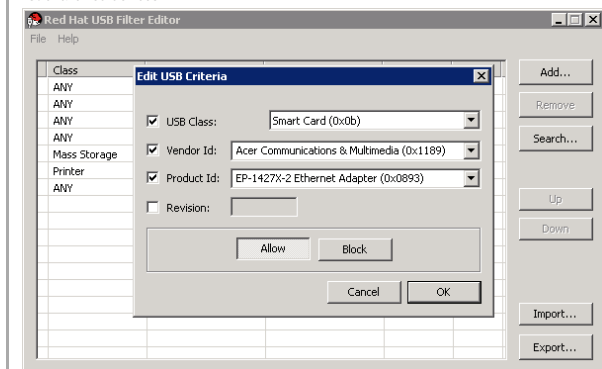
2. Use the **USB Class**, **Vendor ID**, **Product ID**, and **Revision** check boxes and lists to specify the device.

Click the **Allow** button to permit virtual machines use of the USB device; click the **Block** button to prohibit the USB device from virtual machines.

Click **OK** to add the selected filter rule to the list and close the window.

#### Example 15.11. Adding a Device

The following is an example of how to add USB Class **Smartcard**, device **EP-1427X-2 Ethernet Adapter**, from manufacturer **Acer Communications & Multimedia** to the list of allowed devices.



3. Click **File** → **Save** to save the changes.

#### Result

You have added a USB policy to the USB Filter Editor. USB filter policies need to be exported to the Red Hat Enterprise Virtualization Manager to take effect.

[Report a bug](#)

### 15.4.4. Removing a USB Policy

#### Summary

Remove a USB policy from the USB Filter Editor.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

#### Procedure 15.4. Removing a USB Policy

1. Select the policy to be removed.

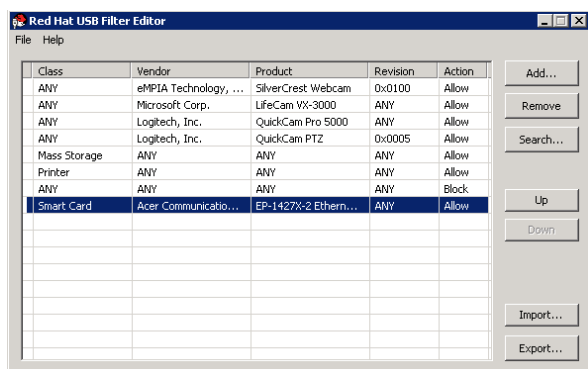


Figure 15.3. Select USB Policy

- Click **Remove**. A message displays prompting you to confirm that you want to remove the policy.

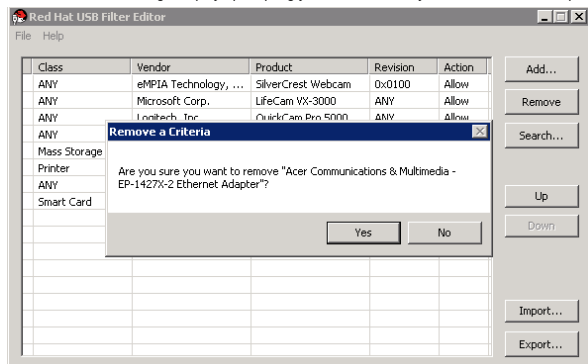


Figure 15.4. Edit USB Criteria

- Click **Yes** to confirm that you want to remove the policy.
- Click **File** → **Save** to save the changes.

#### Result

You have removed a USB policy from the USB Filter Editor. USB filter policies need to be exported to the Red Hat Enterprise Virtualization Manager to take effect.

[Report a bug](#)

### 15.4.5. Searching for USB Device Policies

#### Summary

Search for attached USB devices to either allow or block them in the USB Filter Editor.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

#### Procedure 15.5. Searching for USB Device Policies

- Click **Search**. The **Attached USB Devices** window displays a list of all the attached devices.

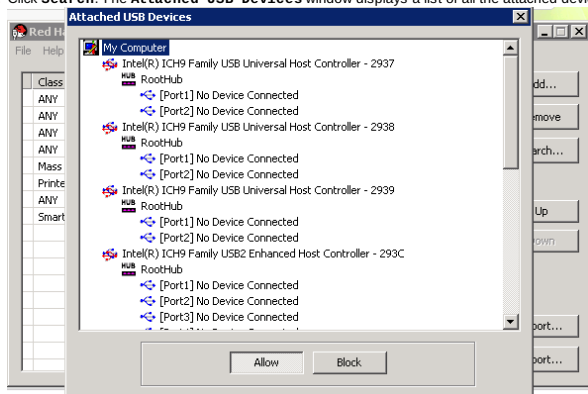


Figure 15.5. Attached USB Devices

- Select the device and click **Allow** or **Block** as appropriate. Double-click the selected device to close the window. A policy rule for the device is added to the list.
- Use the **Up** and **Down** buttons to change the position of the new policy rule in the list.
- Click **File** → **Save** to save the changes.

#### Result

You have searched the attached USB devices. USB filter policies need to be exported to the Red Hat Enterprise Virtualization Manager to take effect.

[Report a bug](#)

### 15.4.6. Exporting a USB Policy

**Summary**

USB device policy changes need to be exported and uploaded to the Red Hat Enterprise Virtualization Manager for the updated policy to take effect. Upload the policy and restart the **ovirt-engine** service.

Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.

**Procedure 15.6. Exporting a USB Policy**

1. Click **Export**; the **Save As** window opens.
2. Save the file with a file name of **usbfilter.txt**.
3. Using a Secure Copy client, such as WinSCP, upload the **usbfilter.txt** file to the server running Red Hat Enterprise Virtualization Manager. The file must be placed in the following directory on the server:  
**/etc/ovirt-engine/**
4. As the **root** user on the server running Red Hat Enterprise Virtualization Manager, restart the **ovirt-engine** service.

```
# service ovirt-engine restart
```

**Result**

The USB device policy will now be implemented on virtual machines running in the Red Hat Enterprise Virtualization environment.

[Report a bug](#)

**15.4.7. Importing a USB Policy****Summary**

An existing USB device policy must be downloaded and imported into the USB Filter Editor before you can edit it.

**Procedure 15.7. Importing a USB Policy**

1. 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:  
**/etc/ovirt-engine/**
2. Double-click the USB Filter Editor shortcut icon on your desktop to open the editor.
3. Click **Import** to open the **Open** window.
4. Open the **usbfilter.txt** file that was downloaded from the server.

**Result**

You are able to edit the USB device policy in the USB Filter Editor.

[Report a bug](#)

**15.5. Collecting Logs with the Log Collector Tool****15.5.1. Log Collector**

A log collection tool is included in the Red Hat Enterprise Virtualization Manager. This allows you to easily collect relevant logs from across the Red Hat Enterprise Virtualization environment when requesting support.

The log collection command is **rhev-log-collector**. You are required to log in as the **root** user and provide the administration credentials for the Red Hat Enterprise Virtualization environment. The **rhev-log-collector -h** command displays usage information, including a list of all valid options for the **rhev-log-collector** command.

[Report a bug](#)

**15.5.2. Syntax for rhev-log-collector Command**

The basic syntax for the log collector command is:

```
rhev-log-collector [options] list [all, clusters, datacenters]
rhev-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 able to filter the log collection based on the listed objects.
- » The **collect** parameter performs log collection from the Red Hat Enterprise Virtualization Manager. The collected logs are placed in an archive file under the **/tmp/logcollector** directory. The **rhev-log-collector** command assigns each log a specific file name.

Unless another parameter is specified, the default action is to list the available hosts together with the data center and cluster to which they belong. You will be prompted to enter user names and passwords to retrieve certain logs.

There are numerous parameters to further refine the **rhev-log-collector** command.

**General options****--version**

Displays the version number of the command in use and returns to prompt.

**-h, --help**

Displays command usage information and returns to prompt.

**--conf-file=PATH**

Sets **PATH** as the configuration file the tool is to use.

**--local-tmp=PATH**

Sets **PATH** as the directory in which logs are saved. The default directory is **/tmp/logcollector**.

**--ticket-number=TICKET**

Sets **TICKET** as the ticket, or case number, to associate with the SOS report.

**--upload=FTP\_SERVER**

Sets **FTP\_SERVER** as the destination for retrieved logs to be sent using FTP. Do not use this option unless advised to by a Red Hat support representative.

**--log-file=PATH**

Sets **PATH** as the specific file name the command should use for the log output.

**--quiet**

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

**-v, --verbose**

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

## Red Hat Enterprise Virtualization Manager Options

These options filter the log collection and specify authentication details for the Red Hat Enterprise Virtualization Manager.

These parameters can be combined for specific commands. For example, **rhev-log-collector --user=admin@internal --cluster ClusterA,ClusterB --hosts "SalesHost"\*** specifies the user as **admin@internal** and limits the log collection to only **SalesHost** hosts in clusters **A** and **B**.

**--no-hypervisors**

Omits virtualization hosts from the log collection.

**-u USER, --user=USER**

Sets the user name for login. The **USER** is specified in the format **user@domain**, where **user** is the user name and **domain** is the directory services domain in use. The user must exist in directory services and be known to the Red Hat Enterprise Virtualization Manager.

**-r FQDN, --rhev=FQDN**

Sets the fully qualified domain name of the Red Hat Enterprise Virtualization Manager server from which to collect logs, where **FQDN** is replaced by the fully qualified domain name of the Manager. It is assumed that the log collector is being run on the same local host as the Red Hat Enterprise Virtualization Manager; the default value is **localhost**.

**-c CLUSTER, --cluster=CLUSTER**

Collects logs from the virtualization hosts in the nominated **CLUSTER** in addition to logs from the Red Hat Enterprise Virtualization Manager. The cluster(s) for inclusion must be specified in a comma-separated list of cluster names or match patterns.

**-d DATACENTER, --data-center=DATACENTER**

Collects logs from the virtualization hosts in the nominated **DATACENTER** in addition to logs from the Red Hat Enterprise Virtualization Manager. The data center(s) for inclusion must be specified in a comma-separated list of data center names or match patterns.

**-H HOSTS\_LIST, --hosts=HOSTS\_LIST**

Collects logs from the virtualization hosts in the nominated **HOSTS\_LIST** in addition to logs from the Red Hat Enterprise Virtualization Manager. The hosts for inclusion must be specified in a comma-separated list of host names, fully qualified domain names, or IP addresses. Match patterns are also valid.

## SOS Report Options

The log collector uses the JBoss SOS plugin. Use the following options to activate data collection from the JMX console.

**--jboss-home=JBOSS\_HOME**

JBoss installation directory path. The default is **/var/lib/jbossas**.

**--java-home=JAVA\_HOME**

Java installation directory path. The default is **/usr/lib/jvm/java**.

**--jboss-profile=JBOSS\_PROFILE**

Displays a quoted and space-separated list of server profiles; limits log collection to specified profiles. The default is **'rhev-slimmed'**.

**--enable-jmx**

Enables the collection of run-time metrics from Red Hat Enterprise Virtualization's JBoss JMX interface.

**--jboss-user=JBOSS\_USER**

User with permissions to invoke JBoss JMX. The default is **admin**.

**--jboss-logsize=LOG\_SIZE**

Maximum size in MB for the retrieved log files.

**--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**.

**--jboss-appxml=XML\_LIST**

Displays a quoted and space-separated list of applications with XML descriptions to be

retrieved. Default is **all**.

#### SSH Configuration

**--ssh-port=PORT**

Sets **PORT** as the port to use for SSH connections with virtualization hosts.

**-k KEYFILE, --key-file=KEYFILE**

Sets **KEYFILE** as the public SSH key to be used for accessing the virtualization hosts.

**--max-connections=MAX\_CONNECTIONS**

Sets **MAX\_CONNECTIONS** as the maximum concurrent SSH connections for logs from virtualization hosts. The default is **10**.

#### PostgreSQL Database Options

The **pg-pass** parameter includes the Red Hat Enterprise Virtualization Manager database in the log. The database user name and database name must be specified if they have been changed from the default values.

Use the **pg-dbhost** parameter if the database is not on the local host. Use the optional **pg-host-key** parameter to collect remote logs. The PostgreSQL SOS plugin must be installed on the database server for remote log collection to be successful.

**--no-postgresql**

Disables collection of database. Database collection is performed by default.

**--pg-user=USER**

Sets **USER** as the user name to use for connections with the database server. The default is **postgres**.

**--pg-dbname=DBNAME**

Sets **DBNAME** as the database name to use for connections with the database server. The default is **rhev**.

**--pg-dbhost=DBHOST**

Sets **DBHOST** as the host name for the database server. The default is **localhost**.

**--pg-host-key=KEYFILE**

Sets **KEYFILE** as the public identity file (private key) for the database server. This value is not set by default; it is required only where the database does not exist on the local host.

[Report a bug](#)

### 15.5.3. 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.

#### Example 15.12. Log Collector Usage

```
# rhevm-log-collector
Please provide the username for rhevm (CTRL+D to abort):
admin@directory.demo.redhat.com
Please provide the password for rhevm (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
# rhevm-log-collector collect
Please provide the username for rhevm (CTRL+D to abort):
admin@directory.demo.redhat.com
Please provide the password for rhevm (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-26110804121320-ce2a.tar.xz.
The MD5 for this file is 6d741b78925998caff29020df2b2ce2a and its size is
26.7M
```

[Report a bug](#)

## 15.6. Uploading ISO Files with the ISO Uploader Tool

### 15.6.1. ISO Uploader

The ISO uploader is a tool for uploading ISO images to the ISO storage domain. It is installed as part of the Red Hat Enterprise Virtualization Manager.

The ISO uploader command is **rhev**-**iso**-**uploader**. You are required to log in as the **root** user and provide the administration credentials for the Red Hat Enterprise Virtualization environment. The **rhev**-**iso**-**uploader** **-h** command displays usage information, including a list of all valid options for the **rhev**-**iso**-**uploader** command.

[Report a bug](#)

### 15.6.2. Syntax for rhev-iso-uploader Command

The basic syntax for the ISO uploader command is:



```
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 valid ISO storage domains available for ISO uploads. The Red Hat Enterprise Virtualization Manager sets this list on the local machine upon installation.
- The **upload** parameter uploads single or multiple space-separated ISO files to the specified ISO storage domain. NFS is used as default; SSH is available.

The ISO uploader command requires either the **list** or **upload** parameter be included for basic usage. The **upload** parameter requires a minimum of one local file name to upload.

There are numerous parameters to further refine the **rhevms-iso-uploader** command.

#### General Options

- version**  
Displays the version number of the command in use and returns to prompt.
- h, --help**  
Displays command usage information and returns to prompt.
- conf-file=PATH**  
Sets **PATH** as the configuration file the tool is to use.
- log-file=PATH**  
Sets **PATH** as the specific file name the command should use for the log output.
- quiet**  
Sets quiet mode, reducing console output to a minimum. Quiet mode is off by default.
- v, --verbose**  
Sets verbose mode, providing more console output. Verbose mode is off by default.
- f, --force**  
Force mode is necessary when the source file being uploaded has an identical file name as an existing file at the destination; it forces the existing file to be overwritten. Force mode is off by default.

#### Red Hat Enterprise Virtualization Manager Options

- u USER, --user=USER**  
Sets the user associated with the file to be uploaded. The **USER** is specified in the format **user@domain**, where **user** is the user name and **domain** is the directory services domain in use. The user must exist in directory services and be known to the Red Hat Enterprise Virtualization Manager.
- r FQDN, --rhevms=FQDN**  
Sets the fully qualified domain name of the Red Hat Enterprise Virtualization Manager server from which to upload ISOs, where **FQDN** is replaced by the fully qualified domain name of the Manager. It is assumed that the ISO uploader is being run on the same client machine as the Red Hat Enterprise Virtualization Manager; the default value is **localhost**.

#### ISO Storage Domain Options

These options specify the ISO domain to which files are uploaded. They are alternatives; do not use these parameters together.

- i, --iso-domain=ISODOMAIN**  
Sets the storage domain **ISODOMAIN** as the destination for uploads.
- n, --nfs-server=NFSERVER**  
Sets the NFS path **NFSERVER** as the destination for uploads.

#### Connection Options

The ISO uploader uses NFS as default to upload files. These options specify SSH file transfer instead.

- ssh-user=USER**  
Sets **USER** as the SSH user name to use for the upload.
- ssh-port=PORT**  
Sets **PORT** as the port to use when connecting to SSH.
- k KEYFILE, --key-file=KEYFILE**  
Sets **KEYFILE** as the public key to use for SSH authentication. You will be prompted to enter the password of the user specified with **--ssh-user=USER** if no key is set.

[Report a bug](#)

### 15.6.3. Specifying an NFS Server

Example 15.13. Uploading to an NFS Server

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

[Report a bug](#)

### 15.6.4. Basic ISO Uploader Usage

The example below demonstrates the ISO uploader and the list parameter. The first command lists the available ISO storage domains; the user name is requested because it was not provided in the command. The second command uploads an ISO file over NFS to the listed ISO domain: **ISODomain**.

**Example 15.14. List Domains and Upload Image**

```
# 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):
```

[Report a bug](#)**15.6.5. Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain**

The example below demonstrates the command to upload the `virtio-win.iso`, `virtio-win.vfd`, and `rhev-tools-setup.iso` image files to the `ISODomain`.

**Example 15.15. Uploading the VirtIO and Guest Tool Image Files**

```
# rhvm-iso-uploader --iso-domain=ISODomain upload /usr/share/virtio-
win/virtio-win.iso /usr/share/virtio-win/virtio-win.vfd /usr/share/rhev-
guest-tools-iso/rhev-tools-setup.iso
```

[Report a bug](#)**15.6.6. VirtIO and Guest Tool Image Files**

The `virtio-win` ISO and Virtual Floppy Drive (VFD) images, which contain the VirtIO drivers for Windows virtual machines, and the `rhev-tools-setup` ISO, which contains the Red Hat Enterprise Virtualization Guest Tools for Windows virtual machines, are copied to an ISO storage domain upon installation and configuration of the domain.

These image files provide software that can be installed on virtual machines to improve performance and usability. The most recent `virtio-win` and `rhev-tools-setup` images are referred to by the following symbolic links on the file system of the Red Hat Enterprise Virtualization Manager:

- `/usr/share/virtio-win/virtio-win.iso`
- `/usr/share/virtio-win/virtio-win.vfd`
- `/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso`

These image files must be manually uploaded to ISO storage domains that were not created locally by the installation process. Use the `rhvm-iso-uploader` command to upload these images to your ISO storage domain. Once uploaded, the image files can be attached to and used by virtual machines.

[Report a bug](#)**15.7. Guest Drivers and Agents****15.7.1. Installing Guest Agents and Drivers**

Installing Red Hat Enterprise Virtualization guest agents and drivers on virtual machines provides optimized performance and extra features.

**Installing the agents and drivers on Red Hat Enterprise Linux guests**

All of the drivers are included in the base channel for RHN registered Red Hat Enterprise Linux virtual machines. They can be installed using the `yum install rhvm-guest-agent` command.

Your guest must be subscribed to the **Red Hat Enterprise Virt Agent** channel to install the agents.

In Red Hat Enterprise Linux 5, this channel is labelled `rhel-x86_64-rhev-agent-5-server`. In Red Hat Enterprise Linux 6, the channel is labelled `rhel-x86_64-rhev-agent-6-server`.

**Installing the agents and drivers on Windows guests**

The agents and drivers are installed on Windows virtual machines using the `rhev-tools-setup.iso` disk image. The guest tools ISO is distributed using the Red Hat Network as `rhev-guest-tools-iso.rpm`, an RPM file installed on the Red Hat Enterprise Virtualization Manager.

After installing the Manager, the guest tools ISO can be found at `/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso`. When setting up the Manager, if you have created a local storage share for an ISO storage domain, the ISO file is automatically copied to the ISO storage domain. In this case the ISO image is automatically attached to Windows guests when they are created. Otherwise, the ISO must be manually attached to Windows guests for the tools and agents to be installed.

Updated versions of the ISO file must be manually attached to running Windows virtual machines to install updated versions of the tools and drivers. If the APT service is enabled on virtual machines, the updated ISO files will be automatically attached.

**See Also:**

- [Section 15.7.4, "Subscribing to channels"](#)
- [Section 15.7.2, "Automating Guest Additions on Windows Guests with Red Hat Enterprise Virtualization Application Provisioning Tool\(APT\)"](#)

[Report a bug](#)**15.7.2. Automating Guest Additions on Windows Guests with Red Hat Enterprise Virtualization Application Provisioning Tool(APT)**

Red Hat Enterprise Virtualization Application Provisioning Tool (APT) is a Windows service that can be installed in Windows virtual machines and templates. Attach the guest tools ISO file to your Windows virtual machine and `RHEV-Application Provisioning.exe` automatically runs to install the APT service.

When the APT service is installed on a virtual machine, attached ISO files are automatically scanned. When the service recognizes a valid Red Hat Enterprise Virtualization guest tools ISO, and no other guest tools are installed, the APT service installs the guest tools. If guest tools are installed, and the ISO image contains newer versions of the tools, an upgrade is automatically performed.

When the APT service has successfully installed or upgraded guest tools on a virtual machine, the virtual machine is automatically rebooted.

[Report a bug](#)

### 15.7.3. Red Hat Enterprise Virtualization Guest Drivers and Guest Agents

Red Hat Enterprise Virtualization provides customized drivers and guest tools to use with Windows and Red Hat Enterprise Linux guests. The drivers allow guests to use enhanced virtual devices that perform better than emulated devices; the guest agents facilitate communication between the guest and the Red Hat Enterprise Virtualization Manager.

Table 15.3. Red Hat Enterprise Virtualization Guest Drivers

Driver	Description	Works on
<b>virtio-net</b>	Paravirtualized network driver provides enhanced performance over emulated devices like rtl.	Server and Desktop.
<b>virtio-block</b>	Paravirtualized HDD driver offers increased I/O performance over emulated devices like IDE by optimizing the coordination and communication between the guest and the hypervisor. The driver complements the software implementation of the virtio-device used by the host to play the role of a hardware device.	Server and Desktop.
<b>virtio-serial</b>	Virtio-serial provides support for multiple serial ports. The improved performance is used for fast communication between the guest and the host that avoids network complications. This fast communication is required for the guest agents and for other features such as clipboard copy-paste between the guest and the host and logging.	Server and Desktop.
<b>virtio-balloon</b>	Virtio-balloon is used to control the amount of memory a guest actually accesses. It offers improved memory over-commitment. The balloon drivers are installed for future compatibility but not used by default in Red Hat Enterprise Virtualization 3.1 or higher.	Server and Desktop.
<b>qxl</b>	A paravirtualized display driver reduces CPU usage on the host and provides better performance through reduced network bandwidth on most workloads.	Server and Desktop.

Table 15.4. Red Hat Enterprise Virtualization Guest Agents and Tools

Guest agent/tool	Description	Works on
<b>rhev-guest-agent</b>	Allows the Red Hat Enterprise Virtualization Manager to receive guest internal events and information such as IP address and installed applications. Also allows the Manager to execute specific commands, such as shut down or reboot, on a guest.  On Red Hat Enterprise Linux 6 and higher guests, the <b>rhev-guest-agent</b> installs <b>tuned</b> on your virtual machine and configures it to use an optimized, virtualized-guest profile.	Server and Desktop.
<b>spice-agent</b>	The SPICE agent supports multiple monitors and is responsible for client-mouse-mode support to provide a better user experience and improved responsiveness than the QEMU emulation. Cursor capture is not needed in client-mouse-mode. The SPICE agent reduces bandwidth usage when used over a wide area network by reducing the display level, including color depth, disabling wallpaper, font smoothing, and animation. The SPICE agent enables clipboard support allowing cut and paste operations for both text and images between client and guest, and automatic guest display setting according to client-side settings. On Windows guests, the SPICE agent consists of vdservice and vdagent.	Server and Desktop.
<b>rhev-ss0</b>	An agent that enables users to automatically log in to their virtual machines based on the credentials used to access the Red Hat Enterprise Virtualization Manager.	Desktop.
<b>rhev-usb</b>	A component that contains drivers and services for Legacy USB support (version 3.0 and earlier) on guests. It is needed for accessing a USB device that is plugged into the client machine. <b>RHEV-USB Client</b> is needed on the client side.	Desktop.

[Report a bug](#)

## 15.7.4. Subscribing to channels

### 15.7.4.1. Subscribing to Channels Using Subscription Manager

#### Summary

To install packages signed by Red Hat you must register the target system to Red Hat Network. You can then use an entitlement from your entitlement pool to subscribe the system to channels.

#### Procedure 15.8. Subscribing to Channels Using Subscription Manager

1. Run the **subscription-manager register** command to register the system with Red Hat Network. To complete registration successfully you will need to supply your Red Hat Network **Username** and **Password** when prompted.

```
# subscription-manager register
```

2. **Identify available entitlement pools**

To subscribe the system to channels, you must locate the identifiers for the relevant entitlement pools. Use the **list** action of the **subscription-manager** to find these.

For example, to identify available subscription pools for **Red Hat Enterprise Virtualization** use the command:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Virtualization"
```

3. **Subscribe system to entitlement pools**

Using the pool identifiers located in the previous step, subscribe the system to the required entitlements. When a system is subscribed to an entitlement pool, the system is automatically subscribed to the channels in the entitlement. The main channel is automatically enabled, other channels in the entitlement must be enabled manually. Use the **subscribe** action of the **subscription-manager** command, replacing **POOLID** with one of the pool identifiers each time the command is run:

```
# subscription-manager subscribe --pool=POOLID
```

4. **Enable additional subscription channels**

When a system is subscribed to an entitlement with a main channel and some additional channel, only the main channel is enabled by default. Other channels are available, but disabled. The additional channels must be enabled using the **yum-config-manager** command as the **root** user:

```
# yum-config-manager --enable CHANNEL
```

**Result**

The system is now registered with Red Hat Network and subscribed to the channels required.

[Report a bug](#)

**15.7.4.2. Subscribing to channels using RHN Classic****Summary**

To install packages you must first register the target system to Red Hat Network and subscribe to the software channels containing your packages.

**Procedure 15.9. Subscribing to the channels using RHN Classic**

1. Run the **rhnc** command to register the system with Red Hat Network. 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.

```
# rhnc
```

2. **Subscribe to Required Channels**

You must subscribe the system to the required channels using either the web interface to Red Hat Network or the command line **rhnc-channel** command.

**A. Using the rhnc-channel Command**

Run the **rhnc-channel** command to subscribe the system to each of the required channels. The commands which need to be run are:

```
# rhnc-channel --add --channel=CHANNEL
```

**Important**

If you are not the administrator for the machine as defined in Red Hat Network, or the machine is not registered to Red Hat Network, then use of the **rhnc-channel** command will result in an error:

```
Error communicating with server. The message was:
Error Class Code: 37
Error Class Info: You are not allowed to perform administrative tasks
on this system.
Explanation:
    An error has occurred while processing your request. If this
problem
persists please enter a bug report at bugzilla.redhat.com.
If you choose to submit the bug report, please be sure to
include
    details of what you were trying to do when this error occurred
and
    details on how to reproduce this problem.
```

If you encounter this error when using **rhnc-channel** then to add the channel to the system you must use the web user interface.

**B. Using the Web Interface to Red Hat Network**

To add a channel subscription to a system from the web interface:

- a. Log on to Red Hat Network (<http://rhn.redhat.com>).
- b. Move the mouse cursor over the **Subscriptions** link at the top of the screen, and then click the **Registered Systems** link in the menu that appears.
- c. Select the system to which you are adding channels from the list presented on the screen, by clicking the name of the system.
- d. Click **Alter Channel Subscriptions** in the **Subscribed Channels** section of the screen.
- e. Select the channels to be added from the list presented on the screen.
- f. Click the **Change Subscription** button to finalize the change.

**Result**

The system is now registered with Red Hat Network and subscribed to the channels required.

[Report a bug](#)

## Chapter 16. Log Files

### 16.1. Red Hat Enterprise Virtualization Manager Installation Log Files

Table 16.1. Installation

Log File	Description
<code>/var/log/ovirt-engine/engine-cleanup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <code>engine-cleanup</code> command. This is the command used to reset a Red Hat Enterprise Virtualization Manager installation. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist.
<code>/var/log/ovirt-engine/engine-db-install-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <code>rhevm-setup</code> command detailing the creation and configuration of the <code>rhevm</code> database.
<code>/var/log/ovirt-engine/rhevm-dwh-setup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <code>rhevm-dwh-setup</code> command. This is the command used to create the <code>rhevm_history</code> database for reporting. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
<code>/var/log/ovirt-engine/ovirt-engine-reports-setup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <code>rhevm-reports-setup</code> command. This is the command used to install the Red Hat Enterprise Virtualization Manager Reports modules. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
<code>/var/log/ovirt-engine/engine-setup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <code>rhevm-setup</code> command. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
<code>/var/log/ovirt-engine/ovirt-engine-upgrade-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <code>rhevm-upgrade</code> command. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.

[Report a bug](#)

### 16.2. Red Hat Enterprise Virtualization Manager Log Files

Table 16.2. Service Activity

Log File	Description
<code>/var/log/ovirt-engine/engine.log</code>	Reflects all Red Hat Enterprise Virtualization Manager GUI crashes, Active Directory look-ups, Database issues, and other events.
<code>/var/log/ovirt-engine/host-deploy</code>	Log files from hosts deployed from the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

### 16.3. Red Hat Enterprise Virtualization Host Log Files

Table 16.3. Installation

Log File	Description
<code>/var/log/vdsm/libvirt.log</code>	Log file for <code>libvirt</code> .
<code>/var/log/vdsm/spm-lock.log</code>	Log file detailing the host's ability to obtain a lease on the Storage Pool Manager role. The log details when the host has acquired, released, renewed, or failed to renew the lease.
<code>/var/log/vdsm/vdsm.log</code>	Log file for VDSM, the Manager's agent on the virtualization host(s).
<code>/tmp/ovirt-host-deploy-@DATE@.log</code>	Host deployment log, copied to engine as <code>/var/log/ovirt-engine/host-deploy/ovirt-@DATE@-@HOST@-@CORRELATION_ID@.log</code> after the host has been successfully deployed.

[Report a bug](#)

### 16.4. Remotely Logging Host Activities

#### 16.4.1. Setting up a Virtualization Host Logging Server

##### Summary

Red Hat Enterprise Virtualization hosts generate and update log files, recording their actions and problems. Collecting these log files centrally simplifies debugging.

This procedure should be used on your centralized log server. You could use a separate logging server, or use this procedure to enable host logging on the Red Hat Enterprise Virtualization Manager.

##### Procedure 16.1. Setting up a Virtualization Host Logging Server

1. Configure SELinux to allow `rsyslog` traffic.

```
# semanage port -a -t syslogd_port_t -p udp 514
```

2. Edit `/etc/rsyslog.conf` and add below lines:

```
$template TmplAuth, "/var/log/%fromhost%/secure"
$template TmplMsg, "/var/log/%fromhost%/messages"

authpriv.* ?TmplAuth
*.info,mail.none;authpriv.none,cron.none ?TmplMsg
```

Un-comment the following:

```
#$ModLoad imudp
#$UDPServerRun 514
```

- Restart the rsyslog service:

```
# service rsyslog restart
```

#### Result

Your centralized log server is now configured to receive and store the **messages** and **secure** logs from your virtualization hosts.

[Report a bug](#)

### 16.4.2. Configuring Red Hat Enterprise Virtualization Hypervisor Hosts to Use a Logging Server

#### Prerequisites:

- » [Section 16.4.1, "Setting up a Virtualization Host Logging Server"](#)

#### Summary

Red Hat Enterprise Virtualization hosts generate and update log files, recording their actions and problems. Collecting these log files centrally simplifies debugging.

Use this procedure on a Red Hat Enterprise Virtualization Hypervisor host to begin sending log files to your centralized log server.

#### Procedure 16.2. Configuring Red Hat Enterprise Virtualization Hypervisor Hosts to Use a Logging Server

- Log in to your Red Hat Enterprise Virtualization Hypervisor host as **admin** to access the Hypervisors text user interface (TUI) setup screen.
- Select **Logging** from the list of options on the left of the screen.
- Press the **Tab** key to reach the text entry fields. Enter the IP address or FQDN of your centralized log server and the port it uses.
- Press the **Tab** key to reach the **Apply**, and press the **Enter** Key.

#### Result

Your Red Hat Enterprise Virtualization Hypervisor host has been configured to send messages to a centralized log server.

[Report a bug](#)

### 16.4.3. Configuring Red Hat Enterprise Linux Hosts to Use a Logging Server

#### Prerequisites:

- » [Section 16.4.1, "Setting up a Virtualization Host Logging Server"](#)

#### Summary

Red Hat Enterprise Virtualization hosts generate and update log files, recording their actions and problems. Collecting these log files centrally simplifies debugging.

Use this procedure on your Red Hat Enterprise Linux host to configure it to send logs to your centralized log server.

#### Procedure 16.3. Configuring Red Hat Enterprise Linux Hosts to Use a Logging Server

- Ensure that you have **rsyslog** installed:

```
# rpm -qi rsyslog
```

If **rsyslog** is not installed, install it:

```
# yum install rsyslog
```

- Edit **/etc/rsyslog.conf**, and specify the address of your centralized log server by adding the following:

```
*.* @YOURSERVERADDRESS:YOURSERVERPORT
```

For example:

```
*.* @@192.168.20.254:514
```

Or:

```
*.* @@log-server.company.com:514
```

The **@@** specifies the TCP protocol for transmission. Use a single **@** to specify UDP.

- Restart the **rsyslog** service:

```
# service rsyslog restart
```

#### Result

Your Red Hat Enterprise Linux host now sends log messages to your centralized log server.

[Report a bug](#)

## Chapter 17. Updating the Red Hat Enterprise Virtualization Environment

### 17.1. Upgrades between Minor Releases

#### 17.1.1. Checking for Red Hat Enterprise Virtualization Manager Updates

##### Summary

Use the **rhev-m-check-update** command, included in Red Hat Enterprise Virtualization Manager, to check for updates.

##### Procedure 17.1. Checking for Red Hat Enterprise Virtualization Manager Updates

1. Run **rhev-m-check-update** as the **root** user.

```
# rhev-m-check-update
```

- A. Where no updates are available the command will output the text **No updates available**.

```
# rhev-m-check-update
Loaded plugins: product-id, versionlock
Checking for updates... (This may take several minutes)
No updates available
```

- B. Where updates are available the command will list the packages to be updated.

```
# rhev-m-check-update
Loaded plugins: product-id
Checking for updates... (This may take several minutes)
10 Updates available:
* rhevm-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-backend-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-config-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-dbscripts-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-genericapi-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-notification-service-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-restapi-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-tools-common-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-userportal-3.2.0-10.14.beta1.el6ev.noarch
* rhevm-webadmin-portal-3.2.0-10.14.beta1.el6ev.noarch
```

##### Result

You have successfully checked for Red Hat Enterprise Virtualization Manager updates.

[Report a bug](#)

#### 17.1.2. Updating Red Hat Enterprise Virtualization Manager

##### Summary

Updates to the Red Hat Enterprise Virtualization Manager are released via Red Hat Network. Before installing an update from Red Hat Network be sure to read the advisory text associated with it as well as the most recent version of the *Red Hat Enterprise Virtualization — Release Notes*, and the *Red Hat Enterprise Virtualization — Technical Notes*. To complete an upgrade a number of actions must be performed including:

- stopping the **ovirt-engine** service,
- downloading and installing the updated packages,
- backing up and updating the database,
- performing post installation configuration, and
- restarting the **ovirt-engine** service.

A script is included to perform these actions for you in an automated fashion. Active virtualization hosts are not upgraded by this process and must be updated separately. As a result the virtual machines running upon them are not affected. All commands in this task must be run while logged into the system hosting Red Hat Enterprise Virtualization Manager as the **root** user.

##### Procedure 17.2. Updating Red Hat Enterprise Virtualization Manager

1. Run the **yum** command to update the **rhev-m-setup** package.

```
# yum update rhev-m-setup
```

2. Run the **rhev-m-upgrade** command to update the Red Hat Enterprise Virtualization Manager.

```
# rhev-m-upgrade
```

##### Note

Be aware that the upgrade process may take some time. Allow time for the upgrade process to complete and do not stop it once initiated. Once the upgrade has been completed you will also be instructed to separately upgrade the data warehouse and reports functionality. These additional steps are only required if these optional packages are installed.

##### Result

You have successfully updated Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

#### 17.1.3. Updating Red Hat Enterprise Virtualization Manager Reports

##### Summary

The Red Hat Enterprise Virtualization Manager Reports must be updated separately to the Red Hat Enterprise Virtualization Manager. This task provides instructions for updating both the history database and the reporting engine. You must run all commands in this task while logged into the system hosting Red Hat Enterprise Virtualization Manager as the **root** user.

##### Procedure 17.3. Updating Red Hat Enterprise Virtualization Manager Reports

1. Use the **yum** command to update the **rhev-m-reports** and **rhev-m-dwh** packages.

```
# yum update -y rhev-m-reports rhev-m-dwh
```



2. Run the `rhevmdwh-setup` command to update the `rhevmdhistory` database.

```
# rhvmdwh-setup
```

3. Run the `rhevmdreports-setup` command to update the reporting engine.

```
# rhvmdreports-setup
```

#### Result

You have successfully updated the Red Hat Enterprise Virtualization Manager Reports.

[Report a bug](#)

### 17.1.4. Updating Red Hat Enterprise Virtualization Hypervisors

#### Summary

Updating Red Hat Enterprise Virtualization Hypervisors involves reinstalling the Hypervisor with a newer version of the Hypervisor ISO image. This includes stopping and restarting the Hypervisor. Virtual machines are automatically migrated to a different host, as a result it is recommended that Hypervisor updates are performed at a time when the host's usage is relatively low.

It is recommended that administrators update Red Hat Enterprise Virtualization Hypervisors regularly. Important bug fixes and security updates are included in updates. Hypervisors which are not up to date may be a security risk.



#### Warning

Upgrading Hypervisor hosts involves shutting down, deactivating guests, and restarting the physical server. If any virtual machines are running on the Hypervisor, all data and configuration details may be destroyed if they are not shut down. Upgrading Hypervisors must be carefully planned and executed with care and consideration.



#### Important

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

#### Procedure 17.4. Updating Red Hat Enterprise Virtualization Hypervisors

1. Log in to the system hosting Red Hat Enterprise Virtualization Manager as the `root` user.
2. Ensure that:
  - » the system is subscribed to the **Red Hat Enterprise Virtualization** entitlement — if using certificate-based Red Hat Network; or
  - » the system is subscribed to the **Red Hat Enterprise Virtualization Hypervisor (v.6 x86-64)** (labeled `rhel-x86_64-server-6-rhev`) — if using classic Red Hat Network.
3. Run the `yum` command with the `update rhv-hypervisor` parameters to ensure that you have the most recent version of the `rhv-hypervisor` package installed.

```
# yum update rhv-hypervisor
```

4. Use your web browser to log in to the Administration Portal as a Red Hat Enterprise Virtualization administrative user.
5. Click the **Hosts** tab, and then select the host that you intend to upgrade. If the host is not displayed, or the list of hosts is too long to filter visually, perform a search to locate the host.
6. With the host selected, click the **General** tab on the **Details** pane.
  - A. If the host requires updating, an alert message indicates that a new version of the Red Hat Enterprise Virtualization Hypervisor is available.
  - B. If the host does not require updating, no alert message is displayed and no further action is required.
7. Ensure the host remains selected and click the **Maintenance** button, if the host is not already in maintenance mode. This will cause any virtual machines running on the host to be migrated to other hosts. If the host is the SPM, this function will be moved to another host. The status of the host changes as it enters maintenance mode. When the host status is **Maintenance**, the message in the general tab changes, providing you with a link which when clicked will re-install or upgrade the host.
8. Ensure that the host remains selected, and that you are on the **General** tab of the **Details** pane. Click the **Upgrade** link. The **Install Host** dialog box displays.
9. Select `rhv-hypervisor.iso`, which is symbolically linked to the most recent hypervisor image.
10. Click **OK** to update and re-install the host. The dialog closes, the details of the host are updated in the **Hosts** tab, and the status changes.
 

The host status will transition through these stages:

  - » **Installing**,
  - » **Reboot**,
  - » **Non Responsive**, and
  - » **Up**.

These are all expected, and each stage will take some time.
11. Once successfully updated, the host displays a status of **Up**. Any virtual machines that were migrated off the host, are at this point able to be migrated back to it.



#### Important

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

#### Result

You have successfully updated a Red Hat Enterprise Virtualization Hypervisor. Repeat these steps for each Hypervisor in the Red Hat Enterprise Virtualization environment.

[Report a bug](#)

### 17.1.5. Updating Red Hat Enterprise Linux Virtualization Hosts

#### Summary

Red Hat Enterprise Linux virtualization hosts are updated the same way as regular Red Hat Enterprise

Linux systems, using **yum**. It is highly recommended that you use **yum** to update your systems regularly, to ensure timely application of security and bug fixes. All steps in this task must be run while logged into the Red Hat Enterprise Linux virtualization host as the **root** user.

#### Procedure 17.5. Updating Red Hat Enterprise Linux Virtualization Hosts

1. On the administration portal, navigate to the **Hosts** tab and select the host to be updated. Click **Maintenance** to place it into maintenance mode.
2. On the Red Hat Enterprise Linux virtualization host, run the **yum** command with the **update** parameter to update all installed packages.

```
# yum update
```

3. If a package such as the kernel was updated, you must reboot the host to get the new functionality. If a package such as VDSM or libvirt was updated, you must restart that service to get the new functionality.

#### Result

You have successfully updated the Red Hat Enterprise Linux virtualization host. Repeat this process for each Red Hat Enterprise Linux virtualization host in the Red Hat Enterprise Virtualization environment.

[Report a bug](#)

### 17.1.6. Updating the Red Hat Enterprise Virtualization Guest Tools

#### Summary

The guest tools allow the Manager to communicate with running guests, providing information like IP addresses, memory usage, and installed applications for display in the Administration Portal. The guest tools are distributed as an ISO file that can be attached to guests. The ISO file is packaged as an RPM file, that can be installed and upgraded on registered Red Hat Enterprise Manager servers.

#### Procedure 17.6. Updating the Red Hat Enterprise Virtualization Guest Tools

1. On the Manager, as root user, use the **yum upgrade** to upgrade the **rhev-guest-tools-iso** package.

```
# yum update -y rhel-guest-tools-iso*
```

2. When the **rhev-guest-tools-iso** package has been successfully upgraded, use the **rhev-iso-uploader** to upload it to your ISO storage domain. Replace **ISODomain** with the name of your ISO storage domain.

```
rhev-iso-uploader --iso-domain=ISODomain upload /usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso
```

The **rhev-tools-setup.iso** file is actually a link to the most recently updated ISO file. The link is automatically changed to point to the newest ISO file every time you upgrade the **rhev-guest-tools-iso** package.

3. Using the web portal or REST API, attach the **rhev-tools-setup.iso** file to each of your guests, and from within each guest, upgrade the tools installed on each guest using the installer on the ISO.

#### Result

You've updated your **rhev-tools-setup.iso** file, uploaded the updated ISO to your ISO storage domain, and attached it to your virtual machines.

[Report a bug](#)

## 17.2. Upgrading to Red Hat Enterprise Virtualization 3.2

### 17.2.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

#### Prerequisites:

- » [Section 17.2.2, "Upgrading to Red Hat Enterprise Virtualization Manager 3.1"](#)

#### Summary

Upgrading Red Hat Enterprise Virtualization Manager to version 3.2 is performed using the **rhev-iso-upgrade** command. Virtualization hosts, and the virtual machines running upon them, will continue to operate independently while the Manager is being upgraded. Once the Manager upgrade is complete you will be able to upgrade your hosts, if you haven't already, to the latest versions of Red Hat Enterprise Linux and Red Hat Enterprise Virtualization Hypervisor.



#### Important

Users of Red Hat Enterprise Virtualization 3.0 must migrate to Red Hat Enterprise Virtualization 3.1 before attempting this upgrade.



#### Note

In the event that the upgrade fails the **rhev-iso-upgrade** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. Where this also fails detailed instructions for manually restoring the installation are displayed.

#### Procedure 17.7. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

1. **Add Red Hat Enterprise Virtualization 3.2 Subscription**

Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat Enterprise Virtualization Manager 3.2 packages. This procedure assumes that the system is already subscribed to required channels and entitlements to receive Red Hat Enterprise Virtualization 3.1 packages. These must also be available to complete the upgrade process.

#### Certificate-based Red Hat Network

The Red Hat Enterprise Virtualization 3.2 packages are provided by the **rhel-6-server-rhev-3.2-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

```
# yum-config-manager --enable rhel-6-server-rhev-3.2-rpms
```

#### Red Hat Network Classic

The Red Hat Enterprise Virtualization 3.2 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.2 x86\_64)** channel, also referred to as **rhel-x86\_64-server-6-rhev-3.2** in Red Hat Network Classic.

```
rhnc-channel --add --channel=rhel-x86_64-server-6-rhev-3.2
```

Use the **rhnc-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.2 x86\_64)** channel.

## 2. Update the *rhev-setup* Package

To ensure that you have the most recent version of the **rhev-setup** command installed you must update the *rhev-setup* package. Log in as the **root** user and use **yum** to update the *rhev-setup* package.

```
# yum update rhvm-setup
```

## 3. Run the *rhev-upgrade* Command

To upgrade Red Hat Enterprise Virtualization Manager run the **rhev-upgrade** command. You must be logged in as the **root** user to run this command.

```
# rhvm-upgrade
Loaded plugins: product-id, rhnplugin
Info: RHEV Manager 3.1 to 3.2 upgrade detected
Checking pre-upgrade conditions...(This may take several minutes)
```

## 4. If the *ipa-server* package is installed then an error message is displayed. Red Hat Enterprise Virtualization Manager 3.2 does not support installation on the same machine as Identity Management (IdM).

```
Error: IPA was found to be installed on this machine. Red Hat Enterprise
Virtualization Manager 3.2 does not support installing IPA on the same
machine. Please remove ipa packages before you continue.
```

To resolve this issue you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information see

<https://access.redhat.com/knowledge/articles/233143>.

## Result

Your Red Hat Enterprise Virtualization Manager installation has now been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.2 features you must also:

- Ensure that all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.
- Change all of your clusters to use compatibility version 3.2.
- Change all of your data centers to use compatibility version 3.2.

[Report a bug](#)

## 17.2.2. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

### Summary

Upgrading Red Hat Enterprise Virtualization Manager to version 3.1 is performed using the **rhev-upgrade** command. Virtualization hosts, and the virtual machines running upon them, will continue to operate independently while the Manager is being upgraded. Once the Manager upgrade is complete you will be able to upgrade your hosts, if you haven't already, to the latest versions of Red Hat Enterprise Linux and Red Hat Enterprise Virtualization Hypervisor.



### Important

Refer to <https://access.redhat.com/knowledge/articles/269333> for an up to date list of tips and considerations to be taken into account when upgrading to Red Hat Enterprise Virtualization 3.1.



### Important

Users of Red Hat Enterprise Virtualization 2.2 must migrate to Red Hat Enterprise Virtualization 3.0 before attempting this upgrade. For information on migrating from Red Hat Enterprise Virtualization 2.2 to Red Hat Enterprise Virtualization 3.0, refer to <https://access.redhat.com/knowledge/techbriefs/migrating-red-hat-enterprise-virtualization-manager-version-22-30>.



### Note

In the event that the upgrade fails the **rhev-upgrade** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. Where this also fails detailed instructions for manually restoring the installation are displayed.

## Procedure 17.8. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

### 1. Add JBoss Enterprise Application Platform 6 Subscription

Ensure that the system is subscribed to the required channels and entitlements to receive JBoss Enterprise Application Platform 6 packages. JBoss Enterprise Application Platform 6 is a required dependency of Red Hat Enterprise Virtualization Manager 3.1.

#### Certificate-based Red Hat Network

The JBoss Enterprise Application Platform 6 packages are provided by the **JBoss Enterprise Application Platform** entitlement in certificate-based Red Hat Network.

Use the **subscription-manager** command to ensure that the system is subscribed to the **JBoss Enterprise Application Platform** entitlement.

```
# subscription-manager list
```

#### Red Hat Network Classic

The JBoss Enterprise Application Platform 6 packages are provided by the **JBoss Application Platform (v 6) for 6Server x86\_64** channel, also referred to as **jbossappplatform-6-x86\_64-server-6-rpm**, in Red Hat Network Classic. The Channel Entitlement Name for this channel is **JBoss Enterprise Application Platform (v 4, zip format)**.

Use the **rhnc-channel** command, or the Red Hat Network Web Interface, to subscribe to the **JBoss Application Platform (v 6) for 6Server x86\_64** channel.

**2. Add Red Hat Enterprise Virtualization 3.1 Subscription**

Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat Enterprise Virtualization Manager 3.1 packages.

**Certificate-based Red Hat Network**

The Red Hat Enterprise Virtualization 3.1 packages are provided by the **rhel-6-server-rhev-3.1-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

```
# yum-config-manager --enable rhel-6-server-rhev-3.1-rpms
```

**Red Hat Network Classic**

The Red Hat Enterprise Virtualization 3.1 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.1 x86\_64)** channel, also referred to as **rhel-x86\_64-server-6-rhev-3.1** in Red Hat Network Classic.

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.1 x86\_64)** channel.

**3. Remove Red Hat Enterprise Virtualization 3.0 Subscription**

Ensure that the system does not use any Red Hat Enterprise Virtualization Manager 3.0 packages by removing the Red Hat Enterprise Manager 3.0 channels and entitlements.

**Certificate-based Red Hat Network**

Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.0 repositories in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

```
# yum-config-manager --disablerepo=rhel-6-server-rhev-3-rpms
```

```
# yum-config-manager --disablerepo=jb-eap-5-for-rhel-6-server-rpms
```

**Red Hat Network Classic**

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to remove the **Red Hat Enterprise Virtualization Manager (v.3.0 x86\_64)** channels.

```
# rhn-channel --remove --channel=rhel-6-server-rhev-3
```

```
# rhn-channel --remove --channel=jbappplatform-5-x86_64-server-6-rpm
```

**4. Update the *rhev-setup* Package**

To ensure that you have the most recent version of the **rhev-upgrade** command installed you must update the **rhev-setup** package. Log in as the **root** user and use **yum** to update the **rhev-setup** package.

```
# yum update rhev-setup
```

**5. Run the *rhev-upgrade* Command**

To upgrade Red Hat Enterprise Virtualization Manager run the **rhev-upgrade** command. You must be logged in as the **root** user to run this command.

```
# rhev-upgrade
Loaded plugins: product-id, rhnplugin
Info: RHEV Manager 3.0 to 3.1 upgrade detected
Checking pre-upgrade conditions...(This may take several minutes)
```

**6. If the *ipa-server* package is installed then an error message is displayed. Red Hat Enterprise Virtualization Manager 3.1 does not support installation on the same machine as Identity Management (IdM).**

```
Error: IPA was found to be installed on this machine. Red Hat Enterprise
Virtualization Manager 3.1 does not support installing IPA on the same
machine. Please remove ipa packages before you continue.
```

To resolve this issue you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information see <https://access.redhat.com/knowledge/articles/233143>.

**7. A list of packages that depend on JBoss Enterprise Application Platform 5 is displayed. These packages must be removed to install JBoss Enterprise Application Platform 6, required by Red Hat Enterprise Virtualization Manager 3.1.**

```
Warning: the following packages will be removed if you proceed with the
upgrade:

* objectweb-asm

Would you like to proceed? (yes|no):
```

You must enter **yes** to proceed with the upgrade, removing the listed packages.

**Result**

Your Red Hat Enterprise Virtualization Manager installation has now been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.1 features you must also:

- ▀ Ensure that all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.
- ▀ Change all of your clusters to use compatibility version 3.1.
- ▀ Change all of your data centers to use compatibility version 3.1.

**See Also:**

- ▀ [Section 17.1.4, "Updating Red Hat Enterprise Virtualization Hypervisors"](#)
- ▀ [Section 17.1.5, "Updating Red Hat Enterprise Linux Virtualization Hosts"](#)
- ▀ [Section 17.2.3, "Changing the Cluster Compatibility Version"](#)
- ▀ [Section 17.2.4, "Changing the Data Center Compatibility Version"](#)

[Report a bug](#)

**17.2.3. Changing the Cluster Compatibility Version****Prerequisites:**

- ▀ [Section 17.1.4, "Updating Red Hat Enterprise Virtualization Hypervisors"](#)

» [Section 17.1.5, "Updating Red Hat Enterprise Linux Virtualization Hosts"](#)

### Summary

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

#### Procedure 17.9. Changing the Cluster Compatibility Version

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. Click the **Hosts** tab.
3. Select all the hosts from the cluster you wish to change.
4. Click the **Maintenance** button to put the hosts in maintenance mode.
5. Click the **Clusters** tab.
6. Select the cluster that you wish to change from the list displayed. If the list of clusters is too long to filter visually then perform a search to locate the desired cluster.
7. Click the **Edit** button.
8. Change the **Compatibility Version** to the desired value.
9. Click **OK**.

### Result

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

[Report a bug](#)

## 17.2.4. Changing the Data Center Compatibility Version

### Prerequisites:

» [Section 17.2.3, "Changing the Cluster Compatibility Version"](#)

### Summary

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

#### Procedure 17.10. Changing the Data Center Compatibility Version

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. A.
  - a. If you can endure virtual machine downtime, you will put all of your hosts in maintenance mode at once. Click the **Hosts** tab.
  - b. Select all the hosts from the data center you wish to change.
  - c. Click the **Maintenance** button to put the hosts in maintenance mode.
 B. Upgrade each cluster's compatibility level individually before upgrading the data center's compatibility level if you cannot put all of your hosts into maintenance mode at once. This process is described in *Changing the Cluster Compatibility Version*.
3. Click the **Data Centers** tab.
4. Select the data center that you wish to change from the list displayed. If the list of data centers is too long to filter visually then perform a search to locate the desired data center.
5. Click the **Edit** button.
6. Change the **Compatibility Version** to the desired value.
7. Click **OK**.

### Result

You have updated the compatibility version of the data center.

[Report a bug](#)

## Chapter 18. Backups

### 18.1. Backing Up the Engine Database Using the backup.sh Script

#### Summary

The Red Hat Enterprise Virtualization Manager includes a script to automate database backups. Using this script on your Manager server, you can protect yourself against potential data loss.



#### Important

If your environment was upgraded from Red Hat Enterprise Virtualization 3.0, please see [KB337653](#) for additional information about backing up and restoring your database.

#### Procedure 18.1. Backing up the engine database using the backup.sh script

1. Change into the `/usr/share/ovirt-engine/dbscripts/` directory.
2. Invoke **backup.sh** with the `-h` parameter to see the available options.

```
Usage: backup.sh [-h] [-s SERVERNAME] [-p PORT] [-d DATABASE] [-l DIR] -u
USERNAME [-v]

-s SERVERNAME - The database servername for the database (def. localhost)
-p PORT      - The database port for the database (def. 5432)
-d DATABASE  - The database name (def. engine)
-u USERNAME  - The username for the database.
-v          - Turn on verbosity (WARNING: lots of output)
-l DIR       - Backup file directory.
-h          - This help text.

for more options please run pg_dump --help
```

3. Invoke the **backup.sh** command again with parameters appropriate for your environment. If you are backing up the local **engine** database, the `-s`, `-p`, and `-d` parameters are not necessary. Use the `-l` to specify the backup directory. This will cause a `.sql` file to be created in the directory you give.
4. Copy the `.sql` you just created from the directory you specified to a safe remote location.

#### Result

You have used the backup.sh script to backup your **engine** database.

[Report a bug](#)

### 18.2. Restoring the Engine Database Using the restore.sh Script

#### Summary

The Red Hat Enterprise Virtualization Manager includes a script to automate database restoration. Using this script on your Manager server, you can recover from database corruption.



#### Important

If your environment was upgraded from Red Hat Enterprise Virtualization 3.0, please see [KB337653](#) for additional information about backing up and restoring your database.

#### Procedure 18.2. Restoring the Engine Database Using the restore.sh Script

1. Change into the `/usr/share/ovirt-engine/dbscripts/` directory.
2. Invoke **restore.sh** with the `-h` parameter to see the available options.

```
Usage: restore.sh [-h] [-s SERVERNAME] [-p PORT] -u USERNAME -d DATABASE -f
FILE [-r]

-s SERVERNAME - The database servername for the database (def. localhost)
-p PORT      - The database port for the database (def. 5432)
-u USERNAME  - The username for the database.
-d DATABASE  - The database name
-f File      - Backup file name to restore from.
-r          - Remove existing database with same name
-h          - This help text.

for more options please run pg_restore --help
```

3. Invoke the **restore.sh** command again with parameters appropriate for your environment. If you are restoring the local **engine** database, the `-s` and `-p` parameters are not necessary. Use the `-d` to specify name of the database you are creating. Red Hat Enterprise Virtualization expects a primary database named **engine**. Use the `-f` to specify the `.sql` file you are restoring from.

#### Result

You have used the restore.sh script to restore your **engine** database.

[Report a bug](#)

### 18.3. Backing Up 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.



#### Important

If your environment was upgraded from Red Hat Enterprise Virtualization 3.0, please see [KB340903](#) for additional information about backing up and restoring your configuration files.

Table 18.1. Configuration files and directories requiring backup.

Location	What is it?
/etc/ovirt-engine/	Contains Red Hat Enterprise Virtualization Manager configuration files.
/etc/sysconfig/ovirt-engine	It is the ovirt-engine service configuration file.
/etc/yum/pluginconf.d/versionlock.list	Contains version information about currently installed Red Hat Enterprise Virtualization components.
/etc/pki/ovirt-engine/	Security certificates provided by the Red Hat Enterprise Virtualization Manager to clients.
/usr/share/jasperreports-server-pro/buildomatic	Contains files required to build the Red Hat Enterprise Virtualization reports server.
/var/lib/ovirt-engine/backups	A folder where all backup data is stored.
/var/lib/ovirt-engine/deployments	A folder where all deployment data is stored.
/usr/share/ovirt-engine-reports/reports/users/rhev-002dadmin.xml	Contains plain-text, un-encrypted user and password information for the rhev-admin user.
/usr/share/ovirt-engine-reports/default_master.properties	Contains settings to handle the configuration and deployment of JasperServer.
/root/.rnd	Random seed file, used to generate secure certificates.
/var/log/ovirt-engine/engine-setup_ <b>SETUPDATE</b> .log	Contains the answers you gave to the setup configuration questions. Restoring the Manager requires running the setup again, and giving the same answers.

When all the files in the above table have been backed up, you will be able to recover the Red Hat Enterprise Virtualization Manager to a working state after an unforeseen event.

[Report a bug](#)


## 18.4. Restoring Red Hat Enterprise Virtualization Manager Configuration Files

### Prerequisites:

- » [Section 18.3, "Backing Up Manager Configuration Files"](#)

### Summary

Restore a backed up copy of configuration files to the Red Hat Enterprise Virtualization Manager.


**Important**

If your environment was upgraded from Red Hat Enterprise Virtualization 3.0, please see [KB340903](#) for additional information about backing up and restoring your configuration files.

### Procedure 18.3. Restoring Red Hat Enterprise Virtualization Manager Configuration Files

1. Stop the **Engine** service:

```
# service ovirt-engine stop
```

2. Completely remove all previous installations of the Red Hat Enterprise Virtualization Manager:

```
# yum remove rhevm
```

3. Remove /etc/pki/ovirt-engine:

```
# rm -rf /etc/pki/ovirt-engine
```

4. Remove the main rhevm directory:

```
# rm -rf /etc/ovirt-engine
```

5. Install the Red Hat Enterprise Virtualization Manager:

```
# yum install -y rhevm
```

6. Run rhevm-setup, giving the same answers as when you originally installed rhevm:

```
# rhevm-setup
```

Your answers can be found in **/var/log/engine-setup-*SETUP-DATE*.log**, which you backed up.

7. Stop the **Engine** service, which was restarted as a part of the previous command:

```
# service ovirt-engine stop
```

8. Restore the backed up configuration files to their original locations.

9. Make sure the ownership of the .keystore file is correct:

```
# chown ovirt:ovirt /etc/pki/ovirt-engine/.keystore
```

10. Make sure the permissions of the notifier.conf file is correct:

```
# chmod 640 /etc/ovirt-engine/notifier/notifier.conf
```

11. Start the **Engine** service:

```
# service ovirt-engine start
```

**Result**

You have restored a backed up copy of configuration files to the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)



**Part III. Gathering Information About the Environment**

## Chapter 19. Reports, History Database Reports, and Dashboards


### 19.1. Reports

#### 19.1.1. 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.

Red Hat Enterprise Virtualization Manager uses **PostgreSQL 8.4.12** 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 **engine**.

Installing the **rhev-m-dwh** package creates a second database called **ovirt\_engine\_history**, which contains historical configuration information and statistical metrics collected every minute over time from the **engine** 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**

The replication of data in the **ovirt\_engine\_history** database is performed by the Red Hat Enterprise Virtualization Manager Extract Transform Load Service, **ovirt-engine-dwhd**. 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 **engine** database, transforming the data to the history database standard and loading it to the **ovirt\_engine\_history** database.

The **ovirt-engine-dwhd** service must not be stopped.

The **ovirt\_engine\_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.

[Report a bug](#)

#### 19.1.2. JasperReports and JasperServer in Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization provides a customized implementation of JasperServer, which allows web-based access to a range of pre-configured reports and dashboards, plus the ability to create ad hoc reports.

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.

[Report a bug](#)

#### 19.1.3. Online Help for JasperReports

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 right-hand corner of the browser.



**Figure 19.1. Red Hat Enterprise Virtualization Reports online help**

[Report a bug](#)

#### 19.1.4. Jasper Reports System Requirements

The Red Hat Enterprise Virtualization Manager Reports tool supports the following browsers:

- In Red Hat Enterprise Linux 5.7 - Firefox 17 or later
- In Red Hat Enterprise Linux 6 - Firefox 17 or later
- In Windows 7 - Internet Explorer 9
- In Windows Server 2008 - Internet Explorer 9

[Report a bug](#)

#### 19.1.5. Users in the Red Hat Enterprise Virtualization Reports Portal

The Red Hat Enterprise Virtualization Reports Portal does not use your directory server for authentication.

By default, there are two Reports Portal users: **rhev-admin** and **superuser**. The passwords for these users were set during the installation of Red Hat Enterprise Virtualization Reports. Generally, additional users must be added manually.


When a domain user accesses the Reports Portal from within the Administration Portal using right-click reporting, a corresponding user is automatically created in the Reports Portal using the user's domain user name. This user cannot login to the Reports Portal directly, but is able to view all the reports accessible from the Administration portal.

[Report a bug](#)

#### 19.1.6. Logging in to Access the Reports Portal

You were prompted to set a password for the **superuser** and **rhev-admin** accounts when you installed Red Hat Enterprise Virtualization Reports. Red Hat Enterprise Virtualization Reports does not provide default passwords.

To access reports, navigate to the reports portal at: **https://YOUR.MANAGER.URL/rhev-m-reports/login.html**. A login screen for Red Hat Enterprise Virtualization Reports is displayed.

 **Note**

You can also access the reports portal from your Red Hat Enterprise Virtualization landing page.

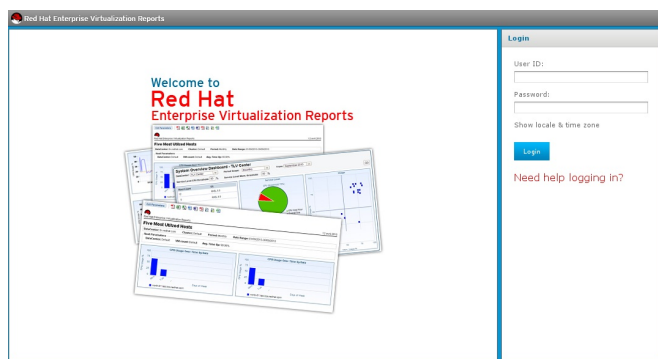


Figure 19.2. Red Hat Enterprise Virtualization Reports login screen

Enter your login credentials. If this is the first time you are connecting to the reports portal, log in as **ovirt-user**. Click the **Login** button.

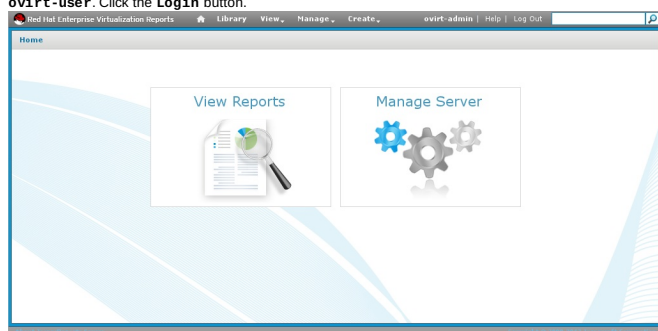


Figure 19.3. Red Hat Enterprise Virtualization Reports main screen

The Reports Portal does not use your directory service for authentication. By default, the Reports Portal includes two users: **rhev-admin** and **superuser**. Generally, additional users need to be created within the Reports Portal.

[Report a bug](#)

### 19.1.7. Accessing the Red Hat Enterprise Virtualization Reports User Management Menu

#### Summary

You can add additional reports users, giving them access to the reports portal. Complete this procedure as a user with sufficient permissions to manage other users, like **rhev-admin**.

1. In to Red Hat Enterprise Virtualization reports portal, hover over the **Manage** button on the top menu bar.
2. Click on **Users** in the drop-down menu that appears to access the **Manage Users** interface. It contains three panes:
  - » **Organizations**
  - » **Users**
  - » **Properties**
3. Select a user in the **Users** pane by clicking on the name of the user. Information about the user displays in the **Properties** pane.
4. Click the **Edit** button at the bottom of the user's **Properties** pane.

The **Properties** pane contains these fields:

- » **User name**,
- » **User ID**,
- » **Email**,
- » **Password (required)**,
- » **Confirm Password (required)**,
- » A **User is enabled** check box,
- » A **The user is defined externally** check box,
- » A list of **Roles Available** to the user, and
- » A list of **Roles Assigned** to the user.

5. Click the **Save** button.

#### Result

You have given more users permissions to access the reports portal.

[Report a bug](#)

### 19.1.8. Reports Portal User Roles

There are three 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.
3. **ROLE\_ANONYMOUS** - Can log in and look at reports and dashboards.

Other roles can be created and assigned. For information on how to create and assign other roles, detailed information about user management, and other system functions, please refer to the JasperServer documentation.




Figure 19.4. JasperReports user roles

[Report a bug](#)

### 19.1.9. Navigating Reports and Dashboards

Select the **View Reports** button on the reports portal home page.





Figure 19.5. Red Hat Enterprise Virtualization Reports home screen

You can use the smaller **Home** (  ) button in the navigation bar at the top of the reports portal to return to this page.

Use the **Filter** pane on the left of the screen to select a subset of reports you would like to view.




Figure 19.6. Red Hat Enterprise Virtualization Reports Filter pane

You can use filters to select from the available reports.

Table 19.1. Navigation Filters

Filter	Description
<b>Available Resources</b>	Select from All, Modified by me, or Viewed by me.
<b>Resource type</b>	Choose from the types of available resources including Reports, Ad Hoc views, Dashboards, and more.
<b>Timeframe</b>	Choose a time frame you'd like to see information from.
<b>Schedule</b>	Filter by data collection schedule.

[Report a bug](#)

### 19.1.10. Report Parameters

Report parameters are user-defined at report run time. Report parameters define the scope and timeframe of the report. When running a report, you are prompted for the parameters applicable to the report you selected.

To view the required parameters for a report, click the report in the reports list.

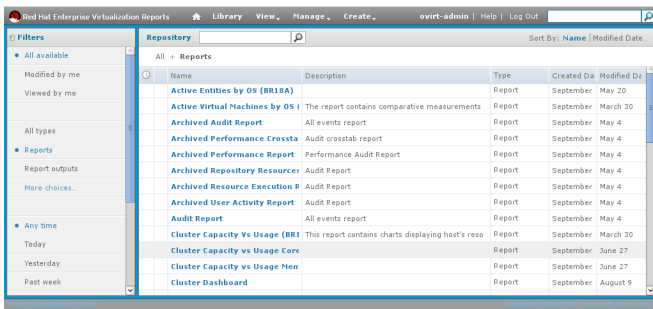


Figure 19.7. Red Hat Enterprise Virtualization Reports - Reports List

Select a report from the list to display the **Input Controls** window. The **Input Controls** window consists of a number of drop-down menus allow you to define the report's parameters.

#### Note

The dialog is contextual and differs from report to report. Parameters marked with an asterisk (\*) are required.

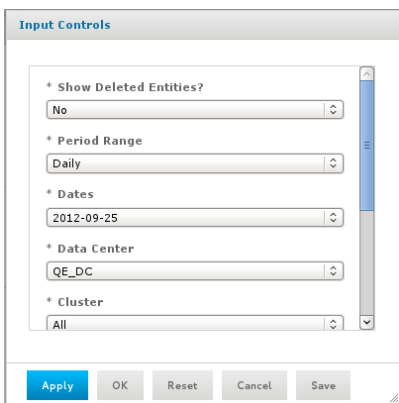


Figure 19.8. Report Parameter Selection

#### Cascading parameters

Many report parameters are cascading input fields. This means the selection made for one parameter changes the options available for another parameter. 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.

You can toggle whether deleted entries are shown in the report using the **Show Deleted Entities?** field in the **Input Controls** window.

[Report a bug](#)

#### 19.1.11. Right-click Reporting Integration with the Red Hat Enterprise Virtualization Administration Portal

The Administration portal provides integrated access to reports on most resources.

To access a report on a given resource, select the resource in the Administration Portal. Right-click the resource to show a context sensitive menu, and select the **Show Report** option. This expands to show all of the available reports on the selected resource.

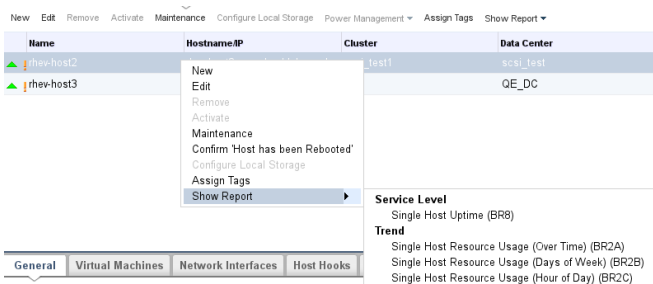


Figure 19.9. Right-click Reporting

Alternatively, you can select a given resource in the Administration Portal. If there are reports on that

resource, the **Show Report** action becomes available above the results pane.




Figure 19.10. Alternative to Right-click Reporting

[Report a bug](#)

## 19.1.12. Executive Reports

### 19.1.12.1. Executive reports: Active Virtual Machines by OS

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 19.2. Active Virtual Machines by OS Parameters

Parameter	Description
<b>Show Deleted Entities?</b>	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
<b>Period Range</b>	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
<b>Dates</b>	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
<b>Data Center</b>	The report includes only virtual machines in the selected data center. The options list shows only data centers that contain virtual machines.
<b>Cluster</b>	The report only includes virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the report includes all virtual machines in the selected data center.
<b>Virtual Machine Type</b>	The report only includes virtual machines of the selected type. Possible types are <b>Server</b> and <b>Desktop</b> . The options list shows only types that exist in the selected data center and cluster. If <b>All</b> is selected, the report includes all virtual machine types.

[Report a bug](#)

### 19.1.12.2. Executive Reports: 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 19.3. Cluster Capacity Vs Usage Parameters

Parameter	Description
<b>Show Deleted Entities?</b>	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
<b>Period Range</b>	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
<b>Dates</b>	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
<b>Data Center</b>	The list of options for the <b>Cluster</b> parameter includes only clusters in the selected data center. The options list contains only data centers that contain clusters.
<b>Cluster</b>	The report only includes the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the report includes all clusters in the selected data center.

[Report a bug](#)

### 19.1.12.3. Executive Reports: Host Operating System 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 19.4. Host OS Break Down Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> is selected, the report includes all hosts in the selected data center.

[Report a bug](#)**19.1.12.4. Executive Reports: 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 19.5. Summary of Host Usage Resources Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> is selected, the report includes all hosts in the selected data center.

[Report a bug](#)**19.1.13. Inventory Reports****19.1.13.1. Inventory Reports: 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 19.6. Hosts Inventory Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> 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 <b>All</b> is selected, the report includes all host types.

[Report a bug](#)**19.1.13.2. Inventory Reports: Storage Domain 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 19.7. Storage Domain Size Over Time Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
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 <b>Dates</b> 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 <b>Storage Domain Name</b> parameter includes only storage domains that were attached during the specified period.
Data Center	The options list for the <b>Storage Domain Name</b> parameter shows only storage domains in this selected data center.
Storage Domain Type	The options list for the <b>Storage Domain Name</b> 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 of options shows only storage domains that were attached to the data center during the selected period.

[Report a bug](#)**19.1.13.3. Inventory Reports: 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 19.8. Virtual Machines Inventory Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> is selected, the report includes all virtual machines in the selected data center.
Virtual Machine 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 <b>All</b> is selected, the report includes all virtual machine types.

[Report a bug](#)**19.1.13.4. Inventory Reports: Cloud Provider Virtual Machine Inventory**

The *Cloud Provider Virtual Machine Inventory* report shows a list of all virtual machines in the selected data center and cluster, and is required by cloud providers to bill customers. The following parameters must be provided to run this report:

Table 19.9. Cloud Provider Virtual Machine Inventory Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> is selected, the report includes all virtual machines in the selected data center.
Virtual Machine 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 <b>All</b> is selected, the report includes all virtual machine types.

[Report a bug](#)



### 19.1.14. Service Level Reports

#### 19.1.14.1. Service Level Reports: 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 19.10. Cluster Host Uptime Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> 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 <b>All</b> is selected, the report includes all host types.

[Report a bug](#)

#### 19.1.14.2. Service Level Reports: 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 19.11. Cluster Quality of Service - Hosts Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> 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 <b>All</b> is selected, the report includes all host types.
CPU Threshold	The report measures the quality of service as the amount of time hosts sustain load above a given threshold. The <b>CPU Threshold</b> 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 <b>CPU Threshold</b> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.
Memory Threshold	The report measures the quality of service as the amount of time hosts sustain load above a given threshold. The <b>Memory Threshold</b> 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 <b>Memory Threshold</b> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.

[Report a bug](#)

#### 19.1.14.3. Service Level Reports: 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 19.12. Cluster Quality of Service - Virtual Machines Parameters

Parameter	Description
<b>Show Deleted Entities?</b>	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
<b>Period Range</b>	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
<b>Dates</b>	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
<b>Data Center</b>	The list of options for the <b>Cluster</b> parameter includes only clusters in the selected data center. The options list shows only data centers that contain clusters.
<b>Cluster</b>	The report includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the report includes all virtual machines in the selected data center.
<b>Virtual Machine Type</b>	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 <b>All</b> is selected, the report includes all virtual machine types.
<b>CPU Threshold</b>	The report measures quality of service as the amount of time virtual machines sustain load above a given threshold. The <b>CPU Threshold</b> 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 <b>CPU Threshold</b> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.
<b>Memory Threshold</b>	The reports measures quality of service as the amount of time virtual machines sustain load above a given threshold. The <b>Memory Threshold</b> 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 <b>Memory Threshold</b> of 60 per cent is a suggested starting point to produce a meaningful quality of service report.

[Report a bug](#)

#### 19.1.14.4. Service Level Reports: 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 19.13. Single Host Uptime Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>Host Name</b> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>Host Name</b> parameter includes all hosts in the selected data center.
Host Type	The list of options for the <b>Host Name</b> parameter includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If <b>All</b> is selected, the list of options for the <b>Host Name</b> parameter includes all host types.
Host Name	The report refers to the host selected. A report is only for a single host and a user must select a host.

[Report a bug](#)

#### 19.1.14.5. Service Level Reports: 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 19.14. Top 10 Downtime Hosts Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> 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 <b>All</b> is selected, the report includes all host types.

[Report a bug](#)

#### 19.1.14.6. Service Level Reports: 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 19.15. High Availability Virtual Servers Uptime Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> is selected, the report includes all virtual servers in the selected data center.

[Report a bug](#)

### 19.1.15. Trend Reports

#### 19.1.15.1. Trend Reports: 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 19.16. Five Least Utilized Hosts (Over Time) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> 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 <b>All</b> is selected, the report includes all host types.

[Report a bug](#)

#### 19.1.15.2. Trend Reports: 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 19.17. Five Least Utilized Virtual Machines (Over Time) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> is selected, the report includes all virtual machines in the selected data center.
Virtual Machine 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 <b>All</b> is selected, the report includes all virtual machine types.

[Report a bug](#)

#### 19.1.15.3. Trend Reports: 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 19.18. Five Most Utilized Hosts (Over Time) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> 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 <b>All</b> is selected, the report includes all host types.

[Report a bug](#)

#### 19.1.15.4. Trend Reports: 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 19.19. Five Most Utilized Virtual Machines (Over Time) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>All</b> is selected, the report includes all virtual machines in the selected data center.
Virtual Machine 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 <b>All</b> is selected, the report includes all virtual machine types.

[Report a bug](#)

#### 19.1.15.5. Trend Reports: 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 19.20. Multiple Hosts Resource Usage (Over Time) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>Hosts List</b> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>Hosts List</b> parameter includes all hosts in the selected data center.
Host Type	The list of options for the <b>Hosts List</b> parameter includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If <b>All</b> is selected, the list of options for the <b>Hosts List</b> parameter includes all host types.
Hosts list	The report includes all hosts selected in the host list. Select any number of hosts up to a maximum of five.

[Report a bug](#)

#### 19.1.15.6. Trend Reports: 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 19.21. Multiple Virtual Machines Resource Usage (Over Time) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>VM List</b> parameter include only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>VM List</b> parameter includes all virtual machines in the selected data center.
Virtual Machine Type	The list of options for the <b>VM List</b> 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 <b>All</b> is selected, the list of options for the <b>VM List</b> parameter includes all virtual machine types.
Virtual Machine 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.

[Report a bug](#)

#### 19.1.15.7. Trend Reports: 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 19.22. Single Host Resource Usage (Days of Week) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>Host Name</b> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>Host Name</b> parameter includes all hosts in the selected data center.
Host Type	The list of options for the <b>Host Name</b> parameter includes only hosts of the selected type. The options list shows only host types present in the selected data center and cluster. If <b>All</b> is selected, the list of options for the <b>Host Name</b> 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.

[Report a bug](#)

#### 19.1.15.8. Trend Reports: 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 19.23. Single Host Resource Usage (Hour of Day) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>Host Name</b> parameter includes only hosts in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>Host Name</b> 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 <b>Host Name</b> parameter. The options list shows only host types present in the selected data center and cluster. If <b>All</b> is selected, the list of options for the <b>Host Name</b> 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.

[Report a bug](#)

#### 19.1.15.9. Trend Reports: 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 19.24. Single Virtual Machine Resources (Days of Week) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>VM Name</b> parameter includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>VM Name</b> parameter includes all virtual machines in the selected data center.
Virtual Machine Type	The list of options for the <b>VM Name</b> 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 <b>All</b> is selected, the list of options for the <b>VM Name</b> parameter includes all virtual machine types.
Virtual Machine 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.

[Report a bug](#)

#### 19.1.15.10. Trend Reports: 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 19.25. Single Virtual Machine Resources (Hour of Day) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>VM Name</b> parameter includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>VM Name</b> parameter includes all virtual machines in the selected data center.
Virtual Machine Type	The list of options for the <b>VM Name</b> 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 <b>All</b> is selected, the list of options for the <b>VM Name</b> parameter includes all virtual machine types.
Virtual Machine 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.

[Report a bug](#)

#### 19.1.15.11. Trend Reports: 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 19.26. Single Virtual Machine Resources (Over Time) Parameters

Parameter	Description
Show Deleted Entities?	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
Period Range	The report is for the period range selected. Daily reports cover a single day. Monthly reports cover a single month. Quarterly reports cover a three-month quarter, beginning on the month specified in the <b>Dates</b> parameter. Yearly reports cover a year, beginning on the month specified in the <b>Dates</b> parameter.
Dates	The report covers the selected period range, beginning on this date. Daily period ranges pass in one day increments. 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. A yearly period range also starts on the selected month.
Data Center	The list of options for the <b>Cluster</b> 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 <b>VM Name</b> parameter includes only virtual machines in the selected cluster. The options list shows only clusters in the selected data center. If <b>All</b> is selected, the list of options for the <b>VM Name</b> parameter includes all virtual machines in the selected data center.
Virtual Machine Type	The list of options for the <b>VM Name</b> 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 <b>All</b> is selected, the list of options for the <b>VM Name</b> parameter includes all virtual machine types.
Virtual Machine 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.

[Report a bug](#)

#### 19.1.16. 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, navigate to the **Create** drop-down menu on the top menu bar and select **Ad Hoc View** to open the **Data Chooser**: **Source** window.




Figure 19.11. Create Ad Hoc Report - Administrator's View

The **Working with the Ad Hoc Editor** section of the online help explains the ad hoc report interface in detail.

[Report a bug](#)

### 19.1.17. Reports Schema: 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.

#### Note

**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.

Tag relations and latest tag relations history views

Tag relations history in the system.

Table 19.27. tag\_and\_ENUM\_VIEWS.table\_v3\_2\_tag\_relations\_history\_viewv3\_2\_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"> <li>» 2 - VM</li> <li>» 3 - Host</li> <li>» 5 - VM pool</li> <li>» 18 - Tag</li> </ul>
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.

Tag details and latest tag details views

Tag details history in the system.

Table 19.28. v3\_2\_tag\_details\_viewv3\_2\_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.

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 19.29. v3\_2\_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.

[Report a bug](#)

## 19.2. History Database Reports

### 19.2.1. Red Hat Enterprise Virtualization History Database

Red Hat Enterprise Virtualization Reports uses data from the Red Hat Enterprise Virtualization History Database (called **ovirt\_engine\_history**) which tracks the **engine** database over time.



#### Important

Sufficient data must exist in the history database to produce meaningful reports. 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.

[Report a bug](#)

### 19.2.2. Tracking Configuration History

The ETL service, **ovirt-engine-dwhd**, tracks three types of changes:

- A new entity is added to the **engine** database - the ETL Service replicates the change to the **ovirt\_engine\_history** database as a new entry.
- An existing entity is updated - the ETL Service replicates the change to the **ovirt\_engine\_history** database as a new entry.
- An entity is removed from the **engine** database - A new entry in the **ovirt\_engine\_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 **ovirt\_engine\_history** database differ from the corresponding tables in the **engine** 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 **engine** database appear in a single table in **ovirt\_engine\_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.

[Report a bug](#)

### 19.2.3. 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.

[Report a bug](#)

### 19.2.4. 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 **ovirt\_engine\_history** database as a new entry.
- A tag is updated - the ETL Service replicates the change of tag details to the **ovirt\_engine\_history** database as a new entry.
- An entity or tag branch is removed from the Administration Portal - the **ovirt\_engine\_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.

[Report a bug](#)

### 19.2.5. Connecting to the History Database

The **ovirt\_engine\_history** database resides within the instance of PostgreSQL that the installer creates during Red Hat Enterprise Virtualization Manager installation.

For local connections to the database, use a PostgreSQL compatible query or reporting tool with the credentials used in Red Hat Enterprise Virtualization Manager installation.

#### Summary

The history data warehouse can also be accessed remotely. All of the PostgreSQL configuration files that you need to change are located in **/var/lib/pgsql/data**

#### Procedure 19.1. Enabling Remote Access to the History Database

1. Edit the **postgresql.conf** file, and add two parameters.

```
# listen_addresses = "*"

```

If you want to allow access by specific hosts, use a comma separated list of IP addresses or hostnames instead of **"\*"**.

```
# ssl=on

```

2. Edit **pg\_hba.conf** to add:

```
hostssl    all             all             <net address/mask>
md5

```

Replace **<net address/mask>** with the IP address and netmask of allowed hosts. For example, **192.168.0.0/24**.

3. A. You can use the certificate and key created by the Red Hat Enterprise Virtualization Manager at installation, **/etc/pki/ovirt-engine/certs/engine.cer** and **/etc/pki/ovirt-engine/keys/engine\_id\_rsa** respectively.
  - a. PostgreSQL requires that the certificate and key be named **server.crt** and **server.key**, respectively. You can use the **ln -s** to create a soft link to put the Manager's certificate and key files where PostgreSQL expects them.

```
# ln -s /etc/pki/ovirt-engine/certs/engine.cer
/var/lib/pgsql/data/server.crt

```

```
# ln -s /etc/pki/ovirt-engine/keys/engine_id_rsa
/var/lib/pgsql/data/server.key

```

- b. Alternatively, you can create a new certificate and key file. Using the commands documented in the [PostgreSQL Manual](#), create **server.crt** and **server.key** files.

4. Stop the engine service.

```
# service ovirt-engine stop

```

## 5. Restart the PostgreSQL service

```
# service postgresql restart
```

## 6. Un-deploy and re-deploy the Jasper WAR files

a. Remove the **.deployed** file:

```
# rm /var/lib/ovirt-engine/deployments/rhev-m-reports.war.deployed
```

This causes an **.undeployed** file to be created: **/var/lib/ovirt-engine/deployments/rhev-m-reports.war.undeployed**

b. When the **.undeployed** file is there, you know the app is down. Remove the **.undeployed** file:

```
# rm /var/lib/ovirt-engine/deployments/rhev-m-reports.war.undeployed
```

## c. Then you can redeploy:

```
# touch /var/lib/ovirt-engine/deployments/rhev-m-reports.war.dodeploy
```

Which causes a **.deployed** file to be created: **/var/lib/ovirt-engine/deployments/rhev-m-reports.war.deployed**

7. Add an **iptables** rule to allow external machines to connect to the Manager and access PostgreSQL. For example, in a default **iptables** configuration, the following will insert a new rule after the SSH rule:

```
iptables -I INPUT 5 -p tcp -m state --state NEW --dport 5432 -j ACCEPT
```

## 8. Start the engine service.

```
# service ovirt-engine start
```

**Result**

You have enabled a remote user with acceptable credentials to access the Red Hat Enterprise Virtualization history database.

[Report a bug](#)

**19.2.6. Allowing Read-only Access to the History Database****Summary**

You want to allow access to your history database without allowing edits. You need to create a read-only Postgres user that can log in and read from the *ovirt\_engine\_history* database. This procedure assumes the Postgres 8 series database, and should be executed locally to your database.

**Procedure 19.2. Allowing Read-only Access to the History Database**1. Create the user that you will grant read-only access to the history database with, called *history-readonly*.

```
# psql -U postgres -c "CREATE ROLE history-readonly WITH LOGIN;"
ovirt_engine_history
```

2. Grant the *history-readonly* user permission to connect to the history database.

```
# psql -U postgres -c "GRANT CONNECT ON DATABASE mydb TO history-readonly;"
ovirt_engine_history
```

3. Grant usage of the public schema to the *history-readonly* user.

```
# psql -U postgres -c "GRANT USAGE ON SCHEMA public TO history-readonly;"
ovirt_engine_history
```

4. Generate the rest of the permissions to grant to the *history-readonly* user and save them to a file.

```
# psql -U postgres -c "SELECT 'GRANT SELECT ON ' || relname || ' TO history-readonly;' FROM pg_class JOIN pg_namespace ON pg_namespace.oid = pg_class.relnamespace WHERE nspname = 'public' AND relkind IN ('r', 'v');" --
pset=tuples_only=on ovirt_engine_history > grant.sql
```

5. Use the file you generated in the previous step to grant permissions to the *history-readonly* user.

```
# psql -U postgres -f grant.sql ovirt_engine_history
```

6. Remove the file you used to grant permissions to *history-readonly* user.

```
# rm grant.sql
```

**Result**

You can now access the *ovirt\_engine\_history* database with the *history-readonly* user name using the **psql -U history-readonly ovirt\_engine\_history** command. Select statements against tables and views in the *ovirt\_engine\_history* database succeed, while modifications fail.

[Report a bug](#)

**19.2.7. History Database Report Examples**

The following examples provide an introduction to reports produced from queries to the *ovirt\_engine\_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.

**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 19.1. 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 host_configuration, host_samples_history
 where host_configuration.host_id = host_samples_history.host_id
 and host_name = 'example.labname.abc.company.com'
 and host_configuration.history_id in (select max(a.history_id)
   from host_configuration as a
  where host_configuration.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 19.30. 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 19.12, "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 19.12. Single host utilization line graph****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 19.2. 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 host_hourly_history
 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 19.31. Resource utilization across all hosts example data**

Hour	CPU	Memory
0	39	40
1	38	38
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 19.13, "Total system utilization line graph"](#) was produced using the **Chart Wizard** tool in **OpenOffice.org Calc**.




Figure 19.13. Total system utilization line graph

### 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 19.3.

```
SELECT vm_name
FROM vm_configuration
  inner join latest_tag_relations_history on (vm_configuration.vm_id =
latest_tag_relations_history.entity_id)
  inner join latest_tag_details on (latest_tag_details.tag_id =
latest_tag_relations_history.parent_id)
WHERE getpathinnames(latest_tag_details.history_id) like '/root/tlv%'
```

This query returns a table of data with all virtual machine names that are attached to this tag:

Table 19.32. Tag Filtering of Latest VM Configuration

vm_name
RHEL6-Pool-67
RHEL6-Pool-5
RHEL6-Pool-6
RHEL6-23

### 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 19.4.

```
SELECT vm_name, vm_type, operating_system
FROM vm_configuration
  inner join enum_translator as vm_type_value on (vm_type_value.enum_type =
'VM_TYPE' and vm_configuration.vm_type = vm_type_value.enum_key)
  inner join enum_translator as os_value on (os_value.enum_type = 'OS_TYPE' and
vm_configuration.operating_system = os_value.enum_key)
```

This query returns a table of virtual machines with OS and VM Type data:

Table 19.33. 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

[Report a bug](#)

## 19.2.8. Statistics History Views

### 19.2.8.1. Statistics History Views

This section describes the statistics history views available to the user for querying and generating reports.

[Report a bug](#)

### 19.2.8.2. Datacenter Statistics Views

Historical statistics for each data center in the system.

Table 19.34. v3\_2\_statistics\_datacenters\_resources\_usage\_dailyv3\_2\_statistics\_datacenters\_resources\_usage\_hourlyv3\_2\_statistics\_datacenters\_resources\_usage\_samples

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	» -1 - <b>Unknown Status</b> (used only to indicate a problem with the ETL -- PLEASE NOTIFY SUPPORT) » 1 - <b>Up</b> » 2 - <b>Maintenance</b> » 3 - <b>Problematic</b>
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 <b>Up</b> and minutes_in_status of 55, the other will have a datacenter_status of <b>Maintenance</b> and a minutes_in_status of 5.
datacenter_configuration_version	integer	The data center configuration version at the time of sample.

[Report a bug](#)**19.2.8.3. Storage Domain Statistics Views**

Historical statistics for each storage domain in the system.

Table 19.35. v3\_2\_statistics\_storage\_domains\_resources\_usage\_dailyv3\_2\_statistics\_storage\_domains\_resources\_usage\_hourlyv3\_2\_statistics\_storage\_domains\_resources\_usage\_samples

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.

[Report a bug](#)**19.2.8.4. Host Statistics Views**

Historical statistics for each host in the system.

Table 19.36. v3\_2\_statistics\_hosts\_resources\_usage\_dailyv3\_2\_statistics\_hosts\_resources\_usage\_hourlyv3\_2\_statistics\_hosts\_resources\_usage\_samples

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).
host_id	uuid	Unique ID of the host in the system.
host_status	smallint	<ul style="list-style-type: none"> <li>» -1 - <b>Unknown Status</b> (used only to indicate a problem with the ETL -- PLEASE NOTIFY SUPPORT)</li> <li>» 1 - <b>Up</b></li> <li>» 2 - <b>Maintenance</b></li> <li>» 3 - <b>Problematic</b></li> </ul>
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 <b>Up</b> and minutes_in_status of 55, the other will have a status of <b>Down</b> 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.
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_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.
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.



[Report a bug](#)

#### 19.2.8.5. Host Interface Statistics Views

Historical statistics for each host network interface in the system.

**Table 19.37. v3\_2\_statistics\_hosts\_interfaces\_resources\_usage\_dailylv3\_2\_statistics\_hosts\_interfaces\_resources\_usage\_hourlylv3\_2\_statistics\_hosts\_interfaces\_resources\_usage\_samples**

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, 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.

[Report a bug](#)

#### 19.2.8.6. Virtual Machine Statistics Views

Historical statistics for the virtual machines in the system.

Table 19.38. v3\_2\_statistics\_vms\_resources\_usage\_dailylv3\_2\_statistics\_vms\_resources\_usage\_hourlylv3\_2\_statistics\_vms\_resources\_usage\_samples

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"> <li>▷ -1 - Unknown Status (used only to indicate problems with the ETL -- PLEASE NOTIFY SUPPORT)</li> <li>▷ 0 - Down</li> <li>▷ 1 - Up</li> <li>▷ 2 - Paused</li> <li>▷ 3 - Problematic</li> </ul>
minutes_in_status	decimal	The total number of minutes that the virtual machine was in the status shown in the status 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.
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.

[Report a bug](#)**19.2.8.7. Virtual Machine Interface Statistics Views**

Historical statistics for the virtual machine network interfaces in the system.

Table 19.39. v3\_2\_statistics\_vms\_interfaces\_resources\_usage\_dailylv3\_2\_statistics\_vms\_interfaces\_resources\_usage\_hourlylv3\_2\_statistics\_vms\_interfaces\_resources\_usage\_samples

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 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.

[Report a bug](#)**19.2.8.8. Virtual Machine Disk Statistics Views**

Historical statistics for the virtual disks in the system.

Table 19.40. v3\_2\_statistics\_vms\_disks\_resources\_usage\_dailylv3\_2\_statistics\_vms\_disks\_resources\_usage\_hourlylv3\_2\_statistics\_vms\_disks\_resources\_usage\_samples

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"> <li>» 0 - Unassigned</li> <li>» 1 - OK</li> <li>» 2 - Locked</li> <li>» 3 - Invalid</li> <li>» 4 - Illegal</li> </ul>
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 an hour, two rows will show for this hour. One will have a status of <b>Locked</b> and minutes_in_status of 55, the other will have a status of OK and a minutes_in_status of 5.
vm_disk_actual_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.
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_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.
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.
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.

[Report a bug](#)

## 19.2.9. Configuration History Views

### 19.2.9.1. Configuration History Views

This section describes the configuration views available to the user for querying and generating reports.

#### Note

`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.

[Report a bug](#)

### 19.2.9.2. Data Center Configuration

Data centers configuration history in the system.

**Table 19.41. v3\_2\_configuration\_history\_datacenters**

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"> <li>» 0 - Unknown</li> <li>» 1 - NFS</li> <li>» 2 - FCP</li> <li>» 3 - iSCSI</li> <li>» 4 - Local</li> <li>» 6 - All</li> </ul>
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.

[Report a bug](#)

### 19.2.9.3. Datacenter Storage Domain Map

A historical map showing the relationships between storage domains and data centers in the system.

**Table 19.42. v3\_2\_map\_history\_datacenters\_storage\_domains**

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.

[Report a bug](#)

### 19.2.9.4. Storage Domain Configuration

Storage domains configuration history in the system.

**Table 19.43. v3\_2\_configuration\_history\_storage\_domains**

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"> <li>» 0 - Data (Master)</li> <li>» 1 - Data</li> <li>» 2 - ISO</li> <li>» 3 - Export</li> </ul>
storage_type	smallint	<ul style="list-style-type: none"> <li>» 0 - Unknown</li> <li>» 1 - NFS</li> <li>» 2 - FCP</li> <li>» 3 - iSCSI</li> <li>» 4 - Local</li> <li>» 6 - All</li> </ul>
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.

[Report a bug](#)

### 19.2.9.5. Cluster Configuration

Clusters configuration history in the system.

**Table 19.44. v3\_2\_configuration\_history\_clusters**

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.

[Report a bug](#)

#### 19.2.9.6. Host Configuration

Host configuration history in the system.

**Table 19.45. v3\_2\_configuration\_history\_hosts**

Name	Type	Description
history_id	integer	The ID of the configuration version in the history database.
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"> <li>» 0 - RHEL Host</li> <li>» 2 - RHEV Hypervisor Node</li> </ul>
fqdn_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.

[Report a bug](#)

#### 19.2.9.7. Host Interface Configuration

Host interface configuration history in the system.

Table 19.46. v3\_2\_configuration\_history\_hosts\_interfaces

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"> <li>» 0 - rt18139_pv</li> <li>» 1 - rt18139</li> <li>» 2 - e1000</li> <li>» 3 - pv</li> </ul>
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.

[Report a bug](#)

#### 19.2.9.8. Virtual Machine Configuration

A list of all virtual machines in the system.

Table 19.47. v3\_2\_configuration\_history\_vms

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"> <li>» 0 - Desktop</li> <li>» 1 - Server</li> </ul>
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"> <li>» 0 - Unknown</li> <li>» 1 - Windows XP</li> <li>» 3 - Windows 2003</li> <li>» 4 - Windows 2008</li> <li>» 5 - Other Linux</li> <li>» 6 - Other</li> <li>» 7 - RHEL 5</li> <li>» 8 - RHEL 4</li> <li>» 9 - RHEL 3</li> <li>» 10 - Windows2003 x64</li> <li>» 11 - Windows 7</li> <li>» 12 - Windows 7 x64</li> <li>» 13 - RHEL 5 x64</li> <li>» 14 - RHEL 4 x64</li> <li>» 15 - RHEL 3 x64</li> <li>» 16 - Windows 2008 x64</li> <li>» 17 - Windows 2008R2 x64</li> <li>» 18 - RHEL 6</li> <li>» 19 - RHEL 6 x64</li> </ul>
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.
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.

[Report a bug](#)**19.2.9.9. Virtual Machine Interface Configuration History**

Virtual interfaces configuration history in the system

**Table 19.48. v3\_2\_configuration\_history\_vms\_interfaces**

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. » 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.

[Report a bug](#)**19.2.9.10. Disks to Virtual Machine Map**

A historical map showing the relationships between virtual disks and virtual machines in the system.

**Table 19.49. disks\_vm\_map table**

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.
detach_date	timestamp with time zone	The date the virtual disk was detached from the virtual machine.

[Report a bug](#)**19.2.9.11. Virtual Machine Disk Configuration**

Virtual disks configuration history in the system.

**Table 19.50. v3\_2\_configuration\_history\_vms\_disks**

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.
vm_disk_description	varchar(4000)	As displayed in the edit dialog.
storage_domain_id	uuid	The ID of the storage domain this disk image belongs to.
vm_disk_size_mb	integer	The defined size of the disk in megabytes (MB).
vm_disk_type	integer	As displayed in the edit dialog. Only System and data are currently used. » 0 - Unassigned » 1 - System » 2 - Data » 3 - Shared » 4 - Swap » 5 - Temp
vm_disk_format	integer	As displayed in the edit dialog. » 3 - Unassigned » 4 - COW » 5 - RAW
vm_disk_interface	integer	» 0 - IDE » 1 - SCSI (not supported) » 2 - VirtIO
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.

[Report a bug](#)**19.3. Dashboards****19.3.1. Dashboards**

A dashboard is a collection of related reports that provide a summary of resource usage in the



virtualized environment. Dashboards feature an active control panel, allowing quick adjustment of the parameters. Though a dashboard cannot be exported or printed, each of the reports in a dashboard can be opened separately to export, print, save, or adjust the data.

Dashboards can be created and configured using the **Designer**, in the Reports Portal. For more information on dashboards, consult the JasperReports documentation by clicking the **Help** in the top menu bar of the Reports Portal.

[Report a bug](#)

### 19.3.2. Header Inventory Dashboard

The **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 can be modified for this dashboard:

Table 19.51. Inventory Dashboard Parameters

Parameter	Description
<b>Show Deleted Entities?</b>	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
<b>Period Range</b>	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 <b>Dates</b> parameter.
<b>Dates</b>	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.
<b>Data Center</b>	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 <b>Cluster</b> parameter includes only clusters in the selected data center.

[Report a bug](#)

### 19.3.3. Header Trends Dashboard

The **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 can be modified for this dashboard:

Table 19.52. Trends Dashboard Parameters

Parameter	Description
<b>Show Deleted Entities?</b>	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
<b>Period Range</b>	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 <b>Dates</b> parameter.
<b>Dates</b>	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.
<b>Data Center</b>	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 <b>Cluster</b> parameter includes only clusters in the selected data center.

[Report a bug](#)

### 19.3.4. Header Uptime Dashboard

The **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 can be modified for this dashboard:

Table 19.53. Uptime Dashboard Parameters

Parameter	Description
<b>Show Deleted Entities?</b>	The report includes deleted objects, such as data centers, clusters, and hosts removed from the environment.
<b>Period Range</b>	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 <b>Dates</b> parameter.
<b>Dates</b>	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.
<b>Data Center</b>	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 <b>Cluster</b> parameter includes only clusters in the selected data center.

[Report a bug](#)

### 19.3.5. Integrated Reporting Dashboard in the Red Hat Enterprise Virtualization Administration Portal

The Administration Portal also features dashboards for data centers, clusters, and the overall

environment. Select the appropriate resource in tree mode and click the **Dashboard** resource tab to display the dashboard information in the results list.

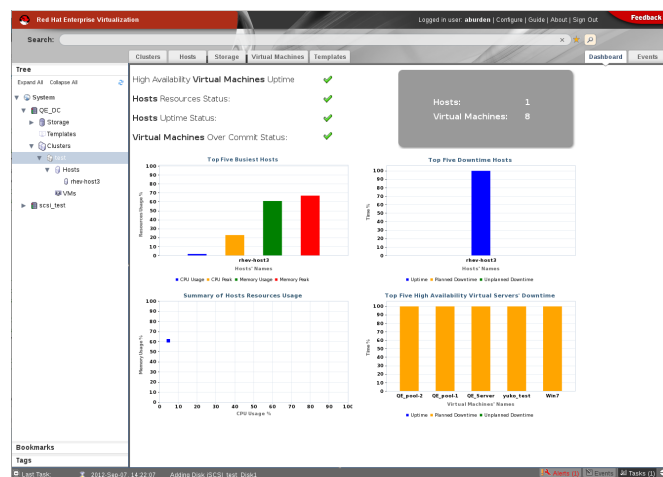


Figure 19.14. Reports Dashboard

The dashboards accessible in the Administration Portal are used for viewing data, as such they do not have an active control panel. Configure these dashboards in the Reports Portal by editing **Datacenter Dashboard**, **Cluster Dashboard**, and **System Dashboard**.

[Report a bug](#)

## Firewalls

### A.1. Red Hat Enterprise Virtualization Manager Firewall Requirements

The Red Hat Enterprise Virtualization Manager requires that a number of ports be opened to allow network traffic through the system's firewall. The `rhevms-setup` script is able to configure the firewall automatically, but this will overwrite any pre-existing firewall configuration.

Where an existing firewall configuration exists the firewall rules required by the Manager must instead be manually inserted into it. The `rhevms-setup` command will save a list of the `iptables` rules required in the `/usr/share/ovirt-engine/conf/iptables.example` file.

The firewall configuration documented here assumes a default configuration. Where non-default HTTP and HTTPS ports are chosen during installation adjust the firewall rules to allow network traffic on the ports that were selected - not the default ports (**80** and **443**) listed here.

Table A.1. Red Hat Enterprise Virtualization Manager Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
-	ICMP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Hypervisor(s)</li> <li>Red Hat Enterprise Linux host(s)</li> </ul>	Red Hat Enterprise Virtualization Manager	When registering to the Red Hat Enterprise Virtualization Manager, virtualization hosts send an ICMP ping request to the Manager to confirm that it is online.
22	TCP	<ul style="list-style-type: none"> <li>System(s) used for maintenance of the Manager including backend configuration, and software upgrades.</li> </ul>	Red Hat Enterprise Virtualization Manager	SSH (optional)
80, 443	TCP	<ul style="list-style-type: none"> <li>Administration Portal clients</li> <li>User Portal clients</li> <li>Red Hat Enterprise Virtualization Hypervisor(s)</li> <li>Red Hat Enterprise Linux host(s)</li> <li>REST API clients</li> </ul>	Red Hat Enterprise Virtualization Manager	Provides HTTP and HTTPS access to the Manager.



#### Important

In environments where the Red Hat Enterprise Virtualization Manager is also required to export NFS storage, such as an ISO Storage Domain, additional ports must be allowed through the firewall. Grant firewall exceptions for the ports applicable to the version of NFS in use:

##### NFSv4

- TCP port **2049** for NFS.

##### NFSv3

- TCP and UDP port **2049** for NFS.
- TCP and UDP port **111** (`rpcbind/sunrpc`).
- TCP and UDP port specified with `MOUNTD_PORT="port"`
- TCP and UDP port specified with `STATD_PORT="port"`
- TCP port specified with `LOCKD_TCPPORT="port"`
- UDP port specified with `LOCKD_UDPPORT="port"`

The `MOUNTD_PORT`, `STATD_PORT`, `LOCKD_TCPPORT`, and `LOCKD_UDPPORT` ports are configured in the `/etc/sysconfig/nfs` file.

[Report a bug](#)

### A.2. Virtualization Host Firewall Requirements

Both Red Hat Enterprise Linux hosts and Red Hat Enterprise Virtualization Hypervisors require that a number of ports be opened to allow network traffic through the system's firewall. In the case of the Red Hat Enterprise Virtualization Hypervisor these firewall rules are configured automatically. For Red Hat Enterprise Linux hosts however it is necessary to manually configure the firewall.

Table A.2. Virtualization Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
22	TCP	» Red Hat Enterprise Virtualization Manager	» Red Hat Enterprise Virtualization Hypervisor(s) » Red Hat Enterprise Linux host(s)	Secure Shell (SSH) access.
5900 - 6411	TCP	» Administration Portal clients » User Portal clients	» Red Hat Enterprise Virtualization Hypervisor(s) » Red Hat Enterprise Linux host(s)	Remote guest console access via VNC and SPICE. These ports must be open to facilitate client access to virtual machines.
5989	TCP, UDP	» Common Information Model Object Manager (CIMOM)	» Red Hat Enterprise Virtualization Hypervisor(s) » Red Hat Enterprise Linux host(s)	Used by Common Information Model Object Managers (CIMOM) to monitor virtual machines running on the virtualization host. If you wish to use a CIMOM to monitor the virtual machines in your virtualization environment then you must ensure that this port is open.
16514	TCP	» Red Hat Enterprise Virtualization Hypervisor(s) » Red Hat Enterprise Linux host(s)	» Red Hat Enterprise Virtualization Hypervisor(s) » Red Hat Enterprise Linux host(s)	Virtual machine migration using <b>libvirt</b> .
49152 - 49216	TCP	» Red Hat Enterprise Linux Hypervisor(s) » Red Hat Enterprise Linux host(s)	» Red Hat Enterprise Linux Hypervisor(s) » Red Hat Enterprise Linux host(s)	Virtual machine migration and fencing using VDSM. These ports must be open facilitate both automated and manually initiated migration of virtual machines.
54321	TCP	» Red Hat Enterprise Virtualization Manager » Red Hat Enterprise Virtualization Hypervisor(s) » Red Hat Enterprise Linux host(s)	» Red Hat Enterprise Virtualization Hypervisor(s) » Red Hat Enterprise Linux host(s)	VDSM communications with the Manager and other virtualization hosts.

**Example A.1. Option Name: IPTablesConfig**

Recommended (default) values: Automatically generated by vdsd bootstrap script

```
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
# vdsd
-A INPUT -p tcp --dport 54321 -j ACCEPT
# libvirt tls
-A INPUT -p tcp --dport 16514 -j ACCEPT
# SSH
-A INPUT -p tcp --dport 22 -j ACCEPT
# guest consoles
-A INPUT -p tcp -m multiport --dports 5634:6166 -j ACCEPT
# migration
-A INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
# snmp
-A INPUT -p udp --dport 161 -j ACCEPT
# Reject any other input traffic
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -m physdev ! --physdev-is-bridged -j REJECT --reject-with icmp-host-prohibited
COMMIT
```

[Report a bug](#)**A.3. Directory Server Firewall Requirements**

Red Hat Enterprise Virtualization requires a directory server to support user authentication. A number of ports must be opened in the directory server's firewall to support GSS-API authentication as used by the Red Hat Enterprise Virtualization Manager.

Table A.3. Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
88, 464	TCP, UDP	» Red Hat Enterprise Virtualization Manager	» Directory server	Kerberos authentication.
389, 636	TCP	» Red Hat Enterprise Virtualization Manager	» Directory server	Lightweight Directory Access Protocol (LDAP) and LDAP over SSL.

[Report a bug](#)

#### A.4. Database Server Firewall Requirements

Red Hat Enterprise Virtualization supports the use of a remote database server. If you plan to use a remote database server with Red Hat Enterprise Virtualization then you must ensure that the remote database server allows connections from the Manager.

Table A.4. Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
5432	TCP, UDP	» Red Hat Enterprise Virtualization Manager	» PostgreSQL database server	Default port for PostgreSQL database connections.

If you plan to use a local database server on the Manager itself, which is the default option provided during installation, then no additional firewall rules are required.

[Report a bug](#)

## VDSM and Hooks

### B.1. 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.

[Report a bug](#)

### B.2. VDSM Hooks

VDSM is extensible via hooks. Hooks are scripts executed on the host when key events occur. When a supported event occurs VDSM runs any executable hook scripts in `/usr/libexec/vdsm/hooks/nn_event-name` on the host in alphanumeric order. By convention each hook script is assigned a two digit number, included at the front of the file name, to ensure that the order in which the scripts will be run in is clear. You are able to create hook scripts in any programming language, Python will however be used for the examples contained in this chapter.

Note that all scripts defined on the host for the event are executed. If you require that a given hook is only executed for a subset of the virtual machines which run on the host then you must ensure that the hook script itself handles this requirement by evaluating the **Custom Properties** associated with the virtual machine.



#### Warning

VDSM hooks can interfere with the operation of Red Hat 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.

[Report a bug](#)

### B.3. Extending VDSM with Hooks

This chapter describes how to extend VDSM with event-driven hooks. Extending VDSM with hooks is an experimental technology, and this chapter is intended for experienced developers. 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.

[Report a bug](#)

### B.4. Supported VDSM Events

Table B.1. Supported VDSM Events

Name	Description
before_vm_start	Before virtual machine starts.
after_vm_start	After virtual machine starts.
before_vm_cont	Before virtual machine continues.
after_vm_cont	After virtual machine continues.
before_vm_pause	Before virtual machine pauses.
after_vm_pause	After virtual machine pauses.
before_vm_hibernate	Before virtual machine hibernates.
after_vm_hibernate	After virtual machine hibernates.
before_vm_dehibernate	Before virtual machine de-hibernates.
after_vm_dehibernate	After virtual machine de-hibernates.
before_vm_migrate_source	Before virtual machine migration, run on the source hypervisor host from which the migration is occurring.
after_vm_migrate_source	After virtual machine migration, run on the source hypervisor host from which the migration is occurring.
before_vm_migrate_destination	Before virtual machine migration, run on the destination hypervisor host to which the migration is occurring.
after_vm_migrate_destination	After virtual machine migration, run on the destination hypervisor host to which the migration is occurring.
after_vm_destroy	After virtual machine destruction.
before_vdsm_start	Before VDSM is started on the hypervisor host. <b>before_vdsm_start</b> 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. <b>after_vdsm_stop</b> hooks are executed as the user root, and do not inherit the environment of the VDSM process.
before_nic_hotplug	Before the NIC is hot plugged into the virtual machine.
after_nic_hotplug	After the NIC is hot plugged into the virtual machine.
before_nic_hotunplug	Before the NIC is hot unplugged from the virtual machine.
after_nic_hotunplug	After the NIC is hot unplugged from the virtual machine.
after_nic_hotplug_fail	After hot plugging the NIC to the virtual machine fails.
after_nic_hotunplug_fail	After hot unplugging the NIC from the virtual machine fails.

[Report a bug](#)

### B.5. The VDSM Hook Environment

Most hook scripts are run as the **vdsm** user and inherit the environment of the VDSM process. The

exceptions are hook scripts triggered by the **before\_vdsm\_start** and **after\_vdsm\_stop** events. Hook scripts triggered by these events run as the **root** user and do not inherit the environment of the VDSM process.

[Report a bug](#)

## B.6. The VDSM Hook Domain XML Object

When hook scripts are started the **\_hook\_domxml** variable is appended to the environment. This variable contains the path of the libvirt domain XML representation of the relevant virtual machine. The exceptions to this rule are **\*\_nic\_hotplug\_\*** and **\*\_nic\_hotunplug\_\*** hooks, in which the **\_hook\_domxml** variable contains the XML representation of the NIC, not the 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.

[Report a bug](#)

## B.7. Defining Virtual Machine 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-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 **ovirt-engine** service must be restarted for it to take effect.

### Example B.1. Defining *smartcard* Custom Property

1. Check the existing custom properties defined by the **UserDefinedVMProperties** configuration key using **rhev-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.

```
# rhev-config -g UserDefinedVMProperties
UserDefinedVMProperties : version: general
UserDefinedVMProperties : version: 2.2
UserDefinedVMProperties : memory=^[0-9]+$ version: 3.2
```

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**.

```
# rhev-config -s UserDefinedVMProperties='memory=^[0-9]+$;smartcard=^(true|false)$' --cver=3.2
```

3. Verify that the custom properties defined by the **UserDefinedVMProperties** configuration key now match your expectations.

```
# rhev-config -g UserDefinedVMProperties
UserDefinedVMProperties : version: general
UserDefinedVMProperties : version: 2.2
UserDefinedVMProperties : memory=^[0-9]+$;smartcard=^(true|false)$
version: 3.2
```

4. Finally, the **ovirt-engine** service must be restarted for the configuration change to take effect.

```
# service ovirt-engine restart
```

[Report a bug](#)

## B.8. Setting Virtual Machine Custom Properties

### Prerequisites:

- » [Section B.7, "Defining Virtual Machine 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 Virtual Machine(s)** dialog box. Custom properties set from the **Run Virtual Machine(s)** dialog box will only apply to the virtual machine until it is next shutdown.

The **Custom Properties** tab provides a facility for you to select from the list of defined custom properties. Once you select a custom property key an additional field will display allowing you to enter a value for that key. Add additional key/value pairs by clicking the **+** button and remove them by clicking the **-** button.

### See Also:

- » [Section 8.4.3.8, "Virtual Machine Custom Properties Settings Explained"](#)

[Report a bug](#)

## B.9. Evaluating Virtual Machine Custom Properties in a VDSM Hook

**Prerequisites:**

- [Section B.7, "Defining Virtual Machine Custom Properties"](#)
- [Section B.8, "Setting Virtual Machine Custom Properties"](#)

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 B.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)
```

[Report a bug](#)

**B.10. Using the VDSM Hooking Module**

VDSM ships with a Python hooking module, providing helper functions for VDSM hook scripts. This module is provided as an example, and is only relevant to VDSM hooks written in Python.

The hooking module supports reading of a virtual machine's libvirt XML into a DOM object. Hook scripts can then use Python's built in **xml.dom** library (<http://docs.python.org/release/2.6/library/xml.dom.html>) to manipulate the object.

The modified object can then be saved back to libvirt XML using the hooking module. The hooking module provides the following functions to support hook development:

**Table B.2. Hooking module functions**

Name	Argument	Description
<b>tobool</b>	string	Converts a string "true" or "false" to a Boolean value
<b>read_domxml</b>	-	Reads the virtual machine's libvirt XML into a DOM object
<b>write_domxml</b>	DOM object	Writes the virtual machine's libvirt XML from a DOM object

[Report a bug](#)

**B.11. VDSM Hook Execution**

**before\_vm\_start** scripts can edit the domain XML in order to change VDSM's definition of a virtual machine before it reaches libvirt. Caution must be exercised in doing so. Hook scripts have the potential to disrupt the operation of VDSM, and buggy scripts can result in outages to the Red Hat Enterprise Virtualization environment. In particular, ensure you never change the uuid of the domain, and do not attempt to remove a device from the domain without sufficient background knowledge.

Both **before\_vdsm\_start** and **after\_vdsm\_stop** hook scripts are run as the **root** user. Other hook scripts that require **root** access to the system must be written to use the **sudo** command for privilege escalation. To support this the **/etc/sudoers** must be updated to allow the **vdsm** user to use **sudo** without reentering a password. This is required as hook scripts are executed non-interactively.

**Example B.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", "root", "/my_file"] )
```

The standard error stream of hook scripts is collected in VDSM's log. This information is used to debug hook scripts.

[Report a bug](#)

**B.12. VDSM Hook Return Codes**

Hook scripts must return one of the return codes shown in [Table B.3, "Hook Return Codes"](#). The return code will determine whether further hook scripts are processed by VDSM.



Table B.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

[Report a bug](#)

### B.13. VDSM Hook Examples

The example hook scripts provided in this section are strictly not supported by Red Hat. You must ensure that any and all hook scripts that you install to your system, regardless of source, are thoroughly tested for your environment.

#### Example B.4. 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)
```

[Report a bug](#)

## Red Hat Enterprise Virtualization User Interface Plugins

### C.1. Red Hat Enterprise Virtualization User Interface Plugins

Red Hat Enterprise Virtualization supports plugins that expose non-standard features. This makes it easier to use the Red Hat Enterprise Virtualization Administration Portal to integrate with other systems. Each interface plugin represents a set of user interface extensions that can be packaged and distributed for use with Red Hat Enterprise Virtualization.

Red Hat Enterprise Virtualization's User Interface plugins integrate with the Administration Portal directly on the client using the JavaScript programming language. Plugins are invoked by the Administration Portal and executed in the web browser's JavaScript runtime. User Interface plugins can use the JavaScript language and its libraries.

At key events during runtime, the Administration Portal invokes individual plugins via event handler functions representing Administration-Portal-to-plugin communication. Although the Administration Portal supports multiple event-handler functions, a plugin declares functions which are of interest only to its implementation. Each plugin must register relevant event handler functions as part of the plugin bootstrap sequence before the plugin is put to use by the administration portal.

To facilitate the plugin-to-Administration-Portal communication that drives the User Interface extension, the Administration Portal exposes the plugin API as a global (top-level) `pluginApi` JavaScript object that individual plugins can consume. Each plugin obtains a separate `pluginApi` instance, allowing the Administration Portal to control plugin API-function invocation for each plugin with respect to the plugin's lifecycle.

[Report a bug](#)

### C.2. Red Hat Enterprise Virtualization User Interface Plugin Lifecycle

#### C.2.1. Red Hat Enterprise Virtualization User Interface Plugin Lifecycle

The basic lifecycle of a User Interface Plugin divides into three stages:

1. Plugin discovery.
2. Plugin loading.
3. Plugin bootstrapping.

[Report a bug](#)


#### C.2.2. Red Hat Enterprise Virtualization User Interface Plugin Discovery

Creating plugin descriptors is the first step in the plugin discovery process. Plugin descriptors contain important plugin metadata as well as (optional) default plugin-specific configurations.

As part of handling administration portal HTML page requests (**HTTP GET**), User Interface plugin infrastructure attempts to discover and load plugin descriptors from your local file system. For each plugin descriptor, the infrastructure also attempts to load corresponding plugin user configurations used to override default plugin-specific configurations (if any exist) and tweak plugin runtime behavior. Plugin user configuration is optional. After loading descriptors and corresponding user configuration files, oVirt Engine aggregates User Interface plugin data and embeds it into the administration portal HTML page for runtime evaluation.

By default, plugin descriptors reside in `$ENGINE_USR/ui-plugins`, with a default mapping of `ENGINE_USR=/usr/share/ovirt-engine` as defined by oVirt Engine local configuration. Plugin descriptors are expected to comply with JSON format specifications, but plugin descriptors allow Java/C++ style comments (of both `/*` and `//` varieties) in addition to the JSON format specifications.

By default, plugin user configuration files reside in `$ENGINE_ETC/ui-plugins`, with a default mapping of `ENGINE_ETC=/etc/ovirt-engine` as defined by oVirt Engine local configuration. Plugin user configuration files are expected to comply with same content format rules as plugin descriptors.

 **Note**

Plugin user configuration files generally follow the `<descriptorFileName>-config.json` naming convention.

[Report a bug](#)

#### C.2.3. Red Hat Enterprise Virtualization User Interface Plugin Loading

After a plugin has been discovered and its data is embedded into the administration portal HTML page, administration portal tries to load the plugin as part of application startup (unless you have configured it not to load as part of application startup).

For each plugin that has been discovered, the administration portal creates an HTML iframe element that is used to load its host page. The plugin host page is necessary to begin the plugin bootstrap process, which (the bootstrap process) is used to evaluate the plugin code in the context of the plugin's iframe element. User interface plugin infrastructure supports serving plugin resource files (such as the plugin host page) from the local file system. The plugin host page is loaded into the iframe element and the plugin code is evaluated. After the plugin code is evaluated, the plugin communicates with the administration portal by means of the plugin API.

[Report a bug](#)

#### C.2.4. Red Hat Enterprise Virtualization User Interface Plugin Bootstrapping

A typical plugin bootstrap sequence consists of following steps:

##### Procedure C.1. Plugin Bootstrap Sequence

1. Obtain `pluginApi` instance for the given plugin
2. Obtain runtime plugin configuration object (optional)
3. Register relevant event handler functions
4. Notify UI plugin infrastructure to proceed with plugin initialization

The following code snippet illustrates the above mentioned steps in practice:

```
// Access plugin API using 'parent' due to this code being evaluated within the
// context of an iframe element.
// As 'parent.pluginApi' is subject to Same-Origin Policy, this will only work
// when WebAdmin HTML page and plugin
// host page are served from same origin. WebAdmin HTML page and plugin host page
// will always be on same origin
// when using UI plugin infrastructure support to serve plugin resource files.
var api = parent.pluginApi('MyPlugin');

// Runtime configuration object associated with the plugin (or an empty object).
var config = api.configObject();

// Register event handler function(s) for later invocation by UI plugin
// infrastructure.
api.register({
  // UiInit event handler function.
  UiInit: function() {
    // Handle UiInit event.
    window.alert('Favorite music band is ' + config.band);
  }
});

// Notify UI plugin infrastructure to proceed with plugin initialization.
api.ready();
```

[Report a bug](#)

### C.3. User Interface Plugin-related Files and their Locations

Table C.1. UI Plugin-related Files and their Locations

File	Location	Remarks
Plugin descriptor files (meta-data)	<code>/usr/share/ovirt-engine/ui-plugins/my-plugin.json</code>	
Plugin user configuration files	<code>/etc/ovirt-engine/ui-plugins/my-plugin-config.json</code>	
Plugin resource files	<code>/usr/share/ovirt-engine/ui-plugins/&lt;resourcePath&gt;/PluginHostPage.html</code>	<resourcePath> is defined by the corresponding attribute in the plugin descriptor.

[Report a bug](#)

### C.4. Example User Interface Plugin Deployment

Follow these instructions to create a user interface plugin that runs a **Hello World!** program when you sign in to the Red Hat Enterprise Virtualization Manager administration portal.

#### Procedure C.2. Deploying a Hello World! Plugin

1. Create a plugin descriptor by creating the following file in the Manager at `/usr/share/ovirt-engine/ui-plugins/helloWorld.json`:

```
{
  "name": "HelloWorld",
  "url": "/webadmin/webadmin/plugin/HelloWorld/start.html",
  "resourcePath": "hello-files"
}
```

2. Create the plugin host page by creating the following file in the Manager at `/usr/share/ovirt-engine/ui-plugins/hello-files/start.html`:

```
<!DOCTYPE html><html><head>
<script>
  var api = parent.pluginApi('HelloWorld');
  api.register({
    UiInit: function() { window.alert('Hello world'); }
  });
  api.ready();
</script>
</head><body></body></html>
```

If you have successfully implemented the **Hello World!** plugin, you will see this screen when you sign in to the administration portal:




Figure C.1. A Successful Implementation of the Hello World! Plugin

[Report a bug](#)

### C.5. Installing Red Hat Support Plugin

The Red Hat Support Plugin provides access to Red Hat Access services from the Red Hat Enterprise Virtualization Administration Portal.

#### Procedure C.3. Installing the Red Hat Support Plugin

#### Note

The Red Hat Support Plugin is installed by default in Red Hat Enterprise Virtualization 3.3. The Red Hat Support Plugin is not installed by default in Red Hat Enterprise Virtualization 3.2. It is not necessary to run this procedure in Red Hat Enterprise Virtualization 3.3.

- Use **yum** to install the `redhat-support-plugin-rhev` plugin:

```
# yum install redhat-support-plugin-rhev
```

[Report a bug](#)

## C.6. Using Red Hat Support Plugin

The Red Hat Access Plugin allows you to use Red Hat access services from the Red Hat Enterprise Virtualization Administration Portal. You must log in using your Red Hat login credentials. The Red Hat Access Plugin detects when you are not logged in; if you are not logged in, a login window opens.

### Note

Red Hat Enterprise Virtualization Administration Portal credentials are not the same as a user's Red Hat login.




Figure C.2. Red Hat Support Plugin - Login Window

After logging in, you will be able to access the Red Hat Customer Portal. Red Hat Support Plugin is available in the details pane as well as in several context menus in the Red Hat Enterprise Virtualization Administration Portal. Search the Red Hat Access database using the **Red Hat Search** field. Search results display in the left-hand navigation list in the details pane.

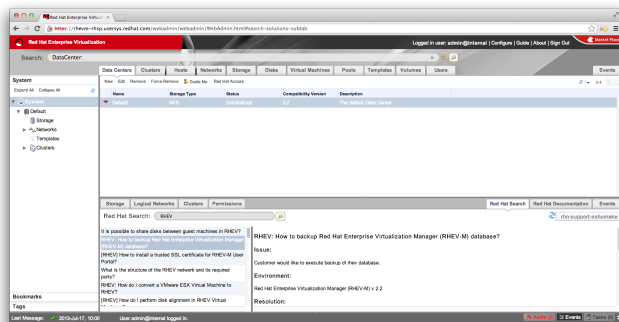


Figure C.3. Red Hat Support Plugin - Query Results in the Left-Hand Navigation List

Right click on context menus in the Red Hat Enterprise Virtualization Administrator Portal to access the Red Hat Support Plugin.

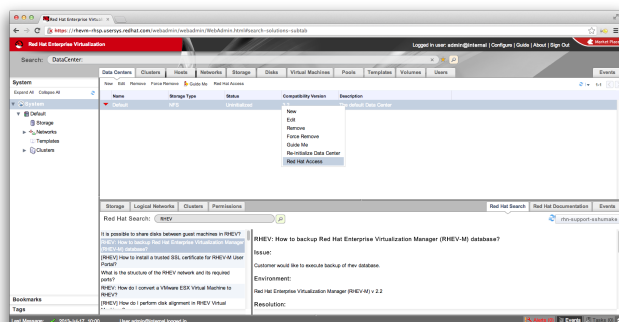


Figure C.4. Right-clicking on a Context Menu to Access Red Hat Support Plugin

Select **Red Hat Access** from the context menu, and a window opens that allows you to search for solutions in the knowledge base.

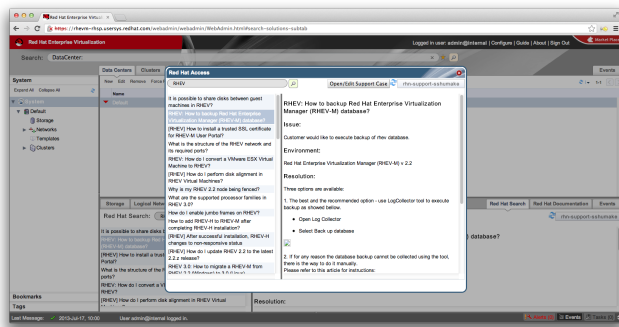


Figure C.5. Red Hat Support Plugin - Solution Search Window

Open a new support case by selecting the **Open New Support Case** radio button.

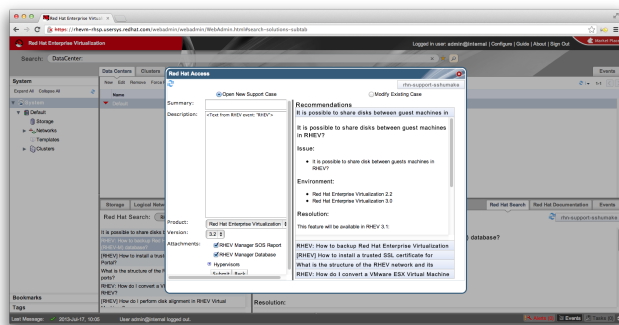


Figure C.6. Red Hat Support Plugin - Opening a New Support Case

Modify an existing support case by selecting the **Modify Existing Case** radio button.

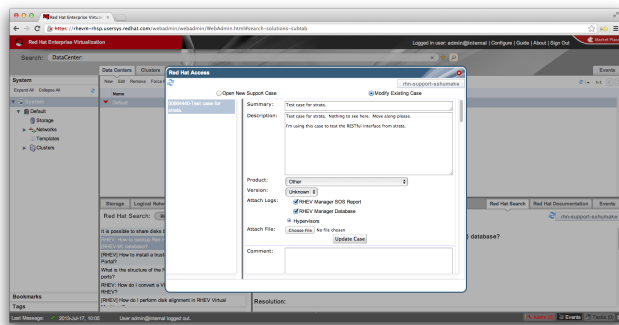


Figure C.7. Red Hat Support Plugin - Modifying a Support Case

Select the **Red Hat Documentation** tab to open the documentation relevant to the part of the Administration Portal currently on the screen.

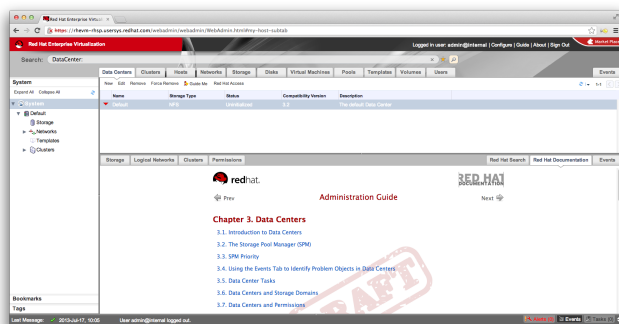


Figure C.8. Red Hat Support Plugin - Accessing Documentation

[Report a bug](#)

## Red Hat Enterprise Virtualization and SSL

### D.1. Replacing the Red Hat Enterprise Virtualization Manager SSL Certificate

#### Summary

You want to use your organization's commercially signed certificate to identify your Red Hat Enterprise Virtualization Manager to users connecting over https.

#### Note

Using a commercially issued certificate for https connections does not affect the certificate used for authentication between your Manager and hosts, they will continue to use the self-signed certificate generated by the Manager.

#### Prerequisites

This procedure requires a *PEM* formatted certificate from your commercial certificate issuing authority, a *.nokey* file, and a *.cer* file. The *.nokey* and *.cer* files are sometimes distributed as a certificate-key bundle in the *P12* format.

This procedure assumes that you have a certificate-key bundle in the *P12* format.

#### Procedure D.1. Replacing the Red Hat Enterprise Virtualization Manager Apache SSL Certificate

1. The Manager has been configured to use `/etc/pki/ovirt-engine/apache-ca.pem`, which is symbolically linked to `/etc/pki/ovirt-engine/ca.pem`. Remove the symbolic link.

```
# rm /etc/pki/ovirt-engine/apache-ca.pem
```

2. Save your commercially issued certificate as `/etc/pki/ovirt-engine/apache-ca.pem`.

```
mv YOUR-3RD-PARTY-CERT.pem /etc/pki/ovirt-engine/apache-ca.pem
```

3. Move your *P12* bundle to `/etc/pki/ovirt-engine/keys/apache.p12`.

4. Extract the key from the bundle.

```
# openssl pkcs12 -in /etc/pki/ovirt-engine/keys/apache.p12 -nocerts -nodes > /etc/pki/ovirt-engine/keys/apache.key.nopass
```

5. Extract the certificate from the bundle.

```
# openssl pkcs12 -in /etc/pki/ovirt-engine/keys/apache.p12 -nokeys > /etc/pki/ovirt-engine/certs/apache.cer
```

6. Restart the Apache server.

```
# service httpd restart
```

#### Result

Your users can now connect to the portals without being warned about the authenticity of the certificate used to encrypt https traffic.

[Report a bug](#)

## Using Search, Bookmarks, and Tags to Find Your Way Around

### E.1. Search

#### E.1.1. Performing Searches in Red Hat Enterprise Virtualization

The Administration Portal enables the management of thousands of resources, such as virtual machines, hosts, users, and more. To perform a search, enter the search query (free-text or syntax-based) in the search bar. Search queries can be saved as bookmarks for future reuse, so you do not have to reenter a search query each time the specific search results are needed.

[Report a bug](#)

#### E.1.2. Search Syntax and Examples

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 E.1. Example Search Queries

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

[Report a bug](#)

#### E.1.3. Search Auto-Completion

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

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

Hosts: Vms.status = down

Table E.2. Example Search Queries Using Auto-Completion

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

[Report a bug](#)

#### E.1.4. Search Result Type Options

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

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

As each type of resource has a unique set of properties and a set of other resource types that it is associated with, each search type has a set of valid syntax combinations. You can also use the auto-complete feature to create valid queries easily.

[Report a bug](#)

#### E.1.5. Search Criteria

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

**<prop><operator><value>**

or

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

##### Examples

The following table describes the parts of the syntax:

Table E.3. Example Search Criteria

Part	Description	Values	Example	Note
prop	The property of the searched-for resource. Can also be the property of a resource type (see <b>obj-type</b> ), or <i>tag</i> (custom tag).	Limit your search to objects with a certain property. For example, search for objects with a <i>status</i> property.	Status	--
obj-type	A resource type that can be associated with the searched-for resource.	These are system objects, like data centers and virtual machines.	Users	--
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	» 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

[Report a bug](#)

### E.1.6. Search: Multiple Criteria and Wildcards

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

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

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

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

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

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

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

[Report a bug](#)

### E.1.7. Search: Determining Search Order

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

For example:

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

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

[Report a bug](#)

### E.1.8. Searching for Data Centers

The following table describes all search options for Data Centers.

Table E.4. Searching for Data Centers

Property (of resource or resource-type)	Type	Description (Reference)
Clusters <i>clusters-prop</i>	Depends on property type	The property of the clusters associated with the data center.
name	String	The name of the data center.
description	String	A description of the data center.
type	String	The type of data center.
status	List	The availability of the data center.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The number of results to display per page.

#### Example

**Datacenter: type = nfs and status != up**

returns a list of data centers with:

- » A storage type of NFS and status other than up



[Report a bug](#)

### E.1.9. Searching for Clusters

The following table describes all search options for clusters.

Table E.5. Searching Clusters

Property (of resource or resource-type)	Type	Description (Reference)
Datacenter <i>datacenter-prop</i>	Depends on property type	The property of the data center associated with the cluster.
Datacenter	String	The data center to which the cluster belongs.
name	String	The unique name that identifies the clusters on the network.
description	String	The description of the cluster.
initialized	String	True or False indicating the status of the cluster.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The number of results to display per page.

#### Example

**Clusters: initialized = true or name = Default**

returns a list of clusters which are:

- » initialized; or
- » named Default

[Report a bug](#)

### E.1.10. Searching for Hosts

The following table describes all search options for hosts.

Table E.6. Searching for Hosts

Property (of resource or resource-type)	Type	Description (Reference)
Vms <i>vms-prop</i>	Depends on property type	The property of the virtual machines associated with the host.
Templates <i>templates-prop</i>	Depends on property type	The property of the templates associated with the host.
Events <i>events-prop</i>	Depends on property type	The property of the events associated with the host.
Users <i>users-prop</i>	Depends on property type	The property of the users associated with the host.
name	String	The name of the host.
status	List	The availability of the host.
cluster	String	The cluster to which the host belongs.
address	String	The unique name that identifies the host on the network.
cpu_usage	Integer	The percent of processing power used.
mem_usage	Integer	The percentage of memory used.
network_usage	Integer	The percentage of network usage.
load	Integer	Jobs waiting to be executed in the <i>run-queue</i> per processor, in a given time slice.
version	Integer	The version number of the operating system.
cpus	Integer	The number of CPUs on the host.
memory	Integer	The amount of memory available.
cpu_speed	Integer	The processing speed of the CPU.
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 number of results to display per page.

#### Example

**Hosts: cluster = Default and Vms.os = rhel6**

returns a list of hosts which:

- » Are part of the Default cluster and host virtual machines running the Red Hat Enterprise Linux 6 operating system.

[Report a bug](#)

**E.1.11. Searching for Networks**

The following table describes all search options for networks.

**Table E.7. Searching for Networks**

Property (of resource or resource-type)	Type	Description (Reference)
Cluster_network. <i>cluster-network-prop</i>	Depends on property type	The property of the cluster associated with the network.
Host_Network. <i>host-network-prop</i>	Depends on property type	The property of the host associated with the network.
name	String	The human readable name that identifies the network.
description	String	Keywords or text describing the network, optionally used when creating the network.
vlanid	Integer	The VLAN ID of the network.
stp	String	Whether Spanning Tree Protocol (STP) is enabled or disabled for the network.
mtu	Integer	The maximum transmission unit for the logical network.
vmnetwork	String	Whether the network is only used for virtual machine traffic.
datacenter	String	The data center to which the network is attached.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

**Example**

**Network:** `mtu > 1500 and vmnetwork = true`

returns a list of networks:

- » with a maximum transmission unit greater than 1500 bytes
- » which are set up for use by only virtual machines.

[Report a bug](#)

**E.1.12. Searching for Storage**

The following table describes all search options for storage.

**Table E.8. Searching for Storage**

Property (of resource or resource-type)	Type	Description (Reference)
Hosts. <i>hosts-prop</i>	Depends on property type	The property of the hosts associated with the storage.
Clusters. <i>clusters-prop</i>	Depends on property type	The property of the clusters associated with the storage.
name	String	The unique name that identifies the storage on the network.
status	String	The status of the storage domain.
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 number of results to display per page.

**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.

[Report a bug](#)

**E.1.13. Searching for Disks**

The following table describes all search options for disks.

Table E.9. Searching for Disks

Property (of resource or resource-type)	Type	Description (Reference)
Datacenters <b><i>datacenters-prop</i></b>	Depends on property type	The property of the data centers associated with the disk.
Storages <b><i>storages-prop</i></b>	Depends on property type	The property of the storage associated with the disk.
alias	String	The human readable name that identifies the storage on the network.
description	String	Keywords or text describing the disk, optionally used when creating the disk.
provisioned_size	Integer	The virtual size of the disk.
size	Integer	The size of the disk.
actual_size	Integer	The actual size allocated to the disk.
creation_date	Integer	The date the disk was created.
bootable	String	Whether the disk can or cannot be booted. Valid values are one of <b>0</b> , <b>1</b> , <b>yes</b> , or <b>no</b> .
shareable	String	Whether the disk can or cannot be attached to more than one virtual machine at a time. Valid values are one of <b>0</b> , <b>1</b> , <b>yes</b> , or <b>no</b> .
format	String	The format of the disk. Can be one of <b>unused</b> , <b>unassigned</b> , <b>cow</b> , or <b>raw</b> .
status	String	The status of the disk. Can be one of <b>unassigned</b> , <b>ok</b> , <b>locked</b> , <b>invalid</b> , or <b>illegal</b> .
disk_type	String	The type of the disk. Can be one of <b>image</b> or <b>lun</b> .
number_of_vms	Integer	The number of virtual machine(s) to which the disk is attached.
vm_names	String	The name(s) of the virtual machine(s) to which the disk is attached.
quota	String	The name of the quota enforced on the virtual disk.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

**Example**

**Disks:** `format = cow and provisioned_size > 8`

returns a list of virtual disks with:

- » Qcow, also known as thin provisioning, format; and
- » an allocated disk size greater than 8 GB.

[Report a bug](#)

**E.1.14. Searching for Volumes**

The following table describes all search options for volumes.

Table E.10. Searching for Volumes

Property (of resource or resource-type)	Type	Description (Reference)
Volume <b><i>cluster-prop</i></b>	Depends on property type	The property of the clusters associated with the volume.
Cluster	String	The name of the cluster associated with the volume.
name	String	The human readable name that identifies the volume.
type	String	Can be one of <b>distribute</b> , <b>replicate</b> , <b>distributed_replicate</b> , <b>stripe</b> , or <b>distributed_stripe</b> .
transport_type	Integer	Can be one of <b>tcp</b> or <b>rdma</b> .
replica_count	Integer	Number of replica.
stripe_count	Integer	Number of stripes.
status	String	The status of the volume. Can be one of <b>Up</b> or <b>Down</b> .
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

**Example**

**Volume:** `transport_type = rdma and stripe_count >= 2`

returns a list of volumes with:

- » Transport type set to RDMA; and
- » with 2 or more stripes.

[Report a bug](#)

**E.1.15. 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.11. Searching for Virtual Machines

Property (of resource or resource-type)	Type	Description (Reference)
Hosts <i>.hosts-prop</i>	Depends on property type	The property of the hosts associated with the virtual machine.
Templates <i>.templates-prop</i>	Depends on property type	The property of the templates associated with the virtual machine.
Events <i>.events-prop</i>	Depends on property type	The property of the events associated with the virtual machine.
Users <i>.users-prop</i>	Depends on property type	The property of the users associated with the virtual machine.
name	String	The name of the virtual machine.
status	List	The availability of the virtual machine.
ip	Integer	The IP address of the virtual machine.
uptime	Integer	The number of minutes that the virtual machine has been running.
domain	String	The domain (usually Active Directory domain) that groups these machines.
os	String	The operating system selected when the virtual machine was created.
creationdate	Date	The date on which the virtual machine was created.
address	String	The unique name that identifies the virtual machine on the network.
cpu_usage	Integer	The percent of processing power used.
mem_usage	Integer	The percentage of memory used.
network_usage	Integer	The percentage of network used.
memory	Integer	The maximum memory defined.
apps	String	The applications currently installed on the virtual machine.
cluster	List	The cluster to which the virtual machine belongs.
pool	List	The virtual machine pool to which the virtual machine belongs.
loggedinuser	String	The name of the user currently logged in to the virtual machine.
tag	List	The tags to which the virtual machine belongs.
datacenter	String	The data center to which the virtual machine belongs.
type	List	The virtual machine type (server or desktop).
quota	String	The name of the quota associated with the virtual machine.
sortby	List	Sorts the returned results by one of the resource properties.
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.

[Report a bug](#)

**E.1.16. Searching for Pools**

The following table describes all search options for Pools.

Table E.12. Searching for Pools

Property (of resource or resource-type)	Type	Description (Reference)
name	String	The name of the pool.
description	String	The description of the pool.
type	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

[Report a bug](#)

### E.1.17. Searching for Templates

The following table describes all search options for templates.

**Table E.13. Searching for Templates**

Property (of resource or resource-type)	Type	Description (Reference)
Vms. <i>Vms-prop</i>	String	The property of the virtual machines associated with the template.
Hosts. <i>hosts-prop</i>	String	The property of the hosts associated with the template.
Events. <i>events-prop</i>	String	The property of the events associated with the template.
Users. <i>users-prop</i>	String	The property of the users associated with the template.
name	String	The name of the template.
domain	String	The domain of the template.
os	String	The type of operating system.
creationdate	Integer	The date on which the template was created. Date format is <i>mm/dd/yy</i> .
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.
quota	String	The quota associated with the template.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

#### Example

**Template:** `Events.severity >= normal and Vms.uptime > 0`

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.

[Report a bug](#)

### E.1.18. Searching for Users

The following table describes all search options for users.

**Table E.14. Searching for Users**

Property (of resource or resource-type)	Type	Description (Reference)
Vms. <i>Vms-prop</i>	Depends on property type	The property of the virtual machines associated with the user.
Hosts. <i>hosts-prop</i>	Depends on property type	The property of the hosts associated with the user.
Templates. <i>templates-prop</i>	Depends on property type	The property of the templates associated with the user.
Events. <i>events-prop</i>	Depends on property type	The property of the events associated with the user.
name	String	The name of the user.
lastname	String	The last name of the user.
username	String	The unique name of the user.
department	String	The department to which the user belongs.
group	String	The group to which the user belongs.
title	String	The title of the user.
status	String	The status of the user.
role	String	The role of the user.
tag	String	The tag to which the user belongs.
pool	String	The pool to which the user belongs.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

#### Example

**Users:** `Events.severity > normal and Vms.status = up or Vms.status = pause`

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.

[Report a bug](#)

### E.1.19. Searching for Events

The following table describes all search options you can use to search for events. Auto-completion is offered for many options as appropriate.

Table E.15. Searching for Events

Property (of resource or resource-type)	Type	Description (Reference)
Vms. <b>Vms-prop</b>	Depends on property type	The property of the virtual machines associated with the event.
Hosts. <b>hosts-prop</b>	Depends on property type	The property of the hosts associated with the event.
Templates. <b>templates-prop</b>	Depends on property type	The property of the templates associated with the event.
Users. <b>users-prop</b>	Depends on property type	The property of the users associated with the event.
Clusters. <b>clusters-prop</b>	Depends on property type	The property of the clusters associated with the event.
Volumes. <b>VoLumes-prop</b>	Depends on property type	The property of the volumes associated with the event.
type	List	Type of the event.
severity	List	The severity of the event: Warning/Error/Normal.
message	String	Description of the event type.
time	Integer	Time at which the event occurred.
usname	String	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.
event_volume	String	The volume associated with the event.
correlation_id	Integer	The identification number of the event.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

#### Example

Events: Vms.name = testdesktop and Hosts.name = gonzo.example.com

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**.

[Report a bug](#)

## E.2. Bookmarks

### E.2.1. Saving a Query String as a Bookmark

#### Summary

A bookmark can be used to remember a search query, and shared with other users.

#### Procedure E.1. Saving a Query String as a Bookmark

1. Enter the desired search query in the search bar and perform the search.
2. Click the star-shaped **Bookmark** button to the right of the search bar to open the **New Bookmark** window.




Figure E.1. Bookmark Icon

3. Enter the **Name** of the bookmark.
4. Edit the **Search string** field (if applicable).
5. Click **OK** to save the query as a bookmark and close the window.
6. The search query is saved and displays in the **Bookmarks** pane.

#### Result

You have saved a search query as a bookmark for future reuse. Use the **Bookmark** pane to find and select the bookmark.

[Report a bug](#)

### E.2.2. Editing a Bookmark

### Summary

You can modify the name and search string of a bookmark.

#### Procedure E.2. Editing a Bookmark

1. Click the **Bookmarks** tab on the far left side of the screen.
2. Select the bookmark you wish to edit.
3. Click the **Edit** button to open the **Edit Bookmark** window.
4. Change the **Name** and **Search string** fields as necessary.
5. Click **OK** to save the edited bookmark.

### Result

You have edited a bookmarked search query.

[Report a bug](#)

### E.2.3. Deleting a Bookmark

#### Summary

When a bookmark is no longer needed, remove it.

#### Procedure E.3. Deleting a Bookmark

1. Click the **Bookmarks** tab on the far left side of the screen.
2. Select the bookmark you wish to remove.
3. Click the **Remove** button to open the **Remove Bookmark** window.
4. Click **OK** to remove the selected bookmark.

### Result

You have removed a bookmarked search query.

[Report a bug](#)

## E.3. Tags

### E.3.1. Using Tags to Customize Interactions with Red Hat Enterprise Virtualization

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.

[Report a bug](#)

### E.3.2. Creating a Tag

#### Summary

You can edit the name and description of a tag.

#### Procedure E.4. Creating a Tag

1. Click the **Tags** tab on the left side of the screen.
2. Select the node under which you wish to create the tag. For example, to create it at the highest level, click the **root** node.
3. Click the **New** button to open the **New Tag** window.
4. Enter the **Name** and **Description** of the new tag.
5. Click **OK** to create the tag.

### Result

The new tag is created and displays on the **Tags** tab.

[Report a bug](#)

### E.3.3. Modifying a Tag

#### Summary

You can edit the name and description of a tag.

#### Procedure E.5. Modifying a Tag

1. Click the **Tags** tab on the left side of the screen.
2. Select the tag you wish to modify.
3. Click **Edit** to open the **Edit Tag** window.
4. Change the **Name** and **Description** fields as necessary.
5. Click **OK** to save the edited tag.

### Result

You have modified the properties of a tag.

[Report a bug](#)

### E.3.4. Deleting a Tag

#### Summary

When a tag is no longer needed, remove it.

#### Procedure E.6. Deleting a Tag

1. Click the **Tags** tab on the left side of the screen.
2. Select the tag you wish to delete.
3. Click **Remove** to open the **Remove Tag(s)** window. The message warns you that removing the tag will also remove all descendants of the tag.

4. Click **OK** to delete the selected tag.

**Result**

You have removed the tag and all its descendants. The tag is also removed from all the objects that it was attached to.

[Report a bug](#)

### E.3.5. Adding and Removing Tags to and from Objects

**Summary**

You can assign tags to and remove tags from hosts, virtual machines, and users.

**Procedure E.7. Adding and Removing Tags to and from Objects**

1. Use the resource tab, tree mode, or the search function to find and select the object(s) you wish to tag or untag.
2. Click the **Assign Tags** button to open the **Assign Tags** window.
3. Select the check box to assign a tag to the object, or deselect the check box to detach the tag from the object.
4. Click **OK**.

**Result**

The specified tag is now added or removed as a custom property of the selected object(s).

[Report a bug](#)

### E.3.6. Searching for Objects using Tags

- » Enter a search query using **tag** as the property and the desired value or set of values as criteria for the search.  
The objects tagged with the specified criteria are listed in the results list.

**See Also:**

- » [Section E.1, "Search"](#)

[Report a bug](#)



## Revision History

<b>Revision 3.2-52</b>	<b>Wed 25 Sep 2013</b>	<b>Jodi Biddle</b>
BZ#1004698 - Fixed issue that was causing text in the PDF to appear tiny/unreadable.		
<b>Revision 3.2-51</b>	<b>Wed 21 Aug 2013</b>	<b>Tim Hildred</b>
Corrected "Virtual Machine Run Once Settings" to note that SPICE is a not supported display protocol for connecting to Windows 8 virtual machines.		
<b>Revision 3.2-50</b>	<b>Mon 19 Aug 2013</b>	<b>Andrew Burden</b>
<a href="#">BZ#996439</a> - Included distinction between clusters that run virtual machines and clusters that run Red Hat Storage servers.		
<b>Revision 3.2-49</b>	<b>Mon 19 Aug 2013</b>	<b>Tim Hildred</b>
<a href="#">BZ#612810</a> - Resolved a confusing typo in "Preparing iSCSI Storage".		
<b>Revision 3.2-48</b>	<b>Fri 16 Aug 2013</b>	<b>Tim Hildred</b>
<a href="#">BZ#921114</a> - Added procedure allowing use of VirtIO devices with Windows virtual machines. <a href="#">BZ#719765</a> - Updated selector to send specific messages to server. <a href="#">BZ#719765</a> - Updated semanage command. <a href="#">BZ#988619</a> - Made it clear that Windows 8 and Server 2012 virtual machines do not have SPICE support. <a href="#">BZ#974515</a> - Updated clone from snapshot content. <a href="#">BZ#965039</a> - Hypervisor hosts cannot be added to Gluster-enabled clusters. <a href="#">BZ#612810</a> - Added "Preparing iSCSI Storage". <a href="#">BZ#985018</a> - The log file created at engine-setup time was incorrectly shown with hyphen rather than an underscore. <a href="#">BZ#977212</a> - Differentiation between hypervisor clusters and storage clusters <a href="#">BZ#986735</a> - Updated host firewall ports to reflect change SPICE range.		
<b>Revision 3.2-47</b>	<b>Fri 16 Aug 2013</b>	<b>Zac Dover</b>
BZ#965337 - adding a full stop, the absence of which was caught by QE		
<b>Revision 3.2-46</b>	<b>Thu 15 Aug 2013</b>	<b>Zac Dover</b>
BZ#978799 - adding Spice proxy configuration information		
<b>Revision 3.2-45</b>	<b>Wed 14 Aug 2013</b>	<b>Zac Dover</b>
No BZ# - test		
<b>Revision 3.2-44</b>	<b>Wed 14 Aug 2013</b>	<b>Zac Dover</b>
No BZ# - restaging to pick up recent changes		
<b>Revision 3.2-43</b>	<b>Thu 01 Aug 2013</b>	<b>Zac Dover</b>
BZ#976169 - documenting Red Hat Support Plugin		
<b>Revision 3.2-43</b>	<b>Tue Jun 25 2013</b>	<b>Tim Hildred</b>
<a href="#">BZ#974366</a> , added cli example of rhn-channel to step 1 <a href="#">BZ#958909</a> , added step requiring user to remove Red Hat Enterprise Virtualization 3.0 channels before upgrading to 3.1		
<b>Revision 3.2-42</b>	<b>Mon 24 Jun 2013</b>	<b>Cheryn Tan</b>
Fixes from localization.		
<b>Revision 3.2-41</b>	<b>Wed 12 Jun 2013</b>	<b>Jodi Biddle</b>
BZ#956777 - Updated out of date information regarding edit template window		
<b>Revision 3.2-40</b>	<b>Tue 11 Jun 2013</b>	<b>Andrew Burden</b>
BZ#954215 - Updated Introduction topic to Templates to be more relevant		
<b>Revision 3.2-39</b>	<b>Tue 11 Jun 2013</b>	<b>Zac Dover</b>
BZ#973010 - the hazards of running services on RHEV nodes		
<b>Revision 3.2-38</b>	<b>Thu 30 May 2013</b>	<b>Jodi Biddle</b>
BZ#957299 - added admonition about stateless virtual machines		
<b>Revision 3.2-37</b>	<b>Mon 27 May 2013</b>	<b>Andrew Burden</b>
BZ#875500 - specified 'user' subject for 'history-readonly' user		
<b>Revision 3.2-36</b>	<b>Mon 20 May 2013</b>	<b>Andrew Burden</b>
Fixed Reports section parameters Corrected instances of 'rhev-hypervisor-advanced' to 'rhev-hypervisor'		
<b>Revision 3.2-35</b>	<b>Fri 10 May 2013</b>	<b>Andrew Burden</b>
Fixed a few minor localisation syncing issues		
<b>Revision 3.2-34</b>	<b>Tue 07 May 2013</b>	<b>Andrew Burden</b>
Corrected several typos		
<b>Revision 3.2-33</b>	<b>Tue 23 Apr 2013</b>	<b>Zac Dover</b>
BZ#871531 - added <guilabel> tags to "Max free Memory for scheduling new VMs"		
<b>Revision 3.2-32</b>	<b>Tue 23 Apr 2013</b>	<b>Zac Dover</b>
BZ#871531 - moved an admonition into a table where it belongs		
<b>Revision 3.2-31</b>	<b>Wed Apr 17 2013</b>	<b>Tim Hildred</b>
BZ#949360 - added program name to the APT topic, and made it clear the APT stands for Application Provisioning Tool		
<b>Revision 3.2-30</b>	<b>Mon Apr 15 2013</b>	<b>Tim Hildred</b>
BZ#950861 - removed unsupported fence agents		
<b>Revision 3.2-29</b>	<b>Thu 11 Apr 2013</b>	<b>Zac Dover</b>
BZ#871531 - documenting "Max free Memory for scheduling new VMs" BZ#889052 - documenting User Spice Session Activity Reports		
<b>Revision 3.2-28</b>	<b>Thu 11 Apr 2013</b>	<b>Cheryn Tan</b>
Changed references of Gluster to Red Hat Storage, minor RHN changes, reports updates.		
<b>Revision 3.2-27</b>	<b>Mon 08 Apr 2013</b>	<b>Tim Hildred</b>
Updated screenshots, moved the utilities content out of appendix, and added part 3 for the reports content.		

<b>Revision 3.2-26</b>	<b>Thu 04 Apr 2013</b>	<b>Tim Hildred</b>
BZ#929350 - added notes to backups content about upgraded environments		
<b>Revision 3.2-25</b>	<b>Tue 02 Apr 2013</b>	<b>Zac Dover</b>
BZ#891732 - information about pinning CPUs to hosts		
<b>Revision 3.2-24</b>	<b>Sun 31 March 2013</b>	<b>Zac Dover</b>
BZ#857789 - one typo		
<b>Revision 3.2-23</b>	<b>Wed 27 March 2013</b>	<b>Zac Dover</b>
BZ#857789 - numerous small issues		
<b>Revision 3.2-22</b>	<b>Tue 26 March 2013</b>	<b>Andrew Burden</b>
BZ#892824 - corrected filepaths for rhvm backup		
<b>Revision 3.2-21</b>	<b>Mon 25 March 2013</b>	<b>Cheryn Tan</b>
Red Hat Storage content for review -- BZ#892437, 887385, 892435, 892376		
<b>Revision 3.2-20</b>	<b>Fri 22 March 2013</b>	<b>Zac Dover</b>
BZ#854247 - s/excced/exceed		
<b>Revision 3.2-19</b>	<b>Fri 22 March 2013</b>	<b>Zac Dover</b>
BZ#924535 - splitting one topic into several to improve granularity and likelihood of reuse		
<b>Revision 3.2-18</b>	<b>Fri 22 March 2013</b>	<b>Zac Dover</b>
BZ#857790 - various minor corrections		
<b>Revision 3.2-17</b>	<b>Thu 21 March 2013</b>	<b>Andrew Burden</b>
BZ#888481 installation update		
<b>Revision 3.2-16</b>	<b>Tue 19 March 2013</b>	<b>Zac Dover</b>
BZ#893210 - Correcting the name of the guest agent package from "rhev-guest-agent" to "rhvm-guest-agent"		
<b>Revision 3.2-15</b>	<b>Mon 18 March 2013</b>	<b>Zac Dover</b>
BZ#889107 - minor fixes		
<b>Revision 3.2-14</b>	<b>Fri 15 March 2013</b>	<b>Zac Dover</b>
BZ#889107 - minor fixes		
<b>Revision 3.2-13</b>	<b>Thu 14 March 2013</b>	<b>Andrew Burden</b>
BZ#912181 - IPA Admin User admonition BZ#912185 - Update VM Resource allocation table BZ#892023 - host installation with tuned BZ#889403 - hotplugging networks		
<b>Revision 3.2-12</b>	<b>Tue 12 March 2013</b>	<b>Cheryn Tan</b>
BZ#889403 - Documented network linking and hotplug options. BZ#889406 - Deleting virtual machines and leaving floating disks. BZ#892390 - Added multi-tier fencing options. BZ#892429 - Documented network permissions.		
<b>Revision 3.2-11</b>	<b>Thu 07 March 2013</b>	<b>Zac Dover</b>
s/UI/ User Interface/g		
<b>Revision 3.2-10</b>	<b>Thu 07 March 2013</b>	<b>Zac Dover</b>
BZ#889107		
<b>Revision 3.2-9</b>	<b>Fri 01 March 2013</b>	<b>Tim Hildred</b>
Staging for Beta 1.		
<b>Revision 3.2-8</b>	<b>Thu 28 Feb 2013</b>	<b>Tim Hildred</b>
Updated Connecting to the History Database content (8281). Added procedure for remote connections. BZ#875500.		
<b>Revision 3.2-7</b>	<b>Tue 26 Feb 2013</b>	<b>Tim Hildred</b>
Changed revision numbers for consistency with other books.		
<b>Revision 3.2-6</b>	<b>Tue 26 Feb 2013</b>	<b>Zac Dover</b>
BZ#889107 - Added topic describing the deployment of ShellnABox User Interface Plugin		
<b>Revision 3.2-5</b>	<b>Fri 22 Feb 2013</b>	<b>Andrew Burden</b>
BZ#891712 - Added VDSM hooks for hot unplug NIC. BZ#889409 - Documented search functions for networks. BZ#889061 - Added iLO2 and iLO4 as fencing options. BZ#889503 - Edited description of vhost custom property. BZ#888515 - Updated rhvm-config values.		
<b>Revision 3.2-4</b>	<b>Wed 20 Feb 2013</b>	<b>Zac Dover</b>
BZ#889107 - RHEV UI Plugins - adding bootstrapping material		
<b>Revision 3.2-3</b>	<b>Fri 15 Feb 2013</b>	<b>Andrew Burden</b>
BZ# 892009 - Added topic to better address events and logs when a VM is restarted due to High Availability		
<b>Revision 3.2-2</b>	<b>Mon 11 Feb 2013</b>	<b>Jodi Biddle</b>
BZ# 881502 - added Windows 8 to supported guest OS list		
<b>Revision 3.2-1</b>	<b>Fri 10 Feb 2013</b>	<b>Zac Dover</b>
BZ#889107 - RHEV UI Plugins Branched from 3.1 docs.		